

Control and Manipulation of Cold Atoms Trapped in Optical Tweezers

Cecilia Muldoon



A thesis submitted for the degree of Doctor of Philosophy
Hilary Term 2012

Supervised by
Dr. Axel Kuhn

The Clarendon Laboratory
Department of Physics
Oxford University

Abstract

The ability to address and manipulate individual information carriers in a deterministic, coherent, and scalable manner is a central theme in quantum information processing^[1]. Neutral atoms trapped by laser light are amongst the most promising candidates for storing and processing information in a quantum computer or simulator, so a scalable and flexible scheme for their control and manipulation is paramount. This thesis demonstrates a fast and versatile method to address and dynamically control the position (the motional degrees of freedom) of neutral atoms trapped in optical tweezers. The tweezers are generated by using the direct image of a Spatial Light Modulator (SLM) which can control and shape a large number of optical dipole-force traps. Trapped atoms adapt to any change in the potential landscape, such that one can re-arrange and randomly access individual sites within atom-trap arrays. A diffraction limited imaging system is used to map the intensity distribution of the SLM onto a cloud of cold atoms captured and cooled using a Magneto Optical Surface Trap (MOST).

Acknowledgements

First and foremost, I would like to thank my supervisor, Dr. Axel Kuhn, for giving me the opportunity to do something I am truly proud of. He took me on as a DPhil student in spite of my minimal lab experience, and gave me the space to learn and grow. He was more helpful and involved than I ever expected a supervisor to be, and I continue to be in awe of the speed with which his formidable intellect finds creative solutions to seemingly intractable problems and plausible explanations for counterintuitive lab mysteries. His patience, trust, and genuine concern have been invaluable to me, both in the lab and on a personal level. Secondly, I would like to thank Dr. Lukas Brandt, for sharing with me the experience of turning an empty room into a working experiment, and Dr. Edouard Brainis, for his patience in teaching me the intricacies of Visual C++ and optimal control theory, and for his foray into isoplanatism. I would also like to thank Jian Dong for all the MOT-moving, fibre-optimising, and general knob-twiddling when I had had enough, and Dustin Stuart for reigniting my excitement in the lab in my last year. Lastly, I would like to thank all the other members of the Atom-Photon-Connection, past and present, who were at some point a sounding board for my ideas and a venting channel for my frustrations.

To my father, who has given me the means, the encouragement, and the reason to follow my dreams.

Contents

1	Introduction	1
2	The Light-Matter Interaction	8
2.1	Two-Level Atom in a Monochromatic Field	8
2.2	The Stark Shift & the Dressed States	12
2.3	The Scattering Rate & the Saturation Intensity	14
2.4	The Scattering Force & the Dipole Force	16
3	Magneto Optical Surface Trap	19
3.1	Rubidium	20
3.2	The Magneto Optical Trap	21
3.3	The Magneto Optical Surface Trap	25
4	Reconfigurable Optical Tweezers	31
4.1	Reconfigurable Optical Tweezers	32
4.2	Trap Depth	36
4.3	Radial and Longitudinal Confinement	40

4.4	Single Atoms	44
4.5	Preliminary Tests of the Optical System	45
4.6	Feasibility Study	58
4.6.1	Dipole trapping of Alkali Atoms	59
4.6.2	The role of the SLM	62
4.6.3	Radial & axial confinement	66
4.7	Ways to single atoms	70
4.8	Single atom detection and transport	72
5	Experimental Setup	75
5.1	MOST Lasers	75
5.2	Laser Stabilisation for the MOST	79
5.3	Acousto Optical Modulators	81
5.4	Around the Vacuum Chamber	84
5.5	The Digital Mirror Device	87
5.6	The DMD Control	92
5.7	The DLX110	96
5.8	Imaging Setup	98
5.9	The Light Sheet	101
5.10	The Electron Multiplying Charge Coupled Device Camera	103
5.11	MOST Computer Control	106
6	Isoplanatism	107

6.1	Introduction	108
6.2	Non-isoplanatism of a pinhole camera	110
6.3	Non-isoplanatism of a thin lens	115
6.4	Non-isoplanatism & large field imaging	118
6.5	Isoplanatic imaging through a thin lens	128
6.6	Experimental investigation	130
6.7	Conclusion	134
6.8	Derivation of condition (6.6)	135
7	Atoms Trapped in Arbitrary Potential Landscapes	136
7.1	Preliminary Results	136
7.2	Ultra-Narrow Light Sheet	141
7.3	Collection Efficiency & Atom Number Calibration	152
8	Deterministic re-arrangement & controlled transport of atoms	154
8.1	Moving Trap Transport	155
8.2	Ballistic Transport	156
8.3	Transport in 3D	163
8.4	Single Ballistic Transport	166
8.5	Double Ballistic Transport	168
8.6	Optimal Transport	169
9	Conclusion	177

A Supplement to The Light-Matter Interaction	179
A.1 The Gauge-Invariant Schrödinger equation	179
A.2 The Dipole Approximation	181
A.3 The Rotating Wave Approximation	182
A.4 Derivation of the Stark Shift	183
A.5 Spontaneous Decay & the Optical Bloch Equations	185
A.6 The Light Force	186
B The Optical Molasses Technique	188
C Doppler Free Saturated Absorption Spectroscopy	192
D Pound Drever Hall Stabilisation	198
E The Electron Multiplying Charge Coupled Device Camera	202
List of Publications	207

Chapter 1

Introduction

Moore's law states that the processing power of a classical computer doubles approximately every two years. This processing power is limited by the number of transistors that can be placed on an integrated circuit, so increasing it is tantamount to the miniaturisation of transistors. In the last decades, this miniaturisation has allowed for ever more powerful microprocessors, and growth in computing power has indeed followed the prediction set forth in Moore's law over forty years ago. However, this miniaturisation is now reaching a limiting value below which quantum mechanical effects will become dominant and transistors will cease to function. Further growth in computing power will thus require a major shift away from transistor technology to a technology which is not hampered by quantum mechanical effects but rather harnesses them. Fortunately, the theoretical framework for a new kind of computing with algorithms based on quantum principles already exists, in the form of quantum computing, and has been the focus of researchers worldwide for several decades.

Seen as esoteric, counter-intuitive, and futuristic from the outside until recent years, quantum computing is increasingly at the centre of attention in the media and its experimental realisation is now clearly not only very pertinent but perhaps also inevitable.

In the simplest of terms, a quantum computer is a computer based on a quantum system, and its advantage is that it should eventually be able to perform certain tasks and computations much faster than a classical computer. As it has been conceived of thus far, a quantum computer at its most basic requires a logic unit, a memory, a bus, and a way to encode and extract information. Thus, the overall structure of a quantum computer does not differ from that of a classical computer. It is the nature of the logic unit that sets the quantum computer apart. Instead of classical bits, which are classical systems that can be in one of two logical states - for instance $|0\rangle$ or $|1\rangle$ - a quantum computer works with quantum bits, which are quantum systems that can once again be in one of two logical states like $|0\rangle$ and $|1\rangle$, but can also be prepared in a coherent superposition of these two states: $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. An example of such a quantum bit is an atom, with $|0\rangle$ and $|1\rangle$ representing two distinct electronic states, and $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ a superposition of these states.

The advantage afforded by the quantum bit, or qubit, is best illustrated by using an example. In a classical computer, a register composed of four bits can be in only one of sixteen possible binary representations of a number at any given time. If this

register were instead composed of four qubits, it could be prepared in a quantum superposition of all sixteen configurations at once. That is to say, it could represent sixteen numbers at the same time. This is what gives the quantum computer what David Deutsch first called an inherent quantum parallelism^[2], for it means that a given mathematical operation can be performed on all the representable numbers in the superposition at the same time, on the same machine. An n -qubit quantum computer needs only one computational step to perform an operation on the 2^n numbers representable by its qubits, whereas a classical computer needs 2^n computational steps or 2^n computers working in parallel to perform the same operation on 2^n numbers. Thus, quantum computing offers a dramatic reduction in the time and processing power necessary to perform computations. This reduction scales with the number of qubits and can therefore be very significant for a large number of qubits.

There are numerous possibilities for the physical realisation of a qubit, such as the polarisation states of a photon, the spin of an electron or nucleus, and the tunnelling of a Cooper pair in a Josephson junction. Each of these physical systems satisfies some but not all of the criteria for the implementation of a real, scalable quantum computer. These criteria, famously summarised by David DiVincenzo^[1], consist of having well-defined qubits, long coherence times, qubit-specific measurements, a universal set of quantum gates, and initialisation of qubits to a pure state. So for instance, entangled photon states are robust against decoherence^[3], but photon creation rates and detection efficiencies are low^[4].

Generally speaking, the ability to address and manipulate qubits in a deterministic, coherent, and scalable manner is the central theme in quantum computing and quantum information processing (QIP)^[1]. Promising examples of matter-based quantum systems that could serve as qubits and have been experimentally explored include trapped ions^[5–8], magnetically trapped atoms^[9,10] and dipole-trapped neutral atoms^[11–13]. The latter are very appealing as the absence of magnetic trapping fields implies no Zeeman shifts and no spin precession, such that the atomic states can be manipulated precisely and preserve coherence. Also, the insensitivity to the electrostatic environment and the largely absent interaction between atoms allow for large-scale atomic arrays close to dielectric surfaces, such as cavity mirrors or hollow-core fibres.

In single optical dipole-force traps, the manipulation of individual atoms has been explored extensively^[14–16], and the scalability of dipole trapping has been shown in optical lattices^[17–19], as well as in dipole-trap arrays created by artificial holograms using spatial phase modulators^[20,21], acousto-optic devices^[22,23], and micro-lens arrays^[24,25]. Our group's approach attempts to combine the scalability of these lattices and arrays with the flexibility of controlling individual trapping sites. It consists of confining neutral atoms in arbitrarily shaped and easily reconfigurable trapping potentials generated using the direct image of an SLM known as a Digital Mirror Device (DMD)^[26,27]. This trapping method has a lot in common with the optical tweezers technique used in biology^[28], where highly focused laser beams are used to

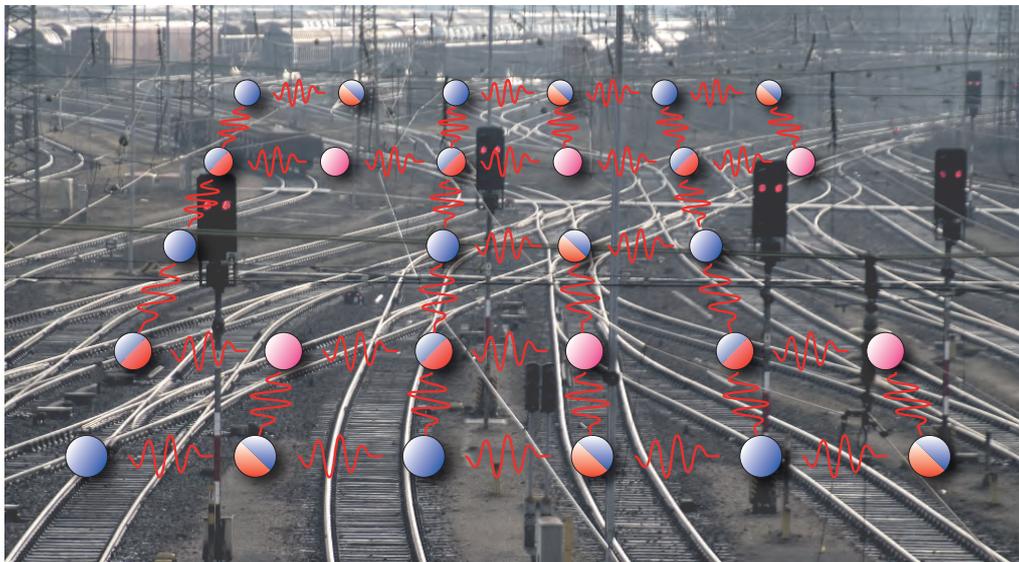


Figure 1.1 Conceptual illustration of the kind of “switchyard” for single atoms that could eventually be implemented using our optical tweezers.

pick up bacteria or uncoil DNA strands, so we refer to it by the same name.

The ideal long term goal of our research is to create a sort of “switchyard” for single atoms, a scalable quantum computing network which uses atoms as matter based stationary qubits and photons as flying qubits connecting the nodes of an array (see Figure 1.1). The otherwise weak interaction between an atom and a photon at each node would be enhanced by the presence of a very small high-finesse optical cavity like a fibre tip cavity^[29]. The idea would be to use our optical tweezers to move atoms in and out of these optical microcavities and use the cavities to realise pairwise entanglement and quantum gate operations through cavity QED^[30], thus mediating the exchange of information between the stationary qubits (the atoms) and the flying qubits (the photons). The tweezers could additionally be used to realise controlled collisions between two atoms in order to implement a two-qubit

gate^[31–33]. In this way, our reconfigurable optical tweezers could be used to implement some of the basic building blocks of QIP and quantum simulation.

So far, we have used these optical tweezers to trap neutral atoms and control their motional degrees of freedom. We have investigated how trapped atoms adapt to changes in the potential landscape, and used this to move them to arbitrary positions. The latter has been accomplished most successfully with a ballistic ‘release and recapture’ scheme in which atoms are transported with minimal heating. In contrast to other trapping schemes, our approach allows for the *fast* and *independent* manipulation and control of a large number of trapping sites. In particular, the DMD attains frame rates which are much faster than our typical trapping frequencies. This is a considerable advance compared to holographic phase modulation techniques that are inherently slow and thus unable to act on the timescale required for successful atom transport.

This thesis gives a description of the work done so far to implement the trapping of atoms with our novel optical tweezers. It is divided into eight chapters. Chapter 1 consists of the introduction and motivation for the work presented in the rest of the thesis, whilst Chapter 2 gives all of the theoretical background necessary to understand it. Chapters 3 and 4 give a conceptual description of the two main components that make up the experiment, the Reconfigurable Optical Tweezers and the Magneto Optical Surface Trap. Chapter 5 describes the entire experimental setup which

is comprised of the setup for both the MOST and the tweezers. Chapter 6 presents a brief digression into isoplanatism, a phenomenon encountered when investigating the imaging system used to generate the tweezers. Chapter 7 shows that atoms can be trapped in single or multiple static traps, and that these traps can have arbitrary shapes, whilst Chapter 8 shows that it is possible to move these trapped atoms to arbitrary positions. It presents the results obtained with our ballistic “release and recapture” scheme, shows how other methods are less effective, and investigates how this transport could be optimised. Finally, Chapter 9 provides a conclusion and our outlook for the future.

Chapter 2

The Light-Matter Interaction

Practically all of the work presented in this thesis relies on the interaction between laser light and atoms, or more specifically, the interaction between a monochromatic light field and a two level atom. This interaction is most easily described by using a semi-classical approach where the atom is treated quantum mechanically and the field is treated classically. The two level description is valid as long as the driving field is resonant or near resonant with the two levels and far detuned from all other levels. By making certain approximations, it is possible to reduce the problem to a set of rate equations with an exact solution which describe the populations of the atomic energy levels. A detailed description of the necessary theoretical concepts referred to in this section can be found in Appendix A.

2.1 Two-Level Atom in a Monochromatic Field

The gauge-invariant Schrödinger equation (see Appendix A)

$$\left\{-\frac{\hbar^2}{2m}[\nabla - i\frac{e}{\hbar}\mathbf{A}(\mathbf{r}, t)]^2 + eU(\mathbf{r}, t)\right\}\psi(\mathbf{r}, t) = i\hbar\frac{\partial\psi(\mathbf{r}, t)}{\partial t}. \quad (2.1)$$

describes the interaction of an electron with an electromagnetic field: the wavefunction $\psi(\mathbf{r}, t)$ represents the electron, and the scalar and vector potentials $A(\mathbf{r}, t)$ and $U(\mathbf{r}, t)$ represent the field. For an electron that is bound by the electrostatic potential $V(\mathbf{r}, t)$ to a nucleus at \mathbf{r}_0 , the interaction Hamiltonian can be simplified by using the *dipole approximation* in which the field wavelength is considered to be larger than the atomic size, and further simplified by applying the gauge transformation $\chi(\mathbf{r}, t) = -\frac{e}{\hbar}\mathbf{A}(\mathbf{r}_0, t) \cdot \mathbf{r}$ and rewriting the electron's wavefunction as $\psi(\mathbf{r}, t) = \exp[\frac{ie}{\hbar}\mathbf{A}(\mathbf{r}_0, t) \cdot \mathbf{r}]\phi(\mathbf{r}, t)$ (see Appendix A). In the radiation gauge in which $U(\mathbf{r}, t) = 0$ and $\nabla \cdot \mathbf{A} = 0$, such that $\mathbf{E} = -\dot{\mathbf{A}}$, this equation takes the simple form

$$i\hbar\dot{\phi}(\mathbf{r}, t) = [\mathcal{H}_0 - e\mathbf{r} \cdot \mathbf{E}(\mathbf{r}, t)]\phi(\mathbf{r}, t) \quad (2.2)$$

where $\mathcal{H}_0 = \frac{p^2}{2m} + V(r)$. The second term in equation (2.2) can be written as $\mathcal{H}_1 = -e\mathbf{r} \cdot \mathbf{E}(\mathbf{r}, t)$. Thus the total Hamiltonian for an atom interacting with a light field can be expressed as the sum of an unperturbed Hamiltonian \mathcal{H}_0 and an interaction Hamiltonian \mathcal{H}_1 :

$$\mathcal{H} = \mathcal{H}_0 + \mathcal{H}_1. \quad (2.3)$$

If the assumption is made that the applied field is nearly resonant with the atomic transition of interest, then only the ground and excited levels involved in

that transition will have a non-negligible probability of being populated, and the atom can be treated as a simple two-level system. Off-resonant couplings lead to a shift in the atomic level structure (the Stark shift), but they do not lead to any population transfer. Additionally, if the difference between the ground and excited levels of the atom is large when compared to its thermal energy $k_B T$, then it is safe to assume that in the steady state, the atom remains in the ground state.

The atomic wavefunction in Dirac notation can be represented by the state vector $|\psi\rangle$. This state vector can be expanded in an orthonormal basis as $|\psi\rangle = \sum_i |n\rangle \langle n|\psi\rangle = \sum_i c_n(t) |n\rangle$, where the probability of the atom being in the state n at time t is given by $P_n(t) = |c_n(t)|^2$. The two-level approximation amounts to truncating this sum to include only two terms:

$$|\psi\rangle = c_1(t)|1\rangle + c_2(t)|2\rangle \quad (2.4)$$

At the position of the atom, the perturbing electromagnetic field $\mathbf{E}(\mathbf{r}, t)$ is given by $\mathbf{E} = |\mathbf{E}_0| \hat{\mathbf{e}}_x \cos(\omega t)$, where \mathbf{E}_0 is its amplitude, and ω is its angular frequency. For the sake of simplicity, it is assumed that this field is linearly polarised along the x axis. The Schrödinger equation to be solved is then

$$i\hbar \frac{d}{dt} (c_1(t)|1\rangle + c_2(t)|2\rangle) = [\mathcal{H}_0 - e\mathbf{r} \cdot \mathbf{E}(\mathbf{r}, t)] (c_1(t)|1\rangle + c_2(t)|2\rangle). \quad (2.5)$$

In order to solve (2.5) more easily, it is instructive to look at the two parts of the

Hamiltonian separately. By using the completeness relation $|1\rangle\langle 1| + |2\rangle\langle 2|$, it is possible to rewrite the unperturbed part of the Hamiltonian \mathcal{H}_0 as

$$\begin{aligned}\mathcal{H}_0 &= (|1\rangle\langle 1| + |2\rangle\langle 2|)\mathcal{H}_0(|1\rangle\langle 1| + |2\rangle\langle 2|) \\ &= \hbar\omega_1|1\rangle\langle 1| + \hbar\omega_2|2\rangle\langle 2|\end{aligned}\quad (2.6)$$

and the interaction Hamiltonian \mathcal{H}_1 as

$$\begin{aligned}\mathcal{H}_1 &= -ex|\mathbf{E}_0|\cos(\omega t) \\ &= -e(|1\rangle\langle 1| + |2\rangle\langle 2|)x(|1\rangle\langle 1| + |2\rangle\langle 2|)|\mathbf{E}_0|\cos(\omega t) \\ &= -(\mu|1\rangle\langle 2| + \mu^*|2\rangle\langle 1|)|\mathbf{E}_0|\cos(\omega t).\end{aligned}\quad (2.7)$$

where $\mu = \langle 1|x|2\rangle$ is the *dipole moment*, and $\langle 1|x|1\rangle = \langle 2|x|2\rangle = 0$ reflects the fact that the operator x changes a symmetric eigenstate into an antisymmetric one and vice versa. By inserting the simplified forms of \mathcal{H}_0 and \mathcal{H}_1 into (2.5) and taking its inner product with $\langle 1|$ and $\langle 2|$ respectively, we can obtain a set of coupled differential equations for the probability amplitudes $c_1(t)$ and $c_2(t)$:

$$\begin{aligned}\dot{c}_1(t) &= -i\omega_1c_1(t) + i\Omega\cos(\omega t)c_2(t) \\ \dot{c}_2(t) &= -i\omega_2c_2(t) + i\Omega^*\cos(\omega t)c_1(t),\end{aligned}\quad (2.8)$$

where $\Omega = \frac{\mu|\mathbf{E}_0|}{\hbar}$ is the *Rabi frequency*. To solve these equations, we use the Rotating Wave Approximation (RWA), in which fast oscillating terms in a Hamiltonian can

be neglected due to their negligible effect on the time evolution of a system (see Appendix A). The solution to these equations is given in equation (A.13). If the atom starts in the ground state, then $c_1(0) = 1$ and $c_2(0) = 0$. Furthermore, when the driving field is on resonance with the atomic transition, $\Delta = 0$, and the populations of the ground and excited states, which are given by the squares of the probability amplitudes, reduce to

$$\begin{aligned} P_1(t) &= |c_1(t)|^2 = \cos^2\left(\frac{\Omega t}{2}\right) \\ P_2(t) &= |c_2(t)|^2 = \sin^2\left(\frac{\Omega t}{2}\right). \end{aligned} \tag{2.9}$$

This means that the atom oscillates between the ground and excited states: at $\Omega t = 0$, the atom is in the ground state, at $\Omega t = \pi$, the whole population has been transferred to the excited state, and at $\Omega t = 2\pi$, it is all back in the ground state. These are called *Rabi Oscillations*. Rabi Oscillations only occur in the presence of a strongly interacting monochromatic field that is on resonance with the atomic transition: a weak field would leave most of the population in the ground state, and a far detuned field would only mix the states, as will be described in the next section.

2.2 The Stark Shift & the Dressed States

In addition to changing the atomic populations, the applied field will have an effect on the atom's level structure; namely, it will cause a shift in its states called the Stark

Shift. This effect is dominant when the detuning from the atomic transition is large and the effect of absorption is negligible. In order to calculate this light shift, we consider again the equations of motion for the probability amplitudes in the RWA, and make a transformation that is equivalent to eliminating the time dependence of the interaction Hamiltonian by switching to a rotating frame. The resulting equations can be written in matrix form (see Appendix A), and the eigenvalues of the matrix can be found by solving the characteristic equation $\det(A-\lambda I)$. This yields $\lambda = \pm \frac{1}{2}\sqrt{\Delta^2 + \Omega^2}$, which means that the eigenenergies of the system are given by

$$\frac{\hbar}{2}(\Delta \pm \sqrt{\Delta^2 + \Omega^2}) \quad (2.10)$$

In the case where $\Omega = 0$, the energy levels are unperturbed, and simply Δ apart; in the far-detuned case where $\Delta \gg \Omega$, the eigenvalues can be written as

$$\lambda \approx \pm \left(\frac{\Delta}{2} + \frac{\Omega^2}{4\Delta} \right) \quad (2.11)$$

which means that if Δ is negative (i.e. the driving field is red detuned), the ground state energy level is shifted down by $\frac{\hbar\Omega^2}{4\Delta}$, whilst the excited state energy level is shifted up by $\frac{\hbar\Omega^2}{4\Delta}$. If Δ is positive (i.e. the driving field is blue detuned), the shifts are in the opposite directions. The total shift in energy is then clearly

$$\Delta E_{Stark} = \frac{\hbar\Omega^2}{2\Delta} \quad (2.12)$$

This shift of the energy levels is known as the ac Stark Shift, and can be used to trap

atoms because it is proportional to the intensity of the driving field. Additionally, if the field is red detuned, then the atoms will be attracted to the regions of highest intensity. This means that it is possible to trap atoms in the focus of a red detuned laser beam, and is the basic concept behind our optical tweezer setup.

In addition to shifting the atomic energy levels, the driving field also mixes the ground and excited states. The eigenstates of the coupled system, known as the *dressed states* are given by

$$\begin{aligned} |\Phi_1\rangle &= \cos\theta|\tilde{1}, n\rangle - \sin\theta|\tilde{2}, n-1\rangle \\ |\Phi_2\rangle &= \cos\theta|\tilde{1}, n\rangle + \sin\theta|\tilde{2}, n-1\rangle \end{aligned} \quad (2.13)$$

where $|\tilde{1}, n\rangle = \tilde{c}_1(t)|1, n\rangle$ and $|\tilde{2}, n-1\rangle = \tilde{c}_2(t)|2, n-1\rangle$, and $\tilde{c}_1(t)$ and $\tilde{c}_2(t)$ have solutions of the form $\exp(-i\lambda t)$, and n is the number of photons. Essentially, in the dressed state picture, the atom and the field are considered jointly, the states $|\tilde{1}, n\rangle$ and $|\tilde{2}, n-1\rangle$ are nearly degenerate, and the ac Stark Shift serves to lift this degeneracy.

2.3 The Scattering Rate & the Saturation Intensity

The treatment so far has not taken into account spontaneous emission from the two-level system. To model the population dynamics of the two-level system including decay, the Optical Bloch Equations must be used (see Appendix A). The solution

to these equations in the steady state gives the population of the excited state as

$$\rho_{22} = \frac{1}{2} \frac{\frac{\Omega^2}{2}}{\Delta^2 + \frac{\Omega^2}{2} + \frac{\gamma^2}{4}} \quad (2.14)$$

From this, we can calculate the Scattering Rate, which is central to describing the cooling mechanism used in a MOT, as well as to calculating the number of photons emitted by trapped atoms in the experiment. The Scattering Rate is given by $R_{Scatt} = \gamma\rho_{22}$, with ρ_{22} as given in (2.14). At very high intensities where Ω is large, it is reduced to $R_{Scatt} = \frac{\gamma}{2}$. In the case where the detuning is large, the denominator is reduced to Δ^2 , and the scattering rate can be given by

$$R_{Scatt} = \frac{\gamma\Omega^2}{4\Delta^2} \quad (2.15)$$

It is instructive to compare the Scattering Rate in this limit with the expression for the ac Stark Shift in the same limit where $\Delta \gg \Omega$. Whilst the scattering rate goes as $\frac{1}{\Delta^2}$, the Stark Shift goes as $\frac{1}{\Delta}$. This is relevant, because the Stark Shift is equivalent to the depth of the dipole trapping potential. Therefore, a larger detuning reduces the Scattering Rate faster than it reduces the depth of the trap, which is advantageous when implementing a dipole trap.

Finally, the Scattering Rate and the ac Stark Shift can also be expressed in terms of the *Saturation Intensity*, which is the intensity at which the atom spends one fourth of the time in the excited state^[34], and is given by

$$I_{Sat} = \frac{I\gamma^2}{2\Omega^2} = \frac{\pi\hbar c\gamma}{3\lambda^3} \quad (2.16)$$

In terms of I_{Sat} , the Scattering Rate and ac Stark Shift are given by

$$R_{Scatt} = \frac{\gamma^3 I}{8\Delta^2 I_{Sat}} \quad (2.17)$$

$$\Delta E_{Stark} = \frac{\hbar\gamma^2 I}{8\Delta I_{Sat}} \quad (2.18)$$

2.4 The Scattering Force & the Dipole Force

The scattering force responsible for the cooling of atoms in a MOT and the gradient force responsible for the trapping of an atom in a dipole trap are simply two facets of the light matter interaction described so far. In fact, the total force exerted on an atom in an electromagnetic field can be expressed as a sum of these two forces. The light force can be written in terms of the expression for ρ_{21} obtained from solving the Optical Bloch Equations as

$$F = \frac{\hbar\frac{\Omega^2}{2}}{\Delta^2 + \frac{\Omega^2}{2} + \frac{\gamma^2}{4}} \left[-q_r\Delta + q_i\frac{\gamma}{2} \right], \quad (2.19)$$

where q_r corresponds to the gradient of the irradiating field's amplitude and q_i corresponds to the gradient of its phase (see Appendix A). The values of q_r and q_i depend on the particular field. A standing wave, for example, has an amplitude gradient but no phase gradient, so q_r is non-zero and $q_i = 0$. A traveling wave, on the other hand, has a phase gradient but no amplitude gradient. In this case $q_i = \vec{k}$ and $q_r = 0$.

The first term in equation (2.19) corresponds to the gradient of the ac Stark Shift derived in section 2.2, and gives rise to a conservative force called the Dipole Force:

$$F_{Dip} = \frac{\partial \Delta E_{Stark}}{\partial z} = \frac{\hbar \Omega}{2\Delta} \frac{\partial \Omega}{\partial z} \quad (2.20)$$

It is proportional to the gradient of the intensity of the applied field as well the detuning of the field from the atomic transition. It disappears when $\Delta = 0$ and when there is no intensity gradient ($q_r = 0$). To obtain the requisite intensity gradient, a standing wave or a strongly focused Gaussian laser beam is needed, as this spatial modulation is not present in a traveling wave. The effect of the Dipole Force on an atom is to drive it to regions of higher intensity if $\Delta < 0$, and to regions of lower intensity if $\Delta > 0$: this is essential to the working of a dipole trap.

The second term in equation (2.19) is directly proportional to the intensity of the applied field, and corresponds to absorption followed by spontaneous emission, as can be seen from the fact that it is proportional to the decay rate γ . Upon absorbing a photon, an atom receives a momentum kick of $\hbar \vec{k}$ along the same direction as the wavevector \vec{k} of the incoming photon. The recoil associated with spontaneously emitted photons, however, is in a random direction, so over many emissions, the recoil averages to zero. This means that the total momentum transfer depends only on the impinging field.

When $\Delta = 0$, equation (2.19) is reduced to the photon momentum $\hbar\vec{k}$ times the scattering rate derived in section 2.3, which in turn is simply the product of the decay rate γ and the population of the excited state ρ_{22} (recalling that in the case of a travelling wave $q_i = \vec{k}$):

$$F_{Scatt} = \hbar\vec{k}\gamma\rho_{22} \quad (2.21)$$

This dissipative force, called the Scattering Force, has a maximum value of $F = \frac{\hbar\vec{k}\gamma}{2}$ ($\frac{1}{2}$ is the maximum population of the excited state) and is crucial to the working of a MOT.

Chapter 3

Magneto Optical Surface Trap

In order to load atoms into our reconfigurable optical dipole force traps, or optical tweezers, it is first necessary to capture them from the background and cool them to a temperature that is lower than the depth of the tweezers. This is because, as was discussed in section 2.4, the dipole force that gives rise to these trapping potentials is a conservative force and cannot be used as a cooling mechanism. The optical molasses technique described in Appendix B is a successful cooling mechanism, but taken alone, it cannot be used to trap atoms, because the Scattering Force it relies on is not a position-dependent force. In order to both cool and localise the atoms, a Magneto Optical Trap (MOT) can be used: this is the most widely used method of trapping atoms in the laboratory. However, the mirror present in our tweezer geometry prevents us from using a standard MOT configuration. We circumvent the problem by using a Magneto Optical *Surface* Trap (MOST).

3.1 Rubidium

To explain the workings of a MOT, it is necessary to move on from the simple two-level atom considered so far, and consider a multi-level atom, in specific, the element used in our experiment, ^{87}Rb . Rubidium is widely used for cold atom experiments because as an alkali atom, it has a closed shell with a single valence electron, and thus a relatively simple electronic structure. ^{87}Rb is the less abundant isotope of Rubidium, but it has a simpler level structure than ^{85}Rb , and is therefore easier to work with. In addition, the transition frequency from the lowest to the first excited state is accessible with near infrared semiconductor diode lasers, which have a fast and relatively simple electronic tunability. Lastly, the room temperature vapour pressure of Rubidium is sufficient to carry out spectroscopy with it.

In its ground state, ^{87}Rb 's one valence electron occupies the $5^2S_{1/2}$ state, and the first excited states, $5^2P_{1/2}$ and $5^2P_{3/2}$, are 795 nm and 780 nm away respectively. The fine structure of ^{87}Rb is relatively simple. The ground state has angular momentum $L = 0$, and spin $S = \frac{1}{2}$, which means it has a total angular momentum of $J = \frac{1}{2}$. The excited state has angular momentum $L' = 1$ and since, again $S = \frac{1}{2}$, the total angular momentum is either $J = \frac{1}{2}$ or $J = \frac{3}{2}$. The level with $J = \frac{1}{2}$ is 15 nm away from the level with $J = \frac{3}{2}$. The transition from the ground state to the $J = \frac{1}{2}$ level is called the D1 line, and has a wavelength of 795nm; the transition from the ground state to the $J = \frac{3}{2}$ level is called the D2 line, and has a wavelength of 780nm. It is the latter transition that we work with.

The hyperfine structure of ^{87}Rb is slightly more complicated. It has a nuclear spin of $I = \frac{3}{2}$, which means that the total angular momentum of the ground state (for the D2 line) is $F = (1, 2)$, whilst that of the excited state is $F' = (0, 1, 2, 3)$. The usual transition rules for optical dipole transitions apply, i.e., $\Delta L = \pm 1$ and $\Delta M_L = 0, \pm 1$, and hence, $\Delta F = 0, \pm 1$ and $\Delta M_F = 0$ unless $\Delta F = 0$ and $M_F = 0$.

The transition between the $F = 2$ and the $F' = 3$ levels of the D2 line is of special interest, as it is a *cycling transition*: once it is excited to the $F' = 3$ state (by σ^+ light), the atom can only decay back to the $F = 2$ state because the transition rules forbid it from decaying to the $F = 1$ state. This transition can therefore effectively be treated as a two level system, which means the optical molasses technique can be used as a cooling mechanism. Of course that off-resonant optical pumping into other levels would lead to atoms being lost from the cycling transition, and would render the cooling ineffective. In this particular case, there is a small possibility that the atom could be excited to the $F' = 2$ level, which *can* decay to the $F = 1$ level. This can be resolved by the inclusion of a second laser that can pump the population from the $F = 1$ level back into the $F = 2$ level (via, for example, the $F' = 2$ level).

3.2 The Magneto Optical Trap

The simple addition of a weak quadrupole magnetic field to the six beam configuration shown in Figure B.2 turns the optical molasses cooling mechanism into an

effective cooling *and* trapping method for neutral atoms, a Magneto Optical Trap (MOT). This magnetic field does not trap the atoms in and of itself; it merely serves to produce a position dependent Zeeman shift in the atomic energy levels that turns the imbalance in the scattering forces into a trapping mechanism.

To create this quadrupole magnetic field, a pair of coils in an anti-Helmholtz configuration must be used (in an anti-Helmholtz configuration, the current in the two coils flow in opposite directions). This configuration can be seen in Figure 3.1. The field produced is zero at the centre of the two coils, and increases linearly in every direction away from this point. The uniform field gradient around this central point induces a Zeeman shift of the excited state that increases linearly with the field strength. Because the field strength is position-dependent, the magnitude of the Zeeman shift is also position-dependent.

Figure 3.2 shows how this position-dependent Zeeman shift can be used to trap atoms. For simplicity, only one dimension is considered, and an excited state with $J = 1$ is used. In this case, the level gets split into the triplet $M_J = 0, \pm 1$. The counter-propagating beams are all circularly polarised, and each pair has opposite handedness, i.e., one is σ^+ and the other σ^- . They are all derived from the same laser, and have a frequency that is slightly red-detuned from resonance, as required for the functioning of an optical molasses. As an atom moves away from the centre in the $+z$ direction, the $M_J = -1$ level shifts closer to the frequency of the incoming

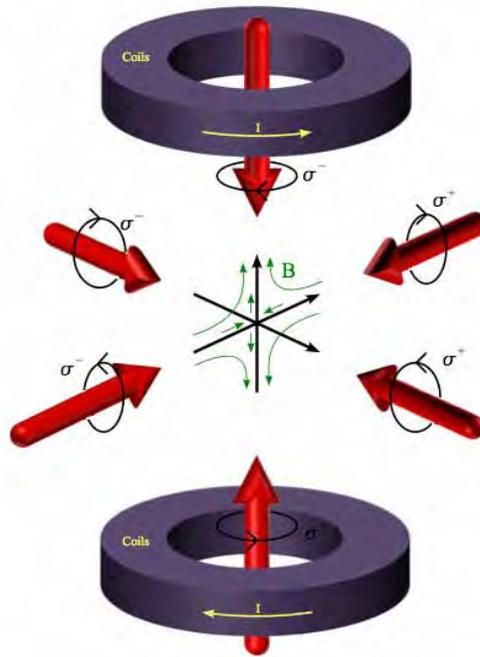


Figure 3.1 A Magneto Optical Trap results from a combination of an Optical Molasses and a pair of Helmholtz coils, which produce a quadrupole magnetic field that increases linearly away from a zero point. The three counter-propagating Optical Molasses beams must have the correct circular polarisation and intersect at the zero-point of the quadrupole field produced by the coils. The Figure is taken from^[35].

laser propagating in the $-z$ direction, which has the right polarisation to excite this transition (in this case σ^-). This leads to a force in the $-z$ direction that pushes the atoms back to the central point. As an atom moves away from centre in the $-z$ direction, the $M_J = 1$ level shifts closer to the frequency of the incoming laser propagating in the $+z$ direction, which has the right polarisation to excite *this* transition (in this case σ^+). This leads to a force in the $+z$ direction that pushes *these* atoms back to the central point.

The force experienced by the atoms in one dimension (in this case along the z direction) can be found by incorporating the Zeeman shift into the same expression

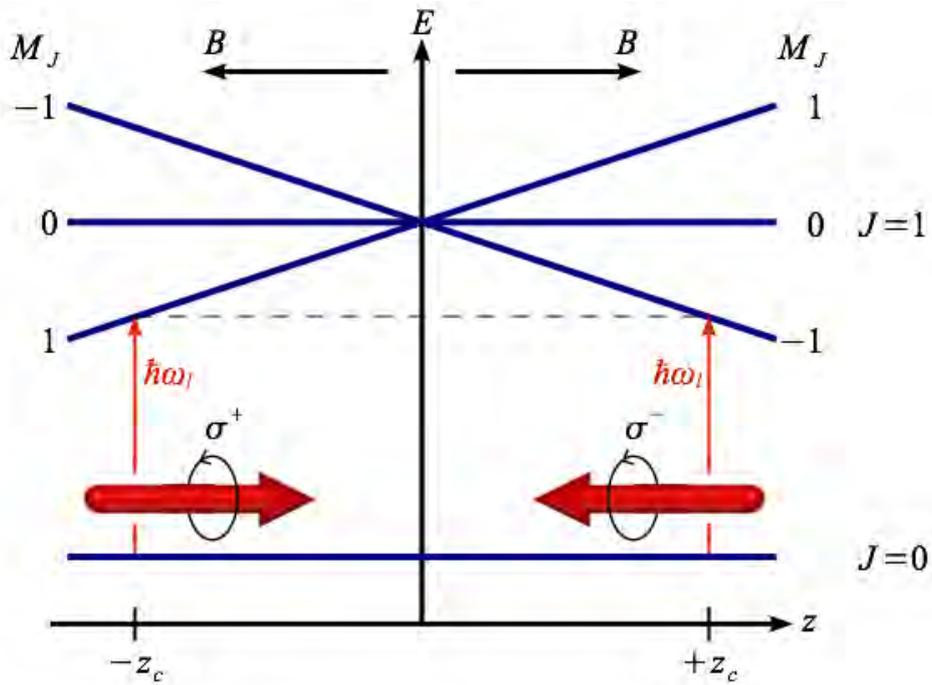


Figure 3.2 In the presence of a magnetic field gradient, the Zeeman splitting of the atomic sublevels depends on the position of the atom^[35]. This means that the radiative force imparted by the circularly polarised Optical Molasses beams becomes position dependent. The selection rules for transitions between Zeeman states lead to an imbalance in this force that pushes the atoms back into the centre of the trap if the polarisation of all the beams is chosen correctly. The Figure is taken from^[35].

used to describe the optical molasses technique (see Appendix B) :

$$F_{MOT} = F_{Scatt}^{\sigma^+}(\omega - (\delta\omega_a + \beta z) - kv) - F_{Scatt}^{\sigma^-}(\omega - (\delta\omega_a - \beta z) + kv) \approx -2\frac{\delta F}{\delta\omega}kv - 2\frac{\delta F}{\delta\omega}\beta z \quad (3.1)$$

where βz is the Zeeman shift at z :

$$\beta z = \frac{g\mu_B}{\hbar} \frac{dB}{dz} z. \quad (3.2)$$

This force takes the form

$$F_{MOT} = -\alpha v - \frac{\alpha\beta}{k}z \quad (3.3)$$

which means that the total force now consists of a cooling part $-\alpha v$ as well as a trapping part $-\frac{\alpha\beta}{k}z$ with spring constant $-\frac{\alpha\beta}{k}$. This same treatment can be extended to three dimensions, with the exception that the field gradient must be different along one of the dimensions in order to satisfy $\nabla B = 0$. More specifically, it should be the case that $\frac{dB_x}{dx} = \frac{dB_y}{dy} = -\frac{1}{2}\frac{dB_z}{dz}$, which is satisfied by the quadrupole field shown in Figure 3.1.

3.3 The Magneto Optical Surface Trap

Unfortunately, the mirror used in our tweezer geometry makes it impossible to trap and cool the atoms using the MOT introduced in the previous section because it would be impossible to implement the requisite free space beam configuration. In order to have our source of cold atoms directly above the mirror used in the tweezer geometry, we must instead use a Magneto Optical *Surface* Trap (MOST)^[36]. The MOST makes use of the fact that upon reflection off a mirror surface at 45° , the helicity of a beam is inverted: an incoming σ^+ beam will result in an outgoing σ^- beam, and vice versa. Due to the finite width of the beams, there is a triangular region directly above the mirror surface where the incoming and outgoing reflected beam overlap. This is illustrated in Figure 3.3. By overlapping such a σ^+/σ^- beam with a counter-propagating σ^-/σ^+ beam and adding a third retro-reflected beam

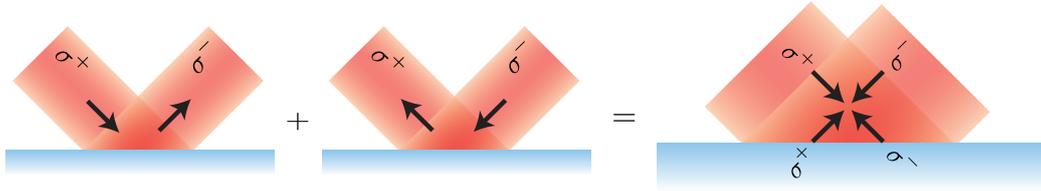


Figure 3.3 Beam configuration for a MOST. The triangular region above the mirror where the incoming and outgoing reflected beam overlap. Adding a third pair of orthogonal counter-propagating beams allows for 3D cooling in this region.

which is orthogonal to the first two and parallel to the mirror surface, it is possible to create a 3D beam configuration in the triangular area just above the mirror surface which is equivalent to the beam configuration present in a free space MOT, but tilted at 45° .

To make a MOT with this tilted beam configuration, the requisite quadrupole field would also have to be tilted by 45° , as the field gradients must be collinear with the beams for the trapping mechanism to work. This means the usual anti-Helmoltz coils would have to be tilted at 45° with respect to the mirror. Due to the nature of our vacuum chamber, this would restrict our optical access and prove to be a great complication. To circumvent this problem, it is possible to create an approximate quadrupole field by superimposing the field created by a square U-shaped rod with a flattened end with a constant bias field (see Figure 3.4). This configuration will give a good quadrupole field near the zero point, but it deteriorates with distance as the field lines start to tilt past the angle where they are collinear with the beams. As will be discussed later, this limits the range in which the trap can be operated.

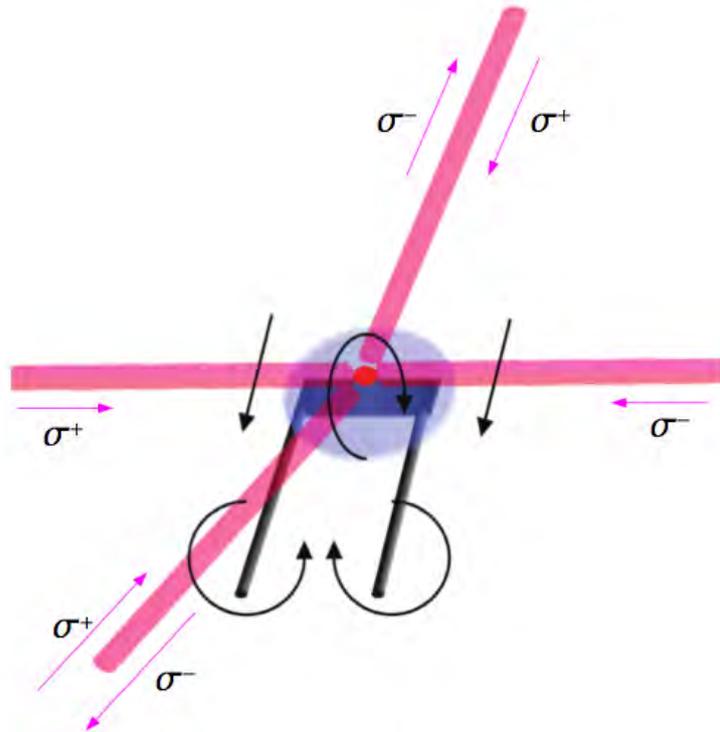


Figure 3.4 The superposition of a U-shaped rod and a constant bias field provides a quasi-quadrupole field whose gradient is approximately collinear with the tilted beam configuration shown in Figure 3.3.

To understand how this quadrupole field is produced, it is necessary to first consider the field produced by superimposing the field generated by a single infinitely long current carrying wire and a constant bias field. Figure 3.5 shows how a zero-point occurs where the bias field and the field of the wire have the same magnitude but point in the opposite direction. This is what occurs in the plane perpendicular to crossbar of the U-shaped rod, which is flattened so as to reduce the circularity of the field which it produces: a more circular field makes the field lines change angle more quickly, contributing to the aforementioned degradation of the quadrupole field. This two-dimensional confinement is extended to three dimensions by the two

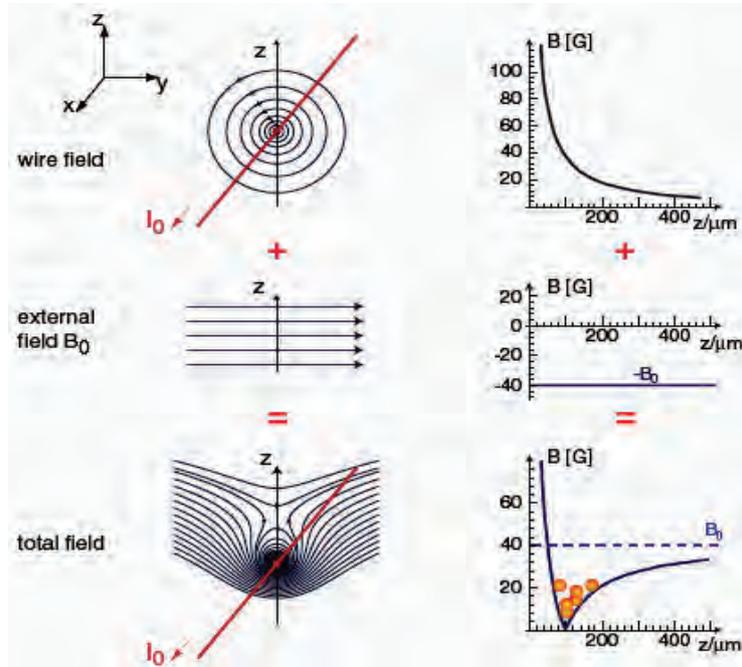


Figure 3.5 Combining the circular field from a straight wire with a constant bias field produces a field which is approximately quadrupole in one dimension. This corresponds to the field gradient along one axis of the U-shaped rod. The Figure is taken from^[37].

sides of the U-shaped rod, which produce field lines perpendicular to the trapping plane of the crossbar. The net effect of all these fields is to produce a zero point above the centre of the crossbar.

One inevitable aspect of the U-shaped rod is that two dimensions of the field will always be coupled. More specifically, the field along the x direction cannot be changed without changing the field along the z direction. This means that it is impossible to move the centre of the MOST (and hence the trapped atoms) along x without also moving them along z. This can be overcome by the use of external Helmholtz coils that produce a constant bias field along one of the two dimensions.

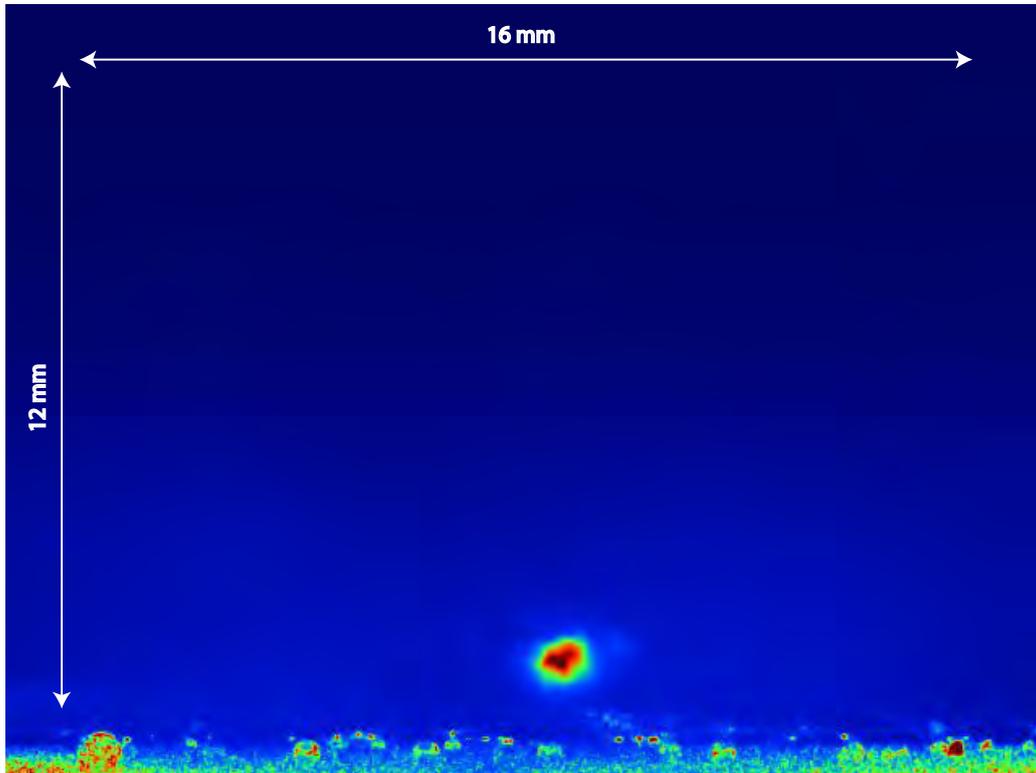


Figure 3.6 False colour absorption image of the MOST. The cloud lies about 1-2 mm above the surface of the mirror, and has a radius of $\sim 200\mu\text{m}$. It contains approximately 10^7 atoms, which have a temperature of $\sim 280\mu\text{K}$.

In fact, extra coils can be used to add bias fields along any of the three dimensions, and as long as these fields remain small, the total field will remain approximately quadrupole and the MOST will function. This extra flexibility allows for the MOST to be moved around significantly.

Figure 3.6 shows a false colour absorption image of the MOST we use as a source of cold atoms in our experiment. We observe the MOST using a CCD camera with a field of view of approximately 12 mm, which means that the MOST in this picture lies about 1-2 mm above the surface of the mirror. The closest the MOST can

be brought to the mirror is about .2 mm; if it is brought any closer, the number of trapped atoms becomes negligible. In order to load cold atoms from this MOST into our optical tweezers, the cloud is positioned $\sim 500\mu\text{m}$ above the surface of the mirror. At that distance, the cloud has a radius of $\sim 120\mu\text{m}$, and contains approximately 10^5 atoms. Its temperature, measured by time of flight (see Section 7.1), is $\sim 280\mu\text{K}$.

Chapter 4

Reconfigurable Optical Tweezers

As outlined in the introduction, the ultimate goal of the experiment is to use the reconfigurable optical tweezers to dynamically control the motional degrees of freedom of cold atoms, with the eventual goal of using these tweezers to move atoms in and out of optical micro-cavities in order to realise pairwise entanglement or quantum gate operations through cavity QED^[30], or using them to realise controlled collisions between two atoms in order to implement a two-qubit gate^[31-33]. The MOST outlined in the previous chapter provides the source of cold atoms used in the experiment. This chapter describes the second part of the experiment, which consists of the Spatial Light Modulator and imaging system used to generate the optical tweezers, as well as the detection system used to observe the trapped atoms. Experimentally, this involved developing and testing the diffraction limited optical

system to image the surface of the DMD onto the atoms, integrating it into the MOST setup, writing the computer program to control the DMD, and performing a careful calibration of the detection system. Most of the experimental details involved in this are left for the following chapter, which deals with the experimental setup, and only a description of the preliminary tests on the optical system is given towards the end of this chapter: the rest of it is mainly conceptual. Further descriptions of the tweezer setup can be found in^[61] and^[40].

4.1 Reconfigurable Optical Tweezers

The reconfigurable optical dipole-force traps, or optical tweezers, at the core of our experiment are produced by imaging the surface of a SLM onto a cloud of cold ^{87}Rb atoms lying directly above a mirror (the MOST described in the previous chapter). The atoms are confined in the potential due to the resulting intensity distribution in the image plane. This potential takes the form

$$U_{dip} \approx \frac{\hbar\Omega^2}{4\delta} \quad (4.1)$$

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 (see Section 2.2), and its gradient gives rise to the conservative dipole force introduced in Section 2.4. 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.

The most widely used type of SLM is the liquid-crystal device (LCD), which has the advantage that it can create grayscales. However, its disadvantage is that it is slow, with refresh rates that are insufficient for the dynamic control of trapped atoms. Our SLM is a Digital Mirror Device (DMD), which, on the contrary, has a high refresh rate but by definition cannot create grayscales. Nevertheless, it is possible to mimic these grayscales, and the refresh rate is sufficient for the real-time manipulation of atoms.

The DMD consists of an array of 1024 by 768 independently addressable micro-mirrors 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 DMD can display either static patterns or movies made from frames of these patterns. The DMD allows us, for example, to move a single trap from one edge of the mirror array to the other by simply switching off a line of mirrors on the trailing edge of the trap, and turning on a new line at the leading edge.

The surface of the DMD is imaged onto the atoms using a two-lens microscope, as shown in Figure 4.1. 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 also experience some degree of longitudinal confinement: atoms that are cold enough are confined to anti-nodal planes stacked along the optical axis. This effect only partially applies to traps smaller than $\sim 20\mu m$ because the reflected beam diverges too quickly. The microscope has a (de)magnification of 57:1, which, given the size of the DMD pixels, 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.

The microscope system consists of a high numerical aperture aspheric lens which is mounted inside the vacuum chamber, and an achromatic doublet which is placed outside the vacuum chamber. We use a high numerical aperture aspheric lens so that our optical tweezers have the smallest possible focii, and hence the best possible axial confinement. An important feature of our microscope system is its confocal geometry: trapped atoms are detected by fluorescence imaging through the same two lenses used to produce the tweezers. Figure 4.1 shows how the trap generation and observation follow the same optical path. A dichroic mirror is used to separate the trapping light from the collected atomic fluorescence.

The DMD is illuminated using a high power diode laser that is red-detuned by 5 nm from the main cycling transition in the D2 line of Rubidium 87. We choose

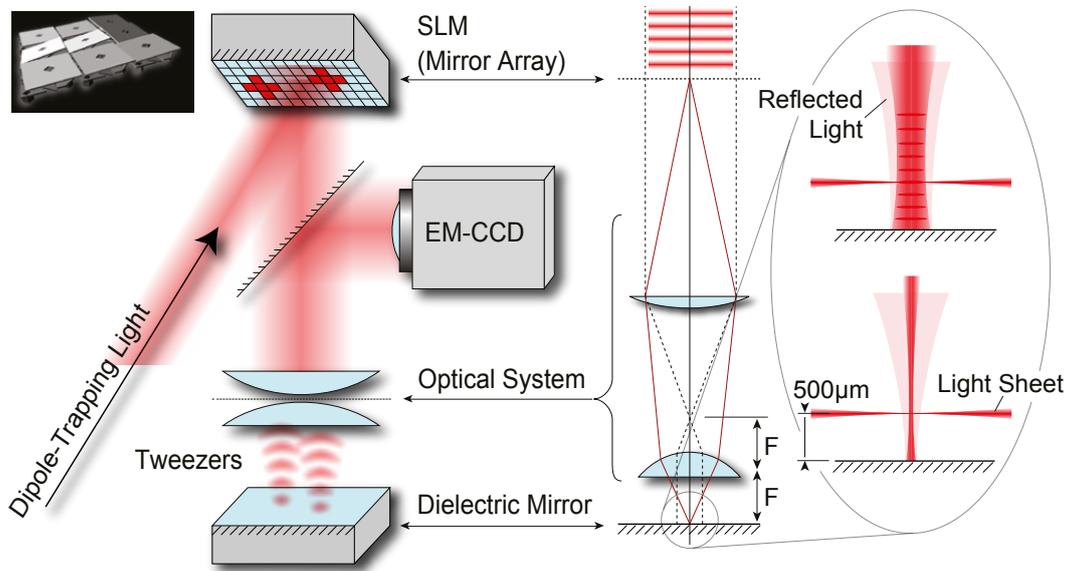


Figure 4.1 Diagram of the setup. Dipole trapping light illuminates the surface of the SLM, which is made up of ~ 1 million independently addressable micro-mirrors. A two lens microscope system images the surface of the SLM onto a cloud of cold atoms lying $500\mu\text{m}$ above the surface of a mirror. The resulting intensity distribution forms a dipole trapping potential in which these atoms are trapped. 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. A thin sheet of resonant light is used to illuminate only the atoms trapped in the image plane of the SLM, and a highly sensitive EMCCD camera is used to observe the atomic fluorescence via a dichroic mirror.

this relatively small detuning to achieve a reasonable trap depth at light intensities remaining below the damage threshold of the DMD. The convergence of the illuminating beam and the lens distances are chosen such that isoplanatism is ensured (see Chapter 6), i.e. the wavefront in the image plane is identical to the wavefront of the light in the plane of the DMD. Ideally, the illuminating beam would have a flat-top profile that homogeneously fills the area of interest on the DMD; in the experiment, a sub-optimal laser profile leads to a loss in laser power, and thus a somewhat lower trap depth.

To detect the trapped atoms, they are illuminated by a very thin sheet of resonant light which is parallel to the mirror surface and causes atoms trapped in the image plane of the tweezers to fluoresce while atoms above or below remain in the dark. The light sheet is retro-reflected to avoid pushing the atoms out of the traps. It is slightly blue-detuned from the main cycling transition in order to account for the dynamic Stark shift induced by the tweezers. As mentioned before, the fluorescence light from the atoms is separated from the reflected trapping light using a dichroic filter and directed to a highly sensitive EMCCD 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.

4.2 Trap Depth

Figure 4.2 shows graphs of the trap depth as a function of wavelength for the two hyperfine levels of the $5^2S_{1/2}$ ground state of ^{87}Rb , obtained from a simulation. Each of the graphs shows the Stark Shift that arises for a different m_F state. These shifts were calculated specifically for the D2 line, and assumed a maximum intensity of $200\mu\text{W}/\mu\text{m}^2$. This is equivalent to $20\text{kW}/\text{cm}^2$, which is just below the damage threshold of the DMD, and given our magnification, corresponds to 1W of power filling between 1/3 and 1/2 of the DMD. As can be seen from the graphs, the trap depth depends not only on the total power but also on the polarisation of the trapping light and the initial state of the atoms, and can even represent an anti-trapping

situation.

In the experiment, the polarisation of the trapping laser is linear, but acquires some birefringence due to the DMD and other optical elements. It is adjusted by rotating a $\lambda/2$ plate in front of the DMD until the fluorescence from the trapped atoms is maximised. If this waveplate is rotated with the MOST running, the trap is seen to have a “destructive” effect on the MOST at certain points, presumably due to the non-trivial way in which the trapping light shifts all of the m_F states.

The actual value for the trap depth in the experiment is also much lower than that predicted in the simulation, as a lot of power is lost due to spatial filtering of the trapping laser’s sub-optimal profile. The trap depth was re-calculated using a measurement of the average intensity and waist of the beam incident on the DMD, as well as the magnification of the imaging system. At the time most of the results in Chapters 7 and 8 were taken, this meant approximately 300 mW of power illuminating an area about 2/3 the size of the DMD. This translated to an intensity in the image plane of $3 \times 10^7 \text{W/m}^2$, which corresponds to a maximum potential depth of $k_B \cdot 100 \mu\text{K}$. The upper limit on the trap depth should thus be $\approx k_B \cdot 400 \mu\text{K}$, for a standing wave with 100% contrast, but none of our traps reached this regime. A fully experimental determination of the trap depth is given in Chapter 7.

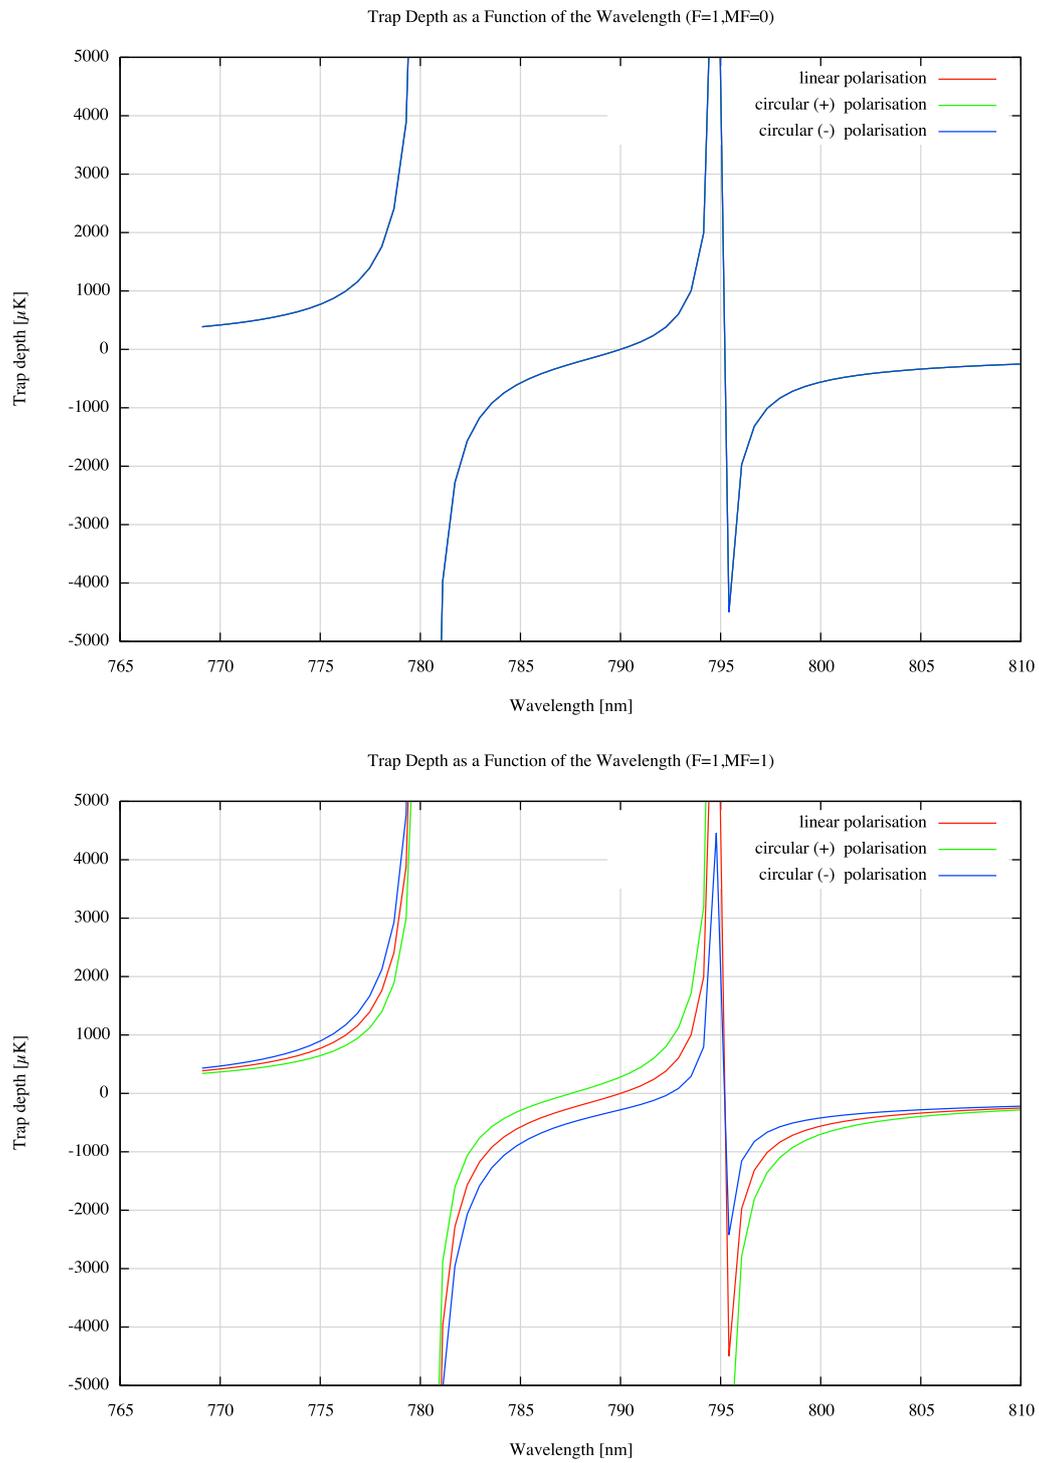


Figure 4.2 Trap Depth as a Function of Wavelength for the F=1 ground state of Rubidium 87. The for the MF=0 (top), and MF=1(bottom) states.

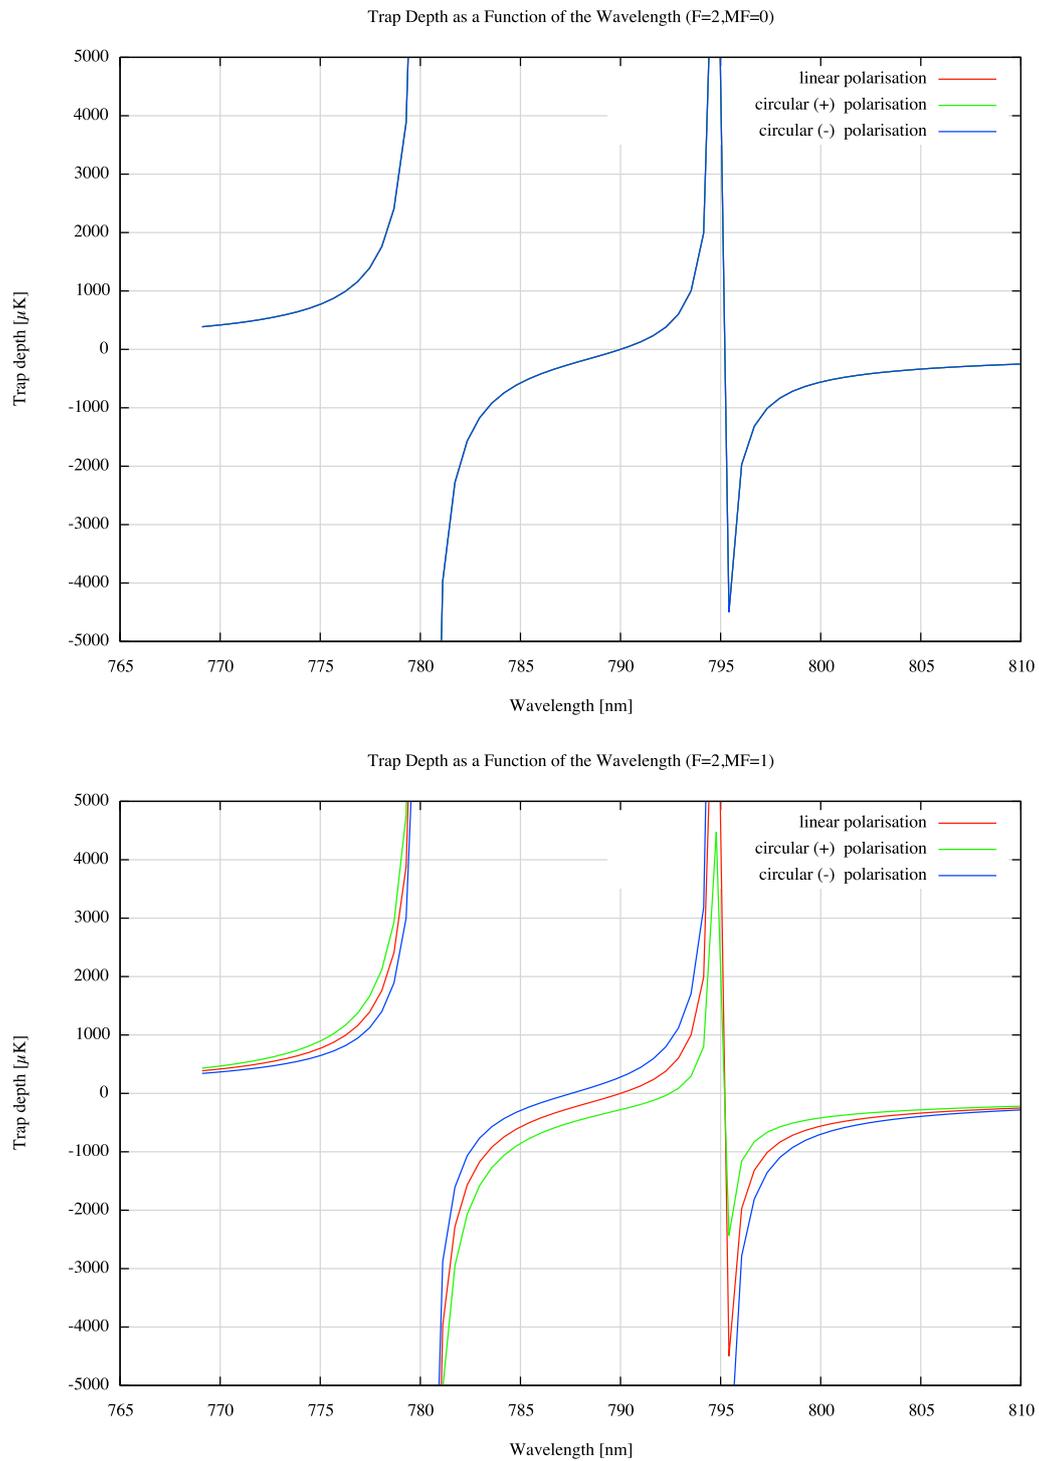


Figure 4.2 Trap Depth as a Function of Wavelength for the F=2 ground state of Rubidium 87, for the MF=0 (top), and the MF=1 (bottom) states.

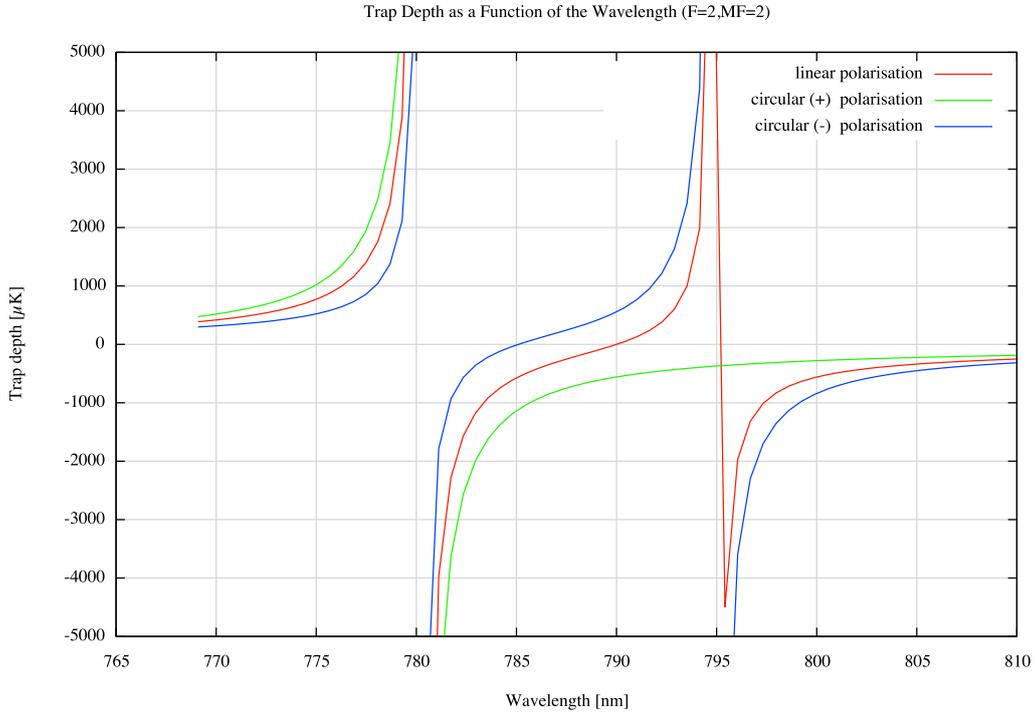


Figure 4.2 Trap Depth as a Function of Wavelength for the F=2 ground state of Rubidium 87, for the MF=2 state.

4.3 Radial and Longitudinal Confinement

In a dipole trap, atoms are more tightly confined in the radial direction than the axial direction. In order to increase confinement in the axial direction, a standing wave can be used. Atoms could then be trapped in the antinodes of the standing wave, leading to a higher net confinement. This was the underlying motivation for the mirror in our setup. Figure 4.3 shows how such a mirror gives rise to a standing wave. When the focus of the beam lies directly on the mirror, a perfect standing wave results, but atoms cannot be trapped in the waist of the beam. When the focus of the trapping beam lies some distance above the mirror, the intensity modulation

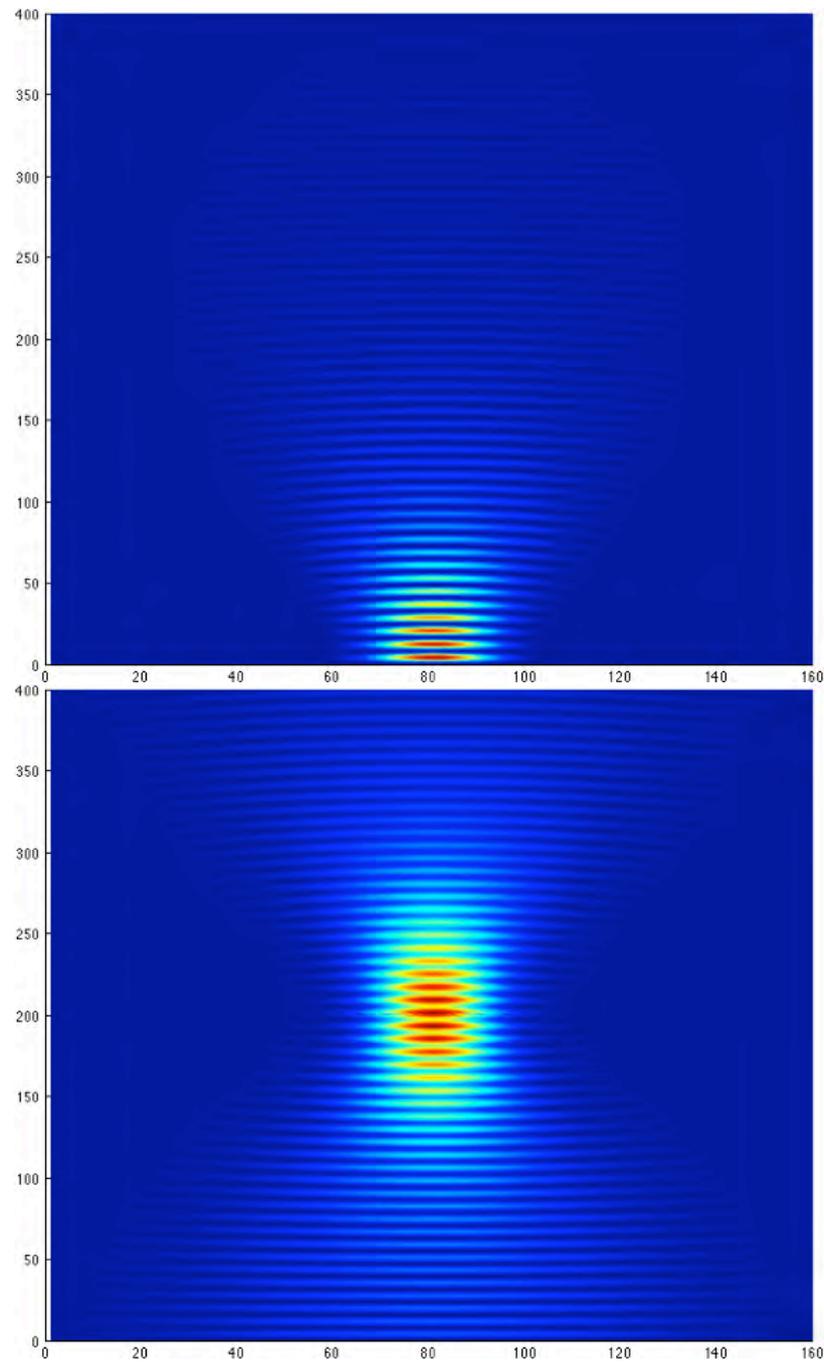


Figure 4.3 A mirror can be used to provide additional confinement in the axial direction of a dipole trapping beam, as atoms can be trapped in the antinodes of the standing wave that it gives rise to. The top half of the figure shows the perfect standing wave that results when the mirror is placed directly in the focus of the beam. This provides the maximum confinement in the axial direction, but atoms cannot be trapped in the waist of the beam. The bottom half of the figure shows the partial standing wave that results when the mirror is placed some distance above the surface of the mirror. The intensity modulation is not as pronounced but still present, and atoms can be trapped in the focus of the beam.

is not as pronounced (i.e. the contrast between the nodes and antinodes is not at a maximum), but atoms can be trapped in the focus of the beam and the partial standing wave that arises still provides more confinement than a simple focussed beam. Additionally, the MOST in our setup cannot function at a non-negligible distance from the mirror, so the focus of the trapping beam cannot lie directly on the mirror, or at least not when loading atoms from the MOST into the tweezers.

For the larger traps in the experiment, the divergence of the trapping beam is small and interference between the incoming and reflected light leads to a partial standing wave whose intensity modulation is pronounced enough to provide the desired longitudinal confinement. In this case, atoms that are cold enough are indeed confined to anti-nodal planes stacked along the optical axis. However, this effect only partially applies to smaller traps (of radii $\lesssim 20 \mu\text{m}$) because the reflected beam diverges too quickly. In this case, the longitudinal confinement is negligible.

To estimate the relative magnitude of the radial and longitudinal confinement, it is instructive to look at the oscillation frequency of the atoms in the trap. The trapping potential that arises due to a simple focussed Gaussian beam is given by

$$U = U_0 \frac{1}{1 + (\frac{z}{z_R})^2} \exp\left(\frac{-2r^2}{w_0^2(1 + (\frac{z}{z_R})^2)}\right) \quad (4.2)$$

where $z_R = \frac{\pi w_0^2}{\lambda}$ is the Rayleigh length, w_0 is the waist at the focus, and U_0 is the maximum trap depth. Around the focal point, this potential can be approximated as

a harmonic potential. This can be seen from the Taylor expansion of the expression in equation 4.2 at $z = 0$, which reads:

$$U \approx U_0 - 2U_0 \frac{r^2}{w_0^2} \quad (4.3)$$

The radial oscillation frequency of the atoms in the trap is thus given by

$$\omega_r = \sqrt{\frac{-4U_0}{mw_0^2}}. \quad (4.4)$$

The longitudinal oscillation frequency can be found in the same way. It is inversely proportional to the Rayleigh length, and is given by

$$\omega_z = \sqrt{\frac{-2U_0}{mz_R^2}} = \sqrt{\frac{-2U_0\lambda^2}{m\pi^2w_0^4}}. \quad (4.5)$$

As can be seen from equations 4.4 and 4.5, the radial trapping frequency is always higher than the axial trapping frequency, and the latter is only really significant for small waists, which have a short Rayleigh length. Even a relatively small trap with a waist of $5\mu m$ already has a Rayleigh length of $100\mu m$, and a corresponding longitudinal confinement which is only .025 as strong as the radial confinement. For a trap depth of $k_B \cdot 100\mu K$, this means a longitudinal oscillation frequency of about 1KHz versus a radial oscillation frequency of about 40KHz. For a trap with a waist of $20\mu m$, this difference increases to 80Hz versus 20KHz, which means that the longitudinal confinement is only .00625 as strong as the radial confinement, and is essentially negligible.

This clearly illustrates the benefit of a standing wave. In a standing wave, the trap depth is 4 times as large (the intensity is the square of the electric field value, which is doubled) and the longitudinal oscillation frequency is even higher than the radial trapping frequency in the antinodes of the standing wave. Of course that the confinement is not as strong in a partial standing wave, but both trapping frequencies will nevertheless be higher than in the case of a simple focussed beam. In the experiment, the longitudinal confinement due to the standing wave can be seen only for traps with radii $\gtrsim 20 \mu\text{m}$.

4.4 Single Atoms

Ultimately, the goal of our experiment is to use reconfigurable optical tweezers to dynamically control the motional degrees of freedom of *single* cold atoms. It is therefore important to have an efficient and reliable way of loading single atoms into a dipole trap. If the number of atoms loaded into a trap is Poisson distributed, then it is necessary to post-select the traps with single atoms: this is possible, but inefficient. Fortunately, in tightly focussed dipole traps, it is possible to circumvent this because of the *collisional blockade* effect that arises. This effect, which has been studied extensively by Grangier et al.^[38], occurs when light-assisted collisions between atoms causes pairs of atoms to be expelled from the trap. This takes place when the two-body loss rate, which is inversely proportional to the trapping volume, is higher than the loading rate in a trap, and occurs only in the presence of a resonant

beam. The mechanism locks the average atom number to a value of 0.5, i.e., the number of atoms is always either one or zero. The reliability of this mechanism has been shown with single atom loading efficiencies of up to 80%^[16]. The trapping volumes and atomic temperatures necessary to make use of the collisional blockade mechanism should be attainable in our experiment.

4.5 Preliminary Tests of the Optical System

In principle, the imaging system used to generate our optical tweezers is very straightforward. It has only two components: an achromatic doublet whose purpose it is to collimate the incident light, and a high numerical aperture aspheric lens whose purpose it is to focus the light down to waist sizes on the order of $1\mu m$. This aspheric lens represents a large simplification from the complicated multiple lens objectives used previously to generate such diffraction limited traps. The idea was to place the aspheric lens inside the vacuum chamber in order to circumvent any wavefront distortions that could arise due to the viewport. By placing the aspheric lens inside the chamber, we would lose the ability to adjust its distance from the MOST mirror, but we could nevertheless adjust the height of the tweezer image plane by shifting the collimating lens, which would sit outside the chamber. By mounting the aspheric lens on a very precisely machined mount, we could guarantee that it would be parallel to the mirror surface and exactly one focal length away from it.

Previous to putting the aspheric lens inside to the vacuum chamber, we decided to test the performance of the imaging system in a stand-alone mock setup. We wanted to ensure, for example, that the aspheric lens performed well off axis, so that we could move a trap from one side of the DMD to the other, and have the quality of the trap generated in the image plane remain the same. From this test setup, we learned that the imaging was very sensitive to very small tilts and to the quality of the aspheric lens itself. We also learned that significant wavefront distortion occurs when imaging with coherent light, and discovered that this was due to a little known effect known as non-isoplanatism (see Chapter 6).

The first test was meant to merely check the performance of the aspheric lens and ensure that our optical system was indeed diffraction limited. It consisted of simply focusing down a laser beam onto a $1\mu m$ pinhole, and propagating through the optical system in reverse, as shown in Figure 4.4. Propagating through the optical system in reverse is in principle no different from propagating through it in the forward direction, and had the advantage that with a single beam, alignment was easier.

The images gathered by the CCD camera were read into Matlab and analysed using a fit routine which fit them (in a mean-square sense) to a model consisting of a two-dimensional Gaussian plus background noise. From these, the value of the waist along the major and minor axis of the spot was determined, as well as the maximum intensity. The off-axis performance of the lens was tested by moving the

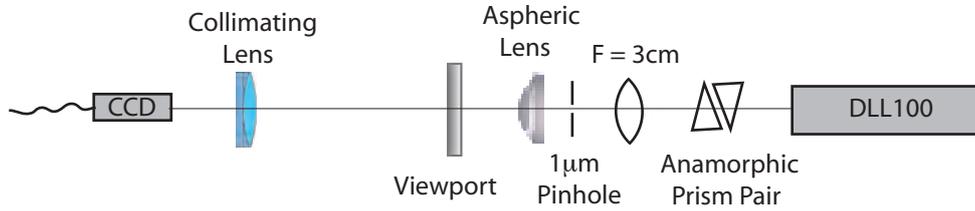


Figure 4.4 Testing with a $1\mu\text{m}$ aperture. The anamorphic prism pair shown was used to make the elliptical beam profile of the laser circular. A viewport identical in composition and coating to the one on the chamber was used to investigate any aberrations arising from its presence in the optical beam path.

$1\mu\text{m}$ pinhole incrementally along x and y respectively, and plotting the variation in the waist and intensity. Although all of the lenses we tested were commercially moulded, designed for 780 nm light, and had their curvature optimised to reduce spherical aberrations, we discovered that there was indeed a marked variation in surface quality and overall performance, even amongst lenses of identical specifications and from the same manufacturer. In the end, we chose a 25-20 HPX-S from Asphericon, with a numerical aperture of .54 and a focal length of 20mm, which was manufactured and tested specifically for us.

The second test consisted of focusing the laser beam onto a $50\mu\text{m}$ pinhole, propagating through the optical system in the forward direction, and then using a mirror and PBS to observe the images on the CCD camera, as shown in Figure 4.5. It was meant to once again test the off-axis performance of the aspheric lens, as well as test whether small displacements of the collimating lens had a significant effect on

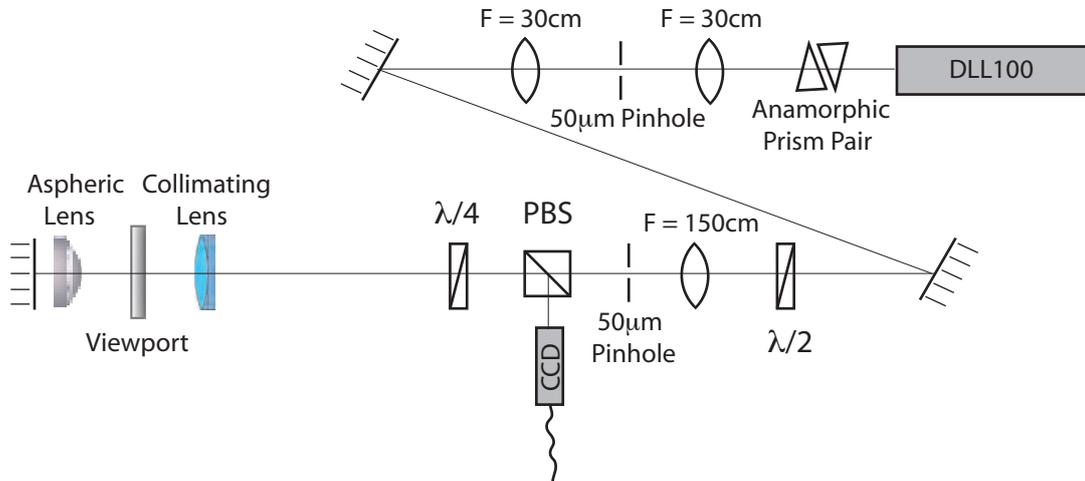


Figure 4.5 Testing with a $50\mu\text{m}$ aperture. A spatial filter was included after the anamorphic prism pair to clean up the profile of the incident beam. A $\frac{\lambda}{4}$ plate and a PBS were used to observe the images on the CCD camera.

the imaging. This test was beneficial in two ways. Firstly, it elucidated the hard parts in the alignment of the optical system to be used in the experiment, and secondly, it's double pass through the optical system served as a sort of amplifier for aberrations. The images gathered were analysed using the same fit routine as before.

One thing we learned from this test was that the collimator had to be displaced by approximately 4 cm to restore the image every time the aspheric lens was moved out of focus by $10\mu\text{m}$. This told us that moving the image plane up or down by even 50 microns inside the vacuum chamber (a distance that would be dictated by the lowest height of the MOST above the mirror surface) would mean having to move the collimating lens accurately by a distance on the order of centimetres. It also gave us an idea what kind of margin of error we had for the length of the lens

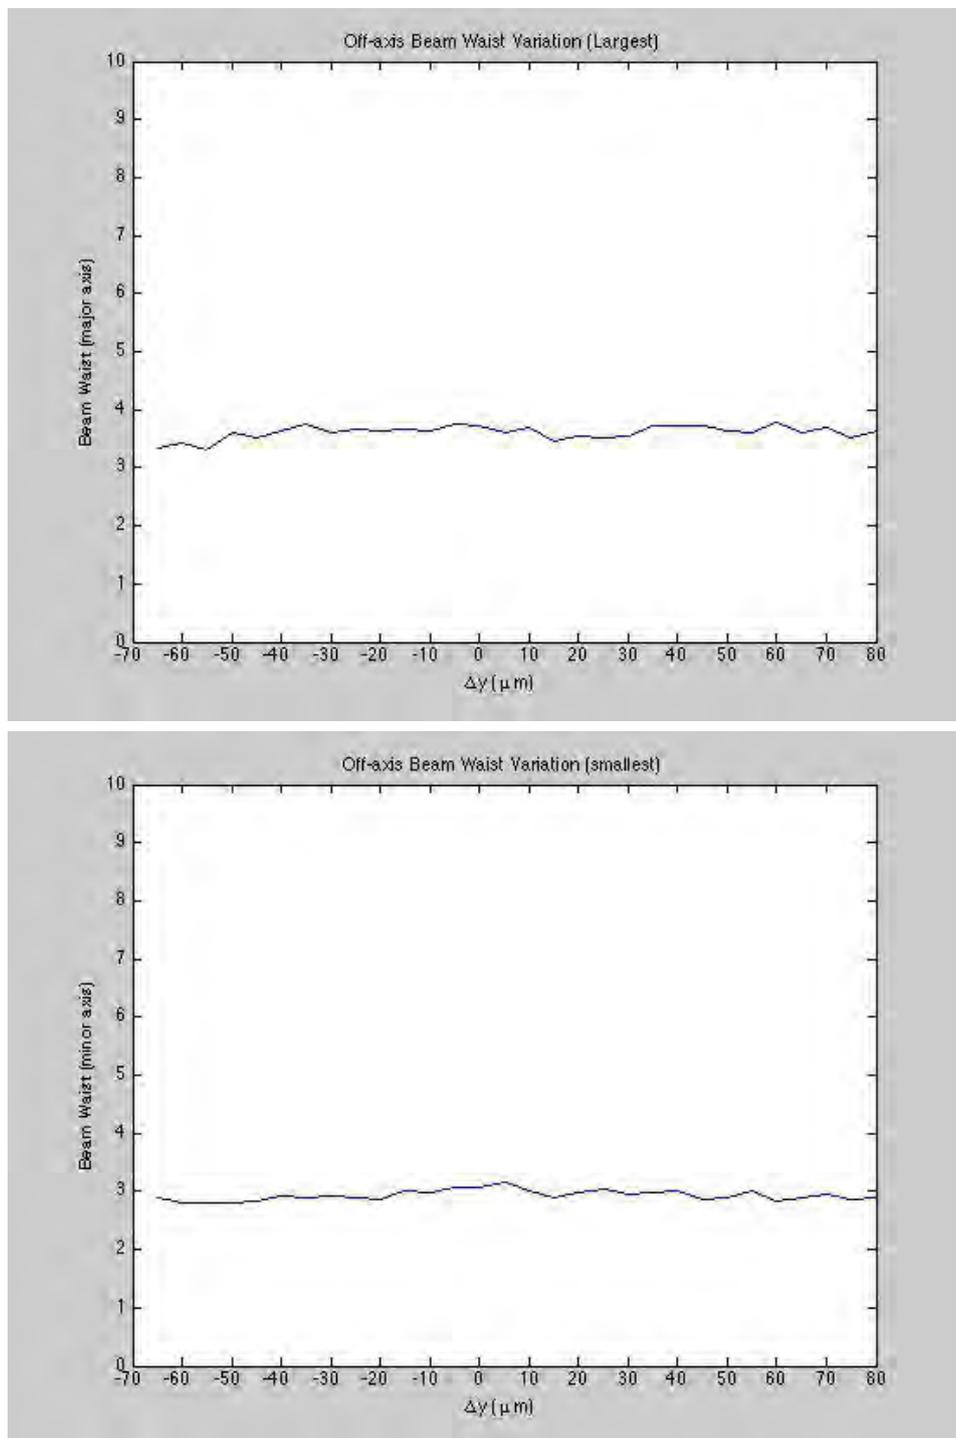


Figure 4.6 Testing the off-axis performance of the aspheric lens using a $50\mu\text{m}$ pinhole. The pinhole was displaced incrementally along the y axis with the aspheric lens $20\mu\text{m}$ out of focus, and the collimator adjusted accordingly. Graphs of the variation in maximum and minimum waist of the imaged spot are shown.

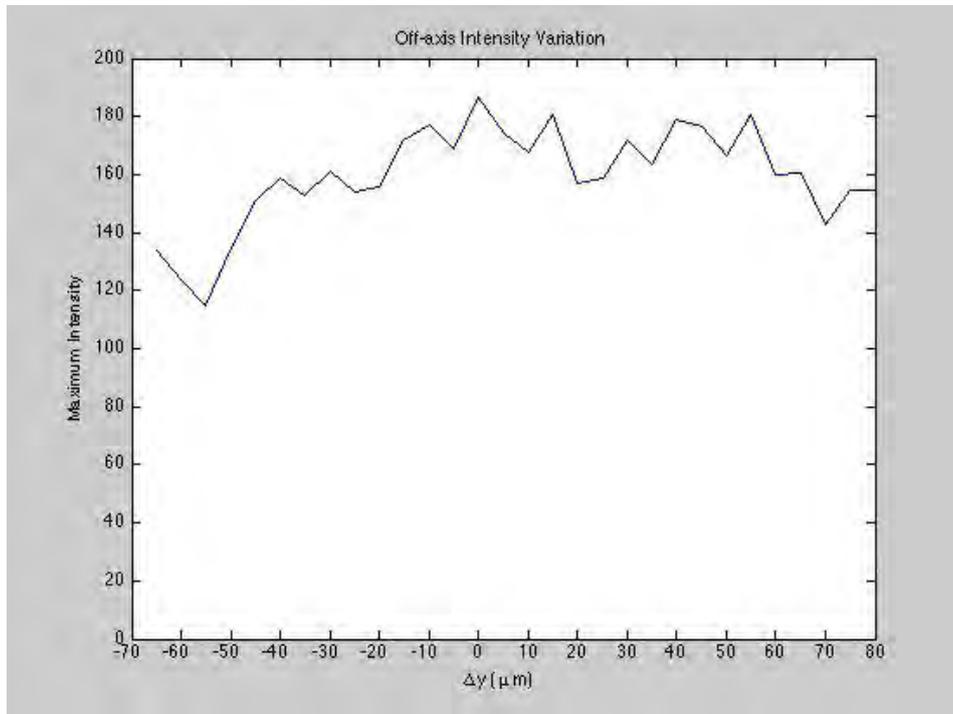


Figure 4.6 Testing the off-axis performance of the aspheric lens using a $50\mu\text{m}$ pinhole. The pinhole was displaced incrementally along the y axis with the aspheric lens $20\mu\text{m}$ out of focus, and the collimator adjusted accordingly. A graph of the variation in intensity of the imaged spot is shown.

mount that would hold the aspheric lens above the mirror inside the chamber. We tested the off-axis performance of the optical system with the aspheric lens $20\mu\text{m}$ out of focus, and the collimator displaced accordingly, and found that the waist and intensity of the image stayed fairly constant as the pinhole was displaced along the y axis. Graphs of these variations in waist and intensity can be seen in Figure 4.6.

The third test consisted of simply replacing the $50\mu\text{m}$ pinhole by the DMD itself, exactly as it would be set up in the experiment. To form a spot with a diameter of

$\approx 50\mu m$ on the DMD, a block of 4 by 4 mirrors was switched on using software that accompanied the device driver. This software allows the user to upload bitmaps and display them as static images on the DMD, but does not allow for mirror-by-mirror or line-by-line addressing. To accomplish the latter, as well as to implement more sophisticated sequences and exploit the full refresh-rate capabilities of the DMD, a customised control program had to be developed (see Chapter 5). The $50\mu m$ spot was shifted across the transverse plane of the DMD by displaying successive bitmaps, and imaged through the optical system as before. This time the aspheric lens was not displaced. As can be seen in Figure ??, the waist variation was once again minimal, although there seemed to be a slightly higher variation in the intensity.

In addition to the $50\mu m$ spot, several other bitmaps were loaded onto the DMD and imaged through the optical system, including the Oxford physics logo, and the Oxford physics ϕ shown on the left hand side of Figure 4.9. As can be seen from this figure, the resulting images showed surprisingly prominent aberrations. This was once again conjectured to be due to a fault in the surface quality of the aspheric lens. Replacing this lens with yet another lens from the same company reduced the effect, but it did not eliminate it entirely. Moreover, the effect did not seem to be present when imaging the $50\mu m$ spot or other small objects, and seemed to be worst for large bitmaps that covered the entire surface of the DMD. Coma and chromatic aberrations were ruled out, and it was speculated that the most likely cause was some form of spherical aberration present when imaging with coherent light. This

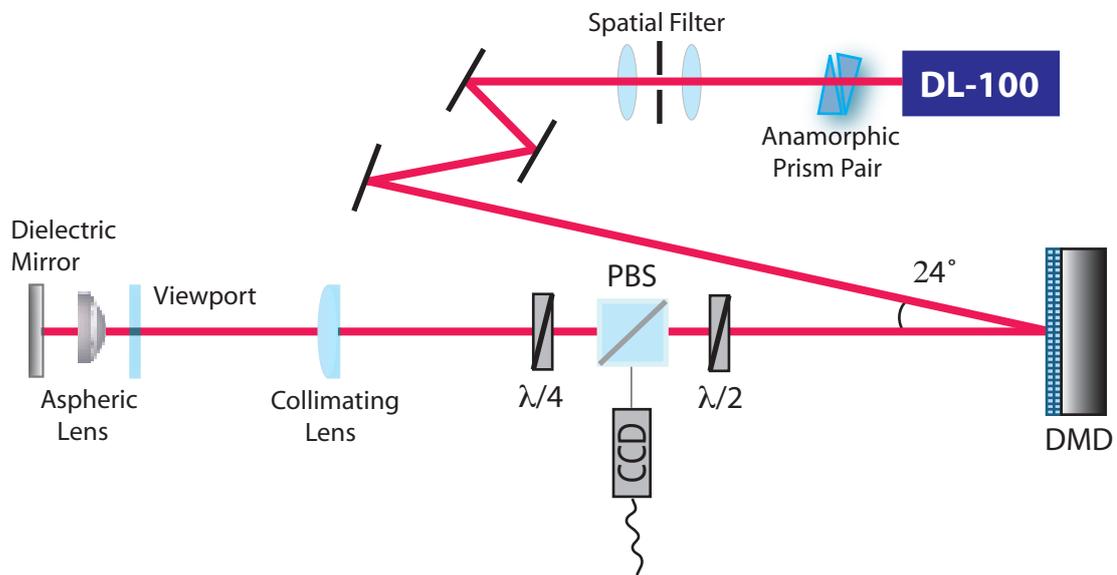


Figure 4.7 Testing with the DMD. As before, a spatial filter was included after the anamorphic prism pair to clean up the profile of the incident beam, and a PBS was used to observe the images on the CCD camera. The angle of incidence of the incoming beam must be 24° for the DMD mirrors to direct light along the optical axis in the “on” position.

eventually led us to conclude that our aberrations were due to non-isoplanatic imaging conditions. Here credit is due to Dr. Edouard Brainis, a former postdoc in our group, now at ULB, who first looked into and suggested the effect, and was the first author of our publication on the topic (see^[39] and Chapter 6).

Non-isoplanatism is an effect that can be explained as follows: a field that propagates through an imaging system will suffer from some phase distortion, even if the imaging system is diffraction limited. This is irrelevant when the light is incoherent, but highly destructive when it is coherent, as this phase distortion translates into a distortion of the field amplitude in the image plane. The effect becomes significant

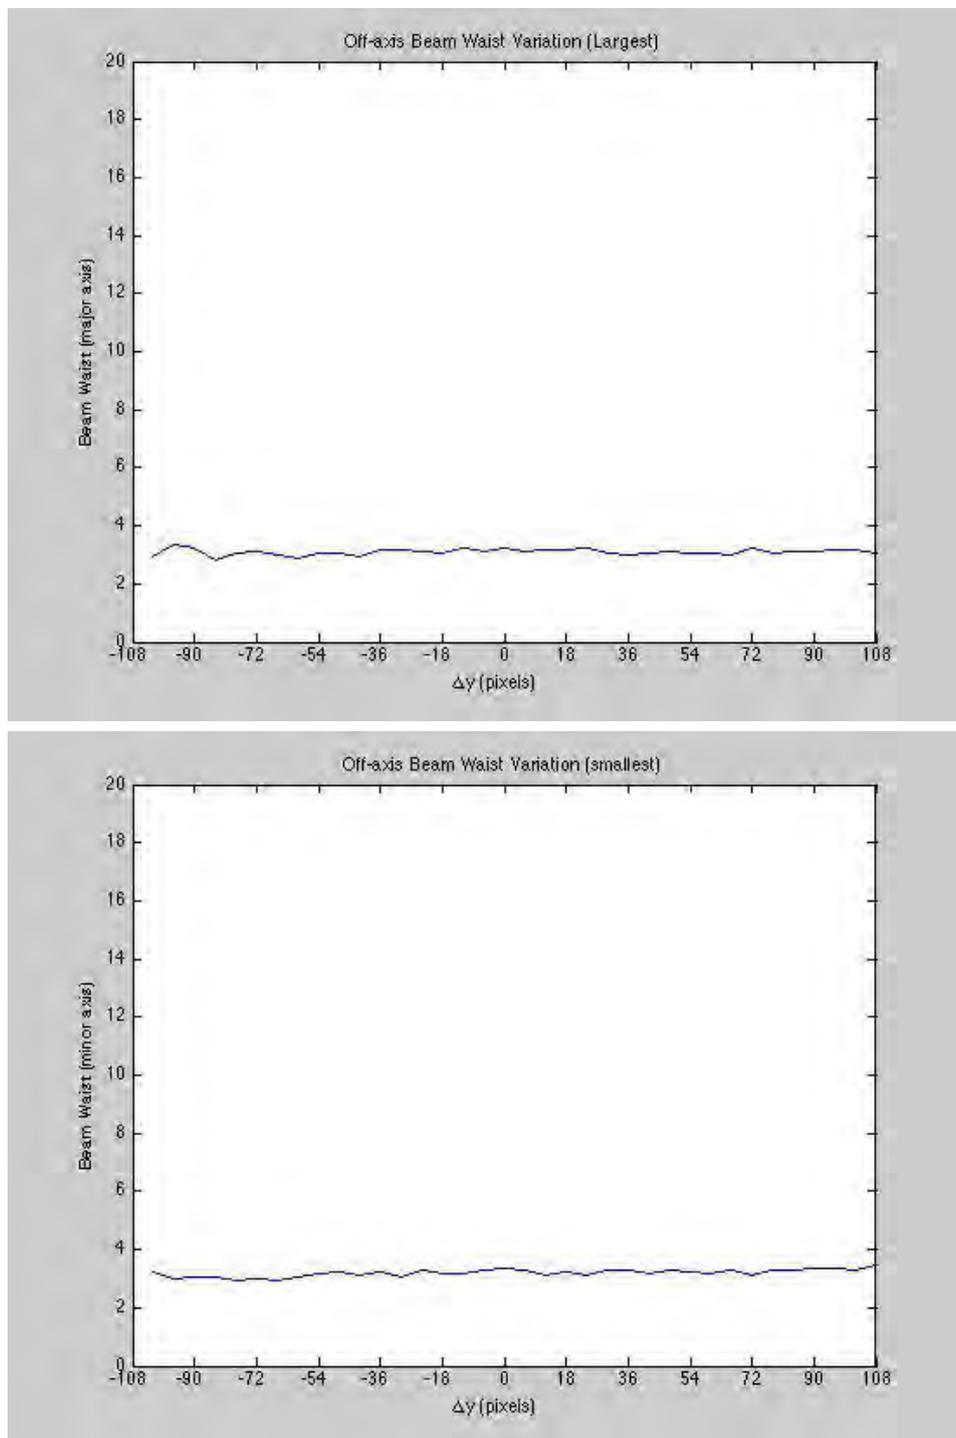


Figure 4.8 A sample of the results obtained by testing with the DMD. Here, the aspheric lens was in focus, and the aperture was displaced along the y axis. Graphs of the variation in the maximum and minimum waist of the imaged spot are shown.

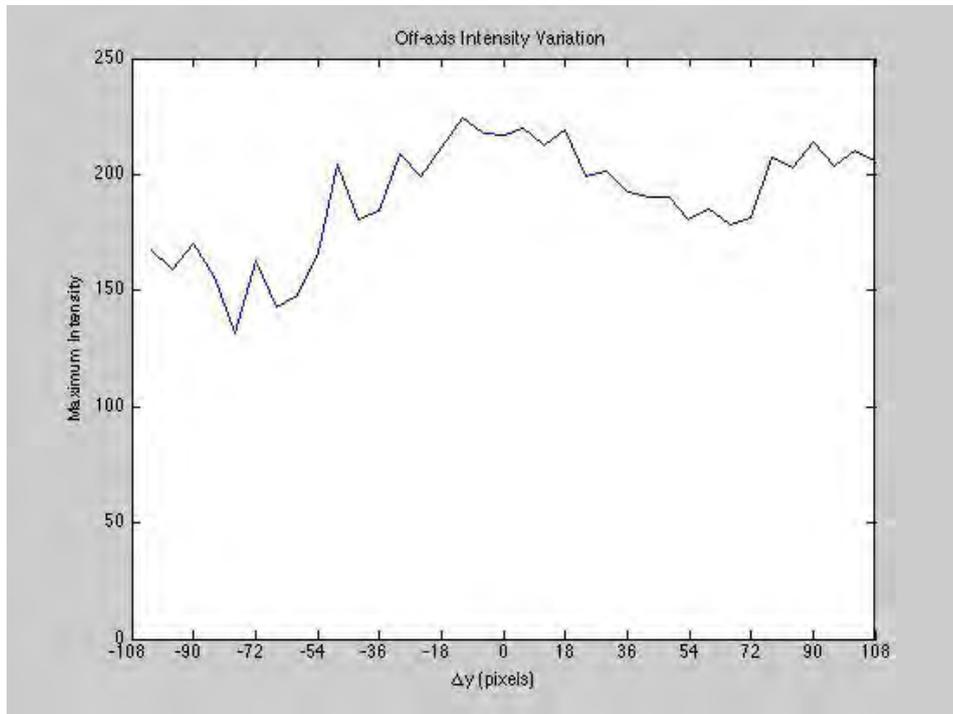


Figure 4.8 A sample of the results obtained by testing with the DMD. Here, the aspheric lens was in focus, and the aperture was displaced along the y axis. A graph of the variation in intensity of the imaged spot is shown.

when imaging extended objects or far off axis points, as we do when imaging the Oxford Physics logo. Usually, diffraction-limited imaging is performed with incoherent light, so the effect goes by unnoticed. Fortunately, it is possible to ensure that even coherent imaging will not suffer any wavefront distortion due to non-isoplanatism by using spherical wave illumination. Figure 4.10 shows images of the same two bitmaps - the Oxford Physics logo and the Oxford physics ϕ - as they appear when using spherical wave illumination on the DMD. As can be seen from the figure, the aberrations disappear entirely.

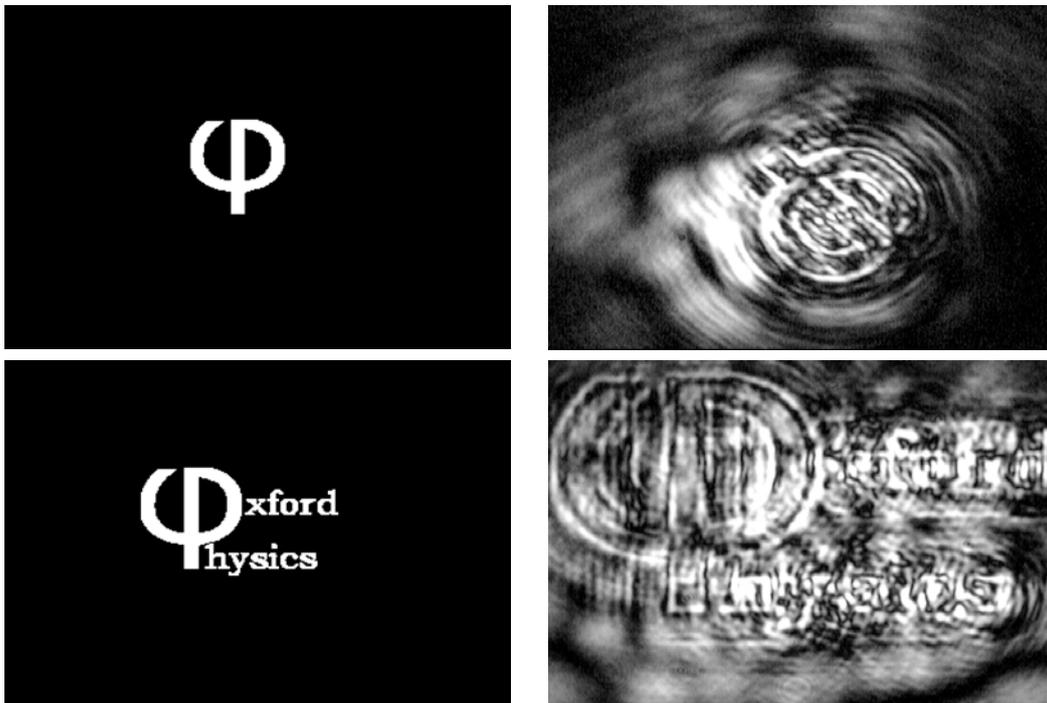


Figure 4.9 Non-Isoplanatic Imaging.

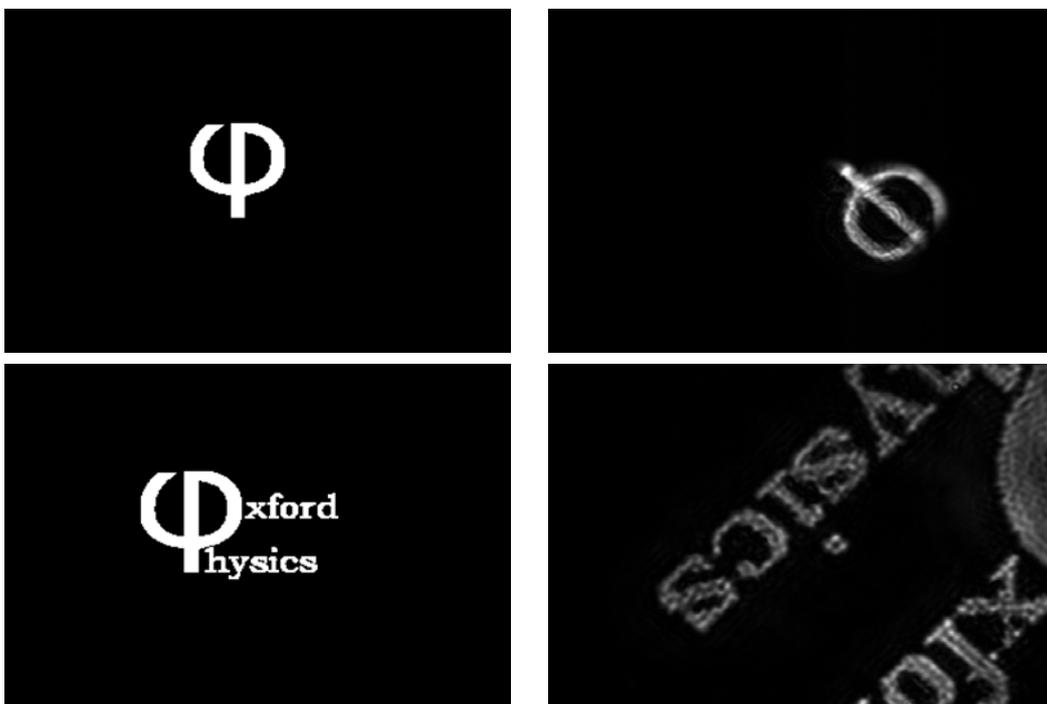


Figure 4.10 Isoplanatic Imaging.

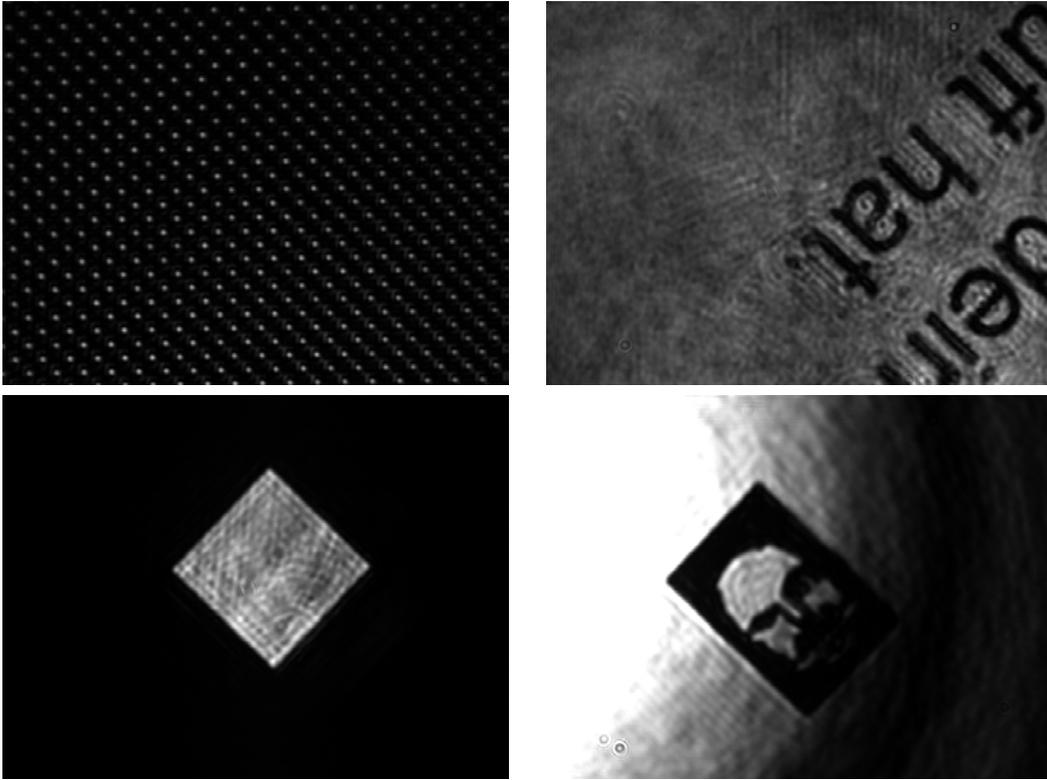


Figure 4.11 Further examples of isoplanatic imaging with the DMD: A large grid of diffraction limited spots, some writing, a large square, and a bitmap adapted from a real photo of a strange German man who kept coming into our lab everyday.

Figure 4.11 shows some further examples of isoplanatic imaging with the DMD. The images on the top right hand corner and the bottom two panels exhibit sharp edges and no aberrations, even though they are all extended objects that cover either all or most of the DMD. The picture on the bottom right hand corner was loaded close to the edge of the DMD, which is why it does not appear fully illuminated: in the test setup, the illuminating beam had a Gaussian profile that just filled the surface of the DMD. The top left hand corner shows a large grid of diffraction limited spots with a waist of $1.2\mu\text{m}$ and a separation of $3\mu\text{m}$. Figure 4.13 shows a false

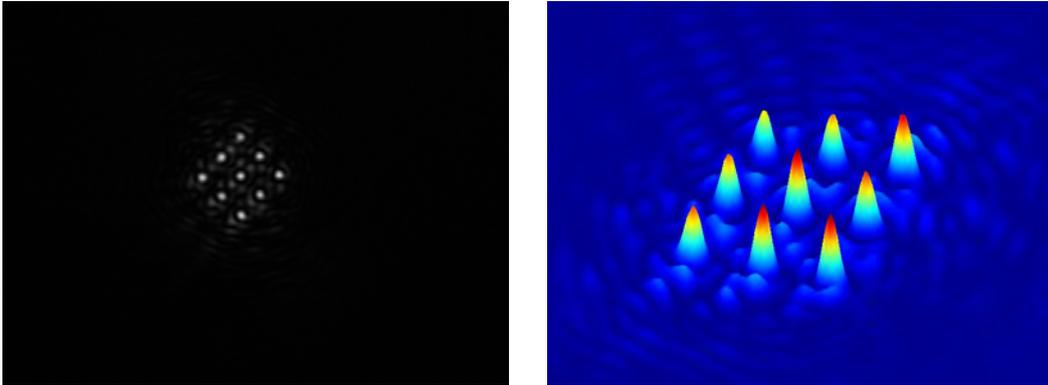


Figure 4.12 Grid of 9 spots with a diameter of $1.2\mu m$ and a separation of $3\mu m$.

colour rotated version of this grid and shows a cut across one of these spots. This cut was obtained from the same fit routine as before (it was adapted to zoom into each spot to find the waist and intensity of each). The waist of these fits gives the size of the spot in the image plane, and in this case, gives the value for the smallest spot we were able to image in the test setup, and this the diffraction limit of our optical system.

Figure 4.12 shows a smaller grid of diffraction limited spots with the same $1.2\mu m$ waist and $3\mu m$ separation as before. The modulation in the background is due to a slight misalignment and is reflected in the 3D fit next to it. This fit clearly shows that the maximum intensity in each spot is not exactly the same. This is partly due to alignment, but also reflects the quality of the laser profile, and means that the depth of each trap would not be exactly the same in the real setup. This should not be a problem if the goal is to trap single atoms in such a grid of tightly focussed

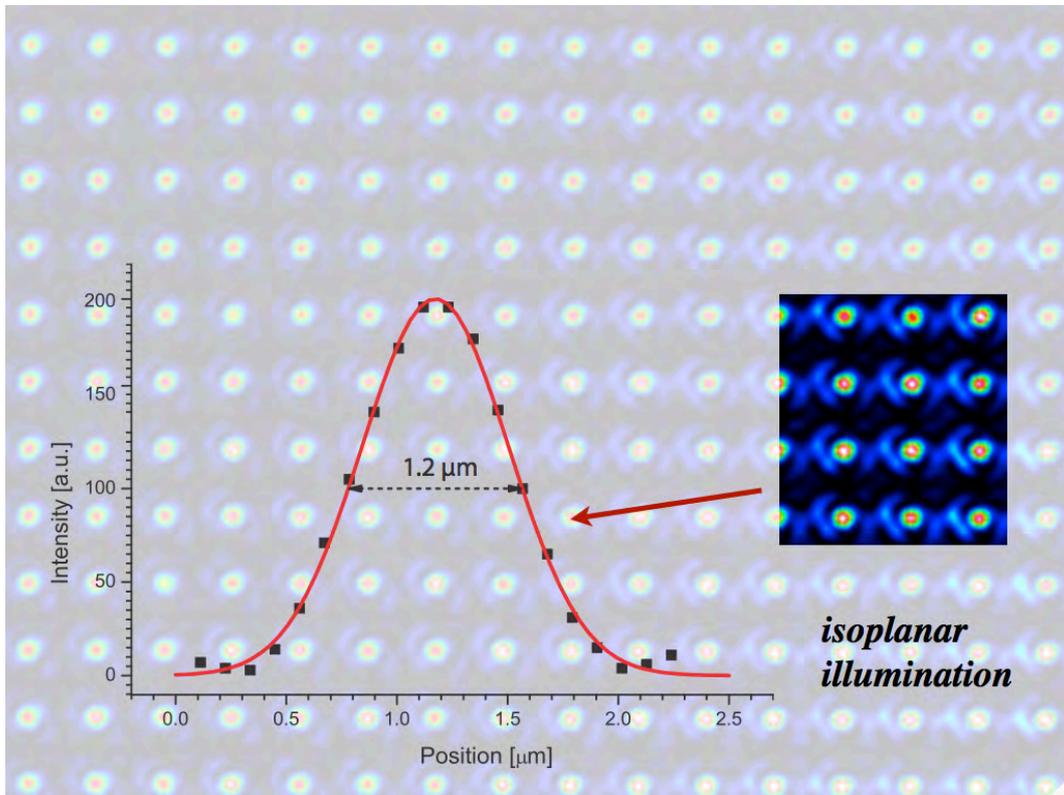


Figure 4.13 False colour image of a large grid of diffraction limited spots.

traps, because the collisional blockade mechanism^[38] would guarantee only one atom per site. However, it could pose problems for bigger trapping geometries in which regions of higher intensity would lead to a higher atom number, and hence a trap that is not filled homogeneously.

4.6 Feasibility Study

Previous to obtaining results, we published a proposal paper presenting the basic outline and feasibility of our experiment^[40]. This paper gives a thorough and de-

tailed analysis of our goals from a theoretical point of view, and contains several instructive calculations that further elucidate the topics introduced in this chapter, so the most relevant sections have been included below. The first author of this paper was Lukas Brandt.

4.6.1 Dipole trapping of Alkali Atoms

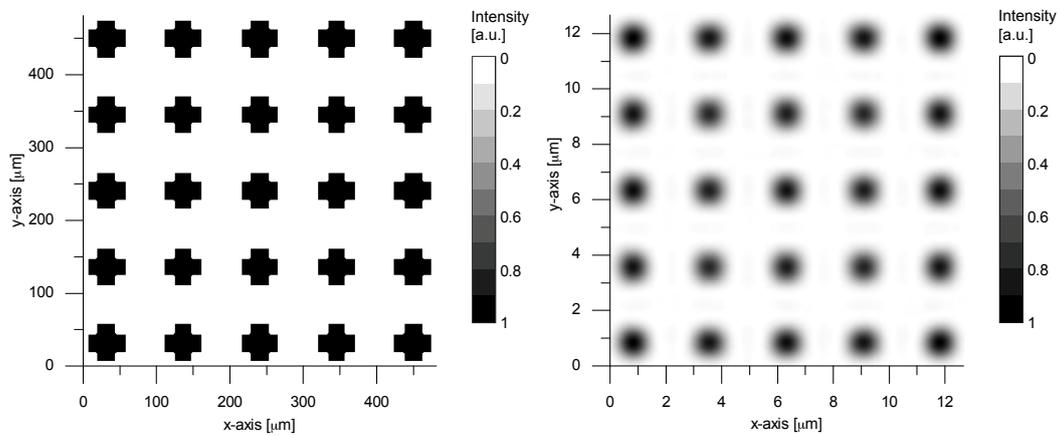


Figure 4.14 The picture on the left shows a pattern on the SLM. The pattern is a repetition of a 4 by 4 mirror block where the corner pixels are missing. Each mirror has a dimension of $14 \mu\text{m}$ by $14 \mu\text{m}$. The picture on the right shows the resulting image when a diffraction-limited lens system with a numerical aperture of $NA = 0.5$ is used to demagnify the pattern by a factor of $1/38$. This picture is obtained by convoluting the demagnified pattern from above with airy spots of $1 \mu\text{m}$ waist.

The majority of experiments cooling and trapping neutral atoms are done with alkali atoms, of which Rubidium is one of the most prominent species. Hence we choose it for the purpose of this feasibility study. The optical tweezers consist of tiny optical dipole-force traps, in which the dynamic Stark shift gives rise to a trapping

potential^[35]

$$U_{dip}(\mathbf{r}) \approx \frac{\hbar\Gamma^2}{8\delta} \frac{I(\mathbf{r})}{I_{sat}} \quad (4.6)$$

in the *far* detuned limit, where δ is the detuning of the trapping laser light with respect to the atomic transition, $I(\mathbf{r})$ is the intensity of the trapping light, I_{sat} is the saturation intensity for the chosen transition, and Γ is the corresponding decay rate. For red detuned light, $\delta < 0$, U_{dip} is negative and hence attractive. The most critical parameter is the trap depth U_0 , which corresponds to the highest intensity of the trapping laser light I_0 . Hence U_0 is the energy required for an atom at rest to escape the trap, neglecting gravity. Obviously, the trap depth has to be at least as large as the energy of the pre-cooled atoms. Therefore, the minimum usable trap depth is determined by the temperature that can be reached by the magneto-optical trap. For rubidium, the Doppler cooling limit would be $T_D = 143 \mu\text{K}$. However, it is desirable to operate deeper traps since the spatial confinement of the atoms increases with depth. For atoms whose kinetic energy is smaller than the potential depth by a sufficient amount, the trap can be treated as harmonic, and the spatial confinement of the atoms can be determined by the harmonic oscillation frequencies ω_r and ω_z for radial and longitudinal motion, respectively. The achievable values of U_0 , ω_r , and ω_z depend on the wavelength λ of the trapping laser, its intensity, and the trap geometry.

In the proposed scheme, the trap geometry is essentially determined by the distance L between the focal plane of the optical system which generates the tweezers

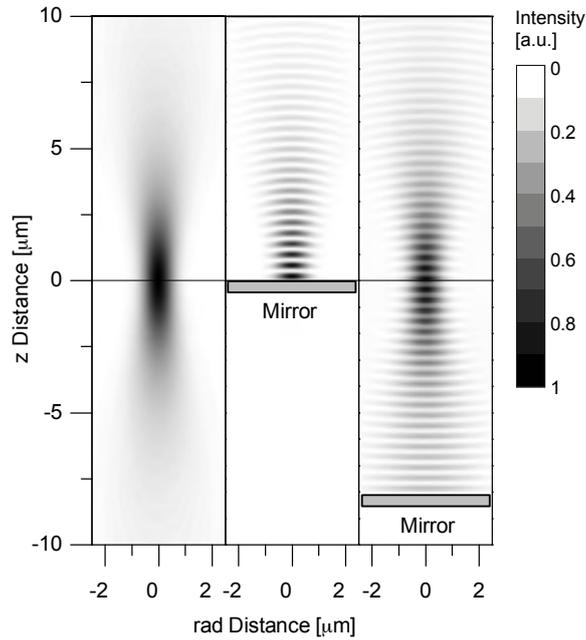


Figure 4.15 Intensity distribution of a focused laser beam (with beam waist $w_0 = 0.78 \mu\text{m}$) in the neighbourhood of the focal plane for different distances L between the dielectric mirror and the focal plane. (Black corresponds to the highest intensities.) The left picture is for $L \rightarrow \infty$, i.e. no mirror being present. The tweezer is a pure forward traveling wave with no longitudinal intensity modulation. The middle picture shows a beam which is reflected in its focal plane ($L = 0$). The tweezer has a perfect standing wave pattern, but atoms cannot be trapped in the beam waist. The right picture shows the intensity of a retro-reflected beam whose focus lies $L = 8 \mu\text{m}$ above the mirror surface. Here the intensity modulation is notably less than in the previous case.

and the dielectric mirror. The intensity modulation induced by the interference of forward and backward travelling waves over this distance can be seen in Fig. 4.15. In the following, we first discuss the limiting case $L \rightarrow \infty$ (left side of Fig. 4.15), which corresponds to having no mirror at all, and then analyse the longitudinal intensity modulation for finite L .

When the dielectric mirror is not present ($L \rightarrow \infty$), the trapping potential has

a single minimum at the beam waist, in the focal plane of the lens system. The maximum achievable intensity at that point is

$$I_0 = I_d \cdot M^2, \quad (4.7)$$

where M is the magnification of the microscope (resulting in a demagnification of the DMD by $1/M$) and $I_d = 10 \text{ W/cm}^2$ is the damage threshold of the DMD. To maximise the spatial control of the atoms, the individual traps should be as small as possible. The smallest achievable trap size is limited by the resolution of the optical system (i.e. the Rayleigh criterion) to

$$D_{min} = 1.22 \cdot \lambda N. \quad (4.8)$$

Here, λ is the wavelength of the light used, and N is the f-number related to the numerical aperture NA of the optical system by $N \approx \frac{1}{2NA}$. If an aspherical singlet lens is used, the NA can be as high as 0.5. For this value, a resolution limit of $D_{min} \approx 1 \mu\text{m}$ can be reached for $\lambda = 785 \text{ nm}$, which is well-suited for trapping Rb.

4.6.2 The role of the SLM

It is important to bear in mind that any experiments involving dynamics require the reconfiguration of the trapping potential. One way atomic transport could be realised is by displacing a pattern of traps. The switching of mirrors would lead to

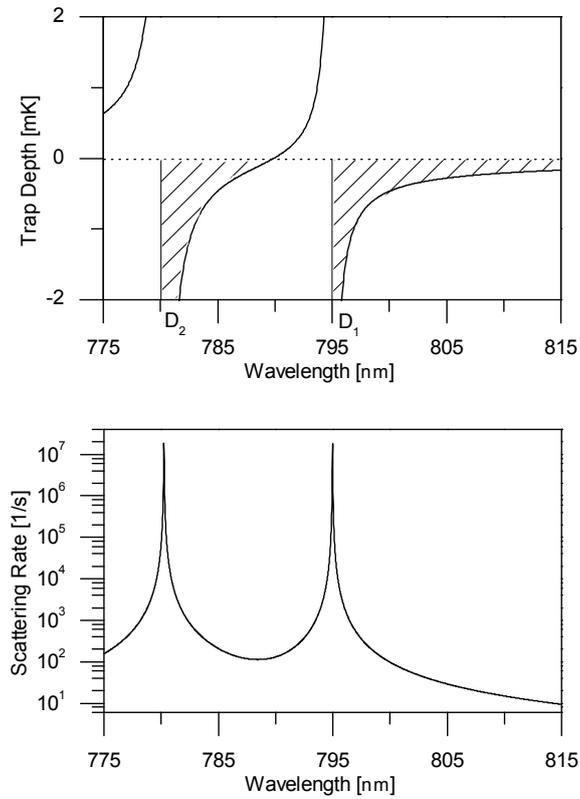


Figure 4.16 The top graph shows the potential depth of the trap with respect to the wavelength of the trapping light, for linearly polarised light with an intensity of $I_0 = 14.4 \text{ kW/cm}^2$. The trapping regions are hatched. The bottom graph shows the corresponding scattering rate.

abrupt changes in the trapping potential if the contribution of light for a point in the image plane of the DMD is dominated by one mirror only. If, however, the optical system does not resolve individual mirrors, the reconfiguration of the traps can be smoother, since the light intensity for every point will have contributions from a number of mirrors. Effectively, this leads to grey scales in the trapping pattern. To take advantage of this, we choose the magnification M such that the diagonal of a 2×2 mirror block is not resolved. As the micro-mirrors of a commercially available

DMD have a size of $14\ \mu\text{m} \times 14\ \mu\text{m}$, a demagnification of $1/38$ is required, i.e. the microscopic magnification should be $M = 38$ (see Fig. 4.14). According to Eq. (4.7), if the DMD is illuminated with light of an intensity at its damage threshold, an intensity I_0 of $14.4\ \text{kW}/\text{cm}^2$ can be obtained at the center of a tweezer. To ensure a high loading efficiency from the MOT into the tweezers, linear polarised trapping light has to be used^[41]. This would saturate the D2 line of Rubidium at an intensity of $I_{sat} = 2.5\ \text{mW}/\text{cm}^2$. Therefore the maximum intensity is $5.8 \cdot 10^6$ times the saturation intensity. Taking the above discussion into account and requiring a trap depth of 1 mK leads to a detuning $\delta \leq 2.1 \cdot 10^5 \Gamma$, according to Eq. (4.6). This corresponds to a wavelength detuning of about 2.6 nm to the red with respect to the D₂ line. However, since Rubidium has a fine structure splitting of 15 nm, this estimation of the required detuning is certainly too crude. The hyperfine transitions within the D₁ and D₂ lines of Rb need to be taken into account. Two possible wavelength regimes present themselves. The first is slightly red detuned from the D₂ line at 780 nm, and the second is red detuned from the D₁ line at 795 nm (top graph of Fig. 4.16). More specifically, for a trapping laser with an intensity of $I_0 = 14.4\ \text{kW}/\text{cm}^2$, for instance, a trap depth of 1 mK can be achieved for the wavelengths of $\lambda_2 = 782.85\ \text{nm}$, and $\lambda_1 = 796.90\ \text{nm}$.

Working with light so close to the D₁ and D₂ resonances gives rise to atomic heating due to photon scattering from the trapping beam. The heating is proportional

to the scattering rate R_{scat} , which, in the far detuned limit^[35], is given by

$$R_{scat} \approx \frac{\Gamma^3}{8\delta^2} \frac{I}{I_{sat}}. \quad (4.9)$$

Here we have to take the scattering from both the D1 and the D2 line into account, as shown in the bottom graph of Fig. 4.16. For both λ_1 and λ_2 , Rb scatters about 600 photons per second. The scattering rate in both detunings is about the same, since in both cases, the detuning to the closest transition is small when compared to detuning from the other transition. If the trapping laser were further detuned, the region to the red from the D₁ line would be preferable, as the potentials of both lines add constructively there, whereas in the region in-between the D₁ and D₂ lines, the potentials add destructively.

Given that each scattering event involves an exchange of two photons, the kinetic energy of an atom increases on average by twice the recoil energy per scattering event. The lifetime τ of a Rubidium atom (initially at rest) in the dipole trap would therefore be limited to $\tau = 1/R_{scat} \cdot U/2E_{rec}$. This yields $\tau = 2.3$ s if photon scattering is the only source of heating. It will be shown below that this time is long enough to allow for many transport and detection cycles.

4.6.3 Radial & axial confinement

We now estimate the radial and longitudinal confinement for a travelling-wave dipole trap ($L \rightarrow \infty$). Since a 2 by 2 mirror block is spatially not resolved in the object plane of the microscope, we may safely assume that the corresponding intensity distribution is basically limited by diffraction and therefore cylindrically symmetric. The intensity of the light field in the radial direction then reads

$$I(q) = I_0 \left(\frac{2 J_1(q)}{q} \right)^2, \quad (4.10)$$

where $q = \frac{\pi r}{\lambda N}$, r is the distance from the axis of the trap, and $N \approx 1/(2 NA)$ is the f-number which is about 1 in this case. $J_1(q)$ is a Bessel function of the first kind. For small q , it expands as

$$J_1(q) = \frac{1}{2}q - \frac{1}{16}q^3 + O(q^5). \quad (4.11)$$

Inserting this into (4.10) and keeping only terms up to second order we obtain the harmonic approximation to the trapping potential in the radial direction:

$$U(r) \approx U_0 - \frac{U_0}{4} \left(\frac{\pi}{\lambda N} \right)^2 r^2. \quad (4.12)$$

Here U_0 is the trap depth, which is plotted in Fig. ?? as a function of the wavelength λ . A particle with mass m oscillates with frequency

$$\omega_r = \sqrt{\frac{-U_0\pi^2}{2m\lambda^2N^2}} \quad (4.13)$$

in this potential. For $-U_0 = 1 \text{ mK} \cdot k_B$, this leads to a radial trapping frequency of $\omega_r = 2\pi \cdot 162 \text{ kHz}$. If we approximate this as a Gaussian distribution, the beam waist would be $w_0 = 2\sqrt{2}\lambda N/\pi = 0.71 \text{ } \mu\text{m}$, which is in agreement with the expected resolution limit. The oscillation frequency along the axial direction ω_z can be calculated using the Rayleigh length $z_0 = \pi w_0^2/\lambda \approx 2.0 \text{ } \mu\text{m}$:

$$\omega_z = \sqrt{\frac{-2U_0}{mz_0^2}}. \quad (4.14)$$

For $-U_0 = 1 \text{ mK} \cdot k_B$, this yields $\omega_z = 2\pi \cdot 35 \text{ kHz}$. The axial trapping frequency ω_z is significantly lower than the radial trapping frequency ω_r , which is generally the case for running wave optical dipole-force traps.

We now consider the dielectric mirror to be in the object plane of the microscope $L = 0$. The backward travelling wave interferes with the forward travelling wave and gives rise to a longitudinal interference pattern, creating an ideal standing wave (see the central picture of Fig. 4.15). The potential has many local minima in which atoms can be trapped. The intensity modulation along the z-axis gives rise to the

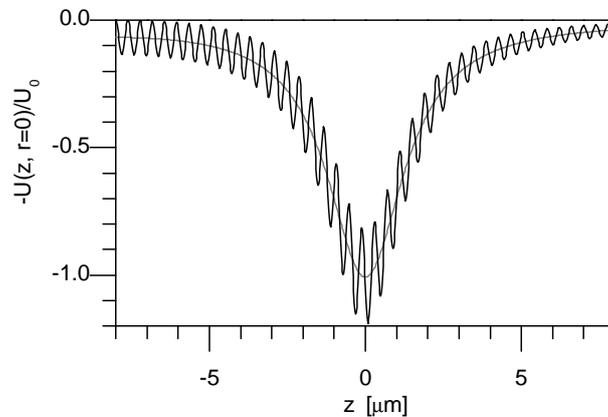


Figure 4.17 Potential depth $U(z)$ along the axis, with the focal point $8 \mu\text{m}$ away from the mirror surface. The slowly varying average of the potential (gray) is a feature of focusing the light, whereas the oscillation (black) is caused by the partial standing wave. The modulation has the highest visibility close to the mirror-surface, becomes less pronounced in the focal spot, and then increases again further away from the focal point.

longitudinal potential

$$U(z) = A(z) \sin^2 \left(\frac{2\pi}{\lambda} z \right). \quad (4.15)$$

The envelope of the modulation, $A(z) = 4U_0/(1 + z^2/z_0^2)$, reaches its maximum at the mirror surface. Here U_0 is the potential depth only due to a single focused beam. Because the contrast of the modulation is one, $A(z)$ is also the trap depth of a local minimum at height z above the mirror. The largest achievable trap depth is $4U_0$, four times higher than without mirror. Therefore the maximum radial trapping frequency in the standing wave configuration is $\omega_r^{st} = 2\omega_r = 2\pi \cdot 324 \text{ kHz}$ and is achieved close to the mirror. The longitudinal oscillation frequency inside a local

trap at height z is given by

$$\omega_z^{st}(z) = \sqrt{\frac{-8A(z)\pi^2}{m\lambda^2}}. \quad (4.16)$$

Close to the mirror, the highest value $\omega_z^{st} = 2\pi \cdot 1100$ kHz is achieved. Note that, strictly speaking, the first anti-node is $\lambda/4$ above the mirror. This slightly reduces the trap depth, but the correction is so small that we may safely neglect it.

When the mirror is not exactly in the image plane of the DMD (finite distance L), the forward and backward propagating waves do not have the same amplitude. Therefore the longitudinal intensity modulation, whose peak-to-peak amplitude is

$$A(z) = \frac{4U_0}{\sqrt{1 + z^2/z_0^2}\sqrt{1 + (2L + z)^2/z_0^2}}, \quad (4.17)$$

where the origin for z is defined to be the beam-waist and the mirror surface is situated at $z = -L$. The amplitude $A(z)$ rapidly fades away as the distance between the mirror and the focal plane increases. Fig. 4.17 shows the longitudinal potential $U(z)$ for $L = 8 \mu\text{m}$. The deepest local minima are found at the focus of the beam. In general the deepest local minima are at found at $d = \sqrt{L^2 - z_0^2} - L$, for $|L| > |z_0|$, otherwise at $d = -L$, i.e. at the mirror surface. An atom trapped in one of these local minima sees a trap depth of about $1.37 U_0$ in the radial direction but only $0.3 U_0$ in the longitudinal direction. Whatever the distance between the focal plane

and the mirror, the trap depth in the radial direction will always be larger than in the travelling wave case ($L \rightarrow \infty$) but smaller than in the standing wave case ($L = 0$): $162 \text{ kHz} < \omega_r/(2\pi) < 324 \text{ kHz}$. Along the z -axis, the trapped atoms are confined either to the slowly varying envelope or to the local minimum, depending on their initial kinetic energy. Only the coldest atoms can be trapped in the local minima. If they are, their confinement shall be strongly improved. In that case the longitudinal frequency can be estimated from Eq. 4.16, with $A(z)$ given by Eq. 4.17. For the example of Fig. 4.17, one finds that the oscillation frequency at the focus of the beam is $\omega_z = 2\pi \cdot 300 \text{ kHz}$. Depending on L , the values of ω_z are intermediate between the travelling wave case ($L \rightarrow \infty$) and standing wave ($L = 0$) case: $35 \text{ kHz} < \omega_z/(2\pi) < 1100 \text{ kHz}$.

4.7 Ways to single atoms

In addition to controlling the size and depth of the traps, it is also important to control and measure the number of atoms. More specifically for the goal of conducting single atom experiments an efficient scheme to prepare single trapped atoms is essential. When loading atoms into moderately sized traps, the number of atoms therein is generally Poisson distributed. In this case, we would need to measure the atom number and then post-select traps with just one atom. Fortunately, for tightly focused dipole traps, the collisional blockade mechanism^[12] should give rise to a significant departure from the Poissonian statistics and favor the loading of only one atom per trap.

The effect occurs when atoms are loaded at a rate R from a magneto optical trap into tightly confined dipole traps. It relies on light-assisted collisions^[41] between atoms in the presence of the MOT cooling beam. Provided the latter is strong enough to saturate the transition, two-body losses occur at a rate $\beta'N(N-1)$. N is the atom number and β' is the rate constant, which is inversely proportional to the trapping volume. There is also a single body loss process γN , which is mostly due to background collisions, so the full rate-equation^[38] for the entire loading process reads

$$\frac{dN}{dt} = R - \gamma N - \beta'N(N-1). \quad (4.18)$$

Collisional blockade will occur when the two-body loss rate dominates over the loading rate of the dipole trap. Hence, it will work very efficiently when β' is large, i.e. if the atoms are confined to a tiny volume. The volume occupied by the atoms depends on the temperature and the trapping frequency. A comparison of the frequencies achievable with the proposed set up, and the trapping frequency $\omega_t = 2\pi \cdot 200$ kHz of a previous experiment demonstrating the collisional blockade^[12] leads us to the conclusion that here this mechanism could be highly efficient as well. Due to the nature of the collisional blockade, the traps would be filled with a random pattern of atoms, with a filling factor of at most 50%. Nonetheless, in connection with an in-situ atom detection and subsequent rearrangement of occupied traps, scaling to very large arrays of single trapped atoms seems feasible.

4.8 Single atom detection and transport

To investigate the loading processes, the atom numbers in the individual traps have to be determined. For well-isolated traps, laser induced resonance fluorescence could be used for atom counting^[42]. To collect the fluorescence, we suggest using the lens system which creates the optical tweezers in backward direction as a microscope, such that a spatially resolved fluorescence image can be obtained by a highly sensitive EM-CCD camera. This technique is well established and nowadays routinely used in many experiments for observing single atoms that are trapped in microscopic dipole traps^[14,43,44].

To discuss the feasibility of single atom detection, the flux of photons impinging on the camera has to be determined. The geometrical collection efficiency η depends on the numerical aperture of the lens system. To evaluate η , we assume a uniform photon emission of the atoms and calculate the ratio of photons passing through the system. The ratio of the surface area of a spherical cap with that of a sphere, yields $\eta = 2\pi (1 - \cos(\frac{\alpha}{2})) / (4\pi)$, where α is the opening angle. For a numerical aperture of 0.5, $\alpha = 60^\circ$, we expect a collection efficiency of $\eta = 0.067$.

If the radiative transition of the atom is completely saturated, the scattering rate on resonance tends to $R_{scat} = \frac{1}{2}\Gamma$, hence $R_{scat} = 19.6 \mu s^{-1}$ for Rubidium. Assuming a further loss of 50% to take the optical system and the limited quantum efficiency of the camera into account, a camera should collect photons from a single atom with

a rate of $R_{cam} = 653 \text{ ms}^{-1}$. In other words, a camera operating with an exposure time of $100 \mu\text{s}$ could count as many as 65 photons. For modern electron multiplying charged coupled device (EM-CCD) cameras, this is well above the signal to noise threshold, so it is possible to use such devices to observe single atoms. We emphasise that this estimate assumes no direct scattering of probe photons off the dielectric mirror into the camera. The latter effect can be eliminated altogether using probe beams traveling parallel to the mirror surface at a sufficient distance. Moreover, in case these beams are counter-propagating and red detuned, the trapped atoms would be subject to molasses cooling during observation, which allows for long observation times. In particular, this way of monitoring dipole-trapped atoms is not subject to broadening effects that otherwise limit the observation time of atoms in free flight^[45–47].

The high frame rate of the DMD should allow for the reconfiguration of the potential landscape, and hence for the transport of confined atoms. A naive method would be to transport atoms adiabatically by changing the potential slowly enough so that the atoms are following without being heated. Since the DMD has only a finite resolution, this must be done using discrete steps. With the aforementioned trap consisting of a four by four block, as shown in Fig. ??, the pattern would have to be shifted by at most one pixel per step. The longest distance the traps could be moved is 1000 pixels, that is from one edge of the DMD to the other. With a refresh rate of 50 kHz, this transport would take 20 ms. This is much shorter than

the anticipated lifetime of the trap, which should be of at least several 100 ms. However a drawback of this method is that discrete switching of mirror elements leads to parametric heating of the atoms. To circumvent this, the atom-transport could also be realised ballistically. The initially trapped atoms would be accelerated through a sudden change in trapping potential, guided through a channel and eventually decelerated and recaptured. This scheme has the advantage that it needs only a few changes in the potential, making it less susceptible to heating, and yet faster at the same time.

Chapter 5

Experimental Setup

There are two main parts to the experimental setup: the setup of the MOST, and that of the tweezer generation and observation. A schematic of the whole setup is shown in Figure 5.1 to be used as reference throughout the chapter.

5.1 MOST Lasers

In order to operate the MOST in the lab, two lasers are needed. The first laser is used to drive the main cycling transition in the D2 line of Rubidium 87, namely the transition between the $F = 2$ and the $F' = 3$ hyperfine levels, which is the transition used to implement the techniques outlined in Chapter 3. We refer to this as the *cooling* laser. The second laser is used to drive the $F = 1$ to $F' = 2$ transition: as was noted at the end of Section 3.1, this can be used to counteract off-resonant optical pumping, which leads to atoms being lost from the cycling transition. We

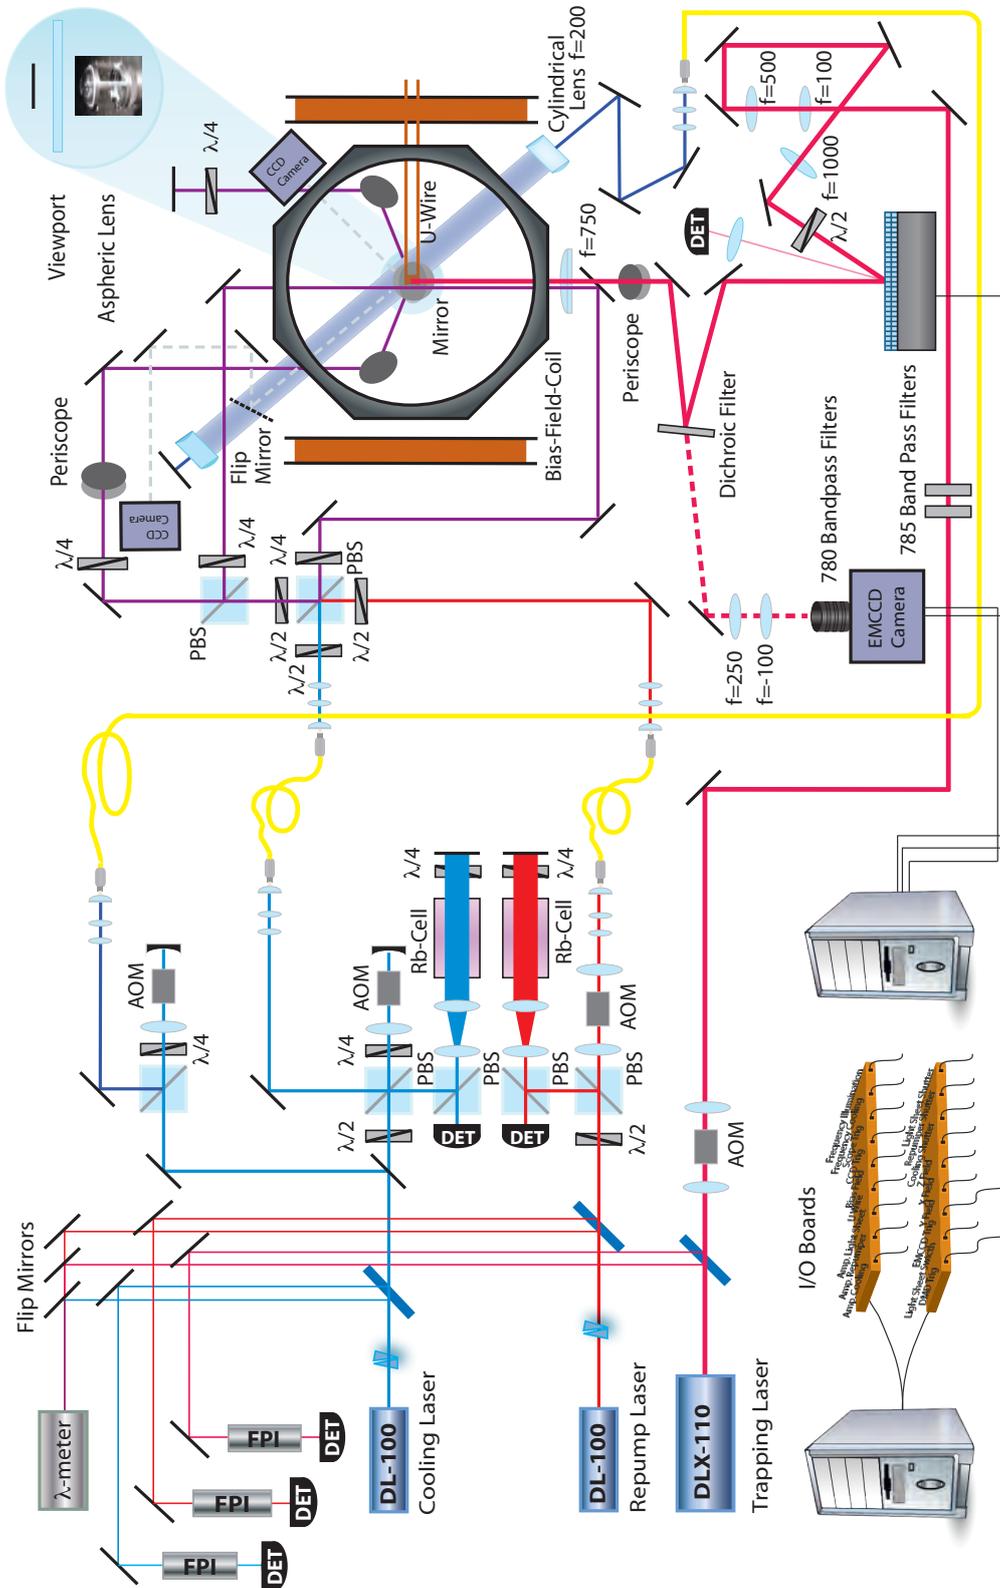


Figure 5.1 Experimental Setup

refer to this as the *repump* laser. The bulk of the experimental effort in setting up the MOST consists of stabilising and tuning the frequency of these lasers very precisely.

The cooling and repump lasers are both Grating Stabilised Tunable Diode Lasers from Toptica - the 780 nm Toptica DL100. These types of lasers, otherwise known as an External Cavity Diode Lasers (ECDLs), offer high stability, a narrow emission linewidth (and thus a long coherence length), precise wavelength selection, and wavelength tunability. The DL100 is specified to have a free running linewidth as low as 100 KHz, and an output power of up to 150 mW. It has a corresponding coherence length of approximately 1km, and coarse tuning of up to 110 nm. The laser head consists of an internal cavity, which is simply the laser diode, and an external cavity, which is formed by putting a collimating lens and a diffraction grating at the output of the laser diode chip. This diffraction grating acts as a wavelength selective element: the first order diffraction from the grating is fed back into the laser diode (whose outer facet is anti-reflection coated), forcing it into single-frequency operation. This configuration has the advantage that there is only one reflection off the diffraction grating, which means that there is no power loss into other diffraction orders and the overall power output of the laser is higher.

The emission wavelength of the laser can be shifted by rotating the diffraction grating. Rough tuning of this wavelength can be achieved by means of a microm-

eter screw, whilst fine tuning is achieved by the use of a piezo actuator. A beam steering mirror after the grating compensates for the slight change in the path of the output beam due to this rotation, and a Faraday isolator protects the laser from back reflections. The overall gain curve of the laser diode, on the other hand, can be shifted by changing the current (which changes the index of refraction of the diode) and the operating temperature. The overall mode structure of the laser is made up of the grating profile, the gain curve, the internal modes of the laser diode, and the narrow comb-like modes of the external cavity. By adjusting both the grating and the current precisely, it is possible to achieve a large mode-hop free tuning range. Feed forward current modulation enhances this tuning range. The current can be modulated up to 200 MHz, which can be used to produce sidebands for the Pound Drever Hall error signal discussed in Section 5.4.

All of the adjustments discussed are implemented via control modules also provided by Toptica. They include the scan control SC110 (which is the voltage controller driving the piezo actuator for the grating), the temperature control DTC 110, the current control DCC110, the proportional-integral-differential controller PID110, and the Pound Drever Hall detector PDD110. The latter two modules are used to lock the laser to a particular ^{87}Rb absorption line, as will be discussed in the next section.

The stability of the lasers is limited on the faster scale by electronic and acoustical

noise, and on a slower scale by temperature and air pressure fluctuations, mechanical vibrations, and piezo drifts. In our experiment, the two biggest problems were vibrations caused by knocking the table, and temperature fluctuations in the room.

5.2 Laser Stabilisation for the MOST

The cooling and repump lasers are both operated at lower than maximum power: at the output, the cooling laser has a power of 97 mW, and the repumper a power of 78 mW. The beam that emerges from the laser head is elliptical, so an anamorphic prism pair is used to compress one of its axes and make it circular. Glass plates are then used to split off a small fraction of the light. The reflection off the front of the glass plate is sent to a wavemeter, which provides a constant reference point for the wavelength. The reflection off the back of the glass plate is sent to a home-made Fabry-Perot interferometer, which allows us to see if the laser is running in single-mode. The length of the Fabry-Perot cavity is adjusted by means of a piezo tube that is glued to the end of one of the two confocal mirrors comprising the cavity.

In order for the MOST to function, the frequency of the cooling and repump lasers must be set very precisely: it must drive the right transitions and the frequency of the cooling laser must be red-detuned from the cycling transition by a precise amount. The frequency of the cooling laser must be stable to within a MHz, as it must be stable to well within the 16 MHz detuning from the cooling transition.

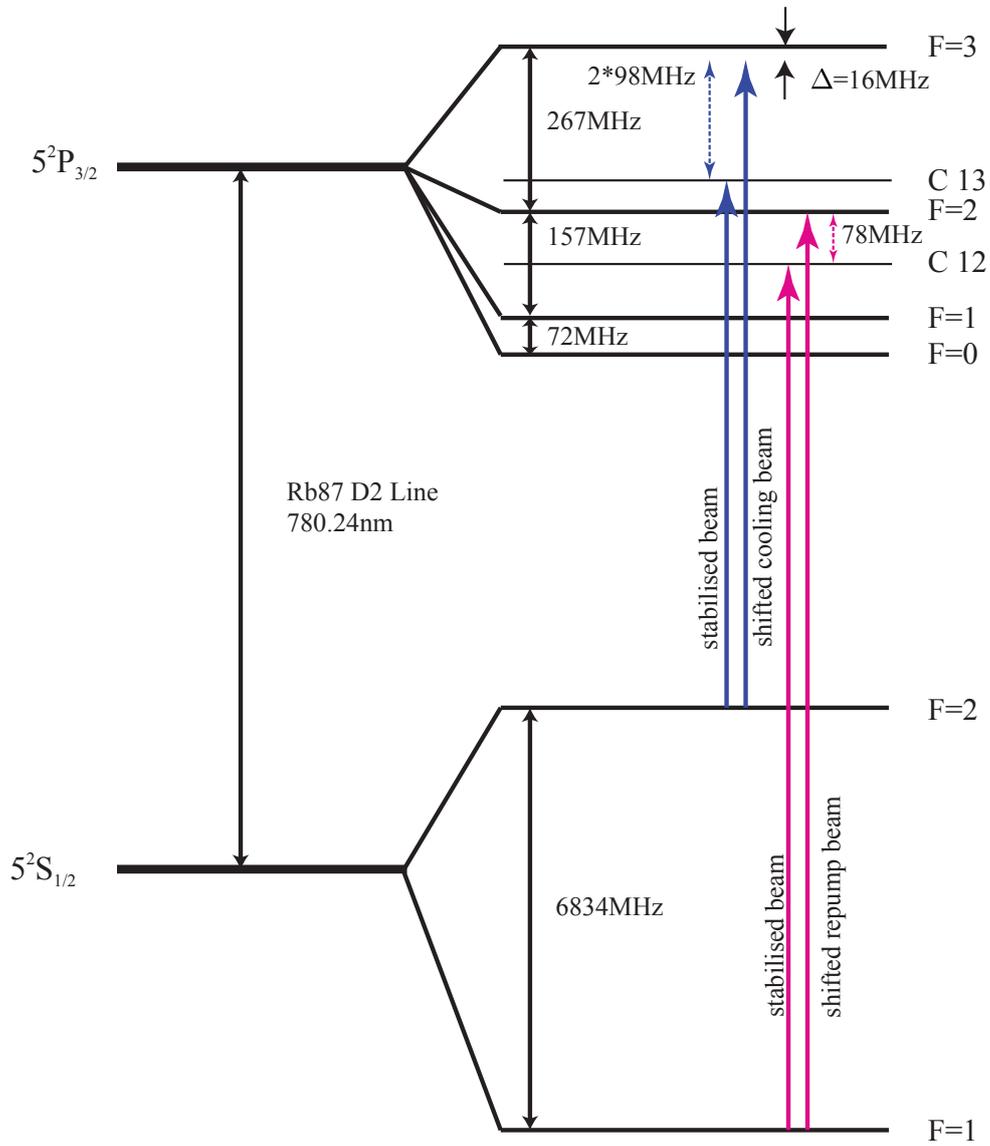


Figure 5.2 Laser Setup

Furthermore, the frequency of the two lasers must not drift away from the required setting. To attain this precision and stability, the lasers are locked to a convenient hyperfine crossover transition in ^{87}Rb by feeding a *Pound-Drever-Hall* (PDH) error signal into a Proportional Integral Derivative regulator (see Appendix D), and then shifted using Acousto Optical Modulators (see Section 5.5). The reference signal for the locking is obtained via Doppler-Free Saturated Absorption Spectroscopy of another sample (see Appendix C). The relevant transitions and detunings are shown in Figure 5.2.

5.3 Acousto Optical Modulators

As was mentioned in Appendix B, for an Optical Molasses or a MOT to function, the cooling beams must be slightly red detuned from the relevant transition. This means that it is necessary to control the frequency of the stabilised cooling laser very precisely. In addition, it is desirable to have precise control of the amplitude of both the cooling and the repump lasers. In the experiment, this requirement is met by the use Acousto Optical Modulators (AOMs), which use acoustic wavefronts propagating in a crystal to modulate the frequency, intensity, and direction of a laser beam. This occurs due to the Bragg diffraction of the incoming light beam off the acoustic wavefronts: by modulating the frequency and amplitude of the wavefronts, modulation of the incoming laser beam is achieved. The change in frequency occurs because momentum is exchanged between the phonons of the acoustic wave and the incident radiation. The condition for constructive interference of the diffracted light

beam is

$$n\lambda_L = \lambda_A(\sin\theta_i + \sin\theta_d) \quad (5.1)$$

where λ_L is the wavelength of the incoming laser beam, λ_A is the acoustic wavelength, and θ_i and θ_d are the angle of incidence and diffraction of the beam respectively. Therefore, the change in frequency of the light beam occurs in integer multiples of the acoustic frequency, and can be either positive or negative. If it is aligned correctly, most of the power coming out of the AOM will be in the first diffracted order. A piezo transducer produces the acoustic waves in the crystal, which are usually at radio frequency, and in the experiment can be tuned around a central frequency of approximately 100 MHz, with a bandwidth of about 30 MHz.

An AOM can operate in either a single pass or a double pass configuration. In the latter, the first order is picked out by a pinhole aperture and retro-reflected back into the AOM such that it experiences a second equivalent shift in frequency. This is accomplished by placing a lens and a mirror (or, as in our case, simply a curved mirror) after the AOM. The advantage of the single pass configuration is that less power is lost, but its disadvantage is the inevitable change in beam pointing that comes with the desired change in frequency. The double pass configuration leads to a higher loss in power, but it circumvents the beam pointing problem because the beam pointing is shifted in opposite directions upon each pass through the AOM.

In the double pass configuration, the beam traveling back from the AOM is collinear with the incoming beam. It is separated by using a polarising beam splitter and a $\frac{\lambda}{4}$ plate which changes the polarisation of the incoming beam such that the twice shifted beam exits the opposite port of the beam splitter. As can be seen in Figure 5.1, the $\frac{\lambda}{4}$ plate sits in front of the mirror used to retroreflect the singly shifted beam, so the beam traverses it twice. Thus, the $\frac{\lambda}{4}$ plate has the same effect as a $\frac{\lambda}{2}$ plate, and the axis of polarisation of the beam changes from vertical to horizontal, allowing the beam to exit the opposite face of the beam splitter.

In the experiment, the cooling beam is stabilised to the crossover between the $F = 2$ to $F' = 1$ and the $F = 2$ to $F' = 3$ transitions, and is shifted up by 196 MHz to run slightly below the $F = 2$ to $F' = 3$ transition using an AOM in double pass configuration. As can be seen in Figure 5.2, this corresponds to a detuning of 16 MHz. The repump beam, on the other hand, is stabilised to the crossover between the $F = 1$ to $F' = 1$ and the $F = 1$ to $F' = 2$ transitions, and is shifted onto the $F = 1$ to $F' = 2$ transition using an AOM in single pass configuration. A further AOM in double pass configuration is used to shift the beam used for the light sheet/absorption imaging (which is split off from the cooling beam) up by $\approx 214\text{MHz}$ to run slightly above the $F = 2$ to $F' = 3$ transition.

The AOMs are controlled by home made drivers consisting of a voltage controlled oscillator, which provides frequency control, a mixer, which provides amplitude con-

trol, a directional coupler to monitor the output frequency, a pre-amplifier, and an RF switch to attenuate the output power completely. Their signal is passed on to a commercial RF amplifier, with an output power of 1.6 W. The AOM in single pass configuration has an efficiency of up to 85%, whilst the two AOMs in double pass configuration have efficiencies of $\approx 75\%$.

After the beams pass through the AOMs, they are coupled into single mode polarisation maintaining fibres. These fibres provide stability as well as a clean Gaussian beam profile, and they allow for the subsequent beam setup around the vacuum chamber to be decoupled from all of the laser stabilisation. A telescope is used at the output of each fibre to set the beam diameters to approximately 6mm. The cooling and repump beams are combined on a polarising beam splitter, which also serves to separate the diagonal and horizontal beams need for the MOST. Further polarising beam splitters and $\frac{\lambda}{2}$ plates are used to set the splitting ratios of the six requisite MOST beams, and $\frac{\lambda}{4}$ plates are used to make the polarisation circular and obtain the correct helicity necessary for the MOST to function. Each MOST beam has an intensity of approximately 8mW, which is equivalent to an intensity of approximately $10I_{Sat}$.

5.4 Around the Vacuum Chamber

The experiment is conducted inside a vacuum chamber in order to avoid collisions with the atomic background. As can be seen in the left hand panel of Figure 5.3,

the chamber is octagonal, with two large viewports at the top and bottom, and six small viewports on the sides. One of these side viewports is connected to an extension that has attached to it a pressure gauge, an Ion Getter Pump, and a Titanium Sublimation Pump, all of which are used to reach the High Vacuum regime. A second viewport provides access for a metal piece holding the U-shaped rod and the mirror needed for the MOST. A third viewport contains electrical feedthroughs for the rubidium dispenser, which releases rubidium gas into the vacuum chamber by means of sublimation when a current of 5A is fed through it. The other three viewports are available for optical access, as is the top viewport. The pressure inside the chamber can be as low as 10^{-8} mbar, rising when the dispenser is switched on. As can be seen in the bottom right hand panel of Figure 5.3, the MOST mirror must be mounted on a ceramic holder to avoid contact between the mirror and the U-shaped rod, which is shown in the top right hand panel of Figure 5.3.

The constant bias field needed for the MOST is provided by two large sets of coils on either side of the vacuum chamber, around the viewports used for the U-shaped rod and the pumps. These have a radius of 8 cm, and contain 40 turns each. Two 20A power supplies in parallel provide a maximum current of 40A for the U-shaped rod, whilst a single power supply provides a maximum current of 20A for the bias coils. The U-shaped rod and the bias field are usually operated at about 24A and 18A respectively. Three extra pairs of coils provide independent control of the MOST along three orthogonal axes, one of which is collinear with the bias field.

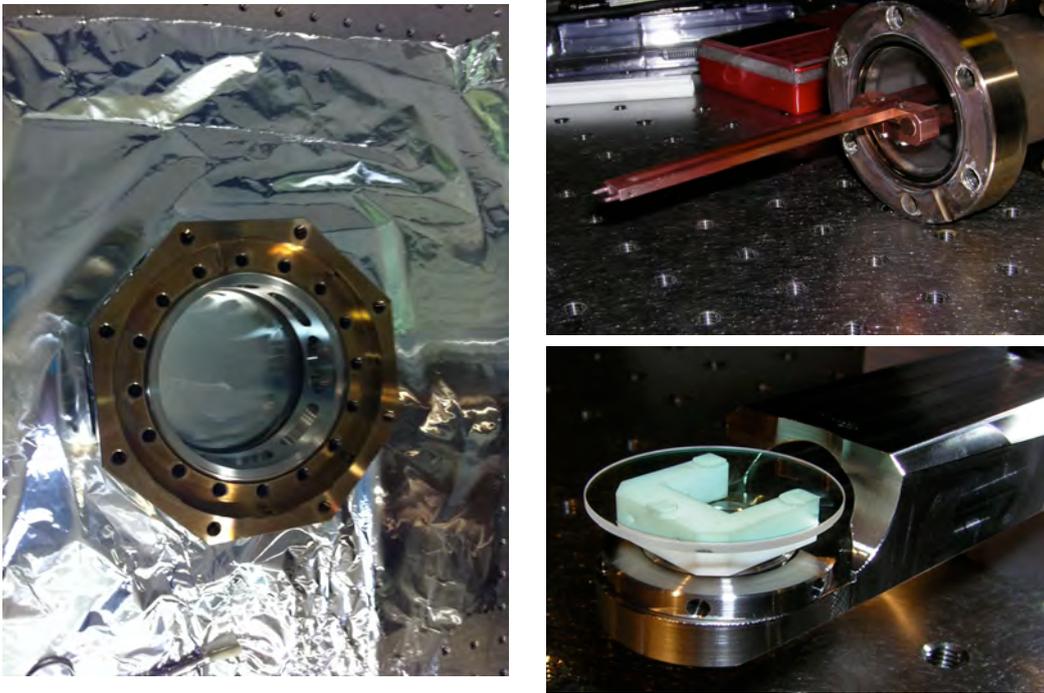


Figure 5.3 The left hand panel shows the octagonal vacuum chamber in which the experiment is conducted. The right hand panel shows the U-shaped rod (top) and the ceramic holder on which the MOST mirror is mounted (bottom).

We refer to these as the x coil, y coil, and z coil, and they are each powered by a 10A power supply. A further 10A power supply gives the 5A necessary to run the dispenser.

As can be seen in Figure 5.4, a stainless steel platform allows for the diagonal MOST beams to enter the vacuum chamber from above. The two CCD cameras that can be seen in the figure are used to monitor the MOST from two perspectives. A flip mirror gives the camera on the right hand optical access when the light sheet, which exits through this viewport, is not on. This camera was originally used to

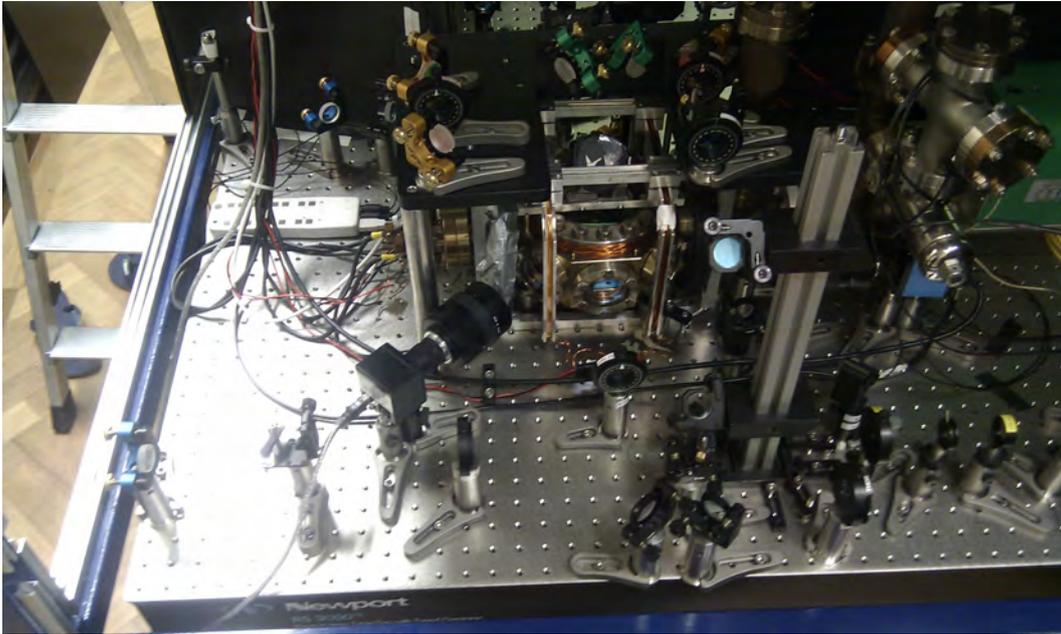


Figure 5.4 Setup surrounding the vacuum chamber.

take absorption images of the atom cloud.

5.5 The Digital Mirror Device

The DMD is a commercially available spatial light modulator (the DMD Discovery 1100 from Texas Instruments) that consists of 1024 by 768 independently addressable aluminium micro-mirrors, $13.7\mu\text{m}$ by $13.7\mu\text{m}$ in size. They 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. The mirrors tilt about a hinged diagonal axis by means of electrodes lying underneath the two corners opposite the axis of

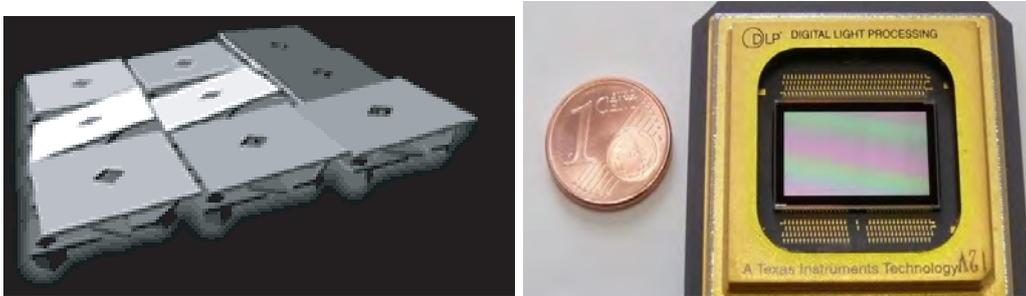


Figure 5.5 The left hand side of the figure shows the DMD micro-mirrors tilting about a diagonal axis. The mirror tilt is 12° , which means that the angle of incidence of the dipole trapping light must be at 24° to the normal when the mirrors are inactive (flat) in order to direct it along the optical axis when they are on. The right hand side of the figure shows the size of the entire mirror array, with a one cent coin for comparison.

rotation(see Figure 5.5). As a consequence, the casing holding the mirror array and the printed circuit controller board must be mounted at a forty-five degree angle (see Figure 5.6. The entire optical surface covered by these mirrors, which have a reflectivity of 89%, has a size of 14 mm by 10.5 mm (see Figure 5.5), and it is protected by a window which is optimised for near infra-red wavelengths. The efficiency of the DMD, which takes into account the window transmittance, the active-area fill factor of the mirrors, the specular reflectivity of the mirrors (89.4%), and the diffraction efficiency, is 68%. The damage threshold of the DMD mirrors is given by a maximum incident power density of $5W/mm^2$.

The DMD can display either static patterns (Figure 5.6 shows what these static patterns look like on the mirror array) or movies made from frames of these patterns. The DMD can be refreshed at a rate of between 4 and 40 kHz, allowing us to display

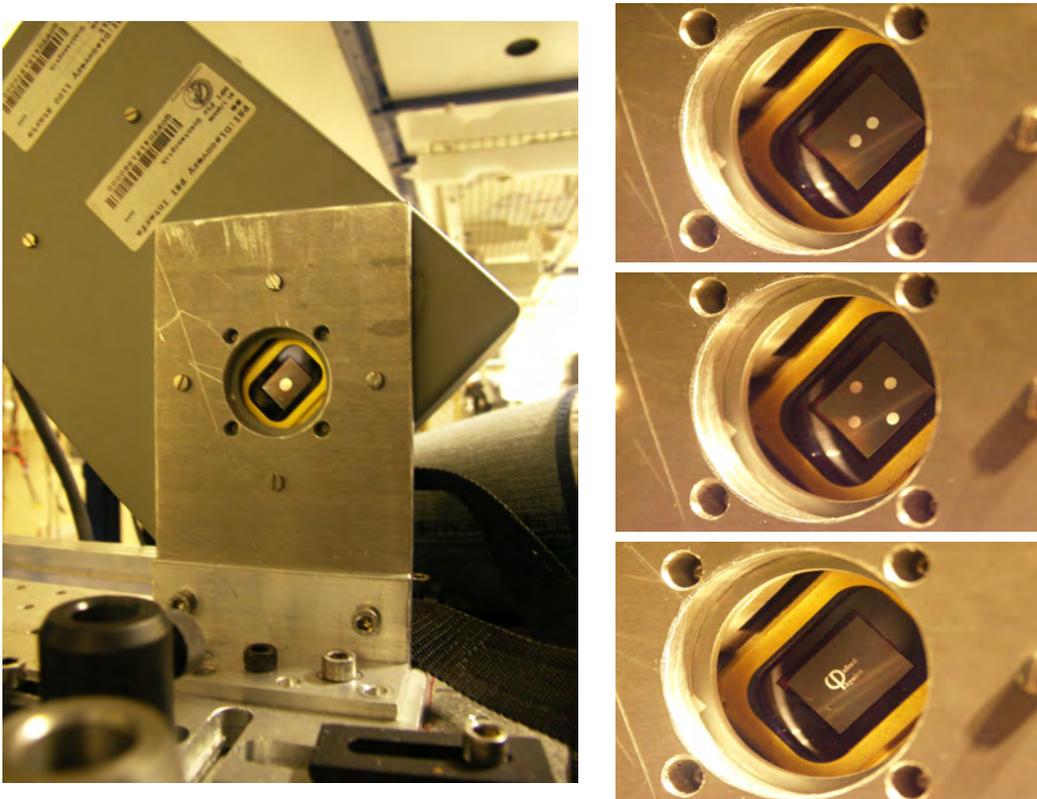


Figure 5.6 The left hand side of the figure shows how the casing housing the mirror array and the controller board must be mounted at a 45° angle. The right hand side of the figure shows several different bitmaps loaded onto the DMD.

high-speed movies, and making it fast enough for a dynamic control of the atoms trapped in the tweezers. The lowest refresh rate occurs when the whole array is loaded with new information (we refer to this as a “reset”) and the entire mirror array is subsequently switched (we refer to this as a “refresh”). The highest refresh rate occurs when only 48 lines of the array are reset and refreshed (as explained in more detail later, the DMD can only be reset by lines and refreshed in increments of 48 lines). These refresh rates do not take into account the amount of time required to transfer the necessary bitmap or mirror loading commands onto the PCI controller

board, but only the amount of time required to pass the former from the PCI board to the DMD onboard memory (the reset action), and to physically switch the mirrors (the refresh action).

The lowest refresh rate quoted by the manufacturer is 10 KHz, for a full array reset and refresh, and they give an upper limit of 40 KHz for a reset and refresh of “less than 190 rows”. Figure 5.7 shows that the actual time required to reset and refresh the whole mirror array is $250\mu s$, whilst the actual time required to do a partial refresh of only 144 rows is $170\mu s$. The former is 2.5 times as long as the $100\mu s$ (10KHz) quoted by the manufacturer, whilst the latter is nearly 7 times as long as the $25\mu s$ (40KHz) quoted for a partial refresh of 190 rows or less. Nevertheless, the measured rates are still on the same timescale as the atomic motion, and allow for transport without heating, so this did not prove to be a limitation.

The fastest full array refresh time achievable with our DMD was $50\mu s$, and as can be seen in Figure 5.8, this was whilst changing only the first 48 lines of the mirror array. As explained in more detail in the following section, 48 lines are referred to as one block, and correspond to the minimum number of lines that must be physically switched, even if the state of only one mirror is changed. Figure 5.9 shows the fastest partial array refresh time achievable with the DMD, which was $26\mu s$ (40KHz). In this case only 48 lines were reset, and only those 48 lines were refreshed. This refresh rate is very fast, but limits us to only a very small area on the

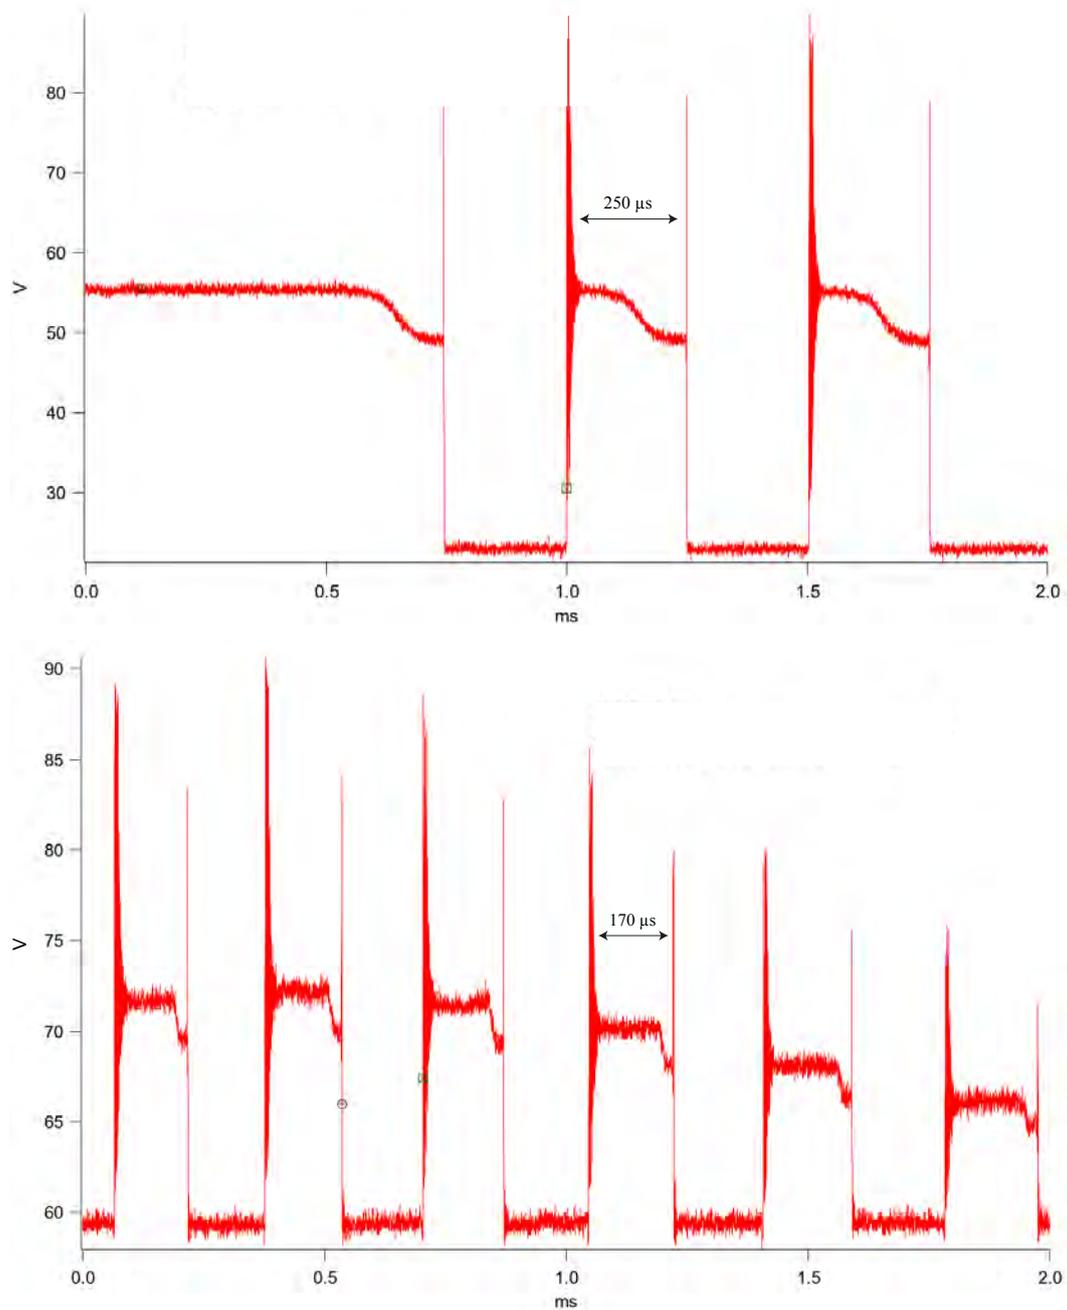


Figure 5.7 One of the diffraction orders of the DMD is monitored on a photodiode in order to determine the refresh time of the DMD. The top graph shows the time required for a full array reset and refresh, whilst the bottom graph shows the time required for a partial reset and refresh of only 144 lines.

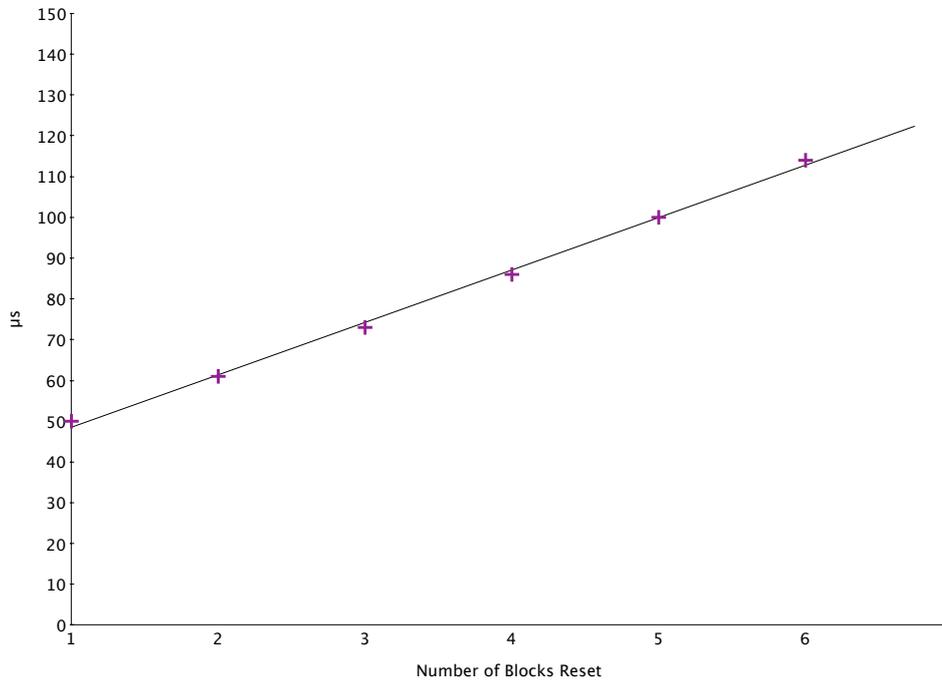


Figure 5.8 Time required for a partial reset followed by a full array refresh as a function of the number of “blocks” reset, where each block is equivalent to 48 lines on the DMD.

DMD. The documentation explains that the maximum refresh rate of the DMD is determined by the mirror transition and settle time, which is essentially the time it takes for the mechanical vibrations to die out after a mirror is switched and is given as $18\mu s$, plus the time required to reload the mirror section with new data, which is given as $6\mu s$. We confirmed this transition and settle time to be approximately $20\mu s$ (see Figure 8.5).

5.6 The DMD Control

Simple tasks such as the display of static patterns on the DMD can be carried out by using the commercial software program provided by Texas Instruments with the

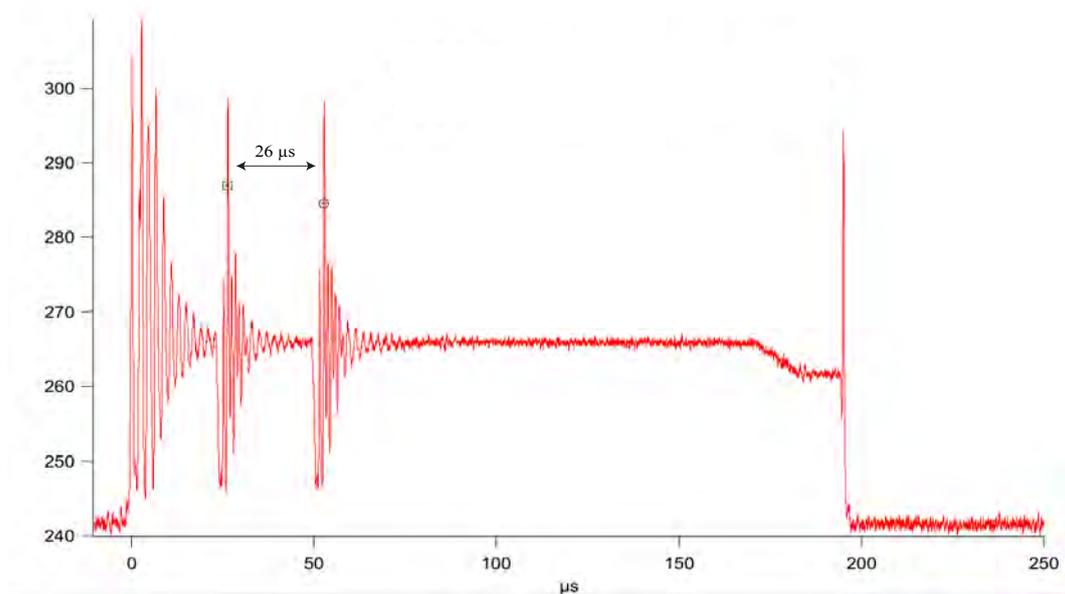


Figure 5.9 Scope trace of the photodiode signal from one of the diffraction orders of the DMD showing the time required for a reset and refresh of only one block of 48 lines in the centre of the DMD. In this case, three consecutive refresh actions follow a single reset command.

DMD and its PCI board. The implementation of more sophisticated sequences, however, called for a customised control program to be written in-house. This control program was written in visual C++, and consists of interactive dialog boxes with which the user can access routines that allow for moving traps, contracting traps, ballistic transport and several other sequences. These routines also implement all of the basic API (Application Programming Interface) functions of the DMD. For instance, the DMD can be set totally black or totally white, and it is possible to load and buffer bitmaps from disk which can then be fully or partially refreshed either manually or by means of a trigger.

The more complicated sequences implemented in the program rely on an API function that allows one to write directly to the PCI board's SDRAM (Synchronised Dynamic Random Access Memory). This function makes it possible to set individual mirrors to either on or off by addressing the correct byte position in the mirror array, and in turn allows for circles, disks, and any other shape to be written to the SDRAM. Addressing a particular byte position within the array is not straightforward. The DMD mirrors are grouped into 16 individually controlled blocks, with each block containing 48 rows of 1024 mirrors, or 49,152 mirrors (each mirror corresponds to one bit, so each block contains 6,144 bytes). Images are controlled and displayed in these blocks, individually or in groups of up to 16 (the whole array). The mirror data for each line are loaded through a 64 bit bus, so the order of the mirrors is (15, 31, 47,, 1023) = group 1, (14, 30, 46,, 1022) = group 2,, (0, 16, 32,, 1008) = group 16, with each group containing 64 bits (8 bytes). Addressing a particular byte position then means addressing the correct block, the correct group, and the correct position within each group in an efficient manner.

In addition to the PCI SDRAM, there is an onboard display memory on the DMD's controller board that holds only the data that is about to be displayed. The SDRAM write function stops short of loading data onto this display memory or switching the mirrors. To actually display an image on the DMD, two functions must be called. The image is first loaded from the PCI SDRAM onto the display memory using a "Refresh" function (like taking an image out of a drawer and putting

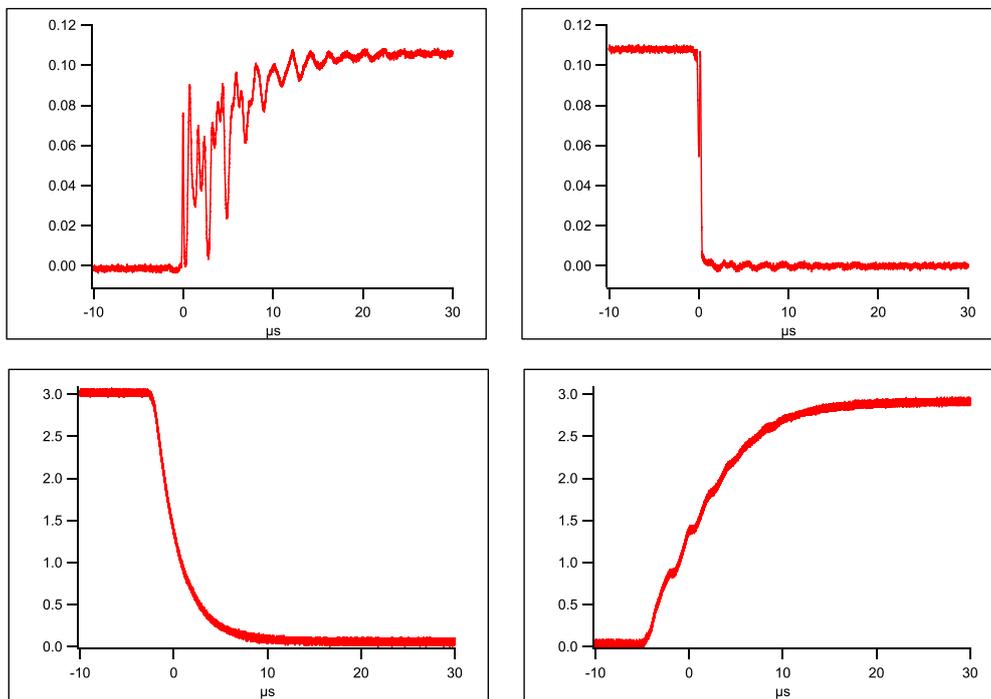


Figure 5.10 Mirror transition and settle time when only one block is reset and refreshed. The top two graphs show the signal from a single 4 pixel spot, with the turn-on time on the left, and the turn-off time on the right. The “ringing” that can be seen on the turn-on signal is the mechanical vibration of the mirrors after they are switched on, and dies out after approximately $20\mu s$. The bottom two graphs show the signal from a larger 48 pixel spot, and show how these mechanical vibrations get “washed out” by the contribution from a greater number of mirrors.

it on the ready to go tray), and the mirrors are then switched to physically display the contents of the display memory using a “Reset” function. The Refresh function loads from one line to all 768 lines of data onto the display memory, whilst the Reset function switches the mirrors in no less than one block of 48 lines. The fastest frame rates are achieved by preloading all of the requisite images on the PCI SDRAM and then just using a sequence of Refresh and Reset commands, with the least number of necessary lines and blocks refreshed and reset each time.

The simplest routine implemented by the control program is the ballistic transport, as it consists of simply preloading a sequence of bitmaps onto the PCISDRAM, and then triggering them at the appropriate times. These bitmaps are 2D harmonic potentials, with the requisite greyscales mimicked by dithering. A round start trap is followed by an elliptical transport trap, and some time later, a round end trap that serves to recapture the atoms. A black screen can be switched on after this to find the temperature after transport by the TOF method. The time in the transport potential can be changed manually, or controlled by an external trigger. If the end trap is switched on immediately after the transport trap, with no trigger or delay, then the minimum transport time is dictated by the full-array refresh rate of $250\mu s$.

The DMD program receives external triggers from the LabView program used to control all of the experimental parameters for the MOST, cameras, light sheet, and mechanical shutters. The latter program runs on a separate computer, and interacts with the experiment via two I/O boards. The DMD computer is connected to one of the I/O boards via its parallel port (the erstwhile printer port), and the triggers are received via pin 10, which is normally one of the status ports.

5.7 The DLX110

The DMD is illuminated using a high power diode laser, the Toptica DLX-110. This laser has a maximal power of 800mW, and is red-detuned by 5 nm ($\lambda=785\text{nm}$) from the $5^2S_{1/2} \rightarrow 5^2P_{3/2}$ transition in ^{87}Rb . We chose this relatively small detuning

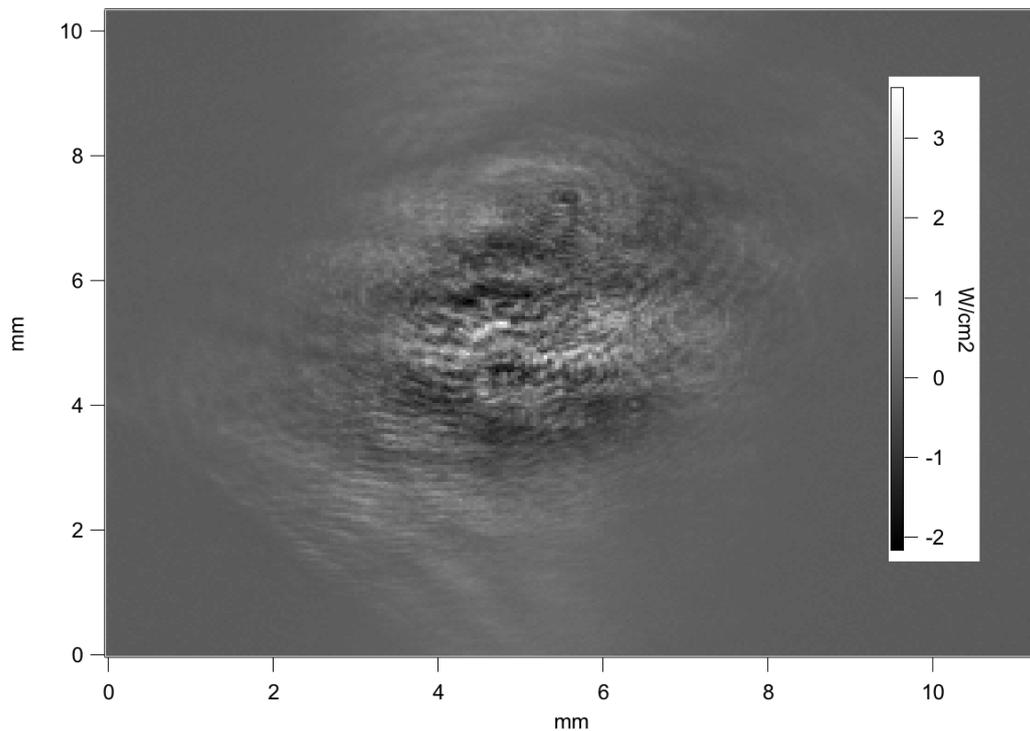


Figure 5.11 Difference between the intensity distribution of the trapping laser on the DMD and a Gaussian. Note that the image is merely meant to illustrate the quality of the beam profile: the intensity represented on the scale in the inset is inverted, and does not represent the actual intensity in the experiment, as several OD filters were used when taking the picture.

to achieve a reasonable trap depth at light intensities remaining below the damage threshold of the DMD. As mentioned previously, the illuminating beam would ideally have a Gaussian profile that homogeneously filled the area of interest on the DMD. Unfortunately, the DLX-110 has a sub-optimal laser profile that can only be made Gaussian by spatial filtering, and this leads to a huge loss in laser power and thus a lower trap depth. In the end, spatial filtering was avoided, and only the wings of the beam were clipped off with an aperture, leaving us with slightly more power

but an uneven intensity distribution on the DMD (see Figure 5.11). The beam was initially controlled by means of an AOM, but this led to even more loss in laser power, and was removed as well. An additional loss in power is accrued due to the diffraction orders of the DMD as well as mirrors in the off state (which direct light away from the optical axis, onto a beam dump), and by the use of two band pass filters, which are used to eliminate any ASE (Amplified Spontaneous Emission), for which $\lambda \neq 785\text{nm}$ (Semrock Maxline LL01-785). Taking into account the damage threshold of DMD and an optical magnification of about 50, the intensity of the laser in the focal plane of the imaging system should be on the order of 25 kW/cm^2 , or 10^7 saturation intensities. However, due to all of the power losses, this figure is actually only 3 kW/cm^2 , which corresponds to a maximum potential depth of $U_0 = k_B \cdot 100\mu\text{K}$.

5.8 Imaging Setup

The microscope system used to image the surface of the DMD consists of a high numerical aperture aspheric lens with $f_1 = 20\text{mm}$ which is mounted inside the vacuum chamber, and an achromatic doublet with $f_2 = 750\text{mm}$ which is placed outside the vacuum chamber. The aspheric lens has a numerical aperture of .54, and is mounted on a stainless steel holder which holds the lens 20.15mm away from the MOST mirror. This custom made mount was designed with the least possible material so as to not obstruct the passage of the MOST beams and the light sheet (see Figure 5.12).

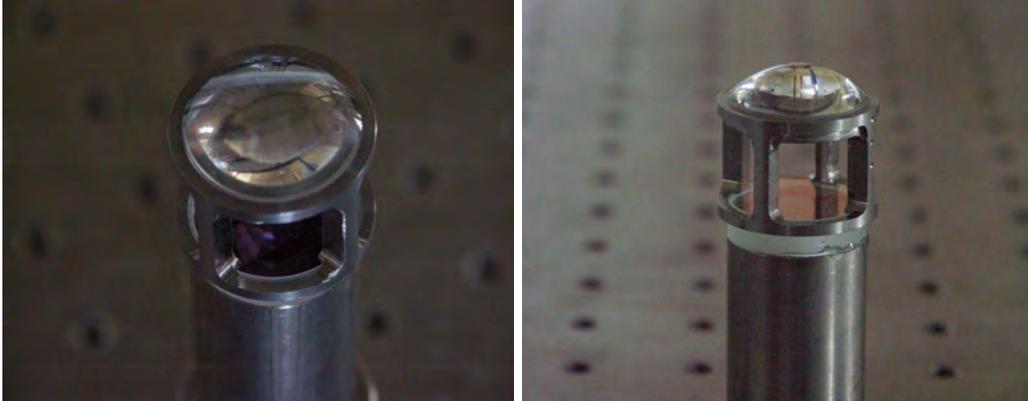


Figure 5.12 The custom made lens mount holds the aspheric lens 20.1mm away from the surface of the MOST mirror inside the vacuum chamber.

In addition to the aspheric lens and the achromatic doublet, there are three lenses placed before the DMD, with $f_3 = 100\text{mm}$, $f_4 = 500\text{mm}$, and $f_5 = 1000\text{mm}$ respectively. These lenses serve to adjust the convergence of the illuminating beam and its size on the DMD. All five lens distances are chosen such that isoplanatism is ensured and such that the image plane of the DMD lies several hundred microns above the surface of the MOST mirror. A dichroic filter in the middle of the optical path splits off the trapping light at 785nm from the collected atomic fluorescence at 780nm, and a further two lenses with $f_1 = 250\text{mm}$ and $f_2 = -100\text{mm}$ respectively are placed between this filter and the EMCCD camera. Four 780nm bandpass filters (Semrock Maxline LL01-780, with an optical depth of 0.4 at 785nm) are used to block out any residual off-resonant trapping light, for although the transmission of the dichroic filter at 785nm is only on the order of 10^{-3} , the amount of trapping light that gets through is still much higher than the collected atomic fluorescence.

The whole region is sectioned off with optical shielding to block off any scattered resonant light from the cooling beam and the repump beam. The full optical beam path is shown in the lower right hand corner of Figure 5.1.

The optical setup used to generate the traps was simulated in Mathematica using ABCD matrix formalism and precise measurements of the distances between all of the optical elements in the setup. Using this simulation, the position of the image plane was calculated to lie $400\mu\text{m}$ above the surface of the mirror (19.6mm away from the focal plane of the aspheric lens), and the whole system was found to have a magnification of 57. The optical setup used to image the atomic fluorescence onto the EMCCD camera (which from the dichroic filter to the mirror is the same as setup for the trap generation, in reverse) was similarly simulated in Mathematica. The position of the object plane was calculated to lie 20.7mm away from the focal plane of the aspheric lens, which means that the EMCCD camera actually looks at the reflected fluorescence light (or the reflected MOST, as the case may be). The magnification in this case was found to be 35.8.

In addition to all the lenses, several mirrors fold the beam as necessary, a $\lambda/2$ plate adjusts the polarisation of the trapping light to make it either horizontal or vertical, and a periscope directs the beam into the vacuum chamber through its top viewport. The mirrors give rise to some optical losses and birefringence, with the major loss in power due to the viewport (9mm thick fused silica). These losses do

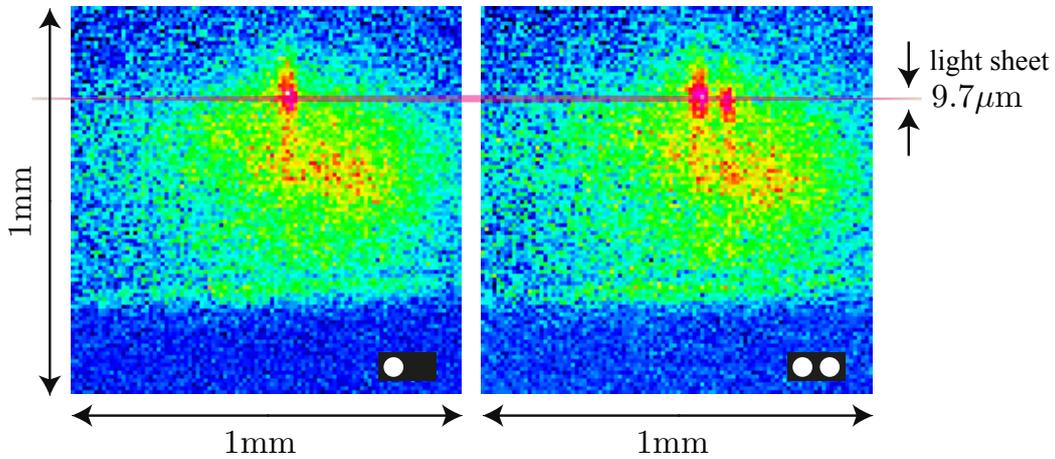


Figure 5.13 Side-view false-colour absorption images of atoms in either one or two dipole traps, with a radius of $\approx 30\mu\text{m}$ each. The snapshots were taken immediately after loading, with a diffuse cloud of untrapped atoms from the MOST still surrounding the traps (blue: no absorption; red: maximum absorption and highest atom density)

not greatly affect the trapping beam, but must be taken into account when determining the number of trapped atoms from the recorded fluorescence.

5.9 The Light Sheet

As mentioned previously, an important feature of our microscope system is its confocal geometry: the same two lenses used to image the DMD are also used to observe the atoms trapped in the tweezers. The top view through the microscope system is well-suited for observing trapped atoms in a two-dimensional plane, but one has to take into account that the trapping geometry is actually three dimensional. Irrespective of whether there is a strong standing wave modulation or not (which depends on the radius of the trap), atoms will be trapped in a vertical column perpendicular

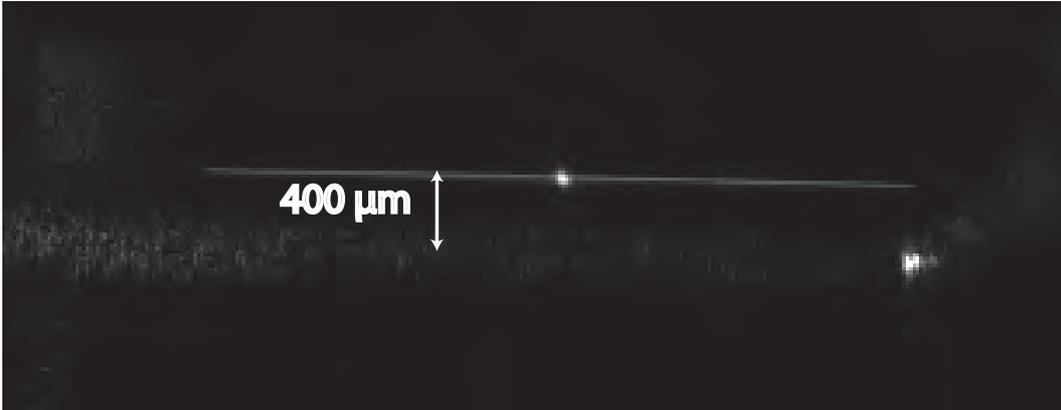


Figure 5.14 Side-view of the light sheet crossing the centre of the MOST. The reflections on the right hand side come from the mount holding the aspheric lens.

to the image plane. We monitor this from the side by absorption imaging: Figure 5.13 shows a particular situation with atoms loaded into either one or two trapping beams with a radius of $30\mu\text{m}$. Observing the atoms by fluorescence imaging through the microscope system would normally give rise to blurry pictures, as there would always be a contribution to the fluorescence signal from out-of-focus atoms lying above and below the plane on which the camera focuses. As we are primarily interested only in those atoms trapped in or near the image plane of the DMD, we focus on this plane and avoid the excitation of the out-of-focus atoms by illuminating the trap with a very thin sheet of resonant light.

The light sheet has a narrow height of $9.7\mu\text{m}$ along z (perpendicular to the mirror surface), and is 5.1 mm wide along y (parallel to the mirror surface), resulting in a Rayleigh length of $\approx 100\mu\text{m}$ along the x direction. To generate the light sheet, a near resonant beam is focussed down in only one direction by using a cylindrical

lens with a focal length of 200mm. This beam is split off from the cooling laser and blue-detuned by 2 MHz from the $5^2S_{1/2} \leftrightarrow 5^2P_{3/2}$ cycling transition using an AOM in order to account for the dynamic Stark shift induced by the presence of the tweezers. The light sheet is retro-reflected to avoid pushing the atoms out of the traps, and its position is controlled by using piezo driven optical mounts. The total power in the light sheet is $7 \mu\text{W}$, which gives rise to a scattering rate of 16 photons/ μs . To detect the trapped atoms at any given moment, the light sheet is flashed on for $90\mu\text{s}$. The timing, duration, and shape of this pulse must be very precise, as it is instrumental in determining the number of fluorescence photons emitted in one loading cycle. It is therefore monitored periodically on an oscilloscope. Figure 5.14 shows the light sheet crossing the centre of the MOST.

5.10 The Electron Multiplying Charge Coupled Device Camera

The atomic fluorescence signal collected through our optical system is very weak, especially when the number of trapped atoms is small, so we use a very sensitive Electron Multiplying Charge Coupled Device (EMCCD) camera to observe it. Structurally, an EMCCD camera is identical to a conventional CCD camera, but with its sensor's shift register extended to include an additional section called the Multiplication or Gain Register (see Figure 5.15). Our camera, the front-illuminated Andor iXon 885 (DV885LC-VP) is an ultra sensitive low light EMCCD camera that

is capable of detecting down to a single photon. Its sensor, an Impactron Frame Transfer sensor from Texas Instruments, is a silicon-based semiconductor chip made up of 1002 x 1004 photo-sensitive pixels which measure $8\mu\text{m} \times 8\mu\text{m}$, for a total image area size of 8mm x 8mm. Appendix E provides a detailed description of how this camera works.

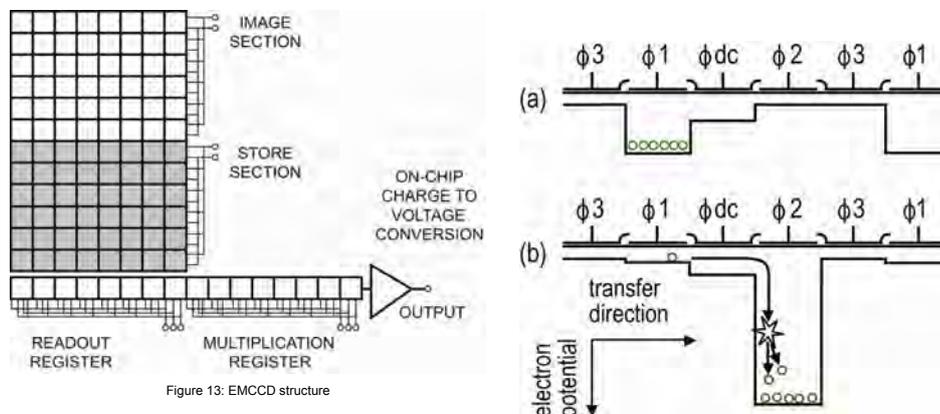


Figure 5.15 The left hand side of the figure shows the additional Multiplication Register that can be found on a typical EMCCD chip. The right hand side of the figure illustrates the basic operation of this Multiplication Register. Both diagrams were taken from the camera’s documentation.

For most of the results presented in Chapters 7 and 8, the binning of the camera was set to 16x16, the EM Gain to 300, the readout rate to 35 MHz, and the pre-amplifier gain to 3.8. The meaning and consequence of these parameters is explained at length in Appendix E. The camera is regularly air cooled down to -80° , and operated in an external exposure mode, with the trigger pulse length set to match the flash of the light sheet so that no light falls on the sensor outside of the

exposure window. We tested the quantum efficiency (QE) of the camera by shining a weak laser beam into it and using a separate powermeter to determine the number of photons incident on the sensor in one second. We found the QE for our camera to have an actual value of 43%, as opposed to the 65% at 780 nm quoted in the documentation (see Appendix E).

To be of any use experimentally, the number of counts output by the camera's ADC must be converted into photons. This will be addressed in the next chapter, but is done in roughly the following way. Firstly, the electronic bias offset on which the counts sit must be subtracted (this electronic bias offset ensures that the displayed signal level always corresponds to a positive number of counts). Secondly, these counts must be multiplied by the number of electrons per count for the chosen pre-amplifier gain, as given in Figure E.2 in Appendix E. Thirdly, this number must be divided by the chosen EM gain. Finally, this number must be multiplied by the QE of the camera. The resulting figure corresponds to the number of photons that is actually recorded by the camera's sensor, and not the number of photons that are actually emitted by the trapped atoms. To estimate the latter, it is necessary to also account for additional losses on the way to the camera, like the collection efficiency of the aspheric lens.

5.11 MOST Computer Control

To perform any experiments with the MOST, and in particular, to load atoms from the MOST into the dipole trap, it is necessary to have rapid control of all the experimental parameters and be able to perform sequences at timescales comparable with the speed of the atoms, which in the Doppler cooling limit is approximately 17 cm per second. This is equivalent to $170\mu\text{m}$ per ms, so given that the spatial dimensions of the traps are on the order of hundreds of microns or less, control of the experimental parameters at a sub-millisecond timescale is required. We accomplish this by means of the two Adlink DAQ-2502 IO-boards illustrated in Figure 5.1, which have a sample rate of 100KHz and are controlled by a LabView program. This program allows for all of the experimental parameters (beam amplitudes and frequencies, magnetic field strengths, camera triggers, and mechanical shutters) to be controlled in real-time, as well as allowing for pre-determined sequences to be loaded.

Chapter 6

Isoplanatism

Whilst conducting the preliminary tests of our optical system, we discovered that significant wavefront distortion occurs when imaging with coherent light. We were eventually led to the conclusion that these distortions were due to a little-recognised effect called non-isoplanatism. Non-isoplanatism can be summarised in the following way. When used with coherent light, even diffraction-limited optical imaging systems are inherently unable to reproduce both the amplitude and the phase of a two-dimensional field distribution because their impulse response function varies slowly from point. For sufficiently small objects, this usually results in a phase distortion and has no impact on the measured intensity, but the intensity distribution can be dramatically distorted when imaging objects of large extension or of special shapes. Foreseeing that an increasing number of experiments and technological applications now call for imaging with coherent light, we thought it interesting to further investigate this effect, and published a paper on the topic^[39]. The first au-

thor of this paper, which is reproduced below, was Dr. Edouard Brainis (now at ULB). The problem is illustrated using two simple examples: the pinhole camera and the aberration-free thin lens. The effects predicted by the theoretical analysis are also verified by experimental observations.

6.1 Introduction

Current technology, especially the ability to manufacture aspherical surfaces, makes it possible to design lenses and mirrors that minimize the most important geometrical aberrations. Such optical elements are nearly ideal instruments obeying the laws of Gaussian optics even for far off-axis points and non-paraxial rays. Self-luminous objects or objects illuminated with *incoherent* light can be imaged with high fidelity. The quality of the resulting image is only determined by the resolution of the instrument, which is related to its numerical aperture. The instrument itself is said to be *diffraction-limited* and can be considered as a linear filter for the intensity of light^[48,49].

When imaging objects with *coherent* light, the conditions for accurate image formation are more severe since both the relative amplitudes and the relative phases of the object points have to be mapped to the corresponding image points (up to the resolution capability of the instrument). This only happens if the effect of the optical instrument on the field from a given point source is independent of its position in the object plane, or in other words, if the coherent impulse response^[48,49] of

the system is *space-independent*. The instrument then acts as a *linear filter* for the complex field amplitude. According to the terminology of^[48], such an instrument is said to be *isoplanatic*. In most cases, even aberration-free optical instruments designed to map some planar object to an image plane under incoherent illumination do not meet this condition. Some spatial phase distortion is unavoidably introduced, which, when combined with a finite resolution, severely modifies the intensity distribution of the image. This was first recognised by Dumontet^[50], and later by Tichenor and Goodman^[51], who showed that a thin lens can only be considered as an isoplanatic system if both the object and the image lie on spherical surfaces \mathcal{S}_o and \mathcal{S}_i , which are tangent to the geometrical-optics object and image planes \mathcal{O} and \mathcal{I} respectively, and have their centre of curvature in the plane of the lens (see Sec. 6.3).

In practice, an aberration-free optical instrument can be treated as an ideal imaging system with coherent light whenever the spherical surfaces \mathcal{S}_o and \mathcal{S}_i can be approximated by their tangent planes \mathcal{O} and \mathcal{I} . We emphasise that this is only viable when the object to be imaged is very small and lies close to the optical axis. A weaker imaging condition has been obtained by Tichenor and Goodman, who showed that the non-isoplanatism of a thin lens has a negligible effect on the *intensity distribution* of the image if the object diameter is smaller than about a quarter of the lens diameter^[51]. In this paper, we investigate the effect of non-isoplanatism on coherent image formation when this condition is not met, and explain, both mathematically and physically, the image distortions that are then observed. These

distortions may be encountered in many fields of optics where large-sized objects are imaged through powerful limited-aperture instruments, as in coherent far-field microscopy, optical lithography^[52], holographic data-storage^[53], and dipole-trapping of neutral atoms^[54]. In particular, arrays of coherently emitting point sources – like individual trapped atoms excited by the same laser beam^[55,56] – will be subject to the phenomena discussed here. Effects similar to those that we described in the context of imaging are also expected when lenses are used to perform the spatial Fourier transform of a two-dimensional field distribution, as in holographic dipole-trapping of atoms^[20].

6.2 Non-isoplanatism of a pinhole camera

The problem of non-isoplanatism only originates in the propagation of light. It is not related to refraction or aberration in the imaging system, regardless of its nature. Because the pinhole camera is the simplest imaging instrument, we use it to introduce the main features of non-isoplanatism. The case of the lens, which is more important in practice but presents some additional subtleties, will be analysed in more details in the following sections.

Consider the following situation (Fig. 6.1): Two mutually coherent point-objects, P_1 and P_2 , are imaged through a small pinhole (of radius a) from an object plane \mathcal{O} to an image plane \mathcal{I} . We call $(x_{o,k}, y_{o,k})$ the transverse coordinates of the point sources P_k , and $(x_{i,k}, y_{i,k})$ the transverse coordinates of their geometrical images P'_k .

For simplicity, we choose $y_{o,k} = y_{i,k} = 0$ and assume that the point sources emit in phase. We also assume that the pinhole is so small that the resolution is limited by far-field diffraction. The “images” of P_1 and P_2 are two Airy patterns centred on the geometrical image points P'_1 and P'_2 . We choose the distance between P'_1 and P'_2 to be equal to the resolution limit, i.e. $x_{i,1} - x_{i,2} = 0.61\lambda z_i/a$. For any given point on \mathcal{I} , the relative phase of the fields originating from P_1 and P_2 is $\phi = 2\pi(r_1 - r_2)/\lambda$. This relative phase ϕ varies as a function of $x_{o,m} \equiv (x_{o,1} + x_{o,2})/2$, the mean distance of the point sources from the optical axis. Therefore, the Airy patterns centered on P'_1 and P'_2 will interfere differently, depending on the location of the P_1 - P_2 pair in the object plane. The phase ϕ is obviously null if $x_m = 0$, but it reaches the value $\pi/2$ when $|x_{o,m}|$ is as small as $a/2.44$. Since a is usually about 1 mm for a pinhole camera, ϕ varies extremely rapidly when the two point sources move away from the optical axis.

The preceding example shows that *different intensity distributions must be expected from identical object patterns depending on their positions in the object plane*. This is due to an incorrect phase mapping between the object and image plane. We refer to this situation as the *non-isoplanatism* of coherent imaging. It is caused by the space variance of the impulse response function of the camera. Note that any linear optical system shows this behaviour^[49].

The complex field amplitudes in the object and image planes, $U_o(x_o, y_o)$ and

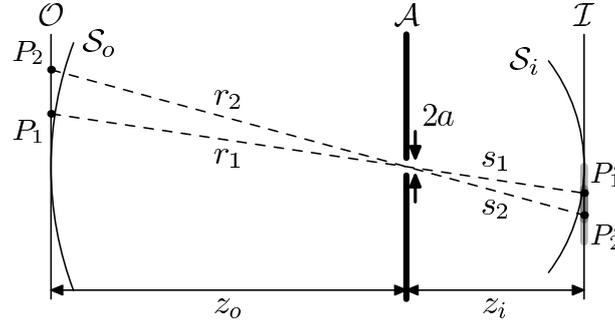


Figure 6.1 Scheme for coherent image formation by a pinhole camera. The points P'_1 and P'_2 in the image plane \mathcal{I} are the geometrical images of the point objects P_1 and P_2 in the object plane \mathcal{O} . The light-gray zones around P'_1 and P'_2 in the \mathcal{I} -plane represent the individual diffraction patterns resulting from illumination of the pinhole by P_1 and P_2 respectively. The dark-gray zone represents the area where the individual diffraction patterns overlap and interfere.

$U_i(x_i, y_i)$, satisfy the integral relation

$$U_i(x_i, y_i) = \iint h(x_i, y_i | x_o, y_o) U_o(x_o, y_o) dx_o dy_o. \quad (6.1)$$

The impulse response of the camera (in the far-field approximation) is given by

$$h(x_i, y_i | x_o, y_o) = |M| \exp(i \frac{2\pi}{\lambda} (r + s)) \times \delta_{\frac{a}{\lambda s}}(x_i - Mx_o, y_i - My_o), \quad (6.2)$$

where (x_o, y_o) are the coordinates of the point source in the object plane, (x_i, y_i) are the coordinates of the “observation point” in the image plane, $M \approx -z_i/z_o$ is the geometrical magnification ratio, and

$$\begin{aligned}
 r &= \sqrt{x_o^2 + y_o^2 + z_o^2}, \\
 s &= \sqrt{x_i^2 + y_i^2 + z_i^2}
 \end{aligned}$$

are the distances from the point source and the observation point respectively to the pinhole center. The point spread function introduced in Eq. (6.2) is given by

$$\delta_N(x, y) = |N| \frac{J_1(2\pi N \sqrt{x^2 + y^2})}{\sqrt{x^2 + y^2}} \quad (N \in \mathbb{R}_0), \quad (6.3)$$

where $J_1(x)$ is a Bessel function of first kind. Its square, $\delta_N^2(x, y)$, is the usual Airy pattern of diffraction-limited optical systems. It is normalised so that $\lim_{|N| \rightarrow \infty} \delta_N(x, y) = \delta(x, y)$ (the Dirac distribution).

The phase factor in Eq. (6.2) plays an important role. Without it, the impulse response function corresponding to any point (x_o, y_o) lying in the object plane would be the same as the impulse response function of the origin $(0, 0)$, but translated to the geometrical image point (Mx_o, My_o) . In that case, the system would be space-invariant, or *isoplanatic*, and the image of a coherent object with field amplitude $U_o(x_o, y_o)$ could be computed by simply convoluting $U_o(x_o, y_o)$ with the impulse response function. In other words, the pinhole camera would act as a linear filter that reproduces the object $U_o(x_o, y_o)$ with a lower resolution. The example analysed in the beginning of this section clearly shows that this is not the case. Because of

the phase factor, the impulse response function is different for different points in the object plane. This creates interference patterns that modify the image more significantly than a simple blur. The image of a field profile $U_o(x_o, y_o)$ is given by

$$U_i(x_i, y_i) = \frac{1}{|M|} \exp(i \frac{2\pi}{\lambda} s) \iint U_o(x_o, y_o) \exp(i \frac{2\pi}{\lambda} r) \delta_{\frac{a}{\lambda z_o}}(\frac{x_i}{M} - x_o, \frac{y_i}{M} - y_o) dx_o dy_o, \quad (6.4)$$

where we made use of the fact that

$$\delta_{\frac{a}{\lambda z_i}}(x_i - Mx_o, y_i - My_o) = \frac{1}{M^2} \delta_{\frac{a}{\lambda z_o}}(\frac{x_i}{M} - x_o, \frac{y_i}{M} - y_o).$$

Seen as a function of (x_i, y_i) , the Airy pattern has a first-zero full-width $e_i = 1.22\lambda z_i/a$ that represents the region of the image plane most influenced by the field originating from the point source at (x_o, y_o) in the object plane. Seen as a function of (x_o, y_o) , its first-zero full-width $e_o = 1.22\lambda z_o/a$ represents the area of the object plane that most contributes to the field at the point (x_i, y_i) in the image plane. In the *non-isoplanatic* situation described by Eq. (6.4), important interferences occur when the phase factor $\exp[i \frac{2\pi}{\lambda} r]$ cannot be considered as constant over the circular region of area $\pi(e_o/2)^2$ centered on $(x_i/M, y_i/M)$ in the object plane.

Isoplanatism is recovered when one considers the imaging problem from \mathcal{S}_o to \mathcal{S}_i (see Fig. 6.1), where \mathcal{S}_o (\mathcal{S}_i) is a spherical surface of radius z_o (z_i) centred

on the pinhole, and having its vertex on \mathcal{O} (\mathcal{I}). The spherical field distributions on \mathcal{S}_o and \mathcal{S}_i are described by the amplitudes $U_{S_o}(x_o, y_o)$ and $U_{S_i}(x_i, y_i)$ respectively (z -coordinates are dependent variables). For slowly varying fields, these are related to the field amplitudes on \mathcal{O} and \mathcal{I} by a phase factor: $U_{S_o}(x_o, y_o) = U_o(x_o, y_o) \exp[i2\pi(r - z_o)/\lambda]$ and $U_{S_i}(x_i, y_i) = U_i(x_i, y_i) \exp[-i2\pi(s - z_i)/\lambda]$. Inserting these expressions into Eq. (6.4) and removing the constant phase factor $\exp[i2\pi(z_o + z_i)/\lambda]$, one gets:

$$U_{S_i}(x_i, y_i) = \frac{1}{|M|} \iint U_{S_o}(x_o, y_o) \delta_{\frac{a}{\lambda z_o}}\left(\frac{x_i}{M} - x_o, \frac{y_i}{M} - y_o\right) dx_o dy_o, \quad (6.5)$$

which is a convolution relation, as expected for an isoplanatic system. The physical reason why isoplanatism is recovered when the object lies on the spherical surface \mathcal{S}_o is simple to understand: since all the point sources are at the same distance z_o from the pinhole, their point spread functions always interfere constructively.

6.3 Non-isoplanatism of a thin lens

The scenario exhibited in the example of the pinhole camera occurs in all optical imaging systems. When imaging with lenses or mirrors, the diffraction effects are less dramatic than with a pinhole camera. However, the problem of phase distortion remains, and may sometimes induce unwanted intensity modulation in the image.

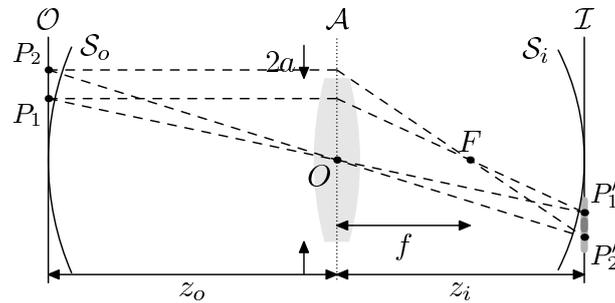


Figure 6.2 Scheme of coherent image formation by an aberration-free thin lens. The points P'_1 and P'_2 in the image plane \mathcal{I} are the geometrical images of the point objects P_1 and P_2 in the object plane \mathcal{O} . The light-gray zones around P'_1 and P'_2 in the \mathcal{I} -plane represent the individual diffraction patterns resulting from illumination of the pinhole by P_1 and P_2 respectively. The dark-gray zone represents the area where the individual diffraction patterns overlap and interfere.

An aberration-free thin lens is an ideal diffraction-limited optical system that acts locally on the impinging field as a pure phase transparency $T(x, y) = \exp[-i\pi(x^2 + y^2)/(\lambda f)]$, where f is the focal length of the lens. It turns out that the impulse response of a thin lens is given by the same formula as for the pinhole camera — Eq. (6.2) — with the understanding that z_i is now related to z_o through the lens formula $1/z_o + 1/z_i = 1/f$. To stress the analogy between the pinhole camera and the thin lens systems, we have drawn on Fig. 6.2 the same information as on Fig. 6.1, and used the same notation. Note that $2a$ now represents the diameter of the lens, which is usually considerably larger than the aperture of the pinhole camera. Therefore, diffraction effects will be weaker, and non-isoplanatism less pronounced. Apart from this comment, the results and discussion of the last section also apply to the thin lens system. As with the pinhole camera, isoplanatism can be recovered when the thin lens is used to image the spherical surface \mathcal{S}_o to \mathcal{S}_i [see Eq. (6.5)]. However, this is less obvious here because, in contrast with the pinhole camera, a lens does

not have an infinite field of view. That \mathcal{S}_i is the image surface of \mathcal{S}_o is therefore questionable (see Sec. 6.5).

Whether non-isoplanatism leads to interference when a plane emitter is imaged to a plane receptor depends not only on the size of the lens, but also on the size of the object and its position with respect to the optical axis. The condition for avoiding any interferences due to non-isoplanatism has been derived in^[51] for one-dimensional imaging. Extended to the two-dimensional case (see Appendix A), it reads:

$$e_o \left(\rho_o^g + \frac{e_o}{4} \right) \ll \lambda z_o, \quad (6.6)$$

where $\rho_o^g = \sqrt{(x_i/M)^2 + (y_i/M)^2}$ is the off-axis distance of the geometrical object point under consideration and $e_o = 1.22\lambda z_o/a$ is the first-zero full width of the Airy pattern in the object plane. The main conclusions are:

- When imaging points very close to the optical axis ($\rho_o^g \ll e_o/4$), interference due to non-isoplanatism does not occur if $a^2/\lambda z_o \gg 1$. This last condition is always satisfied in practice with lenses. However, for a pinhole camera it fails to be satisfied.
- When imaging off-axis points ($\rho_o^g \gg e_o/4$), interference due to non-isoplanatism does not occur if $\rho_o^g \ll a$, i.e. if the object points are not as far off-axis as the edges of the lens.

When criterion (6.6) is satisfied for any point of the object, the phase factor

$\exp[i2\pi r/\lambda]$ can be moved out of the integral in Eq. (6.4), giving

$$U_i(x_i, y_i) = \frac{1}{|M|} \exp(i\pi \frac{x_i^2 + y_i^2}{\lambda(z_i - f)}) \iint U_o(x_o, y_o) \delta_{\frac{a}{\lambda z_o}}(\frac{x_i}{M} - x_o, \frac{y_i}{M} - y_o) dx_o dy_o, \quad (6.7)$$

where the phase factor comes from the second order approximation of $s(x_i, y_i) + r(x_i/M, y_i/M)$ and constant phases have been removed. An intensity detector in the plane \mathcal{I} will record the same image as in the isoplanatic case — Eq. (6.5). It should however be noted that the impulse response function is still *non-isoplanatic* because the phase curvature has not been removed. It is important to keep this in mind when a phase-sensitive detector (hologram) is used or if further optical processing is applied.

6.4 Non-isoplanatism & large field imaging

In practical applications of coherent imaging, the assumption is usually made^[49] that the field mapping from the object space to the image space *is* given by Eq. (6.7) for any lens of a given optical system, i.e. the system is isoplanatic. This assumption is very convenient from a theoretical point of view since it makes the analysis of optical systems much simpler by the use of standard Fourier Optics methods. In this section, we show that the slight difference between Eqs. (6.7) and (6.4) may lead to strong distortions if the object does not fulfill condition (6.6). We explain and interpret these effects.

Let's first point out what is wrong with Eq. (6.7) from a physical point of view. Consider an object field $U_o(x_o, y_o)$ that is varying slowly on the length scale $e_o = 1.22\lambda z_o/a$ of the peaked function $\delta_{a/(\lambda z_o)}(x_i/M - x_o, y_i/M - y_o)$. Then, according to Eq. (6.7), one can write

$$|U_i(x_i, y_i)|^2 = \frac{1}{|M|^2} \left| U_o\left(\frac{x_i}{M}, \frac{y_i}{M}\right) \right|^2. \quad (6.8)$$

This equation implies that the energy is conserved: $\iint |U_i(x_i, y_i)|^2 dx_i dy_i = \iint |U_o(x_o, y_o)|^2 dx_o dy_o$. However, regardless of how slowly the field varies in space, Eq. (6.8) cannot hold for far off-axis points when the lens has a limited aperture. Radiation from far off-axis points (like P_2 on Fig. 6.2) is partially lost, and energy cannot be conserved. For instance, in the case of a plane wave travelling along the optical axis, radiation from the neighbourhood of P_2 will not be transmitted at all. Taking the phase factor $\exp(i2\pi r/\lambda)$ out of the integral in Eq. (6.4) is equivalent to ignoring these losses.

To get a deeper insight into that effect, we use Eq. (6.4) instead of Eq. (6.7) to compute the intensity distribution in the image plane for a *slowly varying* object field. We now obtain

$$|U_i(x_i, y_i)|^2 = \frac{1}{|M|^2} \left| U_o\left(\frac{x_i}{M}, \frac{y_i}{M}\right) \right|^2 \times \left| \iint \exp\left(i\frac{2\pi}{\lambda z_o} \left(\frac{x_i}{M}\xi + \frac{y_i}{M}\eta + \frac{\xi^2}{2} + \frac{\eta^2}{2}\right)\right) \delta_{\frac{a}{\lambda z_o}}(\xi, \eta) d\xi d\eta \right|^2, \quad (6.9)$$

where $\xi = x_i/M - x_o$ and $\eta = y_i/M - y_o$. The integral in Eq. (6.9) has the form of a *Fresnel diffraction integral*. For points (x_i, y_i) which are sufficiently far away from the optical axis, the quadratic phase terms $\xi^2/2$ and $\eta^2/2$ can, as a first order approximation, be neglected, and the integral reduces to the Fourier transform of the Airy pattern in the object space. Using

$$\iint \exp(i2\pi(x\xi + y\eta)) \delta_N(\xi, \eta) d\xi d\eta = \text{circ}\left(\frac{x}{N}, \frac{y}{N}\right), \quad (6.10)$$

Eq. (6.9) becomes

$$\begin{aligned} |U_i(x_i, y_i)|^2 &= \frac{1}{|M|^2} \left| U_o\left(\frac{x_i}{M}, \frac{y_i}{M}\right) \right|^2 \\ &\times \text{circ}^2\left(\frac{x_i}{Ma}, \frac{y_i}{Ma}\right). \end{aligned} \quad (6.11)$$

In Eqs. (6.10) and (6.11),

$$\text{circ}(x, y) = \begin{cases} 1 & \text{if } \sqrt{x^2 + y^2} < 1, \\ 0 & \text{if } \sqrt{x^2 + y^2} > 1. \end{cases} \quad (6.12)$$

Eq. (6.11) shows that the intensity in the image plane exhibits a cut-off. No intensity reaches the image plane at a distance higher than Ma from the optical axis. This can be understood in the following way: since the field has been assumed to be slowly varying, the diffraction in the propagation from the object plane to the

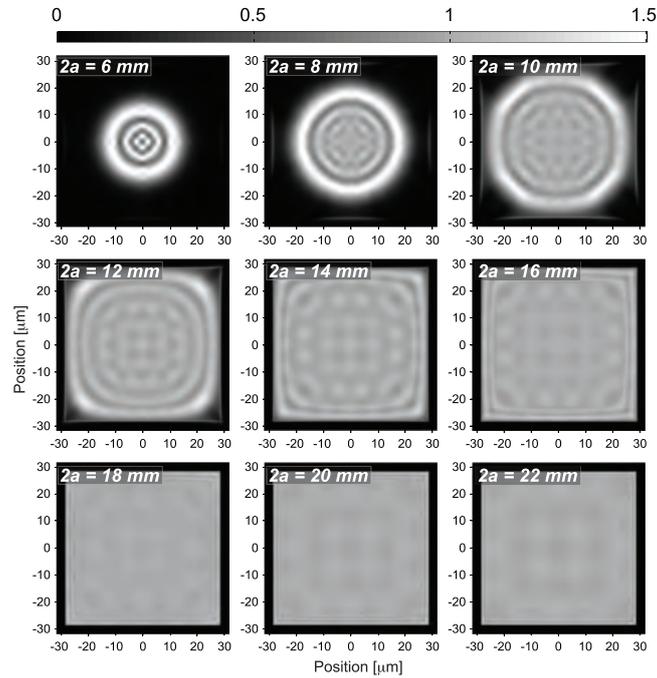


Figure 6.3 Simulation of the intensity in the image plane \mathcal{I} for non-isoplanatic imaging of a large square aperture, centred on the optical axis, with a thin lens. The square aperture is illuminated by a plane wave travelling along the optical axis. The sides of the square are $b = 9.5$ mm. The different panels correspond to different lens diameters. From left to right and top to bottom, the lens diameter is progressively increased in 2 mm-steps from $2a = 6$ mm to $2a = 22$ mm. The relevant parameters are: $\lambda = 780$ nm, $f = 12$ mm, and $z_o = 2$ m. The scale of the figures is expressed in microns. The intensity scale is such that 1 corresponds to the expected uniform intensity in the centre of the square in the limit of an infinite lens.

lens is negligible and the limited aperture of the lens has the same effect as a stop of radius a in the object plane. The function $\text{circ}(x_i/(Ma), y_i/(Ma))$ is the image of that virtual stop and can be interpreted as the shadow of the lens. This is, however, only a first order approximation, since the quadratic phase terms $\xi^2/2$ and $\eta^2/2$ in Eq. (6.9) have been neglected. The effect of these quadratic phase terms is to create radial intensity oscillations in the image.

In Fig. 6.3, the preceding discussion is illustrated with an example. A square object of size $9.5 \times 9.5 \text{ mm}^2$ is imaged using a lens of diameter $2a$ varying from 6 mm to 22 mm. The object field $U_o(x_o, y_o)$ is 1 inside the square and zero outside. Fig. 6.3 shows the intensity distribution in the image plane computed using Eq. (6.4). With small lenses ($2a$ up to 10 mm) the clipping predicted by Eq. (6.11) is observed. One can clearly distinguish the disk (6.12) that limits the observable part of the object, as well as the intensity ripples due to the quadratic phase terms in (6.9) which were neglected when deriving (6.11). Interestingly, the intensity oscillations do not disappear as soon as the lens becomes bigger than the object. For $2a = 22 \text{ mm}$, some residual modulation still remains. We use realistic parameters in the simulation: the object distance z_o has been fixed to 2 m (close to infinite-conjugate ratio imaging) and the focal length f to 12 mm (i.e. the numerical aperture ranges from 0.24 to 0.67). Obviously, energy is lost when imaging large objects through small lenses: For the simulations shown in Fig. 6.3, the percentage of transmitted energy is, from left to right and top to bottom, 31.1%, 56.1%, 83.4%, 96.6%, 99.2%, 99.8%, and nearly 100% for the last three images. For *quickly varying* fields the previous discussion does not hold, but non-isoplanatism still has surprising effects on imaging. No general conclusions can be drawn in this case. Therefore, to get some insight, consider the following instructive example.

A large square grid ($19 \times 19 \text{ mm}^2$) of mutually coherent point sources is imaged through a lens of diameter $2a = 8 \text{ mm}$. The imaging conditions are otherwise the

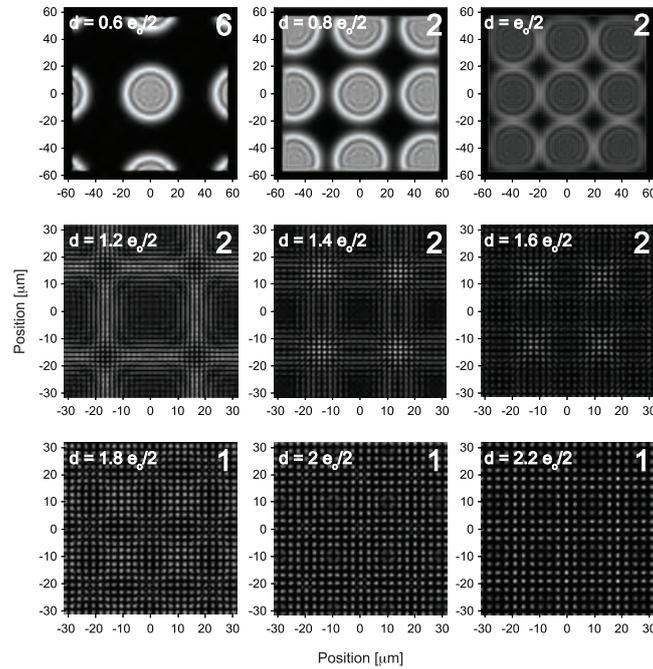


Figure 6.4 Simulation of the intensity in the image plane \mathcal{I} for non-isoplanatic imaging of a large square ($19 \times 19 \text{ mm}^2$) grid of point sources, centred on the optical axis, by a thin lens. The point sources are mutually coherent and in phase. The different panels correspond to different values of the grid period d . From left to right and top to bottom, d is varied from $0.6 e_o/2$ to $2.2 e_o/2$ ($e_o/2$ is the separation corresponding to the Rayleigh resolution criterion). The relevant parameters are: $a = 4 \text{ mm}$, $\lambda = 780 \text{ nm}$, $f = 12 \text{ mm}$, and $z_o = 2 \text{ m}$. The scale of the figures is expressed in microns. The intensity scale is arbitrary but the relative dynamical range is indicated by the number in the upper right corner of each panel. For instance, the number “6” in the first panel means that the intensity range is 6 times larger than in the three panels of the last row.

same as in Fig. 6.3. Let’s consider that all the point sources are in phase and discuss the image formation when the spacing d between the point sources is varied.

If the Airy patterns do not overlap ($d \gg e_o/2$), no interference takes place and non-isoplanatism has no effect on the intensity distribution in the image plane. If the Airy patterns overlap strongly ($d \ll e_o/2$), the object field varies slowly in space, and a fringe pattern similar to the one in Fig. 6.3 is expected. New non-isoplanatism effects appear when d is of the same order of magnitude as $e_o/2$. Fig. 6.4 shows the

intensity distribution in the image plane when d ranges from 0.6 to 2.2 times $e_o/2$. For $d \leq e_o/2$, a fringe pattern similar to the one in Fig. 6.3 is seen in the centre of the field. However, periodic replicas of this pattern are also observed. For $d > e_o/2$, the circular fringe patterns intersect each other, but the sparse sampling due to the grid structure of the image makes this structure barely visible (periodicity, however, remains). For clarity, only the central $30 \times 30 \mu\text{m}^2$ region of the image is displayed in the last six panels of Fig. 6.4. The simulations in Fig. 6.4 show that the intensity distribution in the image plane exhibits two distinct pseudo-periods (in both the x and y directions): the small-scale pseudo-period Md due to the grid structure of the object and the large-scale pseudo-period X associated with the ring patterns due to non-isoplanatism. Strictly speaking the image is periodic only if X is an integer multiple of Md . Hereafter, we use this property to deduce the value of X . The object field is modelled as a two-dimensional Dirac comb:

$$U_o(x_o, y_o) = \sum_{n,m} \delta(x_o - nd, y_o - md),$$

where the integers n and m run from $-\infty$ to $+\infty$. Using Eq. (6.4), we then find that

$$|U_i(x_i, y_i)|^2 = \frac{1}{|M|^2} \times \left| \sum_{n,m} \exp(i\phi_{nm}) \delta_{\frac{a}{\lambda z_o}} \left(\frac{x_i}{M} - nd, \frac{y_i}{M} - md \right) \right|^2,$$

with

$$\phi_{nm} = \frac{\pi}{\lambda z_o} (n^2 + m^2) d^2.$$

In order to find X , we require that $|U_i(x_i, y_i)|^2 = |U_i(x_i + X, y_i)|^2$ when X is a multiple of Md . Using eq. (6.13), we have

$$|U_i(x_i + X, y_i)|^2 = \frac{1}{|M|^2} \times \left| \sum_{n,m} \exp(i\phi_{nm}) \delta_{\frac{a}{\lambda z_o}} \left(\frac{x_i + X}{M} - nd, \frac{y_i}{M} - md \right) \right|^2.$$

The argument $(x_i + X)/M - nd$, can be rewritten as $x_i/M - n'd$ with $n' = n - X/(Md) \in \mathbb{Z}$. Replacing the sum over n by a sum over the values of n' , we obtain

$$|U_i(x_i + X, y_i)|^2 = \frac{1}{|M|^2} \times \left| \sum_{n',m} \exp(i\phi_{n'm}) \delta_{\frac{a}{\lambda z_o}} \left(\frac{x_i}{M} - n'd, \frac{y_i}{M} - md \right) \right|^2.$$

where

$$\phi_{n'm} = \frac{\pi}{\lambda z_o} \left[n'^2 + 2n' \frac{X}{Md} + m^2 \right] d^2.$$

One can note that $\phi_{n'm}$ and ϕ_{nm} just differ by a factor $n' \times 2\pi$, and therefore $|U_i(x_i, y_i)|^2 = |U_i(x_i + X, y_i)|^2$, if X is given by

$$X = M \times \frac{z_o \lambda}{d}. \quad (6.13)$$

When X is not a multiple of Md , formula (6.13) is still valid if X is understood as the large-scale pseudo-period. However, the notion of “large-scale pseudo-period” only holds if $X/(Md) \approx a e_o/d^2 \gg 1$ ($e_o = 1.22 \lambda z_o/a$). Since a/d is usually a large number, this condition can still be valid even if d strongly exceeds the distance $e_o/2$ corresponding to the Rayleigh criterion. The periodic patterns displayed in Fig. 6.4 prove that the field in the image plane corresponding to the Airy pattern of a given point in the grid influences the entire image on a length scale much longer than the usually considered “Airy pattern diameter” e_o . Readers familiar with signal processing will notice that the reason why periodicity appears here is the same one that makes the Fourier spectrum of a sampled signal (Dirac comb) periodic. We however stress that the periodicity that is described here appears in the image plane *and not* in the Fourier plane of the object field. This is a peculiarity of non-isoplanatic imaging. There is also an intuitive way of understanding the periodicity displayed in Fig. 6.4. Since we image a two-dimensional grating, only plane waves with discrete inclination angles $(\alpha_n, \beta_m) = \lambda(n\lambda/d, m\lambda/d)$ impinge on the lens. Because the lens has a finite diameter, each plane wave (α_n, β_m) is clipped by the lens and produces a ring pattern distribution similar to the one in Fig. 6.3 but shifted to the position $(x_n, y_m) = (z_i\alpha_n, z_i\beta_m)$. This explains the periodicity X given in Eq. (6.13). The image of the grid can only be observed properly when all the plane waves originating from the object overlap and interfere in the image plane.

In many applications of coherent optics (like optical lithography), the object is

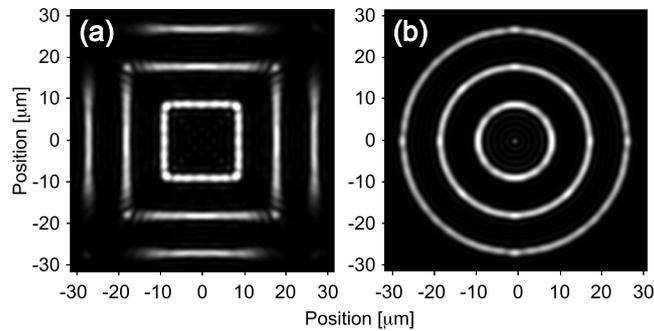


Figure 6.5 Simulation of the intensity in the image plane \mathcal{I} for non-isoplanatic imaging of squared and circular contours by a thin lens. In the object plane, all the point sources forming the contours are mutually coherent and in phase. The imaging conditions are the same as in Fig. 6.3 ($\lambda = 780$ nm, $f = 12$ mm, and $z_o = 2$ m), but the size of the lens is fixed: $2a = 6$ mm, as in the first panel of Fig. 6.3. The intensity scale is arbitrary. Panel (a) shows the image of three squared contours of increasing side $b = a$, $b = 2a$ and $b = 3.2a$. The size of the outer square is the same as in Fig. 6.3. Panel (b) shows the image of three concentric circles, the diameters of which are equal to the sides of the squares in the panel (a).

made of lines instead of single points or filled surfaces. Lines are objects on which the field varies slowly in one direction (the direction tangent to the line) and rapidly in the orthogonal one. Therefore, one can expect that the effects of non-isoplanatism will be intermediate between the two previous cases. Fig. 6.5 helps to understand how lines are imaged through a non-isoplanatic optical system. The imaging conditions are the same as in Fig. 6.3, except that the lens diameter is fixed: $2a = 6$ mm. Panel (a) shows the image of three squared contours. The outer square is exactly the contour of the filled square imaged in Fig. 6.3. The comparison with the upper-left panel of Fig. 6.3 shows that the major part of the contour is now visible; only the corners of the square are clipped. Closer examination shows that, for a straight line, only a segment of length $2a$ in the object plane ($2Ma$ in the image plane) is transmitted. The part of the line that is clipped corresponds to the orthogonal

projection of the shadow of the lens on the straight line. This can be easily understood by analysing the propagation of the cylindrical waves emitted by straight lines through the spherical lens. The smaller square in panel (a) is transmitted because its side is shorter than the diameter of the lens. The intermediate square is at the limit of the cut-off. As shown in the panel (b) of Fig. 6.5, circles are never clipped, whatever their radii, because the orthogonal projection of the lens disk on the circle is the circle itself.

6.5 Isoplanatic imaging through a thin lens

As with the pinhole camera, isoplanatic imaging with a thin lens is possible when the object lies on the spherical surface \mathcal{S}_o and the image is observed on the spherical surface \mathcal{S}_i . The technique for projecting a plane object onto the spherical surface \mathcal{S}_0 has been proposed and demonstrated in^[51] for 1D objects. It consists in placing a convergent thin lens with focal distance z_o immediately before (or after) the object plane. This lens introduces a spherical phase delay in the object field: $U_0 \rightarrow U'_0 = U_0 \times \exp[-i2\pi r/\lambda]$. This extra phase factor compensates the phase accumulated during the propagation from \mathcal{O} to \mathcal{S}_o . Therefore, the field on \mathcal{S}_o is $U_o(x_o, y_o)$ (z_o is now a dependent variable). Similarly, on the image side, a thin lens of focal length z_i placed just after (or before) the image plane can be used to project the image from \mathcal{S}_i onto \mathcal{I} . From a broader point of view, any spherical field distribution in the object space of a centred paraxial optical system can be imaged onto a spherical surface of any curvature using only lenses; some of them will be

imaging lenses, while others will play the role of phase-correction transparencies. This is the basis of the so-called *metaxial optics* theory formulated by Bonnet^[57–59].

Though first proposed for lens imaging, the technique described above is questionable in that case, while fully legitimate for a pinhole camera. The difference between these two imaging systems is that the pinhole camera has an infinite depth of field (the distances z_o and z_i between the pinhole, an object point and its image can be chosen arbitrarily), while for the lens, z_o and z_i are related by the lens law. Therefore, referring to Fig. 6.2, simple Gaussian Optics arguments suggest that if the point P_2 is translated horizontally from the object plane \mathcal{O} to the surface \mathcal{S}_o , its image P'_2 , instead of moving towards \mathcal{S}_i , should move away from the lens. However, this argument relies on the Gaussian approximation that the point P'_2 is initially in the “image plane” \mathcal{I} . In reality, because of the field curvature aberration due to the lens, the stigmatic image of P_2 is closer to the lens than the surface \mathcal{S}_i itself (P'_2 lies on the so-called Petzval surface). Bringing P_2 on \mathcal{S}_o will place its image P'_2 exactly on \mathcal{S}_i ^[60]. Consequently, Eq. (6.5) is exact, while Eq. (6.4) is only valid in the context of Gaussian approximation. If the lens is still diffraction-limited in this regime (no point-aberrations), the isoplanatic imaging geometry provides a nearly perfect transfer of the coherent field from the object to the image space. Only the resolution is reduced because of the finite size of the lens. It should be noted that this approach only works for sufficiently slowly varying fields, because the diffraction from \mathcal{O} to \mathcal{S}_o (and \mathcal{S}_i to \mathcal{I}) has to be negligible for the amplitude of the fields on \mathcal{O}

and \mathcal{S}_o (\mathcal{S}_i and \mathcal{I}) to be the same.

6.6 Experimental investigation

The theoretical discussion of Secs. 6.4 and 6.5 relies on two strong approximations: the paraxial approximation and the thin lens approximation. One may wonder whether our analysis is robust enough to be applied to systems containing powerful lenses, which are usually thick and have a high numerical aperture. The following experiment shows that the previous discussion is also valid for these systems.

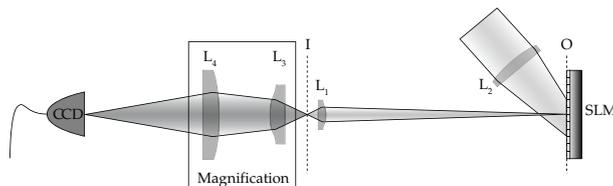


Figure 6.6 Experimental setup: L_1 , aspheric lens ($f = 8$ mm); L_2 , achromatic doublet ($f = 750$ mm); L_3 , aspheric lens ($f = 20$ mm); and L_4 , achromatic doublet ($f = 500$ mm); SLM, spatial light modulator; CCD, coupled-charge camera; O, object plane; I, image plane.

The setup is shown in Fig. 6.6. The test lens L_1 is an aspheric lens having a focal length $f = 8$ mm and a diameter $2a = 8$ mm (LightPath 352240). This lens is used to image the surface of a spatial light modulator (SLM), a 1024×768 micromirror array, that modulates the amplitude of the reflected beam. Using the SLM, we can generate arbitrary patterns. The resolution is set by the size of the micromirrors ($13 \times 13 \mu\text{m}^2$). The object plane is 60 cm away from the lens. We use a double-lens system (L_3, L_4) to magnify the image produced, and project it onto a CCD cam-

era. L_3 is a diffraction-limited aspheric lens and L_4 a long-focal achromatic doublet. The numerical aperture of this double-lens system is large enough to prevent any possible clipping or diffraction during the magnification process. The SLM is either illuminated with a plane wave or a spherical wave converging on L_1 ($\lambda = 780$ nm). In the first case, the imaging system is non-isoplanatic. As explained in Sec. 6.5, it becomes isoplanatic when a spherical-wave illumination is used. The spherical wave is obtained from the impinging plane wave by inserting the additional lens L_2 (75-cm focal length, achromatic doublet) in front of the SLM. We use this setup to verify the theoretical predictions of Sec. 6.4.

Fig. 6.7 shows the pictures recorded by the CCD camera when we image squares of different sizes through the 8-mm diameter aspheric lens L_1 . Let's first consider the case when the SLM is illuminated by a parallel beam (non-isoplanatic imaging, left column of Fig. 6.7). For a square side larger than 8 mm, the clipping effect described in Sec. 6.4 is clearly observed. Circular fringes similar to those of Fig. 6.3 are seen in the image plane. For a square side smaller than 8 mm, the region of non-zero intensity is limited by the size of the square. However, intensity modulation due to non-isoplanatism is still noticeable for a square side as small as 2.7 mm. In the case of the 1.4×1.4 mm² square, non-isoplanatism has a negligible effect. Note that, in the simulations of Fig. 6.3, the lens size was varied while the object size was kept constant. Here, the lens is always the same, but the square size is varied instead. For this reason, the ring pattern is the same for every picture in

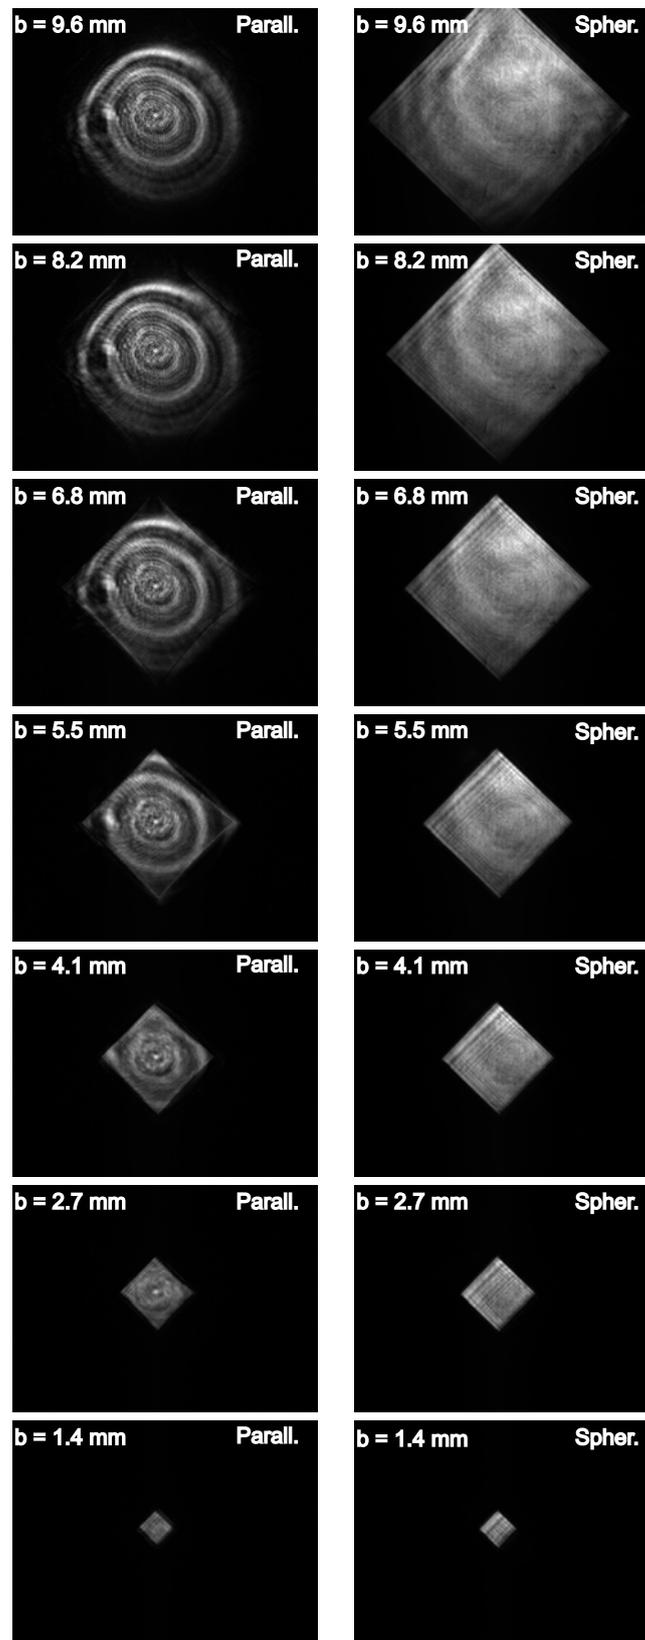


Figure 6.7 Images of different-sized squares (the side b ranges from 1.4 mm to 9.6 mm) recorded using the setup in Fig. 6.6. The object plane is either illuminated with a parallel beam (non-isoplanatic case, left column) or with a spherical one converging on L_1 (isoplanatic case, right column).

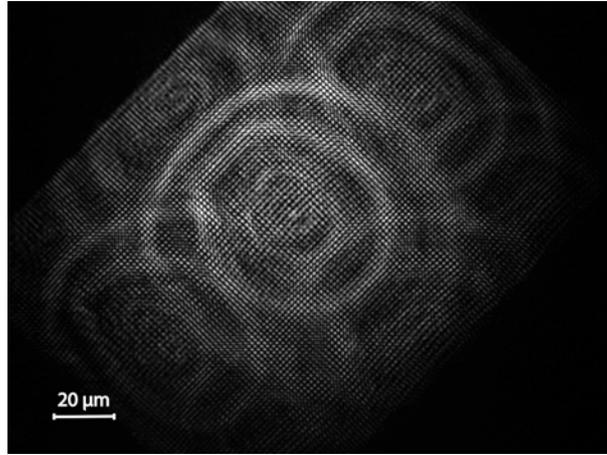


Figure 6.8 Images of a rectangular grid of points recorded using the setup in Fig. 6.6. The object plane is illuminated with a parallel beam (non-isoplanatic case). The period of the grid in the object plane, $d = 109 \mu\text{m}$, corresponds to 1.5 times the Rayleigh criterion separation ($e_o/2$).

the left column of Fig. 6.7. When the SLM is illuminated with a spherical wave converging on L_1 (isoplanatic imaging, right column of Fig. 6.7) no clipping effect occurs and there is, in principle, no limit to the size of the objects that the system can image (the slight variations in intensity are due to amplitude inhomogeneities in the illumination beam).

In a second experiment, we imaged rectangular grids of points illuminated by a plane wave (non-isoplanatic illumination), a situation that we described theoretically in Sec. 6.4. Fig. 6.8 shows the recorded intensity distribution in the image plane for a grid with a period $d = 109 \mu\text{m}$, which corresponds to 1.5 times the Rayleigh criterion separation ($e_o/2$). We distinguish one central ring pattern, similar to the one shown in Fig. 6.7 and four replicas intersecting it. The intensity distribution

is very similar to the simulations displayed in Fig. 6.4. The distance between the centers of the ring patterns is $57 \mu\text{m}$. It is in very good agreement with the value $X = 58 \mu\text{m}$ calculated using formula (6.13). Since the experimental point-spread function is slightly broader than the theoretical Airy pattern, the overlap of the rings is clearly visible for $d = 1.5 e_o/2$, while it is not easy to observe in Fig. 6.4. This is because of the insufficient sampling caused by the sparsity of the imaged grid.

The experiments that we performed show how non-isoplanatism affects image formation for slowly and fast varying fields. We emphasise that these effects are always present when imaging extended objects with small lenses and should not be confused with aberrations.

6.7 Conclusion

We have shown that even diffraction-limited imaging systems distort the phase of the processed fields. This is of no relevance when working with incoherent light, but has a tremendous effect on coherent imaging. In combination with Fraunhofer diffraction from the finite instrumental aperture, the phase distortion that arises leads to a severe degradation of the field amplitude in the image plane. We analysed this phenomenon for the pinhole camera and the thin lens. We showed that substantially different effects arise depending on whether the field varies slowly or rapidly on the length scale of an Airy pattern. However, the degradation of the field amplitude can be overcome, or at least minimised. The main aspects of our analysis have been

confirmed experimentally with a powerful thick aspheric lens, demonstrating that the phenomena that we describe are also present beyond the paraxial and thin lens approximations.

6.8 Derivation of condition (6.6)

The condition for avoiding non-isoplanatism effects is that the phase $2\pi r/\lambda$ in Eq. (6.4) varies by less than π when the point corresponding to coordinates (x_o, y_o) explores the Airy pattern $\delta_{\frac{a}{\lambda z_o}}(x_i/M - x_o, y_i/M - y_o)$ in the object plane. To express this condition mathematically, we rotate the axes to get $y_i = 0$. From the centre of the Airy pattern at $x_o = x_i/M$, the phase $2\pi r/\lambda$ changes most rapidly when x_o moves radially off-axis. The phase change between the centre and the most off-axis first zero of the Airy pattern is $\delta\phi = 2\pi/\lambda(r_1 - r_0)$, where $r_0 \approx z_o(1 + (x_i/M)^2/(2z_o^2))$ and $r_1 \approx z_o(1 + (x_i/M + e_o/2)^2/(2z_o^2))$ in paraxial approximation. Asking that $\delta\phi \ll \pi$ leads to Eq. (6.6).

Chapter 7

Atoms Trapped in Arbitrary Potential Landscapes

The results presented in this Chapter show that atoms can be trapped in single or multiple static traps, and that these traps can have arbitrary shapes. A summary of these results can be found in^[61].

7.1 Preliminary Results

As a first step, we loaded atoms from our MOST into a dipole trap consisting of a single focused Gaussian beam with a $50\mu\text{m}$ waist. An absorption image of the atoms in this dipole trap can be seen in Figure 7.1. The lifetime of the atoms in this trap was of approximately 300ms, and the loading efficiency from the MOST to the dipole trap was 12.5%. The temperature of the atoms was measured by time of

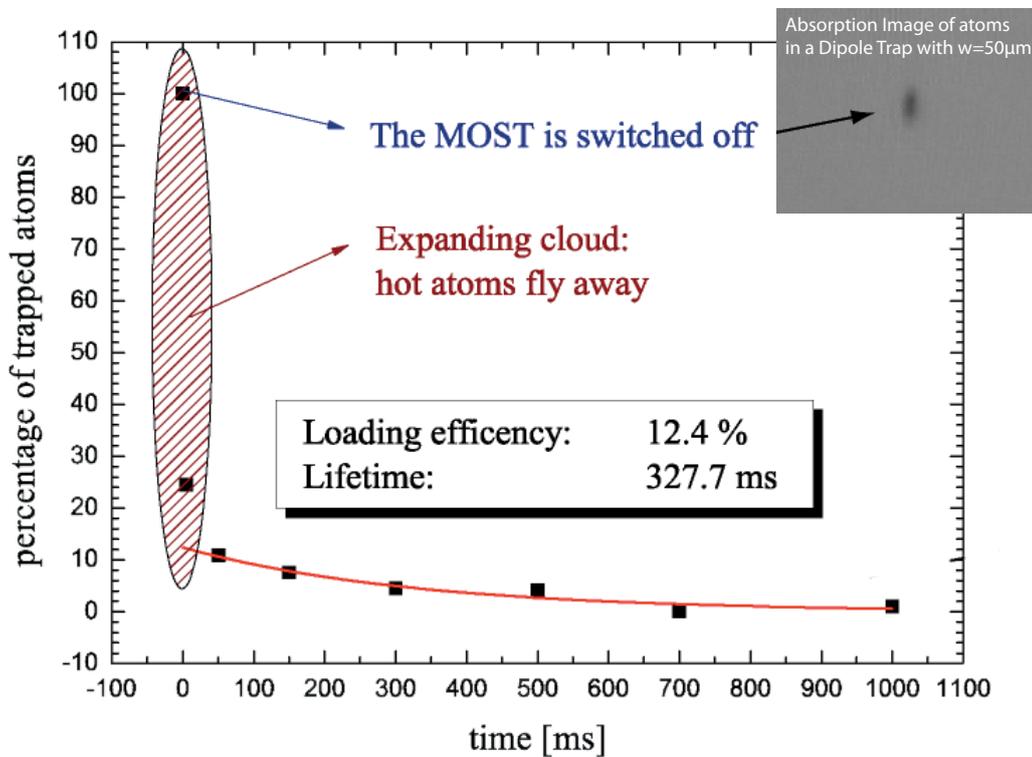


Figure 7.1 Graph of the relative atom number in the $50\mu\text{m}$ dipole trap after being loaded from the MOST. The loading efficiency refers to the percentage of atoms left in the dipole trap with respect to the number of atoms initially in the MOST. The inset shows an absorption image of the atoms in the Dipole Trap.

flight (TOF), and was found to be $42\mu\text{K}$. Generally speaking, in a TOF measurement, the relevant trap is switched off and the atomic cloud is allowed to expand. The velocity of the expanding cloud depends on both the velocity distribution of the atoms within it, which is given by the Maxwell Boltzmann distribution in equation 5.1, and the initial velocity spread in the trap. If the assumption is made that the trapped atoms have a spatial distribution which is approximately Gaussian before expansion, then the velocity of the expanding cloud is given by a convolution of this distribution with the velocity distribution of the expanding atoms. The $1/e$ radius

of the cloud after a time t is thus given as $r(t) = \sqrt{r_0^2 + (v_e t)^2}$, where r_0 is the initial radius, and $v_e = \sqrt{\frac{2k_B T}{m}}$ is the most probable speed at temperature T . Hence, by measuring $r(t)$, the initial temperature of the atoms in the trap can be estimated. The temperature of all the traps presented in this and the following chapter, as well as the temperature of the MOST given in Chapter 3, were measured using this TOF method.

The first observation of atoms trapped in our optical tweezers was made by absorption imaging. Loading atoms from the MOST into the tweezers the first time was not straightforward, as the aspheric lens mount greatly restricted optical access for the MOST beams and the CCD cameras used to monitor the position of the MOST. For this reason, we first tried to load atoms into traps with rather large diameters which could be monitored easily by absorption imaging. The two traps shown in the absorption image in Section 5.12, for instance, were generated by switching on two large rectangles on the DMD (these measured approximately $20\mu\text{m}$ by $50\mu\text{m}$ in the image plane). Nevertheless, absorption imaging is only feasible when the density of trapped atoms is high, so, given that the goal of the project is to eventually trap single atoms, it is clearly not an adequate detection technique. Moreover, the geometry of the trapping setup necessitated that the absorption imaging be perpendicular to the trapping plane, a perspective from which individual traps are not necessarily visible.

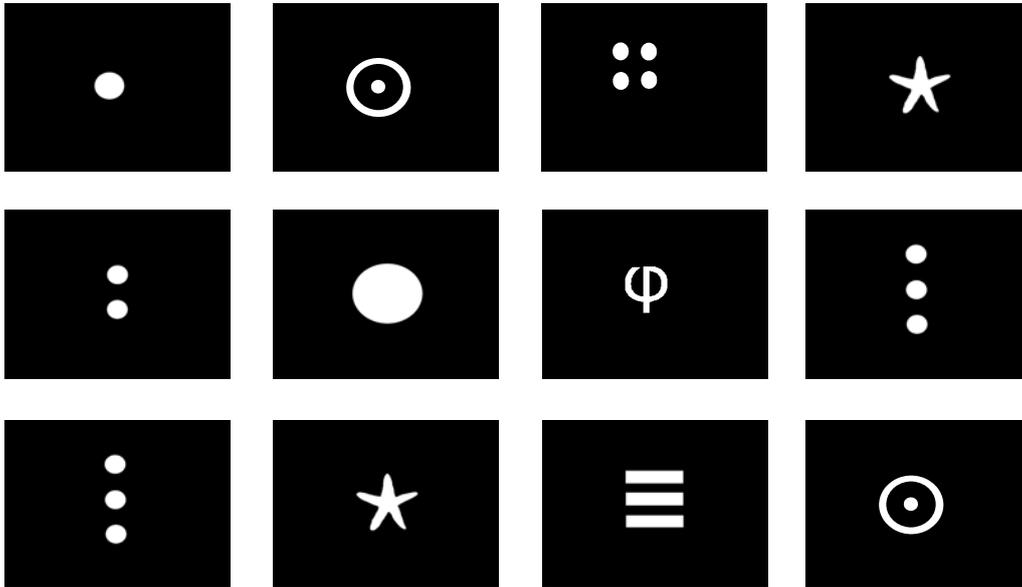


Figure 7.2 Some of the bitmaps used to generate structured trapping geometries on the DMD. These bitmaps form static traps into which atoms are loaded.

The first structured fluorescence images were taken by illuminating the atoms with a single focussed beam having a waist of $25\mu\text{m}$ and $13\mu\text{W}$ of power. We started by loading atoms into large circular traps with radii on the order of $20\mu\text{m}$ to $30\mu\text{m}$. We then moved on to smaller circular traps with radii between $10\mu\text{m}$ and $15\mu\text{m}$, and tried to load atoms into several of these traps simultaneously. Finally, we tried to load atoms into more complicated structures like the phi, the bullseye, the seastar, and the stripes that can be seen in Figure 7.2. Figure 7.3 shows the resulting fluorescence images. The blurry nature of the images is due primarily to fluorescence from out-of-focus atoms lying above and below the plane on which the camera focuses, and, as was explained in section 5.12 was eventually circumvented by using a very thin light sheet instead of a focussed beam. The first light sheet we employed had a

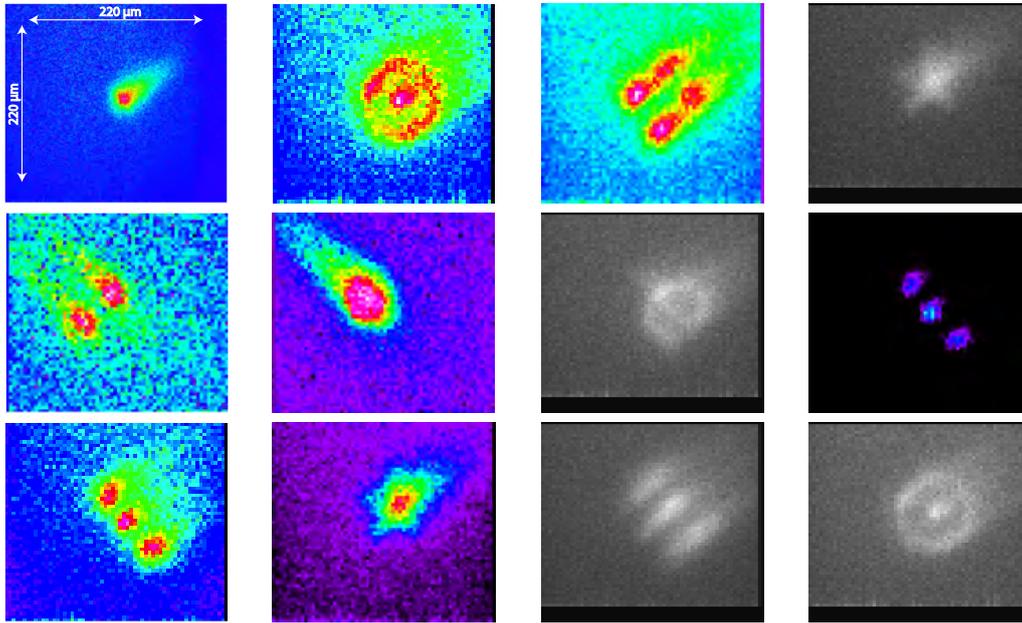


Figure 7.3 These images show atoms in dipole-force trapping potentials formed from the corresponding bitmaps shown in Figure 7.2. The blurry nature of these first fluorescence images can be attributed to fluorescence from atoms above and below the image plane, and to atoms pushed out of the trap by the illuminating light. The colour scale is different for each image, and is undefined because the images were too blurry to conduct an accurate atom-number calibration. The binning was set to 16x16 in all but the first image, which was binned at 4x4.

height of $30\mu\text{m}$ and a Rayleigh length of almost 1mm, and only partially reduced the blur. The clearer pictures in Figure 7.2 were taken with this light sheet. The diffuse “tail” of atoms that can be seen in several of these images arose due to radiation pressure from the illuminating beam, which pushed atoms out of the trap. For this reason, we eventually retro-reflected the light sheet. The inevitably lower power in the back-propagating beam means that the effect can never be fully eliminated, but by adjusting the retro-reflection very precisely, it is possible to reduce the effect significantly.

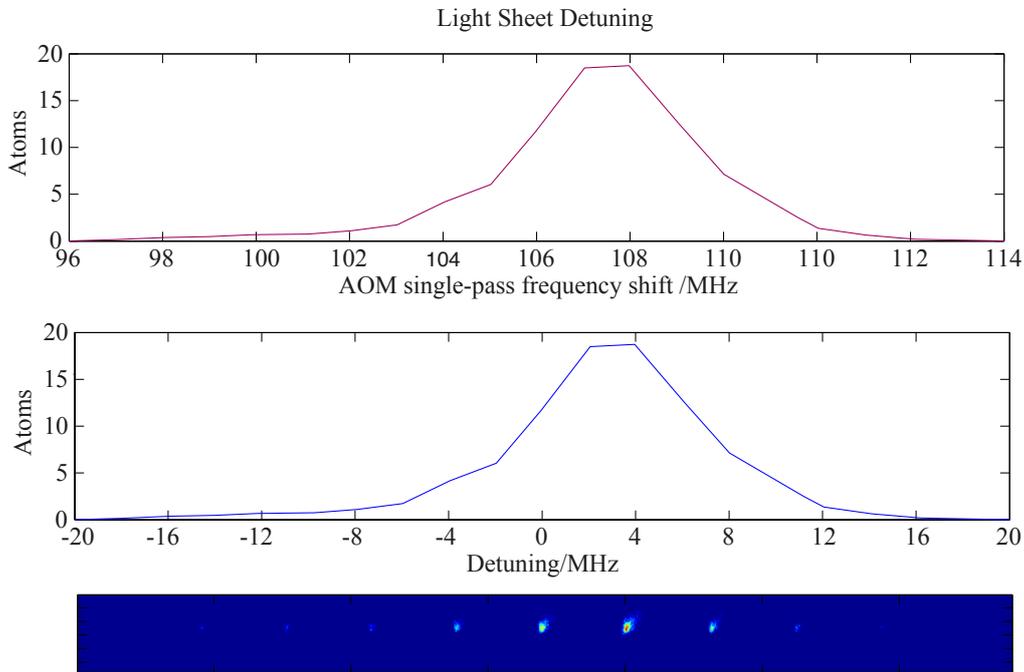


Figure 7.4 The atomic fluorescence plotted as a function of frequency displays a more or less Lorentzian lineshape that peaks at a detuning of +4 MHz from the cycling transition. This detuning gives an experimental measure of the trap depth, as it corresponds to the Stark shift induced on the atomic levels by the presence of the tweezers.

7.2 Ultra-Narrow Light Sheet

The introduction of a much thinner retro-reflected light sheet reduced the blur in the fluorescence images dramatically, and allowed us to see atoms trapped in much smaller tweezers with considerable clarity. It also allowed us to actually characterise the traps with some accuracy. For instance, it would be useless to attempt an atom number calculation with the images shown in Figure 7.3, as it would be impossible to distinguish the atoms constituting the blur from the atoms actually in the trap. As described in Chapter 5, this light sheet has a height of $9.7\mu\text{m}$, a Rayleigh length

on the order of $100\mu\text{m}$, and $7\mu\text{W}$ of power. In all the experiments, it is blue detuned by 2 to 4 MHz from the main cycling transition in order to account for the Stark shift induced by the presence of the tweezers. This detuning was chosen to maximise the atomic fluorescence, but also serves as an experimental measure of the trap depth. Figure 7.4 shows how the collected atomic fluorescence varies as a function of the frequency detuning of the AOM. As expected for homogeneous broadening, the lineshape is more or less Lorentzian. The calculated trap depth corresponds to a value of $\sim 2.5\text{MHz}$, so the measurement gives a fairly good estimate of the trap depth.

The height, focus, and tilt of the light sheet had to be adjusted very precisely for the light sheet to coincide with the image plane of the tweezers. Figure 7.5 illustrates how the fluorescence image becomes sharper and brighter as the light sheet approaches the image plane. The height of the light sheet is adjusted empirically by looking at the fluorescence signal. The retro-reflection of the light sheet is adjusted in a similar way. Figure 7.6 shows the effect of the light sheet on the running MOST when it is not retro-reflected. The power in the light sheet is increased until the radiation pressure has a visible effect on the atoms (this is somewhat higher than the power normally used to illuminate the traps). The light sheet is then retro-reflected (a flip mount blocks normally blocks the retro-reflection, instead allowing access for one of the MOST cameras), and the back-propagating light sheet's height adjusted until the "push" on the atoms is minimised.

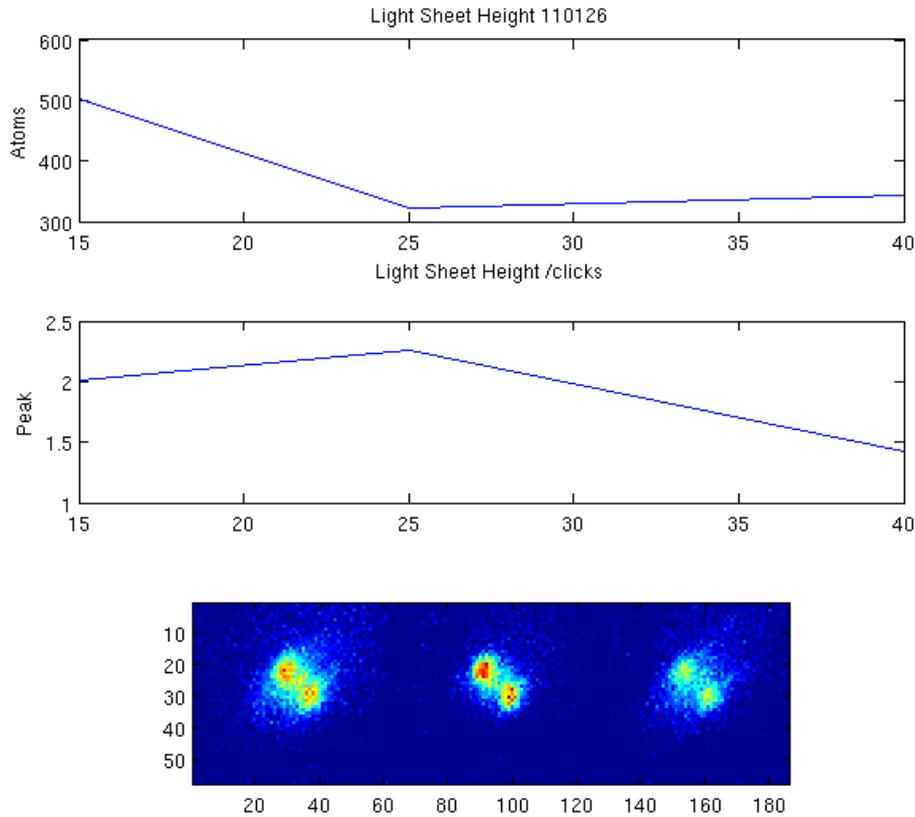


Figure 7.5 The height of the light sheet is shifted until it coincides with the image plane of the tweezers. As the light sheet is moved away from this plane, it becomes harder to resolve the two traps, and the fluorescence signal drops. The “clicks” alluded to in the graph refer to the piezo driven mount holding the mirror which steers the light sheet into the vacuum chamber.

The fluorescence images taken with this thin light sheet reveal that the shape of the atom cloud adapts to the imposed trapping geometry. Figure 7.7 shows atoms confined to a series of circular flat-bottom traps created by displaying disks of different radii on the DMD. This simple trapping geometry allowed us to characterise the traps before moving on to more complex shapes. Normally, one would expect

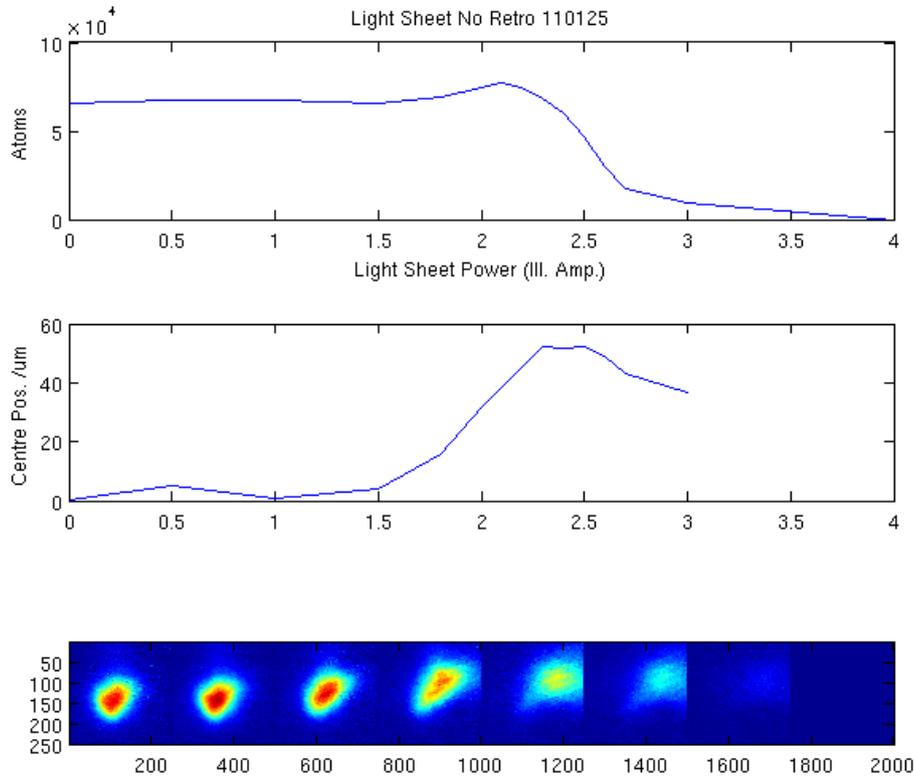


Figure 7.6 When it is not retro-reflected, radiation pressure from the incoming light sheet pushes atoms out of the traps, or in this case, the MOST. The effect is negligible when the power in the light sheet is low, but becomes appreciable as it is increased. The illumination amplitude plotted in the top graph corresponds to the calibrated value used in LabView. The power normally used to illuminate the atoms in the trap ($7\mu\text{W}$) corresponds to a setting of 0.5.

the atoms to fill the trap, such that the radius of the atom cloud equals the trap radius. This is indeed the case for traps of a larger 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. Independent optical characterisation of our imaging system (see Chapter

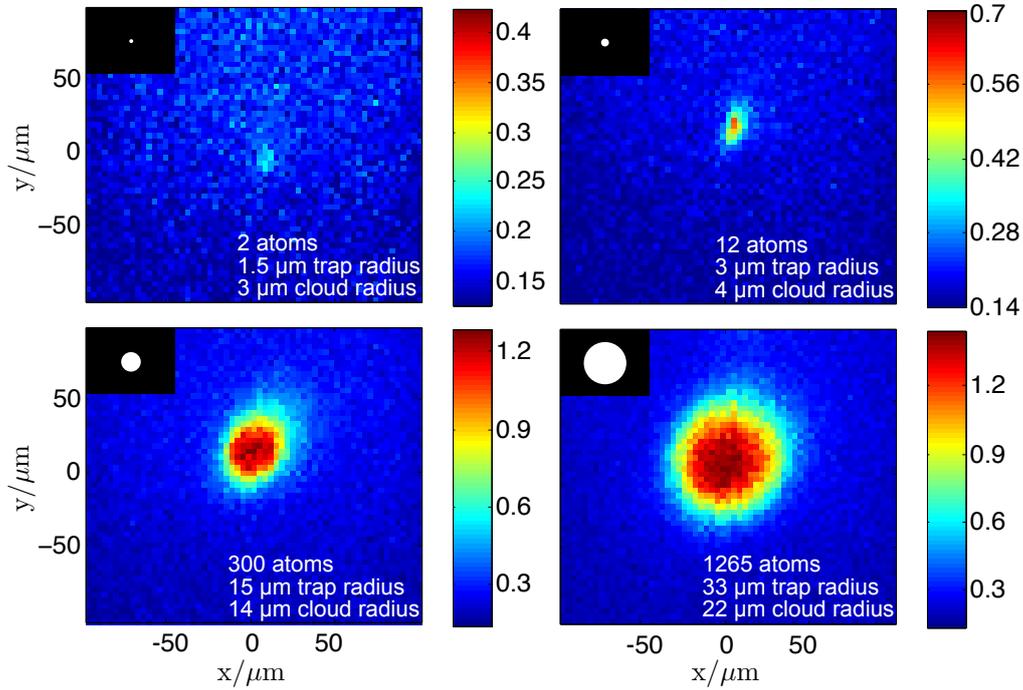


Figure 7.7 Fluorescence images of trapped atoms in single static round traps of different sizes. Each image is an average over 20 consecutive trap loading cycles. The colour scale has units of atoms per μm^2 . The actual patterns on the SLM used to produce the traps are shown in the respective insets using the same length scale.

4) shows that the imaging system has a resolution of $1 \pm 0.2 \mu\text{m}$, which also limits the edge sharpness of any trapping potential, and hence the minimum amount of blur that can be seen in the fluorescence images.

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 $90 \mu\text{s}$, within which a single atom emits ~ 1400 photons; of these, 21 photons are detected. This corresponds to an overall photon-detection efficiency of $\eta_{tot} = (1.5 \pm 0.2)\%$ (see section 7.3). We found that the trap with $d=65 \mu\text{m}$ con-

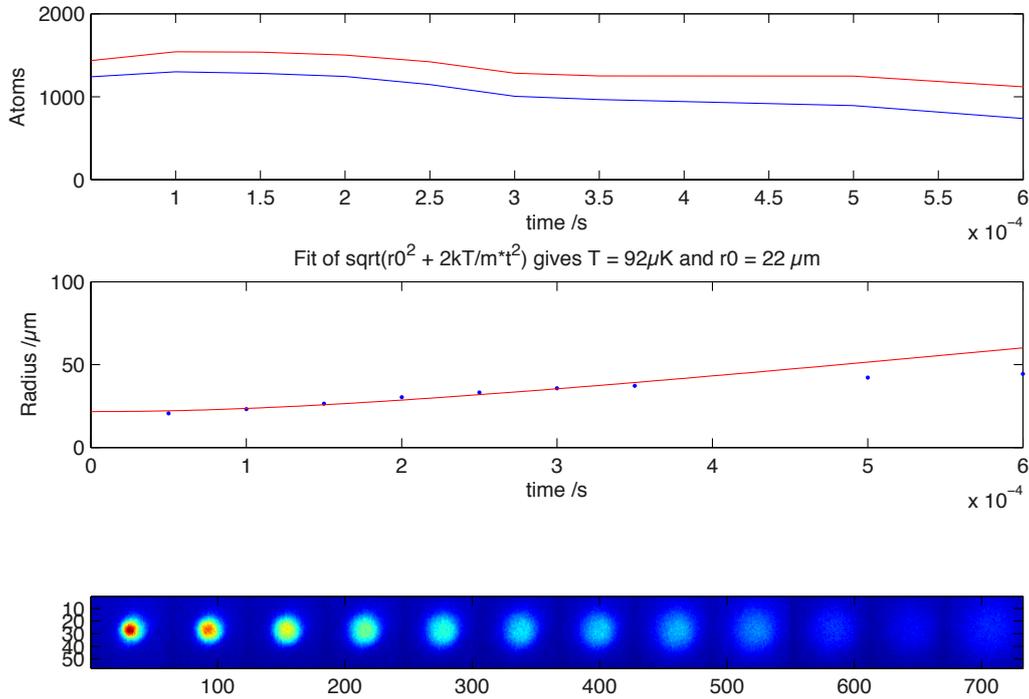


Figure 7.8 Time-of-flight measurement for a trap with $r=22\mu\text{m}$ ($d=44\mu\text{m}$).

tained 1270 ± 170 atoms, the trap with $d=30\mu$ contained 300 ± 40 atoms, the trap with $d=6\mu\text{m}$ contained 12 ± 2 atoms, and the smallest observed trap with $d=3\mu\text{m}$ contained 2 ± 0.3 atoms. This is very close to the single-atom regime, which could be reached by application of the collisional blockade technique described in Chapter 4^[12,41].

For all traps with a radius $\geq 3\mu\text{m}$, the atom number was sufficiently large for time-of-flight measurements, which yielded a consistent temperature of around $90\pm 10\mu\text{K}$. This is on the same scale as the calculated trap depth. Figures 7.8 and 7.9 show time-of-flight measurements for a trap with $d=44\mu\text{m}$ and a trap with $d=30\mu\text{m}$

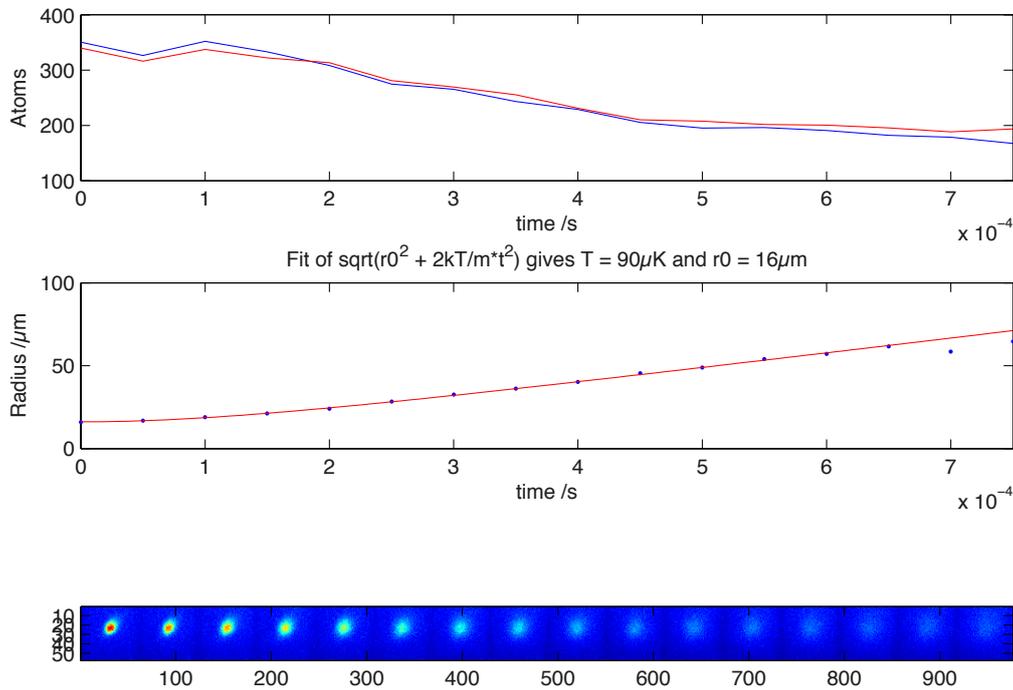


Figure 7.9 Time-of-flight measurement for a trap with $r=15\mu\text{m}$ ($d=30\mu\text{m}$).

respectively. Figure 7.10 shows how this measurement breaks down for a trap with $d=6\mu\text{m}$. The fit yields an initial radius r_0 of $9\mu\text{m}$, and a temperature of $180\mu\text{K}$. This inaccuracy could once again be due to the fact that the contribution from out-of-focus atoms trapped away from the image plane makes for a proportionally larger margin of error in measuring the waist of very small traps. Another explanation is that the smaller traps exhibit a higher temperature because they have a much shorter Rayleigh length, and a shorter Rayleigh length means that it is harder to load the atoms at the right height along z , i.e., at the height that coincides with the bottom of the dipole trapping potential. Additionally, molasses cooling during the loading phase does not work as well for smaller traps.

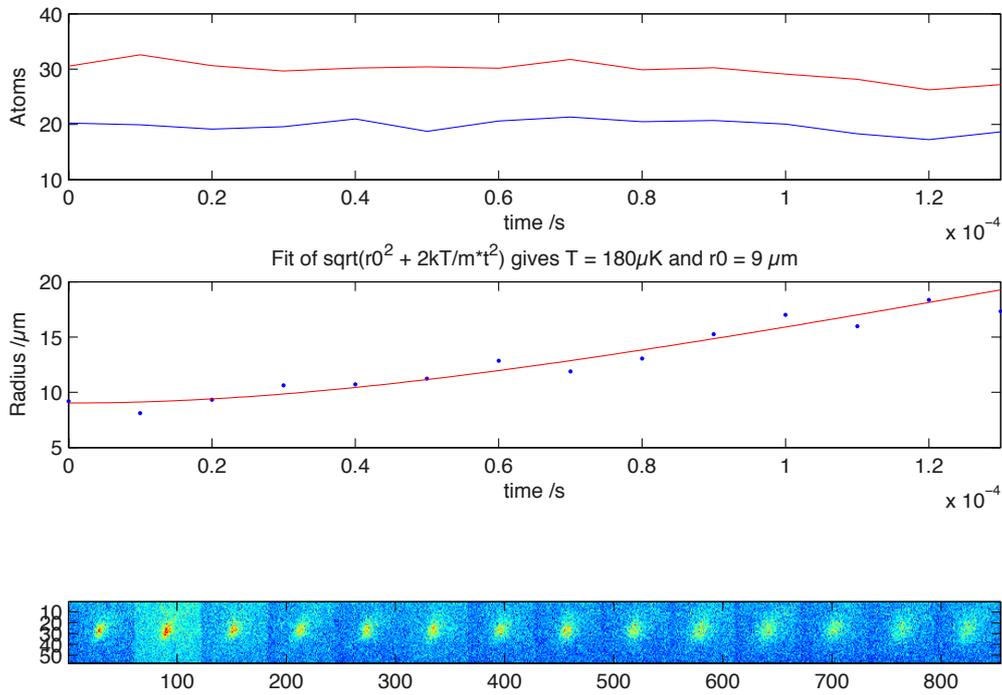


Figure 7.10 Time-of-flight measurement for a trap with $r=3\mu\text{m}$ ($d=6\mu\text{m}$). For traps this size, the measurement is no longer accurate.

The lifetime of atoms in all traps is approximately 50ms, and does not change significantly with the size of the trap. The specific lifetimes are 44ms for the trap with $d=6\mu\text{m}$, 58.8ms for the trap with $d=30\mu\text{m}$, and 37.6 for the trap with $d=65\mu\text{m}$. These lifetimes were found by fitting an exponential decay to the incrementally decreasing atom number within the trap region. The extent of this region was determined either by a centre-of-mass calculation or by a Gaussian fit, with similar results. Figures 7.11 - 7.13 show a graph of the lifetime for each of these three traps.

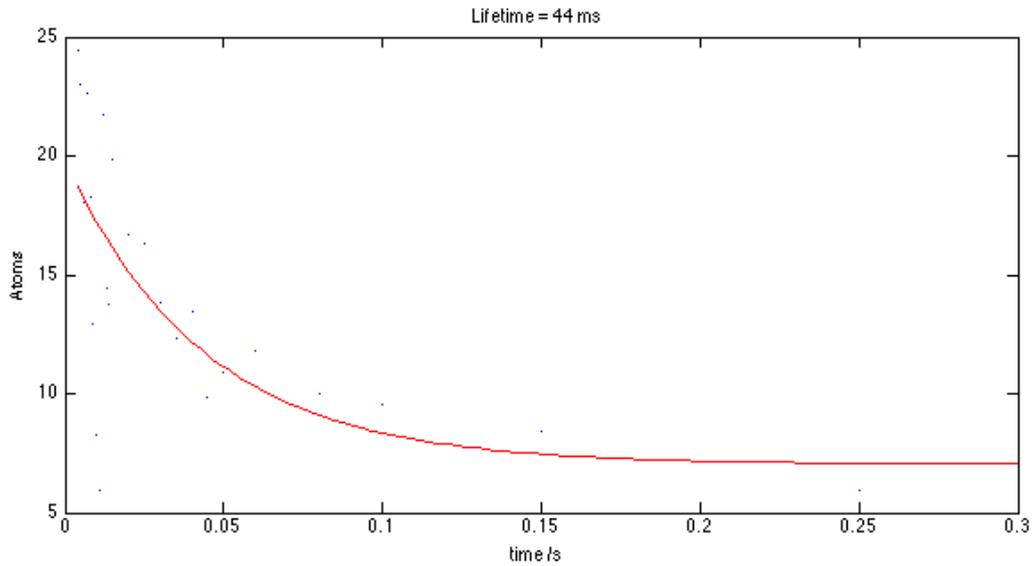


Figure 7.11 Lifetime measurement for a trap with $d=6\mu\text{m}$.

We next verified the versatility of the tweezers by trapping atoms in arbitrary shaped potentials. The trap arrays shown in Figure 7.14 underline the scalability of the scheme: such arrays could be used to implement a large register of atomic qubits. The star shows that it is possible to just have fun with atoms. In the first two images, the radius of each individual trap is $3\mu\text{m}$, and the average atom number in each is 12 ± 2 . The third image shows atoms trapped in a bullseye-shaped trap. The $\text{\O}6\mu\text{m}$ 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 Figure 7.14 shows 300 ± 40 atoms in a star-shaped trap which is affected by an inhomogeneity in the trapping light. The intensity variation of

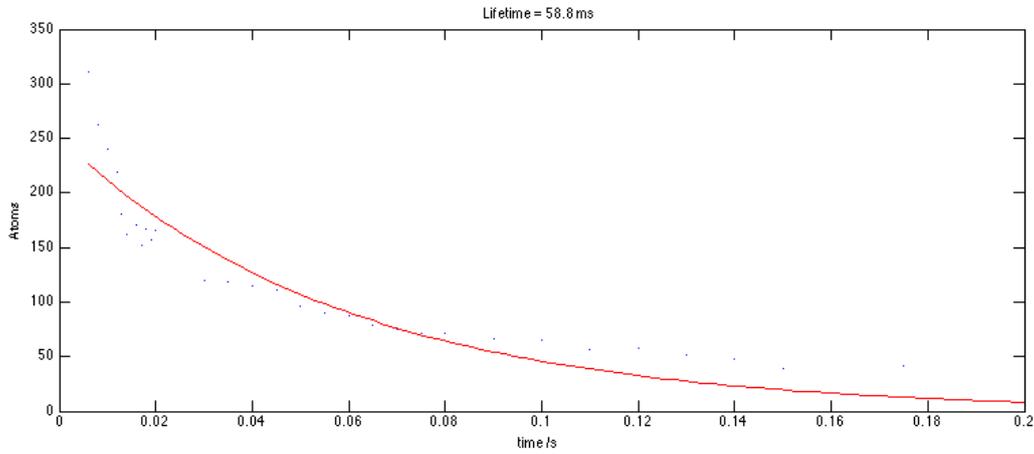


Figure 7.12 Lifetime measurement for a trap with $d=30\mu\text{m}$.

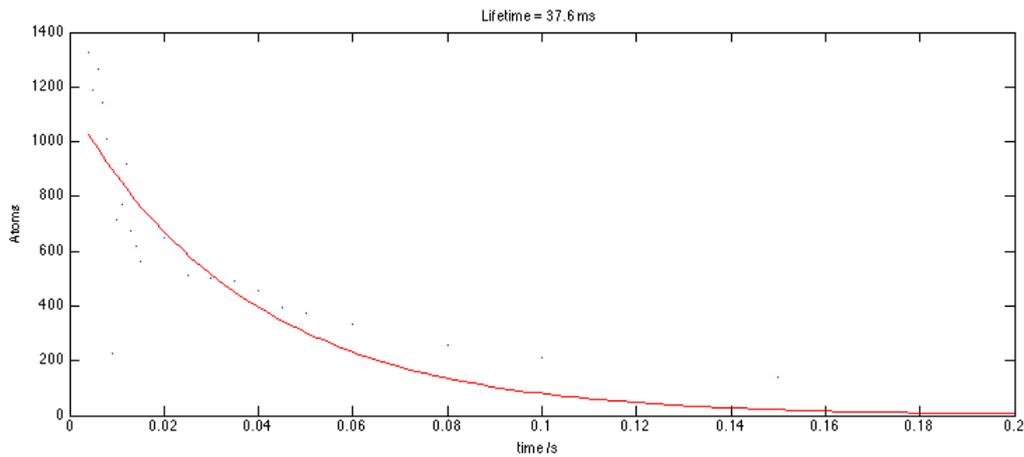


Figure 7.13 Lifetime measurement for a trap with $d=65\mu\text{m}$.

the laser over the whole area of the SLM is about 30% (see Figure 5.14). 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, the large variety of different shapes highlights the flexibility of trap generation with the SLM. Trapped atoms can be arranged into arbitrary patterns

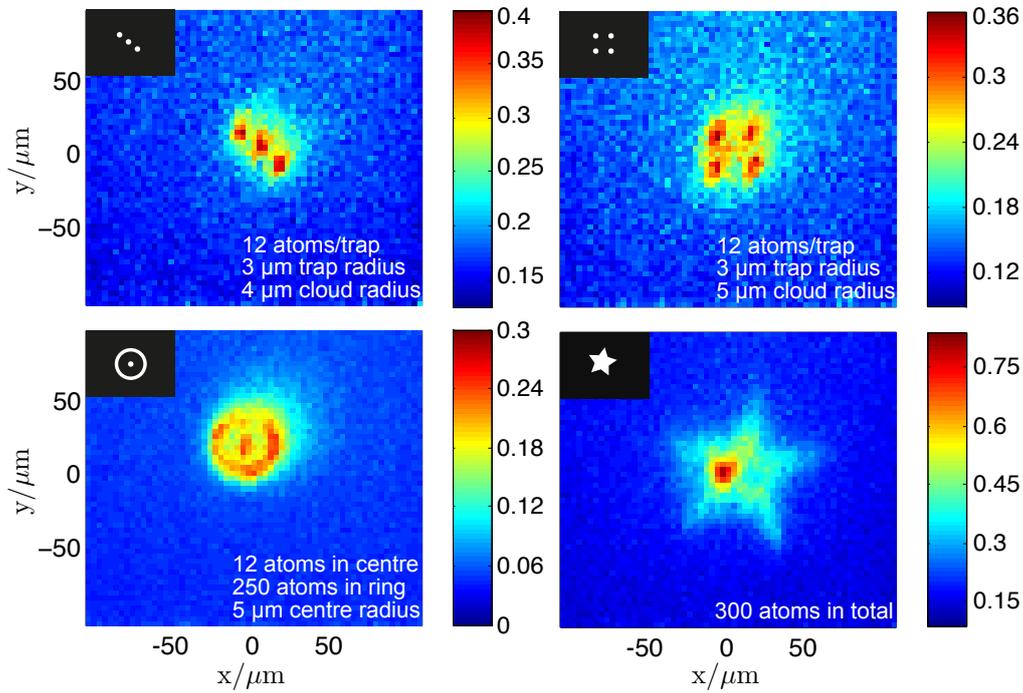


Figure 7.14 Fluorescence images of atoms in various arbitrary trapping geometries: a line, a grid, a bullseye and a star. Each image is an average over 20 consecutive trap loading cycles. The colour scale has units of atoms per μm^2 . The actual patterns on the SLM used to produce the traps are shown in the respective insets using the same length scale. The halo seen in the first three fluorescence images can be attributed to diffusion of atoms in the light sheet.

using the SLM to impose any desired shape of the trapping potential.

For the larger traps depicted in Figure 7.7 and 7.14, 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 $\leq 3 \mu\text{m}$, 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 reabsorp-

tion, and hence does not lead to an underestimation of the observed atom number.

7.3 Collection Efficiency & Atom Number Calibration

We estimated the approximate atom number in each trap from the total number of counts present within a chosen region of a fluorescence image. In order to determine the number of fluorescence photons indicated by these counts, it was first necessary to calibrate the imaging system in order to determine its overall photon-detection efficiency. This 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 $\eta_{lens}=7.6\%$, the losses along the the optical path $\eta_{loss}=51\%$, and the quantum efficiency of the EMCCD camera $\eta_{QE}=41\%$ (see Section 5.12). This photon-detection efficiency means that of the ~ 1400 photons emitted by an atom in the $90\mu s$ the light sheet is flashed on for, only 21 of them are detected. The collection efficiency of the lens system η_{lens} is determined by the numerical aperture of the aspheric lens, and the greatest contribution to the losses along the optical path η_{loss} comes from the presence of the viewport.

Roughly speaking, the procedure used to convert the number of counts on the EMCCD camera to photons is as follows. The first step is to make a background subtraction. The background can either be an image taken with the repump laser blocked off (thus ensuring no atomic fluorescence), or the mean value of a small square neighbouring the region of interest on the fluorescence image itself. The second step is to multiply this adjusted number of counts by the sensitivity of the

camera. The sensitivity is given by the number of electrons per count divided by the quantum efficiency of the camera. As explained in Section 5.12, the number of electrons per count depends on the preamplifier gain used and is also slightly different depending on the readout rate chosen (due to electronic noise). The last step is to divide this number by the EM gain of the camera as well as by η_{lens} and η_{loss} .

To deduce the number of atoms the resulting photons correspond to it is then necessary to calculate the scattering rate using the intensity of the light sheet in the trapping region. Multiplying this scattering rate by the exposure of the camera (which is equivalent to the $90\mu s$ flash of the light sheet) gives the number of photons that should be collected from a single atom within that time. The number of atoms is then obviously given by the number photons obtained from the camera counts divided by the number of photons emitted by a single atom in the exposure window.

Chapter 8

Deterministic re-arrangement & controlled transport of atoms

So far, we have shown that it is possible to use our tweezers to hold atoms 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 also have the ability to move atoms independently from one trapping site into another. This chapter shows how to accomplish this task and deterministically transport atoms between two well defined positions. A summary of these results can be found in^[61].

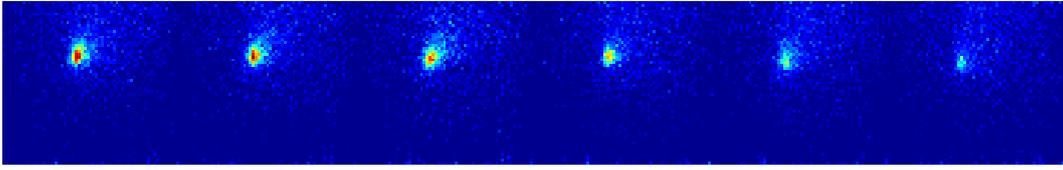


Figure 8.1 Moving Trap Transport

8.1 Moving Trap Transport

One would imagine that the most intuitive way of transporting atoms would be to simply move the trap of interest by gradually changing the image on the DMD such that the centre of the trap gets displaced. In practice, this is done by simply switching off a line of mirrors on the trailing edge of a round trap, and turning on a new line at the leading edge of the same trap. However, this approach of moving a point-like atom trap is hampered by the discrete switching of the DMD's mirrors, which results in the trap hopping along its trajectory rather than in a smooth motion. In turn, the atoms get lost due to excess heating. The filmstrip shown in Figure 8.1 shows how most of the atoms are lost after only a few frames. For each frame, only the necessary lines on the DMD are refreshed – i.e. the ones containing the trap whose leading and trailing edge are being modified – so the overall refresh rate of the DMD is less than it is for a full array refresh. For the movie shown in Figure 8.1, the duration of each frame was of $180\mu\text{s}$. Slower refresh rates led to an even greater loss of atoms from the trap.

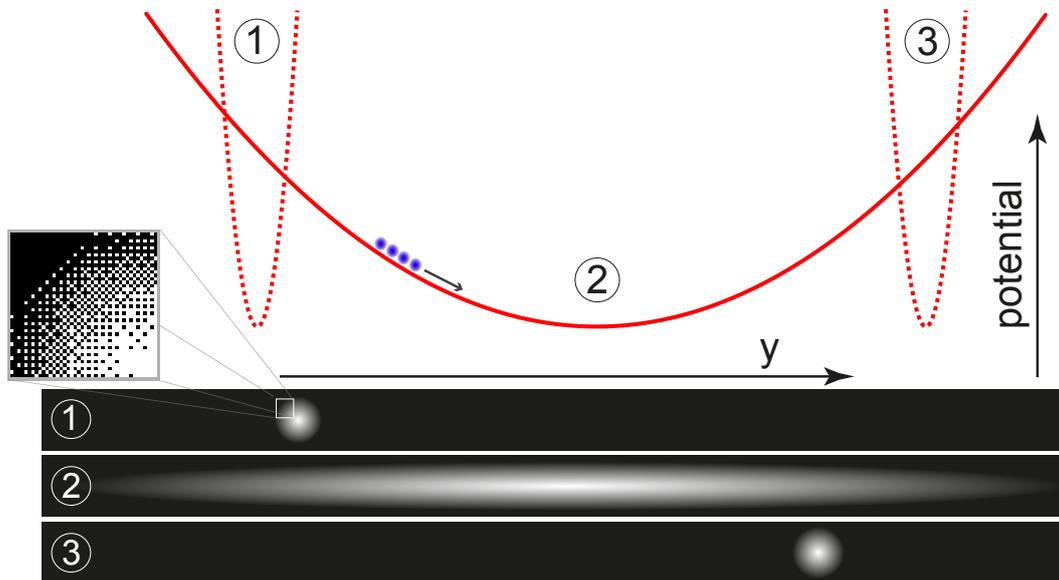


Figure 8.2 Illustration of the concept of the atom transport via ballistic “release and recapture” (not to scale). The atoms are first confined in trap (1), then move along the harmonic transport trap (2), and refocus at the destination when the end trap (3) is switched on. The close-up view shows the dithered bitmaps we use to approximate grayscales on the DMD and hence produce the requisite harmonic potentials. Smoothing results from the optical system not resolving individual mirrors.

8.2 Ballistic Transport

Our solution to this problem was to release the atoms from an initial trap and then let them roll along a harmonic potential before recapturing them on the other side, at 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. Figure 8.2 illustrates this procedure and also shows the dithered mirror pattern used to obtain a smooth trapping potential in the image plane. The smoothing results from the optical system not resolving in-

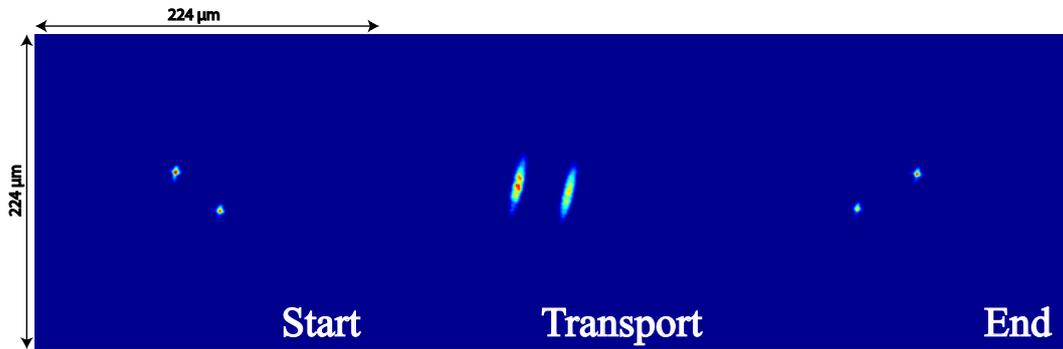


Figure 8.3 False colour image of the transport potentials formed in the image plane of the DMD. The small region of higher intensity present in the leftmost transport trap is a result of the sub-optimal profile of the trapping laser.

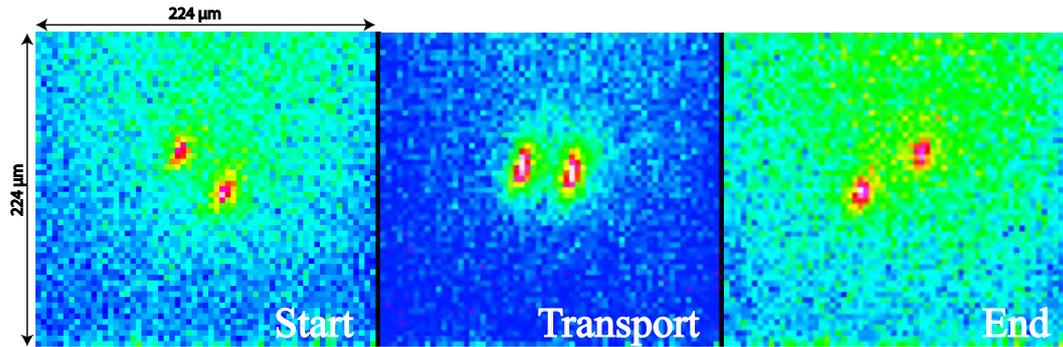


Figure 8.4 Static fluorescence images of atoms trapped in the start, transport, and end potentials used in our ballistic “release and recapture” scheme.

dividual mirrors, and allows us to approximate the grayscale needed to produce a harmonic potential using the DMD. The harmonic transport traps have the form $U \propto -(1 - (x/w_x)^2 - (y/w_y)^2)$, which means that they have a width of w_x by w_y and corresponding Gaussian field-waists of $\sqrt{2}w_x$ and $\sqrt{2}w_y$ (we usually make reference to the width).

Figure 8.3 shows a false colour image of the trapping light taken with the EM-

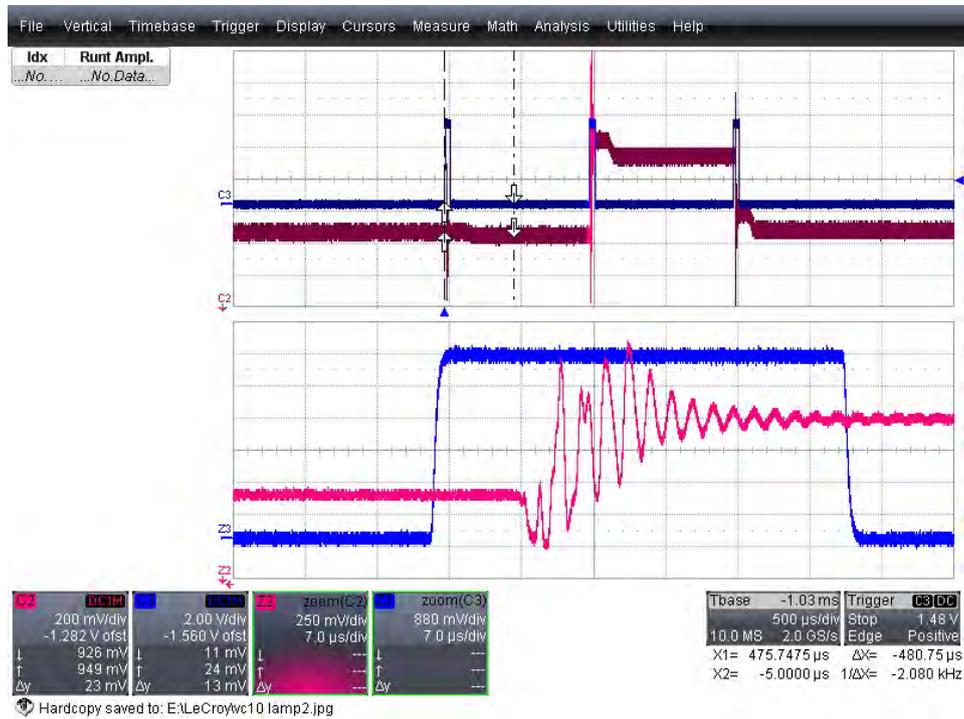


Figure 8.5 A typical oscilloscope trace from a ballistic transport sequence, with a zoom-in of the turn-on time for one of the traps. The mechanical vibrations that arise after the mirrors are switched die out after approximately $20\mu\text{s}$. The sequence was monitored constantly using one of the higher diffraction orders of the DMD because triggers were periodically missed, thereby altering the repeating start-transport-end sequence necessary when conducting several loading cycles.

CCD camera (this required removing merely two of the four 780 bandpass filters blocking off the trapping light at 785). The uneven intensity distribution on the DMD caused by the sub-optimal profile of the trapping laser leads to some degradation in the quality of the optical trapping potentials produced in the image plane. This can be seen clearly in the central panel of Figure 8.3, where a small region of higher intensity can be seen in the middle of the leftmost transport potential. For illustrative purposes, static fluorescence images of atoms trapped in these ballistic

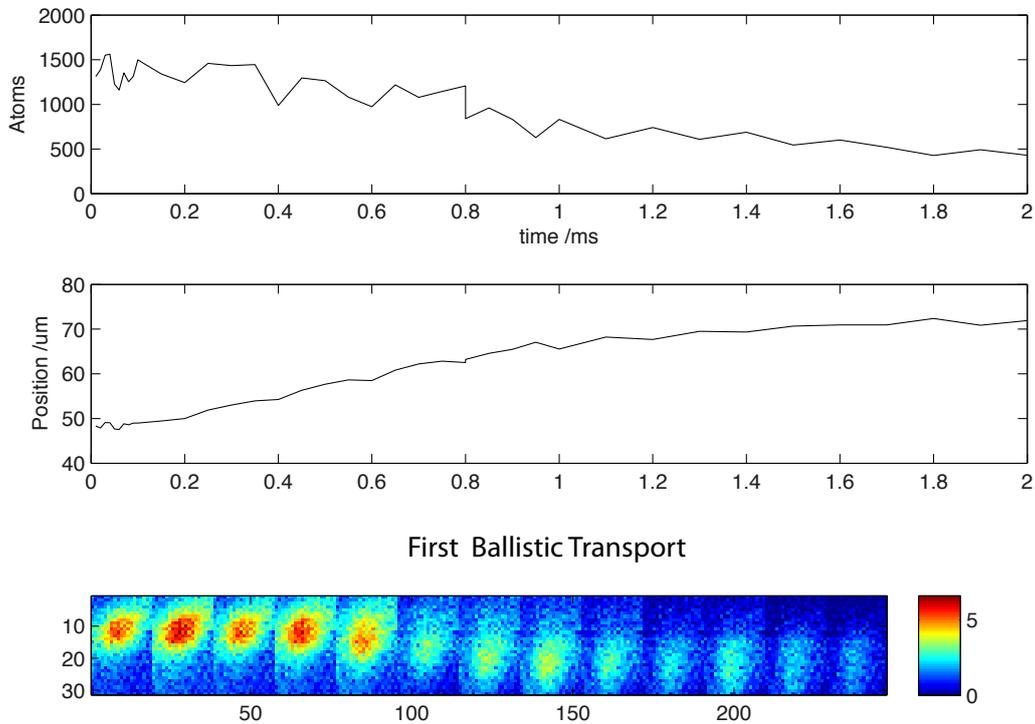


Figure 8.6 Fit of atoms oscillating in a transport potential with $w_y = 12\mu\text{m}$ and $w_x = 47\mu\text{m}$. The filmstrip below the fit shows a sequence of fluorescence images taken at various instants of the transport. The end trap was switched on after 1 ms.

“release and recapture” transport potentials are shown in Figure 8.4.

Figure 8.5 shows a typical oscilloscope trace from a ballistic transport sequence. The sequence was monitored by looking at one of the higher diffraction orders of the DMD. Channel 2 shows the mirrors switching into the start, transport, and end potentials. As can be seen in the zoom below, the mechanical vibrations that arise after the mirrors are switched die out after approximately $20\mu\text{s}$. It was necessary to monitor this sequence in order to ensure that no triggers were missed and that the sequence thus proceeded in the requisite order over several loading cycles.

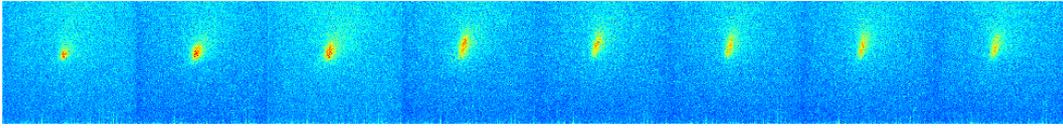


Figure 8.7 Single ballistic transport, 8x8 binning

The first ballistic transport was carried out with the larger light sheet in place (height= $30\mu\text{m}$), and employed relatively large traps. The start and end traps had a waist of $15\mu\text{m}$, and were placed $30\mu\text{m}$ away from the centre of the transport trap (i.e., the distance between the centre of the start trap and the centre of the transport trap was $60\mu\text{m}$). The transport trap had a width of $47\mu\text{m}$ along x (the transport axis), and a width of $12\mu\text{m}$ along y . The filmstrip in Figure 8.6 shows the atoms rolling along this transport potential before being re-captured by the end trap on the other side. The fit above it shows that it took approximately 1 ms for the atoms to reach their outer turning point, at which time the end trap was switched on.

As with the static traps, the introduction of a much thinner retro-reflected light-sheet reduced the blur in the fluorescence images significantly, and allowed us to use much smaller transport potentials. In the transport sequence shown in Figure 8.8, the start and end traps consist of solid circles with a diameter of $6\mu\text{m}$, and the transport potential has dimensions $w_y = 6\mu\text{m}$ by $w_x = 15\mu\text{m}$. The slightly grainier appearance of the fluorescence images shown underneath the graph is due to the fact that the binning on the EMCCD camera was set to 8x8 instead of 16x16. In

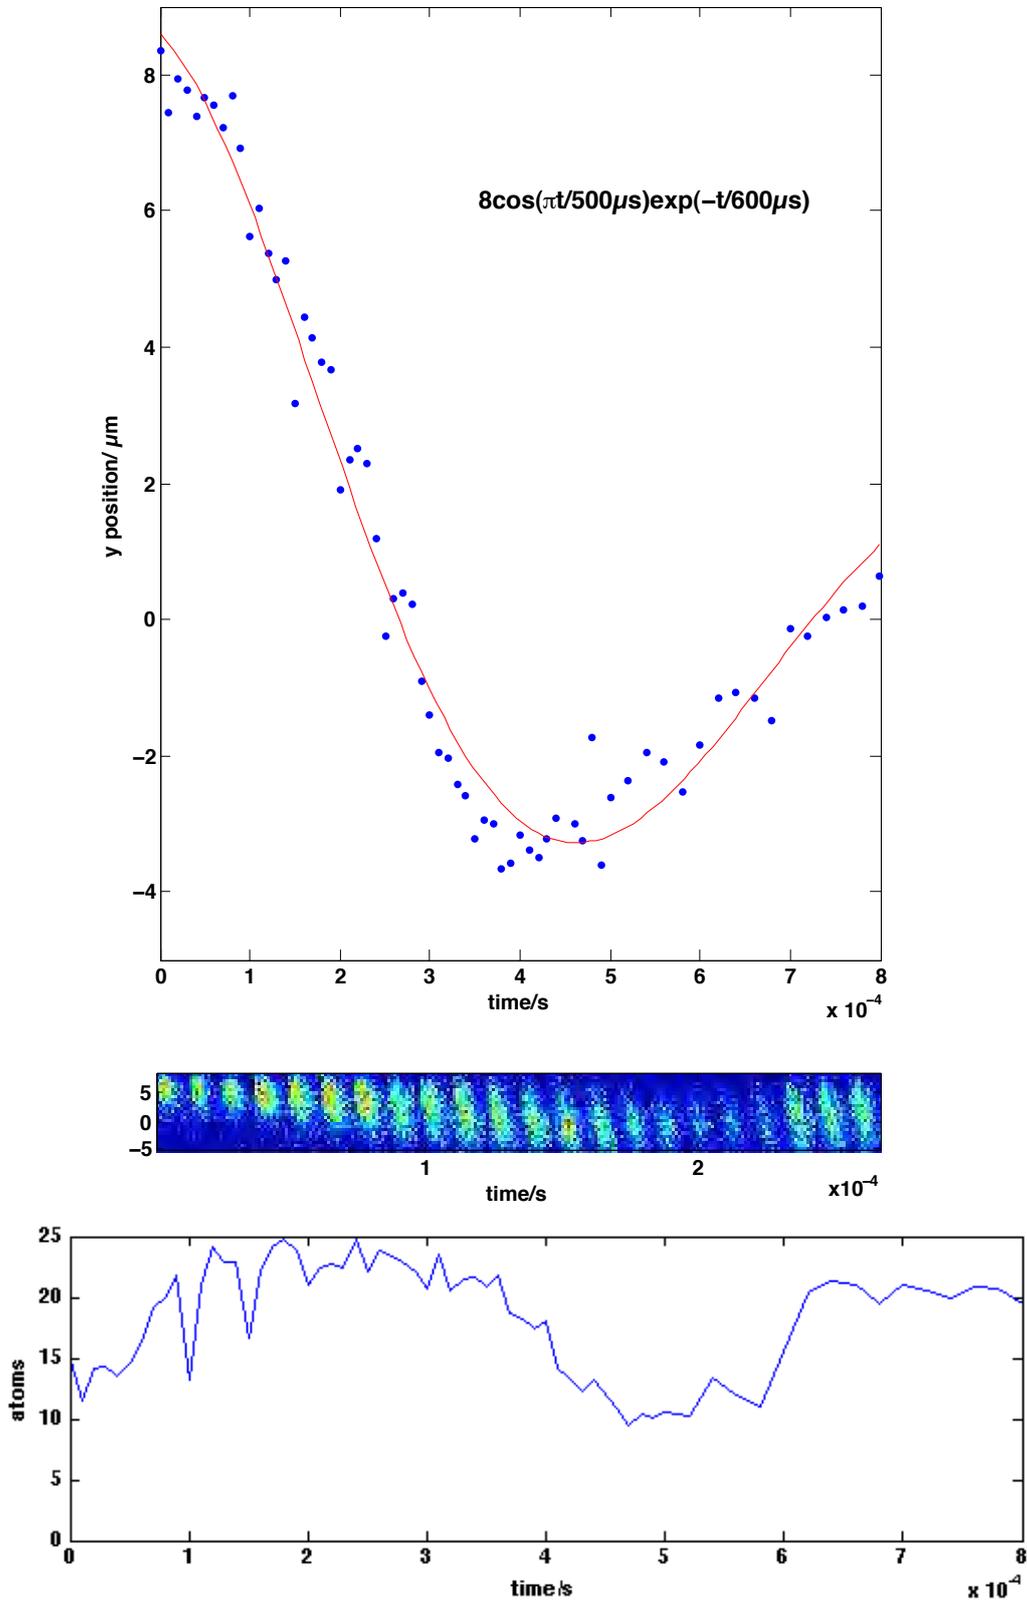


Figure 8.8 Fit of atoms oscillating in a transport potential with $w_y = 6\mu\text{m}$ and $w_x = 30\mu\text{m}$. The filmstrip below the fit shows a sequence of fluorescence images taken at various instants of the transport. The images appear to be slightly grainier due to the binning of the EMCCD camera, which was set to 8x8 instead of 16x16.

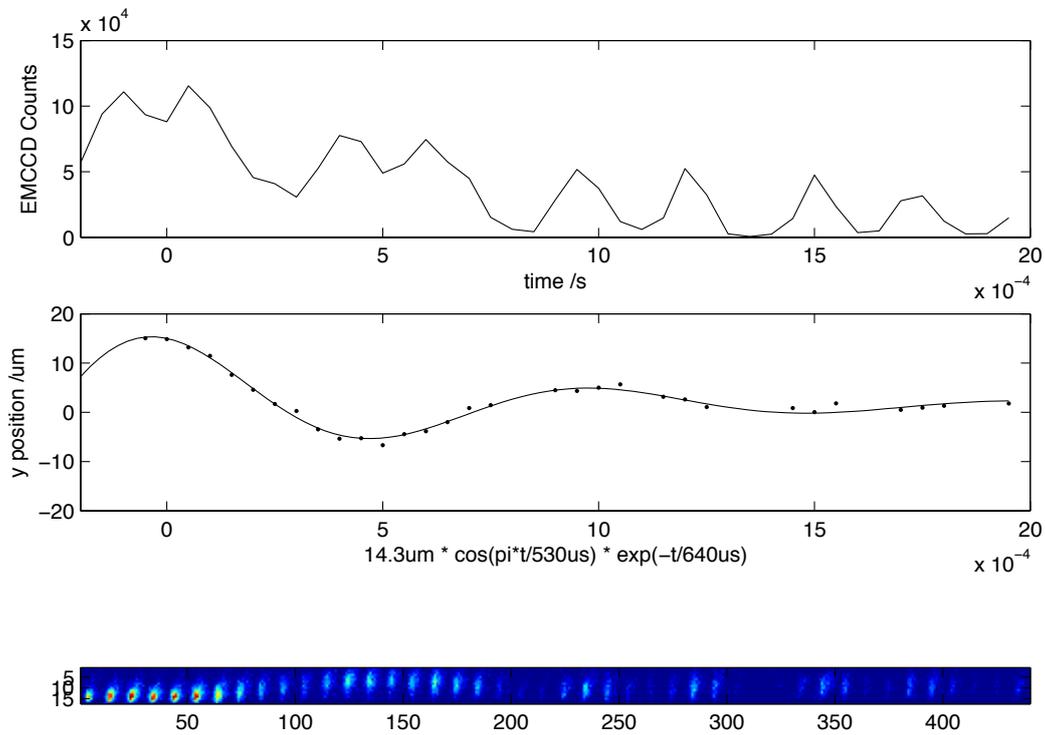


Figure 8.9 Fit of atoms oscillating in a transport potential with $w_x = 27\mu\text{m}$. The filmstrip below the fit shows a sequence of fluorescence images taken at various instants of the transport. The graph of EMCCD counts shown at the top shows that the atom number dips periodically throughout the transport.

this transport, the atoms were allowed to oscillate freely, without turning on the end trap. As can be seen from the shortened filmstrip in Figure 8.7 and the graph in Figure 8.8, the atoms rolled along the potential as before, but did not manage to complete a full oscillation. Rather, they began to spread out in the harmonic transport well approximately one third of the way back to their starting position. More interestingly, the atoms number seemed to dip significantly soon after the atoms reached their turning point.

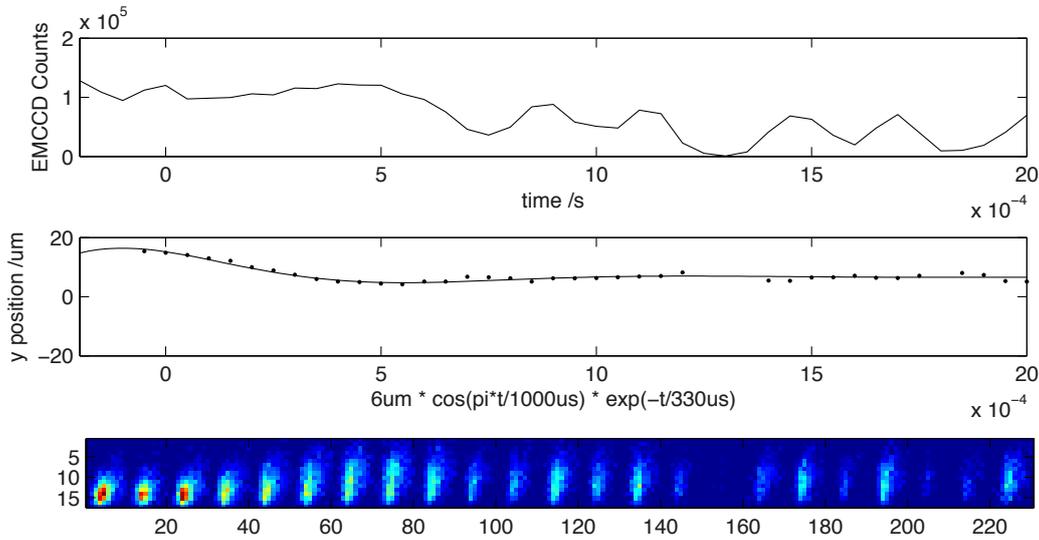


Figure 8.10 Fit of atoms oscillating in a transport potential with $w_x = 40\mu\text{m}$. The film-strip below the fit shows a sequence of fluorescence images taken at various instants of the transport. The graph of EMCCD counts shown at the top shows that the atom number dips still dips periodically throughout the transport, irrespective of the length of the transport potential.

8.3 Transport in 3D

A first attempt at explaining the observed dip in the atom number was to consider that the atoms were most likely not only oscillating in the xy plane, along the transport potential, but that they were also oscillating along the z axis. Thus, atoms might be moving out of the region illuminated by the light sheet, in turn leading to dips in the atom number recorded in the fluorescence images. We therefore made an attempt to make the trapping frequency ω_x along the direction of transport commensurate with ω_z , the trapping frequency perpendicular to the trapping plane. We repeated the transport with three different transport potentials having a width along y of $7\mu\text{m}$ and widths along x of $27\mu\text{m}$, $40\mu\text{m}$, and $54\mu\text{m}$ respectively. We left

the start and end traps unchanged and in the same position relative to the centre of the transport trap, such that the atoms effectively saw a shallower transport potential as the width along x increased. Figures 8.9 and 8.10 show results for the transport potentials with $w_x = 27\mu\text{m}$ and $w_x = 40\mu\text{m}$ respectively, and demonstrate that the dips in the atom number persisted irrespective of the value of w_x . Our conclusion was therefore that these dips could not be explained by the coincidence of ω_z with ω_x . In addition, we realised that even the shortest of these transport potentials would have a Rayleigh length of $400\mu\text{m}$, making the oscillation frequency of the atoms along the z axis less than 100 Hz, and thus much slower than the oscillation frequency along the direction of transport.

Our intuition was not totally incorrect, as in general one has to consider that the atoms actually move in three dimensions. The DMD controls the potential in the trapping plane, whilst 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. To investigate this further, we calculated the near-field intensity distribution away from the image plane for the transport potential with $w_x = 27\mu\text{m}$ using the Fresnel-Kirchhoff diffraction integral. The resulting plots are shown in Figure 8.11. From the running wave shown in the first five panels, it is clear that there is some refocussing about $100\mu\text{m}$ above the image plane. The divergence of the transport potential is moderate enough to yield a standing-wave modulation visibility of 75% at the height

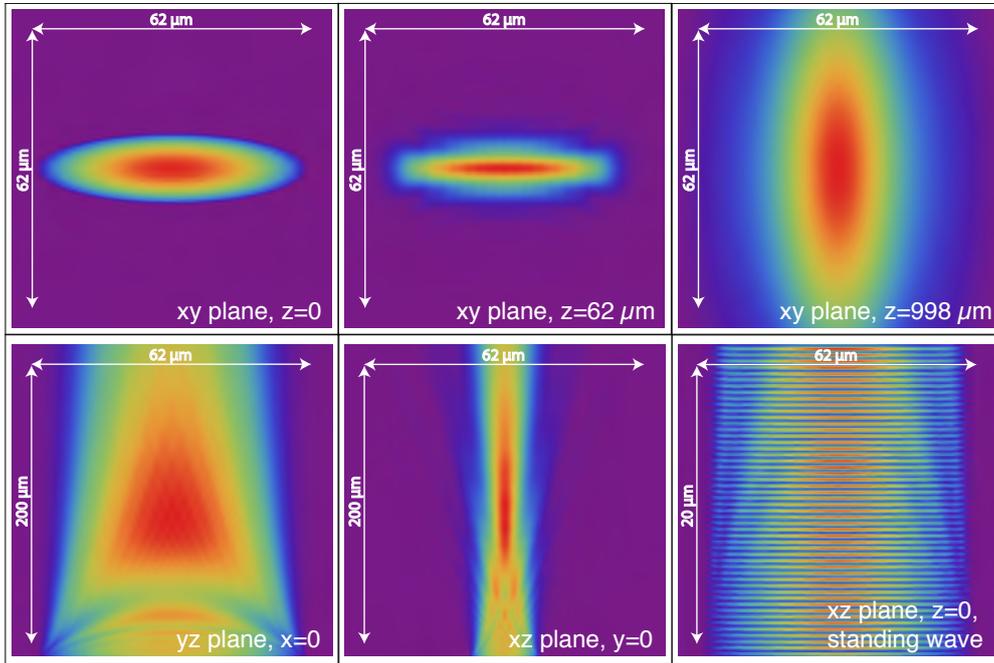


Figure 8.11 Near field intensity distribution of a transport potential with $w_y = 7\mu\text{m}$ and $w_x = 27\mu\text{m}$, away from the image plane of the DMD. The first five panels show the intensity distribution for a running wave in different planes and at different heights above the image plane. In the middle of the transport potential, the reflected beam ($\sim 100\mu\text{m}$ away) still has one fifth of the intensity, which means that standing wave effects are significant. The last panel shows the intensity distribution in the xz plane, along the direction of transport, for a standing wave with 75% visibility.

of the image. The last panel in the figure shows this partial standing wave. It is clear from this panel that the transport probably occurs along several channels at the same time: a modulation of 75% is very strong, and most atoms would be confined to the resulting anti-nodal planes. Therefore, we effectively have 20 identical two-dimensional trapping planes within the width of the light sheet. This could go some way towards explaining the dips in the atom number and the overall losses observed: the atoms are not lost due to heating, but merely because they leave the

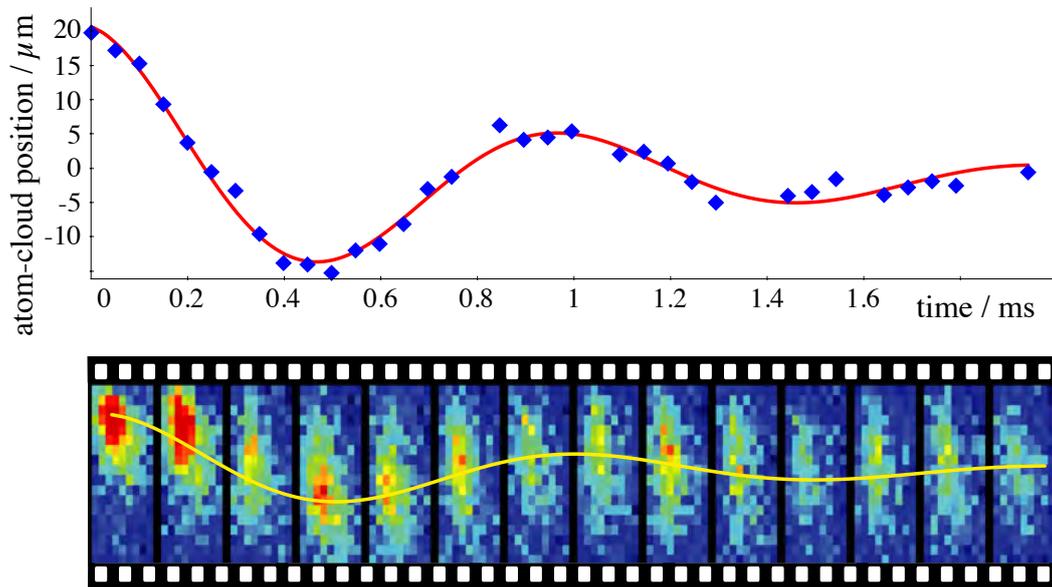


Figure 8.12 The filmstrip below the fit shows a sequence of fluorescence images taken at various instants of the transport.

area illuminated by the light sheet.

8.4 Single Ballistic Transport

In spite of these limitations, it was possible to optimise the loading sequence and the dimensions of the trap in order to successfully recapture the atoms and measure both their lifetime and their temperature. Figure 8.12 shows atoms oscillating along a transport potential with widths $w_y = 7\mu\text{m}$ and $w_x = 27\mu\text{m}$ respectively after being released from a start trap with a width of $7\mu\text{m}$. The half-oscillation period of the atoms in this potential was found to be $520 \pm 10\mu\text{s}$, and the damping time extracted from the fit was found to be 1.3 ms. Figure 8.13 shows the atoms being recaptured after $520\mu\text{s}$, at the turning point of their first oscillation. The atoms travel a dis-

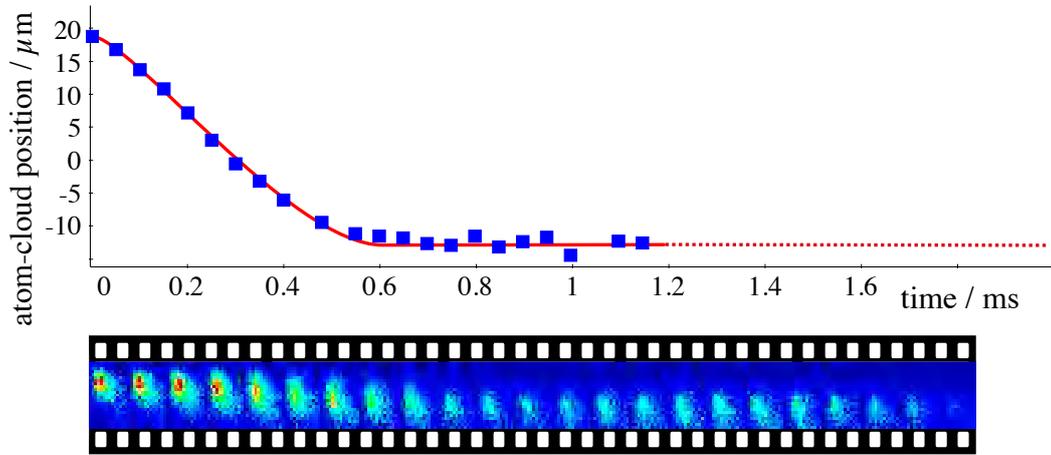


Figure 8.13 The filmstrip below the fit shows a sequence of fluorescence images taken at various instants of the transport.

tance of $27 \mu\text{m}$ before being recaptured, and the lifetime of the atoms in the end trap was found to be 1.5 ms.

The displayed atom-cloud position is the centre-of-mass of the recorded fluorescence images, which were taken at a constant rate. Of the 22 atoms initially in the trap, 12 atoms are recaptured. Most surprisingly, this cannot be attributed to heating-induced losses. Time-of-flight measurements reveal a temperature rise of only $10 \mu\text{K}$. Instead, the reduction in fluorescence is to some extent explained by atoms not exactly regrouping at their destination. The primary reason for this limitation is the small anharmonicity of the transport trap brought about by the profile of the trapping laser, which results in the 1.3 ms damping time of the harmonic motion referred to before (see the yellow trace in Figure 8.12). The secondary loss mechanism is a diffusion of the most energetic atoms along the standing wave, i.e.

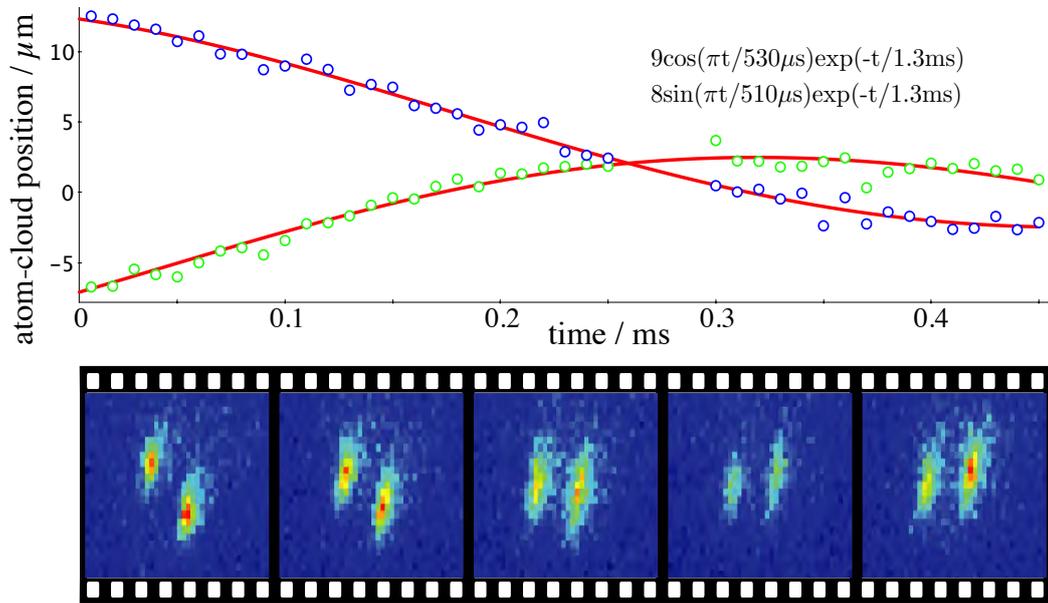


Figure 8.14 Atoms oscillating past each other in two separate harmonic wells. The film-strip below the plot shows a sequence of fluorescence images taken at various instants of the transport.

away from the image plane. Both losses could be easily reduced using a more powerful trapping laser with a better laser profile. Hence, disregarding the spreading, this ballistic transport in harmonic traps illustrates that the DMD is an ideal tool for positioning atoms in a deterministic way.

8.5 Double Ballistic Transport

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 each other. 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

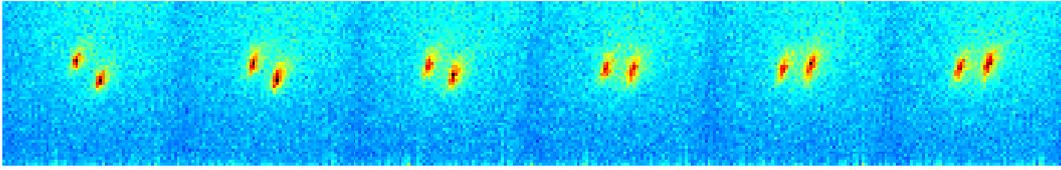


Figure 8.15 Double Transport

above. Figure 8.14 and the movie in Figure 8.15 show atoms oscillating past each other in these harmonic potentials. The atoms move to destinations $8\ \mu\text{m}$ and $9\ \mu\text{m}$ away from their respective starting points, with the two destinations chosen independently. As before, the atoms reach their destinations within $520 \pm 10\ \mu\text{s}$, and there is a $1.3\ \text{ms}$ damping time of the harmonic motion. This individual transport of randomly selected trapping sites shows that DMD-controlled atom traps are capable of regrouping arrays of trapped atoms to nearly any arbitrary pattern.

8.6 Optimal Transport

Prior to testing our ballistic transport scheme in the lab, we considered how we could optimise the shape of the transport potentials so as to optimise the atom transport. In an ideal world, the transport would look something like the filmstrip in Figure 8.16. Put most simply, the goal was to use the DMD to transport atoms from point A to point B without heating them. The solution to this problem is straightforward and deterministic for one atom rolling across a harmonic potential in 1D, but the situation is much more challenging for a thermal cloud rolling along an optical potential produced by the DMD. Even more challenging are the restrictions imposed

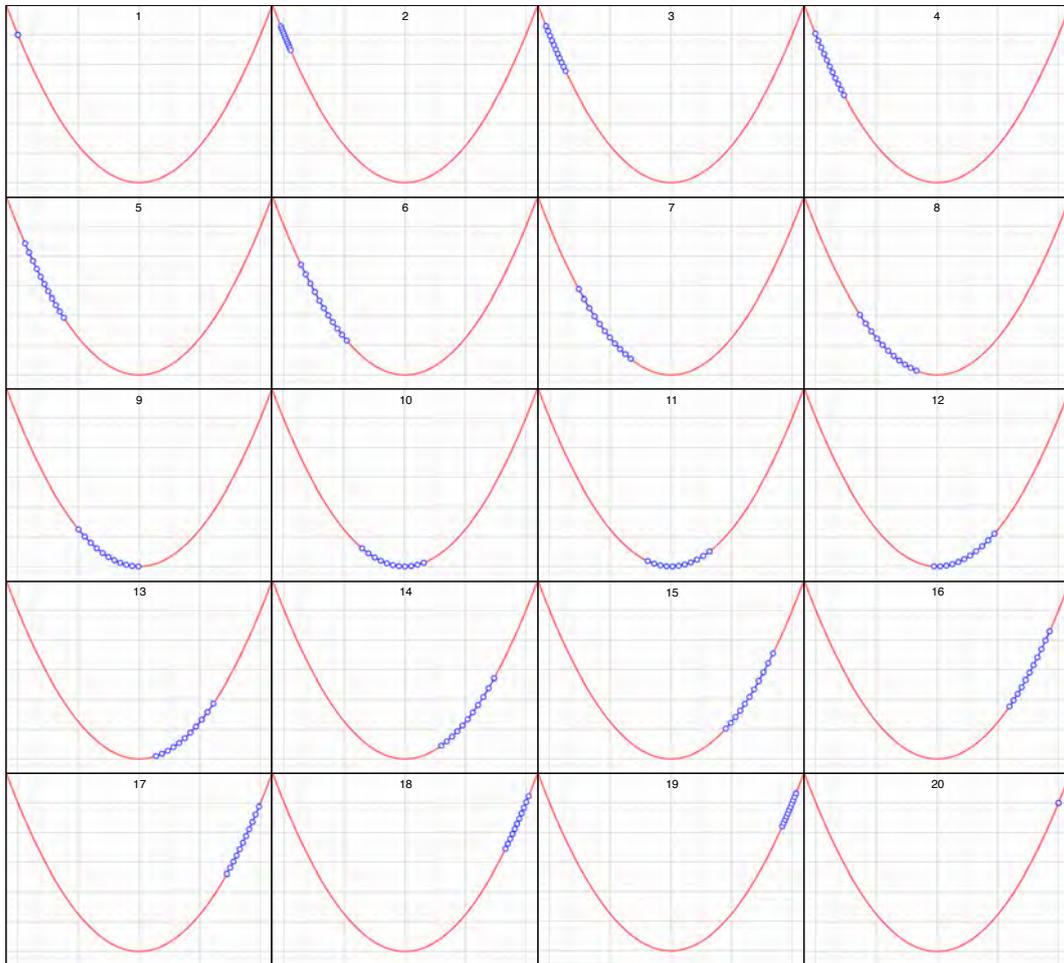


Figure 8.16 Ideal harmonic transport

by the DMD: a limited potential depth, a small dynamic range, and the discrete switching of the mirrors.

We considered the question from the point of view of optimal control theory, an optimisation method whose goal is to find a control rule which guarantees that a given optimality criterion is achieved. In an optimal control problem, there are

variables, controls, and a dynamical law, and the physical constraints consist of admissible states, admissible goals, initial states, and final states. In our case, the states are the contributions from each mirror to the electric field distribution in the image plane of the DMD, the controls are whether the mirror is on or off, and the dynamical law is simply Newton's law, which describes how the atoms move in the potential generated by the mirrors. The constraints are given by the desired final position and final velocity of the atoms, and by the discrete switching of the mirrors. The problem was formulated in more detail with help from Edouard Brainis, a postdoc no longer in our group, and Michael Murphy, a collaborator from Ulm University who specialises in optimal control theory.

More formally, we considered the state of an atom in an optical potential $V(x, t) = |E(x, t)|^2$, where E is the electric field distribution in the image plane of the DMD:

$$E(x, t) = \sum_{m=1}^M h_m(x) u_m(t). \quad (8.1)$$

Here, m is the mirror index, $h_m(x)$ is the field contribution at point x from of the single mirror m , and the u s are the *control parameters*. If the mirror m is “on”, $u_m = 1$. If it's “off”, $u_m = 0$. An atom moving in this potential follows Newton's law: $m \frac{\partial^2 x}{\partial t^2} = -\frac{\partial}{\partial x} V(x, t)$. The state of an atom is described by the dynamical variables $x(t)$ and $v(t)$, which correspond to its position and velocity. The notation $\mathbf{x}(t) = [x(t), v(t)]^T$ is used for the state variables, and the notation $\mathbf{u}(t) = [u_1(t), \dots, u_M(t)]^T$ is used for the control parameters. The equations of motion

for these state variables are given by

$$\frac{\partial}{\partial t}x(t) = v(t) \quad (8.2a)$$

$$\frac{\partial}{\partial t}v(t) = -\frac{1}{m} \frac{\partial}{\partial x} \left| \sum_m h_m(x) u_m(t) \right|^2. \quad (8.2b)$$

or more generically,

$$\frac{\partial}{\partial t}\mathbf{x} = \mathbf{a}(\mathbf{x}(t), \mathbf{u}(t)) \quad (8.3)$$

The idea is to find an *admissible control* $\mathbf{u}^*(t)$ which causes the system to follow an *admissible trajectory* $\mathbf{x}^*(t)$ that minimises some *cost functional* J . We gave our cost functional the following form:

$$\begin{aligned} J = & w_1 (x(t_f) - x_f)^2 + w_2 (v(t_f) - 0)^2 + w_3 \int_{t_0}^{t_f} \sum_m k_1(u_m(t)) \partial t \\ & + w_4 \int_{t_0}^{t_f} k_2(t) \sum_m \left(\frac{\partial u_m}{\partial t}(t) \right)^2 \partial t \end{aligned} \quad (8.4)$$

This cost functional consists of four terms, each term representing a different “optimality criterion”. The importance of each criterion can be weighted according to the weight parameters w_i . The first and second terms correspond to the requirement that the particle must arrive as close as possible to a given final position x_f and that it must do so with the smallest possible final velocity. The third term gives extra penalty to those trajectories that require that the controls $u_m(t)$ take values different from 0 and 1. The continuous function $k_1(u)$ is null for $u \in \{0, 1\}$, has a

broad plateau (value 1) for u between 0 and 1 and takes even higher values (let's say 3) outside the $[0, 1]$ interval. The fourth term insures that controls are changed at some definite times $t_k = t_0 + kT$, where T is the switching period. For this to work, the continuous function $k_2(t)$ must be null at times t_k and have broad plateaus everywhere else. This will give extra penalty to trajectories obtained by switching the controls at times that are not allowed.

The trajectory $\mathbf{x}(t)$ that minimizes J is found by applying variational methods, with the additional constraint that $\mathbf{x}(t)$ must satisfy equations 8.2. The solution is found by applying the method of Lagrange multipliers and minimising the *augmented functional*

$$J_a = h(\mathbf{x}(t_f)) + \int_{t_0}^{t_f} [g(\mathbf{u}(t), \dot{\mathbf{u}}(t), t) + \mathbf{p}^T(t) [\mathbf{a}(\mathbf{x}(t), \mathbf{u}(t)) - \dot{\mathbf{x}}(t)]] \, dt, \quad (8.5)$$

where $\mathbf{p}(t) = [p_1(t), p_2(t)]^T$ are the Lagrange multipliers. For an extremal solution

$$\begin{aligned} \delta J_a = & \left[\frac{\partial g_a}{\partial \dot{\mathbf{x}}} \Big|_{t=t_f} \right]^T \delta \mathbf{x}_f + \left[g_a \Big|_{t=t_f} - \left[\frac{\partial g_a}{\partial \dot{\mathbf{x}}} \Big|_{t=t_f} \right]^T \dot{\mathbf{x}}(t_f) - \left[\frac{\partial g_a}{\partial \dot{\mathbf{u}}} \Big|_{t=t_f} \right]^T \dot{\mathbf{u}}(t_f) \right] \delta t_f \\ & + \int_{t_0}^{t_f} \left\{ \left[\frac{\partial g_a}{\partial \mathbf{x}} \right]^T - \frac{\partial}{\partial t} \left[\frac{\partial g_a}{\partial \dot{\mathbf{x}}} \right]^T \right\} \delta \mathbf{x}(t) + \left[\frac{\partial g_a}{\partial \mathbf{u}} \right]^T - \frac{\partial}{\partial t} \left[\frac{\partial g_a}{\partial \dot{\mathbf{u}}} \right]^T \delta \mathbf{u}(t) \\ & + \left[\frac{\partial g_a}{\partial \mathbf{p}} \right]^T \delta \mathbf{p}(t) \Big\} \, dt = 0. \end{aligned} \quad (8.6)$$

By considering the different variations one by one, it is possible to find a set of Lagrange equations to solve and find the required boundary conditions. These equations are then solved by numerical methods. The idea behind optimal control is that, starting from an initial guess for the control history $u_m^{(0)}(t)$ ($t \in [0, t_f]$, $m \in \{1, \dots, M\}$), better control histories are found after some optimisation steps. Eventually, after enough optimisation steps, the value of $u_m^{(s)}(t)$ will converge to the optimal control: $\lim_{s \rightarrow \infty} u_m^{(s)}(t) = u_m^*(t)$. At each optimisation step, the requirement is that

$$\Delta J_a = J_a(\mathbf{u}^{s+1}) - J_a(\mathbf{u}^s) < 0. \quad (8.7)$$

The δJ_a presented in equation 8.6 is a linear approximation to ΔJ_a , and valid as long as $\delta \mathbf{u}(t) = \|\mathbf{u}^{s+1}(t) - \mathbf{u}^s(t)\|$ is small enough.

The numerical simulation to solve the optimal control problem was implemented by Michael Murphy and later extended by Edouard Brainis. Unfortunately, it did not succeed in finding an optimal solution to the ballistic transport problem. Surprisingly, the optimal solution was an analytical solution, and was in fact an initial guess for the optimisation program. We proposed a potential $Vhs(x, t) = Vh(x) - Vh(Xh(t))$, where $Vh(x) = \frac{F_{max}}{D} x^2$ is an ideal harmonic potential, $Xh(t) = \frac{D}{2} \sin(\omega t)$ corresponds to ideal harmonic motion, and $\omega = \sqrt{\frac{2F_{max}}{DM}}$ is in turn the oscillation frequency in $Vh(x)$. $Vhs(x, t)$ is hence is an ideal potential with a time dependent energy shift, such that one always has $Vh(x)$ at the position of the atom.

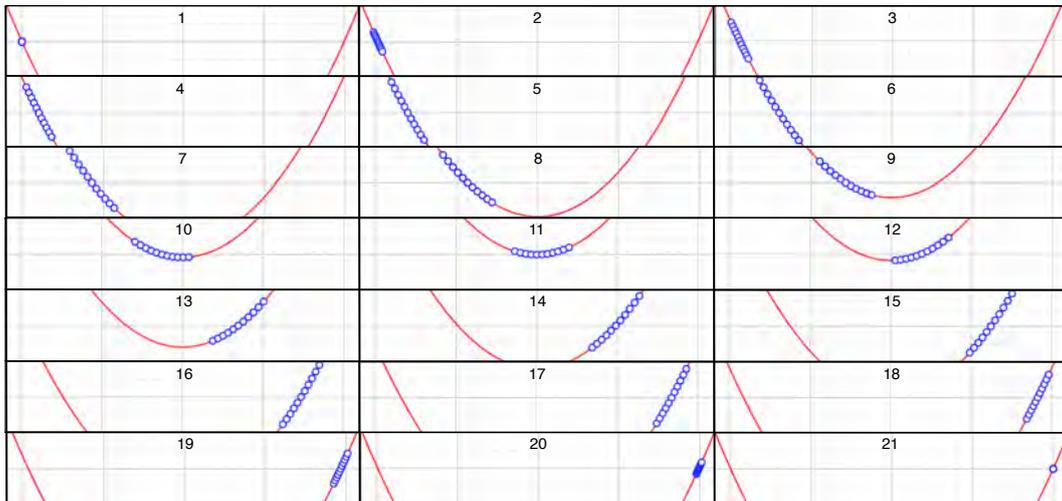


Figure 8.17 Time-dependent transport potential

In other words, the potential follows the motion of the atom. The filmstrip in Figure 8.17 shows how the atoms would “surf” along on this time shifting potential.

In the experiment, such a time-dependent potential would be realised by using a sequence of frames containing rasterised transport bitmaps of increasing/decreasing length and depth onto the DMD. The filmstrip in Figure 8.18 shows the sequence of bitmaps we actually produced to this end. Unfortunately, the refresh rate of the DMD, even at its minimum, was too slow to allow for this transport to be implemented: with the atoms reaching the opposite side of a $27\mu\text{m}$ transport potential in $500\mu\text{s}$, we would only be able to display about 4 or 5 of these frames within the transport window and would not be able to change the potential fast enough to follow the atomic motion. In general, the optimisation of the ballistic transport,

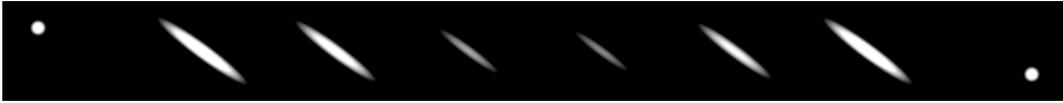


Figure 8.18 Shifting transport bitmaps

whether by analytical methods like this or by methods of optimal control, fails to take into account the three-dimensional standing wave nature of the transport, and was therefore a poor predictor of the actual atomic behaviour we observed in the lab.

Chapter 9

Conclusion

A DMD-based optical-tweezers setup constitutes 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 DMD is fast enough for a dynamic control of trapped atoms in response to the observed fluorescence, such that eventually, a feed-back scheme could be realised to control the atom number in individual sites. Furthermore, we emphasise that the actual trapping parameters would allow us to reach the collisional blockade regime, and at present, only technical limitations leading to an excess background fluorescence prevent us from observing individual atoms. With our ballistic transport applied to individual atoms, one might for instance move two atoms into fibre-tip micro cavities and realise pair-wise entanglement^[30], or use controlled collisions between two atoms to implement a two-qubit gate^[31-33]. This means that scalable atomic arrays for quantum computing or simulation are now in reach of current technology.

Appendix A

Supplement to The

Light-Matter Interaction

This Appendix contains detailed derivations of the concepts referred to in the description of the light-matter interaction presented in Chapter 2. It is based loosely on the derivations that can be found in^[62],^[63], and^[64].

A.1 The Gauge-Invariant Schrödinger equation

The motion of an electron with the wavefunction $\psi(\mathbf{r}, t)$ can be described using the Schrödinger equation

$$-\frac{\hbar^2}{2m}\nabla^2\psi(\mathbf{r}, t) = i\hbar\frac{\partial\psi(\mathbf{r}, t)}{\partial t}. \quad (\text{A.1})$$

If $\psi(\mathbf{r}, t)$ is a solution to this equation, then the probability density of finding the

electron at position \mathbf{r} , at time t , is given by

$$P(\mathbf{r}, t) = |\psi(\mathbf{r}, t)|^2. \quad (\text{A.2})$$

The addition of an arbitrary constant phase χ such that $\psi_1(\mathbf{r}, t) = \psi(\mathbf{r}, t) \exp(i\chi)$ has no effect on $P(\mathbf{r}, t)$, and if $\psi(\mathbf{r}, t)$ is a solution, then so is $\psi_1(\mathbf{r}, t)$. Addition of a phase that varies with \mathbf{r} and t , such that $\psi_2(\mathbf{r}, t) = \psi(\mathbf{r}, t) \exp(i\chi(\mathbf{r}, t))$ also leaves $P(\mathbf{r}, t)$ unaffected, but $\psi_2(\mathbf{r}, t)$ is no longer a solution to the Schrödinger equation. To satisfy this so called local phase (or *gauge*) invariance, it is necessary to modify the Schrödinger equation by introducing the gauge-dependent functions

$$\mathbf{A}(\mathbf{r}, t) = \mathbf{A}(\mathbf{r}, t) + \frac{\hbar}{e} \nabla \chi(\mathbf{r}, t) \quad (\text{A.3})$$

$$U(\mathbf{r}, t) = U(\mathbf{r}, t) - \frac{\hbar}{e} \frac{\partial \chi(\mathbf{r}, t)}{\partial t}. \quad (\text{A.4})$$

These functions correspond to the vector and scalar potentials of the electromagnetic field. They are gauge-dependent, whereas the electric and magnetic fields they define

$$\mathbf{E} = -\nabla U - \frac{\partial \mathbf{A}}{\partial t} \quad (\text{A.5})$$

$$\mathbf{B} = \nabla \times \mathbf{A} \quad (\text{A.6})$$

are gauge-independent. The gauge-invariant Schrödinger equation is then given by:

$$\left\{ -\frac{\hbar^2}{2m} [\nabla - i\frac{e}{\hbar} \mathbf{A}(\mathbf{r}, t)]^2 + eU(\mathbf{r}, t) \right\} \psi(\mathbf{r}, t) = i\hbar \frac{\partial \psi(\mathbf{r}, t)}{\partial t}. \quad (\text{A.7})$$

This equation describes the interaction of an electron with an electromagnetic field: the wavefunction $\psi(\mathbf{r}, t)$ represents the electron, and the scalar and vector potentials $A(\mathbf{r}, t)$ and $U(\mathbf{r}, t)$ represent the field.

A.2 The Dipole Approximation

Equation (A.7) can be further modified so that instead of a free electron, it describes an electron that is bound by the electrostatic potential $V(\mathbf{r}, t)$ to a nucleus at \mathbf{r}_0 . The Hamiltonian for this interaction can be simplified by using the *dipole approximation* in which the field wavelength is considered to be larger than the atomic size. In the dipole approximation, the vector potential $A(\mathbf{r}_0 + \mathbf{r}, t)$ can be simplified using the fact that $\mathbf{k} \cdot \mathbf{r} \ll 1$, so that

$$\begin{aligned} A(\mathbf{r}_0 + \mathbf{r}, t) &= \mathbf{A}(t) \exp[i\mathbf{k} \cdot (\mathbf{r}_0 + \mathbf{r})] \\ &= \mathbf{A}(t) \exp(i\mathbf{k} \cdot \mathbf{r}_0)(1 + \mathbf{k} \cdot \mathbf{r} + \dots) \\ &\simeq \mathbf{A}(t) \exp(i\mathbf{k} \cdot \mathbf{r}_0). \end{aligned} \quad (\text{A.8})$$

In other words, we can replace $A(\mathbf{r}_0 + \mathbf{r}, t)$ by $A(\mathbf{r}_0, t)$. We can further simplify the problem by applying the gauge transformation $\chi(\mathbf{r}, t) = -\frac{e}{\hbar} \mathbf{A}(\mathbf{r}_0, t) \cdot \mathbf{r}$ and rewriting the electron's wavefunction as $\psi(\mathbf{r}, t) = \exp[\frac{ie}{\hbar} \mathbf{A}(\mathbf{r}_0, t) \cdot \mathbf{r}] \phi(\mathbf{r}, t)$. Applying all of these changes to (A.7) gives

$$i\hbar \left[\frac{ie}{\hbar} \dot{\mathbf{A}} \cdot \mathbf{r} \phi(\mathbf{r}, t) + \dot{\phi}(\mathbf{r}, t) \right] \exp\left(\frac{ie}{\hbar} \mathbf{A} \cdot \mathbf{r}\right) = \exp\left(\frac{ie}{\hbar} \mathbf{A} \cdot \mathbf{r}\right) \cdot \left[\frac{p^2}{2m} + V(r) \right] \phi(\mathbf{r}, t) \quad (\text{A.9})$$

After some rearrangement, and noting that we are working in the radiation gauge in which $U(\mathbf{r}, t) = 0$ and $\nabla \cdot \mathbf{A} = 0$, such that $\mathbf{E} = -\dot{\mathbf{A}}$, this equation takes the simple form

$$i\hbar\dot{\phi}(\mathbf{r}, t) = [\mathcal{H}_0 - e\mathbf{r} \cdot \mathbf{E}(\mathbf{r}, t)]\phi(\mathbf{r}, t) \quad (\text{A.10})$$

where $\mathcal{H}_0 = \frac{p^2}{2m} + V(r)$.

A.3 The Rotating Wave Approximation

In order to solve the coupled differential equations for the probability amplitudes $c_1(t)$ and $c_2(t)$ presented in Section 2.1, we use the Rotating Wave Approximation (RWA), in which fast oscillating terms in a Hamiltonian can be neglected due to their negligible effect on the time evolution of a system. This is the same as saying that fast oscillating terms have little effect on the atomic populations – they average out over the interaction time. To see this more clearly, we define the slowly varying amplitudes

$$\begin{aligned} \tilde{c}_1(t) &= c_1(t)\exp(i\omega_{a1}t) \\ \tilde{c}_2(t) &= c_2(t)\exp(i\omega_{a2}t) \end{aligned} \quad (\text{A.11})$$

Substituting for $c_1(t)$ and $c_2(t)$ in (2.8) and simplifying gives

$$\begin{aligned} \dot{\tilde{c}}_1(t) &= i\frac{\Omega}{2}\tilde{c}_2(t)[\exp(i(\omega - \delta\omega_a)t) + \exp(-i(\omega + \delta\omega_a)t)] \\ \dot{\tilde{c}}_2(t) &= i\frac{\Omega^*}{2}\tilde{c}_1(t)[\exp(-i(\omega - \delta\omega_a)t) + \exp(i(\omega + \delta\omega_a)t)] \end{aligned} \quad (\text{A.12})$$

where $\delta\omega_a = (\omega_{a1} - \omega_{a2})$ is the transition frequency of the atom. If the detuning of the driving field from the transition frequency ($\omega - \delta\omega_a$) is small, as is usually the case, then $(\omega + \delta\omega_a) \simeq 2\delta\omega_a$. This means that terms with this quantity in the exponent oscillate much more quickly and can be dropped. Integrating the differential equations in (A.12) gives the solutions for the probability amplitudes as

$$\begin{aligned}
c_1(t) &= \{c_1(0)[\cos(\frac{\Omega_{eff}t}{2}) - \frac{i\Delta}{\Omega_{eff}}\sin(\frac{\Omega_{eff}t}{2})] \\
&\quad + i\frac{\Omega_{eff}}{\Omega}c_2(0)\sin(\frac{\Omega_{eff}t}{2})\}\exp(\frac{i\Delta t}{2}) \\
c_2(t) &= \{c_2(0)[\cos(\frac{\Omega_{eff}t}{2}) + \frac{i\Delta}{\Omega_{eff}}\sin(\frac{\Omega_{eff}t}{2})] \\
&\quad + i\frac{\Omega_{eff}}{\Omega}c_1(0)\sin(\frac{\Omega_{eff}t}{2})\}\exp(\frac{-i\Delta t}{2})
\end{aligned} \tag{A.13}$$

where $\Delta = (\omega - \delta\omega_a)$ and $\Omega_{eff} = \sqrt{\Omega^2 + \Delta^2}$.

A.4 Derivation of the Stark Shift

In order to calculate the Stark shift, we consider again the equations of motion for the probability amplitudes in the RWA:

$$\begin{aligned}
i\dot{c}_1(t) &= \frac{\Omega}{2}\tilde{c}_2(t)\exp(i\Delta t) \\
i\dot{c}_2(t) &= \frac{\Omega^*}{2}\tilde{c}_1(t)[\exp(-i\Delta t)]
\end{aligned} \tag{A.14}$$

where, as before, $\Delta = (\omega - \delta\omega_a)$. In order to simplify the formulation of these equations, it is convenient to make the transformation

$$\begin{aligned}
c'_1(t) &= c_1(t)\exp(-i\Delta t) \\
c'_2(t) &= c_2(t)\exp(i\Delta t)
\end{aligned}
\tag{A.15}$$

This is equivalent to eliminating the time dependence of the interaction Hamiltonian by switching to a rotating frame. Substituting for $\tilde{c}_1(t)$ and $\tilde{c}_2(t)$ in (A.14) yields

$$\begin{aligned}
i\dot{c}'_1(t) &= \frac{1}{2}(\Delta c'_1(t) + \Omega c'_2(t)) \\
i\dot{c}'_2(t) &= \frac{1}{2}(\Omega c'_1(t) - \Delta c'_2(t))
\end{aligned}
\tag{A.16}$$

The convenience of making the transformation in (A.15) is that all the oscillating terms have now vanished. These equations can be written in matrix form as

$$i \begin{pmatrix} \dot{c}'_1 \\ \dot{c}'_2 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} \Delta & \Omega \\ \Omega^* & -\Delta \end{pmatrix} \begin{pmatrix} c'_1 \\ c'_2 \end{pmatrix}
\tag{A.17}$$

The eigenvalues of the matrix can be found by solving the characteristic equation $\det(A-\lambda I)$. This yields $\lambda = \pm \frac{1}{2}\sqrt{\Delta^2 + \Omega^2}$, which means that the eigenenergies of the system are given by

$$\frac{\hbar}{2}(\Delta \pm \sqrt{\Delta^2 + \Omega^2})
\tag{A.18}$$

which, for the far-detuned case where $\Delta \gg \Omega$, gives a total shift in energy of

$$\Delta E_{Stark} = \frac{\hbar\Omega^2}{2\Delta} \quad (\text{A.19})$$

A.5 Spontaneous Decay & the Optical Bloch Equations

The treatment considered so far can only be used to describe one part of the light-matter interaction, namely, absorption and stimulated emission. To describe spontaneous emission, a different treatment is needed. In spontaneous emission, the emitted photon has a random polarisation and direction. It is emitted into the vacuum field, which is outside of the atom plus field system (in other words, there is a loss of photons from the system), so the interaction cannot be described by the coherent evolution of states described in the previous section. It must instead be treated by using the density matrix formalism^[65], which considers a statistical mixture of states, and can be used to describe both pure and mixed states. Using this formalism, it is possible to derive the Optical Bloch Equations, which serve to model the population dynamics of the two-level system including decay.

The density matrix of a two level atom takes the form:

$$\rho = \begin{pmatrix} \rho_{11} & \rho_{12} \\ \rho_{21} & \rho_{22} \end{pmatrix} = \begin{pmatrix} c_1 c_1^* & c_1 c_2^* \\ c_2 c_1^* & c_2 c_2^* \end{pmatrix} \quad (\text{A.20})$$

The Optical Bloch equations are simply the evolution equations for the terms in the density matrix. They are derived by simply adding on a damping term to the evolution equations of the probability amplitudes considered before. They are given by:

$$\dot{\rho}_{11} = \gamma\rho_{22} + \frac{i}{2}(\Omega^* \tilde{\rho}_{21} - \Omega \tilde{\rho}_{12}) \quad (\text{A.21})$$

$$\dot{\tilde{\rho}}_{12} = -\left(\frac{\gamma}{2} + i\Delta\right)\tilde{\rho}_{12} + \frac{i}{2}\Omega^*(\rho_{22} - \rho_{11}) \quad (\text{A.22})$$

$$\dot{\tilde{\rho}}_{21} = -\left(\frac{\gamma}{2} - i\Delta\right)\tilde{\rho}_{12} + \frac{i}{2}\Omega(\rho_{11} - \rho_{22}) \quad (\text{A.23})$$

$$\dot{\rho}_{22} = -\gamma\rho_{22} + \frac{i}{2}(\Omega\tilde{\rho}_{12} - \Omega^*\tilde{\rho}_{21}) \quad (\text{A.24})$$

where $\tilde{\rho}_{ij} = \rho_{ij} \exp(-i\Delta t)$. It is noted that these equations only take into account decay due to spontaneous emission and do not take into account collisions or Doppler broadening. The solution to these equations in the steady state, i.e. taking $\dot{\rho} = 0$, gives the population of the excited state as

$$\rho_{22} = \frac{1}{2} \frac{\frac{\Omega^2}{2}}{\Delta^2 + \frac{\Omega^2}{2} + \frac{\gamma^2}{4}} \quad (\text{A.25})$$

A.6 The Light Force

The force exerted on an atom in an electromagnetic field can be found by calculating the expectation value of the force operator F . This is given (in one dimension) by

$$\langle F \rangle = \frac{d\langle p \rangle}{dt} = \frac{i}{\hbar} \langle [H, p] \rangle = \frac{i}{\hbar} \langle i\hbar \frac{dH}{dz} \rangle = -\langle \frac{dH}{dz} \rangle \quad (\text{A.26})$$

Inserting the interaction part of the Hamiltonian into (A.26) yields

$$F = e \frac{\partial}{\partial z} (\langle \mathbf{r} \cdot \mathbf{E}(\mathbf{r}, t) \rangle) \quad (\text{A.27})$$

where the gradient and the expectation value can be interchanged because the spatial variation of the electric field over the extent of the atom can be neglected in the dipole approximation. The expectation value of this matrix, which has only off-diagonal elements, can be written in terms of the optical coherences by recalling that $\langle \mathcal{O} \rangle = \text{Tr}(\rho \mathcal{O})$, so that

$$F = \hbar \left(\frac{\partial \Omega}{\partial z} \rho_{21}^* + \frac{\partial \Omega^*}{\partial z} \rho_{21} \right) \quad (\text{A.28})$$

where, as before, Ω is the Rabi frequency and the RWA has been used. $\frac{\partial \Omega}{\partial z}$ can be written in terms of the two quadrature components q_r and q_i , where $q_r + iq_i$ is the logarithmic derivative of Ω ^[66]. This gives

$$\frac{\partial \Omega}{\partial z} = (q_r + iq_i)\Omega \quad (\text{A.29})$$

where q_r corresponds to the gradient of the irradiating field's amplitude, whilst q_i corresponds to the gradient of its phase. The light force can be written in terms of these quadratures and the expression for ρ_{21} that was obtained from solving the Optical Bloch Equations, yielding

$$F = \frac{\hbar \frac{\Omega^2}{2}}{\Delta^2 + \frac{\Omega^2}{2} + \frac{\gamma^2}{4}} \left[-q_r \Delta + q_i \frac{\gamma}{2} \right] \quad (\text{A.30})$$

Appendix B

The Optical Molasses Technique

The Scattering Force introduced in section 2.4 is the principal mechanism at work in all laser cooling experiments. Lasers provide well-collimated monochromatic beams, so they are an ideal source for the radiation pressure needed to slow atoms. When an atom absorbs a photon from a counter-propagating laser beam, it receives a momentum kick in the opposite direction to its motion. As was discussed in section 2.4, spontaneously emitted photons go in all directions, so the net effect of many photon scattering events is to slow the atom down. The full expression for the Scattering Force in terms of the population of the excited state ρ_{22} is given by

$$F_{Scatt} = \hbar k \frac{\gamma}{2} \frac{\frac{\Omega^2}{2}}{\Delta^2 + \frac{\Omega^2}{2} + \frac{\gamma^2}{4}} \quad (\text{B.1})$$

or in terms of the Saturation Intensity,

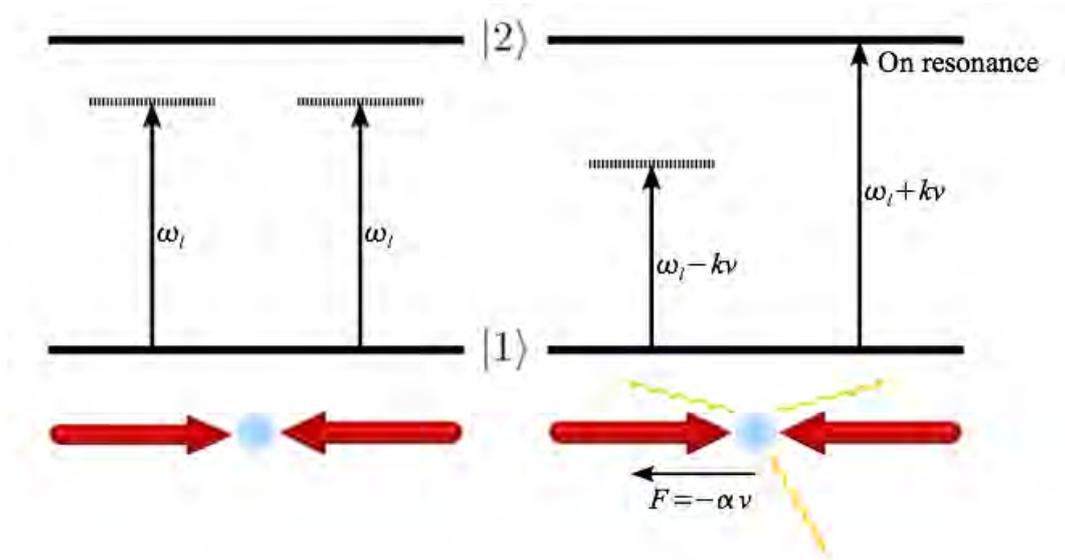


Figure B.1 A pair of slightly red-detuned counterpropagating laser beams incident on a stationary atom exert no net force, because the scattering is the same for each beam. If they are incident on a moving atom, however, the Doppler effect leads to more scattering in the direction opposite to the atom's velocity^[35]. The beam propagating in this direction is effectively brought into resonance with the moving atom. The Figure is taken from^[35].

$$F_{Scatt} = \hbar k \frac{\gamma}{2} \frac{\frac{I}{I_{sat}}}{1 + \frac{I}{I_{sat}} + 4 \frac{\Delta^2}{\gamma^2}} \quad (\text{B.2})$$

where now the detuning Δ is the difference between the laser frequency and the atomic transition taking into account the Doppler shift, i.e., $\Delta = \omega - \delta\omega_a + kv$. This is because in the rest frame of an atom moving with velocity \vec{v} along \vec{k} , photons moving towards it in the $-\vec{k}$ direction will have a frequency of $\omega - kv$ (they will be slightly red detuned), whilst photons moving towards it in the $+\vec{k}$ direction will have a frequency of $\omega + kv$ (they will be slightly blue detuned). As illustrated in Figure B.1, if an atom is moving to the right in a pair of counter-propagating beams that are red-detuned from the atomic transition in the first place, then in the atom's rest

frame, the effective frequency of the beam incident from the right will be closer to resonance, whilst the effective frequency of the beam incident from the left will be even further red detuned. This means that the atom will absorb more photons from the beam propagating in the direction opposite to its motion, and will consequently experience a force that slows it down. The net force on the atom can be calculated by taking the difference between the force from the left propagating beam and the force from the right propagating beam:

$$F_{Molasses} = F_{Scatt}(\omega - \delta\omega_a - kv) - F_{Scatt}(\omega - \delta\omega_a + kv) \approx -2\frac{\delta F}{\delta\omega}kv. \quad (\text{B.3})$$

This takes the form of a damping force with $\alpha = 2\frac{\delta F}{\delta\omega}k$:

$$F_{Molasses} = -\alpha v \quad (\text{B.4})$$

In order to cool a gas of atoms, three orthogonal pairs of such counterpropagating beams are needed, as atoms in a gas move in all directions. This configuration is shown in Figure B.2. This cooling technique, called Optical Molasses in reference to the damping effect of the beams, can slow a gas of atoms down to a limiting temperature called the Doppler Cooling Limit. This is the temperature at which the heating rate equals the cooling rate (heating arises due to fluctuations in the absorption and spontaneous emission process), and it is the lowest temperature to which atoms can be cooled using the Optical Molasses technique. It is given by

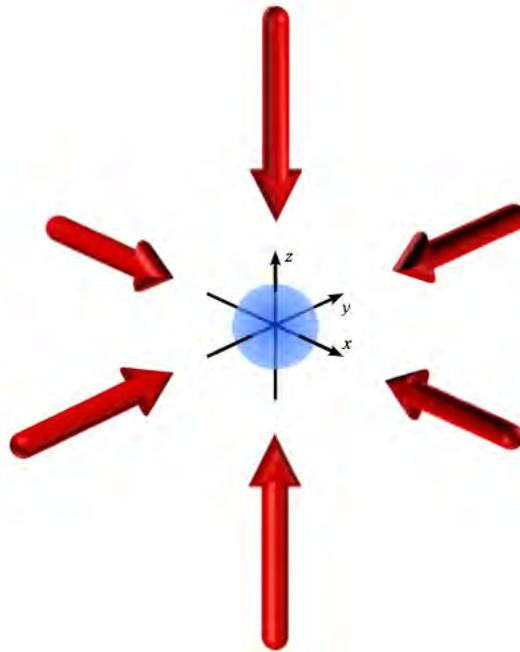


Figure B.2 Three counter-propagating pairs of red-detuned laser beams lead to an overall damping of the atomic motion. This is called an *Optical Molasses*. The Figure is taken from^[35].

$$T_D = \frac{\hbar\gamma}{2k_B} \quad (\text{B.5})$$

Counterintuitively, experiments found that cooling below this limit is indeed possible using an optical molasses; however, the two-level model of the atom discussed so far is not sufficient to explain this *sub-Doppler cooling*, as the effect relies on population transfer between different atomic sublevels.

Appendix C

Doppler Free Saturated Absorption Spectroscopy

In order to lock the lasers to the right hyperfine transitions, it is first necessary to have a reference signal in which the hyperfine structure of ^{87}Rb is resolved. Spectroscopy of a room temperature sample inside a glass cell is a relatively simple way to obtain the requisite atomic spectrum. Unfortunately, at room temperature, the linewidth of atomic spectra are limited by Doppler broadening. The atoms in the sample will have a distribution of velocities given by the Maxwell Boltzmann velocity distribution

$$f(v) dv = \sqrt{\frac{M}{2\pi k_B T}} \exp\left(\frac{-Mv^2}{2k_B T}\right) dv, \quad (\text{C.1})$$

and this velocity distribution leads to an inhomogeneous, or Doppler, broadening

of the atomic linewidth. It is inhomogeneous because it is different for each atom: atoms traveling in different velocity classes see radiation at different effective frequencies in their own rest frame. Radiative broadening, on the other hand, leads to a homogeneous broadening of the atomic linewidth, because the spontaneous decay rate is the same for all atoms of the same species. The total line shape function of the atom is a convolution of the Gaussian function $f(v)$ describing the inhomogeneous broadening brought about by the motion of the atom, and the Lorentzian function describing the homogeneous broadening brought about by spontaneous decay.

The widening of the atomic spectrum brought about by inhomogeneous Doppler broadening can be overcome using the technique of Doppler-Free Saturated Absorption Spectroscopy. In this technique, a weak probe beam and a strong pump beam derived from the same source and having the same frequency ω go through a sample cell in opposite directions. Assuming that all the atoms in the cell start out in the ground state, the total number of atoms N will equal the number of atoms in the ground state N_1 . The number of atoms in a particular velocity class will then be given by $N_1(v) = f(v)N_1$ with $f(v)$ given by equation C.1. Interaction with the weak probe beam leaves this distribution practically unchanged, as almost all of the atoms stay in the ground state. Interaction with the strong pump beam, however, excites a large number of the atoms in the appropriate velocity class, i.e. the atoms with velocity $v = (\omega - \delta\omega_a)/k$, to the excited state. This pump beam saturates the transition, i.e. it makes the difference in population between the ground and excited

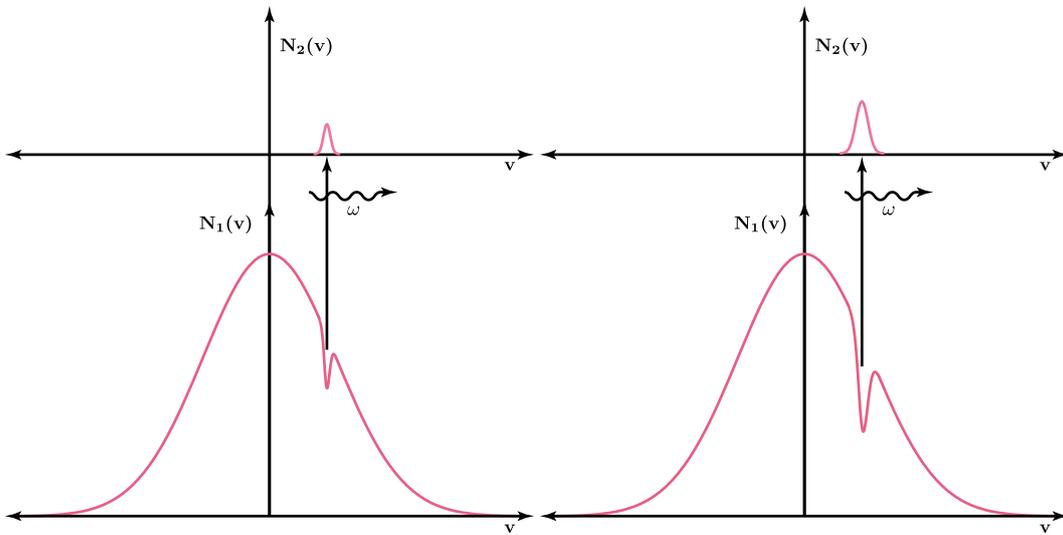


Figure C.1 Saturation of Absorption.

state $N_2 - N_1$ tend to zero. This is shown in Figure C.1.

As long as the pump and probe beams are detuned from the transition frequency, they interact with atoms in different velocity classes: the pump beam interacts with atoms having velocity $v = (\omega - \delta\omega_a)/k$, whilst the probe beam, which propagates in the opposite direction, interacts with atoms having velocity $v = -(\omega - \delta\omega_a)/k$. If the two beams are at the transition frequency, however, then the saturation caused by the pump beam prevents the atomic population from absorbing any of the probe beam. As shown in Figure C.2, this leads to a peak at $\delta\omega_a$ in the Doppler broadened absorption profile that appears when the transmitted probe intensity is plotted as a function of ω .

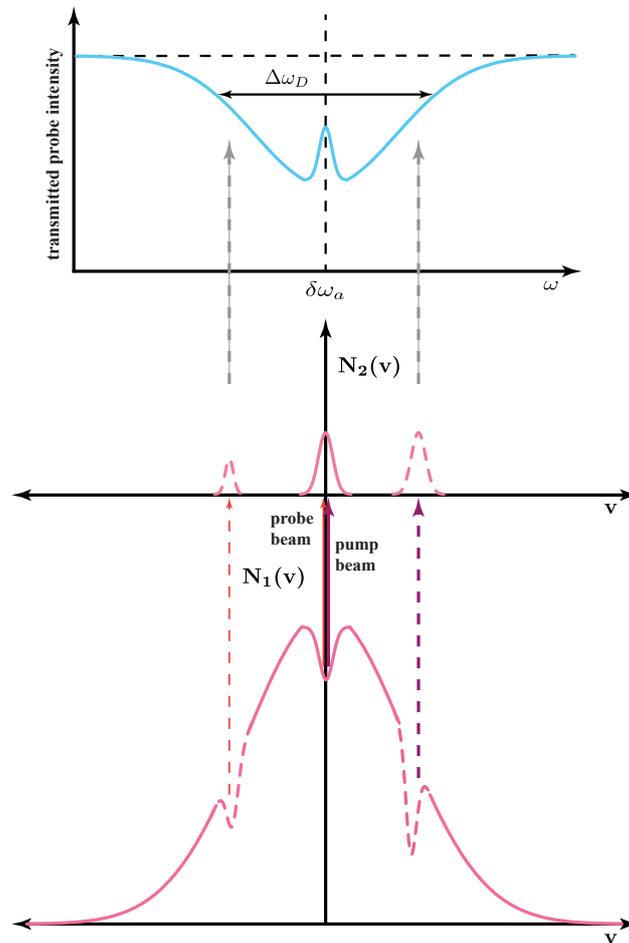


Figure C.2 Saturated Absorption Spectroscopy

For a hyperfine spectrum, peaks like this occur at every allowed transition as well as at frequencies halfway between transitions with shared energy levels. These so-called *crossover resonances* occur when the pump beam reduces the probe absorption at two such transition frequencies simultaneously. As illustrated by Figure C.3, this can only occur when the detuning of the pump/probe beam is midway between the two transition frequencies. A pump/probe laser with frequency ω_{12} saturates the

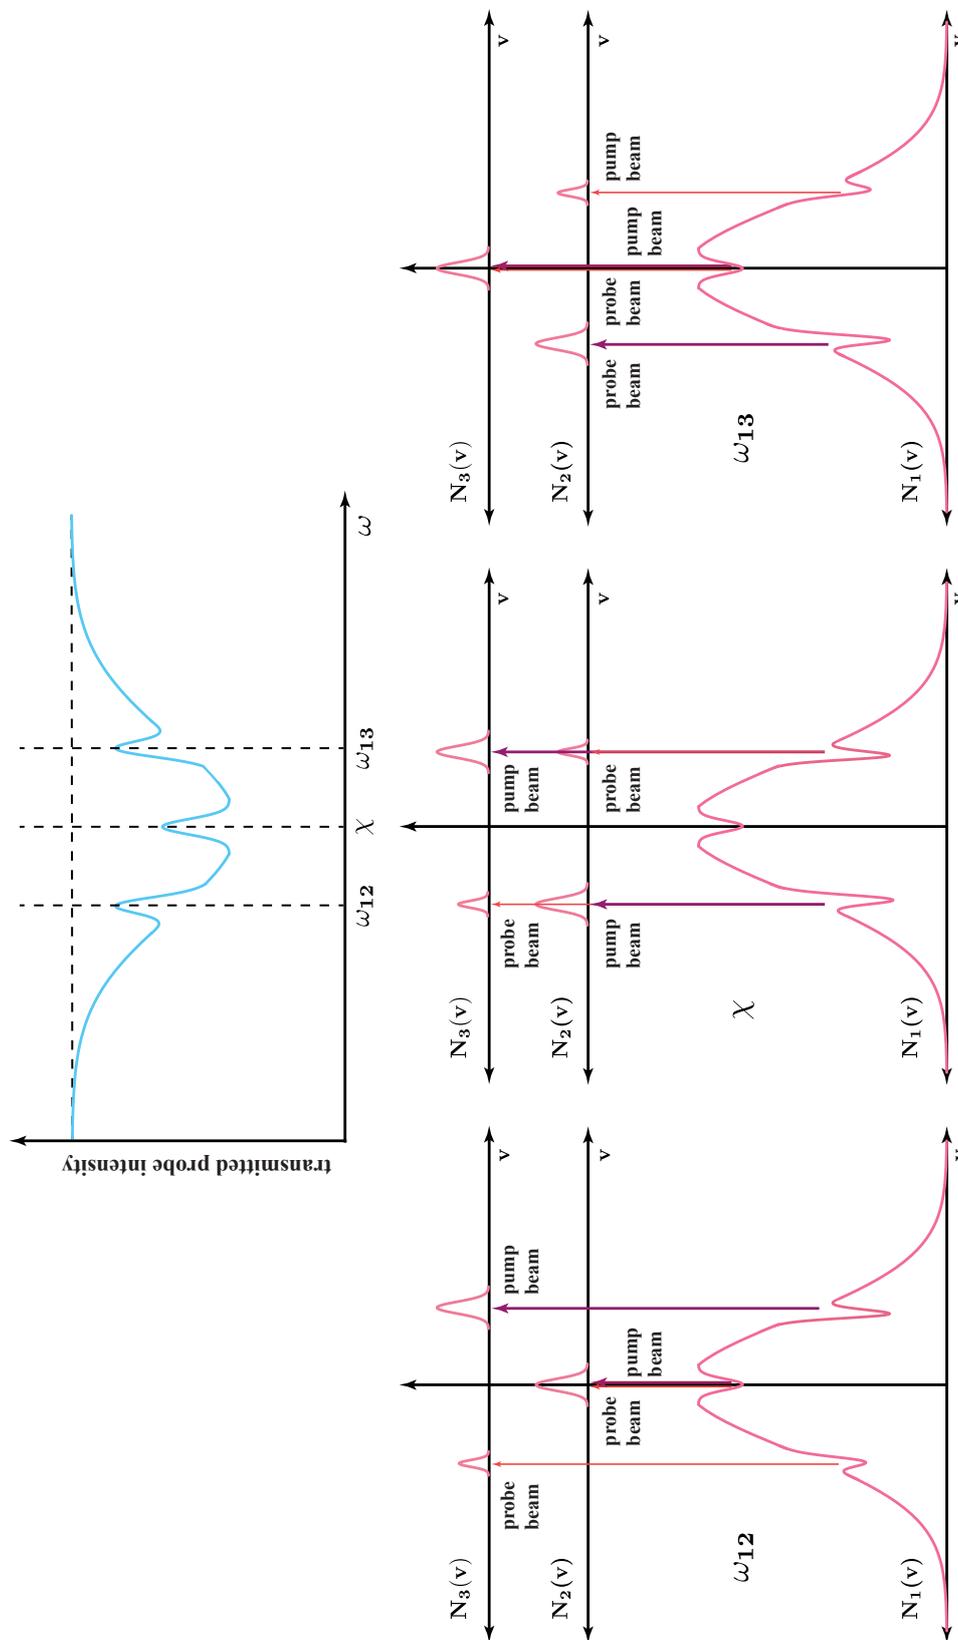


Figure C.3 crossovers

transition between level 1 and level 2, but leaves the transition between level 1 and level 3 unaffected. Conversely, a pump/probe laser with frequency ω_{13} saturates the transition between level 1 and level 3, but leaves the transition between level 1 and level 2 unaffected. If the pump/probe beam has frequency $\chi = (\omega_{12} + \omega_{13})/2$, however, then the pump beam saturates both transitions. This leads to peak in the transmitted probe intensity at χ which can serve as a stable reference signal for locking.

In our setup, we use a mirror, a polarising beam splitter and a $\frac{\lambda}{4}$ plate to split the probe beam off from the pump beam and send it onto a detector. As can be seen in Figure 5.1, we use a $\frac{\lambda}{4}$ plate to rotate the polarisation of the beam before and after reflection from the mirror, such that the returning probe beam has the opposite polarisation to the incoming pump beam and is therefore transmitted through the polarising beam splitter, allowing it to reach the detector. By scanning the laser frequency, a trace of transmitted intensity as a function of frequency can be monitored on an oscilloscope. We use the crossover between the $F = 2$ to $F' = 1$ and the $F = 2$ to $F' = 3$ transitions as a reference to lock the cooling laser, and the crossover between the $F = 1$ to $F' = 1$ and the $F = 1$ to $F' = 2$ transition as a reference to lock the repump laser.

Appendix D

Pound Drever Hall Stabilisation

A small change in a laser's frequency will produce a small change in its measured intensity. Therefore, the most intuitive way to lock a laser would be to set a target level on one side of a reference signal (like an atomic crossover resonance, or a Fabry-Pérot resonance), measure any intensity deviations from that point, and feed them back into the laser in order to hold the intensity constant. One problem with this approach is that it is impossible to distinguish between drifts in the frequency of the laser and drifts in its intensity, because both are measured as changes in intensity. Also, frequency fluctuations can push the laser to the opposite side of a chosen resonance, where the measured deviation would indicate a correction in the wrong direction (this is because the measured intensity is symmetric about zero).

This can be resolved by looking at a signal which is proportional to the derivative of the measured intensity, which is antisymmetric about resonance: $\frac{dI}{d\omega}$ is positive

above resonance, and negative below it. By varying the frequency (or phase) of the laser sinusoidally, it is possible to tell whether the reference point is above or below the resonance: above resonance, the measured intensity will be in phase with the sinusoidal oscillation, whilst below resonance, it will be 180° out of phase. This is the basic idea behind the Pound Drever Hall (PDH) technique.

In practice, it is the phase of the laser that is modulated (at 20MHz in our case). This modulation produces sidebands that have a definite phase relationship to the measured probe beam. The modulated field can be written as

$$E_{inc} = E_0 \exp(i(\omega t + \beta \sin(\Omega t))) \quad (\text{D.1})$$

and can be expanded using Bessel function as

$$\begin{aligned} E_{inc} &\approx [J_0(\beta) + 2iJ_1(\beta)\sin(\Omega t)]\exp(i\omega t) \\ &= E_0[J_0(\beta)\exp(i\omega t) + J_1(\beta)\exp(i(\omega + \Omega)t) - J_1(\beta)\exp(i(\omega - \Omega)t)] \end{aligned} \quad (\text{D.2})$$

The three terms in equation (D.2) correspond to a carrier beam with frequency ω and two sidebands with frequency $\omega + \Omega$, where Ω is the modulation frequency and β is the amplitude of the modulation. Only the first two terms in the Bessel expansion have been kept: this approximation is valid because when β is small, almost all of the power is in the carrier and the first order sidebands. Interfering these sidebands with the measured probe beam leads to a beat pattern whose phase

can be measured, in turn giving a measure of the probe beam's phase. In the experiment this is accomplished by using a mixer to compare the output of the detector in the Saturated Absorption Spectroscopy setup with the local oscillator. The signal at the output of a mixer is a product of its inputs, and takes the form

$$\sin(\Omega t + \phi_1)\sin(\Omega' t + \phi_2) = \frac{1}{2}[\cos((\Omega - \Omega')t + (\phi_1 - \phi_2)) - \cos((\Omega + \Omega')t + (\phi_1 + \phi_2))]. \quad (\text{D.3})$$

In this case, $\Omega = \Omega'$, so the first term on the right hand side of equation (D.3) is reduced to a term that depends on the phase difference: $\cos(\phi_1 - \phi_2)$. A phase delay is used to single out the in phase dc term where $\phi_1 = \phi_2$. A low pass filter is then used to pick out the subsequent dc term at the output of the mixer: this term is proportional to the derivative of the probe beam's intensity, and is the signal that is used to lock the laser.

The lock is maintained by a Proportional Integral Derivative (PID) regulator, which derives an error signal from the difference between the PDH signal and a set value to one side of the desired PDH peak, and uses it as feedback for the grating. The PID feedback signal is actually composed of three parts: a Proportional part, an Integral part, and a Derivative part, hence the name. The Proportional part is simply proportional to the magnitude of the error. The Integral part is proportional to the magnitude and duration of the error, and has the slowest response. The Derivative part is proportional to the rate of change of the error, and has the fastest

response. In the experiment, the relative magnitude of these corrections are adjusted empirically. The PDH and PID control modules used are both from Toptica, the PDD 110 and the PID 110.

Appendix E

The Electron Multiplying Charge Coupled Device Camera

In a normal CCD acquisition, the electronic charge that arises after the image area is exposed to light is shifted down to a masked region and then read out by first shifting the charge vertically into a readout register, and then shifting it horizontally into the output node of the amplifier. The extra gain register in an EMCCD camera serves to amplify the charge exiting the readout register by means of impact ionisation. This occurs when a charge has enough energy to create a secondary electron, and is accomplished by placing a clocked electrode next to a fixed voltage electrode, and setting the voltage of this second electrode much higher than necessary for charge transfer alone (see Figure E.1). The impact ionisation of the transferred electrons results in more electrons, thus allowing for the signal to be amplified on the sensor itself before reaching the output amplifier and being digitised by the Analog to Dig-

ital Converter (ADC). This means that the signal is amplified *before* the addition of the noise generated by the readout electronics (this Read Noise is mainly produced by the the output amplifier, but there is also a small contribution from the ADC).

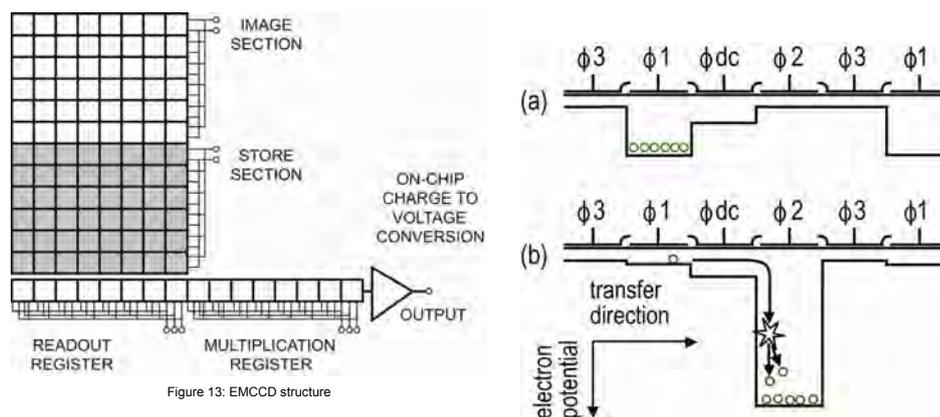


Figure E.1 The left hand side of the figure shows the additional Multiplication Register that can be found on a typical EMCCD chip. The right hand side of the figure illustrates the basic operation of this Multiplication Register. Both diagrams were taken from the camera's documentation.

The gain register of an EMCCD camera does introduce an additional form of noise called Multiplicative Noise which arises due to the statistical variation in the number of electrons generated from an initial charge in the register. However, because multiplicative noise increases the deviation of the signal around the mean value but does not reduce the average signal detected, it can be considered as a form of shot noise, and simply multiplied by the latter when calculating the overall noise of the sensor. Air cooling the camera down to -80° helps to eliminate darkcurrent noise from thermal electrons. This cooling is especially important in EMCCD cameras, as

thermally generated electrons will be amplified in the same way as photon-generated electrons. Temperatures as low as -100° can be achieved by water cooling.

The EMCCD camera can be read out at a rate of 35, 27, or 13 MHz. This readout rate corresponds to the horizontal readout rate at which charge is readout from the shift register. Obviously, the higher the readout rate, the higher the frame rate that can be achieved with the camera. The EM gain can be set from 0 to 300, and the EMCCD pixels can be binned to give a better S/N ratio, as fewer pixels are read out and there are therefore less contributions to noise. 1x1, 2x2, 4x4, 8x8, and 16x16 binning can be chosen. The Andor documentation gives a value of 65% for the quantum efficiency (QE) of the camera at 780 nm (the QE is defined as the probability that an electron will be produced on the sensor given that a photon arrives at its surface). We tested this value by shining a weak laser beam onto the camera and using a separate powermeter to determine the number of photons incident on the camera in one second, and found the QE for our camera to have an actual value of 43%. We attenuated the beam by using several OD filters, using the least attenuation possible so as to avoid the lower limit on sensitivity of the powermeter and thus ensure an accurate reading.

The pre-amplifier gain of the camera can be set to one of three values: 1, 1.9, or 3.8. Table E.2 shows the number of electrons per count for each of these settings. The electrons per count for unity gain is obtained by calibrating the camera signal

Readout Rate	Amplifier Gain	Electrons per Count
35 MHz	1	5.22
35 MHz	1.9	2.52
35 MHz	3.8	1.41
27 MHz	1	4.89
27 MHz	1.9	2.59
27 MHz	3.8	1.26
13 MHz	1	4.42
13 MHz	1.9	2.48
13 MHz	3.8	1.09

Figure E.2 Electrons per count for different values of the pre-amplifier gain. This value varies slightly depending on the readout rate chosen.

against its own shot noise, and corresponds to the constant of proportionality between the number of electrons and the value generated by the ADC. Traditionally, pre-amplifier gain selection in CCD cameras controls the trade off between the dynamic range of the sensor, and its S/N ratio. A higher pre-amplifier gain means fewer electrons/count, resulting in less noise and therefore a lower S/N ratio, whilst a lower pre-amplifier gain allows for a higher dynamic range. Most low-light applications are not limited by well capacity (i.e. they don't need a large dynamic range) as long as adequate EM gain settings are used, so it sufficient to use the highest value pre-amplifier gain setting. A high pre-amplifier gain setting is desirable because it restricts how much charge is built up on the sensor, and thus helps to minimise EM ageing.

The camera can be operated in an Internal Trigger mode, in an External Trigger mode, and in an External Exposure mode in which the start and end of exposure

is defined completely by the trigger pulse. We use mostly this latter option, and set the pulse length to match the flash of the light sheet so that no light falls on the sensor outside of the exposure window. In order to ensure that any remaining scattered light falling on the image area before the start of the exposure does not contribute to the measured signal, the CCD stays in a “Keep Clean” mode until the start of an acquisition is signalled by the rising edge of the external trigger pulse. In the Keep Clean mode, the image and storage area are kept clear by constantly shifting and reading out any accumulated charge.

List of Publications

Bacterial Metapopulations in nanofabricated landscapes. Juan E. Keymer, Peter Galajda, Cecilia Muldoon, Sungsu Park, and Robert H. Austin. Proceedings of the National Academy of Sciences, Volume 103, Issue 46, Pages 17290-17295 (2006)

Coherent imaging of extended objects. Edouard Brainis, Cecilia Muldoon, Lukas Brandt, and Axel Kuhn. Optics Communications, Volume 282, Issue 4, Pages 465-472 (2009)

Spatial light modulators for the manipulation of individual atoms. Lukas Brandt, Cecilia Muldoon, Tobias Thiele, Jian Dong, Edouard Brainis, and Axel Kuhn. Applied Physics B, Volume 102, Issue 6, Pages 443-450 (2011)

Control and Manipulation of Cold Atoms in Optical Tweezers. Cecilia Muldoon, Lukas Brandt, Jian Dong, Dustin Stuart, Edouard Brainis, Matthew Himsforth, and Axel Kuhn. arXiv:1109.0657 (2011)

Towards a Scalable Dipole-Trapping Scheme for Neutral Atoms. Lukas Brandt, Cecilia Muldoon, Edouard Brainis, and Axel Kuhn. Conference on Lasers and Electro-Optics/Quantum Electronics and Laser Science Conference and Photonic Applications Systems Technologies 2008 Technical Digest (Optical Society of America, Washington, DC, 2008), JThA105

Optical Tweezers for manipulating Single Atoms. Lukas Brandt, Cecilia Muldoon, Edouard Brainis, and Axel Kuhn. Conference on Lasers and Electro-Optics Europe and the European Quantum Electronics Conference (CLEO/Europe - EQEC 2009), Conference digest JSS2.1.

Implementation of Atom-Photon Interfaces for Quantum Networking. Lukas Brandt, Cecilia Muldoon, Tobias Thiele, Jerome Dilley, Peter Nisbet, Gunnar Langfahl-Klabes, and Axel Kuhn. Conference on Lasers and Electro-Optics (CLEO) and the Quantum Electronics and Laser Science Conference (QELS) (Optical Society of America, Washington, DC, 2010), QThA2.

Bibliography

- [1] D. P. DiVincenzo. The physical implementation of quantum computation. *Fortschr. Phys.*, 48:771–783, 2000.
- [2] D. Deutsch. Quantum theory, the church-turing principle and the universal quantum computer. *Proc. R. Soc. Lond. A*, 400:97–117, 1985.
- [3] A. Politi, M. J. Cryan, J. G. Rarity, S. Yu, and J. L. O’Brien. Silica-on-silicon waveguide quantum circuits. *Science*, 320(5876):646–649, 2008.
- [4] P. Kwiat and G. Milburn. Optical approaches to quantum information processing and quantum computing. Technical report, US Government Advanced Research and Development Activity, 2004.
- [5] J. Benhelm, G. Kirchmair, C. F. Roos, and R. Blatt. Towards fault-tolerant quantum computing with trapped ions. *Nat. Phys.*, 4(6):463–466, 2008.
- [6] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T. C. Allcock, M. J. Curtis, G. Imreh, J. A. Sherman, D. N. Stacey, A. M. Steane, and D. M. Lucas. High-fidelity readout of trapped-ion qubits. *Phys. Rev. Lett.*, 100(20):200502–200506, 2008.
- [7] K. Singer, U. Poschinger, M. Murphy, P. Ivanov, F. Ziesel, T. Calarco, and F. Schmidt-Kaler. Colloquium: Trapped ions as quantum bits: Essential numerical tools. *Rev. Mod. Phys.*, 82(3):2609–2632, 2010.
- [8] D. Kielpinski, C. R. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417(6890):709–711, 2002.
- [9] J. Reichel and V. Vuletić. *Atom Chips*. Wiley-VCH, 2011.
- [10] R. Folman, P. Krüger, J. Schmiedmayer, J. Denschlag, and C. Henkel. Microscopic atom optics: From wires to an atom chip. *Adv. At. Mol. Opt. Phys.*, 48:263–356, 2002.
- [11] W. Alt, D. Schrader, S. Kuhr, M. Müller, V. Gomer, and D. Meschede. Single atoms in a standing-wave dipole trap. *Phys. Rev. A*, 67(3):033403–033410, 2003.

- [12] N. Schlosser, G. Reymond, I. Protsenko, and P. Grangier. Sub-poissonian loading of single atoms in a microscopic dipole trap. *Nature*, 411(6841):1024–1027, 2001.
- [13] J. Volz, M. Weber, D. Schlenk, W. Rosenfeld, J. Vrana, K. Saucke, C. Kurtsiefer, and H. Weinfurter. Observation of entanglement of a single photon with a trapped atom. *Phys. Rev. Lett.*, 96(3):030404–030408, 2006.
- [14] J. Beugnon, C. Tuchendler, A. Gaetan, Y. Miroshnychenko, H. Marion, A. M. Lance, M. P. Jones, Y. R. Sortais, G. Messin, A. Browaeys, and P. Grangier. Two-dimensional transport and transfer of a single atomic qubit in optical tweezers. *Nat. Phys.*, 3(10):696–699, 2007.
- [15] M. Karski, L. Förster, J. Choi, A. Steffen, W. Alt, D. Meschede, and A. Widera. Quantum walk in position space with single optically trapped atoms. *Science*, 325(5937):174–177, 2009.
- [16] T. Grünzweig, A. Hilliard, M. McGovern, and M. F. Andersen. Near-deterministic preparation of a single atom in an optical microtrap. *Nat. Phys.*, 6(12):951–954, 2010.
- [17] 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(7311):68–72, 2010.
- [18] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner. A quantum gas microscope for detecting single atoms in a hubbard-regime optical lattice. *Nature*, 462(7269):74–77, 2009.
- [19] C. Weitenberg, M. Endres, J. F. Sherson, M. Cheneau, P. Schausz, T. Fukuhara, I. Bloch, and S. Kuhr. Single-spin addressing in an atomic mott insulator. *Nature*, 471(7338):319–324, 2011.
- [20] S. Bergamini, B. Darquié, M. Jones, L. Jacubowicz, A. Browaeys, and P. Grangier. Holographic generation of microtrap arrays for single atoms by use of a programmable phase modulator. *JOSA B*, 21(11):1889–1894, 2004.
- [21] V. Boyer, R. M. Godun, G. Smirne, D. Cassettari, C. M. Chandrashekar, A. B. Deb, Z. J. Laczik, and C. J. Foot. Dynamic manipulation of bose-einstein condensates with a spatial light modulator. *Phys. Rev. A*, 73(3):031402–031406, 2006.
- [22] F. K. Fatemi, M. Bashkansky, and Z. Dutton. Dynamic high-speed spatial manipulation of cold atoms using acousto-optic and spatial light modulation. *Opt. Express*, 15(6):3589–3596, 2007.
- [23] 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(4):043030–043039, 2009.

- [24] 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(6):060308–060302, 2010.
- [25] 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(17):170502–170506, 2010.
- [26] C. MacAulay and A. Dlugan. Use of digital micro mirror devices in quantitative microscopy. *Proc. SPIE*, 3260:201–206, 1998.
- [27] Q. S. Hanley, P. J. Verveer, M. J. Gemkow, D. Arndt-Jovin, and T. M. Jovin. An optical sectioning programmable array microscope implemented with a digital mirror device. *J. of Microscopy*, 196(3):317–331, 1999.
- [28] A. Ashkin and J. M. Dziedzic. Optical trapping and manipulation of viruses and bacteria. *Science*, 235(4795):1517–1520, 1987.
- [29] 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(6):065038–065061, 2010.
- [30] C. Marr, A. Beige, and G. Rempe. Entangled state preparation via dissipation-assisted adiabatic passages. *Phys. Rev. A*, 68(3):033817–033828, 2003.
- [31] 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.
- [32] C. Weitenberg, S. Kuhr, K. Mølmer, and J. F. Sherson. Quantum computation architecture using optical tweezers. *Phys. Rev. A*, 84(3):032331, 2011.
- [33] 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 gate. *Phys. Rev. Lett.*, 104(010503-010507), 2010.
- [34] C. S. Adams and E. Riis. Laser cooling and trapping of neutral atoms. *Prog. Quant. Electr.*, 21(1):1–79, 1997.
- [35] C. J. Foot. *Atomic Physics*. Oxford University Press, 2005.
- [36] 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(3):030901–030905, 2004.
- [37] J. Reichel. Microchip traps and bose-einstein condensation. *Appl. Phys. B*, 74(6):469–487, 2002.

- [38] N. Schlosser, G. Reymond, and P. Grangier. Collisional blockade in microscopic optical dipole traps. *Phys. Rev. Lett.*, 89(2):023005–023009, 2002.
- [39] E. Brainis, C. Muldoon, L. Brandt, and A. Kuhn. Coherent imaging of extended objects. *Opt. Comm.*, 282(4):465–472, 2009.
- [40] L. Brandt, C. Muldoon, T. Thiele, J. Dong, E. Brainis, and A. Kuhn. Spatial light modulators for the manipulation of individual atoms. *Appl. Phys. B*, 102(3):443–450, 2011.
- [41] 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(1):013406–013419, 2000.
- [42] 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(18):3777–3780, 2000.
- [43] I. Dotsenko, W. Alt, M. Khudaverdyan, S. Kuhr, D. Meschede, Y. Miroshnychenko, D. Schrader, and A. Rauschenbeutel. Submicrometer position control of single trapped neutral atoms. *Phys. Rev. Lett.*, 95(3):033002–033006, 2005.
- [44] M. Mücke, E. Figueroa, J. Bochmann, C. Hahn, K. Murr, S. Ritter, C. J. Villas-Boas, and G. Rempe. Electromagnetically induced transparency with single atoms in a cavity. *Nature*, 465(7299):755–758, 2010.
- [45] T. Bondo, M. Hennrich, T. Legero, G. Rempe, and A. Kuhn. Time-resolved and state-selective detection of single freely falling atoms. *Opt. Comm.*, 264(2):271–277, 2006.
- [46] A. Fuhrmanek, Y. R. P. Sortais, P. Grangier, and A. Browaeys. Imaging of a single atom in a time of flight experiment. *New J. Phys.*, 12(5):053028–053031, 2010.
- [47] A. Fuhrmanek, Y. R. P. Sortais, P. Grangier, and A. Browaeys. Measurement of the atom number distribution in an optical tweezer using single-photon counting. *Phys. Rev. A*, 82(2):023623–023629, 2010.
- [48] M. Born and E. Wolf. *Principles of Optics*. Cambridge University Press, seventh ed. edition, 2006.
- [49] J. W. Goodman. *Introduction to Fourier Optics*. Roberts & Company, third ed. edition, 2005.
- [50] P. Dumontet. Sur la correspondance objet-image en optique. *Optica Acta* 2, 2, 1955.
- [51] D. A. Tichenor and J. W. Goodman. Coherent transfer function. *J. Opt. Soc. Am.*, 62(2):293–295, 1972.

- [52] T. Ito and S. Okazaki. Pushing the limits of lithography. *Nature*, 406(6799):1027–1031, 2000.
- [53] J. F. Heanue, M. C. Bashaw, and L. Hesselink. Volume holographic storage and retrieval of digital data. *Science*, 265(5173):749–752, 1994.
- [54] L. Brandt, C. Muldoon, E. Brainis, and A. Kuhn. Towards a scalable dipole-trapping scheme for neutral atoms. In *CLEO/QELS-Phast Technical Digest*, volume JThA1005. OSA, 2008.
- [55] D. Meschede and A. Rauschenbeutel. Manipulating single atoms. *Adv. At. Mol. Opt. Phys.*, 53:75–104, 2006.
- [56] Y. Miroshnychenko, I. Dotsenko, W. Alt, L. Foerster, M. Khudaverdyan, D. Meschede, D. Schrader, and A. Rauschenbeutel. Quantum engineering: An atom-sorting machine. *Nature*, 442(7099):151, 2006.
- [57] G. Bonnet. Introduction to metaxial optics i. *Ann. Télécommunic.*, 33:143–165, 1978.
- [58] G. Bonnet. Introduction to metaxial optics ii. *Ann. Télécommunic.*, 33:225–243, 1978.
- [59] G. Bonnet. La diffraction metaxiale : Doctrine unitaire pour l’optique, l’acoustique et la radio-electricité, dans le langage des signaux. In *Septième colloque sur le traitement du signal et ses applications*, volume 54, pages 1–8. GRETSI, Groupe d’Etudes du Traitement du Signal et des Images, 1979.
- [60] E. Hecht. *Optics*. Addison Wesley, fourth ed. edition, 2002.
- [61] C. Muldoon, L. Brandt, J. Dong, D. Stuart, E. Brainis, M. Himsworth, and A. Kuhn. Control and manipulation of cold atoms in optical tweezers. *arXiv*, arXiv:1109.0657, 2011.
- [62] M. O. Scully and M. S. Zubairy. *Quantum Optics*. Cambridge University Press, 1997.
- [63] A. Beige. Introductory quantum optics. Lecture Notes, 2007.
- [64] M. Lukin. Modern atomic and optical physics ii. Lecture Notes, 2009.
- [65] L. Allen and J. H. Eberly. *Optical resonances and two-level atoms*. Wiley, 1975.
- [66] H. J. Metcalf and P. van der Straten. *Laser Cooling and Trapping*. Springer, 1999.