

ULTRACOLD ATOMS IN FLEXIBLE HOLOGRAPHICS TRAPS

David Bowman

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews



2018

Full metadata for this item is available in
St Andrews Research Repository
at:
<http://research-repository.st-andrews.ac.uk/>

Please use this identifier to cite or link to this item:
<http://hdl.handle.net/10023/16293>

This item is protected by original copyright

Ultracold Atoms in Flexible Holographic Traps

David Bowman



University of
St Andrews

This thesis is submitted in partial fulfilment for the degree of
Doctor of Philosophy (PhD)
at the University of St Andrews

April 2018

Abstract

This thesis details the design, construction and characterisation of an ultracold atoms system, developed in conjunction with a flexible optical trapping scheme which utilises a Liquid Crystal Spatial Light Modulator (LC SLM). The ultracold atoms system uses a hybrid trap formed of a quadrupole magnetic field and a focused far-detuned laser beam to form a Bose-Einstein Condensate of 2×10^5 ^{87}Rb atoms. Cold atoms confined in several arbitrary optical trapping geometries are created by overlaying the LC SLM trap on to the hybrid trap, where a simple feedback process using the atomic distribution as a metric is shown to be capable of compensating for optical aberrations.

Two novel methods for creating flexible optical traps with the LC SLM are also detailed, the first of which is a multi-wavelength technique which allows several wavelengths of light to be smoothly shaped and applied to the atoms. The second method uses a computationally-efficient minimisation algorithm to create light patterns which are constrained in both amplitude and phase, where the extra phase constraint was shown to be crucial for controlling propagation effects of the LC SLM trapping beam.

Candidate's declaration

I, David Bowman, do hereby certify that this thesis, submitted for the degree of PhD, which is approximately 30,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree.

I was admitted as a research student at the University of St Andrews in August 2013.

I received funding from an organisation or institution and have acknowledged the funder(s) in the full text of my thesis.

Date 26/04/18

Signature of candidate

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date

Signature of supervisor

Permission for publication

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand, unless exempt by an award of an embargo as requested below, that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, David Bowman, confirm that my thesis does not contain any third-party material that requires copyright clearance.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

Printed copy

No embargo on print copy.

Electronic copy

No embargo on electronic copy.

Date 26/04/18

Signature of candidate



Date

Signature of supervisor



Underpinning Research Data or Digital Outputs

Candidate's declaration

I, David Bowman, understand that by declaring that I have original research data or digital outputs, I should make every effort in meeting the University's and research funders' requirements on the deposit and sharing of research data or research digital outputs.

Date 26/04/18

Signature of candidate

Permission for publication of underpinning research data or digital outputs

We understand that for any original research data or digital outputs which are deposited, we are giving permission for them to be made available for use in accordance with the requirements of the University and research funders, for the time being in force.

We also understand that the title and the description will be published, and that the underpinning research data or digital outputs will be electronically accessible for use in accordance with the license specified at the point of deposit, unless exempt by award of an embargo as requested below.

The following is an agreed request by candidate and supervisor regarding the publication of underpinning research data or digital outputs:

No embargo on underpinning research data or digital outputs.

Date 26/04/18

Signature of candidate

Date

Signature of supervisor

Acknowledgements

I would first like to thank my supervisor Dr Donatella Cassettari who gave me the opportunity to work with her on this exciting project, even though I had very little prior knowledge of atomic physics. You have been a wonderful mentor to me and I have learnt a great deal from you. Remarkably, my often pessimistic attitude never seemed to effect you, and I had a fantastic time over the past four years thanks to you and your ever-vibrant personality.

A huge amount of guidance and support also came from Dr Graham Bruce, the previous post-doc of our group. You have helped me immensely throughout the project and have made me not only become a better physicist, but a better person. Even after you had left the group, you were always quick to offer any advice or support (especially with the writing of this thesis!) and I would not have enjoyed this project anywhere near as much as I have if it weren't for you.

Working in a physics lab can bring much hardship and annoyance, but thankfully I've had the pleasure of being able to share them with Phil Ireland, my fellow PhD student. Despite the hardships, you were always calm and friendly with me and together we pushed through many problems, always helping each other out and achieving much. You have provided some much needed company in what would have been a far lonelier project without you.

Whatever successes I may have in the future will largely be thanks to the three group members I have mentioned, whom I have learnt so much from and have become my very good friends. But there are also many others who I must offer my thanks to. My friends who are fellow members of the Condensed Matter Centre for Doctoral Training have helped make the PhD thoroughly enjoyable (especially with the weekly badminton!), even managing to make the rather stressful moments become far more tolerable. There are also the many students our group has taken on as part of their degrees, all of whom were a joy to work with and often provided a much needed expansion to our little cold atoms group.

I would also like to thank Professor Kishan Dholakia for lending us an IPG laser for use in our lab (for what I must admit is quite a long time now), as well as Toby Scrivener and Peter Collins at Laser2000 for loaning us the P1920 SLM which we used for one of our papers. It was also a joy to collaborate with Dr Tiffany Harte on the conjugate gradient minimisation work, as well as Dr Elmar Haller whose imaging and data analysis code we used in our cold atoms system. Finally, the build of our experiment would not have been achieved without the help of the in-house technicians, whom I would like to thank for their excellent craftsmanship.

My final thanks goes to my family who have given me nothing but love and support. I can only hope that you understand how important it has been knowing that you have always been there for me.

List of Publications and Presentations

Journal Publications

- D. Bowman, T. L. Harte, V. Chardonnet, C. D. Groot, S. J. Denny, G. L. Goc, M. Anderson, P. Ireland, D. Cassettari, and G. D. Bruce, “High-fidelity phase and amplitude control of phase-only computer generated holograms using conjugate gradient minimisation,” *Optics Express*, vol. **25**, p. 11692, May 2017.
- D. Bowman, P. Ireland, G. D. Bruce, and D. Cassettari, “Multi-wavelength holography with a single spatial light modulator for ultracold atom experiments,” *Optics Express*, vol. **23**, p. 8365, Apr 2015.

Poster Presentations

- D. Bowman, G. D. Bruce, T. L. Harte, P. Ireland, J. Keeling and D. Cassettari, “Phase-Engineered Light Patterns for Ultracold Atom Experiments,” International Conference on Atomic Physics XXV, COEX, Seoul, Korea, 2016
- D. Bowman, P. Ireland, L. A. Lee, G. D. Bruce and D. Cassettari, “Towards a BEC Rotation Sensor”, Young Atom Opticians Conference XXI, ETH Zürich, Zürich, Switzerland, 2015
- D. Bowman, P. Ireland, L. A. Lee, G. D. Bruce and D. Cassettari, “Towards a BEC Rotation Sensor”, Workshop on Many-Body Dynamics and Open Quantum Systems, Grand Central Hotel, Glasgow, Scotland, 2014

Funding

This work was supported by the EPSRC (Grant Number EP/G03673X/1).
This work was supported by the University of St Andrews.

Contents

1	Introduction	9
1.1	Ultracold Atoms	9
1.2	Novel Trapping Techniques and Manipulation of Ultracold Atoms	11
1.3	Thesis Overview	15
2	Phase-Only Spatial Light Modulator Investigations	16
2.1	Liquid Crystal Spatial Light Modulators	16
2.2	Phase Retrieval Algorithms	17
2.2.1	Iterative Fourier Transform Algorithms	18
2.2.2	Mixed Region Amplitude Freedom	20
2.2.3	Minimisation Algorithms	21
2.2.4	Conjugate Gradient Minimisation	23
2.2.5	Algorithm Initialisation	28
2.3	SLM Experimental Implementation	29
2.4	Multi-Wavelength Light Profiles Using MRAF	31
2.4.1	Calibrating the SLM for Different Wavelengths	33
2.4.2	Sub-Diffraction Limited Rings	35
2.4.3	Multi-Wavelength Feedback Algorithm	36
2.5	Phase and Amplitude Control with Conjugate Gradient Minimisation	40
2.5.1	Cost Function Choice	41
2.5.2	Far-Field Results	46
2.5.3	Angular Spectrum Wave Propagation Results	50
2.5.4	Experimental Verification	53
2.5.5	Multi-Wavelength with Simultaneous Phase and Amplitude Control	58
3	Cold Atoms Setup	60
3.1	Overview	60
3.2	Double Chamber Vacuum Setup	61
3.2.1	Loading Chamber	63
3.2.2	Science Chamber	66
3.3	Optical Setup	68
3.3.1	Laser Scheme	70
3.4	2D MOT, 3D MOT and Optical Molasses	73
3.5	Imaging	77
3.6	Magnetic Trap and RF Evaporative Cooling	79
3.7	Hybrid Trap	84
3.7.1	Bose-Einstein Condensation	88
3.8	Experimental Sequence Summary	89

4	Ultracold Atoms in Holographic Traps	93
4.1	Initial Cross Dipole Trap	93
4.2	MRAF vs CG Optical Traps	98
4.3	Absence of Parametric Heating in the SLM Trap	101
4.4	Arbitrary Trapping Patterns	102
4.5	Feedback Improvement	105
5	Conclusion	108
5.1	Summary	108
5.2	Future Prospects	108

List of Figures

2.1	NLC SLM diagram	16
2.2	Inverse Fast Fourier Transform Example	18
2.3	Diagram of the IFTA calculation process	19
2.4	MRAF Regions of Interest	20
2.5	GS and MRAF Comparison	22
2.6	Steepest Descent Minimisation Cartoon	24
2.7	Conjugate Gradient Minimisation Example	25
2.8	Conjugate Gradient Hologram Calculation Process	26
2.9	Conjugate Gradient Calculations of a Second-Order Power-Law Trap	27
2.10	Padding Example	28
2.11	Guess Phase Construction	30
2.12	MRAF Feedback Process	31
2.13	Multi-Wavelength Principle	32
2.14	Experimental Setup for the Multi-Wavelength Holograms	34
2.15	Look-Up-Table Generation	35
2.16	Diffraction Limited Spots Using Multi-Wavelength MRAF	36
2.17	Entropy Removal Scheme	37
2.18	Multi-Wavelength Regions	38
2.19	Multi-Wavelength Entropy-Separation Result	39
2.20	Multi-Wavelength Pattern Gallery	40
2.21	Phase and Amplitude Conjugate Gradient Hologram Calculation Process	42
2.22	Error Metrics from the Cost Function Investigation	43
2.23	Intensity and Phase Results from Cost Function Investigation	44
2.24	Target Size and Initial Beam Size Investigation for Far-Field Calculations	47
2.25	Conjugate Gradient Far-Field Results	48
2.26	Analysis of Conjugate Gradient Calculation	49
2.27	Target Size and Initial Beam Size Investigation for ASWP Calculations	51
2.28	ASWP Pattern Optimisation	52
2.29	Conjugate Gradient ASWP Results	53
2.30	Conjugate Gradient Light Propagation Investigation	55
2.31	Fourier Transform Fringe Analysis Method	56
2.32	Experimental Results of Conjugate Gradient Patterns	57
2.33	Conjugate Gradient Multi-Wavelength Pattern	59
3.1	Double Chamber Vacuum Setup	62
3.2	Vacuum Chamber Bakeout	63
3.3	Glass Cell and Low Conductance Tube along with 2D MOT Beams	63
3.4	Atomic Beam	65
3.5	Vacuum Chamber with 2D MOT Optics	66
3.6	Recessed Viewports in Octagon	67

3.7	Cut-Through of the Octagon	68
3.8	Flow Diagram of the Optics Table Setup	69
3.9	Laser Scheme	71
3.10	^{87}Rb Transition Spectra	72
3.11	MOT Optimisation	74
3.12	MOT Saturation Behaviour	76
3.13	Optical Absorption Imaging	78
3.14	Imaging Setup	80
3.15	Optical Pumping	81
3.16	Lifetime and Gravity Measurements in the Magnetic Trap	82
3.17	RF Evaporative Cooling	83
3.18	PSD Plot of RF Evaporation	84
3.19	Light Shift for a Two-Level Atom	85
3.20	Horizontal ODT Scheme	85
3.21	ODT Lifetime	86
3.22	PSD of Full Evaporation	87
3.23	Formation of the Bose-Einstein Condensate	88
3.24	Diagram of the Sequence and Analysis Procedure	90
3.25	Full Experimental Sequence	92
4.1	Vertical ODT Scheme	94
4.2	Cross Dipole Alignment	95
4.3	Astigmatism in Vertical Imaging due to Dichroic	96
4.4	Double Spot SLM Trap	97
4.5	Spot Width Investigation with MRAF	99
4.6	Spot Width Investigation with Conjugate Gradient	100
4.7	Focus Test with Conjugate Gradient and MRAF	101
4.8	SLM Noise	102
4.9	Arbitrary SLM Trapping Patterns	103
4.10	Signal Region Investigation	104
4.11	Simple Feedback Examples for SLM Atom Trapping	107

List of Tables

2.1	Cost Function Investigation Error Metrics	46
2.2	Error Metrics for the Far-Field Calculated Patterns	50
2.3	Error Metrics for the ASWP Calculated Patterns	54
2.4	Error Metrics for the Conjugate Gradient Experimental Results	58
3.1	Laser Power Output of Optical fibres	73
3.2	MOT Parameters	76
3.3	Atom Number and Temperature for Each Experiment Stage	89

Chapter 1

Introduction

1.1 Ultracold Atoms

In 1995 an astounding breakthrough in experimental physics occurred: the first production of a Bose-Einstein Condensate (BEC) [1]. This achievement has since stimulated a unique area of investigation in physics, often referred to as the field of ultracold atoms. For over twenty years this field has vastly grown, with developments in experimental precision and control allowing the remarkable properties of ultracold atoms to be studied and utilised. These uses are summarised in a variety of excellent text books and review articles [2, 3, 4, 5].

A BEC can be considered as one of the fundamental examples of a macroscopic quantum state and can exhibit a variety of intriguing properties due to the effects of quantum phenomena. The prediction that such a state would occur was first deduced by Albert Einstein [6, 7, 8] in 1925 after Satyendra Nath Bose [9] sent him a paper which introduced a new method to determine the black-body spectrum. The work of Einstein and Bose showed that a BEC transition is unique in that it occurs without the presence of interactions and is fundamentally only due to the quantum statistical nature of particles which can occupy the same quantum state, known as bosons.

Bose-Einstein statistics is the result of considering an ensemble of non-interacting, indistinguishable bosons and is described by the Bose-Einstein distribution

$$f(\epsilon) = \frac{1}{e^{\beta(\epsilon-\mu)} - 1}, \quad (1.1)$$

where $f(\epsilon)$ is the average number of particles occupying a single quantum state of energy ϵ , β is $1/k_B T$ and μ is the chemical potential of the gas which is determined to conserve the total number of particles. At low temperatures, μ can be considered to affect the system negligibly, except for any atoms in the lowest energy state (ground state). The distribution given in Equation 1.1 shows that the number of atoms in the ground state is given as

$$N_0 = \frac{1}{e^{\beta(\epsilon_0-\mu)} - 1}. \quad (1.2)$$

If we consider the ground state population to be large then, below threshold

$$\frac{|\epsilon_0 - \mu|}{k_B T} \simeq \frac{1}{N_0} \quad (1.3)$$

and the difference between the ground state energy and chemical potential is negligible in comparison to the thermal energy. Since we are considering low temperatures, the

contribution of μ to any excited states is also negligible:

$$\epsilon_1 - \mu = (\epsilon_1 - \epsilon_0) + (\epsilon_0 - \mu) \simeq \hbar\omega + \frac{k_B T}{N_0} \simeq \hbar\omega. \quad (1.4)$$

The excited states of the system can then be described by neglecting μ and the conservation of the total particle number is achieved by considering the excited states as a distribution of photon-like bosons exchanging particles with the ground state:

$$N = N_0 + \int_0^\infty f(\epsilon) D(\epsilon) d\epsilon. \quad (1.5)$$

Here N corresponds to the total number of atoms, $f(\epsilon) \simeq \frac{1}{e^{\beta\epsilon}-1}$ is the distribution of excited particles, similar to photons, and $D(\epsilon)$ is the density of states. For a system of photons, as the temperature decreases, the number of photons also decreases. The same behaviour is seen here, except that a decreasing temperature corresponds to an increase in the number of ground state atoms. At $T = 0$ the entire system is in the single particle ground state.

The process of Bose-Einstein condensation can be thought of as the onset of quantum mechanical effects and increased coherence due to proximity of the particles. Atoms can be described as quantum mechanical wavepackets that have a characteristic wavelength given by $\lambda_{dB} = (2\pi\hbar^2/mk_B T)^{\frac{1}{2}}$ (where m is the atomic mass), known as the de Broglie wavelength. This describes the uncertainty in position of the atom due to its thermal momentum. If the inter-atomic separation becomes comparable to λ_{dB} then the individual atomic wave packets will become indistinguishable and start to overlap, forming a macroscopic wave packet. This can be achieved by either increasing the density or reducing the temperature (or both) of an atomic cloud such that the phase space density is $n\lambda_{dB}^3 \gtrsim 1$ [10], where n is the peak atomic density.

The argument for the BEC presented by Einstein however does not include interactions between the atoms, and the inclusion of interactions reveals the difficulty in experimentally realising such an exotic state. If the density of the atomic cloud is too high then the rate of three body collisions becomes problematic as they can cause losses and the formation of molecules. A BEC however can be created in a metastable phase if the density is reduced such that these molecular formation times are long, but in reducing the density the temperature needed to achieve the necessary phase space density also needs to be reduced. This results in the temperatures required for a low density BEC to be in the sub-microkelvin range. Such a low temperature scale is ultimately what made the BEC so elusive for so long.

The first step to reaching such temperatures was, and still is, the Magneto-Optical Trap (MOT), first developed in 1987 [11]. Capable of trapping and cooling atoms from room temperature down to the sub-millikelvin range, the MOT was a crucial breakthrough for ultracold atoms. This remarkably effective method uses a combination of magnetic and optical fields, requiring the high finesse of a tunable laser to cool atoms with radiative pressure. The crucial laser cooling technique is now the first step to forming a BEC in any experiment with alkali atoms, but the MOT itself is unable to create the phase space densities required for a BEC. In order to reduce the temperature even further, evaporative cooling with either an optical dipole trap or magnetic trap (a magnetic trap was used for the first BEC) is conventionally used¹. In evaporative cooling, the hottest atoms are simply expelled from the trap whilst the cloud is left to thermalise to a cooler ensemble. Such a technique, however, results in the loss of a large proportion of the atoms initially

trapped in the MOT. A usual procedure would see a typical MOT with 10^9 atoms become reduced to around 10^5 - 10^6 atoms using evaporative cooling, before the BEC is formed.

Despite the small sizes of these BECs and their relatively short lifetimes (~ 10 s), a plethora of remarkable experiments have been conducted throughout the field and many macroscopic quantum mechanical effects have been observed. Matterwave interferometry [14], superfluid vortices [15], quantum localisation effects [16] and quantum simulation, including condensed matter systems [17] and black holes [18], are just some of the many achievements realised with BECs since their first realisation.

Furthermore, the field of ultracold atoms is not limited to bosons. Fermions can be cooled to a degenerate Fermi gas using the same methods, achieved a few years after the first BEC [19]. Although the atoms do not follow Bose-Einstein statistics and therefore do not condense in to the same ground state (due to the Pauli exclusion principle), they do collapse in to a Fermi sphere, occupying all of the lowest possible states. It is also possible to have a BEC of molecules made up of weakly bound fermions [20], a phenomena associated with the Bardeen-Cooper-Schrieffer (BCS) superfluid to BEC crossover [21]. Such investigations, along with developments with ultracold atoms in optical lattices [22] have prompted highly promising investigative work with ultracold fermions for understanding high temperature superconductivity [23].

The macroscopic quantum behaviour achievable with ultracold atoms is therefore of high interest for both investigations of fundamental physics and applications of quantum instruments. In addition, ultracold atoms are one of the most controllable experimental systems today: experiments can provide exquisite control over interaction strength, particle number, temperature, dimensionality and trap geometry. My own work has concentrated on the development of the latter and I will now review some of the many novel trapping and manipulation techniques.

1.2 Novel Trapping Techniques and Manipulation of Ultracold Atoms

Novel trapping methods have a wide range of applications in ultracold atomic experiments. Different trapping configurations can drastically change the atoms' behaviour, hence flexibilities in both dynamic and static trap choices have led to further advances in quantum simulation, quantum computing and metrology.

The confinement and manipulation of atoms can generally be categorised into three groups: magnetic traps, optical traps, or a combination of the two known as hybrid traps. The basic principle behind these traps is that applying a magnetic or optical field on to the atoms will cause an energy shift in their ground state (either by the Zeeman effect for magnetic fields or the Stark effect for electric fields [24]) such that shaping the fields appropriately can create trapping potentials.

Flexible techniques with magnetic trapping come in the form of Time-Averaged Adiabatic Potentials (TAAPs) [25, 26, 27, 28] or atom chips [29]. The novel TAAP method incorporates three independent magnetic fields: a static trapping field, an RF oscillating field and an oscillating bias field. The RF field is able to adiabatically deform the static trapping potential experienced by the atoms whilst the oscillating field allows for time-averaging of the potential, creating a variety of both static and dynamic trap geometries. Trapping and rotation of a BEC has been demonstrated in TAAPs [27] as well as flexible ring geometries [26] and double-well potentials [28]. The atom chip meanwhile is a

¹NB There has been recent work in creating BECs by laser cooling alone (see [12, 13]).

nano-fabricated device, where arbitrary magnetic geometries are achievable with different configurations of wires on a substrate [30]. Capable of large magnetic field gradients and fast BEC production [31], the atom chip has become a promising development in the pursuit of compact atomtronic devices [32]. One-dimensional magnetic lattices have been achieved with the atom chip [33], as well as two-dimensional square and hexagonal lattices [34] along with double well configurations [35].

With regards to optical traps, the list of trapping methods is far more varied due to the multitude of techniques that have been developed for controlling laser light. Optical lattices for example utilise the interference of coherent light to form standing waves, where the atoms are attracted (for red-detuned light) to the high intensity regions allowing for a periodic lattice of atoms to be formed in either one, two or three dimensions. The standing waves are created by simply retro-reflecting laser light, but a variety of complex lattice configurations have still been achieved through the use of various tilt angles, multiple wavelengths and overlapping multiple lattices. Examples include the Kagome lattice [36], hexagonal lattice [37], bipartite lattice [38] and double-well superlattice [39]. These trapping geometries, along with various other lattice structures, have become one of the main foci of quantum simulation for modelling solid state systems due to the fine tunability of the lattice parameters available. Investigations of Dirac points [40], the superfluid to Mott insulator transition [17] and the observation of spatial charge and spin-correlations in fermions [41] are amongst the many achievements in optical lattices so far [5].

Highly flexible and arbitrary trapping schemes are also possible with light through the use of various novel optical devices, such as the Acousto-Optic Deflector (AOD) which can “paint” potentials [42]. In an AOD, an RF transducer is used to apply a time-varying compression along one axis of an optical crystal, causing spatially-periodic modulations of the refractive index. A laser beam propagating approximately orthogonal to the compression wave experiences a time-varying diffraction grating. Adjustments in the amplitude of the RF wave change the diffraction efficiency of the AOD, whilst frequency adjustments alter the diffraction angle. The response time of the AOD is much faster ($\lesssim 1\mu\text{s}$) than the typical response time of the atoms in the trap ($\sim 1\text{ms}$), such that two AODs used together can create an arbitrary 2D time-averaged potential by quickly adjusting the angle of the laser beam in both axes (where the third axis is confined via a light sheet). This technique has managed to trap BECs in a range of static and dynamic spatial patterns such as rings and small asymmetric lattices [42], exhibiting more variability than the previously mentioned TAAPs, and has also been used to create a segmented ring for an atom SQUID potential [43]. A potential limitation of the technique however is in the nature of the laser scanning, where more complex geometries may prove difficult if the beam cannot raster fast enough to hold the atoms.

A more recently utilised instrument for manipulating the spatial profile of light for cold atom experiments is the Digital Micromirror Device (DMD). These devices have a large array of micromirrors (typically > 1 million mirrors, with a pixel pitch of $\sim 10\mu\text{m}$), where each mirror can be switched between two angles at a rate of ~ 10 kHz using a computer. An arbitrary intensity profile can thus be created by sending a binary image to the device, where the reflected light from the DMD forms the image. The fast micron scale angular control is achieved by having each mirror attached to a suspended platform which has a hinge running centrally beneath it, where the platform and mirror are connected to a bias voltage. Beneath this, in the corners either side of the hinge, are two electrodes which can be individually addressed. When a bias voltage is applied to the suspended mirror it will tilt towards one of the addressed electrodes via electrostatic attraction, causing the fast angular switch. Despite the mechanical nature of the device, over 100,000 operating

hours and more than 1 trillion mirror cycles are capable with a DMD [44].

The DMD, similar to the AOD painting method, has shown to be capable of trapping BECs in a variety of arbitrary shaped potentials [45]. Whilst a microscope objective was required to directly image the patterns on to the atoms, the range of traps created by the DMD is more extensive than the AOD results. Furthermore, due to the speed of the DMD, grey scale images and time-averaging potentials are possible by modulating the light with a duty cycle on the mirrors, where the authors of [45] found that > 3 kHz cycles had a negligible effect on the atom temperature. Grey scale images are also possible with the DMD via error diffusion [46], where such a technique has been used to create an entropy-extracting trapping profile for fermions in an optical lattice which was critical in the recent key observation of long range anti-ferromagnetic ordering [23].

Not only controlling the amplitude, but controlling the phase of light has also been utilised for novel cold atom traps. Static phase masks [47], for example, are formed from etched transparent plates which have a refractive index different from air, where the designed etchings correspond to a desired phase modulation which is imparted to the incident light beam. Arbitrary atom traps with Fresnel holograms have been investigated for phase plates [48], but typically the masks are used to create Laguerre-Gaussian (LG) modes, which are solutions to Maxwell's equations and can be formed by applying a spiral-shaped phase to a Gaussian beam. The $LG_{0,l}$ (where l is the winding number) modes in particular have a doughnut shaped intensity profile (for $|l| > 0$), ideal for ring traps for BECs and investigations of superfluid circulation [49].

Phase manipulation can also be used to create more arbitrary light patterns, similar to DMDs or AODs. One such technique is Generalised Phase-Contrast (GPC) which directly maps a phase profile to an intensity profile [50]. Here the phase modulated beam is focused through a Phase-Contrast Filter (PCF), which imparts a π phase change to the low frequency components of the beam, whilst leaving the high frequency components (including the pattern) unaffected. These two portions then interfere and the result is an intensity profile that takes the same form as the modulated phase.

Typically however, phase manipulation is used to create arbitrary patterns via diffraction. This does away with the need for a PCF and alignment thereof, but requires computational modelling for the propagation of the phase modulated beam in order to determine the phase mask required to realise the desired pattern. This adds a level of complexity compared to the previously mentioned DMDs, but phase controlling the light also offers more efficiency compared to the DMD, which is naturally a lossy device as it redirects light away from the imaging axis.

Finally, whilst phase plates are able to impart the desired phase to the light, they are rather inflexible as each plate needs to be made separately. Fortunately, there is a dynamic programmable device that can create arbitrary phase masks known as the phase-only Liquid Crystal Spatial Light Modulator (LC SLM). This device is the focus of the work in this thesis. A phase-only LC SLM is an electronic device comprised of an array of pixel electrodes beneath a liquid crystal. Each electrode can be computer controlled and, with an applied voltage, the liquid crystal molecules above the pixel can be rotated. This changes the refractive index along the propagation axis for that pixel, allowing for a controllable phase modulation of the incident light. The LC SLM has additionally flexibility in that it can also be used as an amplitude modulator, in conjunction with a polariser, as the pixels can act as adjustable waveplates. Such a method has managed to trap a 2D degenerate gas in potentials such as a ring and dumbbell, as well as even more arbitrary geometries [51].

The LC SLM is slower than a DMD or AOD however, due to the response time

of the liquid crystal molecules, with the fastest LC SLMs capable of kHz refresh rates and the slower more flexible models capable of only 10-100 Hz. Despite this, there have been numerous propositions and investigations with phase-only SLMs with regards to cold atoms [52, 53, 54, 55, 56, 57, 58, 59] and the faster devices have been used to successfully trap a BEC and split it into multiple components [53]. Previous work in our group has also shown that even the slower models are potentially capable of dynamic atom trapping [55].

The use of both light and magnetic fields is not just limited to trapping atoms either [3, 24]. The interaction strength of the atoms for example can be controlled both magnetically and optically via Feshbach resonances [60, 61], allowing for investigations in to such phenomena as the BEC-BCS crossover [62] or molecular formation [63]. The previously mentioned LG modes also exhibit further manipulation as they are able to optically control the angular momentum of a BEC. An LG beam carries with it an orbital angular momentum due to its helical phase shape, and this momentum can be transferred to a BEC via a two-photon process [64, 65, 49], leading to superfluid circulation. Disorder on the atoms can also be investigated by applying speckled light to the trapping profile, leading to localisation effects [66] where a DMD has also been used to control such phenomena [67]. Novel atom addressing schemes have also been exhibited with light manipulation, such as the single-site spin flip technique for optical lattices [68]. Here, a “magic” wavelength of light is used on individual lattice sites which shifts one of the ground state energy levels, whilst leaving another ground state unaffected, bringing the atoms on that site closer to resonance with a separate spin-flip field. This method has also been used with a DMD to create a line of spin-impurities in a lattice, allowing the investigation of simultaneous propagation in multiple 1D chains [69].

Light fields have also been used to create artificial gauge fields with atoms [70]. Notable experiments include the use of laser assisted tunnelling in an optical lattice to induce cyclotron orbits, mimicking the behaviour of charged particles in magnetic fields with strengths that are unobtainable by current magnets [71, 72]. Spatial control of the phase of ultracold atoms is possible as well. A pulsed hard-edged light profile, for example, can be used to phase imprint on to the atoms and create topological defects known as solitons [73, 74] which are long-lasting excited states highly studied with regards to phase transition phenomena.

The techniques for creating bespoke light patterns in combination with atomic manipulation has opened up a plethora of exciting experiments, where further development may see new avenues of investigation ahead. Our group, for example, has made use of the aforementioned, phase-only LC SLM and previously proposed a variety of experiments with the device, such as an efficient production method to create a BEC [56]. Here, a Gaussian beam profile typically used for focused optical dipole traps is instead replaced with dynamic holographic power law traps created with the SLM. A usual evaporation process would involve simply ramping down the laser power of the Gaussian beam, resulting in a shallower trap with reduced trapping frequencies which leads to lower densities as the cloud cools. Here, however, the profile is tunable with the SLM and the ideal evaporation parameters can be met as the trap depth is reduced, allowing for a more efficient creation of a BEC. Another proposed experiment using the SLM regards topological Kondo devices [75]. The proposal is to emulate a solid-state system which has a superconducting central island that interacts with a gas of fermions in one-dimensional external wires. Such a system exhibits Majorana modes and the capability to control such modes has crucial uses in topological quantum computation. An experimental realisation of the solid-state topological Kondo device has yet to be achieved, but a cold atom system offers a clean

and controllable approach to simulate the device. The trapping geometry required for such a cold atom experiment is a set of one-dimensional waveguides joined together in a Y-junction with a junction in the centre acting as the impurity. Our group has shown that such a geometry can be generated and controlled with an SLM and that the topological Kondo device may be possible to realise experimentally. Finally, our group has proposed a smooth dynamic holographic trap for superfluid investigations [55]. The light pattern here is annular with two bright spots which act as stirring rods. Rotations can be induced on the BEC by moving the bright spots around the circumference of the ring, where it was estimated that a sequence of 70 separate light patterns could be used to induce superflow on to the atoms. The stirring rods could then be slowly removed to leave a uniform ring for the atoms to flow through. Such flexibility with rings traps have potential applications to create an atomic analogue of a superconducting quantum interference device (SQUID) which can potentially be used as a rotation sensing device [76].

Further development of light control with the SLM is therefore of great interest for atomic physics. For example, simultaneous control of multiple wavelengths of light is highly desirable not only because of the possible aforementioned atomic manipulation, but because further control of trapping geometries are also attainable by using combinations of blue- and red-detuned light (repulsive and attractive potentials respectively). Using an SLM to control both the amplitude and phase of light incident on the atoms as well can not only be used to impart momentum, but is also a crucial step in fully realising bespoke three-dimensional trapping geometries, where such control may prove to be crucial in future experiments.

1.3 Thesis Overview

The main achievements discussed in this thesis are the observation of the first BEC in St Andrews, and the developments towards manipulating these atoms in flexible, SLM-generated, optical potentials. The thesis is outlined as follows: In Chapter 2, the principles of phase-only modulation using SLMs is covered, along with hologram calculation techniques for bespoke light patterns. My work regarding SLM investigations is presented here and results thereof for a multi-wavelength control technique, as well as high-fidelity simultaneous phase and amplitude shaping of light. In Chapter 3, my work regarding the design, construction and optimisation of our complete ultracold atoms experimental setup is presented, along with our achievement of forming a BEC. Chapter 4 is then effectively the combination of the two previous chapters, presenting the successful trapping of cold atoms in a variety of holographic traps with the SLM methods I have worked on, before Chapter 5 finally concludes the thesis by pointing to the near-future successes of the methods developed here.

Chapter 2

Phase-Only Spatial Light Modulator Investigations

In this chapter, I will present a number of methods for achieving flexible optical traps by manipulating the phase of light. Before doing so I will first outline the key technology we use to realise this: the liquid crystal spatial light modulator.

2.1 Liquid Crystal Spatial Light Modulators

The principle of operation behind the LC SLM is based on the anisotropic properties of the liquid crystals. The liquid crystal material in the SLM is birefringent, where there are two refractive indices which are dependent on the orientation of the molecules. Varying the orientation of the molecules can therefore cause different optical path lengths for an incident light field, resulting in phase modulation (both amplitude and phase modulations can also be created by changing the incident light field's polarisation, although this results in a reduced light efficiency and some coupling between the two parameters). The orientation of the molecules can be controlled electronically by applying voltages to an array of pixel electrodes which are in proximity to the liquid crystal. The diffraction efficiency of

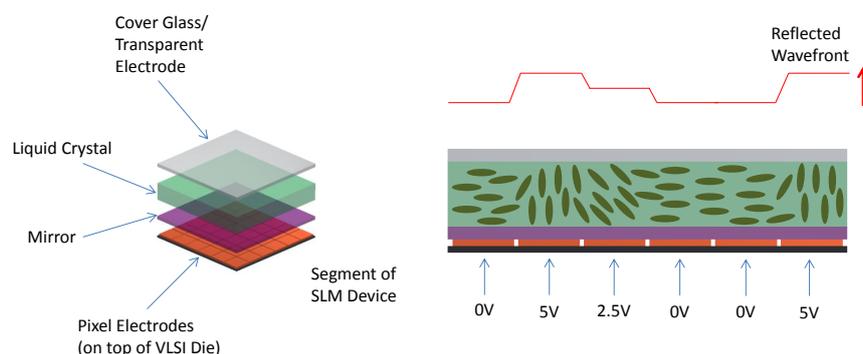


Figure 2.1: Principle of operation of an electronically addressed reflective NLC SLM. The voltages applied to the electrodes can take multiple values and can be controlled dynamically via a computer.

such devices is limited by effects such as the dead space between pixel electrodes, which is most prominent in transmissive SLMs as this space contains the addressing wires. Reflective SLMs overcome this problem, but still have a small amount of interpixel dead space ($\gtrsim 90\%$ fill factor). A sketch of the operation for an electronically addressed reflective LC SLM can be seen in Figure 2.1.

There are two categories of LC SLMs based on the type of liquid crystal they use: ferroelectric (FLC) or nematic (NLC). FLCs are the faster choice as their refresh rates can be in the kHz regime whilst NLCs achieve only 10-100 Hz, although there has been work on improving the speed of NLCs [77]. The limiting factor of FLCs is that they can only produce binary phase modulation, resulting in lower diffraction efficiencies and limiting potential complexities for desired light patterns. NLCs however can achieve multiple levels of phase modulation (typically 256).

For our investigations, we use an electronically addressed reflective NLC SLM as a phase-only modulator for high light efficiencies and the ability to form complex atom traps, at the cost of reduced frame rate. The following sections will first introduce phase retrieval algorithms which are used to calculate the phase masks for the SLM (Section 2.2). Experimental implementation of the technique will then be discussed (Section 2.3), followed by developments that I have made towards holographic trapping in the form of a multi-wavelength technique (Section 2.4) and simultaneous phase and amplitude control (Section 2.5).

The SLM model we use for the multi-wavelength investigations in Section 2.4 and the atom trapping in Chapter 4 is a Boulder Nonlinear Systems (BNS) XY-series P256 phase-only SLM, which has a pixel array size of 256×256 with a pixel pitch of $24\mu\text{m}$. Each pixel can apply 256 different phase shifts between 0 and 2π to an incident light beam. For the conjugate gradient investigation in Section 2.5 we used a Meadowlark XY-series P1920 phase-only SLM. The P1920 has a pixel pitch of $9.2\mu\text{m}$ and an array size of 1920×1152 . Both SLMs provide phase-only modulation when the input light has vertical linear polarisation, where the output polarisation is left unaffected (amplitude modulation can be achieved with 45 degree linear polarisation, although some phase-coupling is still present).

2.2 Phase Retrieval Algorithms

Creating arbitrary light patterns with a phase-only NLC SLM can be fairly challenging. If we consider the far-field diffraction limit for our patterns such that we work in the Fraunhofer diffraction regime, then the problem is as follows: if we know both the amplitude and phase of an input light field (e.g. a laser with a Gaussian beam profile), what is the phase modulation (hologram) that we need to apply in order to produce the target light field in the far-field limit? Far-field diffraction can be modelled via a Fourier transform and can be realised experimentally either by viewing the diffracted light at a long distance from the SLM, or by focusing the beam with a lens [78]. The problem can then be summarised mathematically by writing the input beam as $A_0(x', y') \exp(i\theta(x', y'))$, where we want to find the phase modulation $\phi(x', y')$ such that $\mathcal{F}[A_0(x', y') \exp(i(\theta(x', y') + \phi(x', y')))] = \tau(x, y) \exp(i\Phi(x, y))$, where $\tau(x, y) \exp(i\Phi(x, y))$ is the target light field. For clarity in the following, I assume $\theta(x', y') = 0$, i.e. the input laser beam has a flat phase that we normalise to zero. I will also be frequently omitting spatial co-ordinates (x, y) for brevity and clarity.

Since we can realise the far-field with a Fourier transform, then the naive approach would be to take an inverse Fourier transform of our target light field to determine the

necessary phase at the SLM plane. This, however, will yield both a phase *and* amplitude modulation to be applied to the input light, whereas we only wish to modulate the phase. An example of such structure can be seen in Figure 2.2 for a target amplitude with a flat phase.

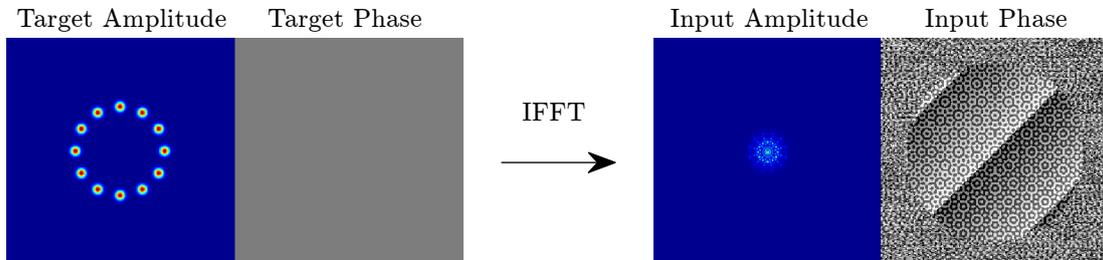


Figure 2.2: Inverse Fast Fourier Transform (IFFT) of a ring lattice with a flat phase. The resulting input amplitude and phase are modulated, and whilst the phase can be created via an LC SLM, the amplitude cannot and must take the form of the incident laser beam profile (typically Gaussian). As can be seen, the input amplitude is not suitably Gaussian and therefore the target light field cannot be achieved with a laser shone on a phase-only SLM.

Thankfully, in many cases the constraints of the problem can be somewhat alleviated due to the nature of cold atom trapping. For far-detuned optical dipole traps the phase of the light does not affect the trapping profile (if the dipole-approximation is valid). Whilst there are some useful capabilities for atomic manipulation using traps with phase structure, for now we will consider traps which have unconstrained phase profiles.

Without a constraint on the target phase, we can attempt to find a phase modulation that optimises the target amplitude. In general, this must be calculated numerically and the quickest and most common approach for doing this is a method known as Iterative Fourier Transform Algorithms (IFTAs) which will be described in the next section. Following this, I will review the most-relevant example of an IFTA: the Mixed Region Amplitude Freedom (MRAF) [79] method, which I have used to create smooth multi-wavelength traps [80]. I will then review alternative calculation methods based on minimisation routines in 2.2.3, followed by our chosen minimisation routine of Conjugate Gradient Minimisation (2.2.4) which I used to create high-fidelity patterns [81]. This section will then conclude with a review of the initialisation of these algorithms.

2.2.1 Iterative Fourier Transform Algorithms

IFTAs, first developed by Gerchberg and Saxton (GS) in 1972 [82], are capable of efficiently computing holograms through an iterative calculation process. The principle of the calculation utilises the phase freedom in two planes, Fourier transforming back and forth between the SLM plane and Fourier plane (when using a lens placed at a distance of one focal length f away from the SLM, the Fourier plane corresponds to the plane at a distance $2f$ away from the SLM) and replacing the amplitudes of the results each time with the target or the laser profile as appropriate. The phase information is maintained throughout and thus evolves towards an acceptable solution.

The IFTA begins the calculation with an initial light field

$$E_{\text{in}}^{(1)} = A_0 e^{i\phi_0}, \quad (2.1)$$

formed of the incident beam profile, A_0 , and a guess phase modulation, ϕ_0 . The guess phase can be chosen to help improve the initial conditions of the calculation [79]. $E_{in}^{(1)}$ is then Fast Fourier Transformed (FFT) to realise the far-field diffraction of the SLM, giving a light field which can be written generally as

$$\begin{aligned} E_{out}^{(1)} &= FFT \left(E_{in}^{(1)} \right), \\ E_{out}^{(1)} &= A_{out}^{(1)} e^{i\varphi^{(1)}}, \end{aligned} \quad (2.2)$$

where $A_{out}^{(1)}$ is the output field amplitude and $\varphi^{(1)}$ is the output field phase. $A_{out}^{(1)}$ is then replaced with the target amplitude, τ , to form

$$G_{out}^{(1)} = \tau e^{i\varphi^{(1)}}. \quad (2.3)$$

An inverse FFT is then performed on $G_{out}^{(1)}$ to determine the field back at the SLM plane, leading to

$$\begin{aligned} G_{in}^{(1)} &= IFFT \left(G_{out}^{(1)} \right), \\ G_{in}^{(1)} &= A_{in}^{(1)} e^{i\phi^{(1)}}, \end{aligned} \quad (2.4)$$

where $A_{in}^{(1)}$ is the resulting amplitude and $\phi^{(1)}$ is the phase. We once again replace the amplitude, but now with the incident beam profile A_0 leading to

$$E_{in}^{(2)} = A_0 e^{i\phi^{(1)}}. \quad (2.5)$$

The first iteration of the calculation is now complete and further iterations of the same process occur, until the number of iterations reaches a specified limit, or the error between

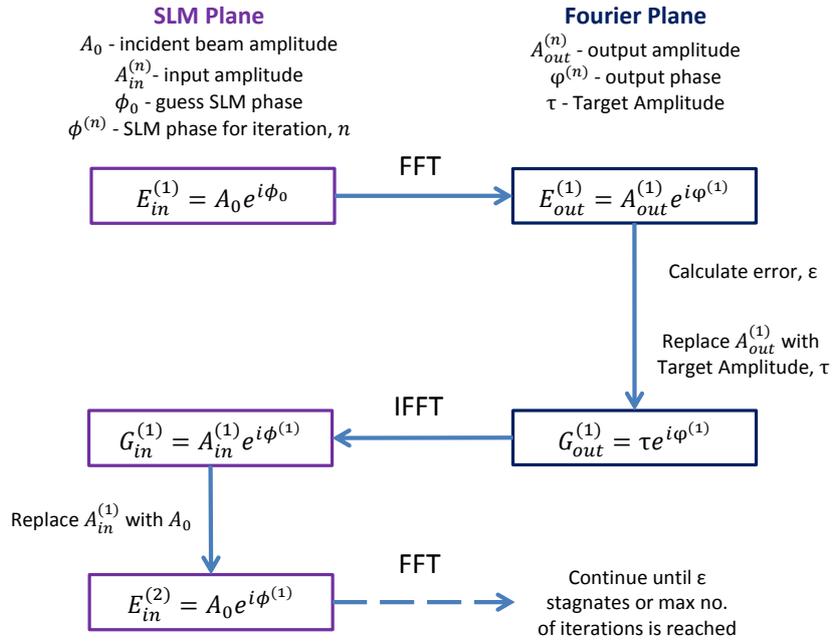


Figure 2.3: Diagram of the IFTA calculation process.

the output electric field amplitude and the target is no longer reduced (stagnates). For our work the error used is the RMS error, defined as

$$\epsilon^{(n)} = \sqrt{\frac{1}{N} \sum_{x,y} \left(A_{\text{out}}^{(n)}(x,y)^2 - \tau(x,y)^2 \right)^2} \quad (2.6)$$

and the stagnation criterion is

$$\epsilon^{(n)} - \epsilon^{(n-1)} < 10^{-k}. \quad (2.7)$$

Here N is the number of pixels, x and y are the pixel indices in the output plane and k is the user defined stagnation parameter. By the end of the calculation, the hologram $\phi^{(n)}$ (where n is the number of iterations the calculation took) will have been determined and can be applied to the incident laser beam to form the pattern in the far-field. A diagram of the IFTA process is shown in Figure 2.3.

Whilst holograms can be generated through this method, the algorithm tends to struggle for continuous patterns and produces a characteristic speckle which would lead to rough trapping potentials for the atoms, causing potential heating and fragmentation as seen in atom chips for magnetic potentials [83].

2.2.2 Mixed Region Amplitude Freedom

A variant of the IFTA algorithm which leads to smoother patterns is MRAF, developed by Pasienski and De Marco in 2008 [79]. This technique takes the concept of phase freedom in the standard GS IFTA and extends it to amplitude as well. By considering that, in general, it is not necessary for the entire output plane to have a well defined amplitude, one can construct a region away from the desired pattern where amplitude is unconstrained. Whilst this will lead to a loss in light efficiency as unwanted light will be dispersed to this region, the MRAF method reduces the constraints on the IFTA process, allowing for smoother patterns to be generated over the standard GS algorithm.

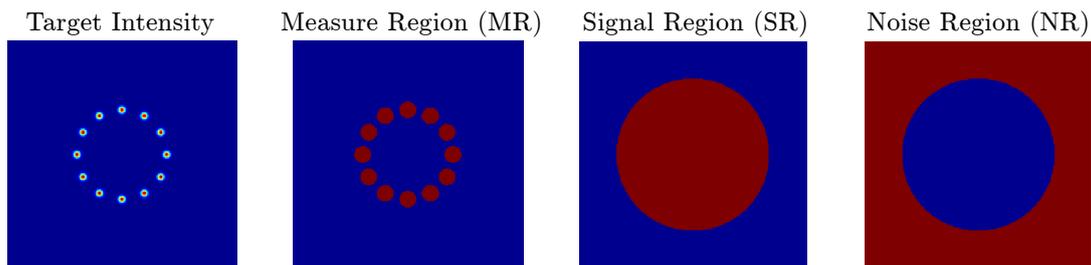


Figure 2.4: Example of ring lattice target pattern and associated regions of interest for an MRAF calculation. The regions of interest are shown in red.

The MRAF process is similar to that of the GS algorithm, except now there are regions of interest in the output plane, shown in Figure 2.4. Three regions are defined in the process, the measure region (MR), signal region (SR) and noise region (NR). Along with these, a mixing parameter (m) is introduced. The measure region is defined for the purpose of error metrics and occupies the high intensity regions of the target pattern (typically regions where the intensity is $> 1\%$ of the maximum intensity of the target). In the signal region the amplitude is kept constrained and the phase is left unconstrained.

Finally, in the noise region *both* amplitude and phase are unconstrained. The relationship between these regions is that $\text{MR} \subset \text{SR}$, while NR is the complement of SR ($\text{NR} = \text{SR}^c$). The mixing parameter is used to control the percentage of light retained within the signal region. With these parameters, $G_{\text{out}}^{(n)}$ from Equation 2.3 is changed to

$$G_{\text{out}}^{(n)} = \left(m \tau|_{\text{SR}} + (1 - m) A_{\text{out}}^{(n)}|_{\text{NR}} \right) e^{i\varphi^{(1)}}. \quad (2.8)$$

By defining the SR such that it leaves a sizeable gap of zero intensity between the noise region and the edge of the trapping pattern, the noise region light can be considered separate from the trap profile. If there is still a possibility for the atoms to be affected by the light in the noise region, an iris can be used on the pattern such that only the desired trap profile is allowed to pass through. Then, the clean resulting pattern can be imaged on to the atoms instead. A slight alternative to the MRAF method, known as OMRAF, puts a base level of light around the target pattern in the SR, rather than leaving it as zero intensity, which has shown some improvements on the RMS error [57], although this may cause interference artefacts when attempting to iris the pattern.

With the MRAF method, a clear improvement on the smoothness of the light profile is seen over the GS algorithm, with approximately one third of the light efficiency for most patterns. Smoothness of the trapping potential for cold atoms is crucial to avoid fragmentation or heating in the cloud. A visual comparison between the two algorithms is shown in Figure 2.5 for an example ring lattice pattern.

For the MRAF algorithm a regional error can be defined based on the measure region:

$$\epsilon_{\text{MR}}^{(n)} = \sqrt{\frac{1}{N_{\text{MR}}} \sum_{x,y \in \text{MR}} \left(\tilde{A}_{\text{out}}^{(n)}(x,y)^2 - \tau(x,y)^2 \right)^2}, \quad (2.9)$$

where N_{MR} is the number of pixels in the measure region and \tilde{A}_{out} is the renormalised output electric field such that the maximum intensity of the calculated pattern in the measure region is equal to the maximum intensity of the target.

We have coded both the GS and MRAF algorithms in Matlab. Example patterns, shown in Figure 2.5, were calculated for 100 iterations, which took 10.9 s and 10.3 s for GS and MRAF respectively, resulting in $\epsilon_{\text{MR}}^{\text{GS}} = 10.7\%$ whilst $\epsilon_{\text{MR}}^{\text{MRAF}} = 0.6\%$.

Our work with the MRAF algorithm for creating atom traps using multiple wavelengths is presented in Section 2.4. However, for completeness I will first make a short digression where I review alternative hologram calculation methods previously undertaken by the group, most notably the conjugate gradient method.

2.2.3 Minimisation Algorithms

Whilst IFTAs are effective at producing SLM phase masks for target intensities, there are still other methods available such as direct search [84] or genetic algorithms [85]. These are less common however as they are computationally far more demanding but, unlike the IFTA approach, they calculate the holograms via a directed minimisation process. This can be advantageous as the optimisation can be controlled and specific aspects of the light pattern can be weighted with higher importance in the calculation.

The most brute force approach is the direct search algorithm [84] which optimises an initial guess hologram by changing the phase of randomly-selected individual pixels in order to reduce a chosen cost function (e.g. the difference between the output intensity and target intensity). Pixel changes which lower the value of the cost function are maintained while those which do not are discarded, allowing the algorithm to find a minimum of the

cost function. Whilst this is feasible with binary holograms, for multiple phase levels the number of calculations needed becomes too high for standard computers as the parameter space of the problem is enormous.

Genetic algorithms take an evolutionary approach to the problem instead by treating a range of initial guess holograms as parents which breed a new set of holograms. The new set are then compared against each other with respect to a cost function, where the most promising are kept for the next breeding stage. The breeding process also allows introduction of mutations, which can help the calculation avoid getting stuck in local minima. This method still requires high computing power however. Our group found that this technique proved to be effective for small simple patterns with binary phase levels. However, for larger patterns and 256 phase levels, the calculation run times took several hours and led to memory issues with the computer. Further information can be found in

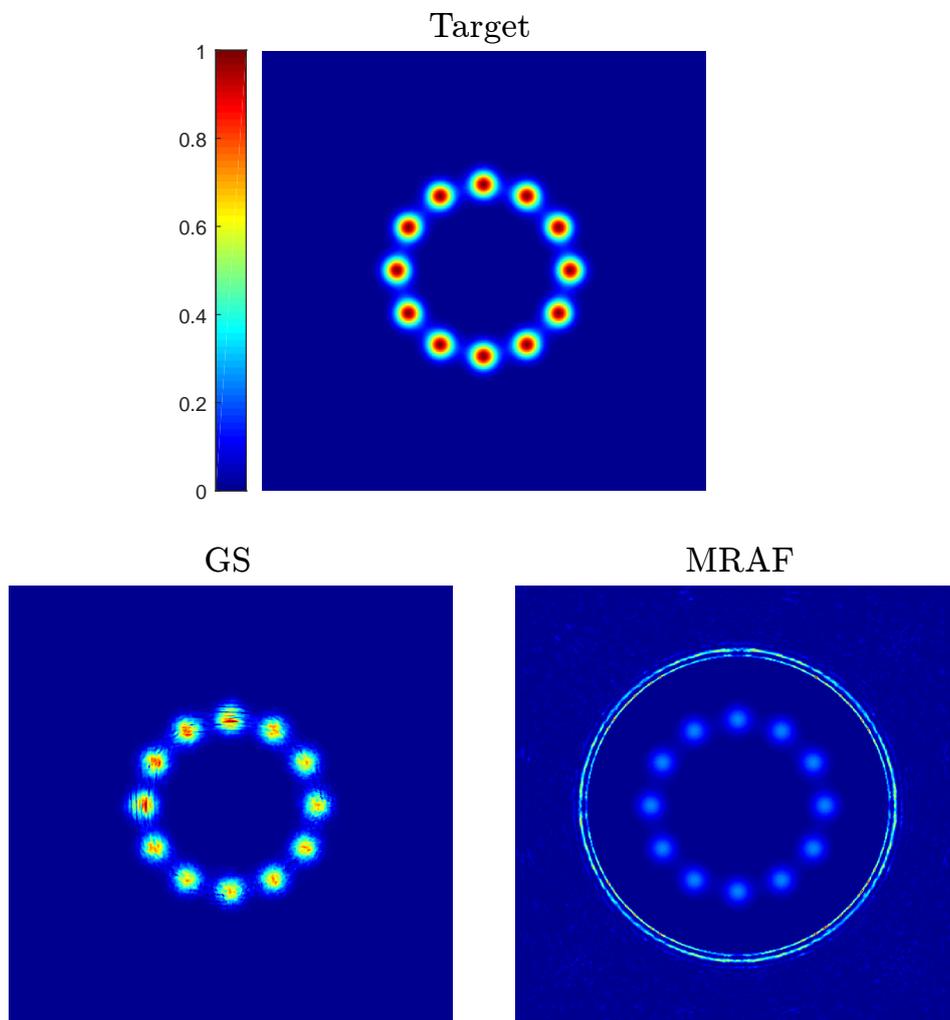


Figure 2.5: Comparison of the intensity output for the GS and MRAF algorithms for an example ring lattice pattern. The spots are fragmented in the GS calculated pattern, whilst the MRAF pattern remains smooth in the region of interest. The mixing parameter used for the MRAF pattern was 0.4 and the light efficiency is 30%, whilst the GS pattern light efficiency is 100%.

Callum MacArthur’s master’s thesis [86].

Another minimisation approach we have investigated was originally developed in our group in 2012 [87, 88] and uses the computationally efficient technique known as conjugate gradient minimisation [89]. Initially this method was also slower than the IFTA approach, requiring over 100 times more iterations to converge than a typical MRAF calculation. However, the minimisation code has now been further developed to allow comparable run times with the MRAF method. We have shown in [81] that this method produces exceptionally high accuracy and versatility beyond state-of-the-art demonstrations of IFTAs for simultaneous phase and amplitude control. This work is presented in Section 2.5, whilst the following section gives an introduction to the conjugate gradient method.

2.2.4 Conjugate Gradient Minimisation

Rather than making calculations throughout the entire parameter space like the direct search method, the conjugate gradient minimisation technique reaches a minimum by utilising the gradient of slopes of the cost function, greatly reducing the number of calculations needed. As an efficient minimisation process this method has seen effective use in a variety of fields, most prominently in energy calculations in quantum mechanics [90, 91, 92], whilst our group is the first to utilise the technique for calculating holograms [88, 81].

The conjugate gradient method is suitable for linear problems of the form $Ax = b$ [89], where b is a vector, A is a matrix and x is an unknown vector to be solved. So, for a scalar quadratic function with a general form

$$f(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T \mathbf{A}\mathbf{x} - \mathbf{b}^T \mathbf{x} + c, \quad (2.10)$$

then the gradient of this function (if A is symmetric) is

$$f'(\mathbf{x}) = \mathbf{A}\mathbf{x} - \mathbf{b}, \quad (2.11)$$

which when equated to zero gives the suitable linear problem and allows us to find a minimum (or maximum) of the function $f(\mathbf{x})$.

The conjugate gradient method is related to the steepest descent approach, which will be described first. In order to simplify the description we can consider a two-dimensional parameter space (for our purposes, this would correspond to a hologram consisting of only two pixels, each capable of multiple phase levels) with an arbitrary cost function which is parabolic in shape in keeping with Equation 2.10. To find the minimum, the calculation starts in an initial location (random or an educated guess) in the parameter space and determines the direction in which the cost function has the steepest descent. A step is then taken along this line such that the cost function is minimised in this direction. At this new location, the orthogonal direction to the previous steepest descent is taken (this also corresponds to the direction of steepest descent at the new location), and a new step is made which minimises the cost function along this line. The steps are then repeated until a local minimum is found, as shown in Figure 2.6.

Whilst this method can find a minimum, the convergence rate can be strongly dependent on the initial starting point of the calculation. Steps are often repeated in the same direction as each descent is orthogonal to its predecessor, as depicted by the zig-zags in Figure 2.6, so if an initial step does not converge well, the latter ones may not be very effective either. To overcome this, instead of the orthogonal search directions being chosen, the conjugate directions can be taken instead i.e.

$$\mathbf{d}_1^T \mathbf{A}\mathbf{d}_2 = 0, \quad (2.12)$$

where \mathbf{d}_1 and \mathbf{d}_2 are the search vectors which are conjugate with respect to the matrix \mathbf{A} . This can be viewed of as taking orthogonal directions, but in a stretched (scaled) space characterised by \mathbf{A} , the principle of which is shown in Figure 2.7(b) (indeed, an alternative name for conjugate directions is A-orthogonal or orthogonal in the space A).

These search directions can be generated through conjugation of the steepest descent directions (using a conjugate Gram-Schmidt process) [89] and take the form

$$\mathbf{d}_{i+1} = \mathbf{r}_{i+1} + \beta_{i+1}\mathbf{d}_i, \quad (2.13)$$

where i denotes the iteration of the calculation, \mathbf{r}_{i+1} is the direction of steepest descent at the starting position of step $i + 1$ and β_{i+1} is defined as

$$\beta_{i+1} = \frac{\mathbf{r}_{i+1}^T \mathbf{r}_{i+1}}{\mathbf{r}_i^T \mathbf{r}_i}. \quad (2.14)$$

An immediate advantage of this approach is that it guarantees convergence in $\leq N$ steps, where N is the dimensionality of the parameter space. Furthermore, in terms of computational storage, each new conjugate search direction can be constructed with storage of only the previous search direction and steepest descent direction whilst still maintaining the minimisation achieved by earlier steps. This means faster convergence in comparison to the steepest descent approach and a highly efficient technique in finding the minimum of a quadratic function.

In a general problem however, the function may be more complicated and will not be purely quadratic. Non-linear conjugate gradient algorithms such as the Polak and Ribière variant [89] help to overcome this and are effective for approximately quadratic functions. However, for such complicated functions the method can also get trapped in local minima rather than finding the global minimum. Despite this, our group has shown that highly accurate hologram calculations are possible with the conjugate gradient method with reasonable robustness to initial conditions of the calculation [88].

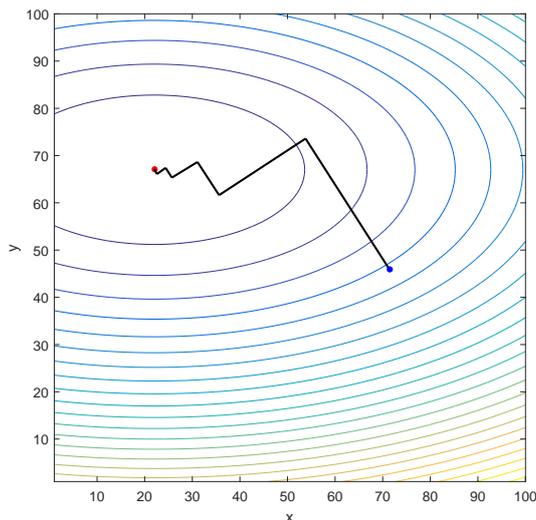


Figure 2.6: Cartoon of the steepest descent minimisation technique for an arbitrary parabolic cost function in a 2D parameter space. The blue point corresponds to the start position of the calculation. The red point is the minimum of the cost function.

For our hologram calculations the cost function in the problem compares the calculated electric field and a chosen target, where the parameter space for the optimisation encompasses all the different phase distributions that the SLM can generate. The simplest cost function (which will be considered here) would be the difference between the target intensity and output intensity. The flexibility of the conjugate gradient technique over the IFTA approach comes about with the freedom to choose an appropriate cost function and in Section 2.5, investigations for an amplitude- and phase-dependent cost function are presented.

A schematic of the conjugate gradient optimisation process is shown in Figure 2.8. The calculation begins by setting the initial parameters as with the IFTA approach, where the output plane is calculated via a Fourier transform (other propagation calculation steps are possible, see Section 2.5.3). This sets the initial position in the parameter space for the algorithm to begin. The cost function and its gradient are then used to determine the direction of steepest descent at this location. For the first iteration, the steepest descent is utilised directly and the minimum along this line is found. At this new location in the parameter space (corresponding to a different hologram with new phase values) the output plane is once again calculated and the cost function and gradient determined as before. For the next iteration (and subsequent iterations), the minimisation step is taken along the conjugate direction, as defined by Equation 2.13. This process then continues and at each iteration the change in gradient compared to the previous step is checked against a stagnation criterion, where if the change is below a chosen threshold, the calculation stops. Alternatively the algorithm can reach a maximum number of iterations as set by the user where the calculation stops. By the end, the optimised hologram for the SLM is

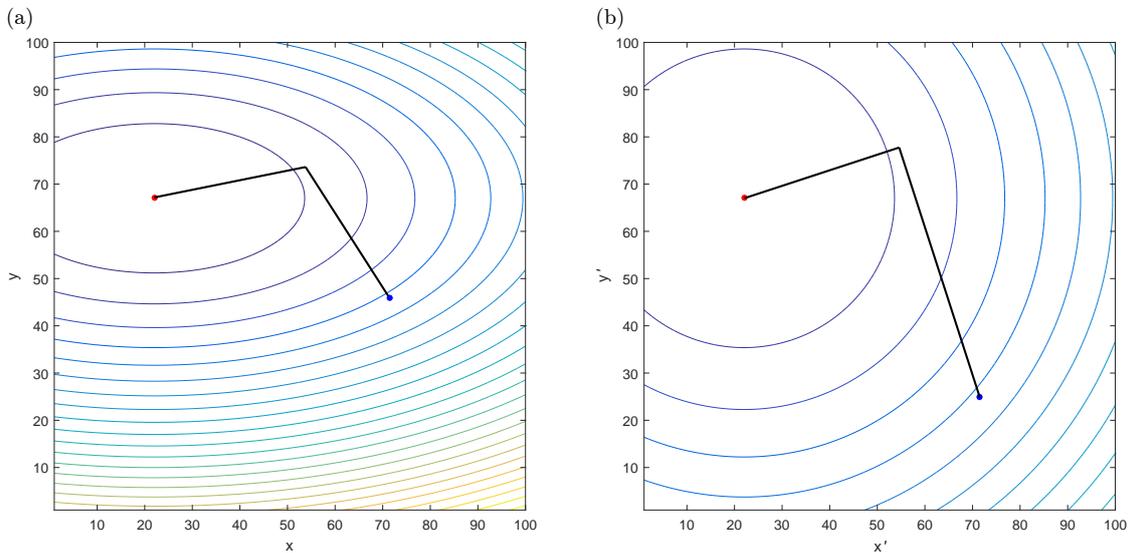


Figure 2.7: Conjugate gradient minimisation example. The blue point corresponds to the start position of the calculation. The red point is the minimum of the cost function. (a) Cartoon of the conjugate gradient minimisation technique for the same function as in Figure 2.6. The method now reaches the minimum in two steps, where the two search directions are conjugate. (b) Visualisation of conjugate directions. The same function as before except viewed in a “stretched” space (x' and y') such that the parabola is symmetric. Here the search directions from (a) are orthogonal.

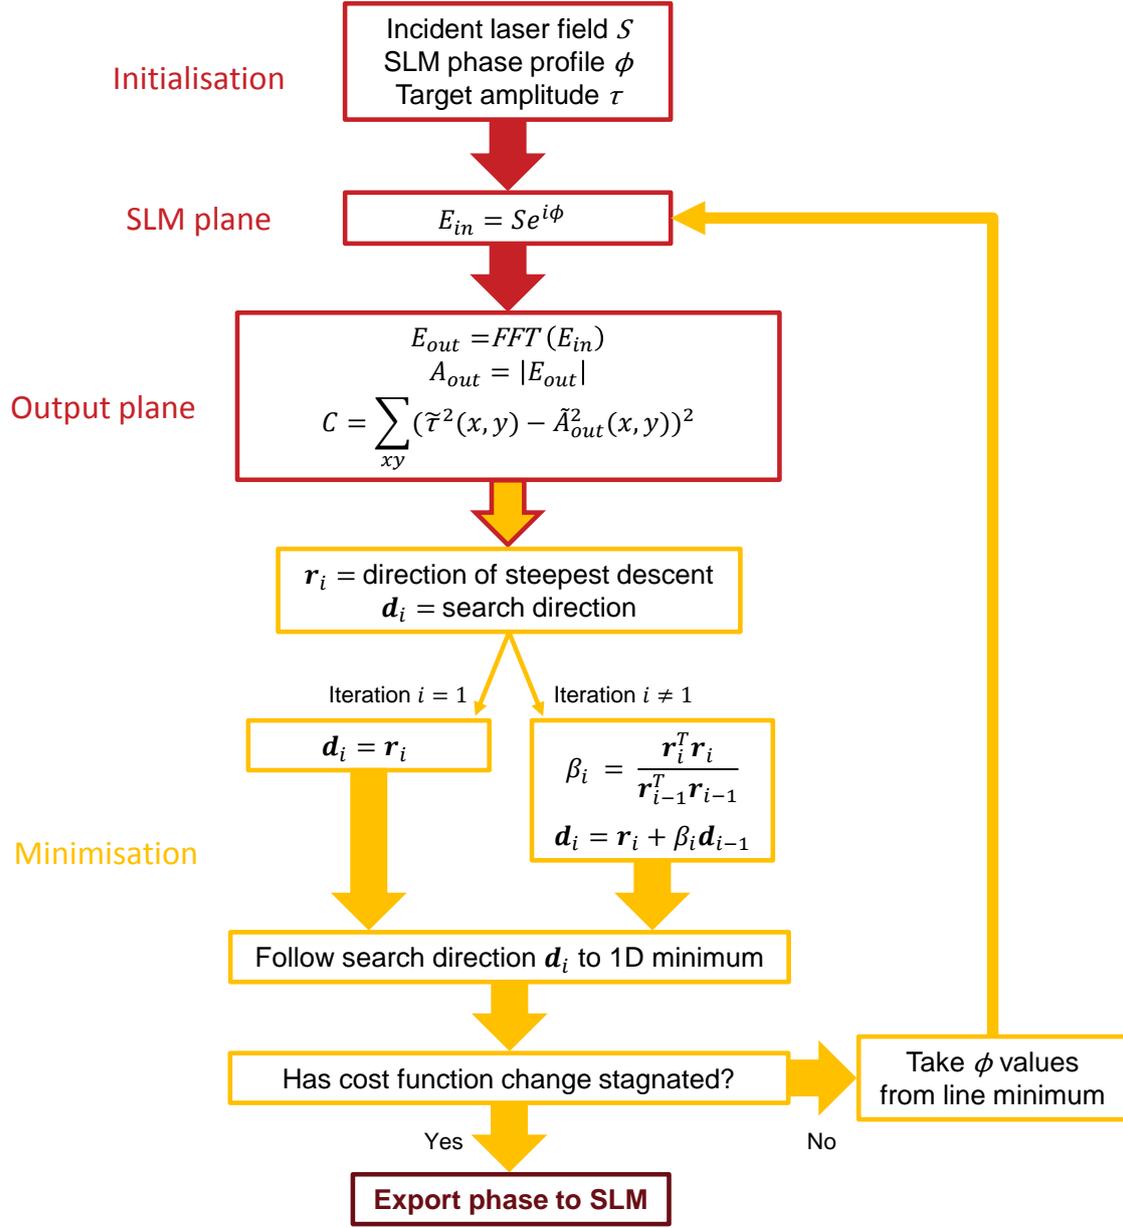


Figure 2.8: Block diagram of the hologram calculation process using conjugate gradient minimisation.

generated and can be exported for use.

The advantage of being able to specify different cost functions can be seen in Figure 2.9 which is taken from [88]. Figure 2.9(a) is the result of using a cost function of the form

$$C_a = \sum_{x,y} \left(\tilde{\tau}^2(x,y) - \tilde{A}_{out}^2(x,y) \right)^2, \quad (2.15)$$

where a multitude of optical vortices are present, greatly distorting the intensity profile such that it has a fractional RMS error of 26%. Optical vortices are points of greatly

reduced intensity which appear due to a phase winding occurring in the output plane and are difficult to remove due to their topologically protected nature. Clearly, the algorithm has been trapped in a local minimum and cannot remove the vortices without increasing C_a . This vortex-stagnation problem is well known in IFTAs as well [79, 93]. Figure 2.9(b) however, uses a cost function which suppresses the vortices by additionally smoothing the intensity over nearest neighbouring pixels

$$C_b = \sum_{x,y} \left\{ \left(\tilde{\tau}^2(x,y) - \tilde{A}_{\text{out}}^2(x,y) \right)^2 + \left[\left(\tilde{A}_{\text{out}}^2(x,y) - \tilde{A}_{\text{out}}^2(x,y \pm 1) \right)^2 + \left(\tilde{A}_{\text{out}}^2(x,y) - \tilde{A}_{\text{out}}^2(x \pm 1,y) \right)^2 \right] \right\}. \quad (2.16)$$

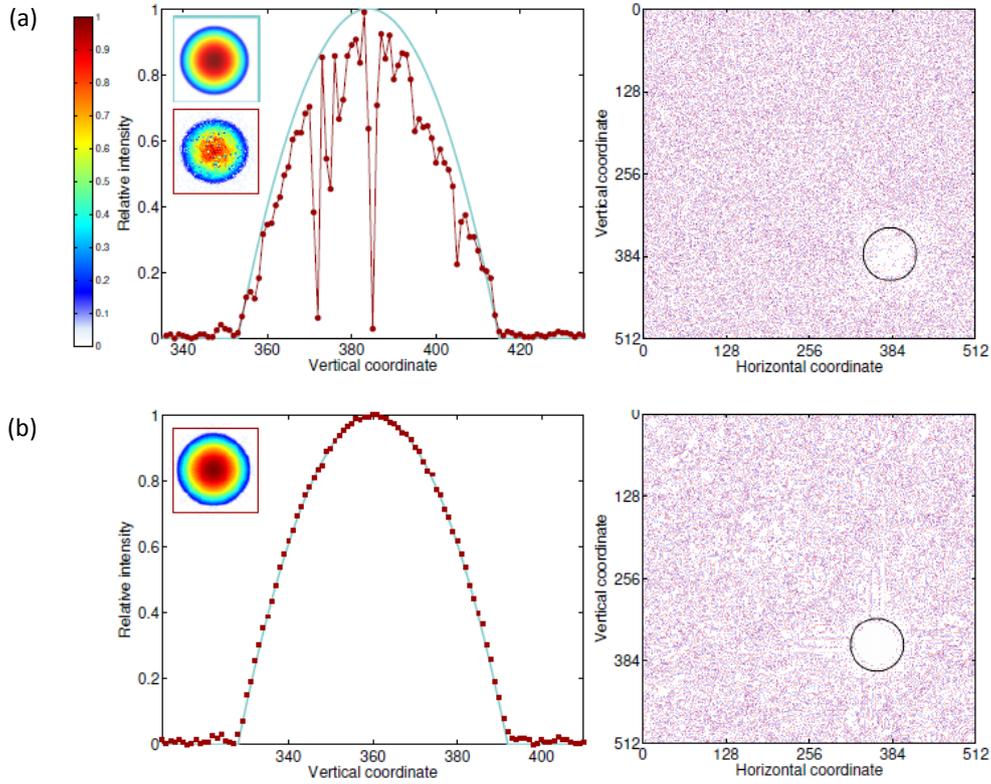


Figure 2.9: Conjugate gradient calculations of a second-order power-law trap. Left: Profiles of the intensity. The colourbar denotes the relative intensity of the insets. Right: Vortex map, where the red pixels in the vortex map are of 2π phase windings, the blue pixels are of -2π phase windings and the black circle corresponds to the trapping (signal) region. (a) An unacceptable intensity profile is generated using the cost function from Equation 2.15. The green inset shows the target intensity and the red inset is the result of 20000 iterations of conjugate gradient minimisation. (b) A far smoother result is achieved using the cost function from Equation 2.16. The reason for this improvement is both a reduction in the number of optical vortices present in the signal region and the migration of the remaining vortices to the very edges of the signal region where the target electric field amplitude is already close to zero. Figures taken from [88].

This removes the low intensity pixels from the trapping region over the course of the calculation (30000 iterations) to create a smooth light pattern, improving the RMS error to 0.43%. This reduction is caused by indirect manipulation of the phase in the output plane. As seen in Figure 2.9, both the number of vortices and their location has improved: a reduction from 232 to 130 vortices within the signal region is beneficial, but of greater importance is the fact that the optical vortices are now present only at the edges of the signal region where the electric field amplitude of the target is intentionally close to zero. This additional specificity that the minimisation technique provides can be used to weight features of interest in the target pattern which the IFTA methods may be unable to address, for example by minimising flicker between consecutive frames of a dynamic pattern [56].

2.2.5 Algorithm Initialisation

The algorithms discussed in the previous sections use Fourier transforms to realise the output plane of the SLM. Computationally, this involves discretising the electric field in to a two-dimensional matrix and using a two-dimensional Fast Fourier Transform (FFT) to calculate the output. However, in order for the FFT not to lose resolution when calculating the output plane, the input electric field needs to be padded. This involves making the matrix larger by introducing zero values to the borders. From the Nyquist theorem, the output plane can be fully resolved by padding an $N \times N$ hologram such that it becomes a $2N \times 2N$ matrix with the $N \times N$ hologram in the centre [57]. An example of padding the initial electric field of a hologram calculation is shown in Figure 2.10. Due to this zero padding, for an SLM of 256×256 pixels, the algorithms actually deal with matrices of 512×512 pixels. Only at the end of the calculation is the unpadded hologram extracted to be used on the SLM. For more details on zero padding see [94].

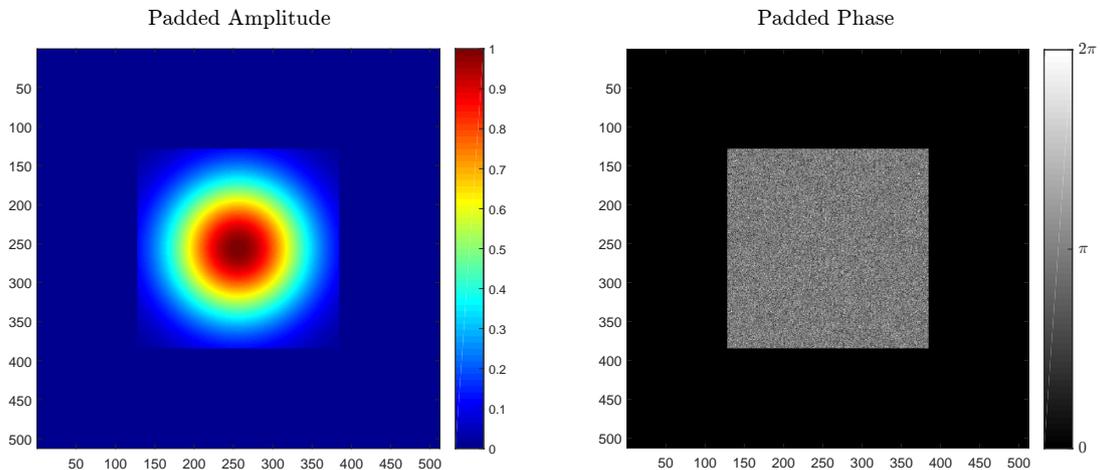


Figure 2.10: Example of padding the initial electric field matrix for a 256×256 SLM in order to fully resolve the resulting output plane calculated by a FFT. In this example the initial phase chosen is a random guess and the amplitude is the input gaussian beam on the SLM plane.

Another important aspect of the initialisation of a hologram calculation is the starting SLM phase. If the algorithm is robust to initial conditions then the initial phase choice

would be arbitrary and could be a random distribution, such as that shown in Figure 2.10. However, random phase values in the output plane can cause speckle noise [95, 96] and it has been shown that an initial guess phase can lead to better convergence for hologram calculation [79] by avoiding early stagnation due to the formation of optical vortices. The guess phases, K_0 , used for both the MRAF and conjugate gradient methods in this work are the same as those in [79] which are a combination of quadratic (K_Q), linear (K_L) and conical (K_C) phase gradients:

$$\begin{aligned} K_Q &= 3R (\alpha x^2 + (1 - \alpha) y^2), \\ K_L &= D (x \cos(\mu) + y \sin(\mu)), \\ K_C &= B \sqrt{(x^2 + y^2)}, \\ K_0 &= K_Q + K_L + K_C. \end{aligned} \tag{2.17}$$

Here, R is the curvature of the quadratic profile and $\frac{\alpha}{1-\alpha}$ is the aspect ratio. For the linear shift, the magnitude of the gradient is D whilst μ is the angle (direction) that the shift takes. Finally, B is the magnitude of the conical gradient. The pixel indices are denoted by x and y .

The quadratic profile, K_Q is a Fresnel lens that changes the spatial extent of the diffracted light, whilst the linear shift K_L is a diffraction grating which moves the light to a desired region of the output plane. The conical phase gradient K_C creates a hole in the centre of the diffracted light and can be useful for an initial guess of a ring shaped target. The combination of these gradients, K_0 can be used as an initial guess phase such that it shapes the light appropriately for a given target. An example of the different phase gradients and an initial guess phase for a ring lattice target is shown in Figure 2.11.

In general, the guess phase has to be optimised manually by adjusting the parameters in Equations 2.17 such that the resulting light in the output plane mimics the target as best as possible, where finer adjustments can be made to further reduce the final error of the calculation.

2.3 SLM Experimental Implementation

When implementing a hologram experimentally with an SLM, the generated light pattern deviates from the expected result due to optical aberrations that arise from imperfections of the optical system [97]. Since an SLM is able to adjust the spatial profile of light, it is also able to compensate for these aberrations by including a correction phase on top of the hologram (summed modulo 2π). This usually involves measuring the wavefront of the beam with a wavefront sensor, although the SLM can also be used to characterise the wavefront itself [98]. One approach for example monitors an interference pattern produced by grating modes that are spatially segmented on the SLM [99, 100]. If the incident wavefront is flat, then the modes should interfere the same way regardless of which region is probed. If there are deviations in the wavefront, however, then the interference pattern will change and a map of the beam's phase can be recorded.

For our holograms we also correct for experimental aberrations; however, we employ the use of a simple feedback process developed previously by our group which bypasses the need to measure the wavefront of the light [59]. Firstly, the technique calculates the target pattern as usual, for example with the MRAF method. The resulting hologram is then applied to the SLM where the diffracted pattern is imaged using a camera. The experimental image is then compared with the original target and a new target pattern

is generated which compensates against any discrepancies. A new hologram is then calculated based on this modified target, which is applied to the SLM. This process is then iterated until a stagnation point or a maximum iteration number is met.

The compensating target is formed by taking the discrepancy in the measure region, $D_i = \tilde{T}_i - \tilde{M}_i$, and adding it to the target for the next iteration with a gain parameter γ (typically chosen to be between 0.3 and 0.6) such that $\tilde{T}_{i+1} = \tilde{T}_i + \gamma D_i$. Here i denotes the iteration of the feedback process, T is the target intensity and M is the measured intensity, where the tildes denote normalisation with respect to the mean value of all pixels in the measured output brighter than 50% of the maximum target value (this normalisation is chosen due to its resistance against low level noise on the camera image).

In order for the discrepancy between the target and measured output to be taken, the camera image is rescaled to match the target. The scaling factors, including angle and magnification, are first determined by applying simple grating phases to the SLM in both the vertical and horizontal directions and then recording the positions of the diffracted spots. The distance and angle between the spots in the measured output is then compared against the target and the scaling factors are taken.

With this feedback approach to correcting aberrations, smoother light patterns have been produced suitable for trapping and RMS errors can be reduced by factors of up to 3 or 4 [59]. An example of such results (taken from [59]) are shown in Figure 2.12 along

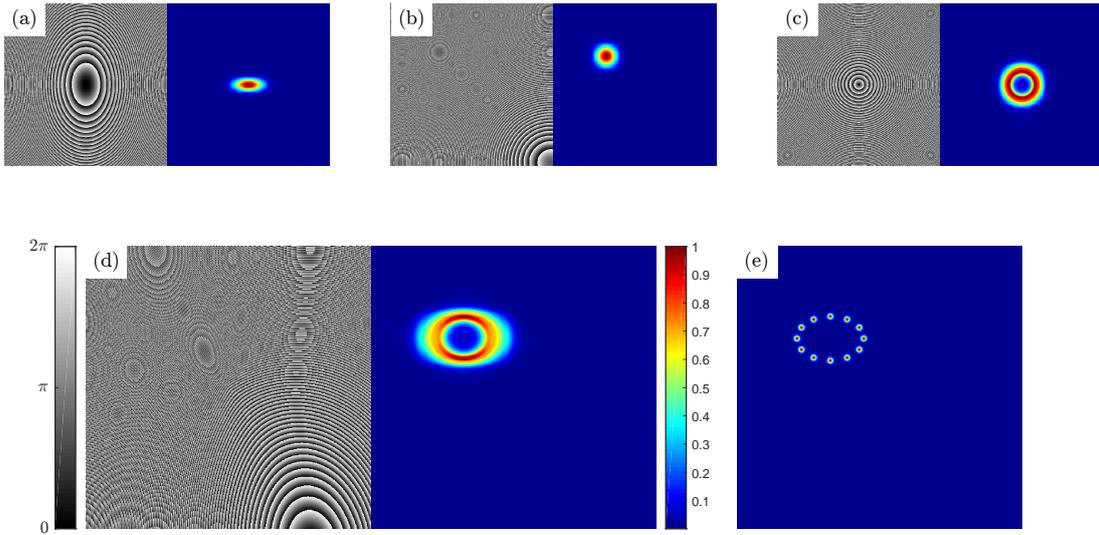


Figure 2.11: Guess phase construction using $K_0 = K_Q + K_L + K_C$ with associated output intensities for an incident gaussian beam. (a) A quadratic phase gradient K_Q is used to spread light across a larger region of the output plane. Here $K_L = 0$ and $K_C = 0$, while $R = 1.5 \text{ mrad px}^{-2}$ and $\alpha = 0.7$ to give the appropriate spread of light in the output plane. (b) The linear phase gradient K_L is used to move the light around the output plane. Now $K_C = 0$, $R = 1.5 \text{ mrad px}^{-2}$ and $\alpha = 0.5$, and the position is set by $D = -\pi/2 \text{ mrad px}^{-1}$ and $\mu = \pi/4$. (c) The conical phase gradient K_C is used to add a hole to the pattern. Here $K_L = 0$, $R = 1.5 \text{ mrad px}^{-2}$ and $\alpha = 0.5$, while $B = 0.35 \text{ mrad px}^{-1}$ adds the hole. (d) Phase guess for an elliptical ring lattice pattern. Here $R = 1.5 \text{ mrad px}^{-2}$, $\alpha = 0.7$, $D = -\pi/2 \text{ mrad px}^{-1}$, $\mu = \pi/4$ and $B = 0.35 \text{ mrad px}^{-1}$. (e) Ring lattice target, for which the guess phase in (d) is suitable.

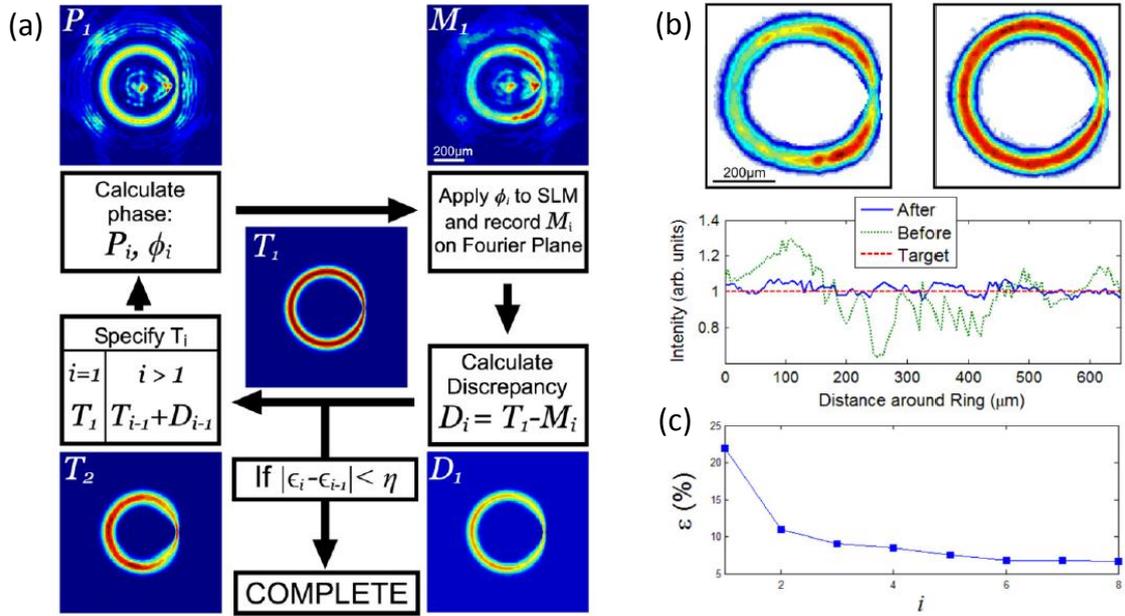


Figure 2.12: (a) The feedback process. Here, P_i denotes the predicted intensity from the MRAF calculation. (b) The measured signal region intensity for a Gaussian ring with a restriction using the initial MRAF-calculated phase profile (left) and after eight iterations of feedback optimisation (right). (Below) the intensity around the circumference showing the target pattern (red), and the measured profile before (green) and after (blue) feedback. (c) RMS error progression for each iteration of the feedback. Figures taken from [59].

with a diagram of the feedback process. In Section 2.4.3, the extension of this method to multiple wavelengths is described along with the experimental results.

2.4 Multi-Wavelength Light Profiles Using MRAF

Previous work in our group has yielded smooth light profiles suitable for atom trapping by using MRAF [55, 56, 59]. Here, I will present an extension to the MRAF method to incorporate different wavelengths, allowing us to produce smooth, multi-wavelength traps through the use of only one SLM [80].

There already exist proposals for cold atom experiments which make use of light fields with multiple wavelengths illuminating the atoms simultaneously but with different spatial distributions, e.g. [101, 102, 103, 104]. Furthermore, multi-wavelength holograms have previously been demonstrated for applications in fullcolour display technology: by using multiple spatially-separated holograms [105, 106]; by time-division [107], spatial-division [108] or depth-division multiplexing [109]; and by illuminating a single phase-pattern with different wavelengths at appropriate angles [108, 110].

We have developed a multi-wavelength method that uses only a single hologram illuminated by co-propagating, overlapped laser beams. This is advantageous to a cold atoms experiment, where optical access is often limited. It also offers the flexibility of adding a new wavelength as the need arises with an easy alignment process and without additional computational demand. The method exploits the regionality of the MRAF technique: we design a target intensity that combines each of the spatial distributions for all wavelengths

within which the position of each feature is set by considering the relevant wavelengths. The resulting calculated SLM phase for such a target will then diffract the multiple wavelengths such that they overlap in the output plane. The MRAF code used in the following investigations was programmed in Matlab.

The diffraction limit in the output plane can be written as

$$\Delta l = 1.22 \frac{\lambda}{2\text{NA}}, \quad (2.18)$$

where λ is the wavelength of the light and NA is the numerical aperture of the lens used to create the light pattern. Since the diffraction limit is linearly dependent on the wavelength of the light, so are a target feature's size and position in the output plane. We use this fact to design target intensity distributions with a distinct feature for each wavelength of light with which the SLM will be illuminated. This is demonstrated by the simple pattern in Figure 2.13, which contains two Gaussians. The positions of the Gaussians in the target pattern are selected such that when the SLM is illuminated with 780 nm and 1064 nm light, the outer Gaussian at 780 nm is overlapped with the inner Gaussian at 1064 nm. To achieve this, the Gaussians in the target are located at a distance $r_{1064} = 208$ pixels (px) and $r_{780} = 1064/780 \times 208 = 283$ px from the centre of the plane. We then calculate the required phase corresponding to monochromatic illumination of the SLM using the MRAF algorithm.

A drawback of the method is that, due to the curvature of the Fourier plane, patterns which extend further away from the zeroth order have aberrations which can effect the desired intensity. Also, the technique requires separable patterns and thus appreciably

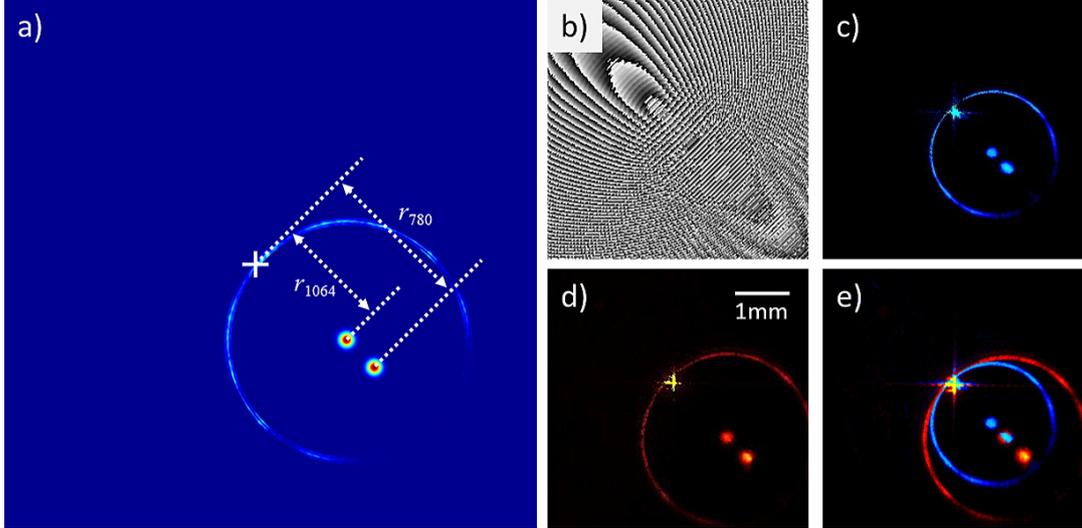


Figure 2.13: a) Target pattern containing Gaussians at $r_{1064} = 208\text{px}$ and $r_{780} = 283\text{px}$, measured from the centre of the output plane, which is the location of undiffracted light in experimental light profiles. The outer ring of light is due to the choice of signal region (light tends to collect at the outer edges of the signal region for our MRAF calculations). b) Phase modulation required to achieve this target pattern. c) Fourier-plane intensity acquired with 780nm (blue) illumination. d) Fourier-plane intensity acquired with 1064nm (red) illumination. e) Fourier-plane intensity acquired with both 780nm and 1064nm illumination (transparency plot).

different wavelengths such that patterns do not need to be diffracted so far from the zeroth order, accentuating the aberrations. However, the technique is easily implemented, in contrast to previous approaches [108, 110], as the incident beams need only be overlapped on top of each other since the different diffraction angles are compensated for by the target itself. Furthermore, whilst the light efficiency is reduced due to wasted light in the output plane, only one SLM is required and the entire SLM plane is utilised by all wavelengths.

Experimentally, the holograms were generated using a single, reflective, phase-only SLM (BNS P256). A total of three diode lasers at different wavelengths were used in the setup, one at 670 nm (Toptica LD-0670-0025-AR-2), another at 780 nm (Sanyo DL7140-201S) and the third at 1064 nm (Roithner RLT 1060-100G), each of which were independently passed through single mode fibres and expanded to a $1/e^2$ radius of 3.5 mm before being overlapped using appropriate dichroic mirrors (Thorlabs DMSP1000L and CVI LWP-45-RP670-TP800-PW-1525-C). After reflection from the SLM (14° AOI), the beams are focused by a single, off-the-shelf, $f = 150$ mm, achromatic doublet. By mounting the camera on a translation stage and observing the focus of the patterns at the different wavelengths, the chromatic shift of the focal plane was measured to be less than $5\mu\text{m}$, where $5\mu\text{m}$ was the smallest increment of translation that the stage was capable of. For alignment, the undiffracted light (which contains $\sim 50\text{-}60\%$ of the incident light power) from each of the three beams is overlapped in the Fourier plane of the SLM, which is imaged onto a Thorlabs DCU224M CCD camera. A diagram of the setup is shown in Figure 2.14.

Returning to our example pattern, when we illuminate the hologram from Figure 2.13(b) with 780 nm light, we measure two Gaussians with $1/e^2$ waist of $91 \pm 3\mu\text{m}$, at 1.34 ± 0.01 mm and 1.93 ± 0.01 mm from the zeroth-order (undiffracted) light. From this and the zeroth-order distances of the target previously mentioned ($r_{1064} = 208$ px and $r_{780} = 283$ px) we calculate the diffraction-limited spot in this optical system as $6.7 \pm 0.2\mu\text{m}$ at 780 nm. With 1064 nm illumination, the Gaussians have $1/e^2$ waist of $130 \pm 9\mu\text{m}$ and are centred at 1.86 ± 0.01 mm and 2.57 ± 0.01 mm, implying a diffraction-limited spot of $9.0 \pm 0.1\mu\text{m}$ at 1064 nm. The farther Gaussian at 780 nm and the nearer Gaussian at 1064 nm are therefore almost overlapped in the output plane, displaced approximately $\delta 70\mu\text{m}$, while both the ratios of the diffraction limits and of the Gaussian widths is equal to the ratio of the illuminating wavelengths, as expected.

2.4.1 Calibrating the SLM for Different Wavelengths

The phase-only SLM changes the phase of the light by altering the orientation and thus the refractive index of the liquid crystals, as mentioned in Section 2.1. The phase response of the SLM can be optimised for a particular illumination wavelength by calibrating the voltages applied to the pixels. The set of voltages to be used on the SLM is referred to as a look-up table (LUT). If multiple colours are incident on the SLM then the LUT cannot be optimised for all of the wavelengths, but a compromise can be found.

In order to calibrate the SLM for different wavelengths we applied a range of checkerboard phase patterns to the SLM with varying phase differences and monitored the undiffracted component of the output intensity. The checkerboard phase pattern creates a grid of spots around the zeroth order (undiffracted light), where the relative phase difference between the blocks of the checkerboard determine the amount of light diffracted. By spatially filtering the diffracted light such that only the intensity of the zeroth order is recorded (using a power meter or photodiode), the phase response of the SLM with respect to the incident light can be determined and a LUT generated.

An example of the LUT generation process is shown in Figure 2.15. The checkerboard

patterns used had blocks with size 4×4 pixels, allowing for a calibration at high spatial frequency suitable for typical holograms. The theoretical response of the zeroth order intensity I_0 with varying phase differences in the checkerboard follows

$$\begin{aligned} \Delta\phi &= 2 \cos^{-1} \left(\sqrt{I_0} \right), & \text{when } \Delta\phi < \pi, \\ \Delta\phi &= 2 \left(\pi - \cos^{-1} \left(\sqrt{I_0} \right) \right), & \text{when } \Delta\phi > \pi. \end{aligned} \quad (2.19)$$

With the known response of the SLM for the checkerboard patterns, the zeroth order can be measured for varying phase levels (using a linear LUT) and compared to the cosine function expected, with a few caveats. The first is the asymmetry between the $\phi < \pi$ and $\phi > \pi$. The intensity of the zeroth order at 2π phase difference between blocks is lower than that of 0 phase difference as the molecules which lie between the pixels experience partial rotation and therefore some diffraction still occurs. Also, the LC molecules have a non-linear response to the voltages applied by the pixels. Rescaling is therefore necessary. Furthermore, the cosine function covers the full 2π range response of the SLM such that not all 256 voltage levels are used. This requires the LUT to be interpolated to fill in the gaps.

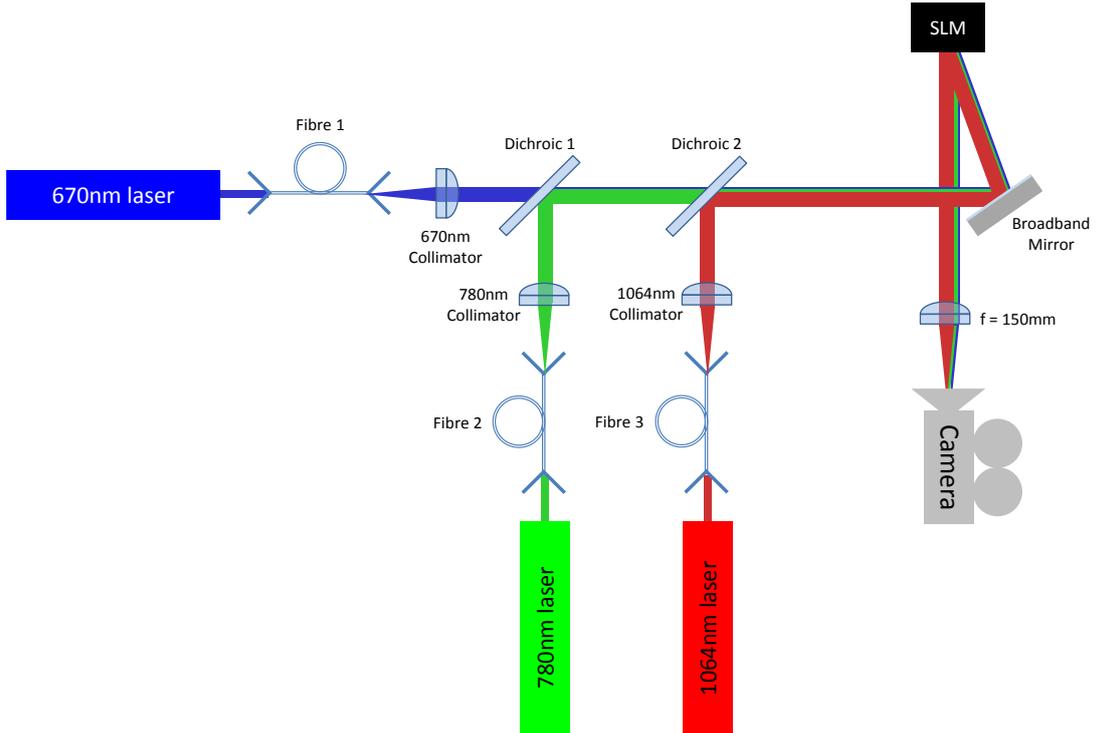


Figure 2.14: Experimental setup for the multi-wavelength holograms. All beams are spatially-filtered by their own single mode fibre, producing a Gaussian profile. For clarity, a single lens is used to represent the beam expansion at the output of the fibres (in the actual setup, more than one lens was used to create the $3.5\text{mm } 1/e^2$ beam radii for each wavelength). Simple beam blocks were used to take images of wavelengths separately. The beams are overlapped on top of each other with the dichroics (the figure shows the beams slightly offset, but this is just for clarity).

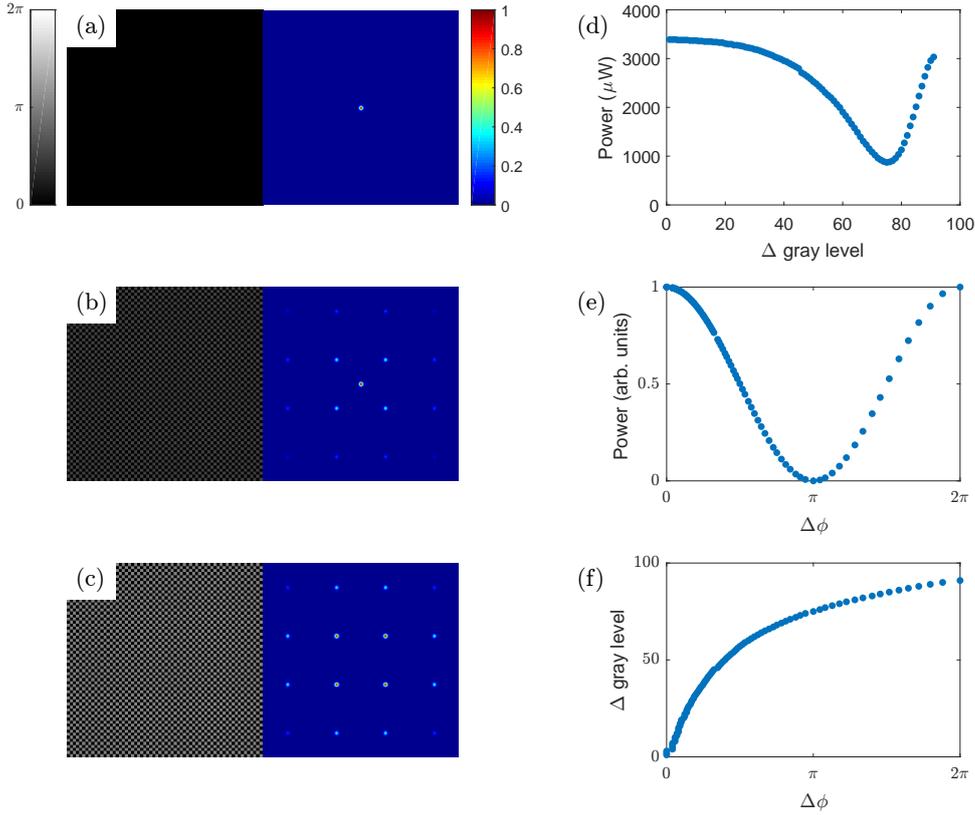


Figure 2.15: LUT generation. Checkerboard phases with simulated resulting diffraction patterns: (a) $\Delta\phi = 0$. (b) $\Delta\phi = \pi/2$. (c) $\Delta\phi = \pi$. (d) Raw zeroth order power as measured by a power meter. The asymmetry of the cosine function is due to the non-linear voltage response of the LC molecules. (e) The rescaled SLM response using Equations 2.19, such that the amplitude is normalised to 1 (for both $\Delta\phi > \pi$ and $\Delta\phi < \pi$) and the phase values range between 0 and 2π . (f) The phase values against gray levels showing the non-linear response of the SLM. These gray levels can now be compared to a phase change and a LUT can be generated (with some interpolation).

With the above method the choice of LUT can thus be optimised. For the three colour patterns in Section 2.4.3 a LUT optimised for 780 nm was chosen as this was the best compromise for all three wavelengths. The light efficiency changes are small, for example, we measured a 7% increase in light efficiency for 1064 nm light diffracted with a 1064 nm LUT as opposed to 1064 nm light diffracted with a 780 nm LUT.

2.4.2 Sub-Diffraction Limited Rings

One application of this multi-wavelength approach is overlapping two *diffraction-limited* spots with wavelengths detuned either side of an atomic resonance. The red-detuned light gives an attractive trapping potential while the blue-detuned light gives a repulsive central potential. As the diffraction limit is smaller for the blue-detuned light, the resultant ring-shaped potential is sub-diffraction limited. Further explanation of the origin of the attractive and repulsive forces for different detunings will be discussed in Section 3.7. While such a simple geometry can also be achieved without holography, an SLM can produce arrays of these small ring traps by tailoring the target profile, while further,

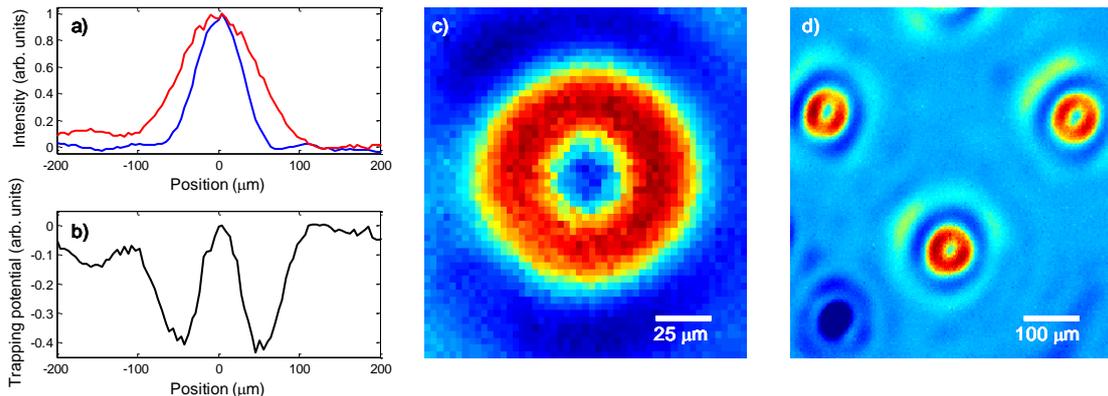


Figure 2.16: a) Intensity profile of diffraction limited spots at 1064 nm (red) and 670 nm (blue). b) and c) Resulting trapping potential, calculated as the difference between the two intensity profiles. The thickness of the ring is $< 50\mu\text{m}$, below the diffraction limit of either colour. d) An array of sub-diffraction limited rings.

more-complicated sub-diffraction-limited features can be designed.

Here we show an example by producing a ring trap with 670 nm and 1064 nm light, which are respectively blue- and red-detuned from the main cooling transitions in rubidium. If we were to design ring traps with only 1064 nm light, their smallest attainable radius would be approximately twice the diffraction-limited spot radius at 1064 nm. By comparison, the radius of the two-colour ring is predicted to be a factor 3.6 smaller than the single-colour ring radius. Compared to a single-wavelength 1064 nm ring, the thickness of the two-colour ring is again predicted to be reduced by a factor of 1.8, giving a higher trapping frequency. Experimentally, we de-magnify the beams impinging on the SLM to a $1/e^2$ radius of 1.1 mm to avoid clipping on the SLM aperture which would aberrate the point-spread function of the diffraction-limited spots. In order to acquire well-resolved images of the light pattern (see Figure 2.16) we also magnify the Fourier plane by a factor of 2.5 using a confocal telescope. The Airy disks of the diffraction-limited spots have $1/e^2$ radii of $53.9 \pm 0.5\mu\text{m}$ for 670 nm and $88 \pm 2\mu\text{m}$ for 1064 nm. When the ratio I/δ (where I is the peak intensity and δ the detuning) is the same for both colours, we measure the resultant ring-shaped trapping potential to have a radius of $53.5 \pm 0.5\mu\text{m}$ and a $1/e^2$ half-width of $44 \pm 3\mu\text{m}$, consistent with our predictions. The generation of arrays of these ring traps, as shown in Figure 2.16d), is accomplished by designing a target profile containing two arrays, with array periodicities and placements being defined by the wavelengths of the illuminating light.

2.4.3 Multi-Wavelength Feedback Algorithm

Our feedback method (as mentioned in Section 2.3) was not applicable for the sub-diffraction limited rings as they were formed from diffraction limited spots. However, in order to correct for aberrations in continuous geometries, a multi-wavelength feedback method was developed by regionalising the algorithm to correct only the relevant output-plane features for each wavelength.

As an example, we consider dressing a standing-wave optical lattice with the trapping potential proposed in [101] for entropy removal of lattice-confined fermionic atoms, which consists of a central attractive dimple potential surrounded by a repulsive barrier to sep-

arate atoms in the dimple from the remainder of the ensemble. In this scheme, which could be implemented in state-of-the-art quantum gas microscopes [111], the dimple and repulsive barrier produce a low-entropy region in the centre, and the high-entropy atoms outside the repulsive barrier can be removed to lower the entropy per particle. A cartoon of the entropy removal scheme is shown in Figure 2.17 (Taken from [101]). This removal could, for example, be achieved using the single-site addressing techniques of Weitenberg, *et al* [68]. In brief, a magic-wavelength light beam is used to impart a differential light shift of the hyperfine ground states of the atoms by ~ 100 kHz. The light-shifted atoms are transferred between the ground states using a microwave pulse. All atoms in one or other state can now be selectively removed by resonant excitation with high efficiency while atoms in the other state remain in the ground state of the potential well [68].

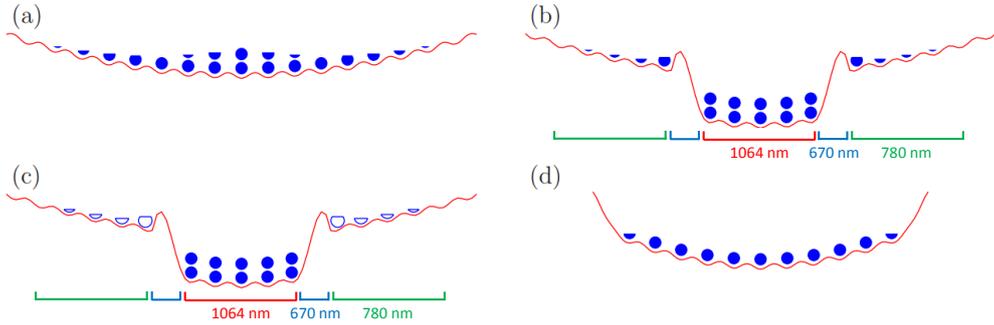


Figure 2.17: Entropy removal scheme taken from [101]. (a) The atoms trapped in a parabolic profile are loaded into an optical lattice. (b) Atoms in a dimple at the centre of the trap form a low entropy band insulator. This core region is isolated from the rest of the system, the storage region, by rising potential barriers. (c) The high entropy atoms in the storage region are removed from the system. (d) The system is relaxed by flattening the dimple and turning off or pushing the barriers outwards.

We have shown that the spatially-varying light patterns in this method can be generated with a single SLM: the central attractive-dimple potential is calculated using 1064 nm light and the repulsive barrier to separate the dimple from the remainder of the ensemble is generated using 670 nm light (which is blue-detuned from resonance for ^{40}K). In addition, we design a tailored profile for magic-wavelength addressing of atoms outside the trapping potential, using our 780 nm light as proof-of-principle (the magic-wavelength of ^{40}K is 768.4 nm and therefore within easy reach of our method). In order to achieve the desired 100 kHz differential light shift between the $4S_{1/2}$ ground states of ^{40}K , 10 mW in the magic-wavelength beam could form, for example, a ring of diameter $120 \mu\text{m}$ with $50 \mu\text{m}$ width. The differential light shift from the 1064 nm light will be very small (~ 40 Hz for 1W of power and a trap radius of $\sim 50 \mu\text{m}$).

Our target intensity profile, shown in Figure 2.18, therefore has three components: a Gaussian with $1/e^2$ radius of 12.4px at $r = 126\text{px}$ from the plane centre; a 27.4px diameter Gaussian ring with $1/e^2$ radial half-width of 5.7px at $r = 172\text{px}$; and a second Gaussian ring of 24px diameter and 5.7px radial half-width at $r = 200\text{px}$. For each of these features we assign a signal region as shown in Figure 2.18, which is where the feedback algorithm compensates aberrations. Converting to output-plane coordinates, all three features will be centred at the same coordinate and their relative sizes will be such that the inner diameter of the 670nm ring overlaps with the edge of the 1064nm Gaussian and the outer

diameter of the 670nm ring overlaps with the inner diameter of the 780nm ring. We calculate an initial phase using the MRAF algorithm, which is displayed on the SLM and illuminated with each of the three lasers in turn. To implement the feedback algorithm, the acquired image is compared to the target profile as in [59]. However, here we compare the acquired image to the target within only the appropriate region for each wavelength, e.g. we ignore the two rings in the image acquired with 1064nm. Within this region the difference between target and acquired pattern is added to the original target to create a new target for a subsequent iteration of MRAF. For all subsequent iterations the same routine applies, except that the difference between target and acquired pattern is added to the *previous* iteration's target. For the example shown, four iterations of feedback reduce the rms error of the 670 nm ring to 8.1%, the 780 nm ring to 9.3% and the 1064 nm Gaussian to 0.5%, from 21.3%, 20.9% and 8.6% respectively. While the algorithm does leave some magic-wavelength light within the region of the central dimple, this will not have a deleterious effect as it will serve to light shift atoms in this region farther from resonance with the microwave source. These feedback-enhanced patterns are shown in Figure 2.19 together with the resulting composite pattern, which could be superimposed on fermions trapped in an optical lattice. If necessary, the light in the remainder of the output plane can be blocked with a pinhole.

To show the generality of our approach, we have created a range of further light patterns which may have use in ultracold atom experiments, as shown in Figure 2.20. Pattern a) consists of a red-detuned elliptical optical trap, which is partially superimposed with a 6th-order super-Lorentzian profile detuned by several tens of GHz from an atomic resonance. The latter would cause the phase of the illuminated and non-illuminated sections of a BEC trapped in the ellipse to evolve at different rates, causing a phase-slip which leads to soliton formation [112]. Our feedback algorithm gives a smooth ellipse with 1.4% rms error and a super-Lorentzian with 1.7% rms variation in the flat-top. Figure 2.20b) shows a red-detuned ring trap with an independent blue-detuned barrier to be used in soliton-interferometry as proposed in [104]. The rms variation around the circumference of the ring

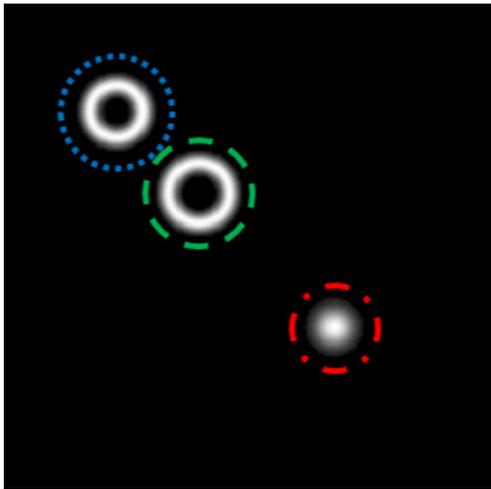


Figure 2.18: Target pattern for lattice-based entropy reduction, showing the subset of the plane containing the signal regions to be used within the feedback algorithm. The target consists of two rings and a Gaussian, designed for 670nm, 780nm and 1064nm. The blue, green and red regions will be optimised for 670nm, 780nm and 1064nm respectively.

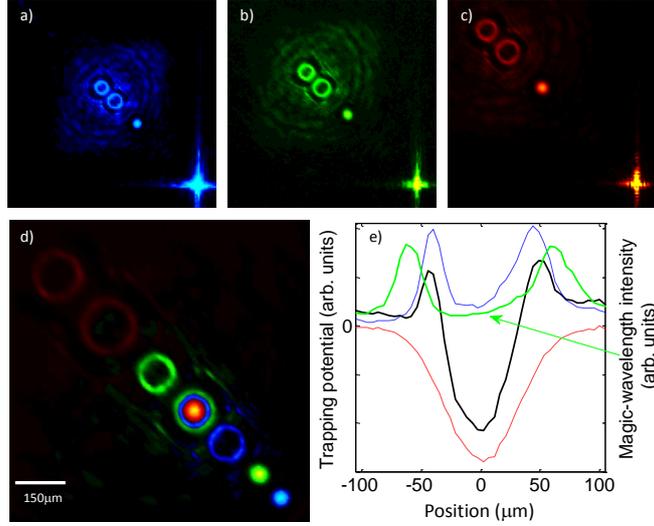


Figure 2.19: Entropy-separation light pattern illuminated with a) 670nm, b) 780nm, c) 1064nm, d) all three wavelengths. e) Trap profile versus position for 670nm (blue) and 1064nm (red) and the resultant combined trapping potential (black) within the region where the three wavelengths overlap. The magic-wavelength light (green) allows spatially-selective microwave state transfer of atoms outside the resultant potential, after which a resonant light pulse (which does not need to be spatially controlled and may be directed orthogonally to the SLM-generated pattern) can selectively remove the transferred atoms.

is 4.5% while the barrier rms error is 4.3%. In previous work [55] we have demonstrated that ring traps of this accuracy are adequate for cold atom experiments.

The study of conduction between two reservoirs along a one-dimensional channel interrupted by repulsive barriers was recently proposed in [113]. An optical potential in this geometry is shown in Figure 2.20c), where the rms variation along the conduction channel is 3.3%. Finally, the authors of [102] have proposed non-equilibrium dynamic behaviour investigations of BECs in attractive potentials with a narrower dimple which is offset from the centre of the trap, with the potential bounded by hard walls. The blue-detuned ring in Figure 2.20d) will produce a steep potential at the edges of the red-detuned trap, which has an rms error of 1.3%.

A test of the success of our approach to multi-wavelength hologram generation is the control of the relative positions of features at different wavelengths. We compare measured and designed multi-wavelength patterns by assuming that one wavelength is well-aligned to its target, and finding the offset of patterns in other wavelengths. When the illuminating beams are well-overlapped both before the SLM and in the zeroth-order of SLM diffraction, after feedback we find that the difference between target and measured positions for the second (and third) wavelengths is less than $5\mu\text{m}$ (one camera pixel) irrespective of the pattern size and distance from the zeroth order. Furthermore, while the chromatic shift of the focal plane of our optical system is small, for more advanced optical systems such as those in [68, 111], the chromatic shift between the wavelengths may be larger. This may be overcome by independently adjusting the initial collimation of each of the laser beams before they impinge on the SLM.

The multi-wavelength scheme is therefore a simple method to produce accurate light

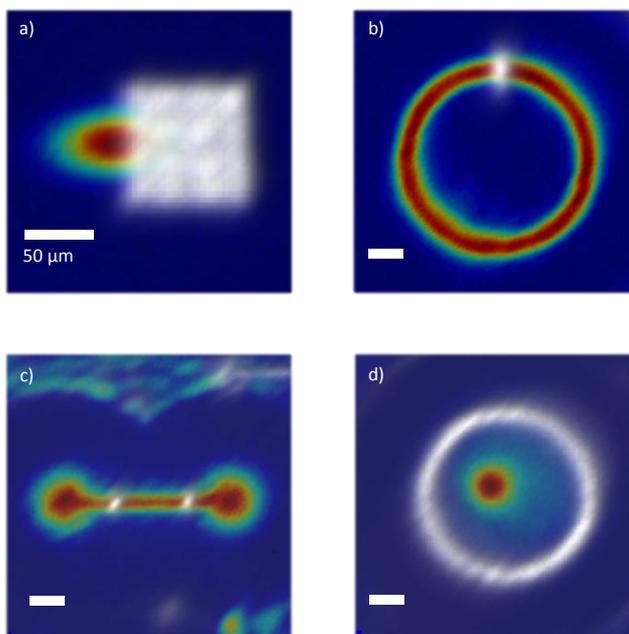


Figure 2.20: A selection of light patterns suitable for cold atom experiments, generated with 1064nm (colour) and 780nm (gray-scale, shown with background transparency for clarity). The scale bar in each image denotes $50\mu\text{m}$. a) Elliptical red-detuned trap, partially illuminated by off-resonant light with a sharp edge for phase manipulation of Bose-Einstein condensates. b) Red-detuned ring trap with blue-detuned barrier. c) Red-detuned double well, connected by a thin channel interrupted by blue-detuned barriers. d) Red-detuned trap with additional offset attractive dimple, bounded by a blue-detuned ring to create hard walls to the trapping potential.

patterns with multiple wavelengths for use in cold atom experiments. In particular, illumination with overlapped beams is desirable from the perspective of cold atoms experiments which often have limited optical access into the experimental vacuum chamber. The method also offers the flexibility to add another wavelength to the experiment with a simple alignment procedure. With only a single phase-only SLM needed, the technique was able to create a variety of trapping patterns with rms errors of only a few percent using a regional feedback scheme, shown to work with 670nm, 780nm and 1064nm wavelengths. Such multi-wavelength schemes would be difficult to achieve with a single DMD or pair of AODs, showing a clear benefit of using a phase-only SLM for atom trapping and manipulation.

2.5 Phase and Amplitude Control with Conjugate Gradient Minimisation

Following on from multi-wavelength control using our phase-only SLM, we also investigated the ability to tailor both the amplitude *and* phase of light in the output plane. As we showed in Section 2.2.4, controlling the output plane phase can give significant improvements in the quality of the trapping potentials. However, there are additional

interesting applications of light fields with control over both amplitude and phase. Simultaneous control over the amplitude and phase of light has allowed significant advances in optical trapping of microscopic objects [114], microscopy [115] and optical communication [116]. Independent spatial control over both the amplitude and phase of trap light is also increasingly desirable in the field of ultracold atoms, for example in the transfer of orbital angular momentum from light to atoms [49], and in the creation of artificial gauge fields [117, 118, 119]. As already mentioned, in the particular case of trapping ultracold atoms in continuous geometries [55, 57, 59, 75, 79, 80], accuracy and smoothness of the intensity are vital to avoid fragmentation.

A variety of methods have been developed which allow arbitrary independent control over both phase and amplitude. Tandem or cascaded approaches sequentially manipulate the amplitude then phase using either two Spatial Light Modulators (SLMs) or two distinct regions of a single SLM [120, 121, 122]. Analytical approaches which calculate a single phase-only modulation to simultaneously sculpt amplitude and phase include the shape-phase method [123] and a variety of methods which spatially control the height, and thus diffraction efficiency, of the applied phase [124]. Recently, IFTA variants have also been demonstrated in controlling both phase and amplitude [125, 126], and a high-fidelity superpixel approach has also been created for DMDs [127].

The following section shows our investigations for an alternative iterative method to creating patterns with independent control over the phase and amplitude profiles, detailing the work in our paper [81]. The method involves the extension of the conjugate gradient minimisation technique described in Section 2.2.4 such that phase and amplitude control is achieved to produce a variety of high fidelity and smooth patterns which are designed primarily for optical trapping. The principle of the calculation is shown in Figure 2.21. The calculation process continues until the cost function stagnates (i.e. when the difference in the value of the cost function between iterations is below 10^{-5}), or when a maximum iteration number of 200 is reached (this maximum iteration number was chosen as it was found that the accuracy would show minimal increase for more iterations, see Figure 2.26a)).

The conjugate minimisation code used for the following investigations was originally written in Fortran by Tiffany Harte and later adapted to Python 2.7, with use of the Theano library [128]. The Python conjugate gradient minimisation uses the Polak and Ribière [89] method. The codes used in this section are freely available online [129].

2.5.1 Cost Function Choice

The main advantage of the conjugate gradient minimisation approach over the IFTAs is the high level of control it gives over any feature of interest in the output plane, provided that the feature can be encapsulated within an analytical cost function C . This defines an effective error to be minimised, and judicious choice of the cost function terms can allow precise guiding of the hologram optimisation process. Therefore, accurately controlling both the amplitude and phase of the light simply requires one to identify a suitable cost function. In order to determine the best choice of cost function, the test pattern used for our initial investigations was a small Gaussian Line pattern in the far-field with a linear phase gradient across it, shown in Figure 2.23(a) (the application of such a light field is discussed in Section 2.5.2). The initial conditions were kept the same for each cost function, with a $1/e^2$ beam radius of 2mm and guess phase parameters of $R = 3\text{mrad px}^{-2}$, $\alpha = 0.5$, $D = -\pi/2\text{mrad px}^{-1}$, $\mu = \pi/4$ and $K_C = 0$.

The error metrics used for initial tests were an intensity RMS ϵ_{MR} , as in Equation 2.9, a relative phase error ϵ_{Φ} , light efficiency η and the fidelity F . The relative phase error is

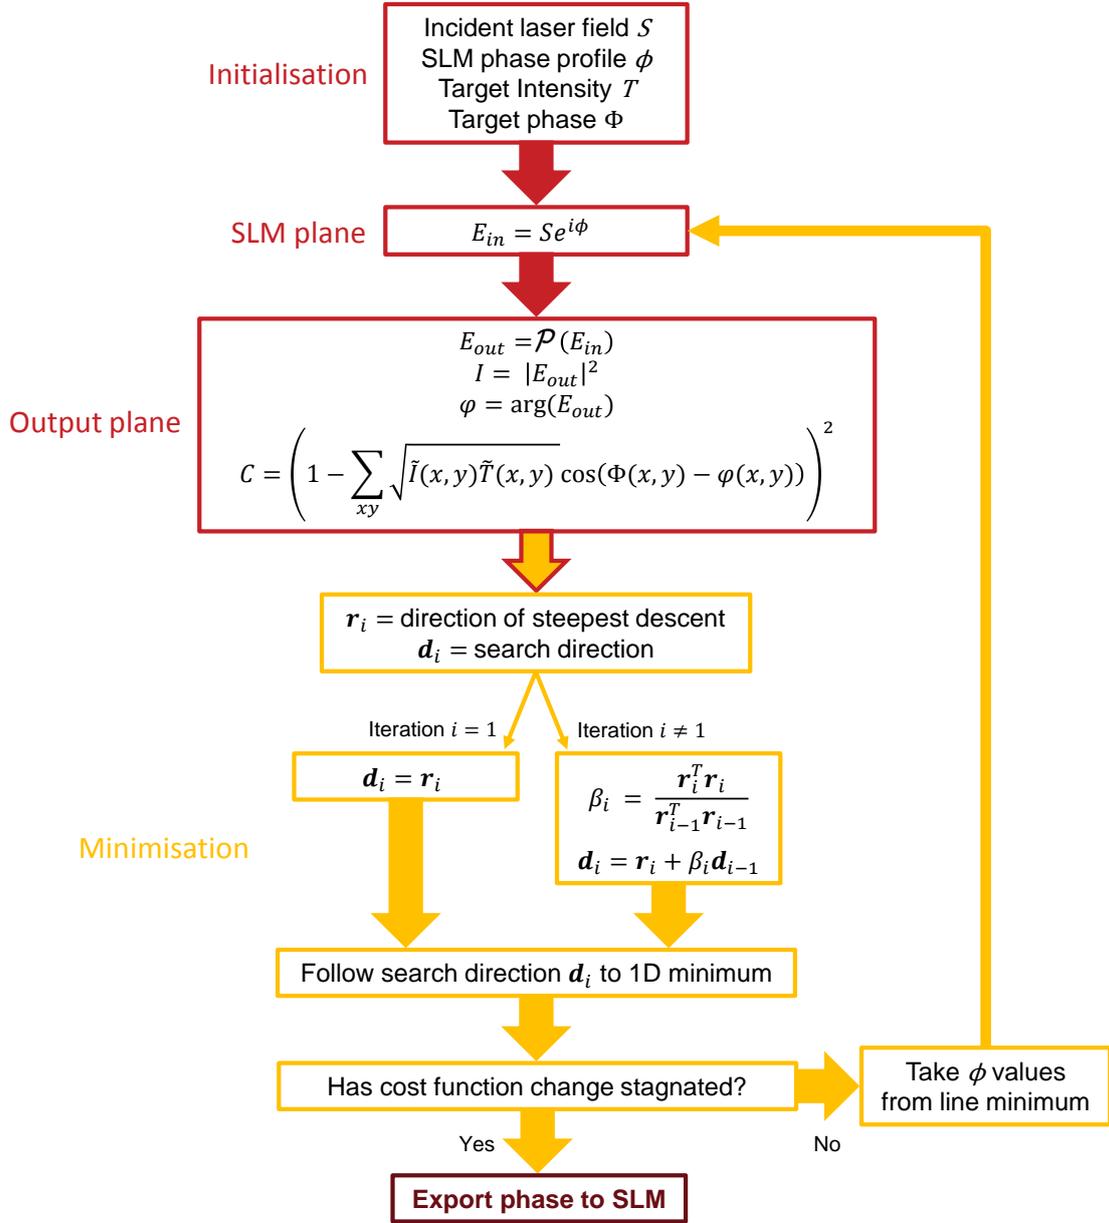


Figure 2.21: Block diagram of the phase and amplitude optimisation process with conjugate gradient minimisation. The \mathcal{P} denotes a light field propagator (which was an FFT in Figure 2.8), where the FFT is used to calculate far-field patterns, whilst ASPW is used for near-field.

given by

$$\epsilon_{\Phi} = \frac{\sum_{x,y} |(\Phi(x,y) - \varphi(x,y) + P(x,y))|^2}{\sum_{x,y} |\Phi(x,y)|^2}, \quad (2.20)$$

where Φ and φ are the target phase and output phase respectively. The P term is a phase correction term which accounts for the cyclical nature of the phase, e.g. a phase target

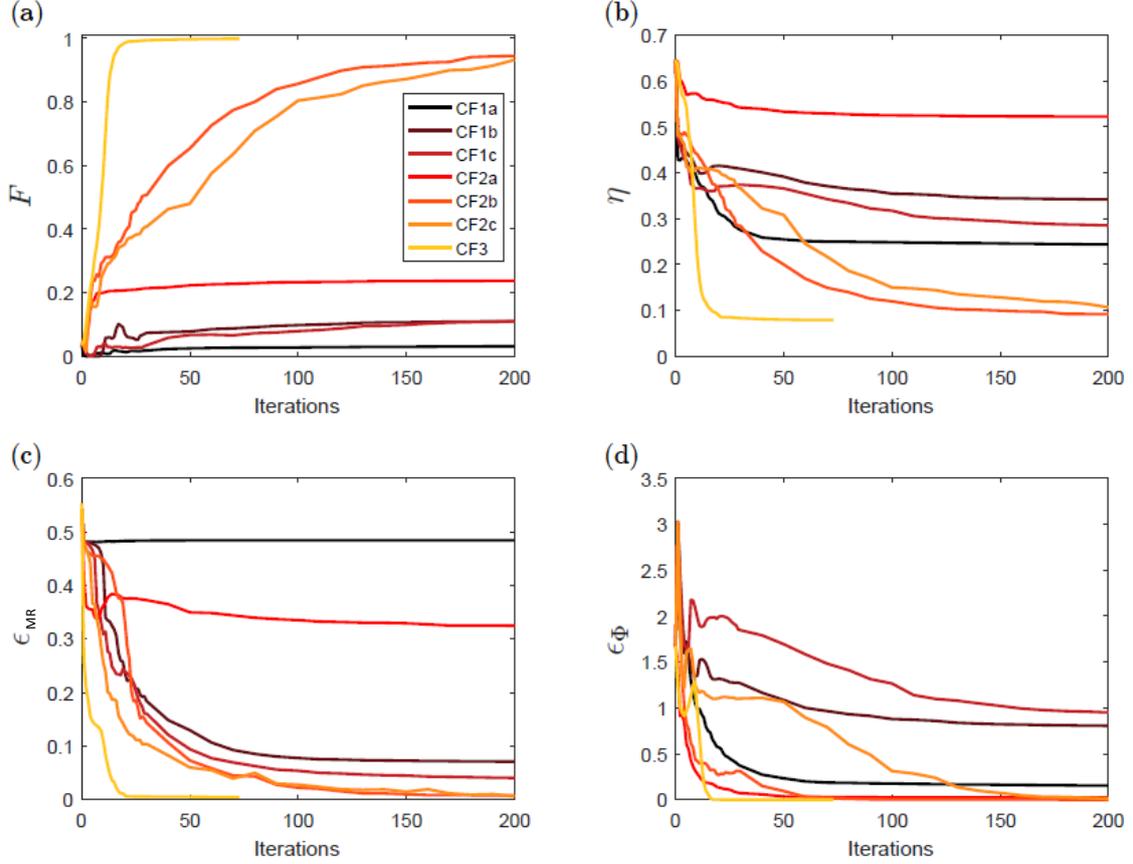


Figure 2.22: Error metrics from the cost function investigation. CF1 corresponds to C_1 from Equation 2.23, where CF1a has $\alpha = 1$ and $\beta = 1$, CF1b has $\alpha = 1$ and $\beta = 0.01$ and CF1c has $\alpha = 1$ and $\beta = 0.005$. CF2 corresponds to C_2 from Equation 2.24, where CF2a has $\alpha = 0.1$ and $\beta = 1$, CF2b has $\alpha = 1$ and $\beta = 1$ and CF2c has $\alpha = 1$ and $\beta = 0.1$. CF3 corresponds to C_3 in Equation 2.26 (without the multiplicative prefactor, i.e. $d = 0$). The resulting patterns of these investigations are shown in Figure 2.23.

pixel of 2π and an output of 0.1π would produce a large error term corresponding to a difference of 1.9π , when the correct error is in fact 0.1π . The fidelity [127] is given as

$$F = \left| \sum_{x,y} \tilde{\tau}(x,y)^* \tilde{E}_{\text{out}}(x,y) \right|^2, \quad (2.21)$$

where $\tilde{\tau}(x,y) = \sqrt{\tilde{T}(x,y)} \exp(i\Phi(x,y))$ is the target electric field. The over-tilde denotes normalisation over a specified region of interest, which was small compared to the total output plane. Similar to the MRAF method described in Section 2.2.2, we chose this region of interest to encompass regions of non-zero amplitude in the target pattern (MR) plus a surrounding area of zero intensity. This signal region is also used for the phase error calculation. The fidelity is effectively the complex inner product of the target and output electric fields, such that for normalised fields (normalised such that the sum of the total in the signal region is equal to 1), the maximum fidelity is 1 corresponding to a perfect match of the target and result.

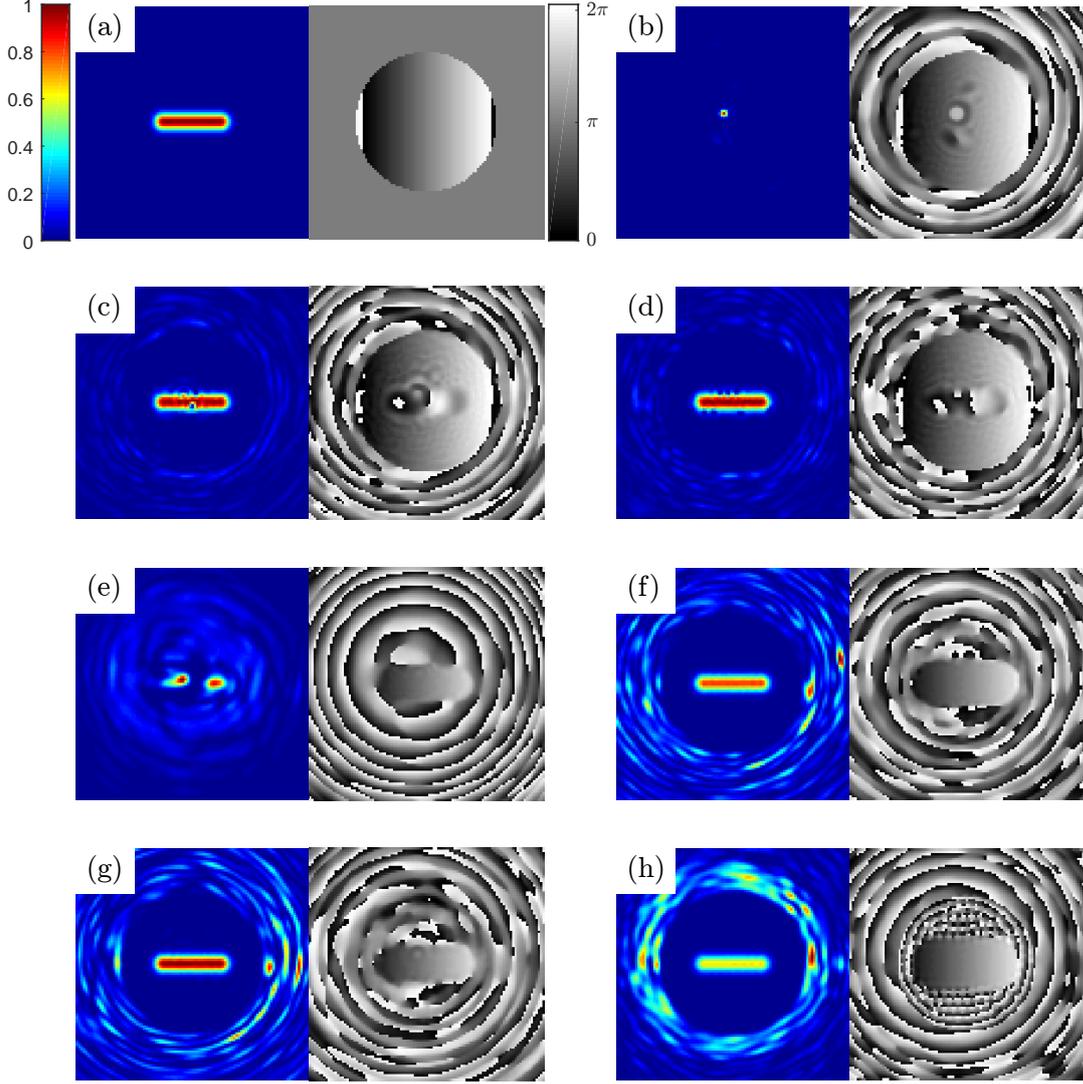


Figure 2.23: Intensity and phase results for different cost functions, corresponding to the calculations in Figure 2.22 (see Table 2.1 for error metrics). (a) Target pattern of a Gaussian line with a linear phase gradient. The predicted output from the hologram calculated in (b) CF1a, (c) CF1b, (d) CF1c, (e) CF2a, (f) CF2b, (g) CF2c and (h) CF3, as defined in Figure 2.22.

The light efficiency η is simply the percentage of light in the region of interest compared to the total output plane. For the final investigations in Sections 2.5.2 and 2.5.3, instead of the intensity RMS being used, we adopted the non-uniformity error from [126] to allow for a comparison with their single SLM method of controlling amplitude and phase. The non-uniformity error is given as

$$\epsilon_{\text{nu}} = \frac{\sum_{x,y} |M(x,y) (\tilde{I}(x,y) - I_a)|^2}{\sum_{x,y} |M(x,y) \tilde{T}(x,y)|^2}, \quad (2.22)$$

where $M(x,y)$ is a binary mask which is equal to one where the target intensity is approximately uniform ($T(x,y) > 0.98$) and zero everywhere else, and $I_a = (1/N) \sum_{x,y} M(x,y) \tilde{I}(x,y)$

is the average output intensity in the uniform region (N is the total number of pixels in the measure region).

Results of the various cost functions investigated for control over the amplitude and phase are shown in Figure 2.22 and Figure 2.23, where the first cost function attempted was a straightforward extension of our previously-used C_a from Equation 2.15, i.e.

$$C_1 = \sum_{x,y} \left[\alpha \left(\tilde{I}(x,y) - \tilde{T}(x,y) \right)^2 + \beta \left(\Phi(x,y) - \varphi(x,y) + P(x,y) \right)^2 \right], \quad (2.23)$$

where α and β are the intensity and phase weighting terms. The difficulty in this cost function is the choice of values for α and β where one must seemingly compromise the other. Investigations with C_1 led to poor results, such that a balance between the phase and intensity terms could not be easily found and high accuracy with respect to the target electric field could not be attained. Furthermore, the phase error also accounts for the low intensity pixels in the signal region, despite the phase constraint being less necessary there.

To combat this unwanted phase constraint, a modification to the cost function C_1 was made by weighting the phase term with the intensity of the target:

$$C_2 = \sum_{x,y} \left[\alpha \left(\tilde{I}(x,y) - \tilde{T}(x,y) \right)^2 + \beta \tilde{T}(x,y) \left(\Phi(x,y) - \varphi(x,y) + P(x,y) \right)^2 \right]. \quad (2.24)$$

The cost function in Equation 2.24 showed a dramatic improvement over C_1 , although it was still dependent on finding the optimal values for α and β , such that a poor choice for the terms would lead to large errors as seen in Figure 2.23(e).

As an alternative to C_2 we also tested an additional cost function which incorporates the fidelity from Equation 2.21. The inner product for the normalised fields has a clear optimal value of 1 which can be used as a minimisation criterion and the form of the function simplifies the equations as the phase correction term P is no longer necessary and the intensity is already weighting the phase constraints. However, a constraint on the chosen cost function became apparent for our technique as a complex function could not be easily differentiated in Python. This meant that the form of the cost function had to take only the real component of the inner product, as follows:

$$C_3 = 10^d \left(1 - \sum_{x,y} \operatorname{Re} \left\{ \left| \tilde{\tau}(x,y)^* \tilde{E}_{\text{out}}(x,y) \right| \right\} \right)^2, \quad (2.25)$$

$$= 10^d \left(1 - \sum_{x,y} \sqrt{\tilde{I}(x,y) \tilde{T}(x,y)} \cos(\Phi(x,y) - \varphi(x,y)) \right)^2, \quad (2.26)$$

where the over-tilde this time denotes normalisation such that the intensity of the entire signal region is equal to 1. The multiplicative prefactor 10^d was later found to increase the steepness of the cost function within the parameter space and improve convergence time and accuracy. The cost function used for the initial investigations shown in Figure 2.22 and 2.23 did not have this prefactor and instead was multiplied by the number of pixels in the signal region.

The various results of the cost function investigation are summarised in Table 2.1. Cost function C_3 proved to be the most promising as it not only optimised the fidelity extremely well as expected, but also did so with the fewest iterations. Whilst the light efficiency is the lowest of all the cost functions (7.9%), the errors are much better, with a

final intensity RMS of 2.7%, a phase error of 0.02% and a fidelity of 0.998. By comparison, the best result from C_2 gave a light efficiency of 9.2%, an intensity RMS of 15.6%, a phase error of 0.03% and a fidelity of 0.944. Even though cost function C_3 stagnated in the fewest iterations, the total time taken for the calculation was actually longer than for C_2 as it required more function and gradient evaluations. The total time taken for the calculation with C_3 was 163 s, whilst the total time taken for the best result from C_2 was 110 s. However, C_3 does not require the added parameters of α and β , therefore no added optimisations are needed on the cost function itself. For these reasons we chose cost function C_3 for our phase and amplitude hologram calculations.

2.5.2 Far-Field Results

To begin with, an investigation of the target size along with initial beam size was undertaken in order to determine the extent of amplitude and phase control possible in the output plane with our method. The results for a Gaussian line, similar to the pattern in Figure 2.23, are shown in Figure 2.24.

As expected, a smaller pattern has both a better fidelity and light efficiency as the smaller signal region results in less pixels being constrained in the algorithm. However, the larger patterns still show promising results. In typical trapping experiments, small, tightly confining light fields are often required anyway, whilst optical setups can adjust the size of images with magnifying or demagnifying telescopes. For this reason the size of the patterns for subsequent investigations with far-field diffraction were kept small, although were sizeable enough to clearly exhibit the desirable amplitude and phase structure.

To exemplify the effectiveness of the technique, we tested our method on a range of target patterns particularly chosen with applications to optical trapping in mind. We calculated a pattern of phase values between 0 and 2π for the SLM plane of 256×256 pixels (with a pixel size of $24\mu\text{m}$). The patterns were diagonally offset from the centre of the plane by 85 pixels to avoid the zeroth order (undiffracted light) that would appear due to the finite efficiency of the SLM. This constrained two of the initialisation parameters to $D = -\pi/2$ and $\phi = \pi/4$. To avoid further initialisation parameters, we also omitted any use of the conical guess phase K_C from Equation 2.17 and kept the aspect ratio of the quadratic phase as 1 : 1.

The region of interest of the calculated intensity and phase for each of our target patterns is shown in Figure 2.25. The pattern similar to a Laguerre-Gaussian (LG) mode provided a good benchmark for our method and such patterns have a wide variety of uses [130], including in ultracold atom experiments to induce circulation states [49]. We could also retain the phase structure of LG modes but with arbitrary amplitude profiles. As

Table 2.1: Cost function investigation error metrics.

	CF	α	β	F	η %	ϵ_Φ %	ϵ_{MR} %
CF1a	C_1	1.0	1.0	0.030	24.4	15	48.4
CF1b	C_1	1.0	0.01	0.110	34.2	81	7.1
CF1c	C_1	1.0	0.005	0.110	28.5	95	4.1
CF2a	C_2	0.1	1.0	0.237	52.2	2.0	32.5
CF2b	C_2	1.0	1.0	0.944	9.2	0.03	0.8
CF2c	C_2	1.0	0.1	0.932	10.7	1.0	0.6
CF3	C_3	-	-	0.998	7.9	0.02	0.5

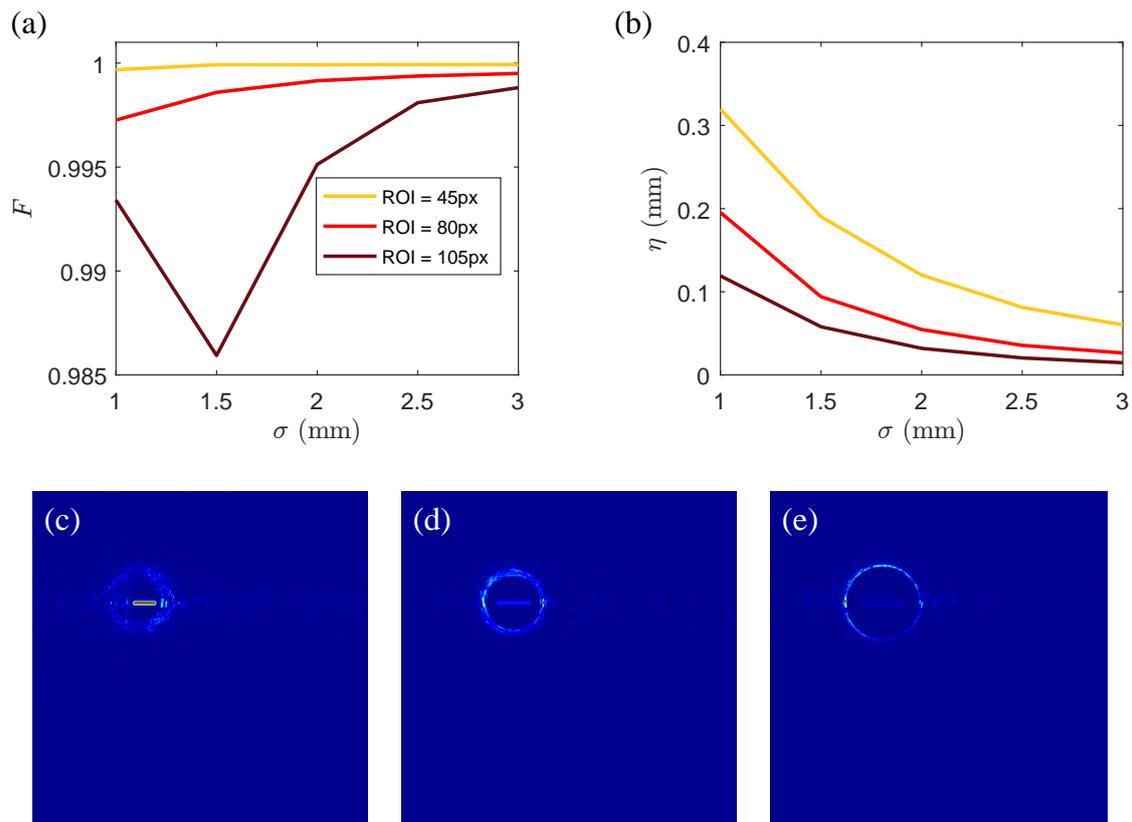


Figure 2.24: Far-field calculations for the Gaussian line pattern with different initial beam sizes and target sizes. The region of interest (ROI) corresponds to the signal region and is circular, centred on the target pattern (see Figure 2.4). The guess phase curvature R chosen for the investigation was 2mrad px^{-2} . (a) Fidelities with different beam sizes and target sizes. (b) Light efficiencies with different beam sizes and target sizes. (c)-(e) Full plane output intensities with ROI = 45 px (Gaussian line size 30 px), ROI = 80 px (Gaussian line size 50 px) and ROI = 105 px (Gaussian line size 70 px) respectively, all with an initial input Gaussian $1/e^2$ half-width of $\sigma = 2$ mm.

examples, ring and square lattices with underlying phase windings have potential applications for quantum simulation of magnetic flux in solid state systems [117]. Ultracold atoms confined in a honeycomb lattice with alternating phase between nearest neighbouring sites have also been shown to experience an artificial gauge field in a graphene quantum simulator [118], while a trapping potential comprising a flat intensity profile and an inverse square power-law phase has been proposed for investigations on sonic horizons and artificial black holes [119]. A Gaussian line with a phase gradient across it can be used to trap particles in optical tweezers, but at the same time cause them to flow [131]. As a test of our method's versatility, we also chose the more arbitrary patterns of a chicken and eggs [132] which have uncorrelated intensity and phase patterns.

For the example of the Gaussian line pattern (with $\sigma = 1.5\text{mm}$ and $R = 3.5\text{mrad px}^{-2}$) Figure 2.26a) shows the evolution of the fidelity through the calculation for different values of the steepness parameter d in Equation (2.26). Lower values of d cause early stagnation of the algorithm into poor quality local minima in comparison to larger values of d . This

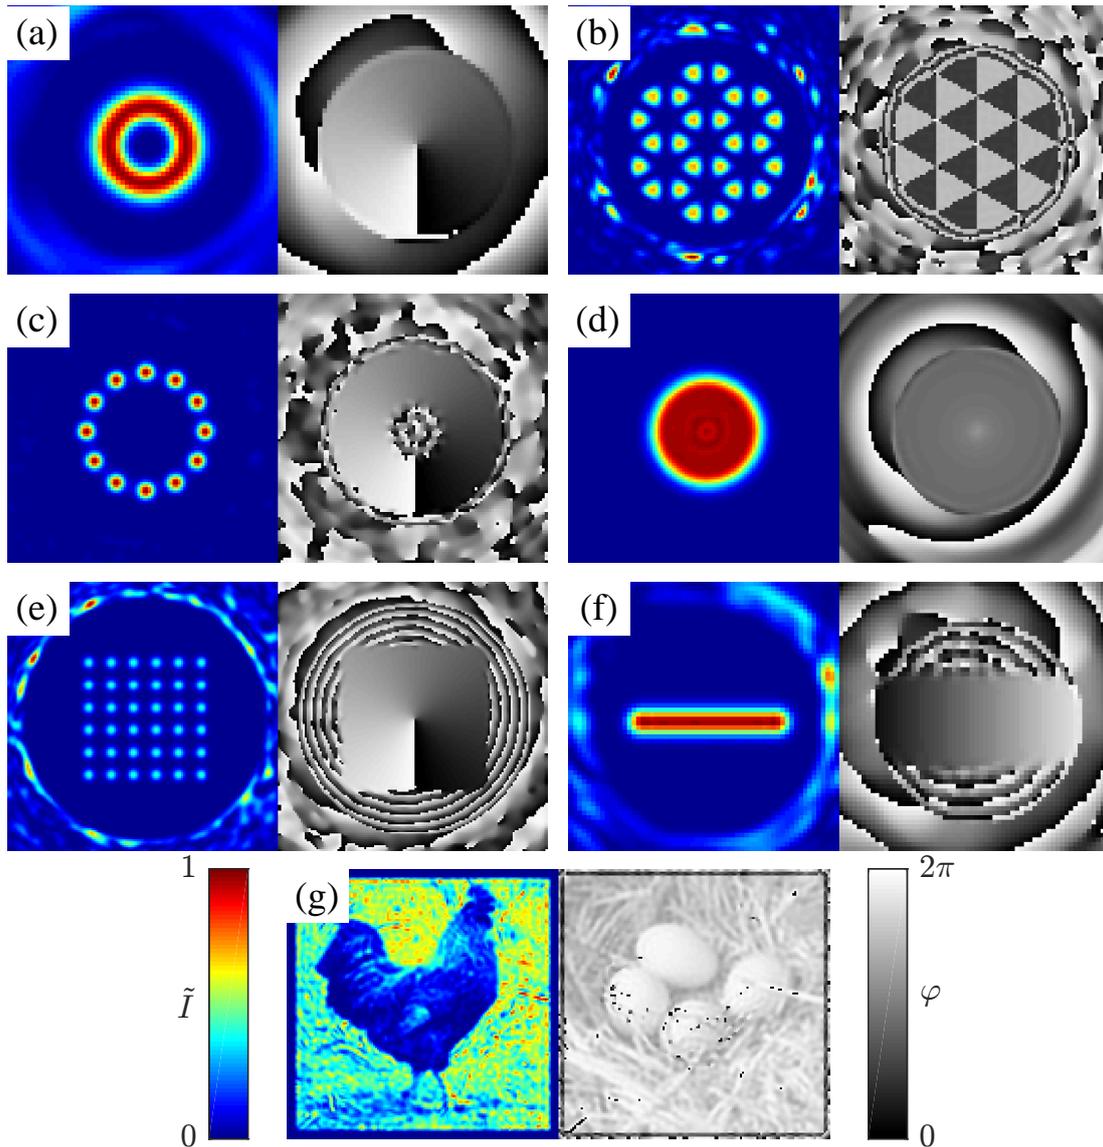


Figure 2.25: The far-field results from the conjugate gradient optimisation showing normalised intensity \tilde{I} (colour) and phase φ (grey) in the region of interest. The flat top pattern (e) has the light outside the measure region removed for clarity whilst the arbitrary pattern (g) does not show the noise region, also for clarity. The error metrics for each pattern are shown in Table 2.2.

is due to the stagnation criterion being kept the same for all calculations (norm of the gradient less than 1×10^{-5}) and thus changing the steepness parameter effectively changes the point at which the calculation stops. For example, a small enough steepness parameter could result in the first gradient calculated to be less than the stagnation criterion, thus immediately stopping the calculation. The maximum iteration number was reached for $d > 6$, whilst the fidelity would increase at approximately the same rate for $d > 4$ (only $d = 1, 2, 3, 6$ and 9 are shown in Figure 2.26 for clarity). It was found that a steeper cost function would not only lead to improved fidelities in the patterns, but also faster calculation times per iteration t (Figure 2.26b)). Multiplying the cost function by a prefactor

initially was thought to only effect the intervention of the stagnation criteria. However the different values of d , especially at low values, take different trajectories regarding error reduction, and also the $d = 1$ result actually reaches stagnation after $d = 2$ or $d = 3$. This may be an indication of floating point rounding errors in the calculation, where more information is lost when the computer deals with very small numbers which have many decimal places. A typical minimisation routine converges in < 200 iterations at a total duration of < 75 s with a standard desktop computer (2.5 GHz processor) with $d = 9$. For all patterns shown in this thesis, we have used $d = 9$ unless stated otherwise.

For each pattern we perform an optimisation over the initialisation conditions σ and R (see Figure 2.26c-d)). It was found that smaller incident laser beam sizes and reduced curvature in the guess phase led to higher light efficiency at a reduced fidelity. The beam size and curvature for the patterns in Figure 2.25 were chosen to provide both good light efficiency whilst maintaining a high fidelity. The optimal values of calculated holograms

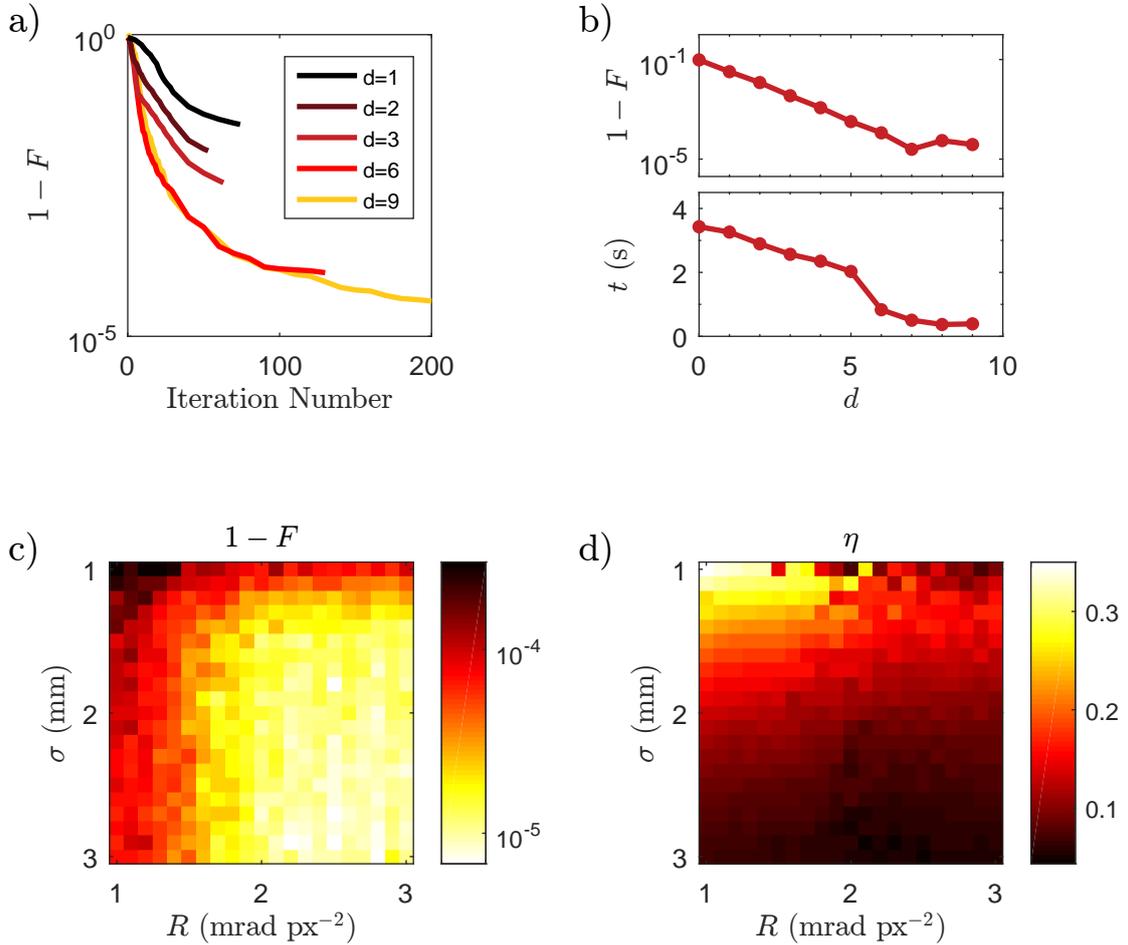


Figure 2.26: a) Evolution of fidelity F for the Gaussian Line pattern shown in Figure 2.25f) with $\sigma = 1.5$ mm and $R = 3.5$ mrad px $^{-2}$. Cost function used was C_3 . At low values of the steepness d of the cost function, the algorithm stagnates earlier and returns a lower fidelity hologram. b) The final fidelity and the time per iteration t as a function of d . c) Fidelity and d) efficiency η as a function of incident laser beam size σ and quadratic guess phase curvature R .

are shown in Table 2.2.

The authors of [126] recently developed an IFTA for full-plane control of amplitude and phase, which they compared to a previous regionally-constrained algorithm [125]. They find that the regionally-constrained algorithm is more accurate at the cost of light-utilisation efficiency, which has also been seen in amplitude-only control algorithms [79, 88] and in the present work. For far-field holograms of lines of continuous intensity with phase gradients, the regional algorithm gives $\epsilon_{\text{nu}} = 0.04\%$, $\epsilon_{\Phi} = 1.63\%$ and $\eta = 3.48\%$, while the full-plane IFTA is less accurate ($\epsilon_{\text{nu}} = 3.48\%$ and $\epsilon_{\Phi} = 3.77\%$) but achieves higher efficiency ($\eta = 77.84\%$). For our chosen cost function in Equation (2.26), the comparable continuous patterns amongst our range of targets (i.e. the Gaussian Line and Flat Top) are significantly smoother: we find ϵ_{nu} is lower by a factor 6-20 and ϵ_{Φ} is lower by one or two orders of magnitude than the regional IFTA. The light-utilisation of the conjugate gradient optimised patterns is a factor 3-11 times higher than the regional IFTA, but between 15-53% of the full-plane IFTA. If light utilisation is of greater importance, the freedom in choice of the cost function terms and their relative weightings could be exploited to prioritise the efficiency at the expense of accuracy or smoothness.

2.5.3 Angular Spectrum Wave Propagation Results

The conjugate gradient method is not only limited to far-field patterns. By changing the propagator for the light fields from a Fourier transform to Angular Spectrum Wave Propagation (ASWP) [78] we are also able to optimise near-field patterns, exemplifying the versatility of the approach.

The ASWP method converts the electric field into an angular spectrum (via Fourier transform) which is built up of plane waves travelling in unique directions. This angular spectrum is then multiplied by a propagation term and converted back to give the electric field in the desired plane. Further details on the propagation method are covered in [78].

A complex field in the $z = 0$ plane can be written in Fourier representation as

$$U(x, y, 0) = \iint A(f_x, f_y, 0) \exp(2\pi i(f_x x + f_y y)) df_x df_y, \quad (2.27)$$

where f_x and f_y are the frequency Fourier pairs of x and y , and $A(f_x, f_y, 0)$ is the angular spectrum which is simply the two-dimensional Fourier transform of $U(x, y, 0)$. The angular spectrum method is able to determine the effects of propagating the field in the z direction via a solution to $U(x, y, z)$ in terms of the original angular spectrum $A(f_x, f_y, 0)$. The

Table 2.2: Error metrics for the calculated patterns in Figure 2.25, with optimal values of σ , R and region of interest diameter ROI.

Pattern	σ mm	R mrad	ROI px ⁻² px	$1 - F$	η %	ϵ_{Φ} %	ϵ_{nu} %
2.25(a) Laguerre Gauss	1.0	4.5	42	3.0×10^{-6}	41.5	0.0003	0.005
2.25(b) Graphene	1.4	2.7	78	4.4×10^{-4}	13.1	0.0003	0.010
2.25(c) Ring Lattice	1.2	3.9	71	1.5×10^{-6}	24.6	0.00006	0.001
2.25(d) Flat Top	1.0	4.5	63	1.8×10^{-4}	11.3	0.2	0.007
2.25(e) Square Lattice	1.2	4.5	124	1.6×10^{-5}	10.6	0.0009	0.02
2.25(f) Gaussian Line	1.4	2.9	45	1.4×10^{-5}	20.4	0.001	0.002
2.25(g) Chicken & Egg	1.6	4.5	128	7.1×10^{-2}	2.0	1.3	-

propagated field can be written as

$$U(x, y, z) = \iint A(f_x, f_y, z) \exp(2\pi i(f_x x + f_y y)) df_x df_y, \quad (2.28)$$

which must satisfy the Helmholtz equation ($\nabla^2 U + k^2 U = 0$), leading to the solution

$$A(f_x, f_y, z) = A(f_x, f_y, 0) \exp(i\mu z), \quad (2.29)$$

where $\mu = \sqrt{k^2 - 4\pi^2(f_x^2 + f_y^2)}$. If μ is real then the propagated angular spectrum is simply the initial angular spectrum with a phase factor and the propagated field can be determined via a sequence of Fourier transforms from the initial field:

$$U(x, y, z) = \mathcal{F}^{-1}[\mathcal{F}[U(x, y, 0)] \exp(i\mu z)]. \quad (2.30)$$

In regards to our setup, k corresponds to the laser light wave vector and f_x and f_y are the spatial frequencies with regards to the dimensions of the SLM. Unlike the far-field calculations which used a single Fourier transform, the ASWP method allows for the propagation distance to be specified and near-field holograms can be calculated.

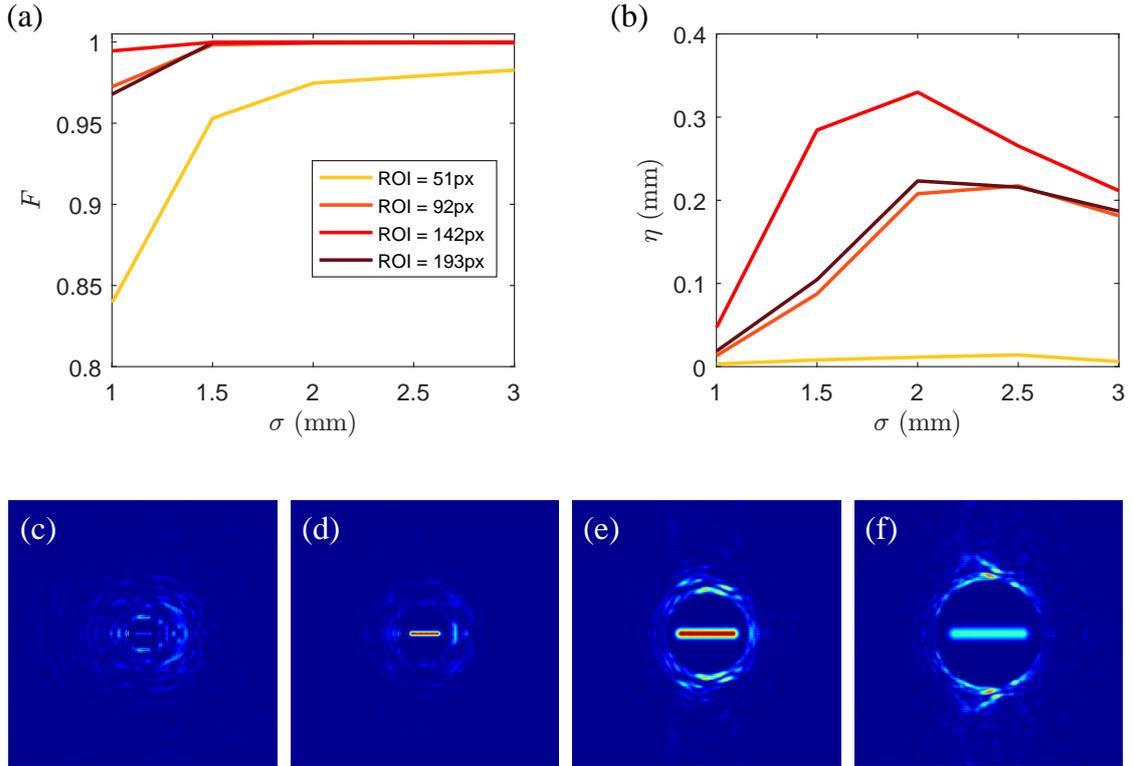


Figure 2.27: Target size investigation with ASWP. The initial guess phase for the calculations is flat. (a) Fidelities for different beam sizes and target sizes. (b) Light efficiencies for different beam sizes and target sizes. (c)-(f) Full plane output intensities for ROI = 51px, ROI = 92px, ROI = 142px and ROI = 193px respectively, all with $\sigma = 2$ mm. Lower light efficiencies and fidelities are seen for the smaller pattern in comparison to the far-field results due to the nature of the ASWP method.

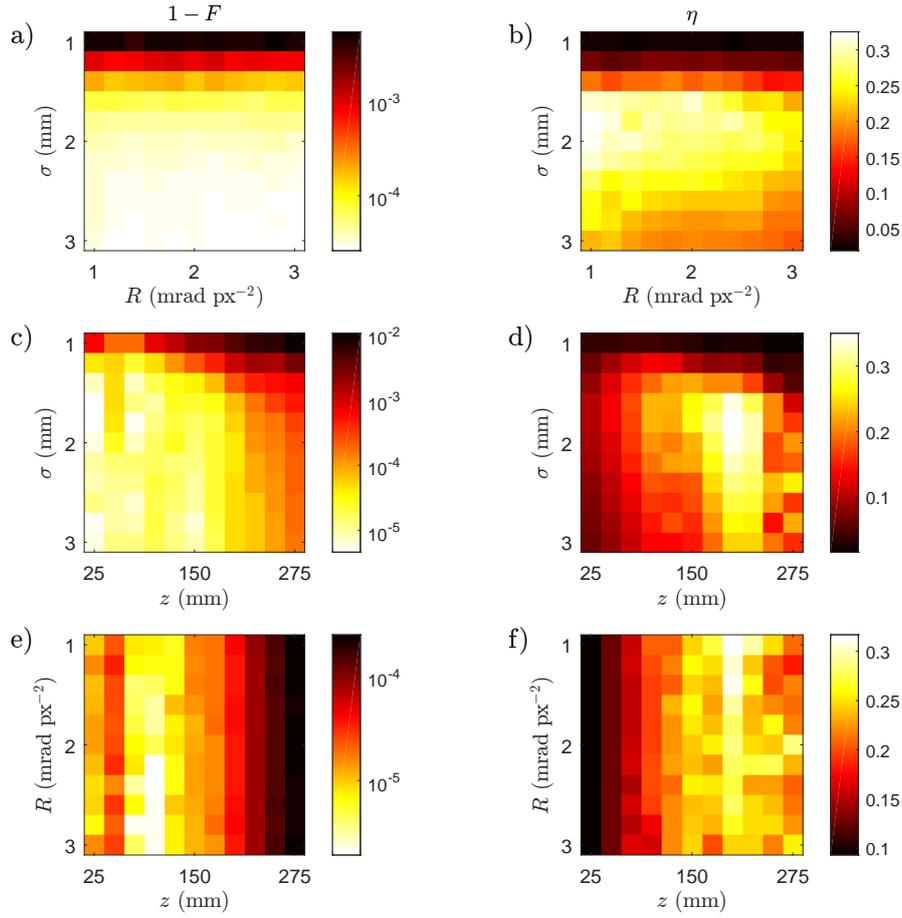


Figure 2.28: ASWP Gaussian line pattern optimisation as a function of guess phase curvature R , propagation distance z and illuminating beam size σ with fidelity results on the left and light efficiency on the right. For a) and b), the propagation distance was kept at 200mm. For c) and d), the guess phase was flat. For e) and f), the beam size was kept at 2mm.

With the ASWP investigation, it was found that larger target patterns corresponded to improved light efficiency and fidelity with the beam sizes chosen for the calculations, differing from the far-field calculations which performed better for smaller sizes. Even though the number of constrained pixels is higher, for the ASWP propagation technique, a Fourier transform and an inverse Fourier transform (along with a small-distance propagation term) are used, such that a target pattern which is comparable in size to the initial beam is most favourable for light utilisation. For all calculations in Figure 2.27, the $1/e^2$ beam waist of the incident beam was 117 px, which is most similar in size to the length of the Gaussian line target (130 px) of Figure 2.27(e) which was indeed both the most accurate and most efficient result. The patterns were therefore chosen to be larger than the patterns for the near-field investigation.

Since the use of the ASWP method has the additional parameter of the propagation distance z , we investigated the effects of this on the Gaussian line target, along with the effects of using the guess phase from the far-field calculations and initial beam size. The results are shown in Figure 2.28.

The results of the parameter scans for the ASWP calculations show that the strongest dependence is on the propagation distance and initial beam size. The guess phase curvature had the least effect on the fidelity and efficiency. For this reason the guess phase was kept flat for further calculations to reduce the free parameters and ease the optimisation process. The optimised results are shown in Figure 2.29 and the corresponding errors and parameters are shown in Table 2.3.

2.5.4 Experimental Verification

We verify the far-field calculated holograms experimentally using the setup shown in Figure 2.32a). The output of a 1070nm fibre laser (IPG Photonics, YLP-5-1070-LP) is expanded

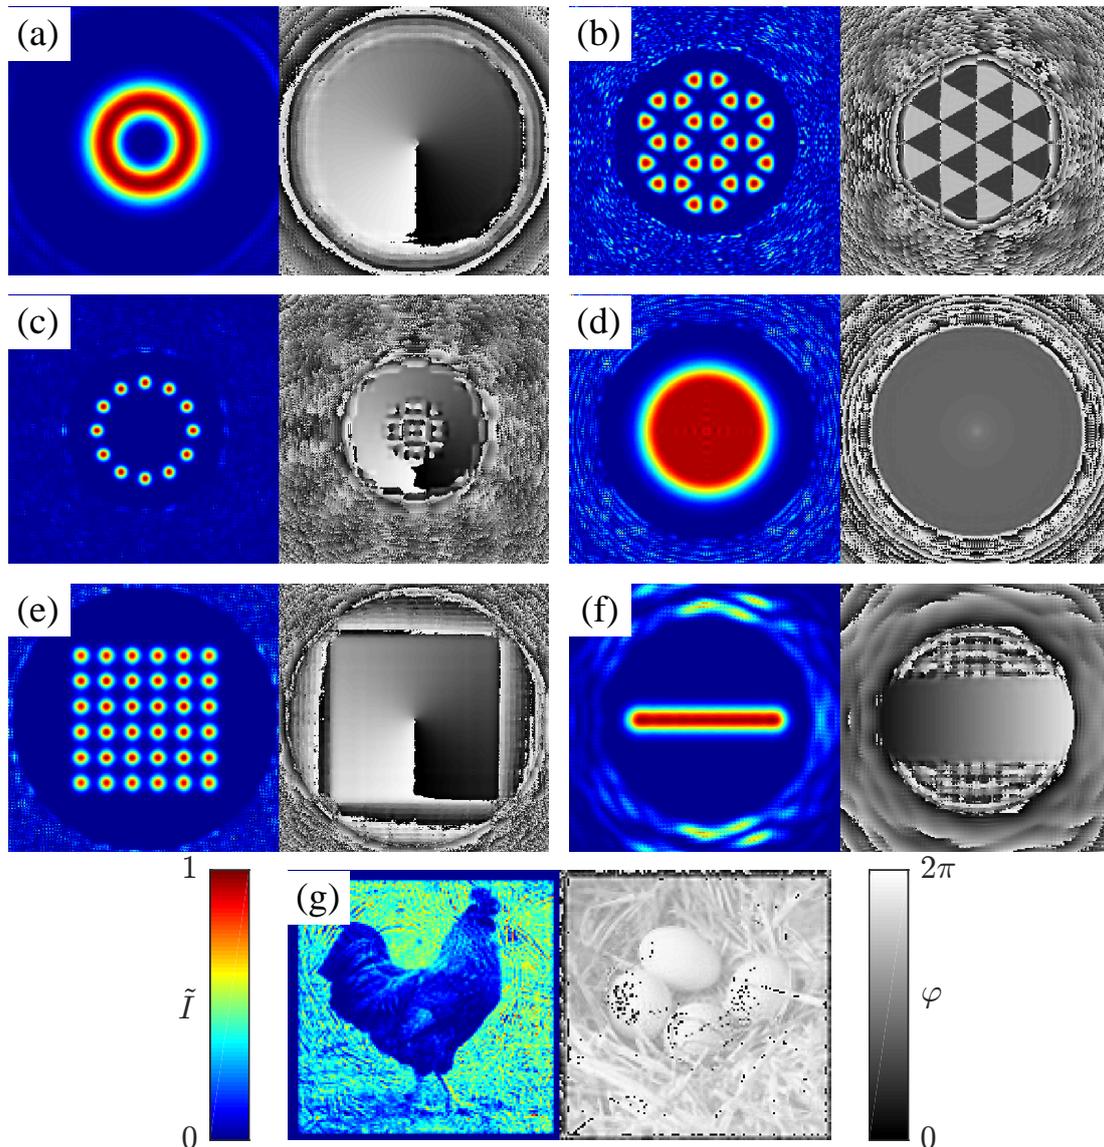


Figure 2.29: The near-field results from the conjugate gradient optimisation showing normalised intensity \tilde{I} (colour) and phase φ (grey) in the region of interest. The error metrics for each pattern are shown in Table 2.3.

Table 2.3: Error metrics for the calculated patterns in Figure 2.29, with optimal values of σ , z and region of interest diameter ROI.

Pattern	σ mm	z mm	ROI px	$1 - F$	η %	ϵ_Φ %	ϵ_{nu} %
a) Laguerre Gauss	1.6	200	196	3.0×10^{-5}	65.7	0.0005	0.004
b) Graphene	3.0	100	198	3.0×10^{-4}	21.8	0.020	0.004
c) Ring Lattice	2.8	200	156	1.0×10^{-3}	23.8	0.030	0.003
d) Flat Top	2.8	125	155	5.6×10^{-5}	27.2	0.014	0.002
e) Square Lattice	3.0	200	266	1.7×10^{-3}	31.4	0.029	0.130
f) Gaussian Line	2.0	200	142	5.0×10^{-5}	33.0	0.002	0.004
g) Chicken & Egg	2.8	50	128	4.3×10^{-2}	8.2	0.600	-

to an experimentally-convenient $1/e^2$ waist of 3.0mm and split using a polarising beam splitter. One path is phase-modulated as it is reflected (14° AOI) by a liquid crystal SLM (BNS P1920) and focused onto a CCD camera (Thorlabs DCU200 Series) using an $f = 150\text{mm}$ achromatic doublet. The other path gives a reference beam which is optionally recombined with the modulated beam after the focussing optic to produce interference fringes which are used to extract the phase of the modulated light via the Fourier transform fringe analysis method [133], described later in this Section.

As an initial test of the patterns however, the reference beam was not used and the propagation of the light field out of the focussing plane was instead observed by moving the camera on a translation stage. Since the hexagonal lattice pattern had a periodic flat phase target, this was ideal for indirectly observing the phase control of the light. A comparison between the hexagonal lattice pattern with phase control against the same pattern without phase control could be made, where the flat phase front should maintain the intensity structure of the light out of the image plane far better than the field with random phase. One pattern was therefore calculated with the cost function C_3 from Equation 2.26, whilst the other was calculated with C_a from Equation 2.15 which leaves the phase unconstrained. Images from this initial investigation are shown in Figure 2.30.

As can be seen from the comparison in Figure 2.30 the phase controlled target does maintain the Graphene pattern structure well, indicating that the light field has the desired phase structure and that the target has been replicated accurately. The noise region however does not have any phase control and if the pattern is observed further out of plane than shown in Figure 2.30, the light in the noise region eventually encroaches on the signal region. However, over the range of ~ 3 mm, the phase controlled Graphene pattern maintains its structure as opposed to the unconstrained pattern which loses its shape if propagated out of the plane by ~ 0.5 mm. Intensity fluctuations throughout propagation are also more apparent without phase control, where a single diagnosis spot in the lattice was found to have a propagation intensity RMS deviation of 13.9% for C_a , as opposed to 4.5% for C_3 over the ~ 3 mm distance.

In order to better diagnose the patterns however, extraction of the phase itself is desirable such that an error can be ascertained with respect to the target. With the Fourier transform fringe analysis method, only one picture of interference fringes is required to extract the phase of a beam (an example of the observed fringes for the Laguerre-Gaussian pattern is shown in Figure 2.31(a)). The technique involves Fourier transforming the interference image formed by overlapping a large flat-phased reference beam and the light pattern of interest, where extraction of the phase is done by isolating the desired spectra via filtering and frequency shifting.

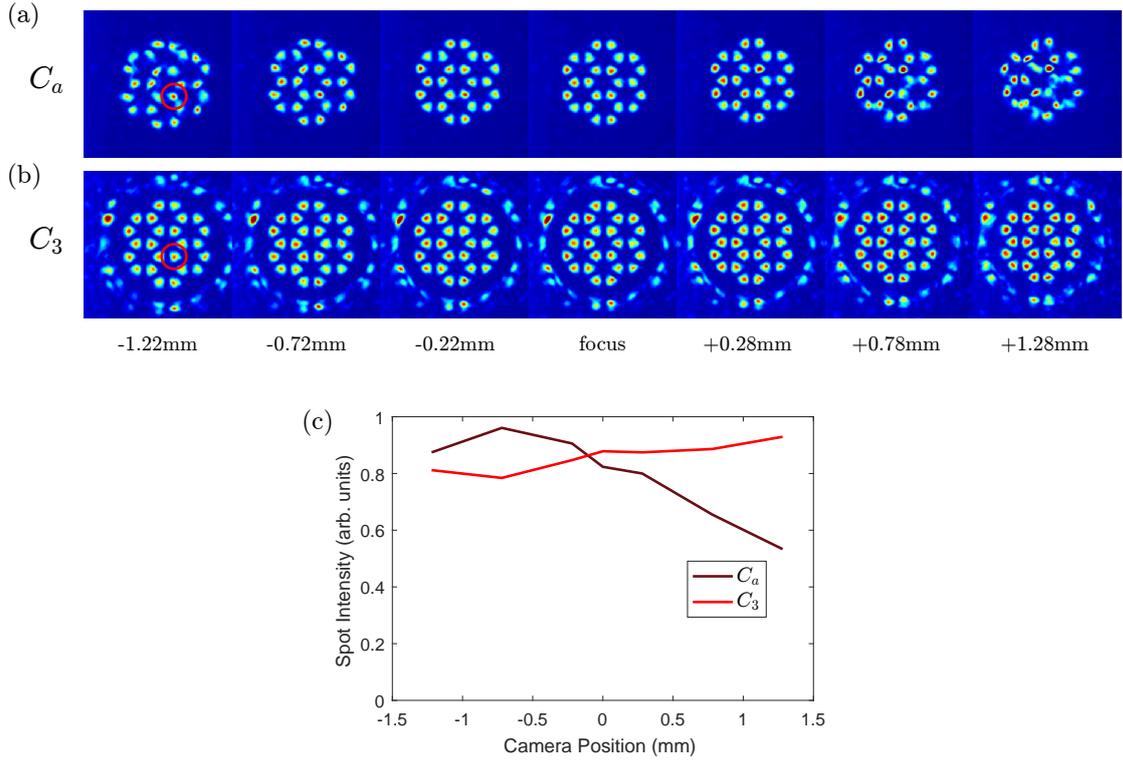


Figure 2.30: Propagation investigation. (a) Graphene pattern without phase control. The light efficiency calculated by C_a was 87.3% and the RMS 19.2%. (b) Graphene pattern with phase control. The light efficiency of the pattern calculated by C_3 was 12.6% and the intensity RMS was 6.6%. Graphene pattern without phase control. The light efficiency calculated by C_a was 87.3% and the RMS 19.2%. The laser power was reduced for the C_a pattern to avoid saturating the camera. (c) Propagating spot intensity fluctuation, where the spot of interest is denoted by a red circle in both (a) and (b).

A general fringe pattern can be written in the form [133]

$$\begin{aligned} g(x, y) &= a(x, y) + b(x, y) \cos(2\pi f_0 x + \varphi(x, y)) \\ &= a(x, y) + c(x, y) \exp(2\pi i f_0 x) + c^*(x, y) \exp(-2\pi i f_0 x), \end{aligned} \quad (2.31)$$

where the $a(x, y)$ and $b(x, y)$ terms contain unwanted variations and $c(x, y)$ is given by

$$c(x, y) = \frac{1}{2} b(x, y) \exp(i\varphi(x, y)). \quad (2.32)$$

Here, $\varphi(x, y)$ is the phase structure of the light pattern which we wish to extract, whilst the f_0 term is the spatial-carrier frequency which is dependent on the relative angle of overlap between the two interfering beams, such that a large angular difference corresponds to a large frequency.

A Fourier transform of $g(x, y)$ with respect to x leads to

$$G(f_x, y) = A(f_x, y) + C(f_x - f_0, y) + C^*(f_x + f_0, y). \quad (2.33)$$

If the spatial-frequency f_0 is large compared to the variations of $a(x, y)$, $b(x, y)$ and $\varphi(x, y)$ then the three spectra in 2.33 will be separated by f_0 and either $C(f_x - f_0, y)$ or

$C^*(f_x + f_0, y)$ can be isolated by filtering out $A(x, y)$. This spectra can then be shifted back to the origin by f_0 and the inverse Fourier transform taken, such that $c(x, y)$ is found and the phase $\varphi(x, y)$ can be extracted. The process of the fringe analysis method for the Laguerre-Gaussian pattern is shown in Figure 2.31. Before Fourier transforming the fringed image, a Hanning window of the form $w(x) = 1 - \cos(2\pi x/D)$ is applied to suppress discontinuities at the edges [133], where D is the size of the image. The filter used to isolate $C(x, y)$ was a simple circular aperture whose size was chosen to encompass the majority of the spectra of interest.

The SLM used in the experiment has more (1920×1152) and smaller $(9.2 \times 9.2\mu\text{m})$

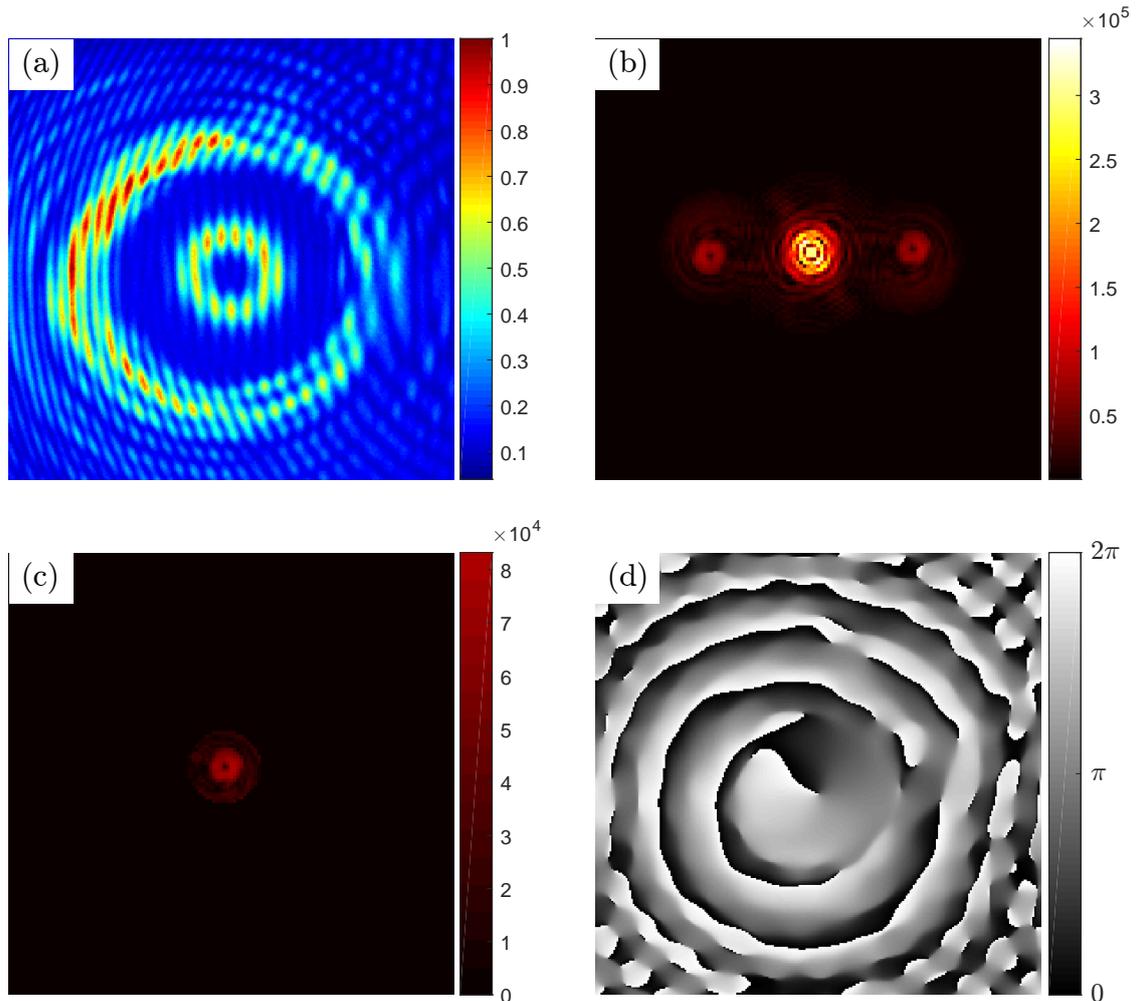


Figure 2.31: Fourier transform fringe analysis method for extracting phase on the Laguerre-Gaussian pattern. (a) Reference beam overlapped with the light pattern, causing interference fringes. The inner ring shape light is the LG pattern, whilst the outer light is the uncontrolled intensity in the noise region. (b) The 2D Fourier transform of the fringe image, where the three separate spectra can be seen (only the real component is plotted). The brightest pixels in the centre have been flattened off for clarity. (c) The spectra of interest, isolated and shifted to the centre (only the real component is plotted). (d) The phase of the inverse Fourier transform of the isolated spectra, corresponding to the phase structure of the light pattern.

pixels than the calculated holograms in Section 2.5.2. In order to minimise calculation time, we chose to keep a small hologram of 256×256 pixels. To do this, we separate the SLM into bins of 3×3 pixels (ultimately using 768×768 pixels in total, where the remaining pixels of the SLM are not used). We calculate new optimal holograms for the effective pixel size of $27 \mu\text{m}$ and an experimentally-convenient beam waist of 3 mm . These are presented in the Theory columns of Table 2.4.

As shown in Figure 2.32b)-d) and detailed in Table 2.4, the measured fidelities are lower than the numerical predictions, but could be improved by the addition of feedback [55, 59] or the characterisation of wavefront aberration in the optical system [99, 100]. For the purpose of comparing the light efficiency with the theoretical result, a rescaled experimental efficiency η^* is also presented in Table 2.4, which is $\eta^* = \eta/0.45$, due to the 45% diffraction efficiency of the SLM. This base efficiency of the SLM was determined by applying a simple blazed grating phase to the device and measuring the amount of light in the first order. Light efficiency of the patterns themselves were measured by placing an appropriate sized pinhole on the field to isolate the signal region and recording the resulting power in comparison to the total input power. In future work, higher diffraction efficiencies could be obtained by replacing the SLM with a micro-fabricated diffractive optical element.

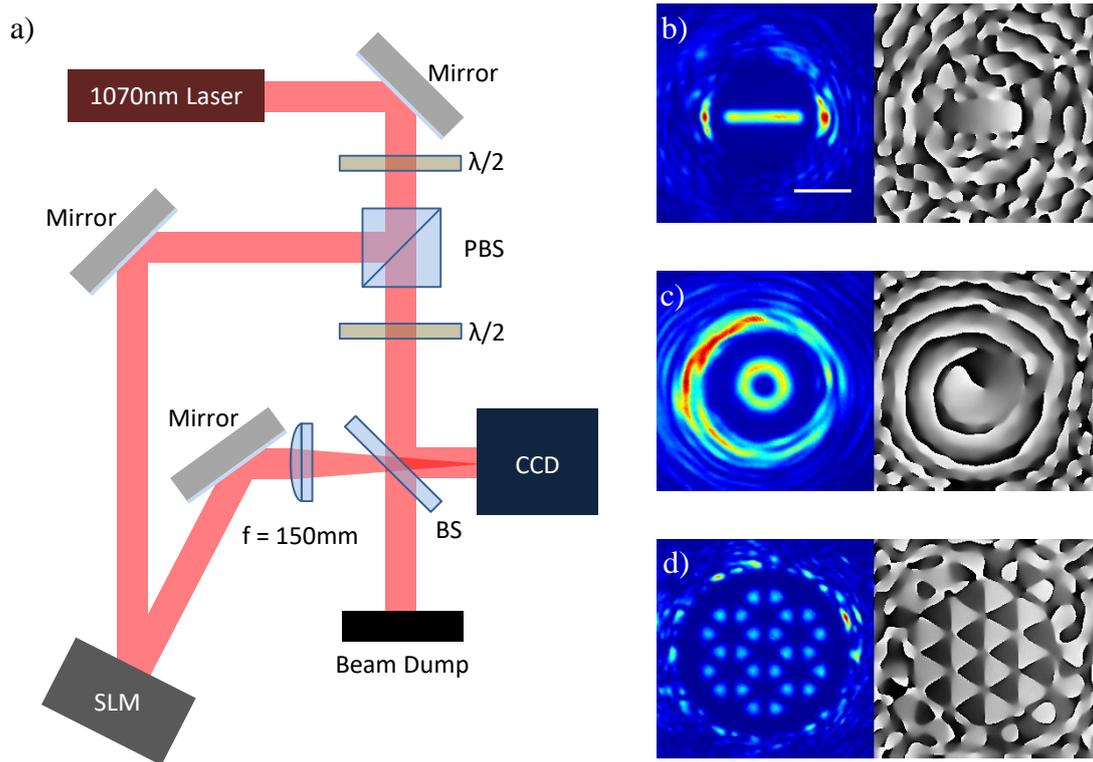


Figure 2.32: a) Experimental Setup. The first $\lambda/2$ waveplate is used to vary the power between the reference and SLM beam, whilst the second waveplate is used to reorientate the polarisation of the reference beam to match the SLM beam for interference. b) - d) Measured intensity (left) and phase (right). The white scalebar in b) denotes $300 \mu\text{m}$, and is common to all images. Color scaling as in Figure 2.25.

Table 2.4: Error metrics for the measured patterns in Figure 2.32.

Pattern	Theory				Experiment			
	$1 - F$	η %	ϵ_Φ %	ϵ_{nu} %	$1 - F$	η^* %	ϵ_Φ %	ϵ_{nu} %
Gaussian Line	7×10^{-5}	8.3	0.005	0.004	3×10^{-2}	7.8	1.85	0.48
Laguerre Gauss	1×10^{-5}	8.4	0.0004	0.004	3×10^{-2}	7.8	2.59	0.52
Graphene	3×10^{-4}	7.0	0.0004	0.015	4×10^{-2}	6.2	2.76	0.42

In summary, we have demonstrated that smooth, high fidelity light patterns with independent control over the amplitude and phase can be generated with a single phase-only SLM. The holograms calculated with the conjugate gradient minimisation approach surpass the accuracy and smoothness of previous IFTA approaches. We note that our approach achieves comparable results in F and η to the super-pixel method for DMDs [127], and improved F for the Laguerre Gaussian mode, at the expense of constraining the pattern to a subset of the output plane. High fidelities were achieved for both far-field and near-field patterns with this method with comparable efficiencies and smoothness. The accurate control over amplitude and phase will be crucial for future research directions in the design of axially-structured light fields.

2.5.5 Multi-Wavelength with Simultaneous Phase and Amplitude Control

The approaches of the multi-wavelength technique from Section 2.4 and the simultaneous phase and amplitude control from Section 2.5 can easily be combined for further flexibility of atomic trapping with an SLM. In order to showcase the compatibility of the two methods, we created a ring and barrier pattern similar to that of Figure 2.20(b), but with a phase winding for the ring. The pattern was calculated for the far-field with an FFT propagator, where the results are shown in Figure 2.33.

The $1/e^2$ half-width of the incident beam σ for the calculation was 1.5mm, whilst the guess phase curvature R was 4.5mrad px^{-2} . The phase was left unconstrained for the barrier and therefore the fidelity was only taken for the ring which had an ROI diameter of 65px. The fidelity gives a result of $1 - F = 3.8 \times 10^{-6}$. The light efficiency η of the pattern was 24.6%, whilst ϵ_Φ was 0.0002 and the non-uniformity ϵ_{nu} was 0.001. These errors are comparable to that of Table 2.2 and exemplify the ease with which a multi-wavelength pattern with a constrained output phase can be created.

The techniques outlined in the previous sections therefore present exciting developments for flexible atom trapping with a single SLM which can be combined where desired. The full realisation of such methods however, requires an experimental setup for ultracold atoms where the SLM traps can be utilised. The next chapter will therefore outline the systems we have built for such a task.

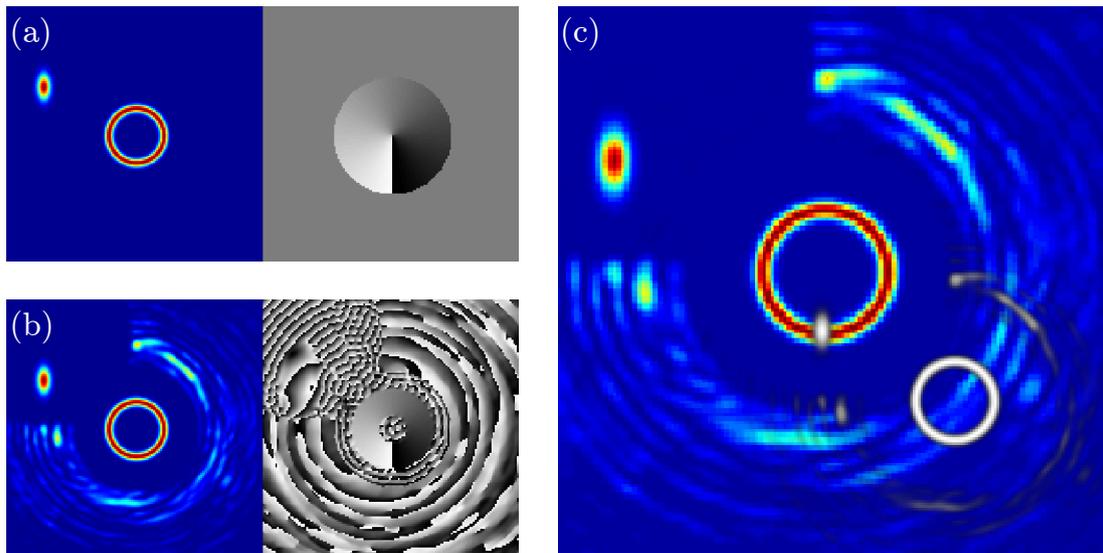


Figure 2.33: Multi-wavelength ring and barrier pattern for 1064 nm and 670 nm light, similar to that of Figure 2.20 b), except there is a phase winding from 0 to 2π now on the ring. The colours denote the same values as in Figure 2.25. (a) Target intensity (colour) and phase (gray), where the phase constraint is only applied to the ring. (b) Resulting intensity (colour) and phase (gray) from the calculation. (c) Simulated intensity if 1064 nm light (colour) and 670 nm light (white) were applied to the SLM. The image was created by resizing the pattern by the ratio 670/1064 and overlaying it on top of the original pattern.

Chapter 3

Cold Atoms Setup

3.1 Overview

In this chapter I will detail the apparatus that we have constructed in order to realise the first BECs in St Andrews. The process of producing an atomic BEC is well-covered in the literature (reviews can be found in [134, 24, 10]). Briefly, atoms are introduced to an Ultra High Vacuum (UHV) chamber, and subjected to initial laser cooling in a magneto-optical trap (MOT). These pre-cooled atoms are neither sufficiently cold nor dense enough to undergo Bose-Einstein condensation. Thus a second phase of evaporative cooling in a conservative trap must be used. The atomic species of interest in our experiment is ^{87}Rb .

Regardless of what trap is employed, background pressure will negatively effect the trapping lifetime of the atoms as collisions with hot background atoms will cause losses. Ultimately, if the pressure is too high we can not condense the atoms as they will leave the trap too quickly for our cooling process. A pressure of the order of 10^{-11} mbar is needed to have a sufficiently long lifetime ($\sim 100\text{s}$) for typical BEC experiments [135].

There are essentially two main vacuum chamber designs for cold atomic experiments: the single chamber setup and the double chamber setup. The single chamber setup has both the BEC experiment and the atomic species loading done in the same environment, giving a compact and simple design. However, the atom loading process contributes considerably to the background pressure, resulting in worse lifetimes. Alternatively the double chamber setup has a first, high pressure “loading chamber” from which the atoms are transferred to a second, lower pressure “science” chamber for the BEC experiment.

The main goal of our experiment is the flexible trapping of BECs using holographic patterns. In a previous incarnation of our experiment we had a single chamber design, which only achieved moderate lifetimes and atom numbers during the MOT stage (6 s and 1×10^8 respectively). We therefore opted for a double chamber design for our new experiment.

Loading atoms into the science chamber is the initial step of the experimental sequence. We accomplish this by using a 2D magneto-optical trap (MOT) [136, 137, 138]. The operating principles of the MOT are comprehensively explained in [24, 139]. In brief, laser beams which are slightly detuned below atomic resonance are shone from multiple directions. Photons which are moving counter to the atomic motion (Doppler shifted higher in frequency) become more favourably absorbed, whilst the spontaneously re-emitted photons are isotropically directed, resulting in a damping force which opposes the atomic motion. A position-dependent force can be added by using a linear magnetic field gradient and circularly polarised beams: the combination of the Zeeman effect and selection rules ensure that atoms are pushed towards a zero-point of the magnetic field.

If pairs of counter-propagating beams are used along three orthogonal directions with a single pair of coils in anti-Helmholtz configuration, the atoms will be cooled and trapped at a point: this configuration is the standard three-dimensional (3D) MOT. However, using counter-propagating beams in only two directions and a spatially varying magnetic field with a zero-field line orthogonal to the beams realises a two-dimensional MOT (2D MOT). This results in an atomic beam where the atoms are free to propagate along the zero-field line, but are confined and cooled in the transverse axes. Such an atomic beam is ideal for transferring a high flux of atoms from the “loading” chamber in to the low-pressure “science chamber”, where we then trap and cool them in a fully confining 3D MOT.

After laser cooling in the 3D MOT, the magnetic field gradient can be increased and the laser beams turned off such that the atoms are kept in a purely magnetic trap. At this point the atoms are evaporatively cooled. This process selects the higher energy atoms in the trap and removes them, the same principle as blowing on the steam from a coffee cup, such that the remaining atoms then thermalise to a colder temperature. We remove the hottest atoms with a sweeping RF knife, which transfers the trapped atoms to an anti-trap state, expelling them out of the ensemble. This cooling technique is naturally a lossy process (and it is necessary to trap as many atoms as possible prior to this stage), but further losses are also taken in our quadrupole configuration magnetic trap due to Majorana spin flips. The coldest atoms near the center of the magnetic trap experience such small magnetic fields that the Zeeman sub levels are near-degenerate, leading to atoms spin-flipping in to an anti-trap state and leaving the system. To minimise such losses we transfer the atoms to a hybrid trap where the final evaporation is performed, following a similar approach outlined in [140].

The hybrid trap consists of a gravity-compensating magnetic trap and single focused-beam optical dipole trap (ODT). The ODT light is sufficiently far red-detuned that photon scattering is negligible, but powerful enough that the AC Stark effect is prominent. A shift in the energy levels of the atoms will then occur, such that the ground state energy lowers in the presence of the light, creating a trapping profile in the shape of the beam. The gravity-compensating magnetic field is created by the same quadrupole field as the magnetic trap (with reduced magnitude), and provides levitation as well as confinement along the ODT axis. Further evaporative cooling can be done during this stage by simply lowering the laser intensity to reduce the trap depth and expel the hottest atoms. With enough cooling, the BEC in our experiment is created. Furthermore, at this stage where the atoms have been cooled sufficiently, a separate, holographic ODT is employed in a crossed configuration with the first ODT to create arbitrary trapping configurations for the atoms.

For diagnostics of the atoms we use absorption imaging, where a pulsed resonant laser beam (“probe beam”) illuminates the cloud and is imaged on to a camera. The atoms absorb the light and re-emit the photons isotropically, resulting in a shadow being cast on the camera. The depth of the shadow corresponds to the density of the atoms.

In the rest of this chapter, the details and implementation of the above experimental scheme are covered.

3.2 Double Chamber Vacuum Setup

A computer model of the double chamber setup is shown in Figure 3.1. Our loading chamber is made up of a cylindrical glass cell (LewVac GA-DE33L-40CF), four rubidium metal dispensers (SAES) and a 40 Ls^{-1} Ion Pump (Varian VacIon StarCell Plus 40). A flexible bellow, spherical octagon (Kimball Physics MCF600-SO200800), titanium subli-

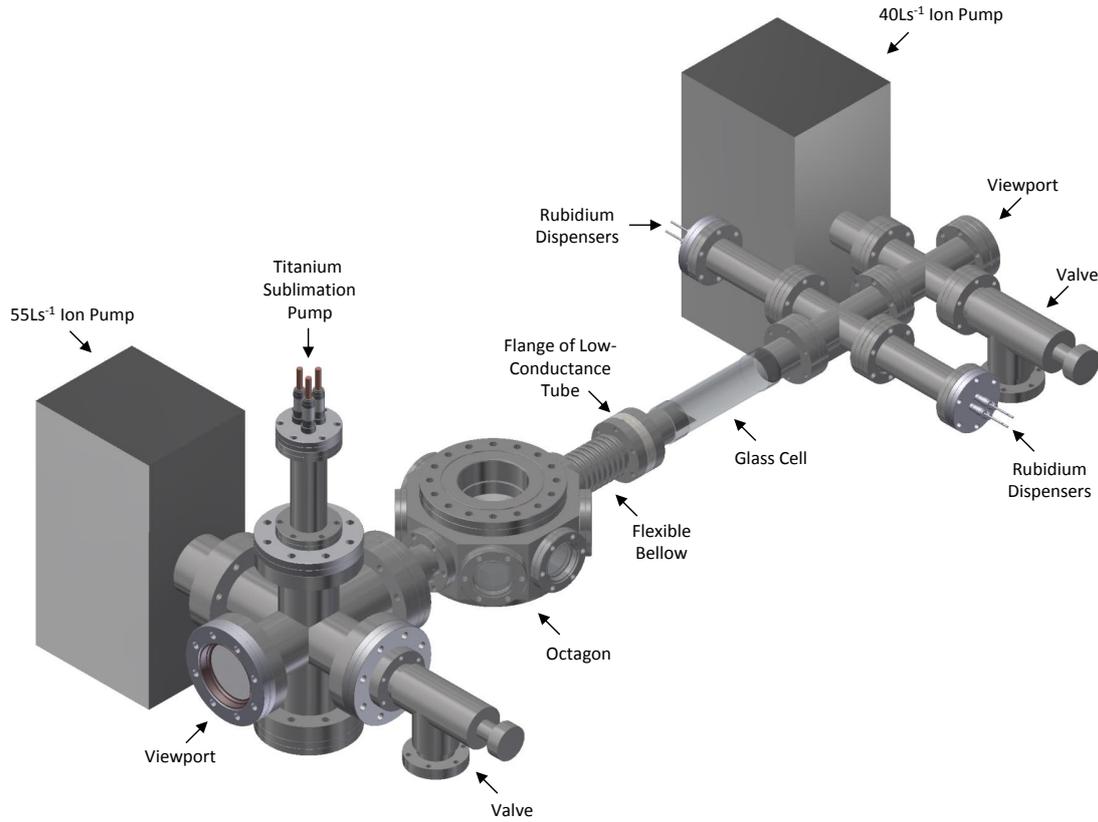


Figure 3.1: Our double chamber vacuum setup. The flexible bellow can be considered as the joining element between the science chamber (octagon side) and loading chamber (glass cell side). For reference, the full length of the glass cell, including flanges, is 203.2mm.

mation pump (VG Scienta ST22) and a 55 Ls^{-1} ion pump (Varian VacIon StarCell Plus 55) form the science chamber. The two chambers are separated by a LewVac custom-built low-conductance tube placed between the glass cell and bellow. Two all-metal angle valves are also present, allowing a turbo pump (Pfeiffer HiPace 80) to be connected to both chambers during the bakeout process (otherwise one of the chambers would need to be pumped through the low-conductance tube).

In order to obtain UHV conditions, we cleaned all components with acetone and isopropanol before baking the assembled chamber at 200°C for two weeks. After bakeout and with both ion pumps and titanium sublimation operational, we achieved a base pressure in the octagon below the 10^{-9} mbar resolution of our ion pump current read-out (see Figure 3.2). With rubidium dispensed in to the loading chamber such that a pressure of 10^{-7} mbar is created, a trap lifetime of $> 60\text{s}$ is accomplished in the science chamber.

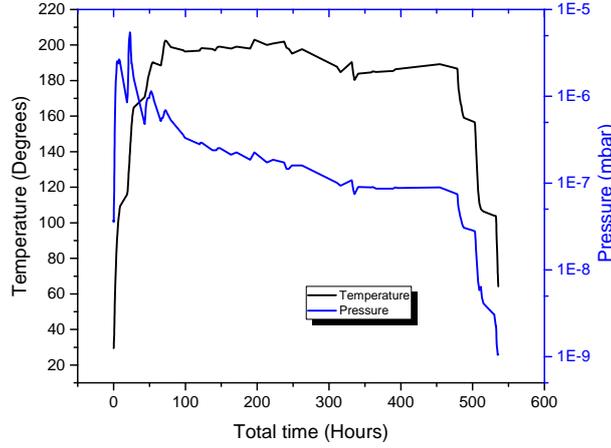


Figure 3.2: Temperature of the octagon and pressure of the chamber as measured by our ion gauge during the bakeout.

3.2.1 Loading Chamber

Our loading chamber houses four dispensers that release rubidium into the 2D MOT where they are cooled into an atomic beam. The atomic beam is then used for loading the 3D MOT in the science chamber. In order to ensure that the pressure of the science chamber is unaffected as the rubidium is dispensed, we use the low-conductance tube to separate the two chambers.

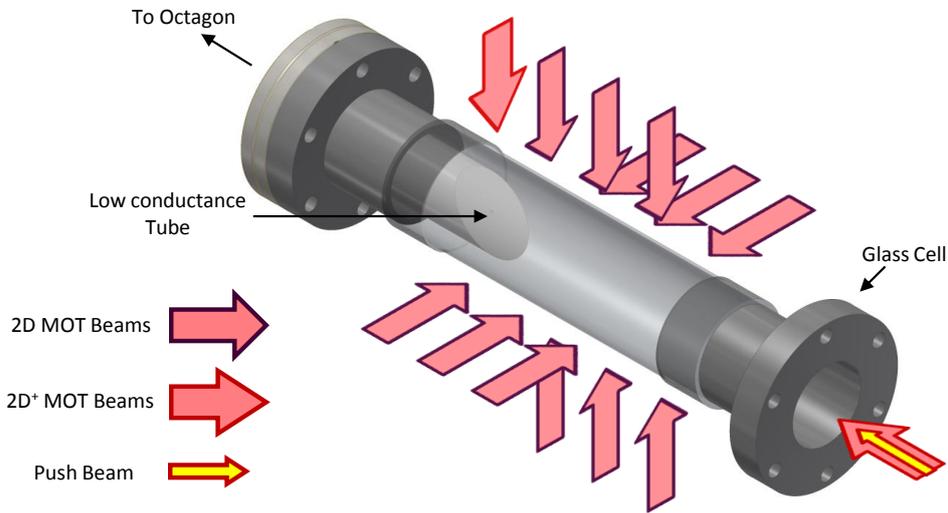


Figure 3.3: A diagram of the glass cell and low conductance tube along with the 2D MOT beams. Six 2D MOT beams are retro-reflected back in to the cell. A longitudinal cooling beam ($2D^+$ beam) reflects off of the 45 degree polished surface of the conductance tube, and is paired with another counter-propagating cooling beam and a push beam which pass through a viewport at the end of the chamber.

Gas flowing through a pipe can produce a pressure differential at the ends of the piping [141]. If we assume the pressure difference between the two chambers is large, the pressure ratio achieved between the ends of the piping can be approximated as

$$\frac{p_{\text{science}}}{p_{\text{loading}}} = \frac{C}{S}, \quad (3.1)$$

where p_{science} is the pressure in the science chamber, p_{loading} is the pressure in the loading chamber, C is the conductance and S is the science chamber pumping speed. For molecular flow conditions, the conductance of air at 20°C in a pipe is given by $C = 12.1d^3/l$, where d is the diameter of the pipe in cm, l is the length of the pipe in cm and C is given in units of Ls^{-1} . To minimise the conductance and therefore maximise the pressure ratio between the two chambers, a small diameter tube with a large length is required. However, there is a limit to the size requirements of the tube set by the atomic beam. Typical atomic beams emerging from a 2D MOT have divergence angles ranging from 10 - 45 mrad [136, 137, 138]. We therefore chose our low conductance tube to have a diameter of 1.2 mm and a length of 12.7 mm, which should allow for an atomic beam with a divergence angle of 47 mrad to pass through. If we take the pumping speed of the science chamber to be the 55 Ls^{-1} of our ion pump, then we get a pressure ratio of $p_{\text{science}} \sim 10^{-4}p_{\text{loading}}$. This lets us reach up to $\sim 10^{-6}$ mbar with the rubidium dispensers in the loading chamber and still maintain good UHV conditions in the science chamber. The flux of 2D MOT atomic beams tends towards a maximum at rubidium vapour pressures of 1×10^{-6} mbar [136, 137], thus our vacuum setup should allow for an optimal atomic beam in this regard. The 3D MOT in our setup benefits from this as a high atomic flux leads to larger and faster MOT loading.

Since a larger cooling volume leads to a stronger atomic beam [137], we chose a cylindrical glass cell for large optical access for our 2D MOT. The glass segment is 152 mm long, although the region of ideal optical access is ~ 110 mm due to distortions of the glass at the ends of the cylinder and the Kovar[®] sleeves. The inner diameter of the glass cell is 33 mm. Large rectangular optics and mirrors could be used to capitalise on this optical access as much as possible, however, such large optics are often expensive and can be quite impractical. We therefore used standard 25.4 mm diameter optics, but arranged them with minimal separation so as to fill the glass cell with three cooling regions approximately 4 mm apart, giving us an effective cooling volume of $\sim 75 \times 25 \times 25$ mm. The three cooling beams are retro-reflected after passing through the cell (see Figure 3.3). Without anti-reflection coating on the glass cell, the power of the retro-reflected beams is 17% lower at the position of the atoms. The cylindrical glass cell also acts as a cylindrical lens with focal length -350 mm. We counteract these two effects with an $f = 300$ mm cylindrical lens placed behind the glass cell to collimate and increase the peak intensity of the retro-reflected beams.

To incorporate longitudinal cooling (the 2D⁺ MOT configuration [136]) and aid our atomic beam further, we have our low conductance tube also act as a mirror. The tube is made from stainless steel and has a 25.4mm outer diameter (and a small 1.2mm inner diameter). The end has a 45 degree polished cut, allowing a beam from above the chamber to be reflected with 60% efficiency along the atomic beam axis (although with a small hole in the centre). A viewport at the end of the loading chamber allows a counter-propagating beam to pass through, forming the cooling pair.

Since the atomic beam has no preferred direction in which to propagate, a “push” beam is used to introduce a radiation pressure imbalance in the 2D MOT which aids the atomic beam propagation through the low conductance tube and in to the science chamber. A diagram of the scheme is shown in Figure 3.3.

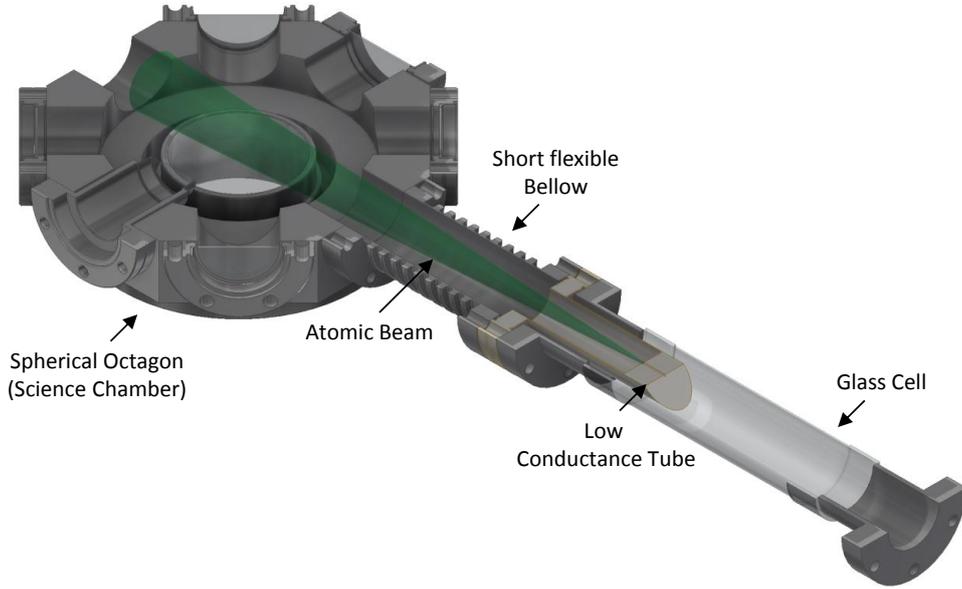


Figure 3.4: A cut through of the glass cell, low conductance tube, bellow and octagon. The atomic beam emerges from the low conductance tube and enters the science chamber.

The linear magnetic field gradients and zero-field line required for the 2D MOT are created by two pairs of rectangular coils in anti-Helmholtz configuration. The coils are approximately 196 mm long and 50 mm wide, where the coil separation is 82 mm. A single coil has 96 turns of 1 mm diameter copper wire, hand-wound around nylon formers (to avoid circulating currents). This geometry gives a field gradient of $6 \text{ Gcm}^{-1}\text{A}^{-1}$ along the two transverse axes of the atomic beam. The operating current used for the coils during the experiment was 2.5 A. As the coils were only used at low currents they did not require water cooling. Background magnetic fields were compensated for by two pairs of shim coils in Helmholtz configuration which were wrapped around the 2D MOT coils. Each shim coil has 20 turns with 0.8 mm diameter wire, creating a field of 2.5 GA^{-1} .

In order to ensure that we could optimally align the atomic beam onto the 3D MOT, we designed the entire 2D MOT and loading chamber to be position-adjustable. The coils and vacuum chamber were mounted on the optics breadboards. The entire assembly was then attached via height-adjustable posts to two “skis”, one for the ion pump, the other for the rest of the chamber. The skis are aluminium plates with a large footprint, allowing them to be slid across the table. The short flexible bellow between the octagon and low conductance tube meant that the assembly could be moved in the two transverse axes of the atomic beam propagation, allowing us to align the atomic beam to the 3D MOT. The maximum distance (short bellow uncompressed) between the 3D MOT and the opening of the low conductance tube on the loading chamber side is 280 mm. Over this distance, an atomic beam with a divergence angle of 47 mrad will spread out to approximately 26 mm. Our 3D MOT beams are 20 mm in diameter, leading to reasonably good overlap. Figure 3.4 shows a diagram of the atomic beam entering the science chamber, whilst the general apparatus of the adjustable loading chamber is shown in Figure 3.5.

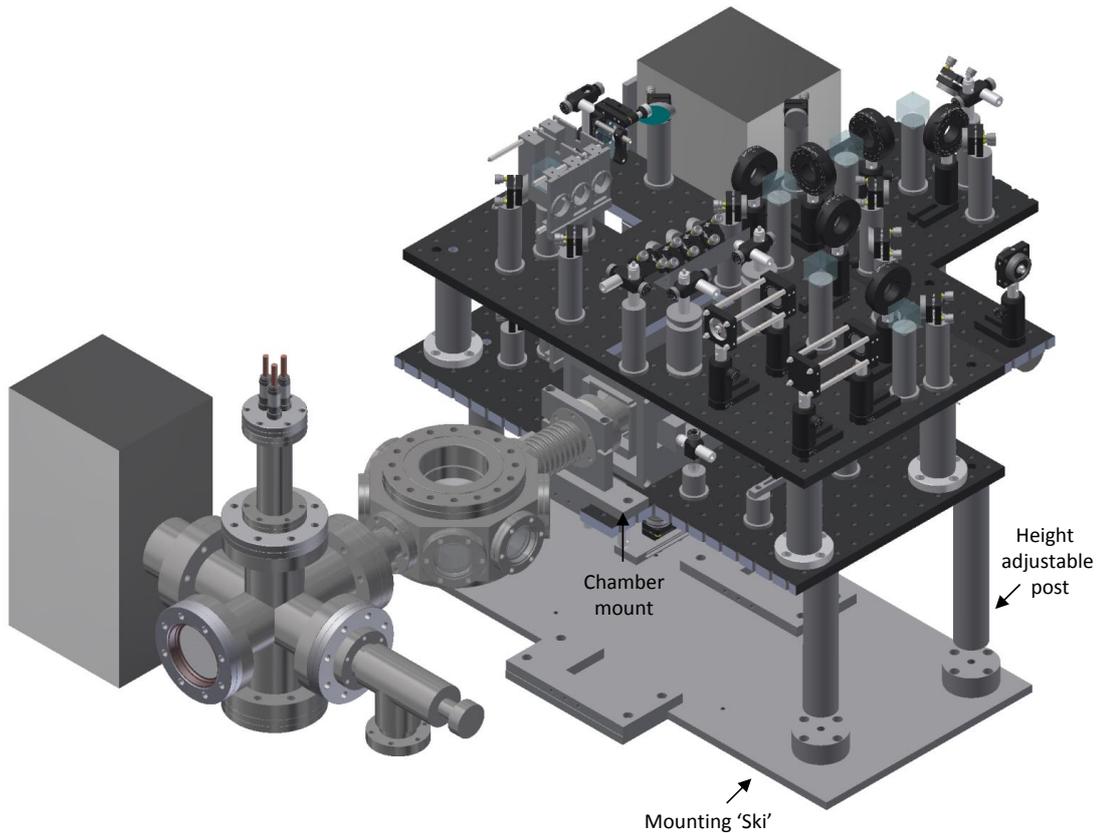


Figure 3.5: The vacuum chamber with the 2D MOT optics in place. The lower mounts for the optics board are height adjustable. The entire optics configuration and loading chamber can be moved as one piece.

3.2.2 Science Chamber

The science chamber is where the BEC formation and main experiment takes place. We ensure that good vacuum conditions are kept in this chamber by using the larger of our two ion pumps and a titanium sublimation pump. We also use a spherical octagon to house the atoms in, instead of a glass cell which typically has a lower conductance and can impede pumping performance. However, the comparatively larger volume of the octagon as opposed to typical glass cells means that optical access is more restricted and it can be harder to place apparatus in close proximity to the atoms. To overcome this we have made use of recessed viewports.

There are a total of three recessed viewports fitted to the octagon. The vertical viewports are recessed for both the RF coils and magnetic coils to be placed closer to the atoms, allowing for stronger fields to be generated at more manageable currents. The recession also allows for a shorter focal length lens to be fitted. For a lens of a fixed diameter, the shorter the focal length the larger the numerical aperture (NA) and the higher the image resolution. This allows for more accurate diagnosis of the atomic cloud as well as finer detailed SLM patterns for our phase-engineered light traps. The third recessed viewport is along the horizontal axis and is also fitted for image resolution purposes (the NA allowed for in the vertical recessed viewport is 0.32, whilst the NA in

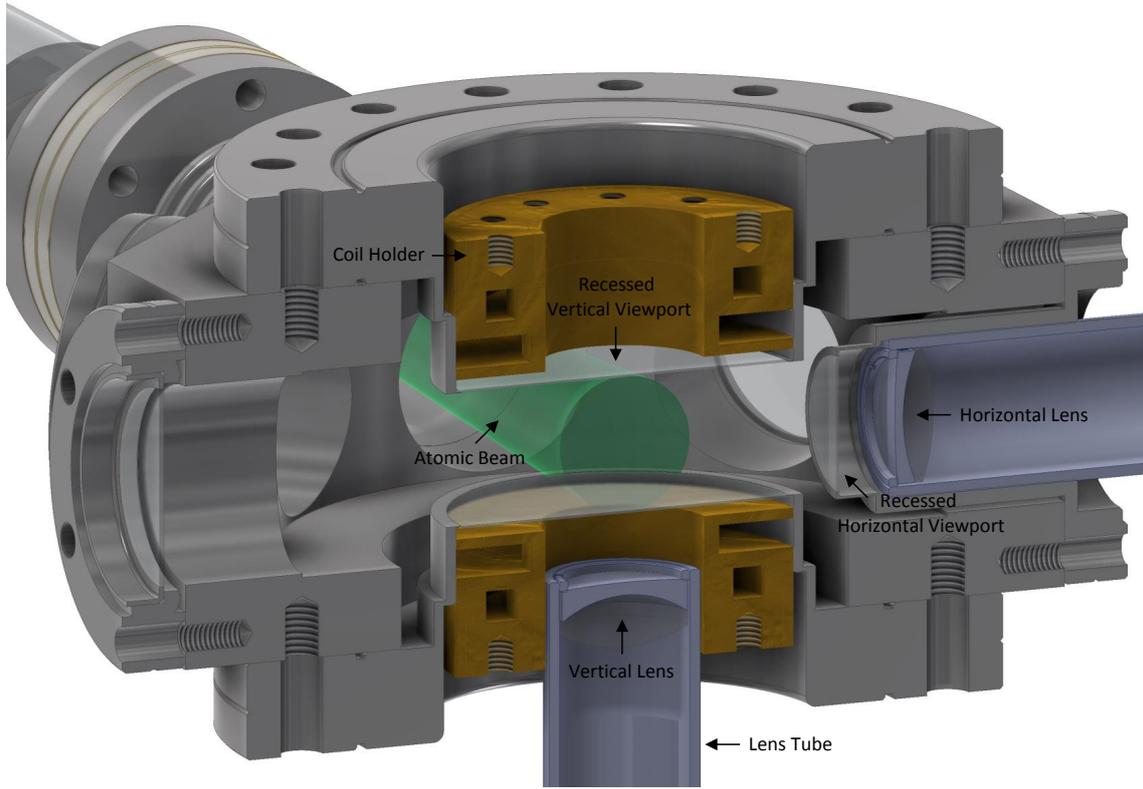


Figure 3.6: A cut of the octagon showing the recessed viewports and the close proximity achieved for the coil holders and imaging/optical trap lenses. For reference the lens diameters are 25.4 mm.

the horizontal is 0.21). The recessed viewports with lenses and magnetic coil holders are shown in Figure 3.6.

The five other viewports fitted to the octagon are standard and are in place for the absorption imaging and 3D MOT beams. A final viewport on the science chamber is at the end of a six-way cross, to which the pumps are fitted. This viewport is used both for the optical pumping beam and for alignment of the push beam from the loading chamber. The vertical recessed viewports are anti-reflection coated such that they have $< 0.6\%$ reflectivity for both 780 nm and 1064 nm light. The rest of the viewports in our experiment are not anti-reflection coated and for these we lose approximately 4% laser power per glass surface. The MOT beams and imaging beams are ≈ 20 mm in diameter and are shown (amongst the other beams) in Figure 3.7.

The magnetic coils for the science chamber are placed as close as possible to the atoms in the recessed viewports. The minimum separation between the coils is 52 mm, where the coils have 48 turns each (4 layers of 12 turns) of 1 mm diameter wire. The mean radius of the coils is 24.4 mm. With this design, a field gradient of 100 Gcm^{-1} is generated with 20 A, suitable for magnetic trapping. This high current however generates a lot of heat, potentially causing damage to the wire. We therefore water cool our coils, where the coil formers have a 6-by-6 mm channel within them which allows 15°C water to be pumped through by our chiller (Kühlmobile, Van der Heijden). The formers are made of brass for good thermal conductivity and have a thin slit cut through them so that they do not form a closed ring which may allow for circulating currents. With this design the coils can be

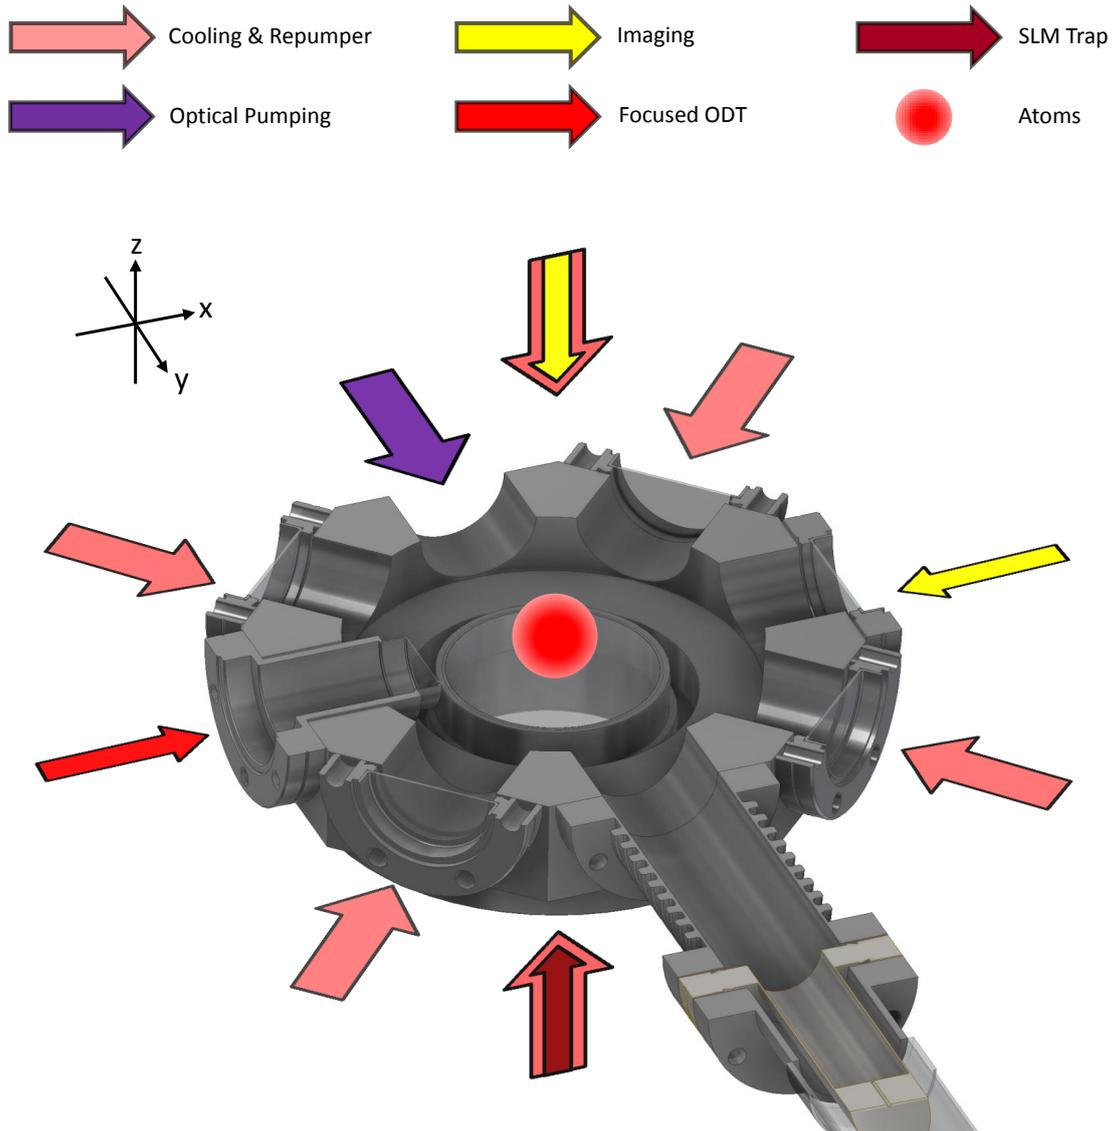


Figure 3.7: Cut of the octagon with laser beams.

operated at 20 A for tens of seconds without their temperature raising above 50° C. A pair of shim coils in Helmholtz configuration in each axis is used to offset the background fields around the octagon. The vertical axis shim coils have ~ 50 turns, a diameter of 160 mm and are separated by 120 mm, creating a field of 1.3 GA^{-1} . The two other axes have square shim coils with 90 turns which have a side length of 110 mm. They are separated by 180 mm, generating a field of 3 GA^{-1} .

3.3 Optical Setup

Careful manipulation of lasers for cooling, trapping, atomic state pumping and imaging are needed for the experiment. The majority of the lasers and optical components which are necessary for this task are placed on a separate work bench from the vacuum system: the optics table. The fundamentals of the laser setup of the optics table is shown in a flow

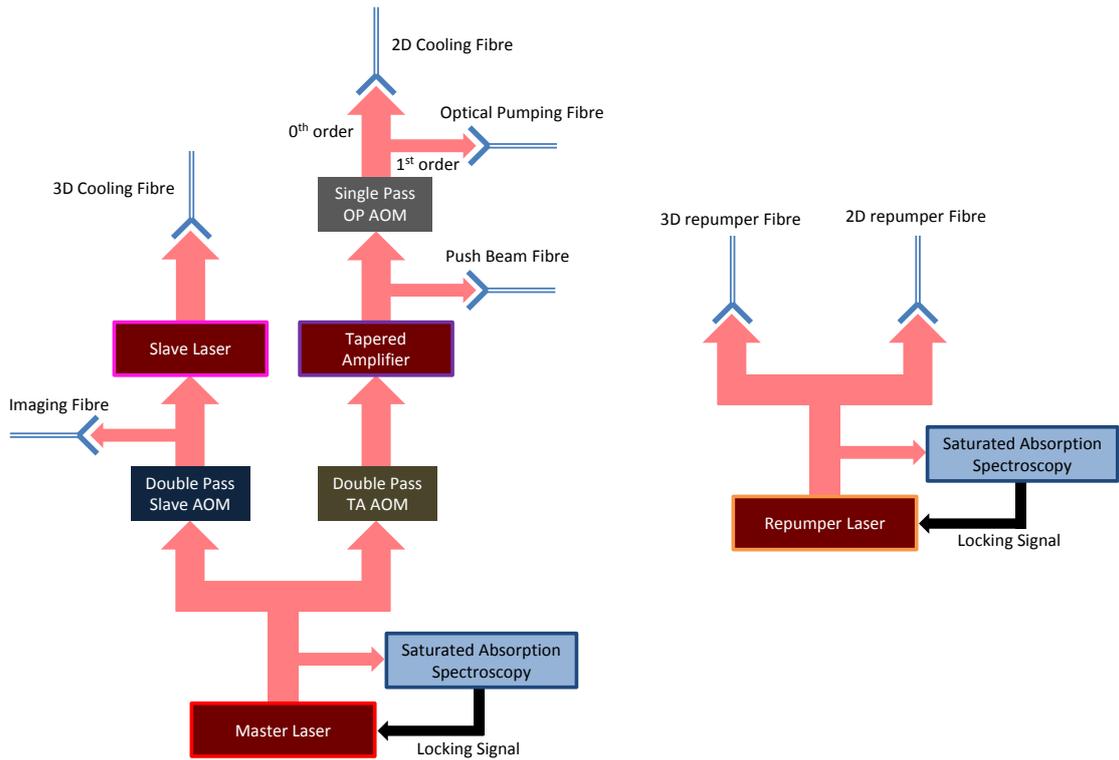


Figure 3.8: Flow diagram for the optics table setup. All fibres also have a mechanical shutter in front of them.

diagram in Figure 3.8.

We have two External Cavity Diode Lasers (ECDLs): the master laser (Toptica DL100 mount with Toptica LD-0785-P220-1 diode) for the laser cooling, and the repumper laser (home-made housing, Sanyo DL7140-201S) which ensures that atoms are kept in the cooling transition. Our master laser operates at a power of 50 mW, whilst the repumper generates 30 mW. Both lasers are frequency locked via Doppler-free saturated absorption spectroscopy of rubidium in vapour cells (the lock points are detailed in section 3.3.1 below). The master seeds a tapered amplifier (TA, Eagleyard EYP-TPA-0780-01000-3006-CMT03-0000) which amplifies the laser power to 900 mW from an 8 mW injection. The master also injects a slave laser (home-made temperature-stabilised mount, Thorlabs LD785-SE400 high-power laser diode) with 8 mW, where the slave generates 300 mW. All lasers are temperature stabilised by thermistors and thermoelectric elements, whilst the TA has additional cooling from a fan. All above lasers also have a 60 dB optical isolator to avoid back reflections disturbing the laser mode. One final laser is used to optically trap the atoms (see section 3.7) and is situated on the vacuum table. This is a 1064 nm fibre laser (IPG Photonics, YLM-5-1070-LP) with a maximum output of 5 W.

A total of three Crystal Technologies Acousto-Optic Modulators (AOMs: Slave AOM model 3080-125, TA AOM model 3110-120, OP AOM model 3200-121) are on the optics table. The slave AOM controls the detuning of the 3D MOT cooling beams and the imaging light, and acts as a fast switch for the latter. The TA AOM only controls the frequency of the 2D MOT cooling beams (and push), whilst the OP AOM is used for frequency shifting and switching of the optical pumping light. Polarisation maintaining single mode optical fibres (seven in total) transport the various beams to the vacuum

table where they are applied to the atoms. In front of each fibre is a home-made shutter to block light appropriately during the experimental sequence. Beam splitting is done via polarising beam splitter cubes (PBS) and rotatable half waveplates. Circular polarisations are set using quarter waveplates for all beams incident on the atoms except the SLM trap, focused ODT and horizontal imaging which are linearly polarised.

3.3.1 Laser Scheme

Laser cooling requires a closed cycling transition such that the excited state of the atoms can only decay back in to the state that it was excited from. As mentioned, the atomic species of interest in our experiment is ^{87}Rb , where the closed cooling cycle occurs on the D_2 line [142], particularly on the transition between the hyperfine ground state $F = 2$ and the excited state $F' = 3$. The D_2 line along with our laser and frequency addressing scheme is shown in Figure 3.9.

The D_2 line for ^{87}Rb occurs at 780 nm, a convenient wavelength which is readily accessible by semiconductor diode lasers. However, the typical linewidth of these diode lasers is usually much larger than the natural linewidth (6.1 MHz) of the ^{87}Rb D_2 line. A broad linewidth laser would lead to prominent excitations to undesired hyperfine transitions. For this reason, diode lasers used for laser cooling are typically ECDLs which have an external cavity, capable of producing linewidths of less than 1 MHz [143, 144]. The external cavity can also be attached to a piezo-electric element, such that a voltage applied to the piezo-electric device will change the cavity length, thereby changing the laser frequency. A voltage ramp will therefore create a frequency scan, suitable for probing the rubidium transitions during spectroscopy.

ECDLs are sensitive devices where slight mechanical noise or thermal drift can change the frequency or mode of the laser. Frequency stabilising is therefore needed to keep the lasers resonant with the atoms. Doppler-free absorption spectroscopy (described in [24]) and lock-in amplifiers are used to frequency lock our ECDL lasers to the correct rubidium transitions. The spectroscopy technique is relatively easy and simple to implement using a frequency scanning laser as it only requires an atomic gas, a photo-diode and two counter propagating beams; a strong pump beam and a weak probe beam. The probe beam is shone on the atoms such that absorption occurs at the rubidium transition, where the signal is picked up on a photo-diode. Due to the motion of the atoms and the Doppler effect, atoms in the rubidium vapour cell show a Doppler broadened spectroscopy profile, unsuitable for frequency locking. To overcome this, a pump beam is also used on the atoms which has the same frequency as the probe. The two beams are counter-propagating, resulting in the probe and pump beam interacting with different atoms if they are Doppler shifted. However, for atoms that aren't shifted (i.e. $v_{atom} \approx 0$), then many of them will be excited out of the ground state by the pump beam, reducing the absorption of the probe beam. This leads to narrow peaks within the Doppler broadened signal at the frequency of the hyperfine transitions, resulting in a Doppler-free spectroscopy signal. Furthermore, if there are two hyperfine transitions within the Doppler broadened signal, an additional peak can be observed halfway between them, known as a cross-over. Here, atoms which are Doppler shifted such that they interact with the pump beam for one transition will also interact with the probe beam for the other transition, resulting in a peak.

The peaks in the Doppler-free spectroscopy signal (shown in Figure 3.10) can be used to frequency lock the laser to the desired transition by using lock-in amplifiers. A small modulating dither is applied to the laser current which allows a differential to be created. Since the top of the peaks will have a zero value for the differential, and either side will be positive or negative, the differential can be used as an error signal as part of a Proportional

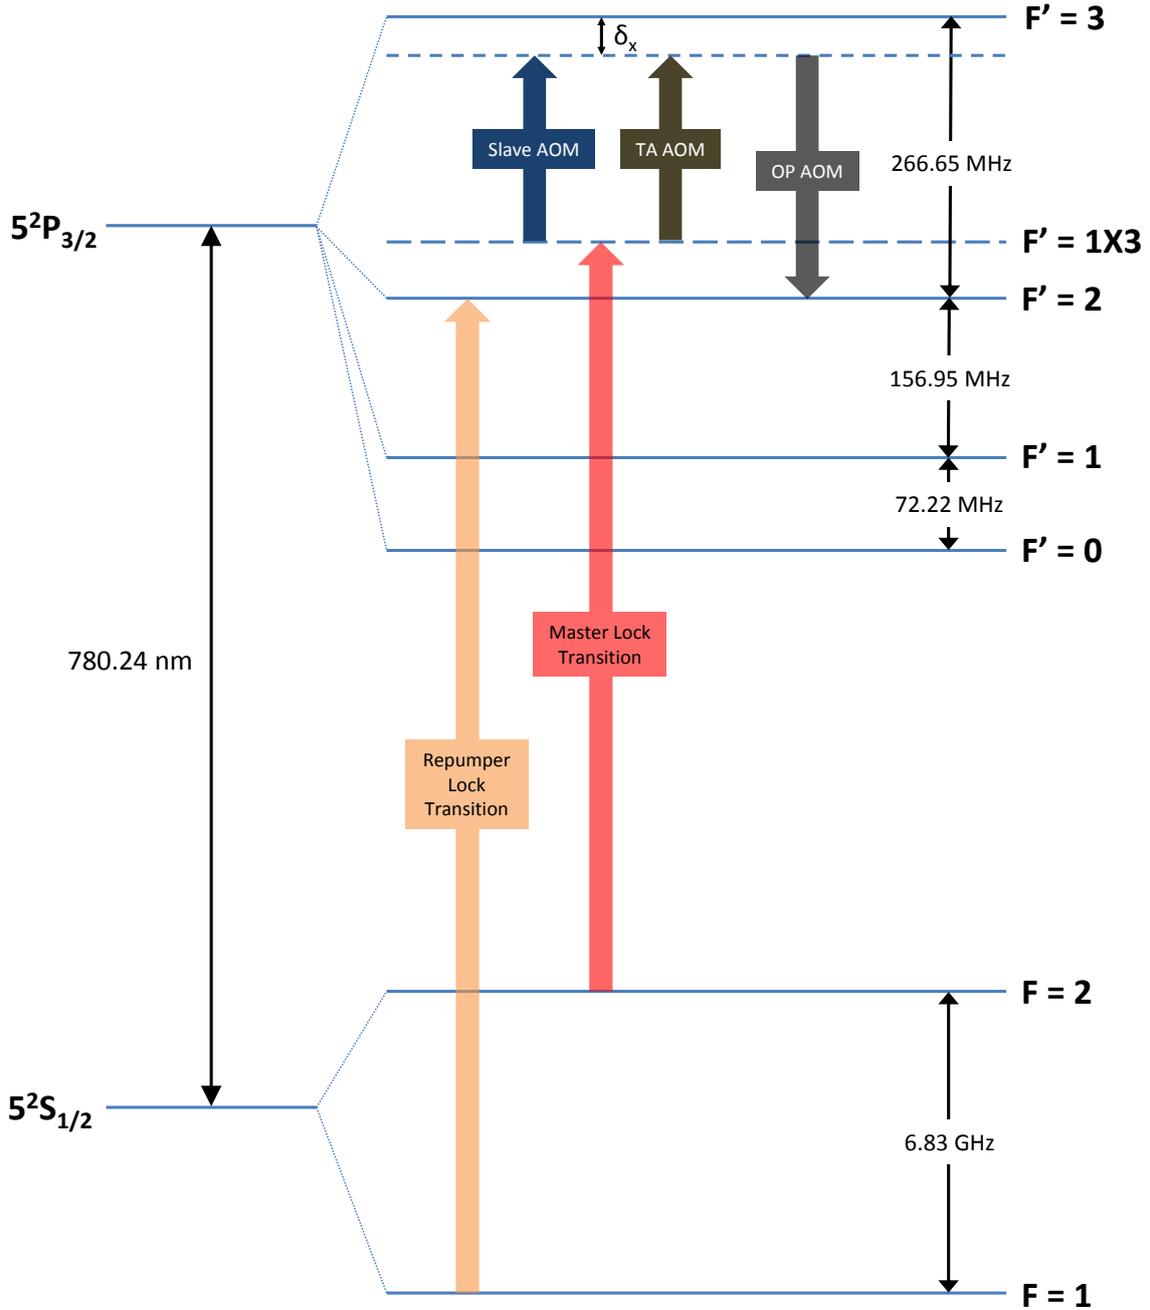


Figure 3.9: The atomic energy levels and relevant transitions. The lockpoint for the master laser is the $F' = 1 \times 3$ crossover. The detunings δ_x take various values during the experiment. The detunings incorporated in to the 3D AOM are $\delta_{3D} = -20\text{MHz}$ for the MOT stage, $\delta_{\text{Img}} = 0\text{MHz}$ for imaging, or $\delta_{\text{Mol}} = -60\text{MHz}$ for optical molasses. The 2D AOM is set such that $\delta_{2D} = -15\text{MHz}$.

Integrable Derivative (PID) regulator. The laser can then be frequency stabilised on the peak, where the laser current and piezo are used to keep the frequency on the transition against thermal drift or mechanical noise.

We do not lock our master laser to the cooling transition itself as a slight detuning of $\sim -15\text{MHz}$ from resonance is required for optimal MOTs, where such a small detuning

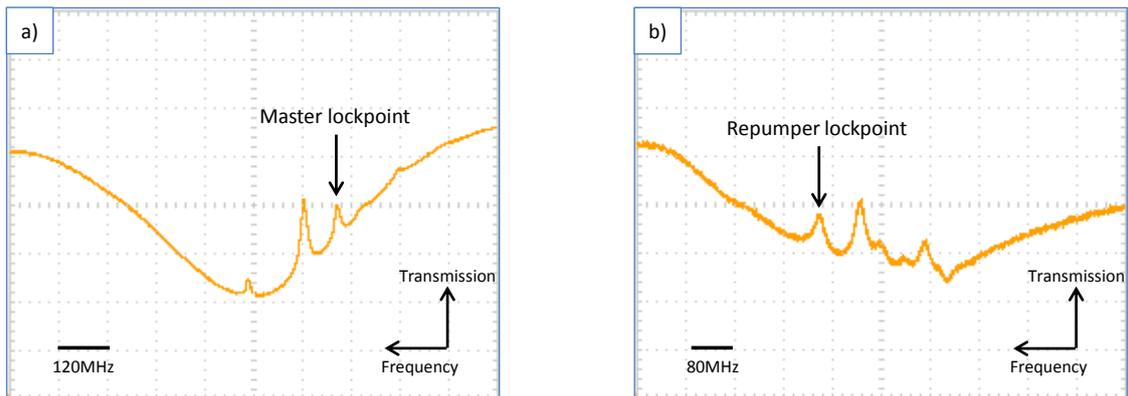


Figure 3.10: The Doppler-free Saturated Absorption Spectrum for the a) $F = 2 \rightarrow F' = X$ and b) $F = 1 \rightarrow F' = X$ transitions of the ^{87}Rb D_2 line from our master and repumper laser respectively. The narrow peaks within the Doppler broadened absorption profiles correspond to hyperfine transitions and cross-overs.

is not easily obtained by readily available devices. Frequency shifts on light of $> 30\text{MHz}$ can be achieved by Acousto-Optic Modulators (AOMs) which use sound waves to create a moving diffraction grating, where the diffracted orders are Doppler shifted by the frequency of the acoustic wave. We therefore lock our master laser 212MHz below the cooling transition at the cross-over between $F = 2 \rightarrow F' = 1$ and $F = 2 \rightarrow F' = 3$ ($F' = 1 \times 3$) and bring the frequency closer to resonance with AOMs (see Figure 3.9).

The Slave AOM has a nominal operating frequency of 80MHz with a bandwidth of 25MHz . We use this AOM in a double pass cat-eye configuration which results in the beam path becoming insensitive to angular shifts that occur when the AOM changes operating frequency. For the 3D MOT, the AOM is set to operate at 96MHz (192MHz in double pass), giving us a detuning from resonance of -20MHz . At this frequency, we manage to achieve 41% efficiency for the AOM in double pass. For the optical molasses stage which follows the MOT stage, a larger detuning is desired to cool the atoms further and we operate at -60MHz detuning, which corresponds to the AOM operating at 76MHz . This reduces the efficiency of the AOM to 23% in double pass, but the seed power for the slave remains sufficient for stable injection locking (4mW). For imaging we bring the light to resonance with the transition whereby the AOM operates at 106MHz . This has a reduced double pass efficiency of 30% , but there is still enough power for imaging.

The TA AOM has a centre line 110MHz and bandwidth 24MHz , and also operates in a double pass configuration. Throughout the experiment the frequency is set at 96MHz with a double pass efficiency of 60% , giving us -20MHz detuning for the 2D MOT. After the TA, there is one final AOM which shifts the frequency closer to the $F' = 2$ transition. This is done to optically pump the atoms in to the magnetically trappable $F = 2, m_F = +2$ state after the molasses stage. This OP AOM has a centre line of 200MHz and a bandwidth of 50MHz , and operates at 268MHz with low efficiency (25% single pass). This is acceptable as the optical power has been amplified by the TA, so that 25% efficiency still gives sufficient power for optical pumping.

The cooling beams in principle should be enough to cool the atoms by themselves. In practice however, due to the proximity between the $F' = 2$ and the $F' = 3$ levels, atoms excited by the cooling laser can also be excited in to the $F' = 2$ state and decay to the $F = 1$ state, thus escaping the closed cycle. To address this, light is needed which pumps

the atoms out of the $F = 1$ state and back in to the cooling transition. Our laser scheme has a separate repumper laser (ECDL) which does this, where the lock point is directly on the $F = 1 \rightarrow F' = 2$ transition.

With the frequencies appropriately set, the laser beams can be applied to the atoms. We fibre couple our laser beams from the optics table to the vacuum table which ensures misalignments on the optics table only affect the available power on the vacuum table, rather than causing misalignments on the atoms.

Whilst each fibre has a shutter before the input to dynamically block the beams, these shutters are mechanical and have a delay of typically 1ms and a switching time of $> 100 \mu\text{s}$. The AOMs however have fast switching times ($< 1 \mu\text{s}$) and can be used as a fast shutter, ideal for sending short pulses. We therefore have the imaging fibre directly after the double pass Slave AOM, where the AOM produces a fast pulse ($\sim 10 \mu\text{s}$) of resonant light for our optical absorption images. We also have the optical pumping fibre after our OP AOM, where only the first order is sent to the fibre, creating the pulse (0.5 ms) necessary to pump the atoms in to the magnetically trappable state, whilst keeping the pulse short so as to reduce heating effects.

All fibres have a 75 mm achromatic doublet lens after their output except for the push beam which has an 8 mm aspheric lens. The push beam has a $1/e^2$ beam radius of 1 mm, comparable to the size of the aperture (1.2 mm) of the low-conductance tube in the loading chamber. The other beams have ~ 9 mm $1/e^2$ radius. The fibres and their associated power outputs are shown in Table 3.1.

Table 3.1: Available laser power after the optical fibres used to transport the light from the optics table to the vacuum table.

Fibre	Output (mW)
3D Cooling	130
2D Cooling	270
3D Repumper	1.8
2D Repumper	8
Push Beam	2
Optical Pumping	5
Imaging	0.5

3.4 2D MOT, 3D MOT and Optical Molasses

As previously mentioned, the initial stage of the experiment is the MOT which utilises optical molasses and the Zeeman effect. MOTs are able to directly capture rubidium atoms at room temperature and cool them to $\sim 100 \mu\text{K}$ [145]. When loaded from an atomic beam, high atom numbers (typically $10^9 - 10^{10}$) are readily attainable, providing an ideal source for cold atoms.

The MOT stage is then followed by a short, purely-optical molasses phase where the magnetic field of the MOT is switched off and the detuning of the laser cooling beams are increased. The optical molasses stage is able to take further advantage of sub-Doppler cooling mechanisms and the cloud can be cooled further before transfer in to a magnetic trap. This sequence takes advantage of the strong trapping attainable by the MOT compared to optical molasses, whilst still exploiting the colder temperatures of the latter.

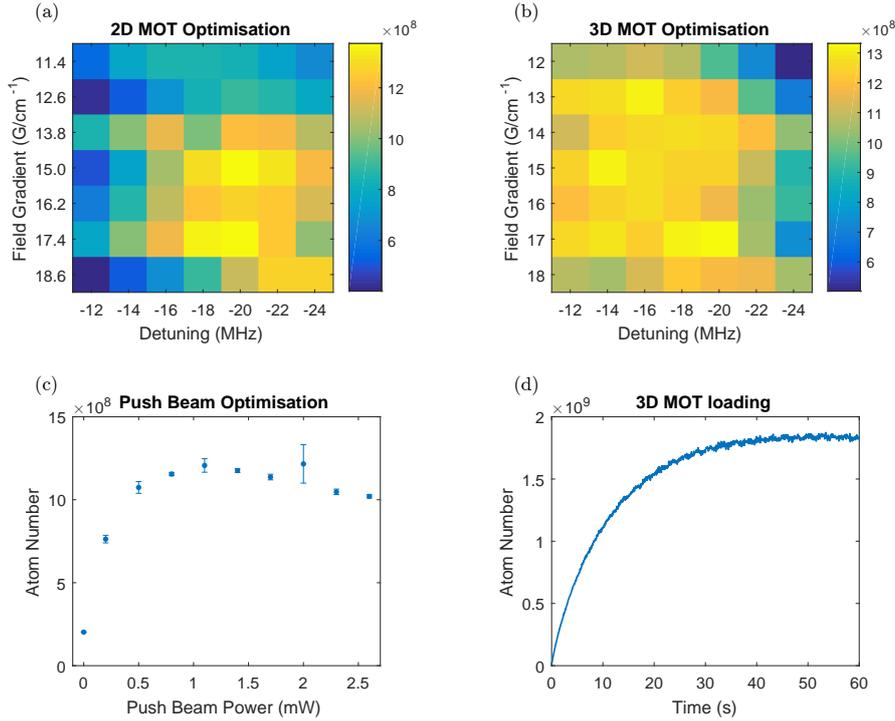


Figure 3.11: MOT optimisation with 10 s loading. (a) 2D MOT scan with $\delta_{3D} = -15$ MHz and 3D field gradient of 15 Gcm^{-1} . (b) 3D MOT scan with $\delta_{2D} = -20$ MHz and 2D field gradient of 15 Gcm^{-1} . (c) Push beam power optimisation with $\delta_{3D} = -15$ MHz, 3D field gradient of 15 Gcm^{-1} , $\delta_{2D} = -20$ MHz and 2D field gradient of 15 Gcm^{-1} . (d) A typical MOT loading curve where $> 1 \times 10^9$ atoms are trapped in 10 s.

A large atom number in the initial MOT stage is sought as latter stages of the experiment involve the lossy process of evaporative cooling (see section 3.6 and 3.7). Fast loading times are also desirable so as to minimise the duration of an experimental run. The optimisation of a MOT requires scanning through various parameters including frequency detuning, beam alignment and magnetic field gradient. Since the atoms are undergoing laser cooling during the MOT stage, spontaneous emission is abundant and the scattered light can be used to characterise the MOT via fluorescence imaging.

We detected the fluorescence of the atoms via a 60mm lens in a 2f-2f configuration on to a photodiode (Thorlabs DET100A/M) which was placed on the horizontal optical absorption imaging path in Figure 3.7 (after the MOT optimisation procedures, the photodiode was removed to make room for our absorption imaging beam). These are the only requirements for fluorescence detection as the atom number is related to the photodiode reading by

$$N = \frac{4\pi}{0.96\kappa\Omega} \frac{P}{h\nu R}, \quad (3.2)$$

where N is the atom number, κ is the number of glass surfaces between collection lens and atoms (0.96 is the transmitted percentage through an uncoated glass surface), Ω is the solid angle subtended by the collection lens, P = photodiode current/photodiode responsivity,

ν is the frequency of the photons and R is the photon scattering rate given by

$$R = \frac{\Gamma}{2} \frac{\frac{I}{I_{sat}}}{1 + \frac{I}{I_{sat}} + 4\frac{\delta^2}{\Gamma^2}}. \quad (3.3)$$

Here, Γ corresponds to the natural linewidth of the ^{87}Rb D_2 line ($2\pi \times 6.1$ MHz), I is the total intensity of all MOT beams on the atoms, I_{sat} is the saturation intensity (1.67 mWcm^{-2} for circularly polarised light) and δ is the frequency detuning from resonance of the beams.

Some examples of optimisation scans using fluorescence detection are shown in Figure 3.11. We optimised both the 2D-MOT and 3D-MOT parameters by using fluorescence of the 3D MOT alone, as the two systems are coupled via the atomic beam. The optimisation procedures in Figure 3.11 were done with 10 s loading time. After the initial alignment procedure of overlapping the MOT beams and ensuring they were well centred on the chamber, alignment was further optimised by using the fluorescence imaging as well. Alignment of the push beam was done in a similar fashion, except for short loading times, where the passage of the beam through the optical pumping viewport at the end of the science chamber was used as an initial guide. The optimised power configuration for the 2D⁺ MOT beams was found to be zero power in the push beam direction (see Figure 3.3) and 1.5 mW in the counter-propagating direction, where the optimal push beam power was 2 mW. The rubidium background pressure in the loading chamber is kept constant (on the timescale of several hours) by the dispensers at $\sim 10^{-7}$ mbar. As can be seen from Figure 3.11, the MOT is robust to various changes in parameters, but regions of optimised loading are still apparent.

Shim coils for the 3D MOT were used to ensure isotropic expansion of the atom cloud. The shim coils for the 2D MOT were more critical for the loading stage as they were used to align the zero field line to the low conductance tube and thus maximise the 3D MOT loading rate.

In future, the atom number and loading rate could be further improved with the addition of more cooling laser power, as can be seen in Figure 3.12 (b) and (d). However, based on Figure 3.12 (a) and (c), the addition of further repumper power would not obviously benefit the experiment.

With the MOT optimised we are able to load 1.3×10^9 atoms in 10 s, whilst a fully loaded MOT reaches 1.8×10^9 atoms in 30 s. We find that loading for 30 s has a minimal effect on the number of atoms transferred in to the magnetic trap, described in section 3.6, compared to loading for 10 s. Our typical experimental run therefore has a 10s MOT loading stage. To summarise, the parameters for our 2D MOT and 3D MOT are shown in Table 3.2.

After the MOT stage there is a short optical molasses phase that achieves colder temperatures ahead of the transfer to our magnetic trap. The sub-Doppler cooling mechanisms which are further exploited in this phase result in the following temperature dependence [146]:

$$k_B T \propto \frac{I}{|\delta|}, \quad (3.4)$$

where T is the temperature of the atoms, I is the intensity of the cooling beams and δ is the detuning. Increasing the detuning during the optical molasses stage is therefore advantageous, and in our experiment we increase the detuning from -20 MHz to -60 MHz in 5 ms. Beyond -60 MHz detuning the AOM efficiency is too low to allow for stable injection of the slave laser. The magnetic field is also ramped down from 17 Gcm^{-1}

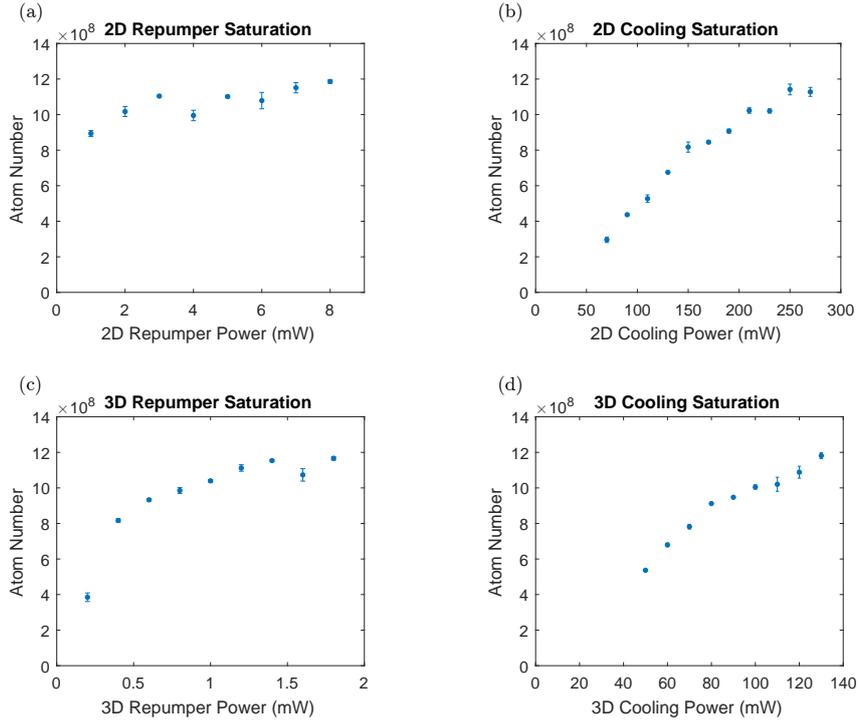


Figure 3.12: The number of atoms trapped in the MOT shows a saturation behaviour. This data indicates that the experiment could be improved by the addition of more cooling power, but additional repumper power is not necessary in the current configuration. Loading times for all data was 10s.

Table 3.2: MOT parameters. For reference for the shim coil axes see Figure 3.7).

Parameter	2D MOT	3D MOT
Detuning (MHz)	-20^1	-20
$\frac{\partial B}{\partial z}$ (Gcm $^{-1}$)	15	17
Power per cooling beam (mW)	37	17.5
Total repumper power (mW)	8	1.8
$1/e^2$ MOT beam radius (mm)	$9(\times 3)$	9
Shim coil x (G)	2	1.5
Shim coil y (G)	-	1.1
Shim coil z (G)	1.5	0

to 0 Gcm^{-1} during this time. The atoms are then left in this optical molasses phase for a further 10 ms before optical pumping and transfer in to the magnetic trap.

Our optical molasses achieves a temperature of $53 \pm 4 \mu\text{K}$. The duration and detuning of the optical molasses stage was optimised by observing the temperature in the magnetic trap after transfer using optical absorption imaging.

¹We later found that -15MHz detuning for the 2D MOT led to slightly colder temperatures in the transfer to the magnetic trap with similar atom number. Subsequent data is therefore taken with -15MHz for the 2D MOT detuning.

3.5 Imaging

We use absorption imaging to accurately diagnose the atoms for all stages beyond the MOT. This technique uses a short pulse ($\sim 10 \mu\text{s}$) of resonant light ($F = 2 \rightarrow F' = 3$) on the atoms which is then imaged on to a camera. Two subsequent images are also taken, one of the imaging beam itself with the same pulse duration, another of background light. With these three images the optical density of the atoms can be determined and the atom number and cloud size can be calculated.

The absorption of light due to the atoms can be described by Beer-Lambert's law where the intensity of the beam decays exponentially as it passes through the cloud [24]. For incident light illuminating the atoms with intensity I_{in} far below the saturation intensity of the transition, the resulting intensity can be described as

$$I_{\text{out}}(x, y) = I_{\text{in}}(x, y) \exp\left(-\sigma(\omega) \int_{z_1}^{z_2} n(x, y, z) dz\right), \quad (3.5)$$

where the light is propagating in the z -direction, $\sigma(\omega)$ is the scattering cross section for absorption, $n(x, y, z)$ is the density of the cloud and the integral spans the cloud size.

For incident light with isotropic light polarisation (or equivalently, atoms without a quantisation axis with Zeeman sub-levels degenerate), the scattering cross section can be written as

$$\sigma(\omega) = \frac{\sigma_{0(\text{iso,eff})}}{1 + \left(\frac{2(\omega - \omega_0)}{\Gamma}\right)^2}, \quad (3.6)$$

where $\omega - \omega_0 = \delta_{\text{Img}}$ and $\sigma_{0(\text{iso,eff})} = 1.36 \times 10^{-13} \text{m}^2$ [142], which is the resonant scattering cross section for isotropic light polarisation between the $F = 2 \rightarrow F' = 3$ transition for the D_2 line.

The integral in Equation 3.5 results in a column density which can be written as

$$\begin{aligned} n_{2D}(x, y) &= \int_{z_1}^{z_2} n(x, y, z) dz \\ &= \frac{1}{\sigma(\omega)} \ln\left(\frac{I_{\text{in}}(x, y)}{I_{\text{out}}(x, y)}\right) = \frac{1}{\sigma(\omega)} \text{OD}(x, y), \end{aligned} \quad (3.7)$$

where $\text{OD}(x, y)$ is the optical density. Using Equation 3.7 we can determine the atom number, N , from our pictures by summing over the pixels:

$$N = \frac{1}{\sigma(\omega)} \left(\frac{d_{px}}{M}\right)^2 \sum_{ij} \text{OD}(i, j) = \frac{1}{\sigma(\omega)} \left(\frac{d_{px}}{M}\right)^2 \sum_{ij} \ln\left(\frac{I_2(i, j) - I_3(i, j)}{I_1(i, j) - I_3(i, j)}\right). \quad (3.8)$$

Here the pixel indices of the images are denoted by i and j , the pixel pitch is d_{px} , M is the magnification of the imaging system, $I_1(i, j)$ is the image of the absorption from the atoms, $I_2(i, j)$ is the picture of the imaging beam without the atoms and $I_3(i, j)$ is the image of the background light with no atoms or imaging beam. The method approximates that the optical density remains constant over each pixel on the camera. Each image in our sequence is separated by 200 ms and the imaging pulses have an 80 μs duration. Figure 3.13 shows an example of the three images and a resulting optical density image taken with atoms in the magnetic trap.

Along with atom number, the temperature of the cloud can also be estimated using time-of-flight (TOF) measurements. Images of atoms which have been released from the trap for a variable time (typically a few milliseconds) can be used to infer the momentum distribution of the atoms.

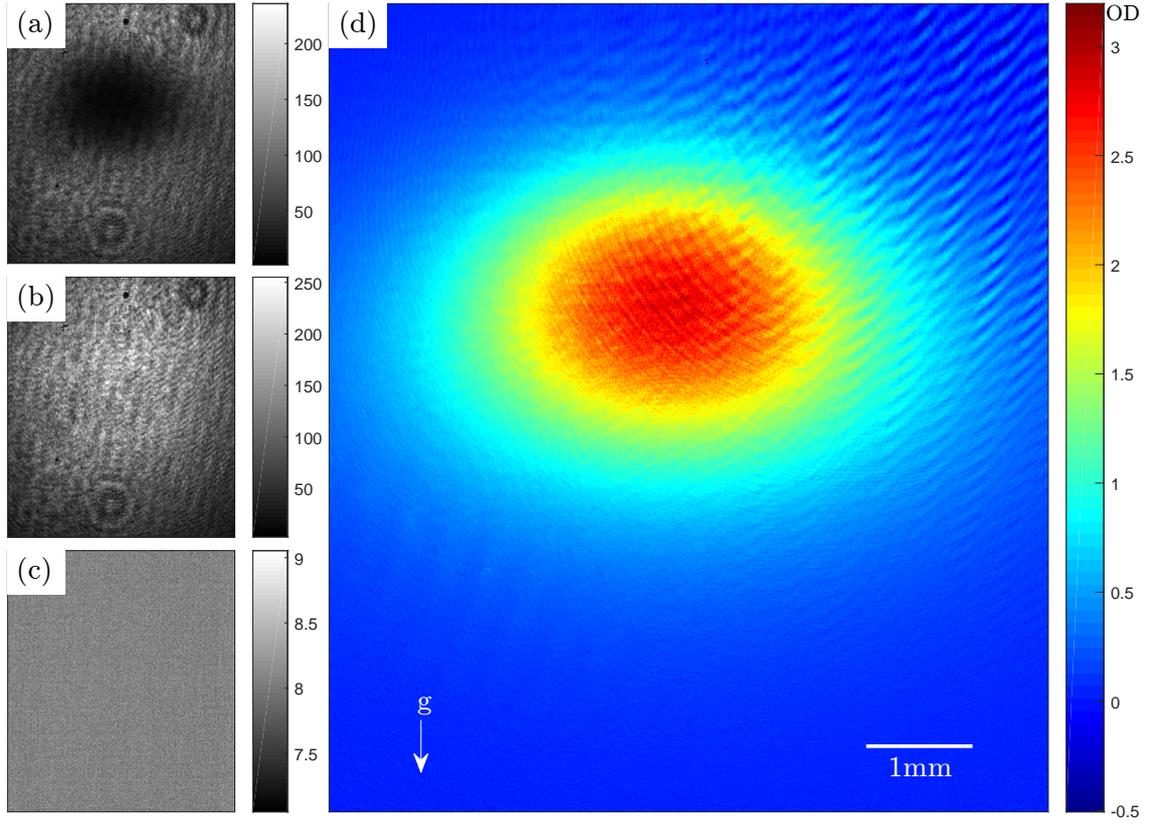


Figure 3.13: Optical absorption imaging of atoms in the magnetic trap (4ms TOF) taking with the horizontal imaging system with a magnification of 1.05. (a) Imaging beam absorbed by atoms, I_1 . (b) Imaging beam only, I_2 . (c) Background light, I_3 . (d) Optical absorption image created from (a),(b) and (c). Due to the stronger confinement in the z-direction created by the quadrapole coils, the cloud is more compressed along this axis. The short TOF also slightly reveals the sag due to gravity of the confined cloud. The OD becomes negative in certain regions due to noise on the images from small vibrations.

For a thermal cloud in a harmonic trapping potential, the position distribution of the atoms is a Gaussian (given in 1D, but extends to 3D)

$$n_{1D}(x) = \frac{1}{\sqrt{2\pi}\sigma_x} \exp\left(-\frac{x^2}{2\sigma_x^2}\right), \quad (3.9)$$

whilst the velocity distribution of the thermal atoms is the Maxwell-Boltzmann distribution

$$f(v_x) = \sqrt{\frac{m}{2\pi k_B T_x}} \exp\left(-\frac{mv_x^2}{2k_B T_x}\right) = \frac{1}{\sqrt{2\pi}\sigma_{v_x}} \exp\left(-\frac{v_x^2}{2\sigma_{v_x}^2}\right), \quad (3.10)$$

where T_x is the temperature of the atoms. For a ballistic expansion of the cloud over a given time t , the final position of an atom is given by $x = x_0 + v_x t$ and the final position distribution can be taken as the convolution of the two functions above. This results in another Gaussian:

$$n_{1D}(x, t) = \frac{1}{\sqrt{2\pi}\sigma_x(t)} \exp\left(-\frac{x^2}{2\sigma_x(t)^2}\right), \quad (3.11)$$

where the width is given by

$$\sigma_x(t) = \sqrt{\sigma_x(0)^2 + \frac{k_B T_x}{m} t^2}. \quad (3.12)$$

Thus, extracting σ_x from images of the cloud at various TOF allows us to extract the temperature of the gas by curve-fitting with Equation 3.12.

The vertical and horizontal imaging schemes are shown in Figure 3.14. We use an Andor Luca CCD camera (1004×1002 pixels with pixel pitch $8 \mu\text{m}$) for either horizontal or vertical imaging, as required. The initial horizontal imaging system used for diagnosis of the magnetic trapping stage was a 1:1 confocal telescope with two 60 mm achromatic doublet lenses. Subsequent imaging in the ODT has used a 1:2.5 confocal telescope for higher magnification, where the last lens before the camera is replaced with a 150 mm doublet and the camera moved appropriately. The vertical imaging path is more complicated. Here, the lens which is closest to the atoms is illuminated by three beams with differing requirements: (1) a MOT beam incident from below, which must be collimated at the point of the atoms; (2) the SLM beam, which will be focused onto the atoms by this lens; (3) the imaging beam incident from above, containing the diverging shadow of the atoms, which must be collected by the camera. To accommodate all these requirements, we use a four-lens imaging system, as shown in Figure 3.14(b). Below the atoms at a distance of 40 mm, the first of a confocal pair of $f=40$ mm lenses focusses the SLM light and collects the imaging beam. This short focal length (high NA, i.e. small diffraction limit) allows us to create the most detailed patterns with our SLM. The lens pair means that there is no net change to the MOT beam, other than a slight power-loss which can be compensated for by rebalancing the MOT beams. However, the resultant image plane (160 mm from the atoms) is too close for convenient combination/separation of the beams before the camera. We therefore add a second set of lenses to image that plane onto the camera at an increased separation, which allows space for both a Thorlabs DMSP1000L dichroic mirror (to overlap the SLM light) and a PBS (to separate the imaging light from the orthogonally-polarised downward MOT beam). This second telescope now only effects imaging light, and so can be used to magnify the image: using a $f=75$ mm doublet and $f=200$ mm singlet gives us a magnification of 2.67. We note that the resolution of our imaging system is now limited by the size of the lenses in this telescope.

3.6 Magnetic Trap and RF Evaporative Cooling

After the optical molasses stage the atoms undergo optical pumping in preparation for transfer in to the magnetic trap. A short $500 \mu\text{s}$ pulse of σ^+ light which addresses the $F = 2 \rightarrow F' = 2$ transition is used whilst a relay switches off the 3D y-axis shim coils. The introduction of the background magnetic field lifts the degeneracy of the Zeeman sub-levels and the σ^+ light pumps the atoms in to the $m_F = 2$ state, as shown in Figure 3.15. This ensures that the majority of the atoms are magnetically trappable in our quadrupole field as the magnetic force experienced by an atom in the state m_F is given by

$$\mathbf{F} = -g_F \mu_B m_F \nabla \mathbf{B}, \quad (3.13)$$

where g_F is the hyperfine Landé g-factor ($g_F = \frac{1}{2}$ for $F = 2$), μ_B is the Bohr magneton and \mathbf{B} is the magnetic field. For the magnetic field gradient produced by a pair of coils in anti-Helmholtz configuration, the resulting trap will have a V-shaped profile in each axis around the zero field point. Using the optical pumping we manage to increase the number of atoms transferred to the magnetic trap by a factor > 2 .

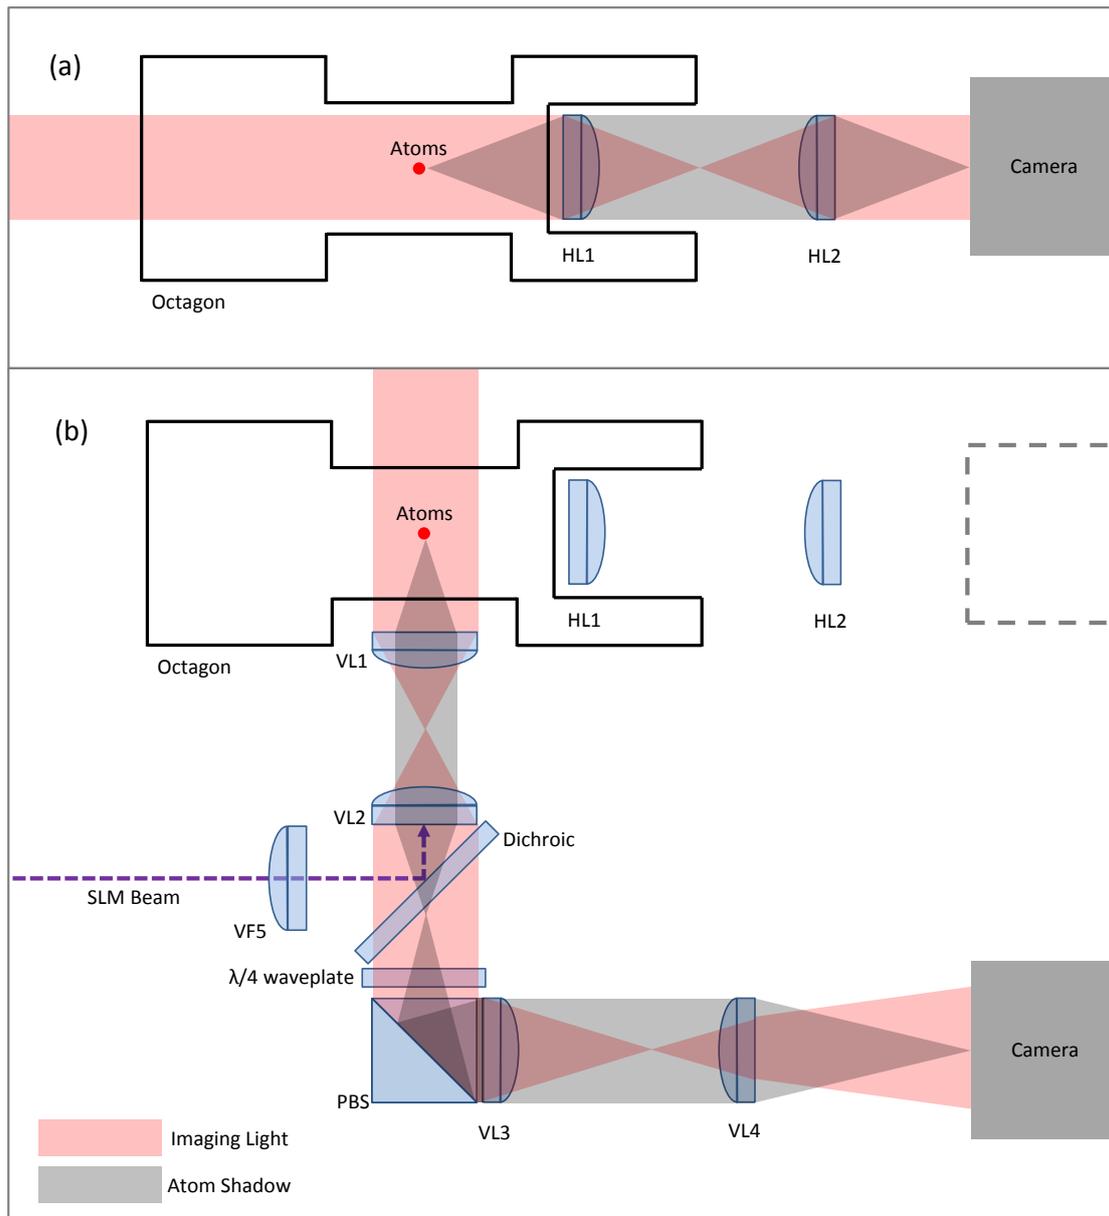


Figure 3.14: Schematic of imaging setup. (a) Horizontal Imaging Setup. (b) Vertical Imaging Setup. The camera is placed on either the horizontal or vertical imaging path. Focal lengths of lenses: HL1=60mm, HL2=60mm or 150mm, VL1 = 40mm, VL2 = 40mm, VL3 = 75mm, VL4 = 200mm, VF5 = 40mm. The dichroic reflects the 1064nm light from the SLM and transmits 780nm light. A quarter waveplate is required for the MOT beam (not shown) arriving from below.

With the atoms in the $m_F = 2$ state, we block off all our laser beams and abruptly switch the coils on to produce a field gradient of 50 Gcm^{-1} . We then compress the atoms (increasing the collision rate further) by ramping the field gradient up to 100 Gcm^{-1} in 100 ms.

At this stage we are able to make a measurement of the magnification of the horizontal

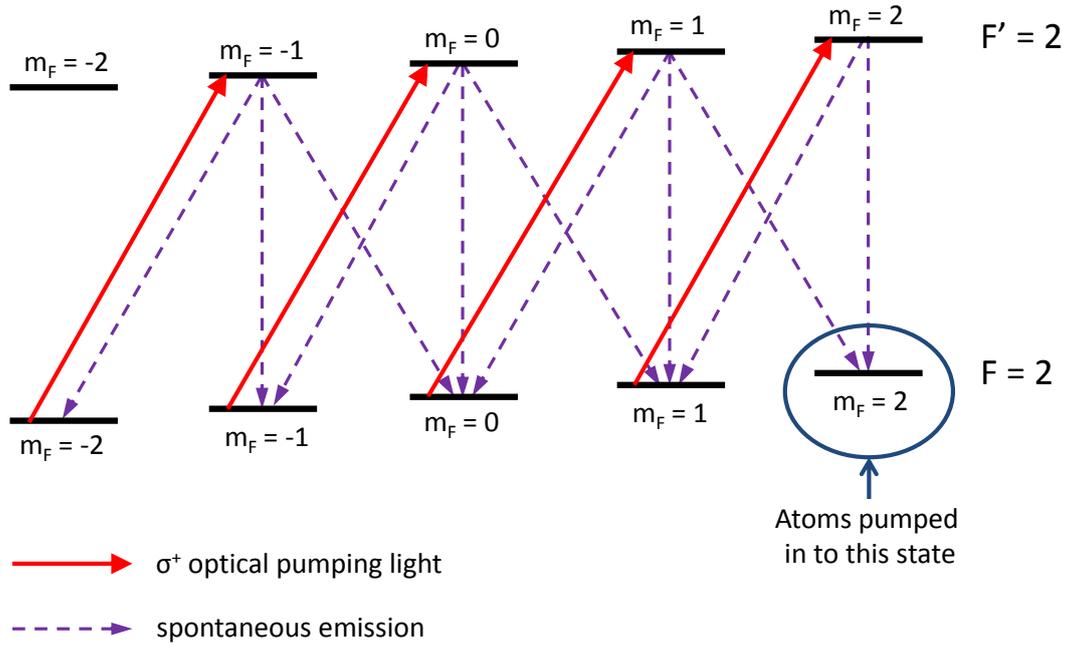


Figure 3.15: Optical pumping. The σ^+ light promotes transitions of $\Delta m_F = +1$ such that atoms accumulate in the $m_F = 2$ magnetically trappable state.

imaging setup (Figure 3.14(a)). This was achieved by releasing the atoms and observing them as they fall, allowing us to extract the acceleration due to gravity. The deviation away from the expected acceleration would then give us the magnification of the imaging system which allows for more accurate diagnostics of the atoms for parameters such as temperature and atom number. An example gravity measurement is shown in Figure 3.16(a).

With the magnification of our images determined, we record $(1.21 \pm 0.05) \times 10^9$ atoms at a temperature of $244 \pm 8 \mu\text{K}$ in the magnetic trap. Varying lifetimes in the magnetic trap have been measured due to dispenser use, but typically the lifetime is > 60 s. An example lifetime plot with $\tau = 180$ s is shown in Figure 3.16 (b).

With the atoms confined in the magnetic trap, we then evaporatively cool them with RF radiation. A schematic of the RF evaporation process is shown in Figure 3.17. The hottest atoms are initially transferred out of the $m_F = 2$ state into anti-trapped states, removing the atoms from the ensemble. The remaining atoms then rethermalise to a colder temperature and with a suitable sweep of the frequency of the radiation, a much colder atomic cloud can be achieved. It should be noted that the V-shaped profile has a degeneracy of the Zeeman sub-levels at the centre of the trap where the magnetic field is zero. This degeneracy allows atoms to spin-flip to anti-trapped states and thus be lost from the ensemble, such that the trap can be thought of as having a hole in the centre. Limiting the effects of these losses is therefore essential, and for this purpose we transfer the atoms from the magnetic trap in to a hybrid trap configuration for a final stage of evaporative cooling (see Section 3.7).

The phase-space density of the atomic cloud is the optimisation criterion for the RF evaporation stage (as well as the hybrid trap stage). The phase-space density for atoms

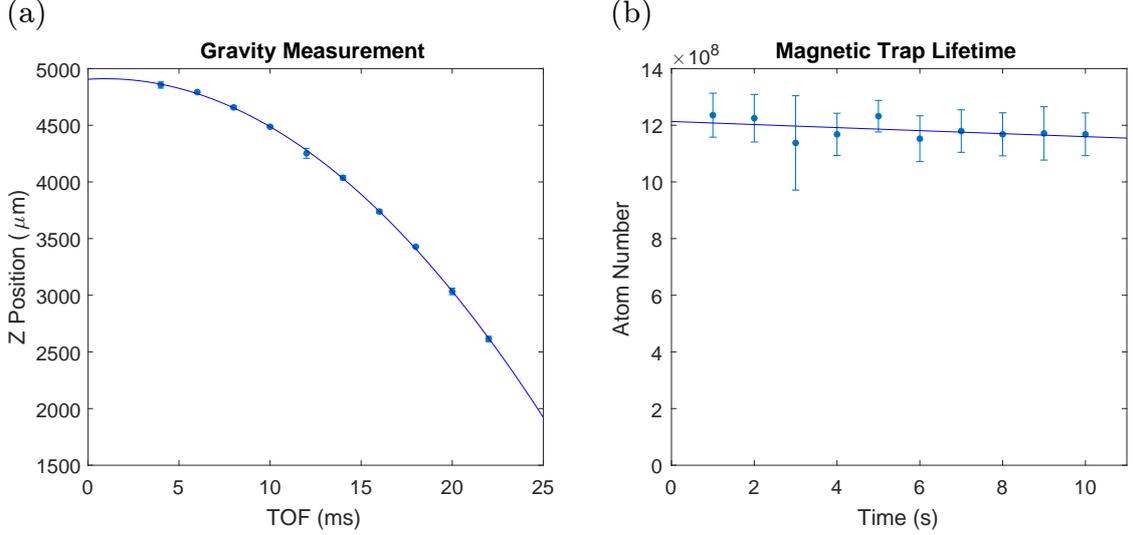


Figure 3.16: (a) Measurement of atoms falling after being released from the magnetic trap. The fit reveals that the magnification of the horizontal imaging system is 1.05 ± 0.05 . The higher magnification used in later stages (see Section 3.7), is measured in the same way. (b) Atom number varying with time held in the magnetic trap. The lifetime is measured as 180 ± 90 s. The lack of losses in the trap over 10 s (data was taken only up to 10 s holding time to avoid over heating the coils) shows that the science chamber has a suitably low vacuum pressure.

in the magnetic trap can be written in terms of the temperature T , atom number N and magnetic field gradient B' as [140]

$$PSD = \frac{N}{32\pi} \left(\frac{\mu_B B'}{k_B T} \right)^3 \left(\frac{2\pi\hbar^2}{mk_B T} \right)^{\frac{3}{2}}, \quad (3.14)$$

where μ_B is the Bohr magneton. With temperature and atom number measurements we are therefore able to infer the phase-space density at different stages during the RF evaporation.

We generate our RF radiation with two single loops of wire in Helmholtz configuration which are wrapped around the magnetic trapping coils and connected to a TTi TG2000 function generator along with a 40 dB gain RF amplifier (Mini Circuits ZHL-5W-1). The function generator has a maximum frequency output of 20MHz and the amplifier has a maximum output power of 5W. The RF coils also have a 50Ω RF terminator connected in series for impedance matching, and the function generator is connected to the amplifier via attenuators and a limiter (Mini Circuits VLM-33W-2W-S+) which ensures the output does not exceed the maximum threshold for the input of the amplifier.

The Zeeman sub-levels have a splitting which scales with the magnetic field as $\Delta E = |\frac{1}{2}\mu_B B|$ in the state $F = 2$. For our z-direction field gradient of 100Gcm^{-1} (50Gcm^{-1} in x and y), RF radiation at a frequency of 20MHz corresponds to an RF knife occurring approximately 2.9mm either side of the trap centre. The $1/e^2$ half-width of the cloud in the trap along z, measured using Equation 3.12, is $2.16 \pm 0.06\text{mm}$ which is therefore within the expected starting radius of the RF knife. For our initial RF sweeps we found that little to no evaporation occurred between 20MHz and 17MHz and thus our initial starting frequency for the sweep is 17MHz.

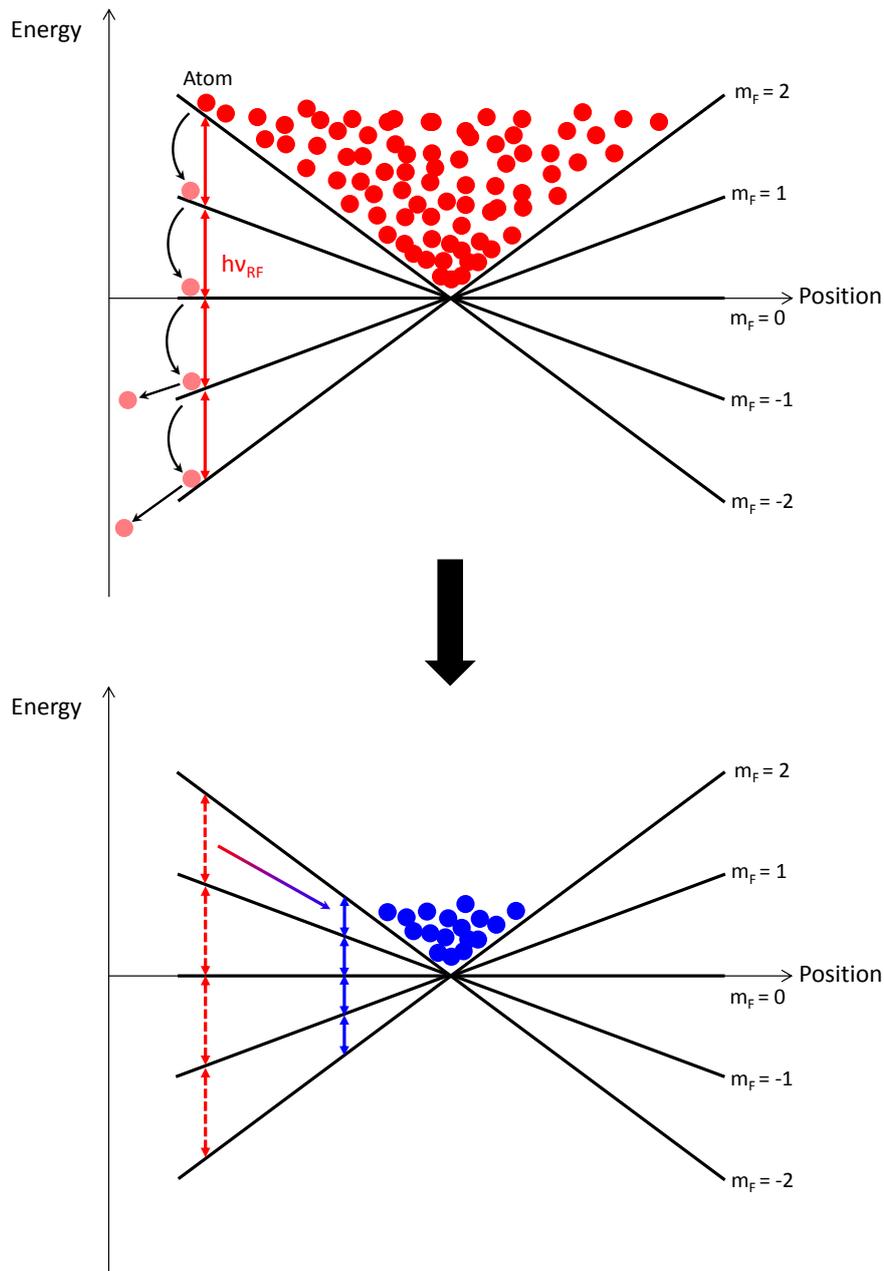


Figure 3.17: A schematic representation of RF evaporative cooling. High energy atoms in the magnetic trap are spin-flipped to anti-trap states via RF electromagnetic radiation. Steadily reducing the frequency of the radiation sweeps the “RF knife” from high to low energies, cooling the ensemble.

With the starting frequency established, the other parameters to optimise for a linear shaped sweep² are the amplitude, final frequency and duration. The amplitude of the RF signal must be large enough to ensure efficient state transfer and thus efficient evaporation. The final frequency is chosen to ensure a low temperature for transfer in to the hybrid trapping stage whilst also maintaining a large atom number against Majorana losses. The optimal duration of the sweep is effected by the collision rate of the atoms, where

a low collision rate must be met with a longer sweep duration to give the atoms time to thermalise during the losses. The collision rate can be improved by stiffening the trap further, but this leads to increased heating during the compression stage of the magnetic trap.

For our system we found that the highest increase in phase-space density for the atoms occurred for a 10 s linear sweep between 17 MHz and 1.4 MHz with the maximum signal output attainable via the amplifier. This sweep leads to a cloud of $(2.3 \pm 0.4) \times 10^7$ atoms with a temperature of $19.4 \pm 0.2 \mu\text{K}$, resulting in a phase space density of $(7.4 \pm 1.2) \times 10^{-4}$. The PSD plot for the RF evaporation stage is shown in Figure 3.18.

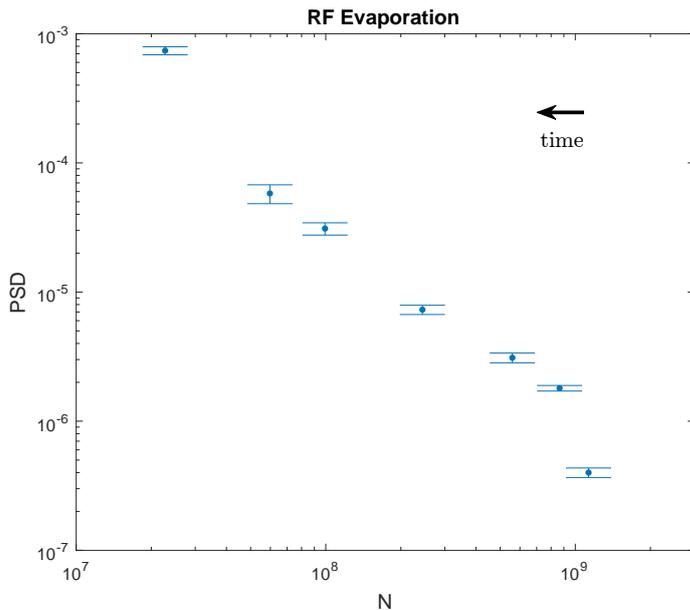


Figure 3.18: PSD plot of RF evaporation. The initial value (right-most) is the PSD after transfer into the magnetic trap. The final value (left-most) corresponds to the atoms 300ms after the 10s RF sweep has finished. Intermediate points are at various stages of the RF sweep.

3.7 Hybrid Trap

After the RF evaporative cooling stage comes the transfer to the focused ODT. In the presence of light, atoms experience a shift in their energy levels due to the A.C. stark effect which can be utilised to form a trap for the atoms [147]. If the polarisation of the light used on the rubidium atoms is linear, and the frequency has a large detuning compared to the hyperfine splitting of the excited state, the energy shift for the ground state can be written as:

$$\Delta E(\mathbf{r}) = \frac{\pi c^2 \Gamma}{2\omega_0^3} \left(\frac{2}{\Delta_{2,F}} + \frac{1}{\Delta_{1,F}} \right) I(\mathbf{r}), \quad (3.15)$$

²Double segmented linear sweeps were attempted but we found the best results from a single linear sweep.

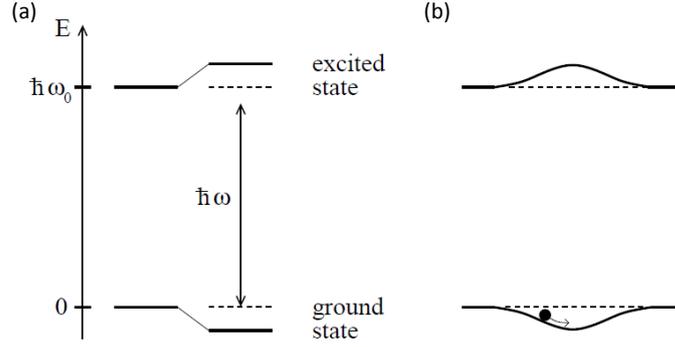


Figure 3.19: (a) Light shift for a two-level atom with red-detuned light. (b) The resulting light shifts for a red-detuned laser with a Gaussian profile, creating a trap for an atom in the ground state. Figure taken from [147].

where ω_0 is the resonant frequency of the D_1 transition (D_1 line resonance is at 795nm for ^{87}Rb [142]), $I(\mathbf{r})$ is the intensity of the light and $\Delta_{2,F}$ and $\Delta_{1,F}$ are the detunings of the light with respect to the D_2 and D_1 line. If the detunings are sufficiently far from resonance that the atoms experience minimal heating from photon scattering (which scales proportionally to I/Δ^2 [24]), then the response of the atoms will be dominated by the light shift. Thus, a red detuned laser beam with a Gaussian profile will create a Gaussian shaped trapping potential for the atoms, as depicted in Figure 3.19. Furthermore, evaporative cooling can be achieved by simply reducing the laser power, although this has shortcomings in that the density is reduced.

To produce the focused optical dipole trap we use a fibre laser (IPG Photonics, YLP-

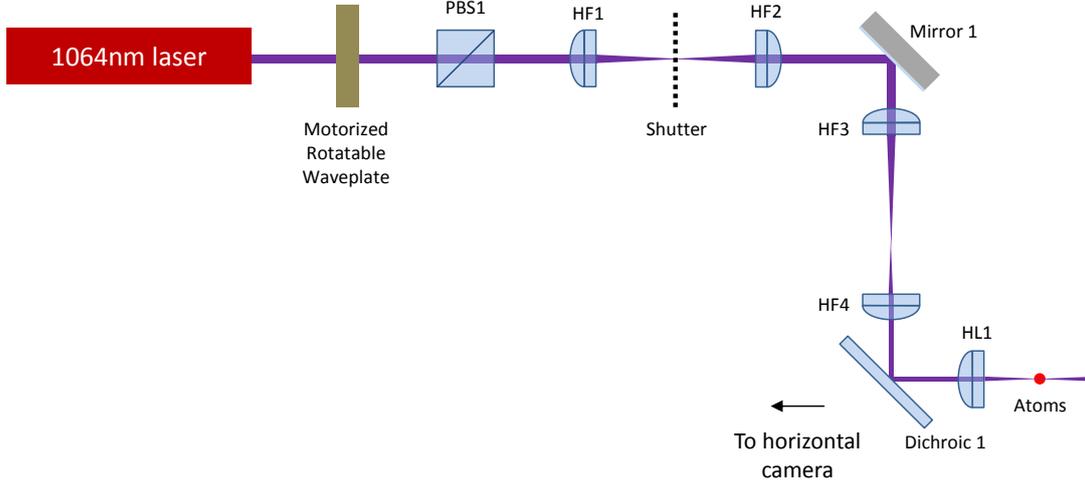


Figure 3.20: Simplified diagram of the 1064 nm laser setup for the horizontal ODT. Periscopes and various mirrors are not present in the diagram for clarity. Focal lengths of lenses: HF1 = 100 mm, HF2 = 100 mm, HF3 = 150 mm, HF4 = 60 mm, HL1 = 60 mm. HF1 is mounted on a translation stage (not shown). The HL label denotes the lens which is common to the imaging axis as well, as shown in Figure 3.14.

5-1070-LP) which generates a maximum beam power of 5 W at 1064 nm. The $1/e^2$ beam width of the laser is 1 mm and we use a combination of three lenses to create a focused spot size of $w = 50 \mu\text{m}$ on the atoms, where w is the $1/e^2$ half-width (see Figure 3.20). Using the peak intensity ($2P/\pi w^2$ where P is the beam power) and Equation 3.15, an estimated trap depth of $120 \mu\text{K}$ is created with 3.5 W of power on the atoms. In order to perform evaporative cooling, the beam power is reduced via a motorised rotatable $\frac{\lambda}{2}$ waveplate (Thorlabs PRM1Z8 and TDC-001) and a PBS where the split light is initially blocked for optimisation of the focused ODT, but is later used for the vertical trapping path (this will be covered in Section 4.1). A 1:1 confocal telescope is used in the path to create a focus for a mechanical shutter to close off the light, where the first lens of the telescope is mounted on a micrometer translation stage. The translation stage allowed for the focus of the beam to be adjusted on the atoms. During the sequence the ODT is on throughout to avoid the shutter from overheating. The shutter blade is a roughened (using sandpaper) aluminium plate which is angled at 45 degrees to the incident beam so as to disperse the high power in to a beam block, and is only used to block the focussed light for the duration of the TOF ($\sim 10 \text{ ms}$). No difference in atom number or temperature was found between switching the ODT on after the RF evaporation or leaving it on throughout the sequence. The schematic of the 1064 nm light for the horizontal path is shown in Figure 3.20.

The initial alignment of the focused ODT was performed by overlapping the 1064 nm light centrally with the counter-propagating imaging beam and releasing the atoms from the trap whilst leaving the ODT on. Due to the relatively weak focussing of the ODT, the beam has a weak trapping potential in the propagation axis, but a confining effect could still be seen on the cloud and the beam could be positioned near the centre of the magnetic trap using this method.

The transfer to the hybrid trap involves ramping down the magnetic field to allow the cloud to expand and fall in to the ODT. The final value of the magnetic field should just

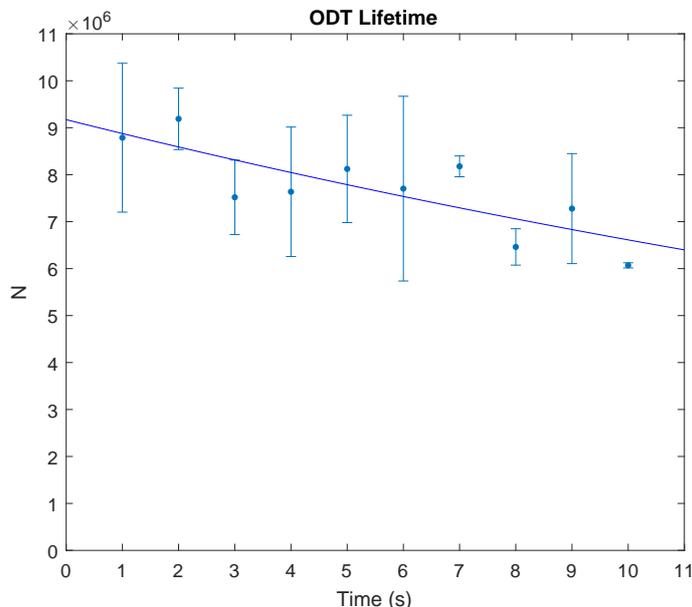


Figure 3.21: Atom number varying with time held in the ODT. The lifetime is measured as $\tau = 32 \pm 9 \text{ s}$.

under-compensate gravity to allow for atoms to evaporate out of the ODT when the laser power is reduced. We found that the optimal magnetic field ramp for the atoms was from 100 Gcm^{-1} to 14 Gcm^{-1} in 1 s. The atoms were largely unaffected for ramp times of > 1 s. For final magnetic field values of $> 15 \text{ Gcm}^{-1}$ it was found that the expanded atoms that were kicked out of the field of view of the camera would eventually come back, indicating that the magnetic field was over-compensating gravity. We found that the largest number of atoms was transferred in to the focused ODT when it was positioned $\sim 60 \mu\text{m}$ below the centre of the magnetic trap in the z-direction (gravity).

The optimised transfer achieved led to $(9.8 \pm 1.6) \times 10^6$ atoms with a temperature of $9.9 \pm 0.1 \mu\text{K}$. The lifetime of the atoms after transfer to the hybrid trap is ~ 30 s, which is reduced from the magnetic trap lifetime due to the finite trap depth that the ODT has, leading to free evaporation. An example lifetime plot is shown in Figure 3.21. This transfer led to a further increase in PSD (factor ~ 20) which could be estimated with [140]

$$PSD \simeq N \left(\frac{\hbar}{k_B T} \right)^3 \frac{1}{\sqrt{\omega_x \omega_y \omega_z}}, \quad (3.16)$$

where $\omega_x \simeq \frac{1}{2} \sqrt{\mu_B B' / m z_{\min}}$ (assuming the magnetic field dominates the propagation axis confinement and where z_{\min} is the distance below the zero point of the magnetic field where the ODT is situated), $\omega_y = \omega_z \simeq 2 \sqrt{U_0 / m w^2}$ (assuming the ODT confinement is dominant over the magnetic field in both axes) and U_0 and w is the ODT trap depth and trap waist respectively.

With the atoms transferred to the hybrid trap, the last stage of evaporative cooling could be optimised. The motorised rotatable waveplate can approximately ramp through angular positions linearly (an acceleration and deceleration stage is present at the beginning and end of the ramps), whilst the transmission and reflection response of the PBS is sinusoidal. When optimising, both a single linear ramp and a double segmented linear

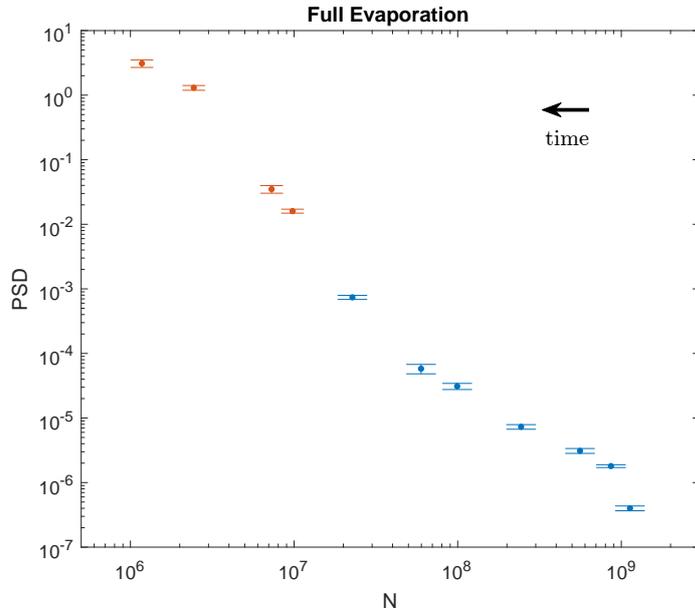


Figure 3.22: Phase space densities for all evaporation stages. The orange points are hybrid trap PSD results, whilst the blue points are magnetic trap PSD results.

ramp were investigated, where the double segmented ramp with a faster segment at the beginning was found to result in higher PSD. A time-scale similar to the RF evaporation stage was needed for optimal PSD, where the total ramp time was 12 s with an initial 3 s ramp from 100% power to 30% power and a final 9 s ramp from 30% to a desired final trapping depth (typically 1 – 2% power). The PBS and waveplate would have a leakage of < 10 mW ($< 0.29\%$) at minimum transmission, however, we found that all the atoms are removed at this power and so full evaporative control was possible.

With the optimised double segmented ramp, the first evidence of the BEC was seen. The PSDs of both the ODT and RF evaporation stages are shown in Figure 3.22, just before condensation is reached. We found that condensation would not occur until estimated PSDs of > 3 were achieved.

3.7.1 Bose-Einstein Condensation

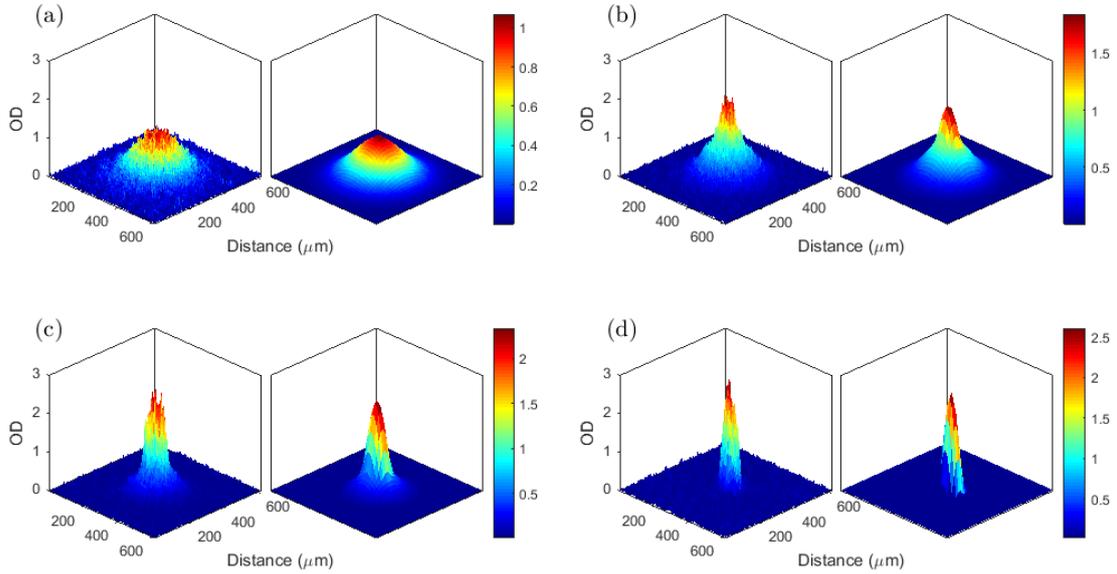


Figure 3.23: Formation of the Bose-Einstein Condensate. All images were taken with -4 MHz imaging detuning at 21ms TOF. The absorption images (left) are accompanied by the 2D fits (right). (a) Trap Depth $2.6 \mu\text{K}$, $N_{\text{Tot}} = 1.25 \times 10^6$, $N_{\text{BEC}}/N_{\text{Tot}} = 0$. (b) Trap Depth $2.3 \mu\text{K}$, $N_{\text{Tot}} = 1.06 \times 10^6$, $N_{\text{BEC}}/N_{\text{Tot}} = 0.06$. (c) Trap Depth $1.8 \mu\text{K}$, $N_{\text{Tot}} = 6.11 \times 10^5$, $N_{\text{BEC}}/N_{\text{Tot}} = 0.29$. (d) Trap Depth $1.2 \mu\text{K}$, $N_{\text{Tot}} = 2.06 \times 10^5$, $N_{\text{BEC}}/N_{\text{Tot}} = 0.87$.

The onset of Bose-Einstein condensation can be observed by a bimodal distribution in the cloud density. The atoms in the condensate take the form of an inverted parabola (in the Thomas-Fermi regime) for a harmonic potential, and a small central peak within the thermal distribution becomes apparent below the critical temperature. The density of the inverted parabola, with a peak density $n_{\text{BEC}}(0)$, takes the form [24]

$$n_{\text{BEC}}(x, y, z) = n_{\text{BEC}}(0, 0, 0) \left(1 - \frac{x^2}{R_x^2} - \frac{y^2}{R_y^2} - \frac{z^2}{R_z^2} \right), \quad (3.17)$$

where the density goes to zero outside of the values $x = \pm R_x$, $y = \pm R_y$ and $z = \pm R_z$. For a harmonic trap the condensate distribution maintains a parabolic shape after being

released [148], and so the TOF absorption images can observe this distribution. Since the absorption images show the column density of the cloud, the above equation needs to be integrated along the imaging axis (taken as z):

$$\tilde{n}_{\text{BEC}}(x, y) = \tilde{n}_{\text{BEC}}(0) \left(1 - \frac{x^2}{R_x^2} - \frac{y^2}{R_y^2}\right)^{\frac{3}{2}}. \quad (3.18)$$

For simplicity, we approximate the thermal component of the cloud below the critical temperature to a standard Gaussian distribution (as opposed to a Bose-enhanced Gaussian [135]). This can underestimate the temperature of the thermal component [149], but is still adequate for initial estimates.

The total column density distribution that we use to model the cloud below the critical temperature is then given by

$$\tilde{n}_{\text{Tot}} = \tilde{n}_{\text{BEC}}(0) \left(1 - \frac{x^2}{R_x^2} - \frac{y^2}{R_y^2}\right)^{\frac{3}{2}} + \tilde{n}_{\text{Th}}(0) \exp\left(-\left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2}\right)\right). \quad (3.19)$$

In order to determine the condensate fraction from our images, the bimodal distribution in Equation 3.19 is first used as a fitting function to the region around the cloud. The total fitted function is then integrated to give N_{Tot} , whilst the thermal function is integrated separately to give N_{Th} . The ratio between the two results is then an estimate of the thermal fraction, and the BEC fraction can be inferred from this. Observation of our BEC is shown in Figure 3.23 along with their respective fits of the bimodal distribution. The imaging detuning is changed to -4MHz for taking pictures of the BEC in order to avoid the atoms fully absorbing the imaging light.

The BEC was observed by reducing the trap depth to below $2.4 \mu\text{K}$. By using the Gaussian component of the fitting procedure, the temperature of the thermal cloud of the bimodal distributions could be approximated and the critical temperature, T_C , estimated. For our system, $T_C \sim 230 \text{ nK}$ and an almost pure BEC with a total atom number of $2.2 \times 10^5 \pm 0.3 \times 10^5$ is achieved with a final trap depth of $1.2 \mu\text{K}$.

3.8 Experimental Sequence Summary

A summary of the atom number and temperatures achieved in the experiment is shown in Table 3.3.

Our experimental sequence is controlled by LabVIEW whilst analysis of the data is performed in Matlab, which we run on separate networked desktop computers. Our LabVIEW code generates analogue and digital waveforms that are sent out via DAQ cards

Table 3.3: Atom number and temperature after each experiment stage, following 10s loading of the MOT.

Stage	Atom Number	Temperature (μK)
MOT	1.3×10^9	-
Optical Molasses	1.3×10^9	53
Magnetic Trap Transfer	1.2×10^9	244
RF Evaporation	2.3×10^7	19.4
ODT Transfer	9.8×10^6	9.9
ODT Evaporation (Pure BEC)	2.2×10^5	$T_C \sim 0.23$

(National Instruments PCI-6723 and PCI-6229) to the experiment devices. A total of twenty digital channels and eight analogue channels are required to run the experiment. All channels have a time resolution of $100 \mu\text{s}$, except for the digital trigger for the Slave AOM (used for imaging pulses), which has a resolution of $10 \mu\text{s}$.

At the start of each experimental run, a data package is sent to our analysis computer from the LabVIEW computer over the network using the User Datagram Protocol (UDP). This package has run details such as sequence number, variable parameters and camera settings. On retrieval of this package, Matlab image analysis code initiates the Andor Luca camera and applies the desired settings. The LabVIEW code then proceeds through the experimental sequence, where the final process is triggering the camera for optical absorption imaging. After the camera is triggered and the three images are taken and saved, an image analysis procedure imports the new image files and extracts the cloud position, width and atom number, and this data passes to a plotting routine. Here, fits such as cloud temperature or lifetime can be automatically performed on the resulting data in essentially real time. The whole analysis procedure takes $< 2 \text{ s}$, as compared to $\sim 30 \text{ s}$ for the experiment to acquire the next datum. A schematic diagram of the sequence and analysis process is shown in Figure 3.24.

An example of a full experimental run is shown in Figure 3.25. In the next chapter, the SLM and its integration to the experiment will be introduced. However, for completeness, a typical sequence with the SLM included will be described here. Firstly, the motorised rotatable waveplate initialises and sets itself to its start value (maximum power in the focused ODT). The ODT is on from the beginning. The program then waits for 10 s for the MOT to load, where the 3D MOT has a magnetic field gradient of 17 Gcm^{-1} . After the MOT stage there is a short 6 ms window for all of the 2D MOT beams to be

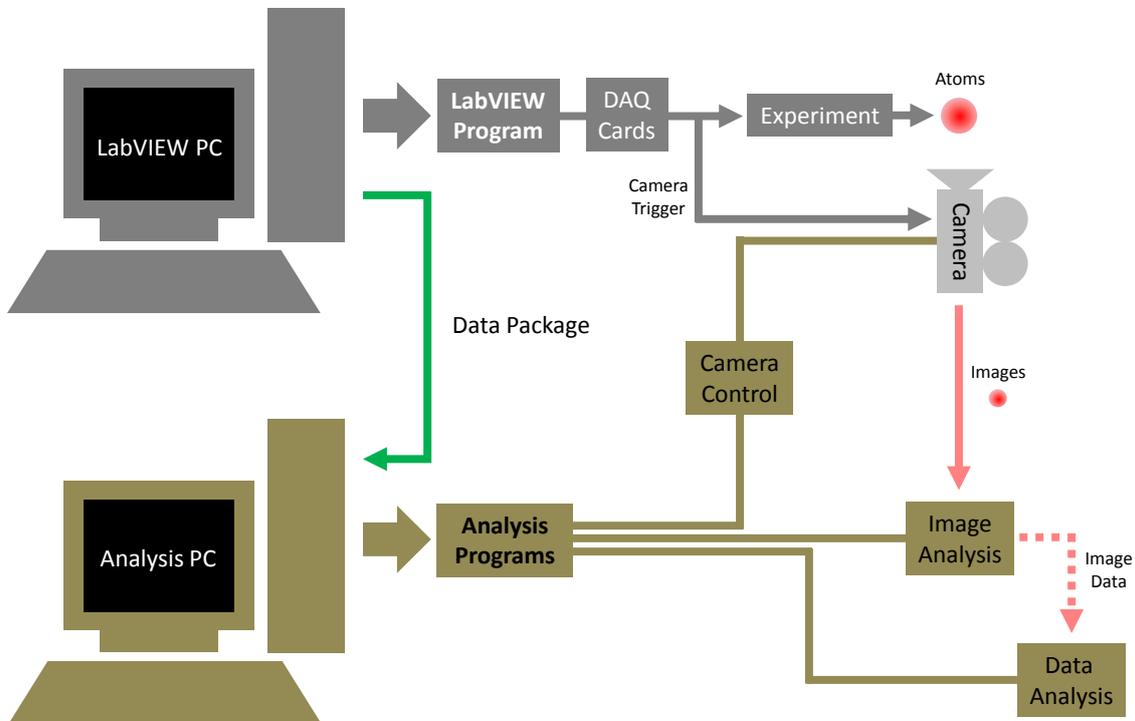


Figure 3.24: Schematic diagram of the sequence and analysis procedure.

switched off via shutters. Then a 5 ms linear ramp of the detuning of the 3D cooling beams occurs for the molasses stage, going from -20 MHz to -60 MHz. Simultaneously, the magnetic field is ramped down from 17 Gcm $^{-1}$ to zero. The molasses then occurs for 10 ms. The 3D cooling beams and the optical pumping axis shim coils are then switched off (3D repumper stays on) and the OP AOM fires a short 0.5 ms pulse of optical pumping light. The repumper shutter then closes off the last of the light on the atoms and the shim coils are switched back on along with the the magnetic trapping coils. The magnetic field is switched on to 50 Gcm $^{-1}$ via a relay and ramped to 100 Gcm $^{-1}$ in 100 ms. The atoms are then held in this trap for 10.3 s during which the RF sweep occurs, going from 17 MHz to 1.4 MHz in 10 s. The magnetic field is then relaxed from 100 Gcm $^{-1}$ to 14 Gcm $^{-1}$ in 1 s, followed by a 1 s hold time for free evaporation in the ODT to occur. The ODT is then ramped from maximum power (3.5 W on the atoms) to a desired final value (typically < 100 mW) using the motorised rotatable waveplate. During the ODT ramp, the power on the SLM increases from a minimum of 10 mW to a maximum of 200 mW. The atoms are held for a further 1 s in the combined SLM and hybrid trap before the ODT and SLM beams are switched off along with the magnetic field, whereupon the TOF imaging occurs.



Figure 3.25: Full Experimental Sequence. Shutter delays and switching times are omitted for clarity. The imaging stage has $80 \mu\text{s}$ pulses of imaging light. The SLM trap timings are used for the experiments in the next chapter, but are shown here for completeness.

Chapter 4

Ultracold Atoms in Holographic Traps

Having achieved a BEC, and worked on flexible optical potentials, the logical next step is to combine the two. The following chapter will detail the implementation of the SLM into the cold atom setup, along with the results of cold atoms in holographic traps.

4.1 Initial Cross Dipole Trap

With the experiment able to produce a BEC, the next step was to introduce the vertical ODT trapping path and SLM. The vertical path, taken from the reflection of the PBS after the motorised rotatable waveplate, first passes through a Crystal Technologies AOM (3110-197) in single pass with a maximum diffraction efficiency of 90%. The configuration for the vertical 1064 nm light is shown in Figure 4.1. The first order from the AOM is used for trapping whilst the zeroth order is blocked which allows for fast switch-off of the beam and further intensity control. Another waveplate and PBS are then used to split the light again. The power transmitted through the PBS will be used in the future for a light sheet (cylindrical beam) which will allow for two-dimensional trapping geometries with the SLM trap. At the time of writing, this path is blocked however and the light sheet will be implemented at a later date. The beam on the reflected path is then expanded with a 1:4 confocal telescope to allow for better coverage on the SLM and to increase the maximum power that can be applied to the device. For all subsequent investigations in this chapter, the AOM is set to its maximum diffraction efficiency, where the transmitted vertical ODT power is adjusted by the rotatable waveplate along with the $\lambda/2$ waveplate after the AOM.

To simplify the configuration and ensure there is plenty of power for alignment of the vertical ODT, we initially use a mirror as a place-holder for the SLM. Once alignment is achieved, irises can be inserted along the path and the mirror can be replaced with the SLM such that the beam can be realigned through the irises. After the mirror (SLM), a 40 mm lens forms a focal plane that is then imaged on to the atoms, after a dichroic, using a 40 mm 1:1 confocal telescope. For diagnosis of the trapping pattern, a flip mirror is present just after the SLM which allows the image to be formed with a separate lensing system on to a Thorlabs DCU224M CCD camera.

In order to align the vertical ODT, we initially used a large beam on the atoms. The expanding telescope comprising lenses VF3 and VF4 in Figure 4.1 was initially 1:2 (VF4 had $f = 100$ mm) to create a beam with $1/e^2$ half width of ~ 1 mm incident on the 40 mm focusing lens, thereby producing a focus at the atoms with a $1/e^2$ half width of

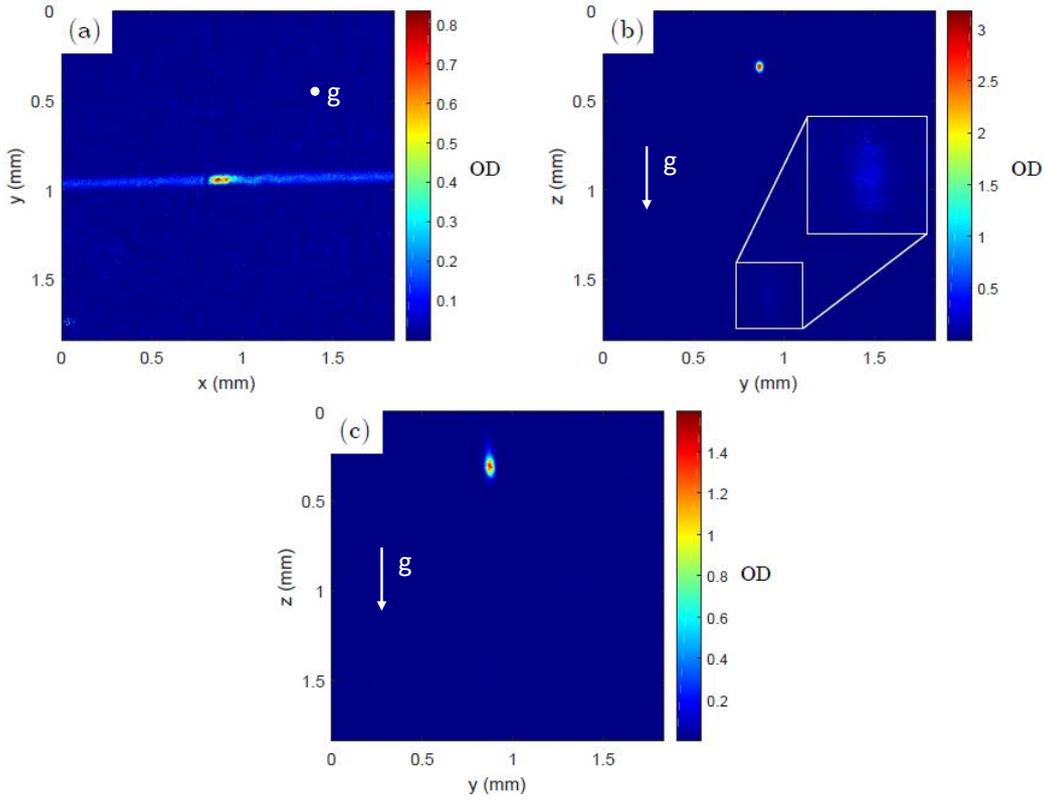


Figure 4.2: Initial cross dipole alignment with images taken after reducing the magnetic field to zero immediately after RF evaporation. (a) An in-situ view along the vertical imaging axis with the vertical ODT aligned to the centre of the horizontal ODT. Here, the vertical ODT is not focused on the horizontal ODT trap. Both ODTs have 800 mW of power. (b) A view along the horizontal imaging axis of the same run from (a) except with 0.5ms TOF. The inset shows a faint secondary cloud due to the focus of the vertical ODT. (c) A 0.5 ms TOF shot with the vertical focus now overlapped with the horizontal beam. Here the horizontal beam has 800mW and the vertical beam has 200 mW.

for in-situ shots. The two clouds due to the two ODTs can be seen in 4.2(b), where the fainter and more disperse cloud is due to the vertical ODT. The larger size of the cloud in comparison to the cloud trapped in the horizontal ODT reveals that they are far hotter in the vertical trap as expected. Figure 4.2(c) shows the vertical ODT focus aligned to the horizontal ODT with a 2 s hold time and 200 mW in the vertical beam. A total of $\sim 1.4 \times 10^5$ atoms are trapped in Figure 4.2(c).

As an initial characterisation of the cross-dipole configuration, a temperature test was undertaken. The sequence for the temperature test was a standard run (see Figure 3.25), where the ODT evaporation segment would transfer light from the horizontal ODT to the vertical ODT via the rotatable waveplate. Following this evaporation, the atoms were held in the cross dipole trap for 1 s before T was measured. For a cross-dipole configuration with 70 mW in the vertical beam and 140 mW in the horizontal beam, the temperature was estimated as $\sim 5 \mu\text{K}$, although due to the small amount of atoms and their relatively quick expansion when released from the trap, the cloud could only be recorded at short TOFs (up to 2.5 ms) before the atoms became too diffuse for the imaging to accurately record. The estimated trap depth for the vertical beam was 41 μK , whilst the horizontal

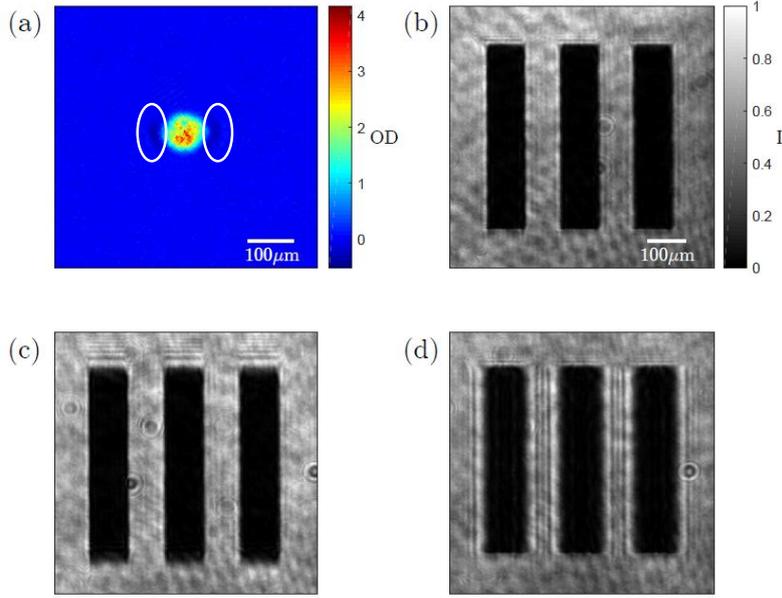


Figure 4.3: Astigmatism introduced by the dichroic when imaging. (a) An absorption image of atoms in the vertical ODT taken using the vertical imaging axis. The dark regions are present only in one axis, exhibiting an astigmatic effect (the negative OD regions are not compensated for in order to exhibit the effect more prominently). The regions are highlighted in the white ellipsoids. (b) Focused image of the USAF 1951 test target taken without a dichroic between the two telescopes. (c) Test target image taken with a dichroic between the two telescopes. Here the camera is positioned such that the horizontal axis is focused. (d) Test target image taken with a dichroic between the two telescopes. Here the camera is positioned such that the vertical axis is focused.

beam had an estimated trap depth of $4 \mu\text{K}$, giving a combined trap depth of $45 \mu\text{K}$ (the two beams are orthogonally polarised to each other so that no interference effects occur). The temperature measured is therefore expected as it corresponds to approximately $1/10$ of the trap depth which is a rule-of-thumb in labs.

Having successfully trapped atoms in the cross dipole configuration, we then increased the $1/e^2$ half width of the beam incident on the mirror/SLM to 2.2 mm to allow for a tighter focus ($\sim 6 \mu\text{m}$) and therefore more detailed holograms at the location of the atoms. This was achieved by changing VF4 in the expanding telescope to a $f = 200 \text{ mm}$ lens. The maximum intensity that the SLM can take before suffering heating damage is 5 Wcm^{-2} . For our 2.2 mm beam, using the peak intensity ($2P/\pi w^2$ where P is the beam power and w is the $1/e^2$ half width), the maximum power that could be applied to the SLM is 350 mW . In order to be conservative with the SLM and stay well within safe boundaries, we decided to ensure that a maximum of only 200 mW of power would be applied. Due to the 35% diffraction efficiency of our SLM (i.e. for a simple blazed grating, only 35% of the light would be diffracted in to the first order spot), this would mean that 70 mW would be available for a simple diffracted spot, resulting in an estimated trap depth of $\sim 164 \mu\text{K}$. Such large trap depths should therefore be plentiful for trapping evaporatively cooled atoms, even when taking in to account the typical $20\% - 40\%$ light efficiencies of

our MRAF holograms and their more complex geometries.

Images taken along the vertical axis showed an astigmatic aberration (see Figure 4.3(a)). In order to diagnose the origin of this astigmatism, an off-line imaging system with two telescopes similar to the vertical imaging axis was set up to image a USAF 1951 test target. These were a 2:3 confocal telescope (40 mm and 60 mm lenses) and a 3:4 telescope (150 mm and 200 mm lenses). It was found that the dichroic inserted in between the two telescopes (see Figure 3.14) caused the astigmatism visible in Figure 4.3, due to the image focussing through the glass tilted at 45 degrees (glass thickness is 5.0 mm). Whilst in-situ shots often had prominent aberrations due to the astigmatism, clouds could still be imaged reasonably well for short TOFs (0.5 ms–1.5 ms) and any trapping geometries created by the SLM inferred. In future a dichroic beamsplitter cube could be fashioned to replace the plate dichroic we have which would remove the effect, or a thinner dichroic could be used instead which would minimise the aberration. Also, whilst space is limited along the optical path, another compensating dichroic plate could be placed, tilted in the orthogonal plane, to correct for the aberration. Furthermore, the astigmatism could potentially be characterised and modelled (e.g. using Zernike polynomials) and compensated for after the images are taken.

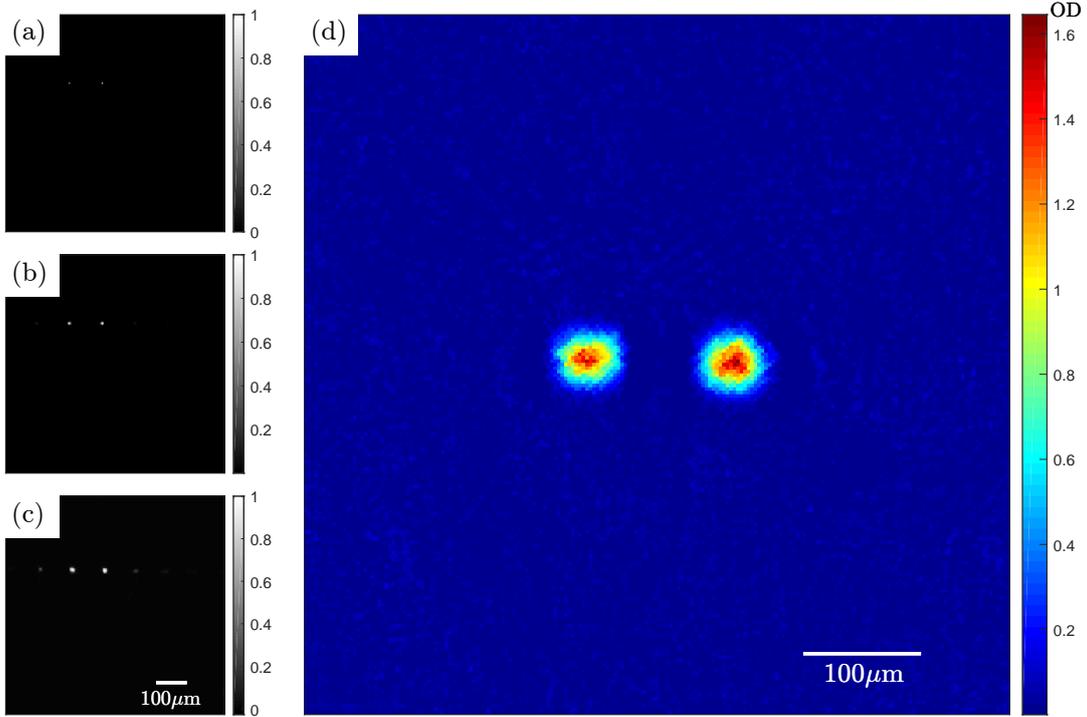


Figure 4.4: The double spot SLM trap. (a) Target pattern for the MRAF algorithm. (b) The predicted intensity result of the calculated hologram. (c) The intensity as seen on the testing camera with the calculated hologram applied to the SLM. (d) Atoms trapped by the SLM pattern, with a total atom number of $\sim 1 \times 10^5$. Image taken at 0.5 ms TOF along the vertical axis.

4.2 MRAF vs CG Optical Traps

With the cross dipole trap working, the mirror was replaced with the SLM and the path realigned such that the SLM first order passed through the irises. As an initial test, the MRAF algorithm was used to create two diffraction limited spots, as shown in Figure 4.4, spaced far enough apart so as to distinguish them in short TOF images. The two spots were spaced 30 pixels away from each other in the MRAF target, and aligned along the propagation axis of the horizontal ODT. This corresponded to a distance of $94 \mu\text{m}$ when using the 40 mm focusing lens such that each pixel in the calculation corresponds to $3.1 \mu\text{m}$. The mixing parameter (see Section 2.8) was $m = 0.7$ and the initial $1/e^2$ beam width for the calculation was 2.5 mm (it was found that this provided the smoothest experimental patterns on the camera for the diffraction limited spots, even though the actual initial beam width was measured to be 2.2 mm). As with all our patterns, the spots were shifted diagonally away from the zeroth order. The light efficiency for the pattern was 73%, leading to individual spot trap depths of $\sim 58 \mu\text{K}$ for 70 mW of diffracted power.

The first image of atoms trapped in a double spot is shown in Figure 4.4(d). The image was taken after transferring the atoms directly from the magnetic trap after RF into the SLM trap, with a constant value of 200 mW of power incident on the SLM from the start of the sequence. This gave sufficient confinement that the horizontal ODT was not required. The atoms were held in the SLM double spot trap (with a background gravity-compensating magnetic field) for 200 ms before being released, where there was a short 0.5ms TOF before the image was taken. A total of $\sim 1 \times 10^5$ atoms were trapped.

Trying to increase the trap volume of the SLM trap provides an instructive illustration of the advantage of our CG approaches over MRAF. A naive approach to increasing spot size would be to simply make a target for our hologram calculations with a larger spot size using MRAF. However, the initial guess phase which allows the spreading of light into larger regions is a Fresnel lens, and some of this phase structure is perpetuated throughout the MRAF calculation (as can be seen in the phase patterns in Figure 4.5). The result of this is to simply change the axial position of the focus of the spot, and due to the confinement achieved in the propagation axis with the vertical ODT at high power (as shown in Figure 4.2(b)), the atoms are dragged towards this new focal position. An example of this is shown in Figure 4.5, where the experimental sequences were that of a standard run (see Figure 3.25).

To avoid this effect we can use the CG calculation method and constrain the phase of the light to be flat (see Figure 4.6). The different propagation characteristics of the flat-phase (CG) and curved-phase (MRAF) beams can be seen in Figure 4.7, where the pictures of the beams were taken using the off-line imaging setup with the camera mounted on a translation stage. The MRAF pattern indeed focuses outside of the desired focal plane by $\sim 250 \mu\text{m}$, whilst the flat-phased CG spot focused where expected and maintained its shape well during propagation. The extra constraint however reduced the power of the CG pattern in comparison to the MRAF pattern, where the MRAF spot had a light efficiency of 84% and the CG spot had a light efficiency of 21%. If light efficiency is an issue, the MRAF algorithm can be used as long as propagation effects are considered, or an additional strong confinement is provided against any propagation effects, e.g. with a light sheet. However, we choose to work predominantly with CG with phase control in the remainder of this thesis, both for the reasons outlined above and to avoid the presence of optical vortices in the trapping light.

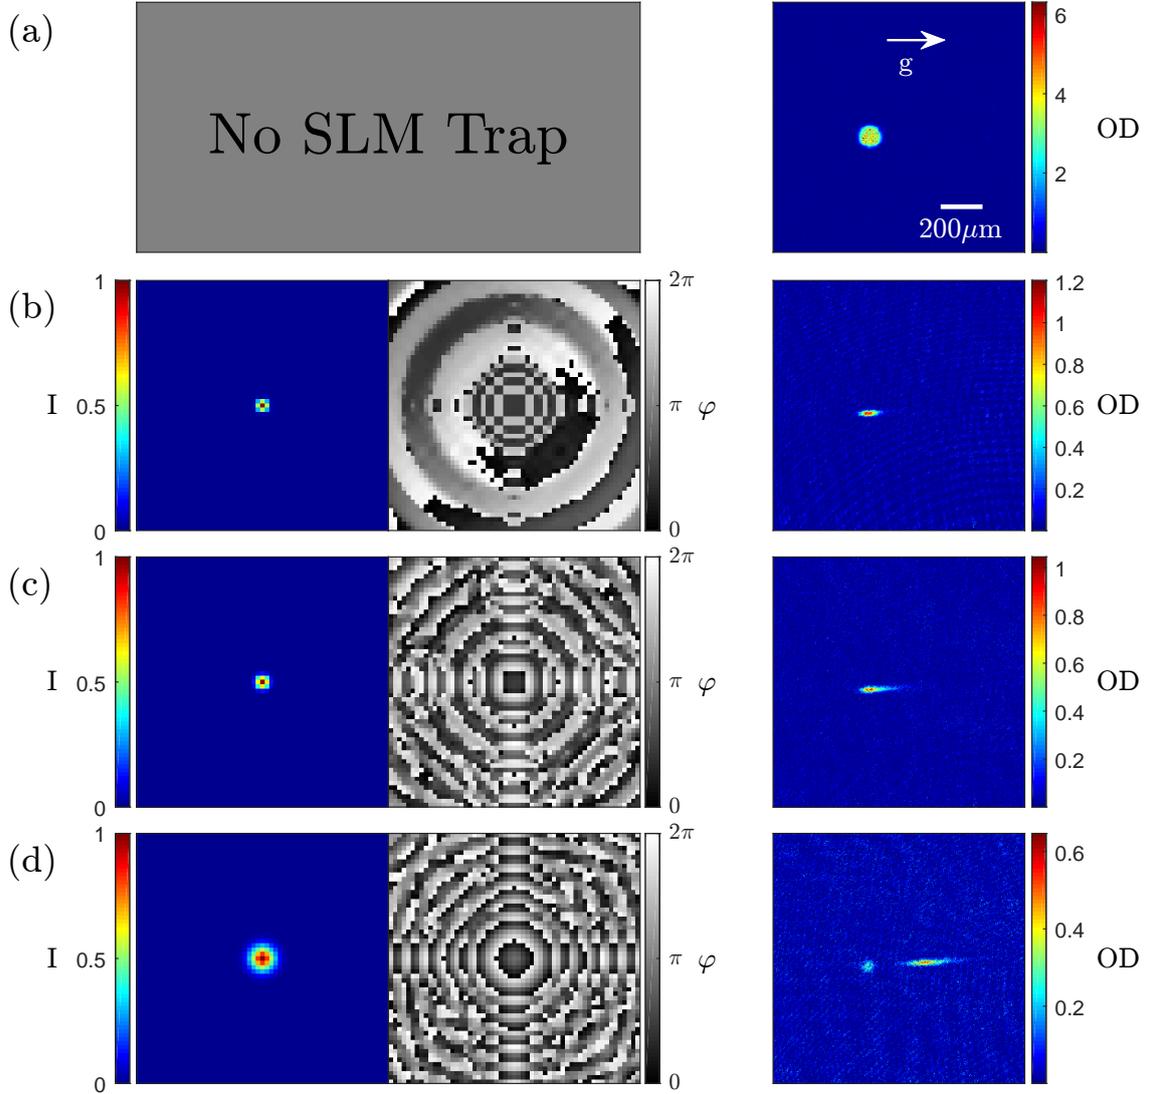


Figure 4.5: Simple spots with different widths generated by the SLM with the MRAF algorithm. The left images show the intensity and phase of the calculated results from the MRAF algorithm. A mixing parameter of $m = 0.7$ was used for all patterns. The right images show the trapped atoms, taken along the horizontal imaging path. (a) Only the horizontal ODT (image saturated). No SLM beam. (b) A gaussian spot with a $1/e^2$ width of $3.1 \mu\text{m}$. Light efficiency of 76%. (c) A gaussian spot with a $1/e^2$ width of $6.2 \mu\text{m}$. Light efficiency of 84%. (d) A gaussian spot with a $1/e^2$ width of $12.4 \mu\text{m}$. Light efficiency of 84%. The cloud on the left is trapped by the horizontal ODT combined with the SLM spot trap, while the elongated cloud on the right is trapped by the real focus of the SLM beam.

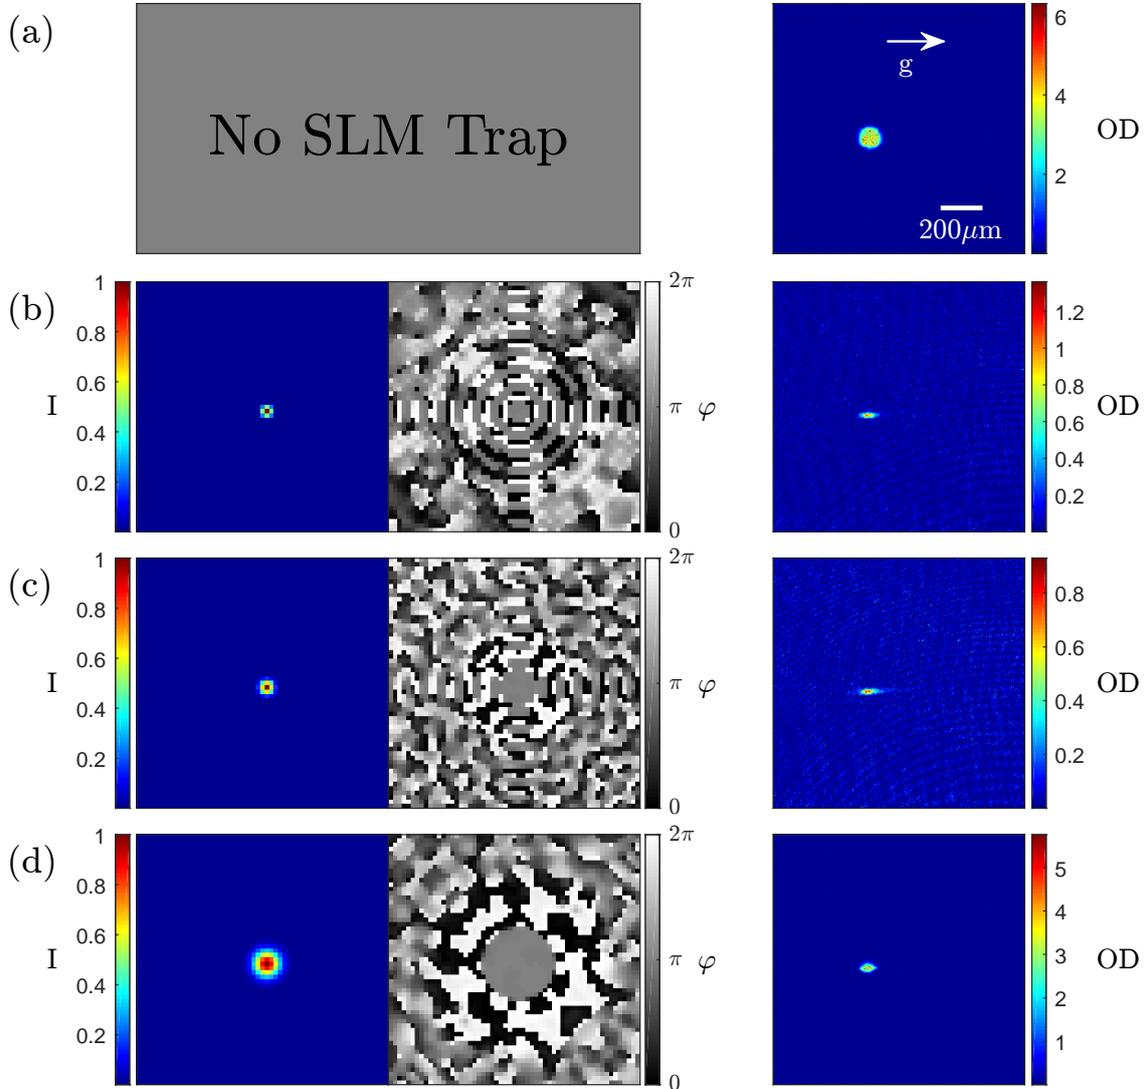


Figure 4.6: Simple spots with different widths generated by the SLM with the CG algorithm. The left image shows the intensity and phase of the calculated result from the CG algorithm. A target flat phase was used for all patterns. The right images show the trapped atoms, taken along the horizontal imaging path. (a) Only the horizontal ODT (image saturated). No SLM beam. (b) A gaussian spot with a $1/e^2$ width of $3.1 \mu\text{m}$. Light efficiency of 39%. (c) A gaussian spot with a $1/e^2$ width of $6.2 \mu\text{m}$. Light efficiency of 69%. (d) A gaussian spot with a $1/e^2$ width of $12.4 \mu\text{m}$ (image saturated). Light efficiency of 21%.

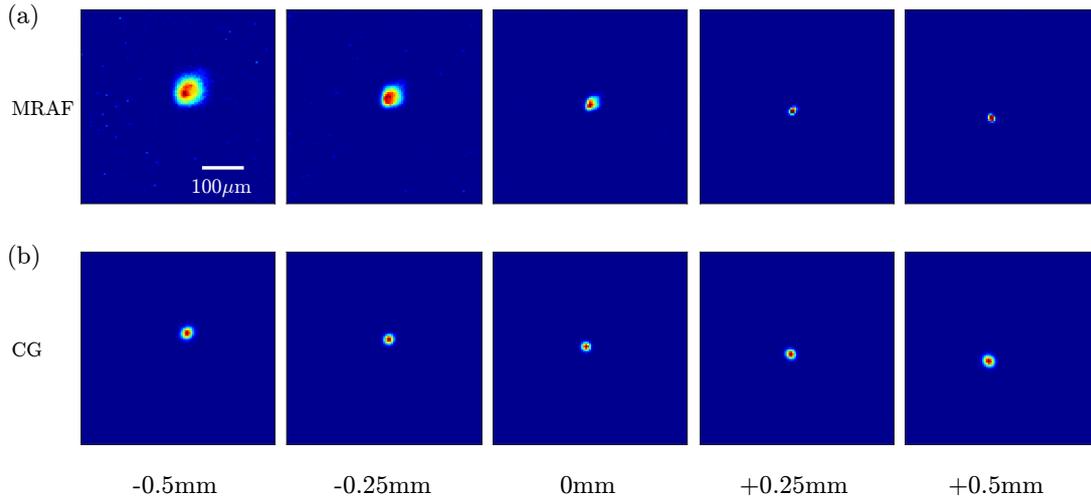


Figure 4.7: Focus test for simple spot patterns for MRAF and CG holograms taken with the off-line imaging setup. Here, 0 mm corresponds to the focal plane when a flat phase is applied to the SLM. (a) MRAF focusing. The light efficiency of the pattern is 84%. (b) CG focusing. The spot is constrained to have a flat phase, as opposed to the MRAF pattern which has an unconstrained phase. The light efficiency of the pattern is 21%.

4.3 Absence of Parametric Heating in the SLM Trap

An important investigation with regards to the initial holographic atom trapping was to determine whether the lifetime may be affected by any noise due to the SLM. Our BNS SLM has a maximum toggle rate of 1014 Hz which determines how quickly frames are loaded to the device. Whilst the response time of the liquid crystal molecules is slower than this (100 Hz), the effect can still be detected as a slight intensity fluctuation in the resulting image which can be measured with a photodiode, as seen in Figure 4.8.

The amplitude of the noise is only $\sim 1\%$ of the signal, but if the frequency of the periodic intensity fluctuations was twice that of the trapping frequency (or other multiples of twice) then the noise can resonantly drive the atomic motion in the trap, causing parametric heating and leading to losses [150, 151]. For our system, this then means that traps with trapping frequencies of 507 Hz may cause heating. However if such a phenomena occurs then the toggle rate of the SLM can be adjusted to a slower value (which will cause slightly larger intensity fluctuations but will overcome any parametric heating) or if the flexibility is allowed, the pattern can be adjusted. If neither of these options are capable of removing the heating, then a feedback loop can be used in conjunction with an AOM inserted on the SLM path to modulate the intensity and flatten off the toggle rate noise. Losses due to parametric heating were not observed in our SLM trap with 1014 Hz toggle rate.

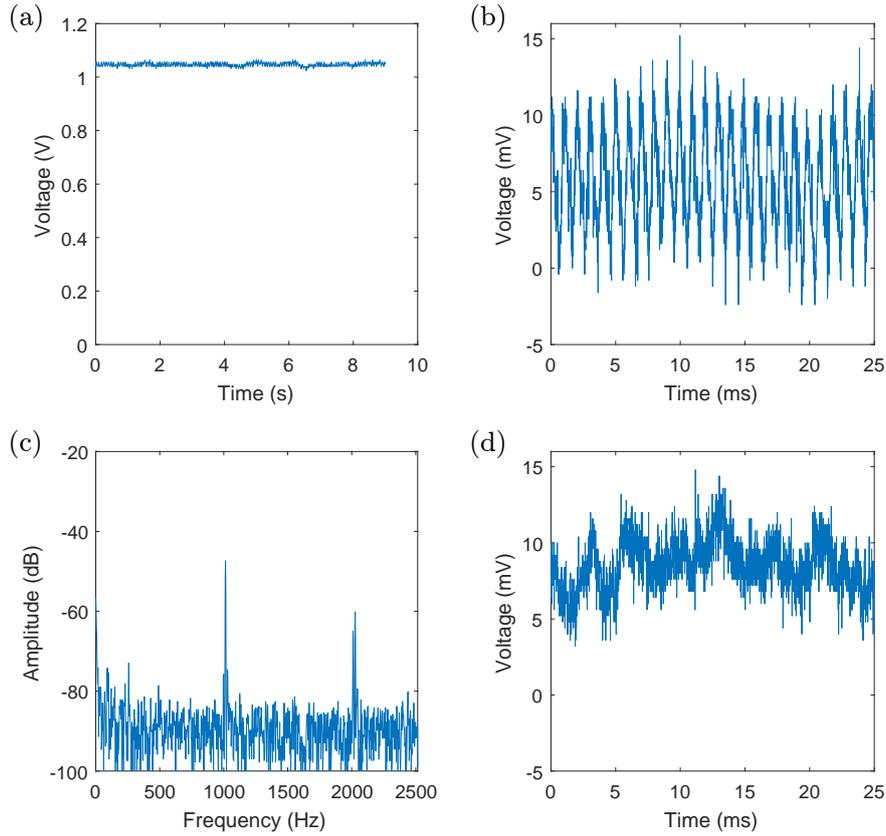


Figure 4.8: Power fluctuations for a simple spot pattern generated by the SLM and 1064 nm light. (a) Long term total power fluctuations. (b) Fast fluctuations due to the toggle rate of the SLM (the scope is set to AC). (c) FFT of the signal in (b), where the dominant peak occurs at 1014 Hz. (d) Fast fluctuations from the laser when the SLM is switched off (the scope is set to AC).

4.4 Arbitrary Trapping Patterns

To further test the capability of our SLM being used for atom trapping, a set of more arbitrary patterns were generated. The following patterns were all calculated with an initial guess phase with a curvature of $R = 2.6 \text{ mrad px}^{-2}$, linear gradient magnitude of $D = -\pi/1.43 \text{ mrad px}^{-1}$ and angular shift of $\mu = \pi/4$. The results are shown in Figure 4.9.

The first pattern was simply an extension of the double spot pattern, where four spots in a line were used instead. Since the spots were chosen to be diffraction limited, the MRAF algorithm was employed for better light efficiency. Individual spots had the same separation as the double spot pattern and were spaced along the long axis of the horizontal ODT. The resulting light efficiency was calculated to be 78%. The signal region of the calculation was quite large as this was found to improve light efficiency and kept unwanted light in the noise region away from the atoms without requiring an iris. The signal region encompassed 24% of the output plane, centred around the pattern. For atom trapping a sequence similar to that shown in Figure 3.25 was used, where 200mW incident on the SLM resulted in $\sim 13\text{mW}$ of power in each spot, corresponding to $\sim 30\mu\text{K}$ trap depth, while the horizontal ODT trap depth was $8.4\mu\text{K}$ trap depth. The atoms were held in

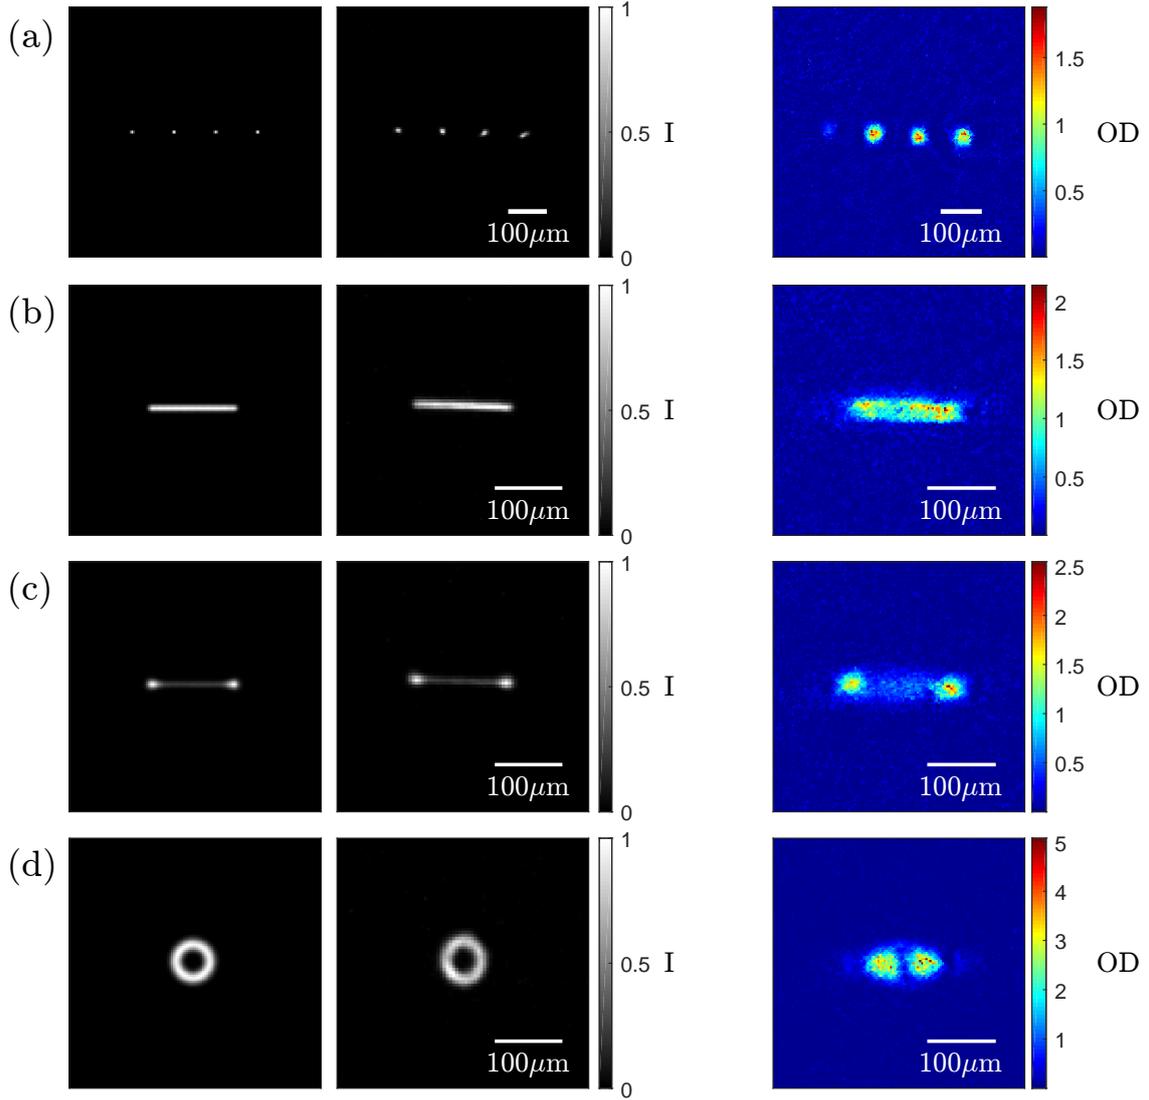


Figure 4.9: Arbitrary SLM trapping patterns. The left-most images are the calculated output intensities for four different patterns, whilst the middle images are the patterns generated by the SLM. The signal regions for these images fully encompasses the field of view. The right-most images are short TOF (0.5 ms–1 ms) pictures taken of the trapped atoms. (a) Four spots in a line. The hologram was calculated with MRAF and 9×10^4 atoms were trapped. (b) Gaussian line. Calculated with CG with a flat phase, where 1.7×10^5 atoms were trapped. (c) Dumbbell. Calculated with CG with flat phase, where 1.3×10^5 atoms were trapped. (d) Gaussian ring. Calculated with CG with flat phase, where 2.7×10^5 atoms were trapped.

this configuration for 1s before a 1ms TOF image was taken. A total of 9×10^4 atoms were trapped. However, whilst three of the spots held good portions of the total atoms (between 25% and 35%), the left-most spot held far less (5%). The results are shown in 4.9(a). An investigation of the pattern taken with the off-line imaging revealed that the left-most spot has a comparable peak intensity to the other spots (0.83 compared to 1, 0.80 and 0.87), however there are other factors along the actual trapping beam path, such

as the viewport, that may effect the pattern which the off-line imaging does not account for and the left-most spot that the atoms experience is likely underpowered.

Three continuous patterns with flat phase were generated using the CG algorithm. A Gaussian line trap with flat phase was created with length $120\mu\text{m}$ along the propagation axis of the horizontal ODT with an orthogonal $1/e^2$ half-width of $6\mu\text{m}$ using the CG method. In a similar fashion to the MRAF calculation for the four spots, slightly improved light efficiencies were observed for large signal regions, where the fidelities in the CG calculations were only slightly reduced (see Figure 4.10). The signal region encompassed 66% of the output plane around the pattern (this was used for all CG patterns in this section). The light efficiency of the calculated Gaussian line was 9.5%, resulting in 6.7mW of trapping power on the atoms and an estimated trap depth of $\sim 1.0\mu\text{K}$. The calculated fidelity (taken in the top 99% of the pattern) was 0.9985. The RMS error for the intensity (top 99%) of the resulting pattern taken on the off-line imaging was 7.1%. The horizontal ODT was reduced to 70mW corresponding to a trap depth of $2.4\mu\text{K}$, leading to a combined trap depth of $\sim 3.4\mu\text{K}$. The TOF for the image shown in Figure 4.9(b) was 0.5ms and a total of 1.7×10^5 atoms were trapped. Some fragmentation in the atomic distribution is present with the middle section of the line having a slightly reduced atom number, but the Gaussian line geometry is still prominent. In future, once the astigmatism of the imaging (section 4.1) is corrected for, fine-tuning of the atomic distribution can be achieved by adjusting the light pattern with feedback.

A pattern combining the features of spots and a line is the “dumbbell” shape which could be used for studies of transport in narrow channels between reservoirs. The dumbbell had two Gaussian spots with $1/e^2$ half-widths of $12\mu\text{m}$, connected by a channel with the same size as the Gaussian line ($120\mu\text{m}$ long, $6\mu\text{m}$ $1/e^2$ half-width). The channel peak intensity is 25% of the peak intensity of the spots. Similar to the Gaussian line, the

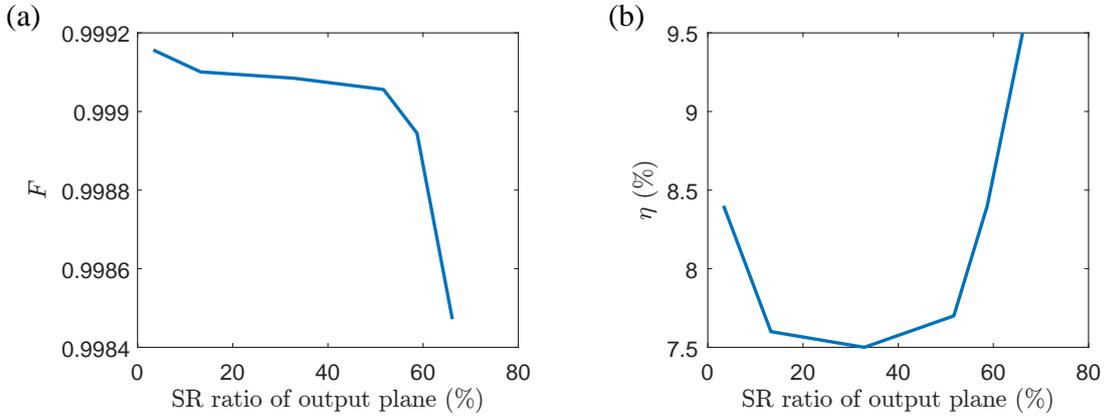


Figure 4.10: The effect of changing the SR size for the Gaussian line pattern with flat phase. An SR encompassing $> 66\%$ of the output plane would not converge to an acceptable light pattern due to the extensive constraint on the output plane. The benefit of a large SR is to ensure that light in the NR is kept far away from the atoms. (a) Fidelity. An increase from an SR which encompasses 3% of the output plane to an SR which encompasses 66% reduces the fidelity by only 0.0007. (b) Light efficiency. An increase from an SR which encompasses 3% of the output plane to an SR of 66% increases the light efficiency by 1%.

pattern was calculated with the CG method keeping a flat output phase. The calculated light efficiency was 9.1% resulting in 6.4mW of trapping power on the atoms and an estimated maximum trap depth of $\sim 1.8\mu\text{K}$ (trap depth in the spots), whilst the fidelity was 0.9994. The RMS error for the intensity (top 99%) of the resulting pattern taken on the off-line imaging was 7.0%. The horizontal ODT also had a trap depth of $2.4\mu\text{K}$ leading to a maximum combined trap depth of $\sim 4.2\mu\text{K}$. The result is shown in Figure 4.9(c), where the TOF of the absorption image was 1ms and a total number of 1.3×10^5 atoms were trapped. The two spots held $\sim 60\%$ of the total number of atoms, with a comparable number in each, whilst the channel held $\sim 40\%$. The width of the atomic cloud in the channel was similar to the two spot widths, despite the bar being thinner in the light pattern, however this is likely due to the finite TOF.

The final trap is a ring with $50\mu\text{m}$ diameter and a $1/e^2$ half-width of $12\mu\text{m}$. A fidelity of 0.9980 fidelity was achieved, along with a light efficiency of 6.1% corresponding to a trapping power of 4.3mW and a trap depth of $\sim 250\text{nK}$. The RMS error for the intensity (top 99%) of the resulting pattern taken on the off-line imaging was 11.5%. A trap depth of $3\mu\text{K}$ was used in the horizontal ODT, where the configuration lead to 2.7×10^5 atoms being held in a combined trap depth of $\sim 3.25\mu\text{K}$. The absorption image in Figure 4.9(d) was taken at 1ms TOF, which revealed that the pattern is too wide for the horizontal ODT. Since the width of the horizontal ODT is comparable to the diameter of the ring pattern, the trapping geometry is skewed such that two ‘D’ shaped atom clouds are formed, with a faint circular trail of atoms connecting them, whilst a clear dark region between the ‘D’s can be seen, indicating the central hole of the ring.

A loose estimate of the temperatures associated with the continuous geometrical patterns can be made. We image the Gaussian line pattern along the horizontal axis using 2ms TOF intervals between 6ms and 14ms, giving a temperature of $\sim 100\text{nK}$. Since the trapping power and size of the Gaussian line is comparable to the other continuous geometries attempted, it is likely that the other traps have a similar magnitude of temperature.

4.5 Feedback Improvement

Both the line of spots and Gaussian Ring patterns showed clear deviations from the expected atomic distributions, prompting an attempt to compensate for such discrepancies by adjusting the intensity distribution of the target patterns via feedback. For the four spots, the target intensity of the left-most spot was redesigned to be 4 times more intense than the other spots. This was well achieved in the MRAF calculation and although the resulting image on the camera had some deviation, it still had a prominent left-most spot, where the relative values of peak intensity from left to right were 1, 0.28, 0.24 and 0.33. The resulting light pattern had a similar light efficiency (78%) as the uncorrected four spots pattern and trapped a total of 9.5×10^4 atoms (Figure 4.11(a) shows the distribution, taken under the same conditions as Figure 4.9(a)). The distribution did see more atoms in the left-most spot, but the largest cloud was the one beside it. The ratios between the four spots now differed from each other, where the atom numbers split between the separate traps from left to right as a percentage were 25%, 35%, 18% and 22%. An even split between the four spots would therefore appear to need several steps of adjustment/feedback using the atomic distribution as the metric, most ideally suited for a feedback algorithm (similar to Section 2.3) rather than manually adjusting the trap light distribution as I have done.

For the Gaussian ring, an attempt to compensate for the horizontal ODT profile was made by subtracting a Gaussian line from across the ring which had the same estimated

width as the ODT beam. The compensation was done empirically by creating several patterns with differing Gaussian line amplitudes along with attempting varying horizontal ODT powers. The result most resembling a ring was found for a relative Gaussian line amplitude of 0.7, where the ring amplitude was 1.0. Such a pattern had a 10.4% light efficiency and a 0.9990 fidelity, leading to 7.3mW of power in the trap. The RMS error for the intensity (top 99%) of the pattern when imaged on the camera was 7.6%. The horizontal ODT was reduced to $3\mu\text{K}$, similar to the uncorrected ring and a total of 2.8×10^5 atoms were trapped. The result is shown in Figure 4.11(b). Whilst the ring shape is more prominent, so too are residual atoms still trapped in the horizontal ODT. Only a slight hole can be discerned at the centre of the ring, whilst some fragmentation is occurring across the geometry. However, when the light sheet is implemented in future, these issues regarding the ring trap (and any other more complex geometries) will no longer be apparent as the large width offered by the cylindrical beam will allow for a more uniform trapping profile in both axes.

In summary, a variety of both continuous and discrete geometries capable of trapping $\sim 1 \times 10^5$ atoms at $\sim 100n\text{K}$ have been created using a phase-only SLM. Compensation of the atomic distribution using basic feedback on the absorption image of the atoms has been shown to work, where a more complete feedback method could yield even smoother continuous geometries. The patterns created with both phase and amplitude control using the conjugate gradient method was found to be the most effective for trapping due to the ability to impose a flat phase on the light, removing out-of-plane focussing effects. The phase-only SLM and the techniques developed for hologram generation in this work have therefore been shown as an effective approach for ultracold atom trapping.

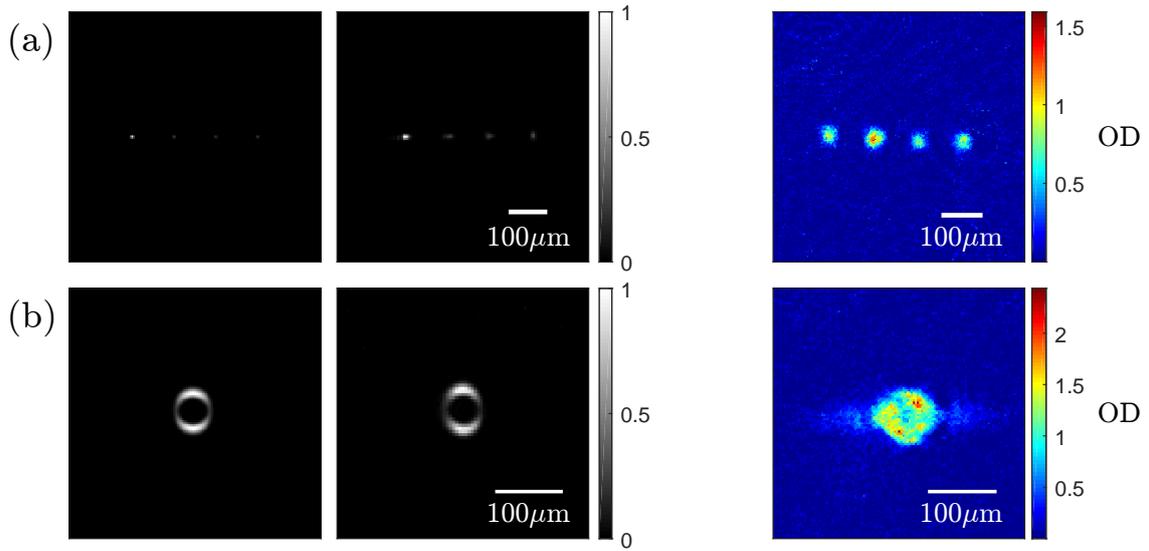


Figure 4.11: Simple feedback examples for SLM atom trapping. The left-most images are the calculated output intensities, whilst the middle images are the patterns generated by the SLM. The signal regions for these images fully encompasses the field of view. The right-most images are 1ms TOF pictures taken of the trapped atoms. (a) The four spots line pattern which has the left-most spot with quadruple the intensity of the other spots. The hologram was calculated with MRAF and 9.5×10^4 atoms were trapped. (b) The gaussian ring pattern with a 0.7 amplitude $50\mu\text{m}$ $1/e^2$ half-width horizontal Gaussian subtracting from one axis. Calculated with CG with flat phase, where 2.8×10^5 atoms were trapped.

Chapter 5

Conclusion

5.1 Summary

To summarise, bespoke trapping and manipulation of ultracold atoms offers a rich vein of investigations for quantum systems and the work in this thesis has shown the creation of the first BEC in St Andrews, where the ultracold atom experiment has been made in conjunction with a phase-only SLM, exhibiting flexibility in both discrete and continuous trapping geometries. The development of compact and flexible holography techniques in multi-wavelength patterns using MRAF, along with high-fidelity phase and amplitude control using our CG method, have been demonstrated both in calculation and experimentally. The latter development allows the calculation of, to our knowledge, the most accurate holographic optical traps, with errors reduced by a factor of 20 compared to the previous state of the art [125]. The output phase manipulation afforded by the CG method was shown to be particularly desirable for optimal control of the holographic traps when trapping atoms in our experiment.

5.2 Future Prospects

For the near future, a correction to the vertical imaging axis is required in order to do in-situ imaging of the atoms in the SLM trap. Such accurate imaging would then afford feedback trap compensation, similar to the feedback method in Section 2.3 and 2.4.3, where the atom distributions are used as the metric for the trap corrections instead of the light itself.

With appropriate imaging and corrections, the efficient evaporation method outlined in [56] can be attempted, where power law traps generated by the SLM can be compared against the standard Gaussian profile in terms of evaporatively cooling the atoms. Beyond this however, it would be desirable to implement the light sheet. For example, future investigations associated with previous theoretical research undertaken by the group, such as investigations of the topological Kondo effect [75] and rotation sensing [55], require patterns that extend beyond the width of the horizontal ODT, and therefore implementation of the light sheet will be necessary here. Investigations of dynamic trapping of the atoms using our SLM will also be required for the rotation scheme in [55].

Regarding holography development, the high fidelity patterns capable by the conjugate gradient minimisation technique, as reported in Section 2.5, may provide an avenue for smooth bespoke three-dimensional optical traps. The additional phase control can be crucial for controlling the light in three-dimensions, as the graphene propagation result from Figure 2.30 and the spot width investigation from Figure 4.7 highlights. Thus,

developing the CG code further in this direction may well be a promising future prospect. Holography to create three-dimensional traps for particles has been developed before, both in discrete crystal lattice geometries [152] and continuous tubes [153]. Atoms taken from a MOT have also recently been trapped in three dimensional discrete arrays using an SLM [154], but no bespoke continuous three-dimensional traps for ultracold atoms have been achieved. Furthermore, configurations with an SLM are capable of manipulating the polarisation of a light beam as well [155]. If a high degree of control over all of the components of light (amplitude, phase and polarisation) are available, this would yield smooth and fully arbitrary three-dimensional patterns, where the CG code could be used to create holograms capable of such control with a single SLM. Our multi-wavelength technique (Section 2.4) could also be incorporated in such future schemes.

Bibliography

- [1] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, “Observation of Bose-Einstein condensation in a dilute atomic vapor,” *Science*, vol. 269, p. 5221, 1995.
- [2] I. Bloch, J. Dalibard, and W. Zwerger, “Many-body physics with ultracold gases,” *Reviews of Modern Physics*, vol. 80, p. 885, 2008.
- [3] C. Cohen-Tannoudji and D. Gury-Odelin, *Advances in Atomic Physics*. World Scientific, 2011.
- [4] T. Langen, R. Geiger, and J. Schmiedmayer, “Ultracold atoms out of equilibrium,” *Annual Review of Condensed Matter Physics*, vol. 6, p. 201, 2015.
- [5] C. Gross and I. Bloch, “Quantum simulations with ultracold atoms in optical lattices,” *Science*, vol. 357, p. 995, 2017.
- [6] A. Einstein, “Quantentheorie des einatomigen idealen Gases,” *Königliche Preußische Akademie der Wissenschaften. Sitzungsberichte*, p. 261, 1924.
- [7] A. Einstein, “Quantentheorie des einatomigen idealen Gases. Zweite Abhandlung,” *Sitzungsberichte der Preußischen Akademie der Wissenschaften (Berlin). Physikalisch-mathematische Klasse*, p. 3, 1925.
- [8] A. Einstein, “Zur Quantentheorie des idealen Gases,” *Sitzungsberichte der Preußischen Akademie der Wissenschaften (Berlin). Physikalisch-mathematische Klasse*, p. 18, 1925.
- [9] S. N. Bose, “Plancks gesetz und lichtquantenhypothese,” *Zeitschrift für Physik*, vol. 26, p. 178, 1924.
- [10] J. R. Anglin and W. Ketterle, “Bose–Einstein condensation of atomic gases,” *Nature*, vol. 416, p. 211, 2002.
- [11] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard, “Trapping of neutral sodium atoms with radiation pressure,” *Physical Review Letters*, vol. 59, p. 2631, 1987.
- [12] S. Stellmer, B. Pasquiou, R. Grimm, and F. Schreck, “Laser cooling to quantum degeneracy,” *Physical Review Letters*, vol. 110, p. 263003, 2013.
- [13] J. Hu, A. Urvoy, Z. Vendeiro, V. Crépel, W. Chen, and V. Vuletić, “Creation of a Bose-condensed gas of ^{87}Rb by laser cooling,” *Science*, vol. 358, p. 1078, 2017.

- [14] M. R. Andrews, C. G. Townsend, H.-J. Miesner, D. S. Durfee, D. M. Kurn, and W. Ketterle, “Observation of interference between two Bose condensates,” *Science*, vol. 275, p. 637, 1997.
- [15] K. W. Madison, F. Chevy, W. Wohlleben, and J. Dalibard, “Vortex formation in a stirred Bose-Einstein condensate,” *Physical Review Letters*, vol. 84, p. 806, 2000.
- [16] G. Roati, C. D’Errico, L. Fallani, M. Fattori, C. Fort, M. Zaccanti, G. Modugno, M. Modugno, and M. Inguscio, “Anderson localization of a non-interacting Bose–Einstein condensate,” *Nature*, vol. 453, p. 895, 2008.
- [17] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, “Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms,” *Nature*, vol. 415, 2002.
- [18] O. Lahav, A. Itah, A. Blumkin, C. Gordon, S. Rinott, A. Zayats, and J. Steinhauer, “Realization of a sonic black hole analog in a Bose-Einstein condensate,” *Physical Review Letters*, vol. 105, p. 240401, 2010.
- [19] B. DeMarco and D. S. Jin, “Onset of Fermi degeneracy in a trapped atomic gas,” *Science*, vol. 285, p. 1703, 1999.
- [20] M. Greiner, C. A. Regal, and D. S. Jin, “Emergence of a molecular Bose–Einstein condensate from a Fermi gas,” *Nature*, vol. 426, p. 537, 2003.
- [21] M. W. Zwierlein, J. R. Abo-Shaeer, A. Schirotzek, C. H. Schunck, and W. Ketterle, “Vortices and superfluidity in a strongly interacting Fermi gas,” *Nature*, vol. 435, p. 1047, 2005.
- [22] 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*, vol. 467, p. 68, 2010.
- [23] A. Mazurenko, C. S. Chiu, G. Ji, M. F. Parsons, M. Kanász-Nagy, R. Schmidt, F. Grusdt, E. Demler, D. Greif, and M. Greiner, “A cold-atom Fermi-Hubbard antiferromagnet,” *Nature*, vol. 545, p. 462, May 2017.
- [24] C. J. Foot, *Atomic Physics*. Oxford University Press, 2005.
- [25] M. Gildemeister, E. Nugent, B. E. Sherlock, M. Kubasik, B. T. Sheard, and C. J. Foot, “Trapping ultracold atoms in a time-averaged adiabatic potential,” *Physical Review A*, vol. 81, p. 031402(R), 2010.
- [26] B. E. Sherlock, M. Gildemeister, E. Owen, E. Nugent, and C. J. Foot, “Time-averaged adiabatic ring potential for ultracold atoms,” *Physical Review A*, vol. 83, p. 043408, 2011.
- [27] M. Gildemeister, B. E. Sherlock, and C. J. Foot, “Techniques to cool and rotate Bose-Einstein condensates in time-averaged adiabatic potentials,” *Physical Review A*, vol. 85, p. 053401, 2012.
- [28] T. L. Harte, E. Bentine, K. Luksch, A. J. Barker, D. Trypogeorgos, B. Yuen, and C. J. Foot, “Ultracold atoms in multiple-radiofrequency dressed adiabatic potentials,” *arXiv:1706.01491*, 2017.

- [29] M. Keil, O. Amit, S. Zhou, D. Groswasser, Y. Japha, and R. Folman, “Fifteen years of cold matter on the atom chip: promise, realizations, and prospects,” *Journal of Modern Optics*, vol. 63, p. 1840, 2016.
- [30] R. Folman, P. Krüger, D. Cassettari, B. Hessmo, T. Maier, and J. Schmiedmayer, “Controlling cold atoms using nanofabricated surfaces: Atom chips,” *Physical Review Letters*, vol. 84, p. 4749, 2000.
- [31] M. Horikoshi and K. Nakagawa, “Atom chip based fast production of Bose–Einstein condensate,” *Applied Physics B*, vol. 82, p. 363, 2006.
- [32] L. Amico, G. Birkl, M. Boshier, and L.-C. Kwek, “Focus on atomtronics-enabled quantum technologies,” *New Journal of Physics*, vol. 19, p. 020201, 2017.
- [33] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, “One-dimensional lattice of permanent magnetic microtraps for ultracold atoms on an atom chip,” *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 41, p. 065301, 2008.
- [34] V. Y. F. Leung, D. R. M. Pijn, H. Schlatter, L. Torralbo-Campo, A. L. L. Rooij, G. B. Mulder, J. Naber, M. L. Soudijn, A. Tauschinsky, C. Abarbanel, B. Hadad, E. Golan, R. Folman, and R. J. C. Spreeuw, “Magnetic-film atom chip with 10 μm period lattices of microtraps for quantum information science with Rydberg atoms,” *Review of Scientific Instruments*, vol. 85, p. 053102, 2014.
- [35] Y. Shin, C. Sanner, G.-B. Jo, T. A. Pasquini, M. Saba, W. Ketterle, and D. E. Pritchard, “Interference of Bose-Einstein condensates split with an atom chip,” *Physical Review A*, vol. 72, p. 021604(R), 2005.
- [36] G. Jo, J. Guzman, C. K. Thomas, P. Hosur, A. Vishwanath, and D. M. Stamper-Kurn, “Ultracold atoms in a tunable optical kagome lattice,” *Physical Review Letters*, vol. 108, p. 045305, 2012.
- [37] P. Soltan-Panahi, J. Struck, P. Hauke, A. Bick, W. Plenkers, G. Meineke, C. Becker, P. Windpassinger, M. Lewenstein, and K. Sengstock, “Multi-component quantum gases in spin-dependent hexagonal lattices,” *Nature Physics*, vol. 7, p. 434, 2011.
- [38] G. Wirth, M. Ölschläger, and A. Hemmerich, “Evidence for orbital superfluidity in the p-band of a bipartite optical square lattice,” *Nature Physics*, vol. 7, p. 147, 2010.
- [39] J. Sebby-Strabley, M. Anderlini, P. S. Jessen, and J. V. Porto, “Lattice of double wells for manipulating pairs of cold atoms,” *Physical Review A*, vol. 73, p. 033605, 2006.
- [40] L. Tarruell, D. Greif, T. Uehlinger, G. Jotzu, and T. Esslinger, “Creating, moving and merging Dirac points with a Fermi gas in a tunable honeycomb lattice,” *Nature*, vol. 483, p. 302, 2012.
- [41] L. W. Cheuk, M. A. Nichols, K. R. Lawrence, M. Okan, H. Zhang, E. Khatami, N. Trivedi, T. Paiva, M. Rigol, and M. W. Zwierlein, “Observation of spatial charge and spin correlations in the 2D Fermi-Hubbard model,” *Science*, vol. 353, p. 1260, 2016.

- [42] K. Henderson, C. Ryu, C. MacCormick, and M. G. Boshier, “Experimental demonstration of painting arbitrary and dynamic potentials for Bose-Einstein condensates,” *New Journal of Physics*, vol. 11, p. 043030, 2009.
- [43] C. Ryu, P. W. Blackburn, A. A. Blinova, and M. G. Boshier, “Experimental realization of Josephson junctions for an atom squid,” *Physical Review Letters*, vol. 111, p. 205301, 2013.
- [44] M. R. Douglas, “Lifetime estimates and unique failure mechanisms of the digital micromirror device (dmd),” *International Reliability Physics Proceedings, IEEE 98CH36173*, p. 9, 1998.
- [45] G. Gauthier, I. Lenton, N. M. Parry, M. Baker, M. J. Davis, H. Rubinsztein-Dunlop, and T. W. Neely, “Direct imaging of a digital-micromirror device for configurable microscopic optical potentials,” *Optica*, vol. 3, p. 1136, 2016.
- [46] J. Liang, R. N. Kohn, M. F. Becker, and D. J. Heinzen, “High-precision laser beam shaping using a binary-amplitude spatial light modulator,” *International Reliability Physics Proceedings, IEEE 98CH36173*, vol. 49, p. 1323, 2010.
- [47] T. Watanabe, M. Fujii, Y. Watanabe, N. Toyama, and Y. Iketaki, “Generation of a doughnut-shaped beam using a spiral phase plate,” *Review of Scientific Instruments*, vol. 75, p. 5131, 2004.
- [48] V. A. Henderson, P. F. Griffin, E. Riis, and A. S. Arnold, “Comparative simulations of Fresnel holography methods for atomic waveguides,” *New Journal of Physics*, vol. 18, p. 025007, 2016.
- [49] A. Ramanathan, K. C. Wright, S. R. Muniz, M. Zelan, W. T. Hill, C. J. Lobb, K. Helmerson, W. D. Phillips, and G. K. Campbell, “Superflow in a toroidal Bose-Einstein condensate: An atom circuit with a tunable-weak link,” *Physical Review Letters*, vol. 106, p. 130401, Apr 2011.
- [50] J. G. Lee and W. T. Hill, “Spatial shaping for generating arbitrary optical dipole traps for ultracold degenerate gases,” *Review of Scientific Instruments*, vol. 85, p. 103106, 2014.
- [51] T. A. Haase, D. H. White, D. J. Brown, I. Herrera, and M. D. Hoogerland, “A versatile apparatus for two-dimensional atomtronic quantum simulation,” *Review of Scientific Instruments*, vol. 88, p. 113102, 2017.
- [52] 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,” *Journal of the Optical Society of America B*, vol. 21, p. 1889, 2004.
- [53] 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,” *Physical Review A*, vol. 73, p. 031402(R), Mar 2006.
- [54] S. Franke-Arnold, J. Leach, M. J. Padgett, V. E. Lembessis, D. Ellinas, A. J. Wright, J. M. Girkin, P. Öhberg, and A. S. Arnold, “Optical ferris wheel for ultracold atoms,” *Optics Express*, vol. 15, p. 8619, 2007.

- [55] G. D. Bruce, J. Mayoh, G. Smirne, L. Torralbo-Campo, and D. Cassettari, “Smooth, holographically generated ring trap for the investigation of superfluidity in ultracold atoms,” *Physica Scripta*, vol. T143, p. 014008, Feb 2011.
- [56] G. D. Bruce, S. L. Bromley, G. Smirne, L. Torralbo-Campo, and D. Cassettari, “Holographic power-law traps for the efficient production of Bose-Einstein condensates,” *Physical Review A*, vol. 84, p. 053410, Nov 2011.
- [57] A. L. Gaunt and Z. Hadzibabic, “Robust digital holography for ultracold atom trapping,” *Scientific Reports*, vol. 2, p. 721, Oct 2012.
- [58] F. Nogrette, H. Labuhn, S. Ravets, D. Barredo, L. Béguin, A. Vernier, T. Lahaye, and A. Browaeys, “Single-atom trapping in holographic 2D arrays of microtraps with arbitrary geometries,” *Physical Review X*, vol. 4, p. 021034, 2014.
- [59] G. D. Bruce, M. Y. H. Johnson, E. Cormack, D. A. W. Richards, J. Mayoh, and D. Cassettari, “Feedback-enhanced algorithm for aberration correction of holographic atom traps,” *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 48, p. 115303, May 2015.
- [60] C. Chin, R. Grimm, P. Julienne, and E. Tiesinga, “Feshbach resonances in ultracold gases,” *Reviews of Modern Physics*, vol. 82, p. 1225, 2010.
- [61] M. Theis, G. Thalhammer, K. Winkler, M. Hellwig, G. Ruff, R. Grimm, and J. H. Denschlag, “Tuning the scattering length with an optically induced Feshbach resonance,” *Physical Review Letters*, vol. 93, p. 123001, 2004.
- [62] M. Randeria and E. Taylor, “Crossover from Bardeen-Cooper-Schrieffer to Bose-Einstein condensation and the unitary Fermi gas,” *Annual Review of Condensed Matter Physics*, vol. 5, p. 209, 2014.
- [63] K. M. Jones, E. Tiesinga, P. D. Lett, and P. S. Julienne, “Ultracold photoassociation spectroscopy: Long-range molecules and atomic scattering,” *Reviews of Modern Physics*, vol. 78, p. 483, 2006.
- [64] M. F. Andersen, C. Ryu, P. Cladé, V. Natarajan, A. Vaziri, K. Helmerson, and W. D. Phillips, “Quantized rotation of atoms from photons with orbital angular momentum,” *Physical Review Letters*, vol. 97, p. 170406, 2006.
- [65] K. C. Wright, L. S. Leslie, and N. P. Bigelow, “Optical control of the internal and external angular momentum of a Bose-Einstein condensate,” *Physical Review A*, vol. 77, p. 041601(R), 2008.
- [66] F. Jendrzejewski, A. Bernard, K. Müller, P. Cheinet, V. Josse, M. Piraud, L. Pezzé, L. Sanchez-Palencia, A. Aspect, and P. Bouyer, “Three-dimensional localization of ultracold atoms in an optical disordered potential,” *Nature Physics*, vol. 8, p. 398, 2012.
- [67] J. Choi, S. Hild, J. Zeiher, P. Schauß, A. Rubio-Abadal, T. Yefsah, V. Khemani, D. A. Huse, I. Bloch, and C. Gross, “Exploring the many-body localization transition in two dimensions,” *Science*, vol. 352, p. 1547, 2016.
- [68] C. Weitenberg, M. Endres, J. F. Sherson, M. Cheneau, P. Schauß, T. Fukuhara, I. Bloch, and S. Kuhr, “Single-spin addressing in an atomic Mott insulator,” *Nature*, vol. 471, p. 319, 2011.

- [69] T. Fukuhara, A. Kantian, M. Endres, M. Cheneau, P. Schauß, S. Hild, D. Bellem, U. Schollwöck, T. Giamarchi, C. Gross, I. Bloch, and S. Kuhr, “Quantum dynamics of a mobile spin impurity,” *Nature Physics*, vol. 9, p. 235, 2013.
- [70] J. Dalibard, F. Gerbier, G. Juzeliunas, and P. Öhberg, “Artificial gauge potentials for neutral atoms,” *Review of Modern Physics*, vol. 83, p. 1523, 2011.
- [71] M. Aidelsburger, M. Atala, M. Lohse, J. T. Barreiro, B. Paredes, and I. Bloch, “Realization of the Hofstadter Hamiltonian with ultracold atoms in optical lattices,” *Physical Review Letters*, vol. 111, p. 185301, 2013.
- [72] H. Miyake, G. A. Siviloglou, C. J. Kennedy, W. C. Burton, and W. Ketterle, “Realizing the Harper Hamiltonian with laser-assisted tunneling in optical lattices,” *Physical Review Letters*, vol. 111, p. 185302, 2013.
- [73] J. Denschlag, J. E. Simsarian, D. L. Feder, C. W. Clark, L. A. Collins, J. Cubizolles, L. Deng, E. W. Hagley, K. Helmerson, W. P. Reinhardt, S. L. Rolston, B. I. Schneider, and W. D. Phillips, “Generating solitons by phase engineering of a Bose-Einstein condensate,” *Science*, vol. 287, p. 97, 2000.
- [74] T. Yefsah, A. T. Sommer, M. J. H. Ku, L. W. Cheuk, W. Ji, W. S. Bakr, and M. W. Zwierlein, “Heavy solitons in a fermionic superfluid,” *Nature*, vol. 499, p. 426, 2013.
- [75] F. Buccheri, G. D. Bruce, A. Trombettoni, D. Cassettari, H. Babujian, V. E. Korepin, and P. Sodano, “Holographic optical traps for atom-based topological Kondo devices,” *New Journal of Physics*, vol. 18, p. 075012, July 2016.
- [76] R. W. Simmonds, A. Marchenkov, E. Hoskinson, J. C. Davis, and R. E. Packard, “Quantum interference of superfluid ^3He ,” *Nature*, vol. 412, p. 55, 2001.
- [77] G. Thalhammer, R. W. Bowman, G. D. Love, M. J. Padgett, and M. Ritsch-Marte, “Speeding up liquid crystal SLMs using overdrive with phase change reduction,” *Optics Express*, vol. 21, p. 1779, Jan 2013.
- [78] J. W. Goodman, *Introduction to Fourier Optics*. McGraw-Hill, 1996.
- [79] M. Pasienski and B. DeMarco, “A high-accuracy algorithm for designing arbitrary holographic atom traps,” *Optics Express*, vol. 16, p. 2176, Feb 2008.
- [80] D. Bowman, P. Ireland, G. D. Bruce, and D. Cassettari, “Multi-wavelength holography with a single spatial light modulator for ultracold atom experiments,” *Optics Express*, vol. 23, p. 8365, Apr 2015.
- [81] D. Bowman, T. L. Harte, V. Chardonnet, C. D. Groot, S. J. Denny, G. L. Goc, M. Anderson, P. Ireland, D. Cassettari, and G. D. Bruce, “High-fidelity phase and amplitude control of phase-only computer generated holograms using conjugate gradient minimisation,” *Optics Express*, vol. 25, p. 11692, May 2017.
- [82] R. W. Gerbergh and W. O. Saxton, “A practical algorithm for the determination of phase from image and diffraction plane pictures,” *Optik*, vol. 35, p. 237, 1972.
- [83] C. Henkel, P. Krüger, R. Folman, and J. Schmiedmayer, “Fundamental limits for coherent manipulation on atom chips,” *Applied Physics B*, vol. 76, p. 173, 2003.

- [84] M. Clark and R. Smith, “A direct-search method for the computer design of holograms,” *Optics Communications*, vol. 124, p. 150, Feb 1996.
- [85] M. Mitchell, *An Introduction to Genetic Algorithms*. MIT Press, 1998.
- [86] C. MacArthur, “Algorithmic optimisation of holographic optical traps,” Master’s thesis, University of St Andrews, April 2015.
- [87] T. Harte, “A conjugate gradient minimisation approach to generating holographic traps for ultracold atoms,” Master’s thesis, University of St Andrews, May 2012.
- [88] T. Harte, G. D. Bruce, J. Keeling, and D. Cassettari, “Conjugate gradient minimisation approach to generating holographic traps for ultracold atoms,” *Optics Express*, vol. 22, p. 26548, Nov 2014.
- [89] J. R. Shewchuk, “An introduction to the conjugate gradient method without the agonizing pain.” <http://www.cs.cmu.edu/~quake-papers/painless-conjugate-gradient.pdf>, 1994.
- [90] M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos, “Iterative minimization techniques for ab initio total-energy calculations: molecular dynamics and conjugate gradients,” *Review of Modern Physics*, vol. 64, p. 1045, October 1992.
- [91] B. Temel, G. Mills, and H. Metiu, “Inelastic scattering with Chebyshev polynomials and preconditioned conjugate gradient minimization,” *Journal of Physical Chemistry A*, vol. 112, p. 2728, February 2008.
- [92] S. J. Thomson, L. S. Walker, T. L. Harte, and G. D. Bruce, “Measuring the Edwards-Anderson order parameter of the Bose glass: A quantum gas microscope approach,” *Physical Review A*, vol. 94, p. 051601(R), 2016.
- [93] P. Senthilkumarana, F. Wyrowski, and H. Schimmel, “Vortex stagnation problem in iterative Fourier transform algorithms,” *Optics and Lasers in Engineering*, vol. 43, p. 43, 2005.
- [94] D. Donnelly and B. Rust, “The fast fourier transform for experimentalists, part i: Concepts,” *Computing in Science & Engineering*, vol. 7, p. 80, 2005.
- [95] H. Pang, J. Wang, A. Cao, and Q. Deng, “High-accuracy method for holographic image projection with suppressed speckle noise,” *Optics Express*, vol. 24, p. 22766, 2016.
- [96] T. Shimobaba and T. Ito, “Random phase-free computer-generated hologram,” *Optics Express*, vol. 23, p. 9549, 2015.
- [97] W. J. Smith, *Modern Optical Engineering: The Design of Optical Systems*. McGraw-Hill, 3rd ed., 2000.
- [98] R. W. Bowman, A. J. Wright, and M. J. Padgett, “An SLM-based Shack–Hartmann wavefront sensor for aberration correction in optical tweezers,” *Journal of Optics*, vol. 12, p. 124004, 2010.
- [99] T. Čižár, M. Mazilu, and K. Dholakia, “In situ wavefront correction and its application to micromanipulation,” *Nature Photonics*, vol. 4, p. 388, May 2010.

- [100] P. Zupancic, P. M. Preiss, R. Ma, A. Lukin, M. E. Tai, M. Rispoli, R. Islam, and M. Greiner, “Ultra-precise holographic beam shaping for microscopic quantum control,” *Optics Express*, vol. 24, no. 13, p. 13881, 2016.
- [101] J. S. Bernier, C. Kollath, A. Georges, L. D. Leo, F. Gerbier, C. Salomon, and M. Köhl, “Cooling fermionic atoms in optical lattices by shaping the confinement,” *Physical Review A*, vol. 79, p. 061601(R), 2009.
- [102] R. R. Sakhel, A. R. Sakhel, and H. B. Ghassib, “Nonequilibrium dynamics of a Bose-Einstein condensate excited by a red laser inside a power-law trap with hard walls,” *Journal of Low Temperature Physics*, vol. 173, p. 177, 2013.
- [103] H. Uncu and D. Tarhan, “Bose-Einstein condensate in a linear trap with a dimple potential,” *Communications in Theoretical Physics*, vol. 59, p. 629, 2013.
- [104] J. L. Helm, S. J. Rooney, C. Weiss, and S. A. Gardiner, “Splitting bright matter-wave solitons on narrow potential barriers: Quantum to classical transition and applications to interferometry,” *Physical Review A*, vol. 89, p. 033610, 2014.
- [105] H. Nakayama, N. Takada, Y. Ichihashi, S. Awazu, T. Shimobaba, N. Masuda, and T. Ito, “Real-time color electroholography using multiple graphics processing units and multiple high-definition liquid-crystal display panels,” *Applied Optics*, vol. 49, p. 5993, 2010.
- [106] M. Makowski, I. Ducin, K. Kakarenko, J. Suszek, M. Sypek, and A. Kolodziejczyk, “Simple holographic projection in color,” *Optics Express*, vol. 20, p. 25130, 2012.
- [107] T. Shimobaba and T. Ito, “Holographic reconstruction system by time division multiplexing with reference lights of laser,” *Optical Review*, vol. 10, p. 339, 2003.
- [108] T. Ito and K. Okano, “Color electroholography by three colored reference lights simultaneously incident upon one hologram panel,” *Optics Express*, vol. 12, p. 4320, 2004.
- [109] M. Makowski, M. Sypek, and A. Kolodziejczyk, “Colorful reconstructions from a thin multi-plane phase hologram,” *Optics Express*, vol. 16, p. 11618, 2008.
- [110] G. Xue, J. Liu, X. L. amd J. Jia, Z. Zhang, B. Hu, and Y. Wang, “Multiplexing encoding method for full-color dynamic 3D holographic display,” *Optics Express*, vol. 22, p. 18473, 2014.
- [111] E. Haller, J. Hudson, A. Kelly, D. A. Cotta, B. Peaudecerf, G. D. Bruce, and S. Kuhr, “Single-atom imaging of fermions in a quantum-gas microscope,” *Nature Physics*, vol. 11, p. 738, 2015.
- [112] C. Becker, S. Stellmer, P. Soltan-Panahi, S. Dörscher, M. Baumert, E. Richter, J. Kronjäger, K. Bongs, and K. Sengstock, “Oscillations and interactions of dark and dark-bright solitons in Bose-Einstein condensates,” *Nature Physics*, vol. 4, p. 496, 2008.
- [113] D. Simpson, D. Gangardt, I. Lerner, and P. Krüger, “One-dimensional transport of bosons between weakly linked reservoirs,” *Physical Review Letters*, vol. 112, p. 100601, 2014.

- [114] M. Woerdemann, C. Alpmann, M. Esseling, and C. Denz, “Advanced optical trapping by complex beam shaping,” *Laser & Photonics Reviews*, vol. 7, p. 839, Nov 2013.
- [115] C. Maurer, A. Jesacher, S. Bernet, and M. Ritsch-Martel, “What spatial light modulators can do for optical microscopy,” *Laser & Photonics Reviews*, vol. 5, p. 81, Jan 2011.
- [116] A. E. Willner, H. Huang, Y. Yan, Y. Ren, N. Ahmed, G. Xie, C. Bao, L. Li, Y. Cao, Z. Zhao, J. Wang, M. P. J. Lavery, M. Tur, S. Ramachandran, A. F. Molisch, N. Ashrafi, and S. Ashrafi, “Optical communications using orbital angular momentum beams,” *Advances in Optics and Photonics*, vol. 7, p. 66, Mar 2015.
- [117] M. Huo, W. Nie, D. A. W. Hutchinson, and L. C. Kwek, “A solenoidal synthetic field and the non-abelian aharonov-Bohm effects in neutral atoms,” *Scientific Reports*, vol. 4, p. 5992, 2014.
- [118] V. E. Lembessis, J. Courtial, N. Radwell, A. Selyem, S. Franke-Arnold, O. M. Aldossary, and M. Babiker, “Graphene-like optical light field and its interaction with two-level atoms,” *Physical Review A*, vol. 92, p. 063833, Dec 2015.
- [119] S. Butera, N. Westerberg, D. Faccio, and P. Öhberg, “Nonlinear synthetic gauge potentials and sonic horizons in Bose-Einstein condensates,” *arXiv:1605.05556*.
- [120] L. G. Netoand, D. Roberge, and Y. Sheng, “Full-range, continuous, complex modulation by the use of two coupled-mode liquid-crystal televisions,” *Applied Optics*, vol. 35, p. 4567, Aug 1996.
- [121] A. Jesacher, C. Maurer, A. Schwaighofer, S. Bernet, and M. Ritsch-Martel, “Full phase and amplitude control of holographic optical tweezers with high efficiency,” *Optics Express*, vol. 16, p. 4479, Mar 2008.
- [122] L. Zhu and J. Wang, “Arbitrary manipulation of spatial amplitude and phase using phase-only spatial light modulators,” *Scientific Reports*, vol. 4, p. 7441, Dec 2014.
- [123] Y. Roichman and D. G. Grier, “Projecting extended optical traps with shape-phase holography,” *Optics Letter*, vol. 31, p. 1675, 2006.
- [124] T. W. Clark, R. F. Offer, S. Franke-Arnold, A. S. Arnold, and N. Radwell, “Comparison of beam generation techniques using a phase only spatial light modulator,” *Optics Express*, vol. 24, p. 6249, Mar 2016.
- [125] S. Tao and W. Yu, “Beam shaping of complex amplitude with separate constraints on the output beam,” *Optics Express*, vol. 23, p. 1052, Jan 2015.
- [126] L. Wu, S. Cheng, and S. Tao, “Simultaneous shaping of amplitude and phase of light in the entire output plane with a phase-only hologram,” *Scientific Reports*, vol. 5, p. 15426, Oct 2015.
- [127] S. A. Goorden, J. Bertolotti, and A. P. Mosk, “Superpixel-based spatial amplitude and phase modulation using a digital micromirror device,” *Optics Express*, vol. 22, p. 17999, Jul 2014.
- [128] T. D. Team, “Theano: A python framework for fast computation of mathematical expressions,” *arXiv:1605.02688*, 2016.

- [129] D. Bowman, T. L. Harte, V. Chardonnet, C. D. Groot, S. J. Denny, G. L. Goc, M. Anderson, P. Ireland, D. Cassettari, and G. D. Bruce, “University of St. Andrews research data: Data underpinning high-fidelity phase and amplitude control of phase-only computer generated holograms using conjugate gradient minimisation,” *doi:10.17630/f30dd10e-9708-4cbf-8fcb-4b91869a1415*, 2017.
- [130] A. M. Yao and M. J. Padgett, “Orbital angular momentum: origins, behavior and applications,” *Advances in Optics and Photonics*, vol. 3, p. 161, 2011.
- [131] Y. Roichman, B. Sun, Y. Roichman, J. Amato-Grill, and D. G. Grier, “Optical forces arising from phase gradients,” *Physical Review Letters*, vol. 100, p. 013602, Jan 2008.
- [132] Images downloaded from Wikimedia Commons, 11/11/2016.
- [133] M. Takeda, H. Ina, and S. Kobayashi, “Fourier-transform method of fringe-pattern analysis for computer-based topography and interferometry,” *Journal of the Optical Society of America*, vol. 72, p. 156, Jan 1982.
- [134] C. J. Pethick and H. Smith, *Bose-Einstein Condensation in Dilute Gases*. Cambridge University Press, 2nd ed., 2008.
- [135] W. Ketterle, D. S. Durfee, and D. M. Stamper-Kurn, *Making, probing and understanding Bose-Einstein condensates. In Bose-Einstein condensation in atomic gases, Proceedings of the International School of Physics “Enrico Fermi”, Course CXL*. Amsterdam: IOS Press, 1999.
- [136] K. Dieckmann, R. Spreuw, M. Weidemüller, and J. Walraven, “Two-dimensional magneto-optical trap as a source of slow atoms,” *Physical Review A*, vol. 58, p. 3891, 1998.
- [137] J. Schoser, A. Batär, R. Löw, V. Schweikhard, A. Grabowski, Y. Ovchinnikov, and T. Pfau, “Intense source of cold Rb atoms from a pure two-dimensional magneto-optical trap,” *Physical Review A*, vol. 66, p. 23410, 2002.
- [138] W. Wohlleben, F. Chevy, K. Madison, and J. Dalibard, “An atom faucet,” *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, vol. 15, p. 237, 2001.
- [139] H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping*. Springer, 1999.
- [140] Y.-J. Lin, A. R. Perry, R. L. Compton, I. B. Spielman, and J. V. Porto, “Rapid production of ^{87}Rb Bose-Einstein condensates in a combined magnetic and optical potential,” *Physical Review A*, vol. 79, p. 063631, June 2009.
- [141] P. V. GmbH, *Vacuum Technology Book Volume II, Part 2 - Know-how Book*. Pfeiffer Vacuum GmbH, 2013.
- [142] D. A. Steck, “Rubidium 87 D line data,” available online at <http://steck.us/alkalidata> (revision 2.1.5, 13 January 2015).
- [143] K. B. MacAdam, A. Steinbach, and C. Wieman, “A narrow-band tunable diode laser system with grating feedback, and a saturated absorption spectrometer for Cs and Rb,” *American Journal of Physics*, vol. 60, p. 1098, 1992.

- [144] L. Ricci, M. Weidemüller, T. Esslinger, A. Hemmerich, C. Zimmermann, V. Vuletic, W. Köinig, and T. Hänsch, “A compact grating-stabilized diode laser system for atomic physics,” *Optics Communications*, vol. 117, p. 541, 1995.
- [145] C. D. Wallace, T. P. Dinneen, K. Y. N. Tan, A. Kumarakrishnan, P. L. Gould, and J. Javanainen, “Measurements of temperature and spring constant in a magneto-optical trap,” *Journal of the Optical Society of America B*, vol. 11, p. 703, 1994.
- [146] J. Dalibard and C. Cohen-Tannoudji, “Laser cooling below the Doppler limit by polarization gradients: simple theoretical models,” *Journal of the Optical Society of America B*, vol. 6, p. 2023, 1989.
- [147] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov, “Optical dipole traps for neutral atoms,” *Advances in Atomic, Molecular and Optical Physics*, vol. 42, p. 95, 2000.
- [148] Y. Castin and R. Dum, “Bose-Einstein condensates in time dependent traps,” *Physical Review Letters*, vol. 77, p. 5315, December 1996.
- [149] J. Szczepkowski, R. Gartman, M. Witkowski, L. Tracewski, M. Zawada, and W. Gawlik, “Analysis and calibration of absorptive images of Bose–Einstein condensate at nonzero temperatures,” *Review of Scientific Instruments*, vol. 80, p. 053103, 2009.
- [150] T. A. Savard, K. M. O’Hara, and J. E. Thomas, “Laser-noise-induced heating in far-off resonance optical traps,” *Physical Review A*, vol. 56, p. R1096, 1997.
- [151] M. E. Gehm, K. M. O’Hara, T. A. Savard, and J. E. Thomas, “Dynamics of noise-induced heating in atom traps,” *Physical Review A*, vol. 58, p. 3914, 1998.
- [152] J. Leach, G. Sinclair, P. Jordan, J. Courtial, M. J. Padgett, J. Cooper, and Z. J. Laczik, “3D manipulation of particles into crystal structures using holographic optical tweezers,” *Optics Express*, vol. 12, p. 220, 2004.
- [153] H. Chen, Y. Guo, Z. Chen, J. Hao, J. Xu, H.-T. Wang, and J. Ding, “Holographic optical tweezers obtained by using the three-dimensional Gerchberg–Saxton algorithm,” *Journal of Optics*, vol. 15, p. 035401, 2013.
- [154] D. Barredo, V. Lienhard, S. de Léséleuc, T. Lahaye, and A. Browaeys, “Synthetic three-dimensional atomic structures assembled atom by atom,” *arXiv:1712.02727*, December 2017.
- [155] X.-L. Wang, J. Ding, W.-J. Ni, C.-S. Guo, and H.-T. Wang, “Generation of arbitrary vector beams with a spatial light modulator and a common path interferometric arrangement,” *Optics Letters*, vol. 32, p. 3549, 2007.