

Development of an apparatus for the rapid production of Rubidium-87 Bose Einstein Condensates

Author:

Sebastien

TEMPONE-WILTSHIRE

Supervisor:

Prof. Kristian

HELMERSON

Dr. Lincoln TURNER

A thesis submitted in total fulfilment of the degree of
Doctor of Philosophy
in the School of Physics and Astronomy, Monash University

February 26, 2020



Abstract

This thesis details the methods used to create arbitrary complex optical fields using a phase only spatial light modulator, as well as a novel method for creating holographic optical elements from these complex optical fields using a photopolymer known as Bayfol HX. As well as covering the method of creating HOEs, this thesis outlines all the stages from design to execution required to build an apparatus that can produce ^{87}Rb Bose–Einstein condensates of 2.2×10^6 atoms in just under 10 seconds. There is a focus in this thesis on distilling the relevant information required to understand and optimise the workings of a Bose–Einstein condensate apparatus from the literature, and provide both a set of benchmarks for anyone building a new apparatus, as well as a set of procedures for optimising the production of BECs. In the final chapter of this thesis I also demonstrate how holographic optical elements can be used to create a uniform potential for a condensate.

Cover picture: Photograph of the Monash University crest, reconstructed via a phase-only hologram using our spatial light modulator.

Declaration

This thesis contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

Sebastien Joseph Tempone-Wiltshire
18/11/2019

©Sebastien Tempone-Wiltshire (2019)

Acknowledgements

The journey of a doctoral degree is rarely completed in vacuo, especially when your path is in experimental physics and rests upon the shoulders of so many who came before you. As such, there are a great number of people I would like to acknowledge who have helped me along my path; motivating, supporting and teaching me.

Firstly, I would like to thank my primary supervisor, Kris, whom I first met when choosing my honours project and have worked with ever since. Kris has been an excellent supervisor to me, his more hands-off approach has allowed me to venture into the unknown and explore by myself, allowing me to develop my own ability to create new ideas and learn to problem solve. Throughout my candidature he has been an endless source of positivity, in the lab he is always available to help when called upon and is very solutions focussed, often presenting simple solutions to problems that I have missed, with Kris by your side it is hard to feel daunted by what is ahead. Outside of the lab he is great at bouncing ideas off, frequently we have come up with new ideas throughout my candidature that have been worthwhile investigating, and his knowledge of other fields outside of cold atoms and his ability to see the role our skills can play in those areas is invaluable.

When I first stepped into the lab at the beginning of my Honours year in 2014, Phil and Shaun were there to introduce me to what research in experimental physics was really like. Phil and Shaun have both been invaluable to me throughout my candidature, teaching me the everyday skills needed in a lab; from fiber coupling to soldering, designing circuit boards through to understanding labview code, and they have both played an enormous role in helping me design and construct the new apparatus, without them I'm sure the design and execution of the new apparatus would not have gone as well, and would have taken considerably longer. It is upon

their shoulders, along with the work of Chris in the design and execution of the control system, that the simplicity in running of the apparatus on a day to day basis sits.

My co-supervisor Lincoln, and unofficial supervisor Russ both played an enormous role in my development as a scientist. To Lincoln, thank you for your endless willingness to spend time having a chat either about science or everyday matters, making me aware of all aspects of academia and professional work in general. In particular your ability to explain science through your enthusiastic ‘hands on’ demonstrations, as well as your determination for a thorough explanation of any observed phenomena are qualities I revere in a scientist. To Russ, your enthusiasm, motivation and friendship has helped me greatly throughout my PhD and in particular in the lab I attribute a large portion of the optimisation of my apparatus to you, your focus on defining the right metrics for the optimisation has been invaluable as has your help in visually demonstrating results with Mathematica and teaching me to simulate with Mathematica.

Outside of my research group I have received a great amount of support; Jean kept the administrative side of my candidature running and is the sole person you would want on your side in an argument with MGE, Brett and Rodd taught me a great number of mechanical skills in the supervised access workshop, as well as building a great number of components of my apparatus for me and Nino was always willing to discuss an electronics problem and offer his insights.

Finally, I would like to thank my family; first to my fiancée Frannie, who has been with me throughout the whole journey, through the days of stress where it feels like I’m achieving nothing, through to the happier moments such as making my first condensate, you have always been there to ground me being the realist that you are, never allowing me to become too pessimistic. To my Mum and Dad as well as brother and sister, thank you for your support and endless confidence in me, at times it has felt unjustified, but it has helped me at every milestone, from high school through to now and even beyond in my further career decisions, its incredibly reassuring to take a step into the unknown and know that your family has trust that you will make it work.

Contents

1	Introduction	1
1.1	Structured potentials for cold atoms	2
1.1.1	Structuring light	3
1.1.2	The uses of structured light	4
1.1.3	Holographically generated structured optical fields	6
1.2	Thesis outline	8
2	Complex control of an optical field	13
2.1	Spatial light modulators for phase control	13
2.2	Methods for full complex control of an optical field	16
2.2.1	Encoding a complex optical field with a single phase only filter	16
2.2.2	Implementing the encoding method	20
2.3	Creating an optical vortex knot	22
2.4	Optical vortex knots, one photon at a time	23
2.4.1	Measuring an optical vortex knot	24
2.4.2	‘Which path’ measurements	28
2.4.3	Summary	30
3	Holographic optical elements for ultracold atom trapping	33
3.1	Problems with SLMs	34
3.2	Bayfol Hx	38
3.2.1	Recording with Bayfol HX	38
3.3	Recording a holographic optical element	41
3.3.1	Quantifying the success of our method	43
3.4	Generating large angular momentum Laguerre-Gaussian modes	45
3.5	Creating other optical elements	48
3.6	Summary	49
4	Building a Bose–Einstein Condensate apparatus	51
4.1	Ultrahigh vacuum for ultra-insulation	52
4.1.1	Vacuum pumps	54

4.1.2	Designing the differential pumping tube	56
4.1.3	Pressure measurement	58
4.1.4	Source and Science chambers	59
4.1.5	Overall system design	60
4.2	Achieving Ultra-High Vacuum	64
4.2.1	Component preparation and construction	64
4.2.2	Bakeout	65
4.3	Laser cooling system	69
4.3.1	Cooling laser	70
4.3.2	Repump laser	75
4.4	Quadrupole coil design	77
4.4.1	Geometry of the coils	78
4.4.2	Switching speed of a coil	79
4.4.3	Power dissipation of a coil	80
4.4.4	Power supply considerations	81
4.4.5	Final quadropole coil design	82
4.4.6	Bias coils and formers	83
4.5	Magneato coil drivers	84
4.5.1	Design of the Magneato	85
4.5.2	IGBT power dissipation	88
4.5.3	Magneato performance	89
4.5.4	Future improvements to the Magneato design	90
4.6	Coil cooling system interlock	91
4.6.1	The interlock system	92
4.6.2	Temperature measurement	92
4.6.3	Water flow measurement	93
4.7	RF generation and shuttter control	94
4.7.1	RF generation	94
4.7.2	Laser shutter control	95
5	Laser cooling of atoms	97
5.1	Optical forces and molasses	98
5.1.1	Limitations of optical molasses	101
5.1.2	Magneto-optical traps	103
5.1.3	Experimentally relevant quantities of a MOT	104
5.1.4	Optimal parameters for two and three dimensional MOTs	112
5.2	Sub-Doppler cooling	118
5.2.1	Polarisation gradient cooling	120
5.2.2	Experimentally relevant qualities of sub-Doppler cooling	122
5.2.3	Sub-Doppler cooling in a MOT	125
5.3	Compression of atoms in a MOT (CMOT)	127
5.4	Experimental implementation of laser cooling	130
5.4.1	Two-dimensional Source MOT	130
5.4.2	Three-dimensional Science MOT	137
5.4.3	Compressed MOT optimisation	144

5.4.4	PGC optimisation	147
6	Evaporative cooling of atoms	155
6.1	Forced evaporative cooling	156
6.1.1	Cooling rates during forced evaporation	160
6.1.2	Maximising PSD	162
6.1.3	Runaway evaporation	164
6.2	Forced evaporation in a magnetic trap	165
6.2.1	Trapping in a magnetic field	165
6.2.2	Majorana losses	166
6.2.3	Radiofrequency evaporation in a magnetic trap	167
6.2.4	Adiabatic compression of a magnetic trap	171
6.3	Forced evaporation in an optical dipole trap	171
6.3.1	The optical dipole force	172
6.3.2	Forced evaporation in an optical trap	175
6.4	Experimental considerations of forced evaporation in combined optical and magnetic potentials	177
6.4.1	Magnetic traps	177
6.4.2	Dipole traps	178
6.4.3	The hybrid trap	179
6.5	Experimental implementation of forced evaporation	186
6.5.1	The catch	186
6.5.2	Moving and compressing the trap	192
6.5.3	Characterising the magnetic trap	194
6.5.4	Forced evaporation in the compressed magnetic trap	197
6.5.5	Loading the hybrid trap	199
6.5.6	Characterising the hybrid trap	207
6.5.7	Forced evaporation in the hybrid trap	211
6.5.8	Bose-Einstein condensation	213
7	Imaging cold atoms	219
7.1	Optical interactions	220
7.2	Absorption imaging	224
7.2.1	Measuring saturation intensity	226
7.2.2	Measuring imaging system magnification	228
7.2.3	Taking an absorption image	229
7.2.4	Interference fringes	231
7.3	Fluorescence imaging	233
7.4	Measuring the temperature of a cold gas	235
7.5	Optical imaging systems	236
7.5.1	Side imaging	237
7.5.2	Bottom imaging	238
8	Holographic potentials for Bose-Einstein condensates	243
8.1	Creating uniform potentials	243

8.1.1	Blue-detuned uniform potentials	248
8.1.2	A dynamic uniform potential	250
8.2	Implementing a rotating potential	253
8.2.1	Axial confinement- The sheet trap	257
8.2.2	Aligning the uniform trap	259
8.3	Characterising the uniform trap	261
8.3.1	Rotating the uniform trap	267
8.4	Future outlook	269
A	Bose-Einstein condensation	271
A.0.1	Thermodynamics of trapped gases	275
B	Vortices	277
C	Transconductance of the Magneato driver	279
D	Phase space density calculations	281
	Bibliography	281

List of Figures

2.1	Scavenging an SLM from a projector a) , b) and c) show the original Hitachi CPA100 projector from which I scavenged an SLM from, contained within the 3RU box with the SLM attached beneath with the connector plate visible and one of the SLMs as removed from the projector.	15
2.2	Creation of an arbitrary complex field. A schematic showing how a complex optical field can be created using a phase only spatial light modulator and a computer generated hologram. A halfwave plate orients the polarisation of the light for maximum diffraction efficiency, followed by a pair of lenses that expand the beam to overfill the SLM. Another pair of lenses follows the SLM, reimaging the SLM in a 4f configuration and allow spatial filtering of the desired mode. .	20
2.3	Implementing computer generated holograms. a) shows how the addition of a carrier frequency to the computer generated hologram spatially separates the Fourier components of the hologram in the Fourier plane, whilst b) shows how this can be implemented optically. Reproduced from [1].	22
2.4	Optical set-up for producing optical vortex knots, one photon at a time. Each photon passes through the spatial light modulator (SLM), which acts as a static phase element generating both the knot and a Gaussian reference beam, which is split upon a non-polarising beam splitter cube (NPBS) and then spatially filtered with a pinhole (SF) such that one path contains the knot, and the other the reference beam. These two paths are then recombined at a small angle, which generates an interferogram that facilitates identification of the vortex structure, and then imaged onto the camera (EMCCD). The double headed arrow indicates the Imaging lens that is translated to image different transverse planes of the knot. Reproduced from [2]	24

- 2.5 **Left: Measurement of the knotted topology a**, A single frame imaged by the camera in photon counting mode. **b-d**, Measurements resulting from the summation of 100, 1000 and 8000 images, respectively, with the contrast of the resultant interference pattern increasing with increasing numbers of images making the forked vortex structure more visible (shown magnified in insets). **Right: Normalised cross correlation as a function of photon number**. The peak value of the NCC between the interferometric measurement of the knotted optical field at $Z=0$ with the expected theoretical intensity pattern of the knot, a Gaussian mode and the modal decomposition of the knot used in the 'which path' measurement, each interfered with a Gaussian mode to reproduce the global fringe pattern as a function of increasing number of summed images, displayed as increasing photon number. Reproduced from [2] 26
- 2.6 **Measuring the optical fields knotted topology**. Different transverse planes of the optical field along the beam propagation direction (indicated by the black arrow) are imaged onto the camera by moving the position of the final imaging lens before the camera (see optical set up in left of Fig. 2.4). The length scales in both the propagation direction and transverse plane are shown. **a**, Vortices within each plane (indicated by red dots) are located by the characteristic forked structure they create in each interferogram, due to the 2π phase winding about each vortex. **b**, The vortices located in 25 separate planes are then 'stitched' together to form the three dimensional trefoil knot nodal structure. Reproduced from [2]. 27
- 2.7 **Left: Optical set-up for a 'which path' measurement**. The SLM displays two holograms side by side. The laser beam is split into two and passes through a chopper blade (CB), which ensures that photon's are present along the optical path from only one of the two holograms at any given time. The two beam paths are later recombined and imaged onto the camera. **Right: 'Which path' measurements of the optical field**. The three columns shown contain the theoretical intensity patterns of the knotted optical field (left), the intensity patterns resulting from the unchopped beams (centre) and from the chopped beams (right). These images are taken at three different transverse planes along the beam path (positions are indicated in centimetres). A normalised cross correlation of the data (centre and right) with simulation (left) shows a higher correlation of the knot with the unchopped optical field, with an average peak value of 0.89, as opposed to when the beam is chopped, which yields an average peak value of 0.62. Reproduced from [2] 29

- 3.1 **Pixellation of an SLM.** **a)** and **b)** show the effect of the pixellation of the SLM upon the transmitted light, with **a)** showing a Gaussian mode that is not passing through the SLM whilst **b)** shows how transmission through the SLM generates a large number of diffracted modes, significantly reducing the power in the undiffracted mode. 35
- 3.2 **The working principle of Bayfol HX.** Bayfol HX consists of two main components, imaging monomers, and matrix precursors in a resin. During the films production at Bayer, the matrix precursors are formed into a cross-linked matrix. The film is then exposed to laser irradiance to record the hologram, initiating photopolymerisation, and the imaging monomers then diffuse away from the dark regions and this, coupled with the polymerisation, creates a refractive index modulation storing the hologram. 37
- 3.3 **Overmodulation of a transmission hologram.** The dependence of the diffraction efficiency, η , upon the grating strength, which increases with exposure time is shown. The graph has three defined regions during exposure. Region 1 corresponds to an initial inhibition in which recording of the hologram has not yet begun as photopolymerisation is still occurring. Region 2 shows the growth of the refractive index modulation whilst the grating strength is still in the weak grating regime, and thus the diffraction efficiency increases with increasing modulation. Region 3 shows the strong grating regime, where the diffraction efficiency begins to fall with further modulation. Reproduced from [3]. 40
- 3.4 **Hologram recording setup.** A schematic of the experimental set up used for recording transmission holograms. M, mirror; NPBS, non-polarising beam splitter cube; P, polariser; $\lambda/2$, half-wave plate; L, lens; SF, spatial filter; SLM, spatial light modulator. For recording reflection holograms the stage on which the hologram is mounted is rotated such that the object and reference beams enter the photopolymer on opposite surfaces. Reproduced from [4]. 42
- 3.5 **The result and use of a $HG_{(0,1)}$ recorded into a hologram.** **a)** Image of the intensity profile of the $HG_{(0,1)}$ mode at the focus of a lens. **b)** A fit (dashed lines) of a $HG_{(0,1)}$ intensity profile to the measured intensity profile (circles) taken along the green line shown in the image. **c)** An image of ultra cold Rubidium-87 atoms trapped within the nodal region of $HG_{(0,1)}$ mode (shown in 3.5a) formed from ≈ 50 mW of 532 nm laser light with a magnetic field providing the transverse confinement of the atoms. Optical density is shown on the right hand side. Reproduced from [4]. 45

3.6	Optical setup for recording large angular momentum modes. The optical set up used to record the hologram with twice the angular momentum per photon achievable by diffraction from just the SLM; all labels are as in Fig 3.4. The difference in the number of reflections each mode undergoes must be an odd integer in order for the modes to be mirror images of each other. Reproduced from [4].	46
3.7	The resulting optical field generated by the method described in section 3.4. a) The intensity distribution of the optical field. The (green) line is the azimuthally averaged, radial profile. b) The interference of the optical field with its mirror reflection, verifying the high angular momentum state. Reproduced from [4].	48
4.1	Nextorr D200. A Solidworks representation of the NexTorr D200. The red box contains the ion pump which is designed to actively pump the noble gases, whilst the cartridge seen extending from the box is the NEG cartridge that performs most of the pumping, but passively.	55
4.2	A Solidworks model of the designed differential pumping tube. The designed differential pumping tube formed by three concentric holes of varying diameters and lengths drilled into a 12mm diameter grade 304 stainless steel rod. The holes were designed to allow divergence of the atomic beam coming from the 2D MOT while also achieving the desired conductance to maintain a pressure differential of approximately four orders of magnitude.	57
4.3	Glass cells. a) A Solidworks representation of the source cell, made from a rectangular cuvette with 4 glass tubes attached to one end in which the getters are installed. b) A Solidworks representation of the octagonal glass cell used as the science cell.	60
4.4	Vacuum system perspectives. a) A Solidworks model of the final vacuum system looking along the length of the system from the source end, when cut through the centre of the spherical octagon, the centrepiece of the system. b) shows the same model but viewed from the side instead and cut along the length vertically to show the path the atoms follow from the 2D MOT formed in the Source chamber on the left hand side, through to the 3D MOT in the octagonal chamber on the right hand side.	61
4.5	Finalised vacuum system. The final vacuum system design with mounts and some optics included. The source end with the 2D MOT optics as well as the coil formers is shown in the forefront, whilst the science side is in the background with the octagonal cell mostly obscured by the quadrupole coil formers.	62

- 4.6 **The bakeout oven.** To achieve ultrahigh vacuum the entire vacuum system must be ‘baked out’ to help remove any water present and reduce the outgassing of hydrogen from the steel. Pictured is the vacuum system once an oven had been constructed about it and the entire system insulated with aluminium foil. On the table are 7 red variacs, that allow independent control of the seven heater tapes used to heat the vacuum system, ensuring the temperature of the vacuum system can be uniformly increased thereby preventing any leaks forming from differential expansion of the materials. 66
- 4.7 **Bakeout temperature and pressure.** During bakeout the temperature and pressure of the vacuum system was measured. The bottom graph displays the temperature measured by the 15 different temperature sensors on a linear scale, the lowest temperature reading (red) is of a thermistor placed upon a set of magnetic field coils, which had to be in-situ during bakeout, visible on the left of figure. 4.6 The top graph displays the pressure read by the RGA during the bakeout on a logarithmic scale. The pressure is seen to increase by 2 orders of magnitude as the temperature is increased, and falls approximately exponentially by over an order of magnitude over the two weeks. On the day we returned from Easter we slightly increased the temperature further, leading to a small increase in pressure, but there was also a significant drop in the pressure when we started backing the turbo pump with a scroll pump. 67
- 4.8 **RGA scans pre and post bake.** a) and b) show measurements of the partial pressure of gases remaining in the vacuum system made with a residual gas analyser on the turbo pump that pumped out the vacuum system, pre and post bake respectively. Clearly after the bakeout the partial pressures of all the gases have reduced considerably, notably the largest value on the scale of the post bake out scan is almost an order of magnitude less than on the pre bakeout scan. 69
- 4.9 **Hyperfine levels of $5^2S_{1/2}$ to $5^2P_{3/2}$ transition.** The hyperfine levels of the $5^2S_{1/2}$ to $5^2P_{3/2}$ transition are shown, as well as the features that are used in this thesis. Arrow a) indicates the lock point for the cooling laser, 47 MHz below the $F=2$ to $F'=2/F'=3$ crossover, which places the lock point 180 MHz below the $F=2$ to $F'=3$ cycling transition. Arrows b) and c) represent the approximate detunings used for the MOT cooling light (in both the source and Science MOT, and close to the value used for the push beam) and the imaging light. Arrow d) represents the lock point of the repump laser locked directly to the $F=1$ to $F'=0/F'=1$ crossover, placing the lock 194 MHz below the repump transition, $F=1$ to $F'=2$, where the repump light for the Source and Science MOT (arrow e)) are fixed to. 70
- 4.10 **Saturated absorption signal for the cooling transition.** The Digilock display of the saturated absorption signal (yellow) and resulting error signal (red) used for locking the cooling laser to the $F' = 2/F' = 3$ crossover, indicated by the dotted white lines. 71

- 4.11 **Schematic of the cooling and repump laser setup.** A schematic of the cooling and repump laser system at the time of writing. Not shown are the saturated absorption setups for locking the lasers. The bottom lefthand side of the image indicates the key, where PBS is a polarising beam splitter cube typically used combined with a half wave plate, $\lambda/2$ to control the amount of power in a given beam line or to combine two beams, $\lambda/4$ is a quarter wave plate, which combines with a lens and a mirror to form a cateye retroreflector for the double passed acousto-optic modulators (AOMs). The thick apertures are optical shutters, used to extinguish light leakage through the AOMs, whilst the thin apertures are irises used to remove the undiffracted mode in the double passed AOM configurations. Fiber couplers consist of a half wave plate used to align the axis of polarisation of the input light to the fast axis of the fiber, an aspheric lens to focus the light onto the fiber and an x-y mount used to align the fiber mount with the focussed spot. 72
- 4.12 **Error signal for the repump transition.** An image of an oscilloscope displaying the error signal (blue) and saturated absorption signal (yellow) used to lock the repump laser to the $F = 1$ to $F' = 0/F' = 1$ crossover transition. The largest feature in the error signal is used as the lock point 77
- 4.13 **Coil formers.** The interlocking coil former design, made from 3 mm PVC to prevent eddy currents persisting after the coils are switched off. The coil former design ensures the alignment and relative position of the coils with respect to one another, and allows for easy construction about the science cell. 84
- 4.14 **Unipolar Magneato schematic.** A schematic of a unipolar Magneato circuit used to control the current through the quadrupole coils. Shaded in green in the sensing circuit, used to sense how much current is passing through the coil and convert it to a voltage. Shaded in red is the control circuit, used to buffer the input control signal, preventing feedback into the NI card. Shaded in blue is the comparison circuit used to compare the signals from the control and sense circuits, and provide a control signal to the IGBT used to modulate the current through the quadrupole coil. 86

- 4.15 **Simetrix simulations of the current switching for the catch** **a)**, **b)** and **c)** show the effect of changing the resistor, R_5 , in the Magneato driver on the resultant catch current. For values of R_5 that are too low, **a)**, some ringing is present as the system is underdamped, for values that are too large, **c)**, the driver can go into oscillation, whilst for a small range of values, **b)**, the system will be close to critically damped and will perform optimally. All simulations show some ringing when suddenly switched off, however, these oscillations are damped by a diode placed across the coil, not included in the simulation. 89
- 4.16 **Quadrupole Magneato simulations** **a)** A Simetrix simulation of the current through the quadrupole coil when suddenly switched on after 5 ms to the catch current value (40 A). **b)** A Simetrix simulation of the transconductance of the Magneato optimised for the Quadrupole coils. The control voltage is linearly ramped from 0 to 10 V over 100 ms, before being held constant for 10 ms and finally being linearly ramped back down to 0 V. The current through the quadrupole coil clearly follows the voltage linearly, however, there is a slight lag initially, this is not observed experimentally. 90
- 4.17 **Shutter control circuit.** The designed PCB circuit board used to switch between digital (TTL) control of the shutters, and manual on/off mode for diagnosis of problems. This board was designed to be mounted directly to the front panel of a 2RU box via the DPDT switches, and also switches the state of two LEDs to allow the user to know whether the shutters are under manual or TTL control, and whether they are on or off. 95
- 5.1 **Force in optical molasses** The force upon an atom due to a pair of low intensity counter propagating plane waves detuned half a linewidth below resonance. The dashed lines indicate the force arising from the two plane waves independently, whilst the solid line indicates the resultant force. There is a clear linear region near $v = 0$, which is utilised for laser cooling. Reproduced from ref. [5]. 100

- 5.2 **Representation of the drunkards walk** A representation of the drunkards walk and how it can lead to cooling in a MOT. On the left hand side we see the effect of a single absorption/spontaneous emission event. The atom initially absorbs a photon from the MOT laser (red headed arrow) which reduces the atoms velocity, and then undergoes spontaneous emission and re-emits the photon in a random direction (blue headed arrow). The right hand side shows how undergoing this process multiple times reduces the atoms initial velocity to some finite velocity. The two concentric circles are representative of the capture velocity (larger circle), v_c , determined by the MOT parameters, and the minimum velocity (small circle) of the atom, whilst the coloured circles represent atoms undergoing multiple absorption re-emission events as they are cooled. Note that once the atoms are within the smaller circle they can still undergo absorption re-emission cycles, these determine the minimum temperature of the MOT. . . . 108
- 5.3 **Schematic of a pure 2D MOT** A schematic of a pure 2D MOT including the rectangular cuvette, the four racetrack MOT coils required to produce the line of field zero in the centre of the chamber, the MOT beams with their respective polarisations as well as the differential pumping tube is shown. The probe-laser, plug-laser and photodiode (PD) were not utilised. Reproduced from reference [6]. . . 113
- 5.4 **2D MOT dependencies. a) Capture velocity of a 2D MOT.** The dependence of the capture velocity of the pure 2D MOT upon the intensity per MOT beam is shown, Schoser *et al.*[6] saw a saturation of the capture velocity, and thus a saturation of the load rate, for intensities on the order of 17 mW cm^{-2} . **b) Velocity dependence on MOT beam length.** Shown above is the dependence of the velocity distribution of the cold atomic beam emerging from the 2D MOT on the MOT beam length. With longer MOT beams the cooling volume is extended and thus atoms with larger axial velocities can be cooled, leading to an increase in the peak axial velocity of the distribution. Circles indicate measured data whilst squares are the results of a model. **c) Flux dependence on MOT beam length.** The dependence of the flux of the 2D MOT on the length of the MOT beams, there is a clear almost linear dependence of the flux on the length of the MOT beams. The squares show the theoretical values whilst the circles show the measured data. All figures reproduced from ref. [6]. 115

- 5.5 **Graphical representation of polarisation gradient cooling** A graphical representation of how polarisation gradient cooling occurs. Along the x-axis the polarisation of light is shown to vary over the course of half a wavelength, represented by the red arrows, from initially linearly polarised, through to circularly polarised, back to linearly polarised and finally circularly polarised but of the opposite handedness. The atom is represented by the coloured circle, with the colour indicating the magnetic substate the atom is in, either $m_f = 1/2$ (blue) or $m_f = -1/2$ (red), and the red and blue curves depict how the light shifts of the different substates varies as a function of the polarisation. An atom initially in $m_f = 1/2$ moving along the x-axis climbs the potential hill that arises due to the spatially varying polarisation before being optically pumped into the $m_f = -1/2$ and the process repeating until the atom no longer has enough kinetic energy to climb the potential hill. 121
- 5.6 **PGC temperature dependence on magnetic fields** The experimental results of Lett *et al.*, in which they demonstrate how the minimum temperature achieved by PGC shows a strong dependence upon external magnetic fields. Two different sets of measurements at different powers are shown. Reproduced from ref. [5] 125
- 5.7 **Light induced atomic desorption** An image of the source chamber whilst illuminated by UV LEDs to perform light induced atomic desorption. When the pressure in the source chamber is low enough LIAD can provide up to a factor of 20 increase in the load rate of the 3D MOT. 133
- 5.8 **2D MOT optimisation** The measured peak fluorescence count of the Science MOT after 3 seconds of loading from the Source MOT while scanning three different parameters; **a)** the current in the bottom source coil, **b)** the current in the source coil on the south side of the chamber and **c)** the detuning of the 2D MOT below resonance. Clearly there is an optimum value for the current through each coil that, when viewed with a camera along the length of the Source MOT, coincides with the 2D MOT aligning with the hole in the differential pumping tube, whilst there is also a clear optimal detuning for the load rate, which is ≈ 13 MHz below resonance. 135
- 5.9 **2D MOT image.** **a)** A 2.37 ms exposure taken of the Source MOT, looking along its length. Above the fluorescence of the 2D MOT, scattered light off the differential pumping tube from the vertical MOT beam can be seen. **b)** A vertical (top) and horizontal (bottom) lineout taken through the centre of the 2D MOT. 137

- 5.10 **Push beam optimisation.** **a)** and **b)** show the peak fluorescence from an image of the MOT, used as a metric for the load rate, as a function of the detuning below and above resonance (positive numbers indicate below) of the push beam, as well as a function of beam power. The colour coding of the 2D graphs are shown to the right of each, respectively. Note that the integrated fluorescence for the red detuned cases is approximately 3 times larger than for the blue detuned case. 138
- 5.11 **3D MOT optimisation.** **a)** and **b)** show the Science MOT fluorescence as a function of the control voltage applied to the bias coils Magneato in the Z (along gravity) and Y (along the optical axis for side imaging) directions respectively, which determine the location of the quadrupole field zero, and hence the centre of the MOT. Each has a clear optimal value for the load rate. **c)** The MOT fluorescence measured for different detunings of the cooling light below resonance, and values of the quadrupole gradient. 140
- 5.12 **Loading of, and the final 3D MOT.** **a)** A fluorescence image of the MOT when fully loaded, taken with the optimal parameters found in the optimisation procedure. It also shows two attempted Gaussian fits to the data, resulting in widths of 2.6 mm and 1.35 mm in the horizontal and vertical directions, however, the MOT clearly has an irregular shape thus these are only rough indicators of size. **b)** The measured fluorescence signal from the MOT while the MOT loads, the signal saturates to its maximal value in just under 4 seconds. . . 141
- 5.13 **MOT Thermometry.** Thermometry of the MOT once fully loaded, with data taken over the first 0.5-7 ms. The vertical axis shows the width of the cloud squared, whilst the horizontal axis shows the time squared, with the resultant temperature of 700 μK found from the gradient of the fitted lines (shown). 142
- 5.14 **3D MOT loading movement** **a), b)** and **c)** show the 3D MOT at three sequential times during its loading. As the MOT loads the beam imbalance caused by the atoms casting a shadow in the MOT beams pushes the MOT back towards the centre of the cell. 143
- 5.15 **CMOT gradient optimisation.** **a)** The dependence of the width of the resulting compressed MOT on the magnetic field gradient. A clear minimum in width occurs for gradients of 80-90 Gcm^{-1} . This data was taken after compressing the gradient over 200 ms. **b)** The integrated fluorescence from the MOT for the same set of experiments as in **a)** is shown, used as a metric for the atom number. This data clearly shows that whilst the width is decreasing the integrated fluorescence is increasing, indicating that the spatial density is increasing and is not being limited by other loss mechanisms. 144

- 5.16 **CMOT detuning optimisation.** **a)** and **b)** show the dependence of the width of final compressed MOT on the detuning of the cooling and repumping light, respectively. Each dataset was taken with the quadrupole gradient set to 85 Gcm^{-1} . Each has clear optimal values for minimising the width of the cloud, and the peak fluorescence for the two data sets shows the expected inverse relationship, being near a peak when the width is minimised. A 2D scan of the detuning of both the repump and cooling light was also performed which showed that each parameter showed little dependency upon the other. 146
- 5.17 **In-situ fluorescence image of the resulting CMOT.** The compressed MOT in situ after the compression sequence. Clearly the MOT now has a more regular Gaussian form with the displayed fits to the line outs fitting the data quite well, indicating waists of 1.3 mm and 0.68 mm in the horizontal and vertical directions respectively. 147
- 5.18 **CMOT thermometry.** **a)** Thermometry performed on the CMOT from 0.5ms to 30ms of expansion. The data is clearly not linear, **b)** and **c)** Temperature fits to the short (0.5-7ms) and long (7-30ms) drop time data. The short drop time data has a temperature found to be quite close to the MOT temperature, whereas the long drop time data indicates that sub-Doppler cooling is occurring in the centre of the CMOT. 148
- 5.19 **Optimisation of polarisation gradient cooling.** **a)** Fluorescence as a function of Z bias control voltage. Note the very small range of control voltages ($\approx 10 \text{ mV}$), outside this range no atoms were observed after 500 ms of PGC. **b)** Cloud width after a 20 ms drop as a function of detuning and intensity. **c)** Peak fluorescence after PGC as a function of repump intensity and detuning. **d)** Peak fluorescence as a function of PGC time and detuning 149
- 5.20 **PGC Thermometry.** Thermometry for the PGC cooled cloud for drop times in the range of 0-30ms. The data is much closer to a linear trend than the thermometry performed on the CMOT, indicating that nearly all the atoms are cooled to this final temperature and thus the two component nature of the cloud has been removed. 151
- 5.21 **Runviewer trace of the laser cooling process** A trace of relevant parameters that change during the laser cooling process is shown. The included parameters, from top to bottom, are the control voltage applied to the quadrupole coils Magneato driver, the RF frequency applied to the science repump AOM and the RF frequency applied to the Science MOT cooling AOM. 152

- 6.1 **Forced evaporation efficiency and runaway evaporation dependence on the potential and ratio of good to bad collisions.** **a)** shows the dependence of the efficiency of the evaporation, characterised by γ , upon the depth at which the distribution is truncated, $\eta k_B T$ for a 3D harmonic potential (solid lines) and a 3D linear potential (dashed lines) for 3 different ratios of good to bad collision, R , 5000, 1000 and 200 respectively from the upper to lower sets of lines. From comparison of the curves, at a given value of R , the peak efficiency in the parabolic trap is always greater than in the linear trap, and peak efficiency increases with increasing R , as expected. **b)** shows how the minimum ratio of good to bad collisions, R , required to achieve runaway evaporation varies as a function of the truncation parameter for a parabolic potential (solid lines) and a linear potential (dotted lines). **a)** and **b)** were both altered from Reference. [7]. 163
- 6.2 **Graphical representation of RF evaporation.** A graphical representation of how RF evaporation occurs in a quadrupole potential. The top part of the figure shows how the velocity distribution changes through the three times considered, before evaporation, during evaporation and after thermalisation, whilst below a graphical interpretation of the process is shown. In the representation atoms are shown as coloured circles with the colour indicating the kinetic energy the atom has, and the potential is drawn as a black curve, with the potential seen by the $(F = 1, m_f = -1)$, $(F = 1, m_f = 0)$ and $(F = 1, m_f = 1)$ indicated. 169
- 6.3 **Mode-matching a hybrid trap.** **Top** shown above are the calculated trap volumes in the hybrid trap as a function of cloud temperature for a range of different field gradients, given in Tm^{-1} , the dashed lines indicate the trap volumes in the dipole and quadrupole traps alone. **Bottom** The bottom figure shows how the entropy per particle changes with temperature in the hybrid trap (solid lines) and quadrupole only trap (dashed lines). Reproduced from reference [8] 183
- 6.4 **Optimisation of the catch.** **a)** Peak fluorescence from the magnetic trap as a function of the Z bias control voltage, clearly more atoms are caught as the bias increases until it reaches a peak at approximately 3V. Interestingly this corresponds to a magnetic trap position above the initial cloud. **b), c), d)** Cloud width measured after a 30 ms drop once the atoms have been allowed to thermalise in the magnetic trap as a function of the bias coils control voltages. This is the metric used to optimise the magnetic trap catch. There is a clear minimum for each bias coil, however, over the ranges investigated, there is only minimal heating seen, at most a 20% increase in the width occurred for the Z bias coil scan. 187

- 6.5 **Uncompressed magnetic trap thermometry and measured coil response** **a)** shows the measured response of the quadrupole coil and Z bias coil during the CMOT stage, PGC and the catch. There is clearly an overshoot of about 15% as the quadrupole coils are suddenly ramped on, but I find this has little effect upon the resulting temperature of the atoms. **b)** Thermometry performed on the uncompressed magnetic trap, 5 seconds after catching the atoms to allow for adequate thermalisation. 191
- 6.6 **Adiabatic movement and compression of the magnetic trap.** **a)** shows how the position of the atoms in the quadrupole field moves as the trap is decompressed for a range of different control voltages for the Z-bias coil. This data is used to minimise the movement of the quadrupole trap whilst loading the hybrid trap. **b)** and **c)** show the dependence of the width of the ballistically expanded cloud on the time over which the magnetic trap is moved and compressed, respectively. Clearly there is a stronger dependence upon the time over which the cloud is moved as opposed to the time over which it is compressed, with the data showing a maximal 50% and 5% increase, respectively. This data was used to determine the rate at which the cloud could be moved and compressed that ensured the cloud did not significantly heat. 193
- 6.7 **Characterising the quadrupole trap gradient and lifetime.** **a)** The measurement used to calibrate the field gradient as a function of current for the quadrupole coils. The data shows the mean fitted position of a cold cloud of atoms for a range of drop times from 10-100 ms, for three different control voltages. The largest variation in position is only 1 pixel $4.5\ \mu\text{m}$ over the range of drop times investigated, suggesting a cancellation of gravity on the order of 1 part in 10^4 . **b)** The fluorescence from the magnetic trap, measured after a 0.5 ms drop using light detuned 20 MHz below resonance, decays as a function of time in the magnetic trap, a reasonable fit to the data (solid line) indicates a lifetime on the order of 73 seconds, this is indicative of a close to vacuum limited lifetime. 196
- 6.8 **RF Antenna and forced evaporation** **a)** The RF antenna, formed by two loops of round 2 mm enamelled copper wire wound about a 65 mm circular former. **b)** To find the initial RF frequency used for evaporation the RF is held on at a constant frequency for 2 seconds and the atom number is measured. 198
- 6.9 **Dipole setup schematic** A schematic of the optical setup used to generate the optical dipole beam used for hybrid trapping and evaporation (left hand side) and the sheet trap (right hand side) used for axial confinement of the atoms for later experiments. All optical elements are as in the cooling and repump laser schematics. 200

- 6.10 **Finding the dipole beam** An absorption image showing how the dipole beam is found through allowing ballistic expansion of the compressed magnetic trap to expand the volume of the cloud whilst keeping the dipole beam on to capture the atoms that are cold enough showing the location of the beam. 202
- 6.11 **Optimisation of dipole location** **a)** and **b)** show the dependence of the number of atoms loaded into the hybrid trap upon the position of the dipole beam. In both **a)** and **b)** there is a clear dip in the number of loaded atoms that arises from precise alignment of the dipole trap with the field zero, greatly increasing the Majorana loss rate. The red dots indicate the position the dipole beam was left at. In **a)**, the horizontal position is scanned whilst the vertical position is fixed at the location of the field zero, and similarly for **b)** the vertical position is scanned whilst the horizontal position is fixed at the location of the field zero. **c)** shows the measured temperature of the atoms in the hybrid trap. As the temperature in the hybrid trap is set by the depth of the dipole trap, translating the lens to maximise the temperature thus maximises the depth at the field zero and thus places the focus at the atoms. 204
- 6.12 **Hybrid load dependence upon dipole power** **a)** and **b)** show how the number of atoms loaded into the hybrid trap and the waist of the cloud after a 25 ms drop time, used as a proxy for temperature, depends upon the power in the dipole beam. **a)** shows a clear increase in number as a function of power until it reaches a value of 0.62, after which it decreases, the decrease is thought to be attributable to an increase in 3 body losses caused by the increased density. **b)** shows a monotonic increase with laser power as expected. 206
- 6.13 **Hybrid trap frequency measurements** **a)** and **b)** shows for 2 different dipole powers the measured (dots) mean position in the vertical direction of a cloud of cold atoms in the hybrid trap after displacing the atoms a small amount from the trap centre, waiting a variable amount of time indicated by the horizontal axis, and then allowing them to ballistically expand for 25 ms. There is a clear sinusoidal sloshing motion of the atoms, the frequency of which is equal to the trap frequency. Solid curves are decaying sinusoidal fits to the data. **c)** shows the measurement of the trapping frequencies of the hybrid trap using parametric resonance, there is a clear decrease in atom number when the frequency of modulation is twice that measured using the sloshing motion, but the decrease is very wide in frequency space. 208

- 6.14 **Hybrid trap lifetime** **a)** shows the atom number decay as a function of time, there is a faster than exponential initial decay. The data in **b)** shows the variation in the waist of the cloud for the same experiments as in **a)**, and show that the waist of the cloud decreases as a function of time indicating some ‘plain evaporation’ occurs. Thermometry after a 7 second hold agrees with this. 210
- 6.15 **Optimal evaporation in the hybrid trap.** **a)** and **b)** show how evaporation in the hybrid trap occurs. In **a)** the evaporation ramp is faster than the optimal rate, leading to evaporation occurring only in one dimension, along gravity, whereas in **b)** evaporation is occurring at the optimal rate, and atoms are seen to evaporate in 3-directions, along the length of the dipole beam as well as along gravity. 211
- 6.16 **First sign of condensation.** Shown are the very first signs of condensation occurring at the end of the 3rd evaporation ramp, indicated by the peaked centre of the cloud, the first signs of a bimodal distribution. 213
- 6.17 **Emergence of bi-modality** **a), b), c)** and **d)** show the emergence of a bimodal distribution in time of flight as the final dipole power is lowered. As the power is lowered the fraction of thermal atoms reduces whilst the condensate fraction increases, although this is difficult to see at this relatively short drop time (35 ms) as the OD becomes very large. 214
- 6.18 **Final BEC and PSD versus N** **a)** shows the final condensate after 100 ms of ballistic expansion. A Thomas-Fermi and Gaussian fit to the data indicate a 90% pure condensate with $\approx 2.2 \times 10^6$ atoms, produced in just under 10 seconds. **b)** shows how the phase space density and temperature varies as a function of atom number during forced evaporation in the magnetic trap and the hybrid trap, the different stages being indicated by the dotted lines. All axes are logarithmic, with the horizontal axis indicating the atom number, the left hand side vertical axis the phase space density, and the right hand vertical axis the temperature. The data shows an almost linear trend throughout the entire process, with an average evaporation efficiency of $\gamma = 3$, indicated by the dashed line. 215

- 6.19 **Final BEC characterisation a)** shows the result of modulating the dipole beam potential for 1 second with a 2% modulation at various frequencies. Only one clear resonance was observed at 29 Hz, which coincides with the trap frequency as measured by a dipole mode oscillation measurement. **b)** shows the voltage measured by a photodiode used to monitor the MOT fluorescence during an experiment. The first 2 seconds show an increase in fluorescence as the MOT loads before suddenly switching off during the PGC stage. After a few hundred milli-seconds the Hybrid dipole beam is switched on and the scatter from the beam is picked up by the photodiode. From ≈ 6.5 s to 9.5 s the power in the dipole beam is seen to ramp down as hybrid evaporation occurs, followed by two pulses that correspond to the imaging beam switching on for absorption imaging. 216
- 7.1 **Interaction of photons with atoms A, B and C** show the three types of interactions photons have with an atom. In **A**, the absorption of a photon by an atom is shown, exciting the electron in the atom into a higher orbital. **B** shows the process of spontaneous emission of photons as an atom in an excited state decays to its ground state. **C** shows how the phase of a photon as it passes through an atom is shifted with respect to a photon passing through vacuum. 221
- 7.2 **Generalised imaging scheme** A generalised imaging scheme is shown above. A probe beam is shown to be incident on the atoms from the left hand side (red arrows), which scatter some of the light (blue arrows). These photons are then captured with a lens and reimaged onto the camera using a 4f configuration. In the fourier plane of the first lens, a grey box represents either an amplitude (t) or phase (ϕ) modulating element. In the case of standard absorption imaging, the element is such that $(t, \phi) = (1, 0)$, in the case of darkground imaging $(t, \phi) = (0, 0)$ and in the case of phase imaging typically $(t, \phi) = (1, \pi/2)$. 223
- 7.3 **Side imaging calibration** A calibration of the effective side imaging camera pixel size is shown above, performed by dropping a cold cloud for various different drop times, and fitting the displacement along the vertical axis. 227
- 7.4 **Taking an absorption image a), b) and c)** show the images captured and used to extract the optical density of a cloud via absorption imaging, with **d)** showing the resulting optical density of the cloud. **a)** shows the first frame taken, the ‘atom’ frame, in which the atoms cast a shadow in the laser. **b)** shows the flat frame, the exact same conditions of the exposure as in the ‘atom’ frame are used, but the atoms have had enough time to disperse. **c)** shows the dark frame, an exposure taken with no laser sources switched on, to capture the background noise, whilst **d)** shows the resultant calculated optical density using frames **a), b) and c)**. 230

- 7.5 **Interline CCD mechanism** The process that allows an interline CCD camera to take very fast back to back images is shown. **A** shows the initial accumulation of charge in the photosensitive pixels of the chip during the exposure. **B** shows how this accumulated charge is then transferred to a non-photosensitive area of the chip to be stored whilst the next exposure, shown in **C** occurs, before the entire chip is read out. 233
- 7.6 **Side imaging a)** A schematic of the side imaging optical system, shown are the two achromatic lenses of 75 mm focal length placed approximately in a 4f imaging configuration. Note the orientation of the achromats changes to minimise spherical aberration and the rectangle in front of the first lens represents the glass of the Science cell. **b)** A Solidworks model of the entire side imaging system, including camera, lenses mounted directly off the camera in 1" lens tube, and the two translation stages used for positioning the camera. 238
- 7.7 **Bottom imaging a)** A schematic of the bottom imaging system, shown are the custom objective lens as well as the three achromats, which overall provide a magnification of $40\times$. Note all pairs of lenses are placed in a 4f configuration, however, this is only particularly critical to the performance for the first pair of achromats. 239
- 8.1 **Metrics for a uniform potential** Shown above is how the shape of a potential, approximated by a power law, varies as a function of the power. Clearly as the power increases the potential becomes more uniform. Indicated in the bottom left hand side of the figure are the widths, R1 and R2, over which the potential increases from 10% to 100% for the X^{10} and X^6 potentials respectively, clearly as the power of the potential increases this width decreases. 245
- 8.2 **Approximating an LG mode** Shown above the fractional error in approximating an $LG_{(l,0)}$ mode by an $x^{|l|}$ mode is shown for $l = 1, 10$ and 100 , for $x \in [0, 1]$ where x represents the scaled variable $\frac{r}{w}$. Clearly, as the azimuthal index increases the Laguerre-Gaussian mode more clearly resembles a power law potential. 249
- 8.3 **A dynamic uniform potential** shown above are the amplitude and phase profiles of two opposite handedness Laguerre-Gaussian modes of $|l| = 40$, and the resulting amplitude profile that arises from their interference. The interference of the two modes produces a ring with corrugations, whose position is determined by the relative phase of the modes. By changing this relative phase, ϕ , the corrugations can be made to rotate about the edge of the trap. 251

- 8.4 **Rotating trap schematic.** A schematic of the optical setup used to create a sharp walled potential for the atoms, with the ability to also rotate corrugations upon the wall. The output of a 1 W, 532 nm laser is split into two beamlines using two polarising beam splitter cubes and half-wave plates that are then double passed through separate AOMs. The resultant diffracted modes are then magnified before diffracting off two separate holograms to generate two $\text{LG}_{(120,0)}$ modes that are then combined upon a non-polarising beam splitter cube before being magnified by a final telescope. 253
- 8.5 **LG hologram and resulting profile a)** The recorded hologram used to generate an $\text{LG}_{(120,0)}$ mode from a Gaussian mode, illuminated by the beams used to record it. The photo shows a square (the photopolymer film) upon which the hologram is recorded, whilst the hologram shows a small hole in the centre that arises due to the dark centre of the $\text{LG}_{(120,0)}$ mode, and alignment of the two beams is performed by ensuring the hole in each beam overlaps. **b)** and **c)** show the resulting intensity pattern produced by the $\text{LG}_{(120,0)}$ hologram, and the intensity resulting from the interference of the $\text{LG}_{(120,0)}$ and the $\text{LG}_{(-120,0)}$ mode. 255
- 8.6 **Aligning the ring trap a)** and **b)** show the effect of passing through the first telescope off axis, showing the ring clearly deteriorating when aberrations are introduced, whilst **c)**, **d)** and **e)** shows the ring at varying stages of defocus, with **e)** showing the final resultant focussed ring. Note that these images were taken with auto-exposure on, so the exposure of each image is scaled to prevent saturation. 256
- 8.7 **Sheet trap ray trace** A ray trace showing how a pair of cylindrical lenses and a pair of spherical lenses are used to tightly focus the beam in the vertical direction whilst collimating the beam along the horizontal axis. 259
- 8.8 **Finding the ring trap a)** and **b)** show time of flight absorption images used to find the uniform ring trap, **a)** shows a condensate that has been loaded into the sheet trap and imaged after 100 ms of ballistic expansion without the ring on whilst **b)** shows the effect of switching the ring trap on, confining the atoms radially whilst they can still expand axially, allowing the shape of the trap to be visualised. 260
- 8.9 **Loading the uniform trap a)** shows an image of the condensate *in-situ*, taken from below before it is loaded into the ring trap whilst **b)** and **c)** shows the condensate imaged *in-situ* when loaded into the ring trap formed by each of the LG-modes separately. 262

8.10 Atom number in the uniform trap a) shows an image of the condensate <i>in-situ</i> in the combined ring and sheet trap when viewed from the side, clearly a large fraction of atoms remain in the sheet but not in the ring. b) shows the atom number as a function of the power in the $LG_{(-120,0)}$ beam, as the power in the second LG mode is increased, the corrugations caused by the interference of the two modes allow atoms to leak out, as the effective potential height of the trap is reduced.	264
8.11 Perturbing the uniform trap a) shows the mean position of the cloud, as measured using the bottom imaging system, when displaced from the centre of the uniform trap by applying a bias field. Clearly the atoms move back to their central position and remain there without sloshing as would be the case in a harmonic trap. b) shows how the atom number in the uniform trap decays as a function of the frequency at which the power in the ring is modulated.	265
8.12 Rotating a condensate a) and b) shows images of the condensate released from the uniform potential after 50 ms of ballistic expansion, imaged from below, which clearly show the presence of one and two vortices respectively, whilst c) shows the effect of rotating the corrugations generated by the interference of the two $LG_{(120,0)}$ modes at different frequencies.	268

Introduction

Bose Einstein condensates (BECs) of weakly interacting, dilute gases are a relatively new phase of matter, with their first observation in 1995. A coherent, macroscopic sample of atoms, such as BECs, are the atomic analogues of lasers, with atoms replacing photons, and have allowed a range of new physics to be investigated. Since their first realisation there has been a flurry of research, as a large range of possibilities have opened up.

From a fundamental physics point of view, BECs allow quantum mechanics to be studied on a macroscopic scale, with recent demonstrations of matter-wave interference of wave-packets separated by up to 54 cm over 1 second[9]. The macroscopic size of a BEC allows almost complete control of the system experimentally, as the relevant time and length scales for condensates are all within experimentally achievable means, with timescales on the order of μs to ms, and length scales on the order of μm to mm. This has given physicists a quantum system in which the potential, typically formed from magnetic or optical fields, can be dynamically manipulated on a local scale, and the results of these changes be observed, providing a testbed for a range of physics. This ability to manipulate BECs has improved our ability to study a range of other quantum systems, as the physics of BECs is analagous to a range of other physical systems which do not have such a high degree of control, thus through studying the behaviour of a BEC as its potential is manipulated we can also study how analogous systems behave.

In a similarly recent amount of time our ability to control optical fields has increased an enormous amount, with the advent of spatial light modulators (SLMs) and digital

micromirror devices (DMDs) allowing us high resolution control over both the phase and amplitude of the optical field. As far off-resonant optical fields can provide a potential in which BECs can be confined, this increased ability to manipulate light provides a promising avenue for studying condensates in a range of different potentials.

1.1 Structured potentials for cold atoms

The first demonstration of optical fields used to trap and manipulate particles was by Arthur Ashkin in 1969[10], where he demonstrated that a focussed laser beam could be used to trap and accelerate micron sized particles. This work led to him showing that structured optical fields could be used to manipulate viruses and bacteria[11], opening up new pathways in biology for the study of these systems. Later, he also demonstrated the first use of the optical dipole force to trap sodium atoms[12], opening up an entire new path to achieving condensation, and the ability to confine atoms non-magnetically. This demonstration of the uses of structured optical fields in such a wide range of disciplines eventually lead to him receiving a Nobel prize in 2018. Before Ashkin could realise the potential of structured optical fields the laser had to be invented, as well as the laser cooling of neutral atoms for trapping neutral atoms. This is due to the very small momentum a single photon carries, and the very weak coupling between neutral atoms and the electromagnetic fields, thus to manipulate and trap particles extremely high intensities of light are required, and atoms of extremely low thermal kinetic energy are required.

The optical dipole potential arises from the interaction of an atomic dipole moment with light far detuned from resonance. This interaction is very weak compared to other interactions used to capture atoms, such as near resonant radiation pressure traps, magnetic traps or electrostatic traps, this makes the traps very weak. One key benefit of optical dipole traps however is that the optical excitation can be kept very low, with the trapping force arising from the dispersive interaction of the induced dipole moment with the light field, as opposed to it arising from the absorption of photons.

For an atom interacting with a far detuned optical field, detuned by an amount Δ from resonance ω_0 , it can be shown that the potential the atom sees is given by[13]

$$U_{dip}(r) = \frac{3\pi c^2}{2\omega_0^3} \frac{\Gamma}{\Delta} I(r) \quad (1.1)$$

where Γ is the excited state linewidth, which determines the rate at which photons can be scattered. The dipole force arises from the gradient of this potential, and is thus a conservative force, which is clearly proportional to the gradient of the intensity and inversely proportional to the detuning.

There are two important consequences of equation. 1.1. The first is that as the potential is proportional to the intensity of the light. This implies we can structure the potential seen by our atoms if we can structure the intensity of the optical field. The second important consequence arises from the dependence on the detuning of the light. As we can detune our light both above and below resonance, the detuning can thus have both positive and negative values, implying that we can create both attractive and repulsive potentials by choosing the frequency of the optical field correctly.

1.1.1 Structuring light

For millennia humankind has attempted to structure optical fields for useful purposes, the most obvious example of which are lenses, which have been used to transform optical fields by focussing the light, creating an image that can then be viewed with magnification. Until recently we have been quite limited in the range of optical fields we can produce however, as objects that manipulate the phase or amplitude of an optical field have mostly been limited by our ability to machine materials that will act as phase or amplitude objects, typically leading to spatial resolutions on the order of 10's of microns, and prototyping with these methods is time consuming and expensive. As well as limitations to the manufacturing of these elements, the ability to simulate the effect that these objects have upon an optical field are typically computationally expensive, requiring ray tracing of a great number of rays to adequately simulate the effect the object has upon the optical field,.

Recently though, with the advent of the spatial light modulator (SLM) and the digital micromirror device (DMD), we have gained the ability to spatially manipulate, with high resolution, both the phase and amplitude of an optical field in real time. This has greatly reduced the cost in both time and money of producing optical filters, as it has allowed the rapid prototyping of different phase and amplitude filters inexpensively, allowing optical filters that generate arbitrary complex optical fields to be quickly realised. This ability to create arbitrary 3-dimensional patterns has been furthered through a number of computational efforts in holography. Techniques

have now been developed that can manipulate both the phase and amplitude of an optical field with a single phase or amplitude only hologram[14][1], useful when a known 3-dimensional optical field is to be generated, but these can suffer from low efficiency in their use of light. Methods which create arbitrary 2-dimensional amplitude patterns from a single phase element in the far-field[15][16] have shown high efficiency in their use of light when patterns are confined to a plane, whilst other techniques have allowed the independent control of both the phase and amplitude along curves in 3-dimensional space using phase only filters[17][18].

Using these advances in both the ability to manipulate optical fields in real time as well as the computational efforts in holography we can thus, almost arbitrarily, structure the potential that atoms see in an optical dipole trap.

1.1.2 The uses of structured light

Arguably one of the most beautiful theorems in physics is Noether's theorem[19], which states that with every symmetry of a system there is a corresponding conservation law; for example the homogeneity of space leads to the conservation of linear momentum whilst the isotropy of space leads to the conservation of angular momentum. The homogeneity and isotropy of the geometry of space thus determines key conservation laws that must be obeyed by the physics of the system. This is one example of how the geometry of a system plays a large role in the physics the system displays.

Another clear example of how the geometries of a system affects the resulting physics is that of either a particle in a uniform box compared to a particle in a harmonic potential. In the case of a particle in a box, the energy of the n^{th} level is proportional to n^2 , whereas in a harmonic oscillator they are proportional to n and are thus equidistantly spaced. Thus the energy spectrum of these two cases is clearly different, indicating that a BEC confined in these different potentials would have different excitation spectra. As atoms can be manipulated with light, structured optical fields can thus tailor the geometry of our cold atom systems allowing different geometries to be realised, allowing the study of new and interesting physics.

A key paper that motivated the work during my candidature was produced by A. Groszek *et al.*[20], in which they theoretically investigated the dynamics of an initially random distribution of quantised vortices in a highly oblate BEC in two different potentials, a harmonic potential and a uniform potential approximated by

a power law potential. Through their computations they saw the formation of large scale vortex clusters, known as Onsager vortex clusters, in the uniform potential, but no evidence of clusters in the harmonic potential. They attributed this as arising due to the energy of the vortex configurations, which depends upon the local density of the atoms in the trap and thus depends upon the confining potential.

The direct observation of the formation of Onsager clustering in a 2-dimensional system has been a goal for many years, as it provides evidence for the realisation of an inverse energy cascade. In the inverse energy cascade energy flows from small length scales to large length scales, from a disordered array of vortices with minimal net flow through to a cluster of vortices about which there is a large net flow. This is in contrast to the usual direct energy cascade in 3D systems where energy flows from large lengths to small lengths before being dissipated. Experimental observation and the theoretical study[21][22][23] of the inverse energy cascade in an enormous range of systems, from very small scale systems such as superfluid helium[24], everyday systems such as soap films[25] through to some of the largest systems[26], has generated a large amount of interest in recent years and thus creating a uniform potential to attempt to demonstrate this phenomenon is one of the key aims of this thesis.

Structured potentials have played a large role in studying a range of other phenomena using ultracold gases. One example is that of a uniform potential as when one wants to use cold atoms to study condensed matter physics. Cold atoms have found a niche within the condensed matter community, as the physics is typically analogous to condensed matter systems, however the cold atom system itself is much more controllable. However conventional condensed matter systems are uniform in density whereas condensates adopt the shape of the potential in which they are confined, and thus are typically not uniform in density. Methods have been made to circumvent this issue, such as by making local density approximations[27], or by only probing a small portion of the cloud[28], however these methods are limiting when studying critical phenomena where the correlation length diverges[29][30]. By using a uniform potential a uniform system of cold atoms can be produced and these problems removed, making the condensate a better test bed for understanding the physics.

Other examples of the use of the physics that can be studied using structured optical potentials such as optical lattices includes; many body physics in periodic potentials, or with reduced dimensionality[31], the Berezinskii-Kosterlitz-Thouless crossover[32], classical magnetism[33] as well as the phase transition from a su-

perfluid to a Mott insulator[34]. Two dimensional structured potentials have also played important roles in cold atom experiments such as in atomtronic circuits where cold atoms are confined to different potentials to make analogs of electrical circuits[35], ring traps, where persistent currents have been observed and the equivalent of a Josephson junction has been made[36] as well as the quantised decay of vortices observed[37] and double well potentials, used to demonstrate an atom interferometer[38]. Three dimensional structured potentials have played similarly important roles in cold atom experiments, having been used to demonstrate Anderson localisation in three dimensions, using disordered potentials[39] as well as creating as exotic structures as the Eiffel tower or atoms trapped in knotted configurations[40], demonstrating unprecedented control of the potential of cold atom systems.

Reference. [41] provides an overview of the use of structured fields in a range of other areas of physics, from its uses in microscopy and metrology through to both classical and quantum optical communications. Structured light currently plays a large role in a range of different areas of physics, and with increasing accessibility to the ability to arbitrarily structure light at a reduced cost, the use of structured light throughout different areas of science will likely only grow.

1.1.3 Holographically generated structured optical fields

Although there are a number of methods by which one can generate an arbitrary complex optical field via an SLM or DMD, these methods have a number of drawbacks. One common problem is the efficient use of the light available, with DMD's and SLMs suffering from the pixellated nature of the device, with dead space between pixels absorbing the incident light and the regular pixellated pattern of the device generating a large number of diffracted orders, much like a grating, which reduces the efficiency¹. These reductions in efficiency also reduce the overall power that can be used to manipulate atoms, as damage thresholds for these devices are typically on the order of a few W cm^{-2} for an SLM, and slightly higher at a few tens of W cm^{-2} for a DMD. Alternatively, fabrication of optical elements from glass or plastics, or the nanofabrication of diffraction gratings can circumvent these efficiency problems, however prototyping is expensive and time consuming.

In this thesis I present an alternative solution, combining the versatility of an SLM

¹These problems can be reduced, but not eliminated, by purchasing better engineered devices, but this leads to costs on the order of \$50k

(or alternatively a DMD could be used) with the high efficiency that a volume diffraction grating can provide. To achieve this I first utilise an SLM to generate the desired optical field, and then record this optical field holographically, by interfering it with a Gaussian beam, into a relatively new inexpensive photopolymer, Bayfol HX[42]. This new photopolymer can achieve large refractive index modulations, $\delta n > 0.03$, and is thick ($\approx 16\mu\text{m}$), allowing volume holograms to be recorded which can achieve diffraction efficiencies greater than 98%. The other key benefit of this film is that it is designed for mass production, with the ability of contact copying, and recording and subsequent production of the hologram is all optical, making it incredibly simple to work with.

The recorded hologram can then be used to generate the desired optical field by simply diffracting a Gaussian beam off the hologram, resulting in the desired optical field, but does not have the inefficiencies associated with the pixellated nature of an SLM or DMD as these can be spatially filtered before recording the hologram. This method of generating a desired optical field via an SLM and recording it in a hologram which can subsequently recreate the optical field has an immense benefit to a range of applications of structured optical fields, as it allows rapid prototyping, followed by inexpensive production of an optical element that can produce the desired optical field². Not only does this method reduce the cost in the optical elements required to generate the desired optical fields, it also reduces the number of optical elements required, with a single hologram replacing a number of lenses required for beam expansion to fill the SLM or DMD and the subsequent spatial filtering of the mode to remove the effects of pixellation, both saving cost, time in alignment, and the ever valuable ‘real estate’ of an experiment. These holograms could thus replace SLMs or DMDs wherever static structured fields are required and would find their greatest use in the commercialisation of ultracold atom research.

Although the field of cold atoms has now been around for several decades, the commercialisation of the research is still in its infancy, however the precision measurements that cold atoms have made as gravimeters[43][44], gyroscopes[45][46], inertial sensors[47][48], magnetometers[49] and pressure sensors[50] are very promising indications that compact cold atom sensors could provide the next generation of commercialised precision sensors. Our holograms would find their application in two key areas of this commercialisation; in the creation of the cold atom samples, and in the final optical traps which may be used.

²A 1 cm^2 piece of photopolymer costs ≈ 10 cents

As already noted by others[51], one issue that arises with commercialising a cold atom apparatus is the large size of the apparatus required, partially due to the magneto-optical trap (MOT) that requires 4+ beams to be accurately aligned. Methods have been developed to decrease the size required, the most popular method being to generate all the required beams from one beam and typically are either reflective[52][53][54] or diffractive[51][55][56]. Use of our method to generate holograms to shape one beam into multiple would fall under the category of a grating MOT, one generated via diffraction, however the use of the SLM to generate the beams would allow the MOT to be rapidly optimised for any given conditions, removing the time and cost associated with prototyping with nanofabricated structures.

In terms of the final optical traps used there is a huge list of possible applications, including; generating ring traps used for rotational sensing[57] or large area atom interferometers[58], confining atoms to hollow beam blue-detuned traps allowing longer interrogation times and reduced perturbations[59][60][61], double well potentials for atom interferometry[38], and spatially periodic lattices for improved measurements of gravity[62], as well as the ability to generate arrays of the above potentials for spatially resolved measurements.

There is clearly a huge range of applications for holographic elements that can generate arbitrary optical fields with high efficiency, and as such it is the aim of this thesis to develop the techniques required to create these holographic optical elements, and demonstrate that they can perform better than an SLM.

1.2 Thesis outline

The initial aim of my PhD has been to create novel optical potentials for manipulating and studying Bose-Einstein condensates with the focus being upon the change in behaviour of the condensate as a result of the potential to which it is confined. However, it seems that things rarely go according to plan in experimental physics, as new ideas are discovered through the process of solving problems whilst trying to achieve your original aim. This has resulted in a greater focus on how novel complex optical fields can be created and recorded into holographic optical elements, as well as how Bose-Einstein condensation is achieved. As such, with particular reference to the chapters focussed on the steps required for achieving Bose-Einstein condensation, I have written this thesis with the mindset of creating a text which I would have desired when I first began designing and building my apparatus, outlining all

the theory I deem necessary for producing a BEC, as well as how this is put into practice, and optimised for the production of BEC.

In this thesis I will present a novel method for the creation of holographic optical elements (HOEs) which can be used to generate arbitrary complex optical fields with both high efficiency and fidelity. I shall also demonstrate how these HOEs can be used to confine a Bose–Einstein condensate to interesting potentials with the benefit of being much less complex than other methods that are currently used. To demonstrate this I have built an apparatus capable of creating ^{87}Rb condensates of 2.2×10^6 atoms in just under 10 seconds, and confined these condensates to optical potentials generated holographically. The optical potential of most interest that has been generated is that of an $\text{LG}_{120,0}$ mode, which creates a very uniform potential for studying condensates in.

This thesis is divided into 8 chapters, the first of which is this introduction in which I have briefly reviewed some of the basics of Bose-Einstein condensation and the thermodynamics of a Bose gas, as well as discussed the role that interesting optical potentials can play in studying a range of different physics. In Chapter 2 I detail how one can use a phase-only SLM to create arbitrary complex optical fields through the use of computer generated holograms (CGH). I then go on to describe how I implemented these methods to create an optical vortex knot and with this interesting optical field I demonstrate the principle of wave-particle duality in 3-dimensions, a result that was published in *Scientific Reports*.

Chapter 3 begins with an outline of some of the issues that arise when using an SLM to create arbitrary optical fields, and then describes how these issues can be avoided through the use of HOEs as opposed to an SLM for the spatial manipulation of the optical field. I then go on to describe the process of recording HOEs and how to maximise the diffraction efficiency of a HOE for given recording conditions. Finally, I demonstrate that these HOEs can achieve close to the highest efficiency possible a single phase only element can achieve for generating an arbitrary optical field from an initial Gaussian beam, this work was subsequently published in *Optics Express*.

At this point during my candidature I had demonstrated the ability to create arbitrary complex optical fields, as well as a method for creating HOEs that could reproduce these fields with similar fidelity but much higher efficiency than the SLM could, and with a much simpler optical setup. The next stage would be to use the holographically generated optical fields to manipulate a condensate, however at the time it was decided that our current BEC apparatus was not up to the task, and

as previous experiments were taking longer than expected to be completed it was decided that I would build a new apparatus. As such, the next three chapters are dedicated to the achievement of Bose–Einstein condensation in this new apparatus.

Chapter 4 begins with a description of the vacuum requirements of a BEC apparatus, and then delves into the details of the design of the vacuum chamber. This chapter has a focus upon the requirements placed upon the vacuum system pressure as well as how this is related to the vacuum pumps, and how regions of different vacuum pressures can be created through low conductance paths leading to differential pumping. I also detail how ultra-high vacuum is achieved in my apparatus, including details of the cleaning procedures involved as well as the process followed during bakeout. The vacuum system section is followed by a description of the laser systems required for the laser cooling of ^{87}Rb , including details of how the lasers are locked to the appropriate transitions as well as the geometric layout used to create all the necessary frequencies of light for cooling, repumping, pushing and imaging the atoms. The chapter ends with considerations of geometry, switching speed and power dissipation on magnetic coil design for confining atoms to quadrupole potential, as well as the necessary electronics for driving the coils as well as other elements of the BEC apparatus.

Next, Chapter 5 describes the theory behind the laser cooling of atoms as well as the experimental implementation of the theory. In the theory section I describe the scattering force that arises from two counter propagating plane waves leading to optical molasses, and extend this theory to include differing magnetic sublevels and polarisation, resulting in both the confinement and cooling of atoms in a magneto-optical trap (MOT), and consider some experimentally relevant quantities. I then consider sub-Doppler cooling mechanisms, as well as the spatial compression of atoms in a MOT arising from large magnetic field gradients, and how these two processes can be used to maximise the initial phase space density (PSD) of the atoms before they are loaded into a pure magnetic trap. Although this theory behind laser cooling might be very familiar to some, I believe there are useful experimental insights within the theory section that will be particularly useful to anyone constructing a new BEC apparatus. In particular I believe that it is easy for the context of some of the equations to be forgotten leading to their incorrect use, and insights within this chapter may help clear some of these issues. The final section of this chapter focusses upon the process of optimising laser cooling for the production of BEC, and thus is particularly useful as both a benchmark and a guide for anyone constructing a new apparatus.

Chapter 6 then describes the process of forced evaporation in a confining potential. I begin with a theoretical model of how forced evaporation is implemented and how key thermodynamic quantities relevant to Bose-Einstein condensation change during evaporation, as well as describing key metrics for the efficiency of evaporation. I then apply this model to forced evaporation in a quadrupole magnetic trap and an optical dipole trap, which also requires consideration of how atoms are confined to these potentials. Next, I describe some of the experimental considerations of forced evaporation in these traps and consider the case of the hybrid magnetic-optical trap first implemented by Lin *et al.*[8], and how this can remove some of the negative aspects of each case. Finally the chapter ends with the results of catching the atoms in the laser cooled cloud in a purely magnetic potential, implementing radio-frequency (RF) forced evaporation in the quadrupole trap, loading of the hybrid trap, and the resulting evaporation to BEC in the hybrid trap.

In Chapter 7 I describe one of the most critical aspects of a BEC apparatus, the imaging of a cold atomic cloud. Accurate quantitative imaging of a cold atoms is critical to the achievement of condensation, as imaging atoms is the only method we have for gaining quantitative information of what is happening. In this chapter I describe the theory behind the interactions between atoms and light, and then go on to consider the implications for absorption and fluorescence imaging. I then detail how imaging is implemented on my apparatus, and how it can be quantitatively calibrated and used to measure properties of the cloud such as atom number and temperature.

In the final chapter of my thesis I show how the HOEs I have created can be used to confine a condensate to a uniform potential, and characterise the resulting potential.

This thesis has been ordered somewhat chronologically with regards to the order in which I carried out the work during my candidature. However where necessary I have included theory with experiment rather than in chronological order, as I believe this makes more sense to the reader. For example, in the considerations of the design for the vacuum system detailed knowledge of laser cooling is required to understand the optical access that is required and the constraints this places upon the vacuum system design, however it does not seem logical to combine the theory of laser cooling with the design of the vacuum system as opposed to with the experimental implementation of laser cooling.

One further exclusion to the ordering of this thesis is the chapter on imaging cold atoms. The results of the imaging of cold atoms is necessary to the optimisation

of a BEC apparatus, as the results of the imaging are the only quantitative metric we can use for the optimisation of the function of the apparatus. However the quantitative calibration of an imaging system is much simpler with very cold dense clouds, thus there is necessarily a recalibration of the imaging system once one has a cold cloud. As such the chapter on imaging is placed after condensation is achieved as the calibrations performed are most accurate with a cold cloud, and as the chapter stands independently of the other chapters within the thesis this seemed most appropriate.

One final note to any reader whom may use this as a guide to creating a BEC apparatus of their own is that, like most processes in life, the progress of the design, construction and resultant optimisation of a BEC apparatus is rarely linear, typically not even monotonic, with one needing to frequently revisit earlier stages even once one has progressed, and that they should not be disheartened by setbacks, but rather use those setbacks to learn more about their unique apparatus and cold atom physics.

Complex control of an optical field

Some of the work presented in this chapter, as part of my PhD research, has previously been published in:

- *S. J. Tempone-Wiltshire, S. P. Johnstone, K. Helmerson, “Optical vortex knots - one photon at a time”, Scientific Reports **6**, 24463 (2016).*

The first step towards trapping a Bose-Einstein condensate in an optical potential requires the ability to arbitrarily control the complex valued optical field. Most commercial devices for manipulating optical fields either manipulate the phase, such as a spatial light modulator (SLM), or the amplitude, such as a digital micromirror device, of the optical field, not both properties. As such, other methods have to be implemented to gain full complex control of the optical field. In this chapter I shall outline the methods I used for achieving full complex control of the optical field by combining an SLM with phase only computer generated holograms (CGHs), and the application of this method to generating an optical vortex knot; an intricate three-dimensional optical field in which optical vortices twist themselves into a non-trivial topological structure, one photon at a time.

2.1 Spatial light modulators for phase control

In the past there were a number of methods to manipulate light for many purposes: the focusing of a beam for optically manipulating of atoms, the magnification of an optical field for viewing, or the creation of non-trivial beam profiles for cold atom

trapping, all which relied upon certain 'standard' optical elements such as lenses, axicons and diffraction gratings, typically formed by machining glass. In recent years with the development of liquid crystal on silicon (LCoS) SLM devices, the ability to control the wavefront of light has increased considerably.

Liquid crystal SLMs have become commonplace amongst society, with most people using at least one device that relies upon them every day, such as a phone, computer or television which utilises a liquid crystal display¹. SLMs consist of an array of pixels, within which is a small 'fluid' of liquid crystals. By applying a small voltage to these crystals, the crystals can be made to align with the electrical potential, which creates a birefringence in that pixel which, in most standard technological configurations, can be used to rotate the polarisation of the incoming light onto that pixel. Coupled with a linear polariser directly after the SLM, this creates amplitude modulation of the light exiting the pixel, and is how most liquid crystal displays work.

As we desire full complex control of the optical field, we must combine the SLM with CGHs to achieve this. Given this, we configure our SLM in a phase only modulation mode by aligning the polarisation axis of the input light with the axis of the liquid crystals. In this configuration applying a voltage to the liquid crystals simply results in a relative phase delay between light passing between neighbouring pixels, and thus spatial modulation of the phase of the optical field can be achieved.

For the experiments performed in this thesis we use an SLM that has been removed from a Hitachi CPA100 projector, with all unnecessary elements of the projector removed (light source, casing etc.) and the projector configured as a second display². The Hitachi CPA100 contains three SLMs, one for each colour channel, red, green and blue. Only a single SLM was required for this work, thus the SLM optimised for green was used and only the green channel of the display addressed the SLM. The SLM in this projector consists of 1024 X 768 pixels, however, typically only the central 768 X 768 pixels are addressed due to the symmetries of our desired modes.

The Hitachi projector and the resultant SLM and electronics contained in a 3RU box are shown in figure. 2.1. To make the projector work for our purposes, the casing was opened and all unnecessary optics and electronics was removed, with what was

¹The underlying technology of an SLM and LCD is the same, but some may draw distinctions between the two.

²i.e. we connect our 'projector' to the computer as a second display via a VGA cable, as well as a third display that duplicates what is shown on the projector so as to allow us to see what is displayed on the SLM, and then simply address the projector as one would a regular computer display.

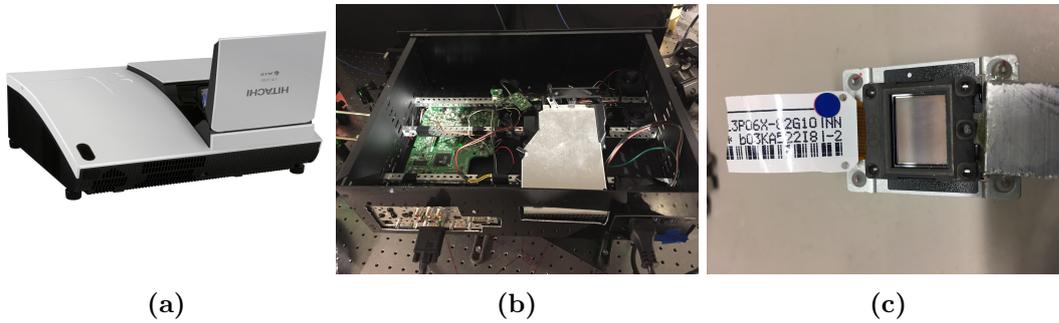


Figure 2.1: Scavenging an SLM from a projector a), b) and c) show the original Hitachi CPA100 projector from which I scavenged an SLM from, contained within the 3RU box with the SLM attached beneath with the connector plate visible and one of the SLMs as removed from the projector.

left over being moved into a 3RU box, and the SLMs attached through a hole cut in the lid of the box.. The connector plate from the projector casing was also removed and fixed to the 3RU box so that addressing the SLM was as simple as addressing an additional display on the computer.

To make the projector work a number of additional sensors had to be bypassed³, such as those that checked to see if the mirror had been moved into place, that the lamp was turning on (when no lamp was present) and that the entire system remained within its thermal limits. These sensors were found one by one and shorted so that the projector would switch on without requiring these features. Once this had been done the projector could then be turned on, connected to a computer via a VGA cable, and the SLMs could then be used as an additional display for the computer.

As the phase response of an SLM depends upon many factors such as; the wavelength of light used, the angle of incidence, and properties of the SLM itself, such as the depth of the pixels, the phase delay as a function of applied pixel value must be measured, yielding a look up table (LUT) of pixel values to corresponding phase shifts, to be able to accurately generate the desired holograms. The LUT of our SLM was measured to have a maximum phase delay of 1.19π at 532nm, and was measured using the methods described in Ref. [63]

³Components such as the lamp, a motorised mirror, and thermistors cannot simply be removed, as the sensors are there to prevent the projector running when a fault is present.

2.2 Methods for full complex control of an optical field

To gain full control of the complex optical field, some technique must be used to allow amplitude modulation of the optical field as well as the phase modulation the SLM can provide. If using a single phase-only SLM this is typically achieved in one of two ways. One method is to split the SLM display in half and use one half of the display to modulate the amplitude of the optical field, using a phase mask that produces the desired amplitude modulation in the Fourier plane, and then refocussing that field onto the second half of the display and using it to generate the desired phase. This method has been demonstrated by A. Jesacher *et al.*[64], however, it suffers from passing through the SLM twice and results in twice the losses from the SLM, which can be as large as 50% on each pass. The method we chose to employ was to use a single, phase only, CGH wthat can perform both phase and amplitude modulation of the optical field, first described by J. Kirk and A. Jones[14] and later extended to the encoding method we employ, which only requires 1.17π of phase modulation, by V. Arrizon *et al.*[1].

2.2.1 Encoding a complex optical field with a single phase only filter

To qualitatively describe how a single phase only filter may generate an arbitrary complex optical field, we analogise the filter to how the radio is transmitted. First, the phase of a diffraction grating is applied to the incident field, which acts to generate a number of diffraction orders from the incident field. These resultant diffracted modes can be thought of as akin to the carrier frequency that the audio signal for the radio is transmitted upon. This modulation depth of this grating is typically π , which diffracts the most light possible into the desired mode. Next, the local phase and amplitude of the diffraction grating is then modulated to provide the desired field, similar to how radio either uses amplitude modulation, or frequency modulation of the carrier to transmit the audio signals. For instance, if we consider a pixel at coordinates (x,y) , where the amplitude of the optical field at that point is 50% of the maximal intensity, then the diffraction grating modulation is reduced to 50% also (i.e. $\frac{\pi}{2}$). Whereas if the phase between pixel (x,y) and $(x+1,y)$ varies by π , then that phase is also added to the grating (modulo 2π) shifting the relative position of that grating from perhaps a maximum to a minimum. The analogy here

most accurately describes the original method of J. Kirk and A. Jones[14] for a single, complex valued phase only filter. V. Arrizon *et al.*[1] developed this method further, generating a number of different methods of encoding complex valued filters in phase only filters, and it is their methods that will be applied in this thesis.

To describe how this complex filter works, we begin by considering the complex optical field after it has transmitted through the SLM, and assume an incident uniform optical field of constant phase. Assuming we are attempting to create a complex optical field, $s(x, y) = a(x, y) \exp[i\phi(x, y)]$, the optical field upon exiting the SLM can thus be described by

$$h(x, y) = \exp[i\psi(a(x, y), \phi(x, y))] \quad (2.1)$$

where ψ is the phase imparted by the SLM which is dependent upon the amplitude, $a(x, y)$, and phase, $\phi(x, y)$, of our desired complex scalar field, both functions of position, (x, y) . The aim is to thus find an encoding function, $\psi(a, \phi)$ that will realise our desired optical field.

Noting the analogy above, we wish to only shape the first diffraction order from our grating into the desired mode, so at this point we shall expand the optical field into a Fourier series in the variable ϕ

$$h(x, y) = \sum_{q=-\infty}^{+\infty} c_q^a \exp(iq\phi) \quad (2.2)$$

where the coefficients of the expansion, c_q^a , indexed by q , are given by

$$c_q^a = (2\pi)^{-1} \int_{-\pi}^{+\pi} \exp(i[\psi(a, \phi) - q\phi]) d\phi \quad (2.3)$$

and where we have included the superscript ‘ a ’ as a reminder that the coefficients of the expansion are also dependent upon the amplitude of the desired optical field.

From equation 2.2 it is clear that the first Fourier term, $q = +1$, already has the desired phase and thus it is left to determine how the spatially dependent coefficients can be used to generate the desired amplitude modulation. We thus need to satisfy the condition

$$c_1^a = Aa(x, y). \quad (2.4)$$

This condition is typically referred to as the encoding condition, and includes a

normalisation constant, A , so as to allow for encodings that do not diffract all the light into the desired mode, but still recreate the desired optical field⁴. To proceed further, we next explicitly calculate the first coefficient,

$$\int_{-\pi}^{\pi} [\cos(\psi(a, \phi) - \phi) + i \sin(\psi(a, \phi) - \phi)] d\phi = 2\pi Aa(x, y). \quad (2.5)$$

This integral form of the coefficient has been arrived at by simply using the standard definition for the coefficients of a harmonic Fourier expansion of a function, and can be found in any mathematical textbook describing Fourier series. With this form we next note that as the amplitude function, $a(x, y)$, is a real scalar function we thus require the imaginary component of the coefficient to vanish. This yields two sufficient and necessary conditions to fulfil the encoding condition;

$$\int_{-\pi}^{\pi} \sin[\psi(a, \phi) - \phi] d\phi = 0 \quad (2.6)$$

$$\int_{-\pi}^{\pi} \cos[\psi(a, \phi) - \phi] d\phi = 2\pi Aa. \quad (2.7)$$

From observation of equation 2.6, if ψ is chosen to be an odd function in the variable ϕ then our first condition will be satisfied, and thus we are left with choosing a function that satisfies equation 2.7⁵.

There are numerous odd functions that will satisfy equation 2.7, it is now up to us to determine an explicit form. Here we shall derive the form for the encoding method used throughout this thesis, originally developed by V. Arrizon *et al.*[1]. The chosen encoding function has the form

$$\psi(a, \phi) = f(a) \sin(\phi) + \phi \quad (2.8)$$

note this function is still odd in the variable ϕ , and thus satisfies our first condition. If we then substitute this form into our original transmittance function, equation.2.10,

⁴As A is spatially invariant, the generated optical field has the same spatial variations of amplitude and phase as the desired optical field, just not the same peak intensity as the incident optical field.

⁵It is noted that there are numerous encoding functions that satisfy these conditions, and thus all realise the same desired optical mode in the first Fourier coefficient, what differs between these encoding methods is the spatial amplitude of the higher order Fourier coefficients

we have

$$h(x, y) = \exp(i(f(a) \sin(\phi) + \phi)) \quad (2.9)$$

which if we separate into two terms, one containing the phase term and the other containing the combination of the phase and amplitude, we have

$$h(x, y) = \exp(i(f(a) \sin(\phi)) \exp(i\phi)). \quad (2.10)$$

Next we expand our first term using the Jacobi-Anger expansion,

$$\exp(iz \sin(\phi)) = \sum_{q=-\infty}^{+\infty} J_q(z) \exp(iq\phi) \quad (2.11)$$

where $J_q(z)$ is the q^{th} Bessel function of the first kind. Upon substitution into equation 2.10 we yield

$$h(x, y) = \sum_{q=-\infty}^{+\infty} J_q(f(a)) \exp(iq\phi). \quad (2.12)$$

From here it is clear that the coefficient corresponding to the first term which has the desired phase is

$$c_1 = J_0(f(a)) \quad (2.13)$$

and thus the encoding condition becomes

$$J_0(f(a)) = Aa. \quad (2.14)$$

We thus have the functional form of $f(a)$ if we can find the inverse of equation 2.14, and we can thus apply our encoding method.

With the above encoding, the Fourier terms in the expansion are all still centred upon the optical axis in the Fourier plane however, and thus the desired Fourier component can not be spatially separated from the other terms. By adding a ‘carrier frequency’ to our desired phase the Fourier terms spatially separate in the Fourier plane and the appropriate Fourier term can be isolated. Thus we add the phase of a plane wave whose k-vector is not aligned with the original optical field,

$$\phi'(x, y) = \phi(x, y) + 2\pi(\mu_0 x + \nu_0 y) \quad (2.15)$$

where μ_0, ν_0 are two chosen wavevectors in the x and y directions respectively. This ‘carrier signal’ then spatially separates all of the Fourier terms, allowing isolation of

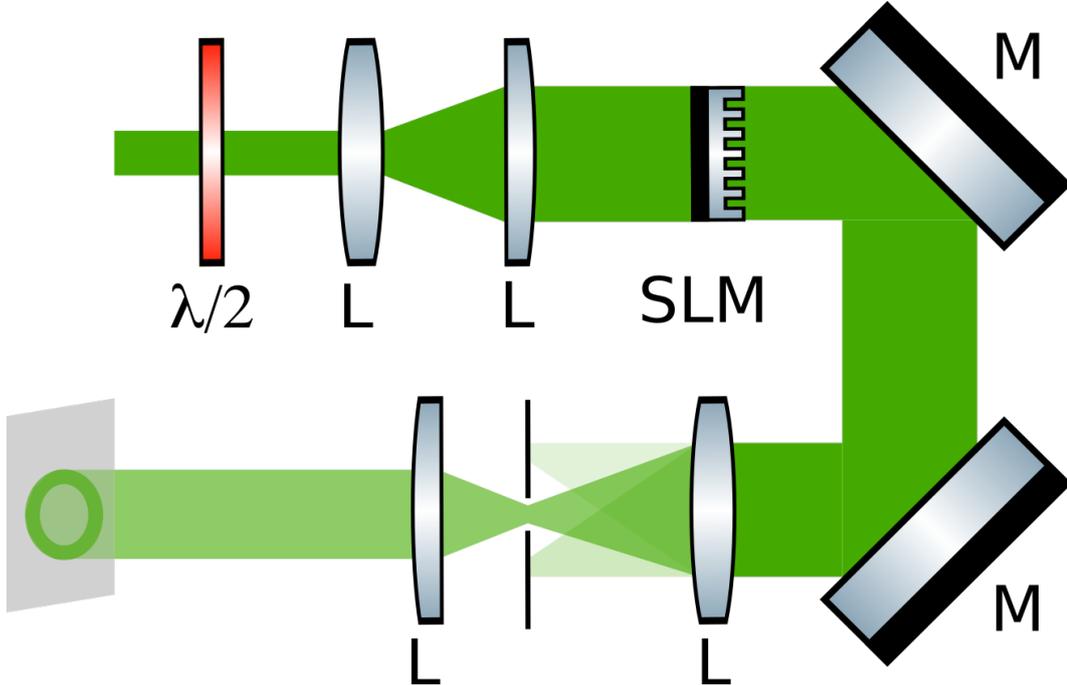


Figure 2.2: Creation of an arbitrary complex field. A schematic showing how a complex optical field can be created using a phase only spatial light modulator and a computer generated hologram. A halfwave plate orients the polarisation of the light for maximum diffraction efficiency, followed by a pair of lenses that expand the beam to overfill the SLM. Another pair of lenses follows the SLM, reimaging the SLM in a 4f configuration and allow spatial filtering of the desired mode.

the desired term.

2.2.2 Implementing the encoding method

To implement the above method we numerically invert equation 2.14 to find the values of $f(a)$ for $a \in [0, 1]$ ⁶, and storing these pairs of values in a LUT for calculation of the CGH. Of particular importance to this work is the fact that numerical inversion of this equation, coupled with a phase range of $[0, 2\pi]$ for the desired complex optical field, requires only a $[0, 1.17\pi]$ phase range for the SLM⁷, just below the limit of 1.19π that our SLM can achieve. Once the desired complex optical field is known and the amplitude and phase stored in an (m,n) matrix of the size of the resolution of the SLM display, the carrier frequency phase is added to the mode and the

⁶This is done by finding the solution to $J_0(f(a)) - a = 0$ numerically

⁷Note that this arises due to our choice of encoding function, different encoding methods shall require different phase ranges.

encoding function, equation 2.8 is then applied to generate the appropriate CGH⁸. The resultant CGH then contains an array of the desired phase for each pixel, and this is then converted to grayscale values using the LUT described in section 2.1, and the result displayed upon the SLM.

So far we have only discussed how an appropriate CGH can be calculated, to generate a desired complex optical field, here I will describe how this is experimentally implemented. A schematic of a typical setup used to generate a desired optical field with this method is shown in figure. 2.2. As can be seen, the output of the laser is first expanded to overfill the SLM, providing a relatively uniform intensity across the SLM. The polarisation of the light is also adjusted here with a half waveplate to ensure maximal diffraction power in the diffracted mode, as only the polarisation that aligns with the axis of the liquid crystals will be phase shifted.

To separate the desired field from the higher order modes generated by the encoding, the optical field needs to be spatially filtered in the Fourier plane. This is achieved by imaging the SLM with a telescope, typically with some level of magnification dependent on the application, and placing an iris at the focal plane of the first lens, where the modes will be smallest and most easily separated, and allowing only the first order mode to pass through the iris unobstructed. This is performed with two lenses in a 4f imaging configuration, with the SLM one focal length in front of the first lens, and an iris placed one focal length behind the first lens. Finally the second lens, placed one focal length away from the iris, collimates the beam and places the desired optical field one focal length behind it.

⁸All this is performed in MATLAB.

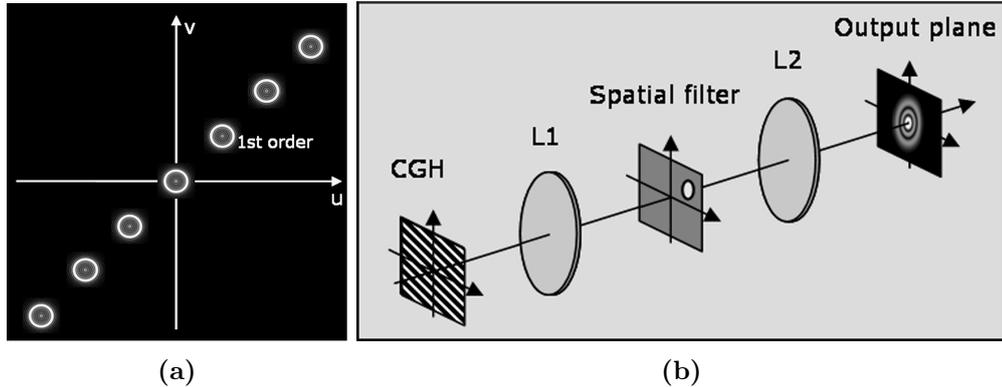


Figure 2.3: Implementing computer generated holograms. a) shows how the addition of a carrier frequency to the computer generated hologram spatially separates the Fourier components of the hologram in the Fourier plane, whilst b) shows how this can be implemented optically. Reproduced from [1].

2.3 Creating an optical vortex knot

Having established a method of generating arbitrary complex optical fields, the method was to be put to the test generating arguably the most intricate of optical fields, an optical vortex knot. An optical vortex knot is an optical field in which, if you image a series of planes transverse to the direction of propagation, optical vortices appear in pairs of opposite handedness from a field that initially contains no topological defects, twist about one another as they propagate, and then recombine to annihilate. If one traces out the trajectory that these vortices follow, they'll find a closed loop tied into a knot. Optical vortex knot's were first realised by Leach *et al.*[65], a complex⁹ feat of work combining the mathematics of knot theory with the paraxial Helmholtz equation to calculate complex paraxial optical fields that upon spatial evolution would realise an optical vortex knot. It is this method that was followed to create a knotted optical field and is briefly summarised below, however, reference [66] or [65] describe the process in detail.

To realise the knotted optical fields of Leach *et al.*[65]¹⁰ a complex-valued scalar polynomial function must first be defined in which the nodal lines realise the desired knotted topology, it should be noted that this complex scalar polynomial function does not obey the Helmholtz equation. Next, this complex scalar polynomial is approximated in the $Z=0$ plane using the Laguerre-Gaussian (LG) modes, of a chosen

⁹No pun intended!

¹⁰An extension to this method that introduces two new parameters allowing the geometry of the knot to be tuned can be found in Ref. [66] or the supplementary information of Ref. [2]

waist, w , which affects the geometry of the knot, as a basis, as these are the cylindrically symmetric solutions of the paraxial Helmholtz equation, which naturally contain the topological defects (vortices) that forms the knot. This superposition of LG modes is truncated at some finite order¹¹, and is then experimentally implemented, realising the desired topological structure.

2.4 Optical vortex knots, one photon at a time

The double slit experiment as described by Feynman, is “a phenomenon which is impossible, absolutely impossible, to explain in any classical way and which has in it the heart of quantum mechanics[67]”. The double slit experiment has been instrumental in demonstrating key features of quantum mechanics, almost simultaneously demonstrating both the wave and particle nature of light and matter; by requiring the propagation of the quantum object of interest through the slits to be described by their wave like properties if an interference pattern is seen, and yet any information that might yield ‘which way information’ about the object as it propagates through the experiment will destroy the resulting pattern, indicating particle behaviour. To date, numerous double slit experiments have been performed with single photons[68][69][70], electrons[71][72][73], atoms[74] and even macroscopic molecules[75], however, each of these experiments has only demonstrated the interference within a given 2D plane and with further propagation the resulting fields have all been self-similar.

Given such a significant statement by such a notable physicist, we set out to realise the first measurements of a three dimensional complex optical field on average¹², one photon at a time, demonstrating that wave-particle duality not only implies that interference will occur between spatially distinct modes, but also through the complex evolution of the superposition of these modes. By measuring the complex three dimensional vortex structure of the resulting optical field we demonstrate that the effects of diffraction upon what can naively be thought of as a single particle of

¹¹With larger numbers of modes, the complex field in the $Z=0$ plane is more accurately realised, however, the interference of the varying Gouy phases of the LG modes can lead to the destruction of the knotted topology

¹²We can make a statement only about the statistical number of photons, as a highly attenuated laser beam is described by a weak coherent state where the distribution of photons is Poissonian and therefore there is a non-zero, but small probability that the beam contains more than single photons. Ideal single photon behaviour, however, can be mimicked using superposed weak coherent states[76]

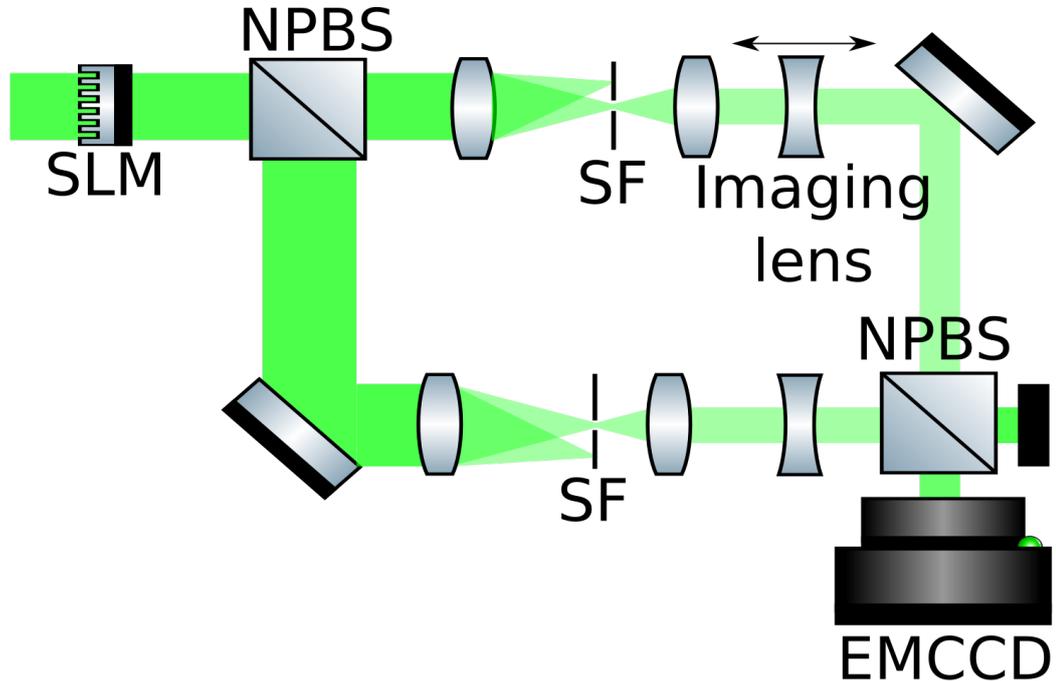


Figure 2.4: Optical set-up for producing optical vortex knots, one photon at a time. Each photon passes through the spatial light modulator (SLM), which acts as a static phase element generating both the knot and a Gaussian reference beam, which is split upon a non-polarising beam splitter cube (NPBS) and then spatially filtered with a pinhole (SF) such that one path contains the knot, and the other the reference beam. These two paths are then recombined at a small angle, which generates an interferogram that facilitates identification of the vortex structure, and then imaged onto the camera (EMCCD). The double headed arrow indicates the **Imaging lens** that is translated to image different transverse planes of the knot. Reproduced from [2]

light, are the same as for the entire optical field passing through the apparatus, even though the effect of diffraction is essentially ‘non-local’, and thus the light particle must have ‘knowledge’ of the entire phase hologram.

2.4.1 Measuring an optical vortex knot

To demonstrate that photons, approximately one at a time, can realise an optical vortex knot we used the experimental setup shown on the left of figure. 2.4. First, light from a continuous wave, 532 nm diode laser is spatially filtered and expanded to overfill the SLM, upon which a complex-valued phase-only hologram, generated as described in section. 2.2, is displayed that generates both a complex optical field

containing a Trefoil knot, and a Gaussian beam, that shall be used to interferometrically measure the resulting complex structure of the knotted field. The resultant field is then split in two via a non-polarising beam splitter cube (NPBS) to create an interferometer with two arms, each arm containing a telescope, which magnifies and spatially filters the optical field, selecting either the knotted optical field or the Gaussian reference beam. In the arm that selects the knotted optical field, an imaging lens is placed after the telescope allowing different planes of the knotted optical field to be imaged onto an electron multiplying CCD (EMCCD) array. Finally, the two fields are recombined on another NPBS and directed onto the EMCCD to measure the resultant field, with the reference beam directed at a slight angle with respect to the knotted optical field creating slanted interference fringes used to identify the locations of the optical vortices.

To ensure that our apparatus only contains on average, a single photon at a time within the device, the output power of the laser is reduced to <10 pW which corresponds to <0.1 photons at a time along the path from the SLM to the camera, ≈ 1 m. If we assume a Poissonian distribution of the emission of photons from the laser [77][78], then the chance of two photons being within the apparatus at any given time is ≈ 20 times less than that of a single photon being within the device, however, due to optical losses from a number of elements along the path, the chance of any given photon making it through the apparatus from the SLM to the camera is 5%, and as such the chance of two photons making it through the device to the camera is much further reduced. The laser beam can be described by a coherent superposition of different photon number states, here I have attenuated the beam such that the average of this superposition is such that <0.1 photon is in the apparatus at any given time. To be able to detect single photons, we use an Andor Ixon EMCCD in high gain, photon counting mode, where each pixel can only return a binary response, either a photon has been detected or not. Figure 2.5a) shows a single image taken by the EMCCD. The exposure of a single image is 0.3 seconds, and this on average records ≈ 600 photons, much less than the total number of pixels of the camera¹³.

From a single image, however, it is not possible to discern the shape of the resulting field, let alone the locations of the vortices that are in the darkest regions of the optical field, thus thousands of images must be taken, and the images summed together to create high enough contrast to pick out the forked structures in the interferogram. Figure 2.5b)-d) show the result of 100, 1000 and 8000 images summed

¹³The camera has 512 by 512 pixels, thus if a photon was equally likely to fall into any given pixel, the chance of two falling into the same pixel is $\approx 5\%$ per image

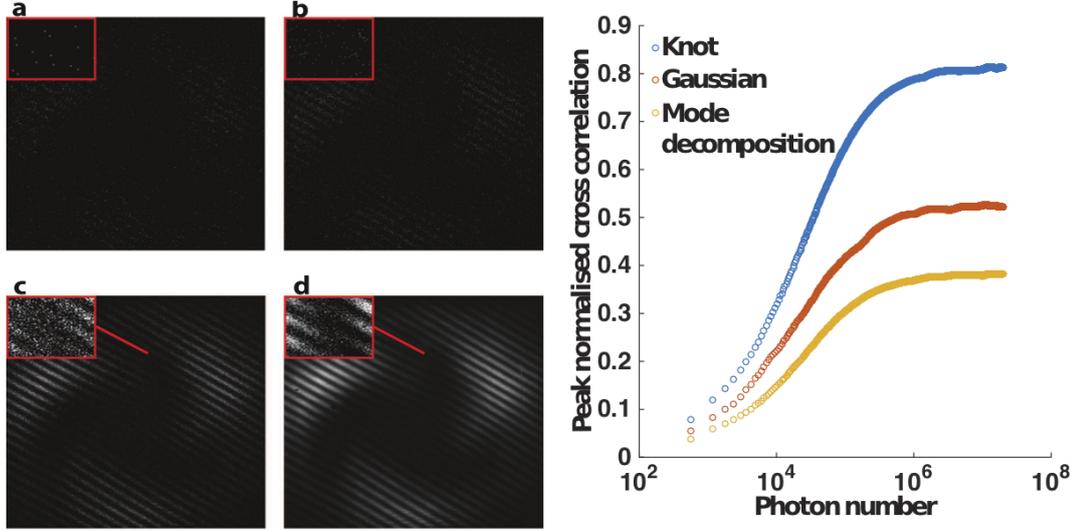


Figure 2.5: Left: Measurement of the knotted topology a, A single frame imaged by the camera in photon counting mode. **b-d**, Measurements resulting from the summation of 100, 1000 and 8000 images, respectively, with the contrast of the resultant interference pattern increasing with increasing numbers of images making the forked vortex structure more visible (shown magnified in insets). **Right: Normalised cross correlation as a function of photon number**. The peak value of the NCC between the interferometric measurement of the knotted optical field at $Z=0$ with the expected theoretical intensity pattern of the knot, a Gaussian mode and the modal decomposition of the knot used in the 'which path' measurement, each interfered with a Gaussian mode to reproduce the global fringe pattern as a function of increasing number of summed images, displayed as increasing photon number. Reproduced from [2]

together, respectively, while the inset shows the location of one of the vortices. From these figures, it is clear that with increasing number of images the contrast of the resultant image improves, and eventually the location of the vortices can be determined. To be able to take thousands of images without the interference fringes shifting, the difference in the optical path length must remain constant, as variation by more than $\frac{\lambda}{2}$ will shift the phase of the fringes by half a wavelength and wash out the resultant interferogram. We found our optical setup to be stable on time scales greater than one hour, allowing us to collect ≈ 8000 images for each position, creating images with high enough contrast for accurate location of the vortices.

To quantitatively evaluate how close the resultant optical field was to the expected theoretical field, a normalised cross correlation (NCC) between the resultant image¹⁴, of different numbers of exposures, and the expected theoretical field, as well as a Gaussian beam of the same waist as used in the knotted optical field and a non-

¹⁴measured in the $Z=0$ plane

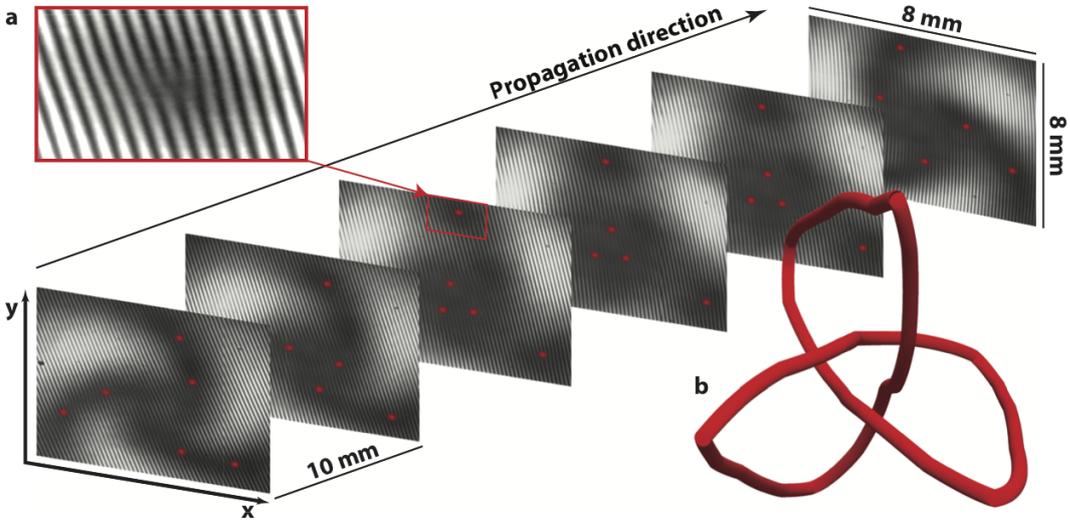


Figure 2.6: Measuring the optical fields knotted topology. Different transverse planes of the optical field along the beam propagation direction (indicated by the black arrow) are imaged onto the camera by moving the position of the final imaging lens before the camera (see optical set up in left of Fig. 2.4). The length scales in both the propagation direction and transverse plane are shown. **a**, Vortices within each plane (indicated by red dots) are located by the characteristic forked structure they create in each interferogram, due to the 2π phase winding about each vortex. **b**, The vortices located in 25 separate planes are then ‘stitched’ together to form the three dimensional trefoil knot nodal structure. Reproduced from [2].

coherent modal decomposition of the knotted optical field¹⁵ was performed. The result of these normalised cross correlations can be seen on the right hand side of figure 2.5, as a function of the total number of photons that had been detected by the camera. From this data, it is clear that even after a single image containing only ≈ 600 photons, corresponding to only 0.002 photons per pixel, the correlation of the resultant image with theory is greater than any of the other choices, with respective values of 0.08, 0.06 and 0.04. Each of the resultant curves show a similar shape, initially rapidly growing approximately linearly followed by a flattening once the number of photons approaches 10^6 , corresponding to ≈ 3 photons per pixel on average.

Once we could measure the resultant optical field in a given plane, the imaging lens in figure 2.4 was then translated to image a different transverse plane of the knot onto the camera, while keeping the Gaussian reference beam the same. This was

¹⁵This is the resultant field if the amplitude of the Laguerre-Gaussian basis used to form the knot were added together, but their phases were not.

then repeated to generate a series of 30 images¹⁶, within which the entire knotted structure was encapsulated. For each of these positions, a series of 8000 images was captured to ensure high enough contrast in each plane to locate the vortices by their characteristic forked structure in the interferogram, and the locations of the vortices were recorded. The resultant locations of the vortices in three dimensions was then known¹⁷ and stitched together, with a linear interpolation used in between locations, resulting in the knotted structure shown in the bottom right hand corner of figure 2.6. The rest of figure 2.6 shows 6 typical images of different transverse planes while the inset shows the characteristic vortex ‘fork’[79] used to locate the positions of the vortices, and the red dots indicate these positions.

This data shows that even when the measurements are performed such that on average only one photon is within the apparatus at any given time, the knotted topology of the optical field remains, showing that the ‘non-local’ effect of diffraction is preserved both in space (within a single transverse plane) and time (throughout the temporal evolution of the knotted field).

2.4.2 ‘Which path’ measurements

Having verified that an optical vortex knot could be realised approximately one photon at a time, we had our main result. However, one key feature of the double slit experiment is that if one of the two paths are blocked throughout the experiment, or alternatively a measurement made of which path the photon traversed throughout the experiment, the interference is lost. Performing this type of experiment ensures that we can produce a null result, that we can ensure that the photon being able to traverse both paths at the same time is a necessary condition for generating an optical vortex knot.

To perform this ‘which path’ measurement, the superposition of 5 LG modes that is used to form the knotted optical field was generated via two complex filters displayed on the SLM side by side, with the first four LG modes¹⁸ generated by the first complex filter, and the final LG mode generated by the second complex filter¹⁹.

¹⁶Only 25 planes of which contain vortices

¹⁷the camera provided the (x,y) coordinates, while the distance the lens was translated and the position of the corresponding plane that was imaged provided the z coordinate

¹⁸The LG modes displayed on one complex filter were of orders (p,m)=(0,0), (1,0), (2,0) and (3,0), and on the other was order (0,3)

¹⁹The chosen separation of modes ensures that the non-coherent summation of the modes retains its rotational symmetry upon propagation, while the knotted optical field does not

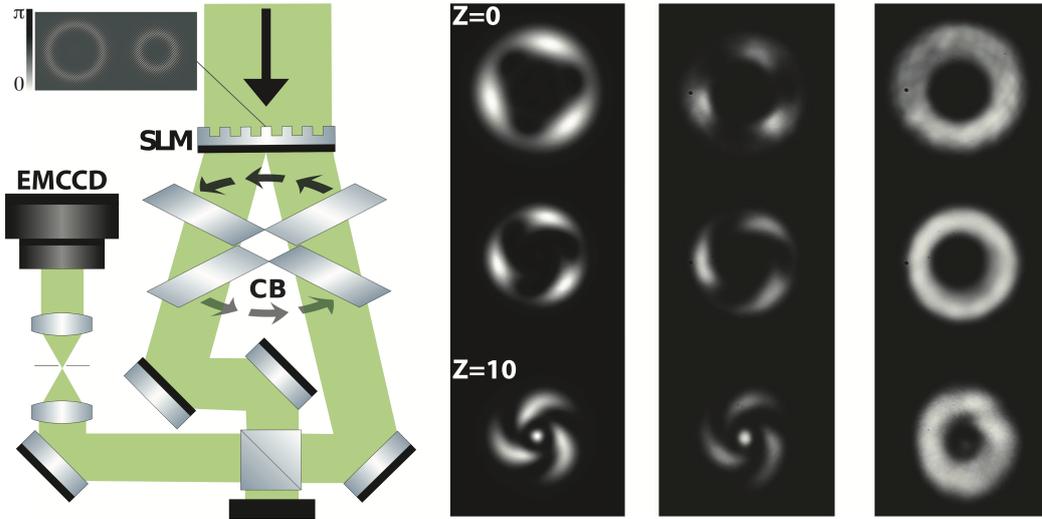


Figure 2.7: Left: Optical set-up for a ‘which path’ measurement. The SLM displays two holograms side by side. The laser beam is split into two and passes through a chopper blade (CB), which ensures that photon’s are present along the optical path from only one of the two holograms at any given time. The two beam paths are later recombined and imaged onto the camera. **Right: ‘Which path’ measurements of the optical field.** The three columns shown contain the theoretical intensity patterns of the knotted optical field (left), the intensity patterns resulting from the unchopped beams (centre) and from the chopped beams (right). These images are taken at three different transverse planes along the beam path (positions are indicated in centimetres). A normalised cross correlation of the data (centre and right) with simulation (left) shows a higher correlation of the knot with the unchopped optical field, with an average peak value of 0.89, as opposed to when the beam is chopped, which yields an average peak value of 0.62. Reproduced from [2]

These complex filters diffracted the resultant modes in different directions, allowing a chopper blade to be inserted that would guarantee that light could only pass through one of the two filters at any given time. The optical setup shown on the left hand side of figure 2.7 shows an inset with the two complex filters side by side, and the resultant diffraction of the two desired modes and the chopper blade (CB) inserted as well. The two modes are then recombined upon a NPBS, and the resultant field imaged onto the camera.

To generate the desired knotted optical structure, light must pass through both slits at the same time so that the resultant modes can coherently interfere with one another, otherwise the resultant field will simply be the summation of the intensity profiles of the two complex filters. To perform these experiments the chopper blade was first removed, and the resultant optical field was measured in a number of planes as described in section 2.4.1, the result of these measurements can be seen in the central column of the right hand side of figure 2.7, which shows three different

transverse planes of the knotted optical field, alongside theory on the left hand side. The chopper blade was then subsequently inserted, and the measurements repeated with the chopper blade alternately blocking the two paths at a rate of 1kHz. This resulted in the intensity profiles shown on the far right column of figure 2.7, which can be seen to approximately retain its rotational symmetry.

The three images were taken at the centre of the knot (0 cm), and after 5 cm and 10 cm of propagation respectively. We then also performed NCC of the resultant profiles with theory, when the photons were able to pass along both paths uninterrupted the NCC resulted in values of 0.91, 0.89 and 0.89 for the three images respectively, while when one of the two paths was alternately blocked the resultant NCC was 0.67, 0.76 and 0.46, all significantly less than that of the unchopped field. We then also looked at the average NCC between the three images in a given data set, i.e. the images within a column of figure 2.7, for the unchopped path the average NCC was 0.54 while for the chopped path the average NCC was 0.75, indicating that the chopped path was more self-similar upon propagation, and thus contained less three-dimensional structure.

2.4.3 Summary

In this chapter we have demonstrated the ability to generate arbitrary complex paraxial optical fields using a phase only spatial light modulator and a complex valued phase only filter.

To demonstrate this ability we generated arguably one of the most intricate paraxial optical fields, an optical vortex knot, and realised this optical field approximately one photon at a time. We then performed an experiment akin to a ‘which path’ measurement of the double slit experiment, in which the knot was decomposed into two separate groups of optical modes, the coherent interference of which resulted in the desired knotted topology. By inserting a chopper blade to prevent the photons from being able to traverse both paths simultaneously, the coherent interference, and the resultant knotted topology was lost. These measurements demonstrate a fundamental physical principle, that particle-wave duality implies both interference in space, between spatially distinct modes, as well as time, through the complex evolution of these modes. The techniques for the complex manipulation of an optical field presented in this chapter allows us to be able to generate a range of optical fields that can be used for trapping and manipulating atoms, however, it has its

drawbacks, mostly associated with the use of the SLM.

In the next chapter we shall describe how these issues can be circumvented by using the SLM to record our own holographic optical elements (HOEs) which can recreate our desired optical fields, but without many of the drawbacks associated with an SLM.

Holographic optical elements for ultracold atom trapping

Some of the work presented in this chapter has previously been published in:

- *S. J. Tempone-Wiltshire, S. P. Johnstone, K. Helmerson, "High efficiency, low cost holographic optical elements for cold atom trapping", *Optics Express*, Vol. 25 ,1, 99. 296-304 (2017).*

Although we have demonstrated the ability to manipulate the complex wavefunction of an optical field, the path we have chosen is not the most experimentally feasible due to problems that arise due to the spatial light modulator. In this chapter we shall introduce a method of recording holographic optical elements (HOEs) from the optical fields that have been generated by the SLM, which when played back with a Gaussian beam can recreate the desired fields with high efficiency and fidelity. This method relies upon a relatively new photopolymer, known as Bayfol HX, which requires no chemical processing and due to a large refractive index modulation coupled with a relatively thick medium, can produce volume holograms capable of reaching close to 100% diffraction efficiency. In this chapter we shall also introduce a method of recording HOEs that can record phase structures that require an SLM with higher resolution than our SLM has. It is these HOEs that will be used on our experiment to confine and manipulate our atoms.

3.1 Problems with SLMs

As mentioned, SLMs have greatly enhanced our ability to manipulate the complex wavefront of light, however, this is not without drawbacks arising from the physical structure of the device, as well as how they are electronically addressed. The problems most relevant to the application of cold atom trapping are; low total diffraction efficiency[80], time varying profiles and low optical damage thresholds.

The low total diffraction efficiency arises from typically two properties of the SLM; the first is the large number of optical modes that are generated when a phase or amplitude only complex filter is used, whilst the second arises from the pixelated nature of the device. Two factors arise from the pixelated nature of the SLM, which act to reduce the total diffraction efficiency of the device; the first is due to the fill factor of the screen, which is never 100%. The fill factor measures the percentage of the screen through which light can pass, with the rest of the space being dead space between pixels that does not interact with the light and thus reduces the diffraction efficiency of the device.

The second factor arising from the pixellation occurs due to the regular structure of the screen, with the dead space between the pixels acting as a periodic diffraction grating, generating a number of diffraction orders from the screen. This effect is clearly shown in figure 3.1, where a Gaussian beam is shown both with and without being transmitted through the SLM, when the beam passes through the SLM the diffraction into the higher order modes is clear. For our SLM, the transmission is $\approx 50\%$ which, coupled with a damage threshold of 1 W cm^{-2} and a display size of $\approx 1 \text{ cm}^2$ means that little laser power will be available for trapping, $< 500 \text{ mW}$, without even considering the efficiency of the phase holograms used to generate the desired modes.

The other major drawback of a spatial light modulator is the time varying profile it generates. This time varying profile is caused by a phenomenon known as ‘phase flicker’, which arises from the pulsed modulation scheme used to address the liquid crystals[81], and occurs at 100Hz, which is close to the typical trapping frequencies for BECs which can thus lead to resonant parametric heating of the atoms. This phenomenon can be reduced by actively cooling the liquid crystal display, however, it has not been fully removed[82].

Finally, as will be discussed in section 8.1 we intend on generating optical fields with large orbital angular momentum (OAM) for confining our atoms as this creates

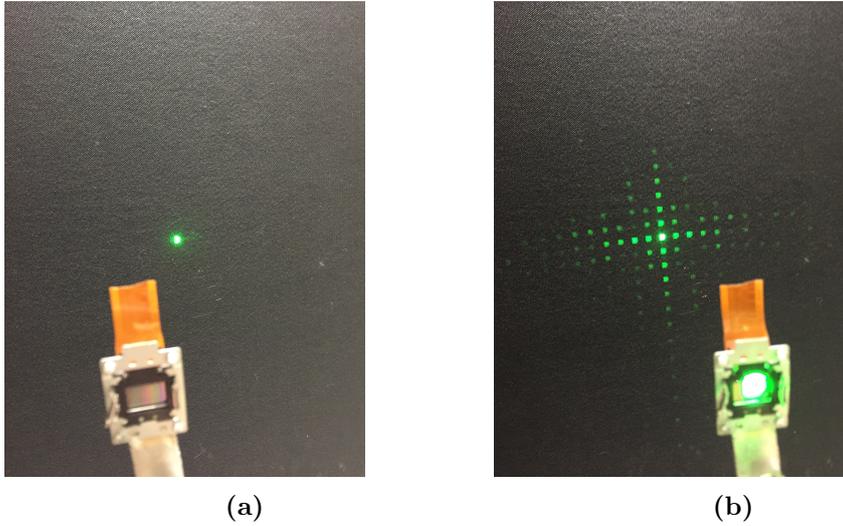


Figure 3.1: Pixellation of an SLM. **a)** and **b)** show the effect of the pixellation of the SLM upon the transmitted light, with **a)** showing a Gaussian mode that is not passing through the SLM whilst **b)** shows how transmission through the SLM generates a large number of diffracted modes, significantly reducing the power in the undiffracted mode.

more uniform potentials. However, due to the techniques employed to generate these modes with an SLM, there is a limit to the maximum OAM that can be imparted to the optical field, limited by the resolution of the SLM. This is because our phase only complex filter requires the different diffraction orders to be separable in the Fourier plane to be able to separate out the desired mode. However, as more orbital angular momentum is imparted to the optical field, for a given Gaussian waist illuminating the display, the size of the mode in the Fourier plane tends to increase until the modes start to overlap, limiting the amount of angular momentum that can be imparted to the optical field.

It may not be immediately clear why it is the resolution of the SLM that is the key factor here. Consider an SLM with a given pixel size, if the resolution of the display is increased (more pixels) then the overall size of the display is increased and a Gaussian mode with a larger waist can be used to illuminate the display creating a smaller beam in the Fourier plane, allowing greater separation of the modes. One could think that this could also be achieved by increasing the size of the pixels of the SLM at a fixed resolution, however, overall this does not help, as although the display size increases and a smaller beam at the focus can be created, the pixel pitch is also increased. This increase in pixel pitch reduces the angular separation achievable by the diffraction grating used to separate the modes, which brings them

closer together at the focus, and these two effects act to cancel one another¹.

We thus seek a solution that can alleviate these issues associated with using an SLM, without sacrificing the high fidelity of the optical fields it generates. Our solution to this problem was to continue using the SLM to generate these fields, but then using these fields holographically record optical elements that could then recreate these high fidelity fields with high efficiency. To be able to achieve this though, a recording medium is required that is ideally simple to work with and can accurately encode our desired fields.

¹The spot size at the focus is proportional to the area of the LCD, whilst the spatial frequency of the pixellation, which is proportional to the separation in the Fourier plane, is inversely proportional to the pixel pitch.

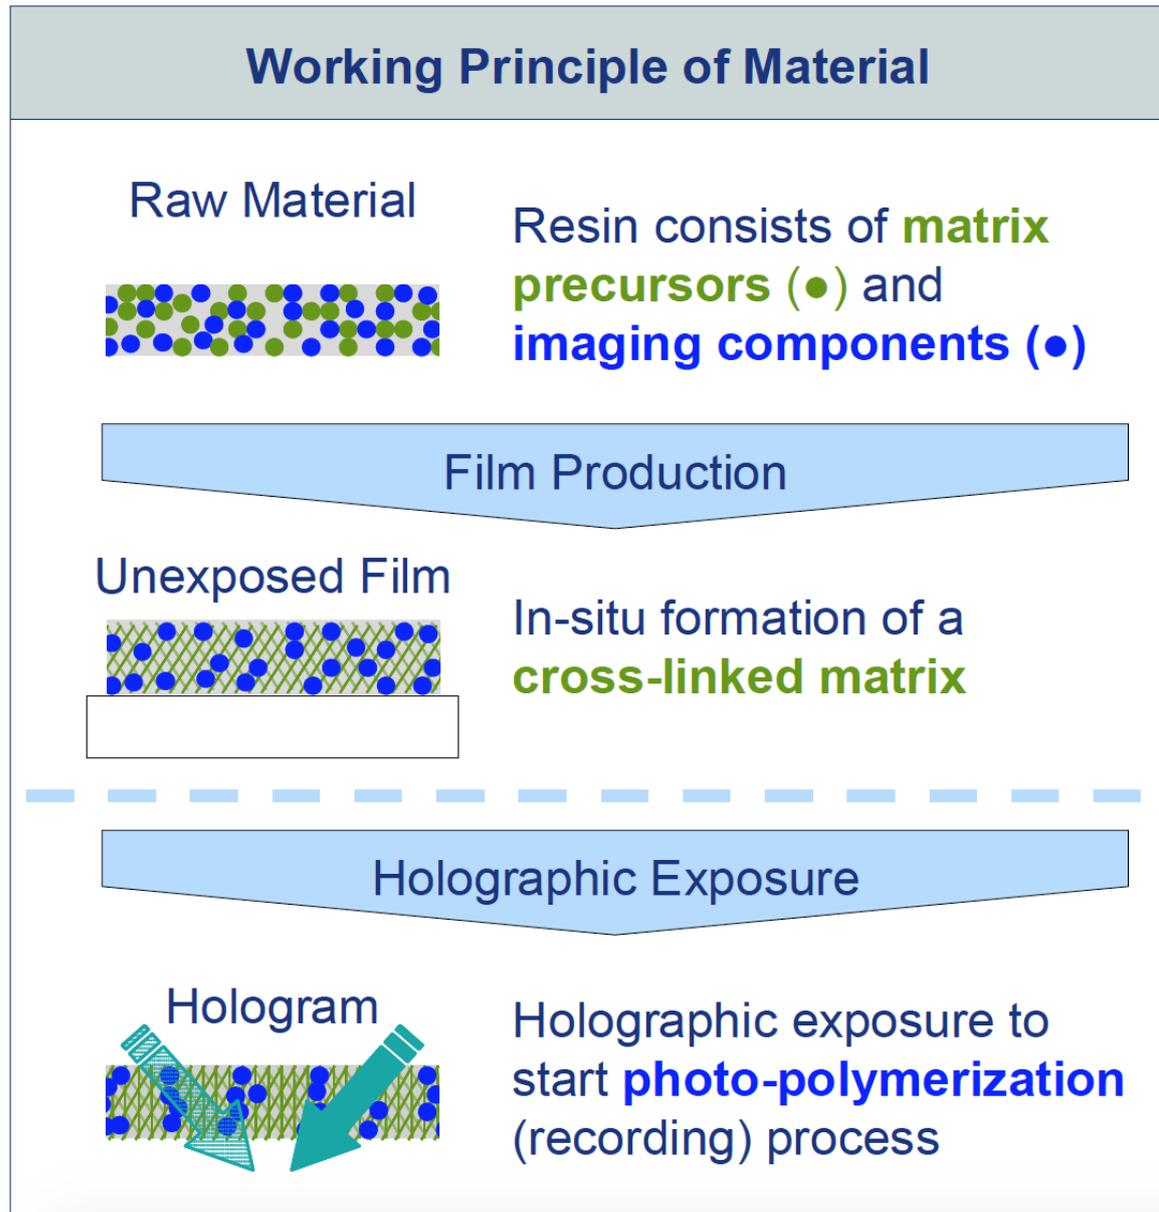


Figure 3.2: The working principle of Bayfol HX. Bayfol HX consists of two main components, imaging monomers, and matrix precursors in a resin. During the films production at Bayer, the matrix precursors are formed into a cross-linked matrix. The film is then exposed to laser irradiance to record the hologram, initiating photopolymerisation, and the imaging monomers then diffuse away from the dark regions and this, coupled with the polymerisation, creates a refractive index modulation storing the hologram.

3.2 Bayfol Hx

The problem of most holographic recording media is that they cannot achieve large refractive index modulations, Δn , which limits the maximum diffraction efficiency they can obtain. Besides this, from a commercial point of view, most holographic recording media require a cumbersome stage of wet chemical or thermal processing after laser exposure, making the process of recording a hologram more complex and susceptible to failure. In 2009 Bayer Material Science introduced a new photopolymer[42] it was developing, which required no chemical or thermal post-processing, could achieve high diffraction efficiency and was durable.

Figure 3.2 shows the working principles of Bayfol HX, which relies upon two different polymer chemistries, one which forms a cross-linked matrix during production, making the film stable, allowing it to be stored and handled in light safe conditions before exposure, and the other which is responsible for the hologram formation. Once the film has been produced, recording of a hologram via laser irradiance initiates photopolymerisation by activating the imaging monomers. Unactivated monomers then tend to diffuse through the material until they react with another activated monomer or polymer chain near a bright interference fringe, which generates a net diffusion from the dark regions to the light regions. It is this net diffusion, coupled with the photopolymerisation that generates the refractive index modulation within the material, and can lead to local refractive index changes of $\Delta n > 0.03$. Once all imaging monomers have reacted, the writing chemistry has been consumed and all that is required is to photo-bleach the material with UV to gain full optical transparency.

3.2.1 Recording with Bayfol HX

Recording with Bayfol HX seems relatively straightforward, expose the material and you're done! However, life is never so simple and there are some considerations that need to be kept in mind to create high diffraction efficiency holograms. The first consideration is simply of recording geometry, whether the hologram will be a transmission or reflection hologram, as this will dictate other recording parameters. It has been demonstrated that Bayfol HX can achieve high diffraction efficiency in both transmission and reflection geometries[83], however, as our holograms are to be used with our experiments, transmission holograms are typically the simplest to

work with since the diffracted mode won't retrace the path of the incident beam making it simpler to isolate.

With the recording geometry chosen we then must determine what recording parameters will give the highest diffraction efficiency, to this end we consider some basic volume hologram theory. When considering a simple non-slanted transmission volume Bragg grating, the diffraction efficiency, ignoring Fresnel reflections off the surfaces and absorption by the media, is given by[3]

$$\eta = \sin^2\left(\frac{\pi}{\lambda \cos(\theta)} \cdot d \cdot \Delta n\right) \quad (3.1)$$

where η is the diffraction efficiency, ranging from 0 (no diffraction) to 1 (100% diffraction), λ is the recording wavelength, θ is the angle of incidence of the recording beam with respect to the normal of the surface of the photopolymer (and hence also the Bragg angle), Δn is the refractive index modulation of the photopolymer, and d is the thickness of the photopolymer.

The argument of \sin^2 in equation 3.1 is sometimes referred to as the grating strength, ϕ , and it is clear that as ϕ ranges from 0 through to $\frac{\pi}{2}$, the diffraction efficiency of the grating will increase from 0 through to 100%, this is known as the weak grating strength regime. As the grating strength is further increased above $\frac{\pi}{2}$ the diffraction efficiency then begins to fall off again approaching zero as $\phi \rightarrow \pi$ this regime is known as the strong grating strength regime. For maximal diffraction efficiency it is therefore clear that we wish to make the grating strength as close to $\frac{\pi}{2}$ as possible. Figure 3.3 shows how this modulation occurs during the recording process of a transmission hologram. There are three distinct regions shown, the first region is a period of inhibition in which there is no change in the refractive index modulation as photo polymerisation is only beginning. After the period of inhibition the refractive index modulation begins to grow, as does the diffraction efficiency, whilst the hologram is still within the weak grating regime. Once the grating strength has reached $\frac{\pi}{2}$ the hologram starts to enter the strong grating regime and the diffraction efficiency begins to fall with further refractive index modulation. As the thickness of the photopolymer film is fixed in production, and the wavelength, with which we record our holograms, is fixed by the laser we will use for trapping our atoms, there are only two parameters left that can be varied, the refractive index modulation and the Bragg angle.

The refractive index modulation of the recorded photopolymer depends upon a great

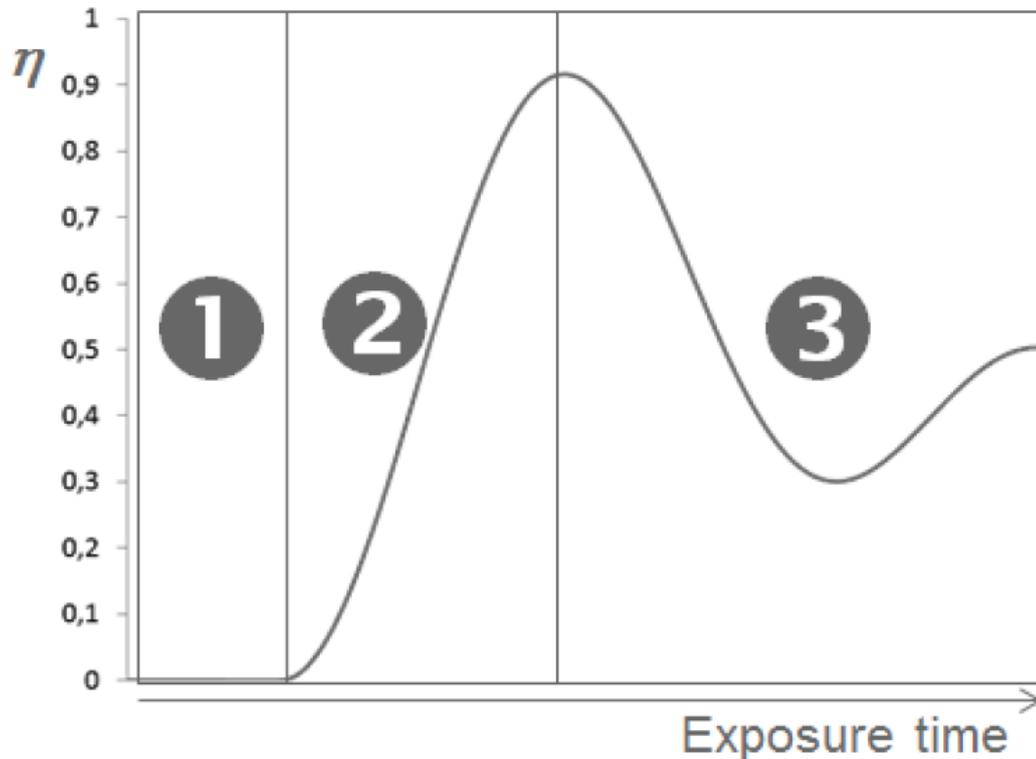


Figure 3.3: Overmodulation of a transmission hologram. The dependence of the diffraction efficiency, η , upon the grating strength, which increases with exposure time is shown. The graph has three defined regions during exposure. Region 1 corresponds to an initial inhibition in which recording of the hologram has not yet begun as photopolymerisation is still occurring. Region 2 shows the growth of the refractive index modulation whilst the grating strength is still in the weak grating regime, and thus the diffraction efficiency increases with increasing modulation. Region 3 shows the strong grating regime, where the diffraction efficiency begins to fall with further modulation. Reproduced from [3].

number of parameters, such as the spatial frequency of the recorded grating², the intensity of light used for the recording, the total exposure dosage[83] as well as the fringe contrast of the resulting interference[3]. However, in reference [83] it can be seen that the refractive index modulation increases rapidly with increasing exposure dosage, until it saturates to a fixed value. We chose to make use of this saturation in order to make the recording of our holographic optical elements more repeatable, by making the diffraction efficiency effectively independent of exposure dosage by reaching this saturation regime. In doing this, however, the refractive index modulation that the photopolymer attains is fixed, and thus there is only one

²The spatial frequency of the grating is determined by the relative angle between the two beams used to record the grating.

parameter left to vary in order to attain the maximal diffraction efficiency, the Bragg angle³. A quick calculation of the required Bragg angle using the values presented in reference [83] for the refractive index modulation, $\Delta n \approx 0.015$, yields;

$$\theta = \arccos\left(\frac{2 \cdot 0.015 \cdot 16 \cdot 10^{-6}}{532 \cdot 10^{-9}}\right) \approx 26^\circ. \quad (3.2)$$

Experimentally we find that $22^\circ < \theta < 28^\circ$ yields diffraction efficiencies greater than 80%

3.3 Recording a holographic optical element

Once the parameters for recording a transmission hologram have been chosen, it is time to experimentally implement the recording process, the experimental setup employed is shown in figure 3.4.

The output of a continuous wave 532 nm laser is split on a NPBS to generate two arms of the setup, the reference arm and the object arm. The light traversing the object arm passes through a polariser and a halfwave plate, allowing the control of the amplitude of the incident light, as well as ensuring the polarisation of the light incident on the SLM is correct for phase only modulation. This light then passes through the SLM displaying the appropriate phase only complex filter, as described in section 2.2, to generate the desired optical field and a telescope then images this field onto the photopolymer at an angle of 25° with a magnification that generates an appropriate sized hologram for utilising in the experiment. The telescope also acts as a spatial filter, with an iris placed in the focal plane of the first lens used to separate out the desired optical field. Meanwhile along the reference arm, the incident beam passes through a telescope that also acts as a spatial filter to produce a clean filtered Gaussian beam, and magnifies the beam again to the appropriate size for use with the experiment. This beam is also then directed onto the photopolymer at an equal angle of incidence as the object beam.

There are a few experimental properties of the recording process that ensure recording and playback of the hologram are as simple as possible. Firstly, the path lengths of each arm of the setup, the incident polarisation and the local intensity of the

³I should stress that as the saying goes "There are many ways to marinate your tofu" and I have chosen but one method to attain maximal diffraction efficiency, however, the reduction in parameter space from four orthogonal parameters to one makes the achievement of high diffraction efficiency relatively simple experimentally.

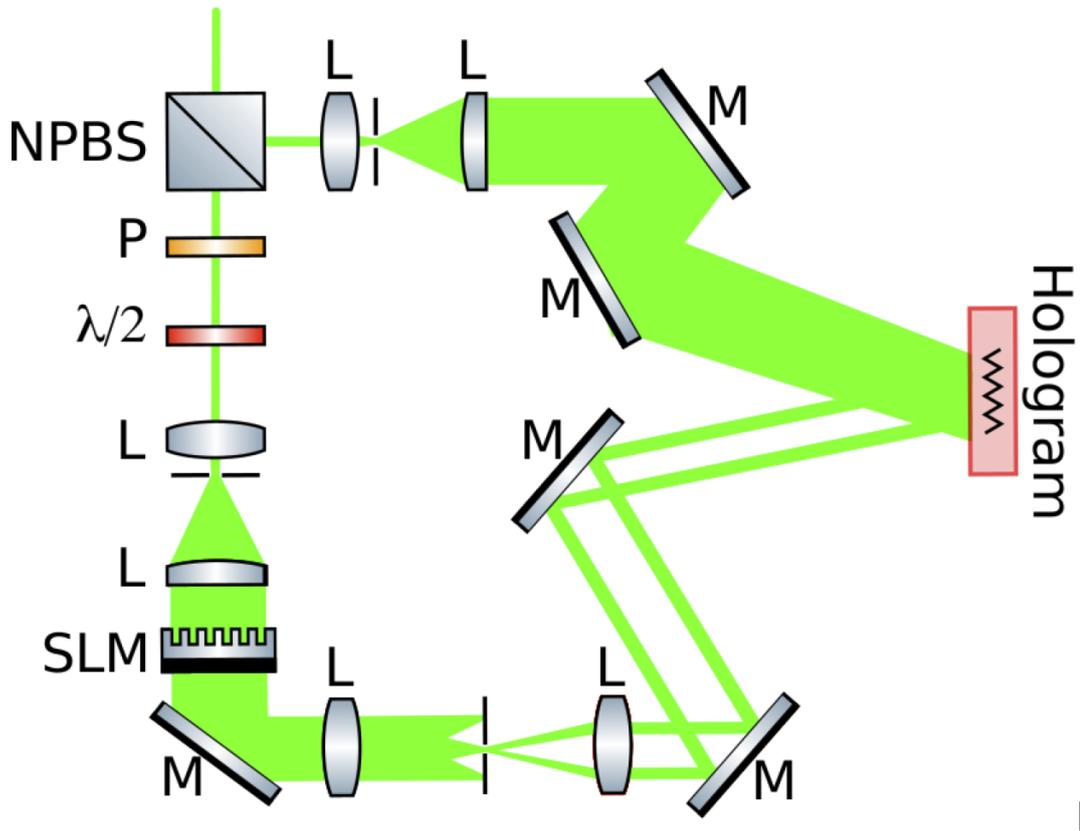


Figure 3.4: Hologram recording setup. A schematic of the experimental set up used for recording transmission holograms. M, mirror; NPBS, non-polarising beam splitter cube; P, polariser; $\lambda/2$, half-wave plate; L, lens; SF, spatial filter; SLM, spatial light modulator. For recording reflection holograms the stage on which the hologram is mounted is rotated such that the object and reference beams enter the photopolymer on opposite surfaces. Reproduced from [4].

two beams at the photopolymer film should be matched as closely as possible, as differences in these three quantities will lead to reduced fringe contrast, resulting in different optimal angles for recording. With regards to the intensity of the object and reference beams, spatially matching the intensities is not always possible due to the spatially varying profiles of the desired optical fields. In this case we typically choose the intensity of the reference beam such that it is everywhere greater than the object beam.

To ensure the object and reference beam are incident upon the hologram at equal angles, and thus ensure a non-slanted hologram is recorded removing any asymmetry between the two Bragg angles, the reflection of one of the beams off the front surface of the hologram film can be traced back along the incident path of the other beam. Both beams should also be made as collimated as possible at the hologram plane, ideally using a shear plate interferometer, as any divergence of the beams will create an additional lensing effect in the resultant hologram.

During the exposure, the recording of the hologram can be visualised by placing a camera in one of the two beam paths, as once the hologram starts to record the object and reference beams will begin to diffract into one another, and phase variations between the two paths then manifest themselves as intensity variations upon the camera. Once the intensity variations on the camera have reached their largest amplitude, the hologram is then at the maximal diffraction efficiency it will achieve⁴. Typical recording intensities for a transmission hologram are $\approx 10 \text{ mW cm}^{-2}$ and exposures are typically on the time scale of 1-2 seconds, with the exact time varied to maximise diffraction efficiency. Once the photopolymer has been exposed, the film can be immediately played back and the efficiency and resulting fidelity measured, however, some optical bleaching may be required to gain maximal transmission through the film.

3.3.1 Quantifying the success of our method

To measure the effectiveness of our method for recording HOEs and their subsequent ability to recreate the desired field, we chose to create a HOE that would convert an incident Gaussian mode into a Hermite-Gaussian mode[84], $HG_{(0,1)}$, to axially confine a BEC.

⁴Visualising the recording process is a good method to ensure that some form of recording is occurring, and can also be used immediately after to check the fidelity of the resultant optical field by blocking the object beam.

Measuring the effective diffraction efficiency of a single holographic element that converts a Gaussian beam into an optical mode that is spatially different is not as simple as measuring the ratio of input to output power. This is because using a single phase-only element to generate a complex optical field limits the maximal obtainable efficiency to the spatial overlap of the amplitude profiles of the two modes. The overall diffraction efficiency can thus be expressed as

$$\eta_{total} = \frac{P_D}{P_T + P_D} \leq \eta_{S.O.} \eta_H, \quad (3.3)$$

where η_{total} is the overall efficiency of the HOE in generating the desired mode, P_D and P_T are the optical power in the diffracted and transmitted modes respectively, and $\eta_{S.O.}$ and η_H are the normalised spatial overlap of the amplitude of the two modes and the diffraction efficiency of the hologram respectively. Here we have chosen to ignore losses due to reflections off the glass substrate the photopolymer is mounted to, as these can be effectively removed via an appropriate anti-reflection coating, and have also ignored absorption losses in the material, which are wavelength dependent.

In the case of two identical beams, the overlap is obviously 100%, and $\eta_{total} = \eta_H$, in this case we have measured diffraction efficiencies in excess of 90%. For a Gaussian mode to a $HG_{(0,1)}$ mode where the waist of the Gaussian is chosen so as to maximise the spatial overlap, the spatial overlap and thus the maximal obtainable diffraction efficiency, is $\approx 54\%$. For the HOE that we have recorded producing the optical field in figure 3.5a, η_{total} was measured to be $\approx 50\%$. In figure 3.5b, the profile of the optical field is shown when focussed with a cylindrical lens, along with a fit along the vertical direction for a $HG_{(0,1)}$, the fit shows close resemblance to the data, indicating high fidelity of the reconstructed mode and that our holograms achieve close to the largest possible efficiency for a single, phase-only hologram. Figure 3.5c shows the use of the optical field in practice: by replaying the hologram with 532 nm light straight out of an optical fiber, and focussing the resultant field with a cylindrical lens, ultracold atoms are confined to the dark centre of the mode, whilst the atoms are prevented from expanding radially by a quadrupole field.

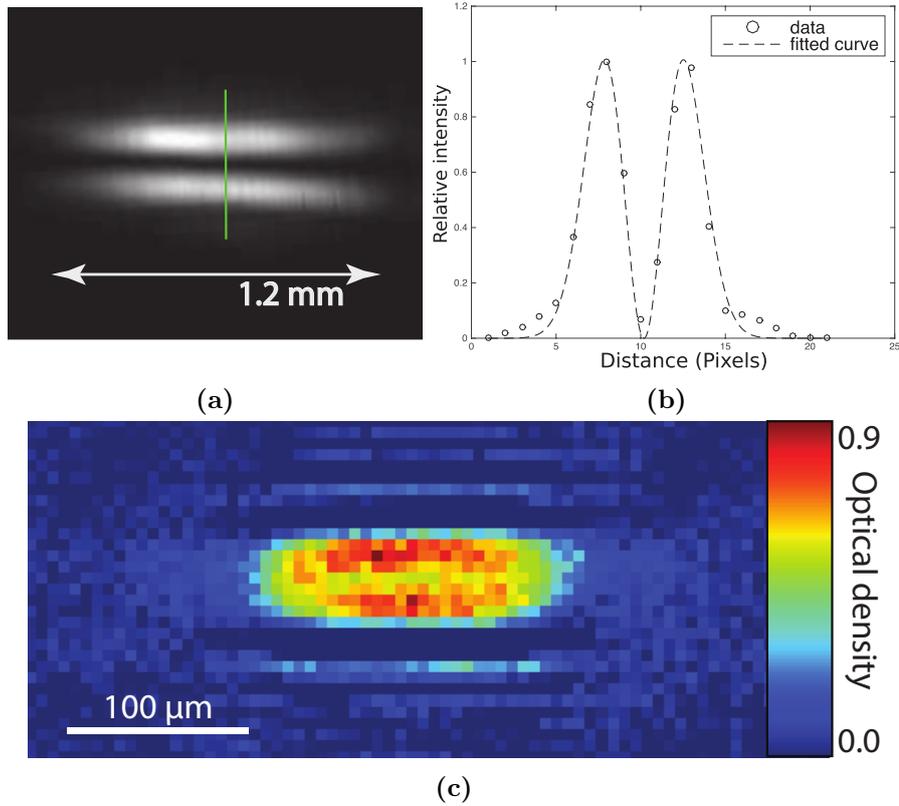


Figure 3.5: The result and use of a $\text{HG}_{(0,1)}$ recorded into a hologram. **a)** Image of the intensity profile of the $\text{HG}_{(0,1)}$ mode at the focus of a lens. **b)** A fit (dashed lines) of a $\text{HG}_{(0,1)}$ intensity profile to the measured intensity profile (circles) taken along the green line shown in the image. **c)** An image of ultra cold Rubidium-87 atoms trapped within the nodal region of $\text{HG}_{(0,1)}$ mode (shown in 3.5a) formed from ≈ 50 mW of 532 nm laser light with a magnetic field providing the transverse confinement of the atoms. Optical density is shown on the right hand side. Reproduced from [4].

3.4 Generating large angular momentum Laguerre-Gaussian modes

One limitation of using an SLM for creating large angular momentum Laguerre-Gaussian modes was discussed in section 3.1, that the largest angular momentum the SLM could impart was limited by the resolution of the SLM. In this section we present a method that allows this upper limit to be increased by a factor of 2 when writing a HOE directly from the SLM⁵.

⁵By iteratively applying this method to the optical mode generated by the HOE that it creates, HOEs that impart greater angular momentum could be written. It is not known what limits this process and what the upper limit to the amount of angular momentum that could be imparted to

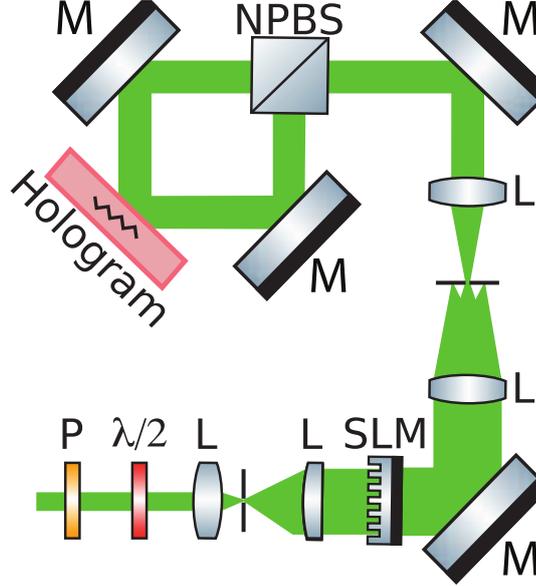


Figure 3.6: Optical setup for recording large angular momentum modes. The optical set up used to record the hologram with twice the angular momentum per photon achievable by diffraction from just the SLM; all labels are as in Fig 3.4. The difference in the number of reflections each mode undergoes must be an odd integer in order for the modes to be mirror images of each other. Reproduced from [4].

Our method of generating $LG_{(0,l)}$ modes with twice the angular momentum of the mode generated by the SLM relies upon interfering the original mode with its mirror image. When an optical field reflects off a surface, the mirror image of the spatial mode is formed, for an $LG_{(0,l)}$, however, this leads the spatial amplitude unchanged under lateral inversion. However, an $LG_{(0,l)}$ mode does have an azimuthal phase $\exp(il\theta)$, where θ is the azimuthal angle and l is the azimuthal order of the mode, which under reflection from a mirror does change. This change simply corresponds to reversing the sign of the azimuthal order, $l \rightarrow -l$.

Utilising this property, we form the interference of a $LG_{(0,l)}$ mode with its mirror image ($LG_{(0,-l)}$) by splitting the $LG_{(0,l)}$ mode generated from the SLM with a NPBS and then recombining the two modes onto the photopolymer with a pair of mirrors, as shown in Fig. 3.6. The critical feature of this setup is that the number of reflections that the two modes undergo differs by an odd number, otherwise the modes would end up the same handedness when they interfere and a simple diffraction grating would be written. As the written HOE must convert the $LG_{(0,l)}$ mode to the $LG_{(0,-l)}$ mode, the HOE must impart $2l$ of angular momentum to the diffracted optical

an optical field is.

field, twice that of either mode. This set up is also particularly stable, as the path length between the NPBS and photopolymer is quite small, which decreases the sensitivity to vibrations. Splitting the modes on a non-polarising beam splitter cube also ensures that the power in each mode is approximately balanced and the polarisation the same.

To demonstrate this technique we generated an optical field with $160\hbar$ of orbital angular momentum per photon, twice the limitation of the SLM, with typical holograms of this type achieving 50% diffraction efficiency⁶. This mode is shown Fig. 3.7a, along with the radial intensity profile, averaged azimuthally. As can be clearly seen, the resultant optical field appears as a thin ring, with a peak to peak diameter of 1.6 mm, while a Gaussian fit to estimate the thickness of the ring results in a $1/e^2$ waist of $20\ \mu\text{m}$ yielding a 80:1 ratio of the diameter to the wall thickness⁷. The intensity variation over the central 1.48 mm is less than 1%⁸. If we fit a power law function to the average radial intensity, a power law exponent of 100 yields an R^2 value of 0.94, further verifying the sharp walls of the optical mode⁹. To verify the angular momentum of the mode, an interferogram of the optical field with its mirror image is shown in Fig. 3.7b, which would show $2l$ azimuthal fringes. The figure shows 320 fringes, as expected for the interference of a $\text{LG}_{(0,160)}$ mode with a $\text{LG}_{(0,-160)}$ mode.

⁶The diffraction efficiency is typically mostly limited by the spatial overlap of the Gaussian and the LG mode.

⁷For creating a uniform flat bottomed potential, this is the most relevant parameter for creating uniform flat bottom potential.

⁸This is the second most relevant parameter.

⁹The greater the exponent, the sharper the walls, for a harmonic potential the exponent is 2.

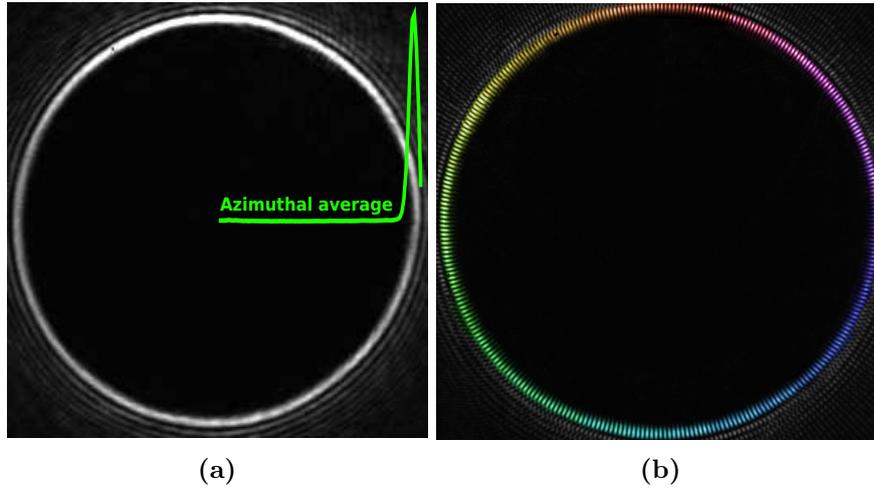


Figure 3.7: The resulting optical field generated by the method described in section 3.4. **a)** The intensity distribution of the optical field. The (green) line is the azimuthally averaged, radial profile. **b)** The interference of the optical field with its mirror reflection, verifying the high angular momentum state. Reproduced from [4].

3.5 Creating other optical elements

During my PhD we experimented with generating a number of other HOEs from physical optics such as lenses. Few were ultimately used on experiments as they offered little advantage over the physical elements themselves, however, they did demonstrate the versatility of Bayfol HX as a recording medium, and some of the information learned could be of use to others. The most interesting optical elements that were recorded into the photopolymer were lenses. When recording a lens into the photopolymer there are a number of interesting properties of the resultant hologram dependent upon how the lens was recorded.

The first thing to note is that when the lens is fully illuminated with light, and is recorded by interfering the focussed light with a collimated Gaussian beam, the results will be a HOE that similarly acts like a lens with the same numerical aperture, however, the focal length of the lens may be varied by altering the distance between the photopolymer and the focus of the beam. When playing back this hologram, the lens can then either function as a positive or negative focal length dependent upon the direction in which the HOE is played.

Next, if the collimation of the Gaussian beam is instead changed to be diverging (converging), then a lens with a larger (smaller) numerical aperture may be written. For example, consider the case of interfering two beams initially of identical size,

passing through two lenses of equal focal length, and thus the resultant Gaussian beams having equal numerical aperture, NA. Now, if one of these lenses was of negative focal length and the other positive, and the two beams were to interfere upon the photopolymer film at an equal size, then the recorded HOE would act as an element that would take a diverging Gaussian beam of numerical aperture, NA, and convert it to a converging Gaussian beam of the same numerical aperture. The HOE would then be behaving identically to a single lens in a $2f$ imaging configuration, which must have twice the numerical aperture of either beam to be able to behave in this fashion. Using the second of these properties, HOEs with greater numerical aperture than physical lenses, which may be available, can be written¹⁰, and the HOE can also be optimised to the type of imaging required, whether it be a $2f$ configuration, imaging to infinity etc.

Finally, one other benefit of volume HOEs is their directionality. As volume holograms are highly angularly selective, HOEs can be formed that will only affect light coming from a particular direction. For example, if one was to want an optical element that could form an image from light collected in one particular direction but not the other, then a HOE could be used for this purpose. This could be especially useful when the desired signal may be especially weak compared to background light levels, allowing the angularly selective nature of the HOE to act as a filter reducing the relative background light level.

3.6 Summary

In this chapter we have described a method for recording holographic optical elements that can reproduce arbitrary complex optical fields that have been generated from an SLM via the methods described in section 2.2. This method utilises Bayfol HX, a photopolymer that has proved both easy to use and versatile, and to date¹¹ we have seen no degradation of a hologram written in late 2015. To verify the effectiveness of this method a HOE was written that could generate a $HG_{(0,1)}$ mode, which achieved 95% of the maximal possible mode conversion efficiency for a single phase only hologram. This mode was subsequently used to confine a condensate axially.

¹⁰There would be a limit to the highest numerical aperture that could be recorded, as at some point one would expect total internal reflection into the film to become a problem, however this limit was not explored.

¹¹April 2019

We also introduced a method for creating HOEs that can generate Laguerre-Gaussian modes that have twice the orbital angular momentum per photon than can be generated with an SLM. This method relies on the symmetry of the amplitude of the mode under reflection from a mirror, and the reversal of the phase gradient under reflection. This method will be critical for generating the large angular momentum modes used in this thesis for ultracold atom trapping.

Building a Bose–Einstein Condensate apparatus

At this point in my PhD I had implemented the creation of arbitrary complex optical fields using a phase only spatial light modulator and had developed the techniques required to record these fields into a HOE, which could then recreate the desired field with high efficiency and fidelity. The creation of high order orbital angular momentum modes to confine and manipulate our condensate was thus achievable and implementable, and the final steps to performing the experiment was to trap the atoms with the optical fields. This was always intended to be performed on the first generation BEC apparatus built at Monash, however, this apparatus was having difficulty producing condensates at the time, and the design of the apparatus was such that high numerical aperture imaging was performed in a glass cell, which the cold atoms had to be transported to, during which a significant loss of atoms occurred.

Given two of our requirements for performing our experiment; reproducible large atom number condensates and high numerical aperture imaging for both creating healing length features in the resulting trap as well as being able to image vortices in-situ, we decided it was time to design and construct a new apparatus that meets these requirements. In the following chapter I will outline the design and construction of all elements of a Bose–Einstein Condensate apparatus; from vacuum system design, construction and the resultant bakeout required to achieve ultra-high vacuum, through to the layout of the laser system required for laser cooling

and trapping, the design of coils for magnetically trapping and manipulating atoms as well as the electronics required for driving the coils, generating RF, shuttering lasers and interlocking the entire system to prevent disaster.

One of the most critical aspects of a BEC apparatus is the vacuum system itself as this is typically the hardest to change once built, and has to meet a number of demanding requirements such as;

- Good optical access for optically manipulating the atoms
- Good optical access for high numerical aperture imaging
- Room for placement of critical components such as the magnetic coils for confining the atoms
- Close proximity and high conductance paths to vacuum pumps to ensure a good vacuum is achieved where evaporative cooling will occur

The design that I settled on utilises a two chamber approach, with a high vacuum side of the vacuum system that acts as a source of cold atoms (the source chamber) for the ultra-high vacuum side (the science chamber) of the apparatus. A 2D MOT is built around the source chamber and provides a beam of pre-cooled atoms that are pushed into the science chamber through a differential pumping tube, allowing a large pressure difference to be maintained. The atoms are then collected in a 3D MOT in the science chamber where they undergo further cooling before they are transferred into a purely magnetic trap for forced evaporation, and eventually into a hybrid magnetic/optical trap in which a condensate is formed and the result can be imaged.

To begin, we will consider the vacuum quality requirements placed upon the system to achieve condensation.

4.1 Ultrahigh vacuum for ultra-insulation

To reach Bose-Einstein condensation and perform experiments on time-scales of the order of seconds, lifetimes on the order of 10's to 100's of seconds must be had to ensure effective evaporative cooling as the effectiveness of evaporative cooling depends greatly upon the ratio of 'good' collisions, which allow rethermalisation between the trapped atoms, and 'bad' collisions, collisions which either heat the atoms or results in atoms lost from the trap. There are typically three mechanisms

through which heat can be transferred, and the vacuum system almost entirely eliminates two of these mechanisms, while greatly reducing the third.

The most common form of heat transfer that is considered is that of mechanical heating, whereby two objects in physical contact exchange heat in order to reach an equilibrium temperature. As the vacuum system is exposed to the laboratory atmosphere, the vacuum system itself is kept at 294 K, nine orders of magnitude hotter than the temperature of the BEC. However, as the BEC is either held in a magnetic or optical trap, the atoms themselves are not in physical contact with the vacuum chamber walls, preventing any heating.

Another method of heating arises from the blackbody radiation emitted by all objects not at absolute zero, this heat is exchanged through photons being emitted from one body and absorbed by the other, and is known as radiative heating. Radiative heating of the BEC, however, is not much of a concern, as the quantised absorption lines of the condensate must be considered, and thus when integrating the power absorbed by the condensate from the blackbody, the power spectrum should only be integrated across the respective line widths of the accessible transitions. This has been calculated[85] to only result in the absorption of single photons every 10^{14} years, causing negligible heating.

The final and most difficult method of heating to remove is due to convective heating. Convective heating is typically considered as the transfer of heat from one object to another through an intermediary fluid, the atmosphere between the objects, however, in this case it is the intermediary fluid itself that causes the heating. In a perfect vacuum this heating mechanism will be prevented, however, even in our ultra high vacuum system of $\approx 5 \times 10^{-12}$ torr, there are still $\approx 10^8$ particles per cm^3 , making convective heating the dominant source of heating for a cold atom experiment. Convective heating is a large problem for producing BEC's, as since only the atomic species of interest are cooled in our apparatus, the dominant gas load on the system, typically hydrogen molecules (H_2 from outgassing of the steel in the vacuum chamber, remains at 294 K and thus a single collision of a cold atom with an atom at this temperature will impart enough energy such that the cold atom is no longer trapped¹. It is possible to estimate the loss rate from a cold atom trap due to the background gas collisions, reference [86] gives a detailed calculation that is not repeated here, but typically a background pressure on the order of 10^{-11} torr

¹As I cannot load much of a MOT in the science cell without the Source MOT turned on, and the lifetime in my magnetic trap is in excess of a minute, this indicates the background pressure of rubidium is quite low.

provides a suitably long lifetime, on the order of a minute.

There is a second, almost contradictory, requirement for creating a BEC, however, a high partial pressure of rubidium is desirable for loading large numbers of atoms into a magneto-optical trap for cooling to reach condensation. As the equilibrium number of atoms in a MOT is limited by the load rate² and the load rate is typically proportional to the partial pressure of the species being cooled[87]³, greater pressure typically results in larger atom numbers when loading from a background vapor. Secondly, the load rate of the MOT also determines the time required to load the MOT to a reasonable size for a starting point for achieving condensation, and the longer this time is, the more time consuming it is to optimise the creation of a BEC.

To satisfy both these requirements with a single vacuum chamber, the chamber is typically split in two; an ultra high vacuum side in which the pressure is kept below 10^{-11} torr for evaporative cooling to reach BEC, typically known as the Science side, and a high vacuum (HV) side where the pressure is kept at $\approx 10^{-7}$ torr⁴, for the fast loading of large numbers of atoms into a MOT, known as the Source side. These two sides of the vacuum system are separated with a low conductance aperture, a differential pumping tube, placed between the two sides, which maintains a pressure gradient of approximately 4 orders of magnitude.

4.1.1 Vacuum pumps

To reach ultrahigh vacuum, the vacuum system must be actively pumped with vacuum pumps. This is as even in an ideally sealed vacuum system with an incredibly low leak rate, there is still a relatively large gas load. This results as most vacuum systems are constructed from 304 or 316 stainless steel, which typically has hydrogen dissolved into it, incorporated undesirably through a number of the techniques used in the manufacturing of the vacuum components, and the refining of the iron from the ore itself. Thus even in an entirely sealed vacuum system, hydrogen can diffuse out of the steel walls of the chamber, creating a gas load upon the system. In our vacuum system we have two vacuum pumps, one is a Nextorr D200 NEG pump, and the other a Titanium Sublimation Pump (TSP), which actively maintain our

²As in equilibrium the loss rate from the MOT, caused by background gas collision and light induced collisions between atoms in the MOT, is balanced with the load rate of the MOT.

³Until the partial pressure becomes large enough to start causing significant loss from the MOT itself.

⁴Note that this chamber is initially pumped out to UHV pressures, but then is filled with the atomic species of interest.

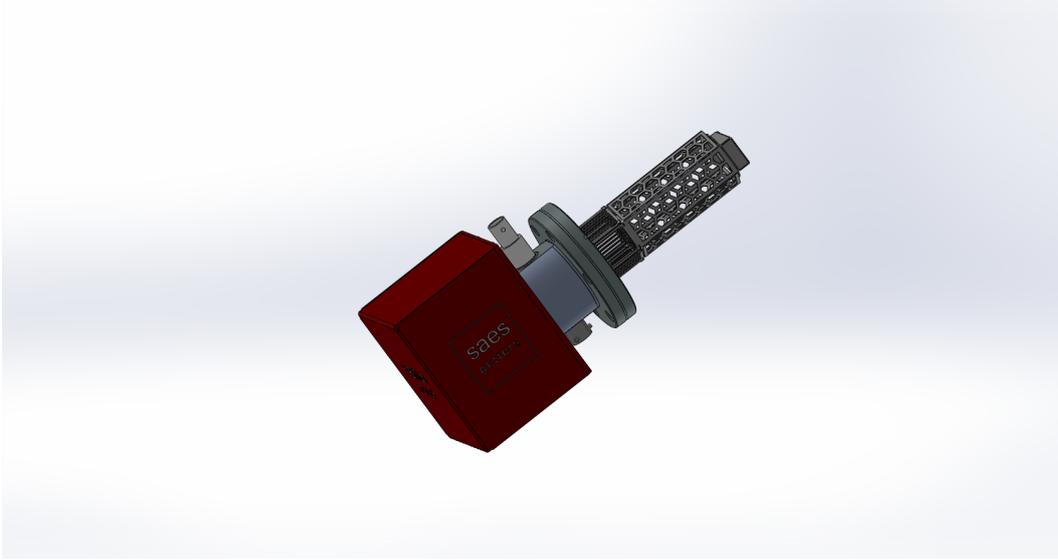


Figure 4.1: Nextorr D200. A Solidworks representation of the NexTorr D200. The red box contains the ion pump which is designed to actively pump the noble gases, whilst the cartridge seen extending from the box is the NEG cartridge that performs most of the pumping, but passively.

vacuum system at $\leq 5 \times 10^{-12}$ torr.

The Nextorr D200 NEG pump is a relatively new vacuum pump on the market, and contains two different types of pump, an ion pump and a non evaporable getter (NEG) pump. Ion pumps are the standard workhorses for producing UHV vacuum chambers, however, they have their drawbacks as they are relatively large, actively consume power and require large magnets to work, which leads to stray magnetic fields in the science chamber. NEG pumps on the other hand are passive and very compact, pumping on the gas load by providing a very high surface area for gases to adsorb onto. However, NEG pumps have their drawbacks as they cannot actively pump noble gases, which is why the ion pump is also incorporated to pump those particles the NEG cartridge cannot. The Nextorr D200 combines these two technologies, providing a pumping speed of $\approx 200 \text{ l s}^{-1}$ for hydrogen[88] which is the main gas load on the chamber, as well as being able to pump the noble gases. Figure 4.1 shows a Solidworks model of the Nextorr D200, the red box contains the ion pump, whilst the cartridge sticking out of the box are the NEG pump elements.

A titanium sublimation pump works very similar to a NEG pump in that it sublimates titanium onto the chamber surrounding the pump when activated, which provides a surface for gas molecules to react with, forming a solid, which then re-

mains there. TSPs have incredibly high pumping rates for reactive gases such as hydrogen making them very useful in UHV applications, however, they are not particularly efficient at pumping much else, which is why the Nextorr is also required. Our TSP is contained in a 4.5” CF flanged, 8.25” long nipple, this provides a surface area of $\approx 750 \text{ cm}^2$ of surface area for the titanium filament to deposit upon, which allows the TSP to provide up to $\approx 1900 \text{ l s}^{-1}$ [89] of pumping speed for hydrogen.

4.1.2 Designing the differential pumping tube

With the expected pumping speed of the vacuum pumps known, we can now turn to calculating the required conductance of the differential pumping tube to maintain the desired pressure difference between our Source and Science chambers. To begin we must first consider the rate at which gas can be pumped out of the system by the Nextorr and TSP. The mass flow rate pumped by the Nextorr and TSP, Q_{pumped} , is given by the total pumping speed, $S_{total} = S_{TSP} + S_{Nextorr}$, and the pressure at these pumps, which we shall denote as P_{UHV} ⁵, and is given by

$$Q_{pumped} = S_{total} P_{UHV}. \quad (4.1)$$

Meanwhile the mass flow rate into the Science side, Q_{sci} is given by the product of the conductance, C , of the impedance and the difference in pressure between the two chambers,

$$Q_{sci} = C(P_{HV} - P_{UHV}). \quad (4.2)$$

In equilibrium the mass flow pumped by the vacuum pumps will equalise with the mass flow through the impedance, yielding the pressure in the Source chamber as

$$\frac{P_{HV}}{P_{UHV}} = \frac{S_{total}}{C} + 1. \quad (4.3)$$

Given we typically desire $\frac{P_{HV}}{P_{UHV}} \approx 10^4$ we can simplify the above to

$$\frac{P_{HV}}{P_{UHV}} = \frac{S_{total}}{C}. \quad (4.4)$$

Thus as $S_{total} \approx 2000$, we thus require $C \approx 10^{-2} - 10^{-3}$ to maintain the desired pressure differential.

⁵Given the high conductance from the Science chamber to the central chamber where the vacuum pumps are located, we shall also take this as the pressure in the Science chamber.

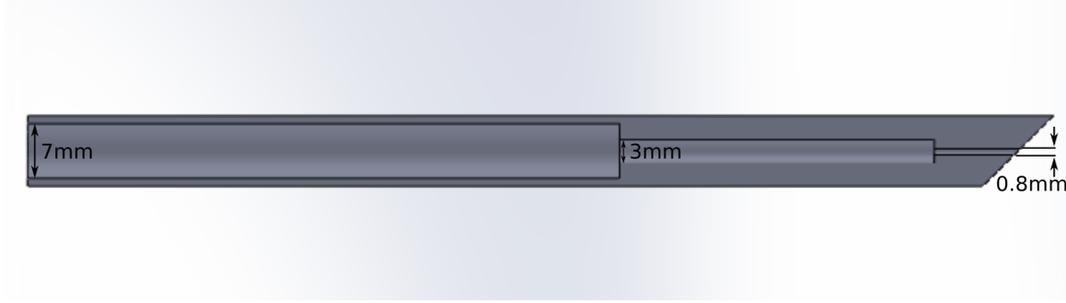


Figure 4.2: A Solidworks model of the designed differential pumping tube. The designed differential pumping tube formed by three concentric holes of varying diameters and lengths drilled into a 12mm diameter grade 304 stainless steel rod. The holes were designed to allow divergence of the atomic beam coming from the 2D MOT while also achieving the desired conductance to maintain a pressure differential of approximately four orders of magnitude.

The differential pumping tube (DPT) between the Science and Source chambers is formed by a long (≈ 120 mm) grade 304 stainless steel rod attached to a 2 3/4" blank flange, with a series of different sized concentric holes drilled into it to form the conductance path for the atoms, as shown in figure 4.2. The DPT consists of three holes, one is 0.8 mm in diameter and is 11 mm long, the second is 3 mm in diameter and is 40 mm long, while the final hole is 7 mm in diameter and is the rest of the length of the tube, 75 mm long. The ratio of the sizes of the concentric holes and the length of each holes was designed to provide a path which would allow a diverging atomic beam of ≈ 60 mrad⁶ to pass through without colliding with the tubes walls, ensuring the flux of atoms travelling from the Source chamber to the Science chamber is not reduced. The total conductance of the DPT can be calculated by first calculating the reciprocal of the conductance of each of the holes, and then calculating the reciprocal of the summation of these reciprocal conductances, as conductance in series adds much like resistors in parallel[91]. The conductance of a hole for a vacuum system in the molecular flow regime⁷ is given by[91]

$$C = \frac{\pi}{6} \sqrt{\frac{2k_b T}{\pi m}} \frac{D^3}{L} \quad (4.5)$$

where m is the mass of the atom, k_b is Boltzmann's constant, D is the diameter

⁶This number is based on references [87], [90] and [6], which typically found a divergence of ≈ 30 mrad, thus 60 mrad will easily allow the beam to traverse the DPT with minimal loss.

⁷The pressure at which molecular flow occurs depends upon the geometry and size of the vacuum system, but occurs when the mean free path of a particle is greater than the diameter of the tube the particle is flowing through, this typically occurs well above high vacuum pressures and is thus guaranteed for our vacuum system.

of the hole, and L is the length of the hole. Using eq. 4.5 the conductance for our differential pumping tube is calculated to be $\approx 3 \times 10^{-3} \text{ ls}^{-1}$

4.1.3 Pressure measurement

Although not crucial to the design of a BEC apparatus, two devices are particularly helpful in diagnosis of vacuum quality issues; a residual gas analyser (RGA) and a Bayard-Alpert ion gauge pressure monitor. An RGA is effectively a mass spectrometer that can provide the user with a measurement of the partial pressure in the vacuum system of a particular mass, whereas an ion gauge provides an accurate measure of the total pressure in the vacuum system, and in particular has a lower x-ray limit, allowing it to measure pressures 10-100 times lower than a typical RGA.

The RGA finds its use throughout the 'bake out' phase, where the vacuum system is heated to 250°C in order to remove as much hydrogen as possible from the steel of the vacuum chamber. During this phase the RGA allows one to look at the partial pressures of the different mass components in the system, allowing one to identify when all of the water is baked out of the system, as well as if there are any leaks through analysing the relative abundance of oxygen to nitrogen as these two gases are typically pumped at different rates, thus finding them in the same relative abundance as in atmosphere can indicate a leak. The RGA is also particularly useful in a final check for any leaks of the vacuum system by squirting helium about the flanges of the vacuum system, as these are most susceptible to leaks, and looking for the appearance of helium on the RGA, hence any possible leak can be localised.

The ion gauge finds its usefulness through continual monitoring of the vacuum system, allowing leaks to be quickly identified and addressed reducing possible downtime of the experiment, as well as providing an accurate measure of the pressure in the vacuum system. However, the pressure in our vacuum system is typically below 5×10^{-12} torr, which is the x-ray limit of the ion gauge, thus there is no accurate measurement for the pressure of the system system below this apart from measuring the lifetime of the atoms held in a trap⁸, therefore the ion gauge only provides an upper limit on the pressure.

⁸This is only accurate if the trap is not limited by other loss mechanisms, such as technical noise from current instability in the magnetic trap or non-resonant scattering of photons in a dipole trap.

4.1.4 Source and Science chambers

The final crucial components to the vacuum design are the Source chamber, which provides the initial source of cold atoms, and the Science chamber, where evaporative cooling to reach condensation occurs and the science begins.

The Source chamber consists of a glass cell fabricated by Precision Glassblowing Inc. in Boulder, Colorado, USA. This cell consists of a 25 mm square profile glass chamber that is 100 mm long, attached to a glass cross that contains two getters in each arm, four of which are rubidium, while the other four are potassium, all procured from SAES. This cell is shown in figure 4.3a. One end of this cell has a window attached to allow the a laser beam to propagate along the length of the cell, to push the atoms loaded into the MOT that is formed around the rectangular cell out towards the Science cell. The other end of the chamber is attached to a minimal length coupler⁹, which allowed the chamber to be easily attached to the vacuum system upon which the DPT was already mounted.

The Science chamber, shown in figure. 4.3b, consists of an octagonal glass cell made from 4 mm thick glass, with 40 mm (external) square optical flats on 6 sides, and glass to metal seals attached to the other two faces that oppose each other. This cell is mounted directly off the main vacuum chamber, a 6" spherical octagon from Kimball Physics, via a custom 6" to 2-3/4" reducer flange. There are two such custom reducer flanges, one attached to either side of the spherical octagon, which are designed with the 2-3/4" flange offset from the centre of the custom reducer by 24.8 mm. These flanges ensure that the cold atomic beam that is emitted from the source side can travel to the Science side unperturbed by the presence of the Nextorr-D200, which protudes into the spherical octagon, maximising the conductance to the Science cell. The other end of the octagonal cell is attached to a bellows¹⁰, before being attached to a gate valve which is capped with a window. This gate valve uses a Kalrez O-ring and is thus bakeable to 200°C, and is in a position such that the apparatus can be extended in the future with an additional chamber.

⁹This allowed the glass chamber, with the end window already attached, to be placed upon an optical rail and a vertical translation stage allowing it to be easily moved into place to prevent the DPT damaging the cell, and the orientation of the cell to be fixed while tightening the bolts.

¹⁰The bellows allows the gate valve to be mounted to the table, as otherwise the rigid vacuum system would be mounted in multiple places and during bakeout could open up leaks due to thermal expansion of the system.

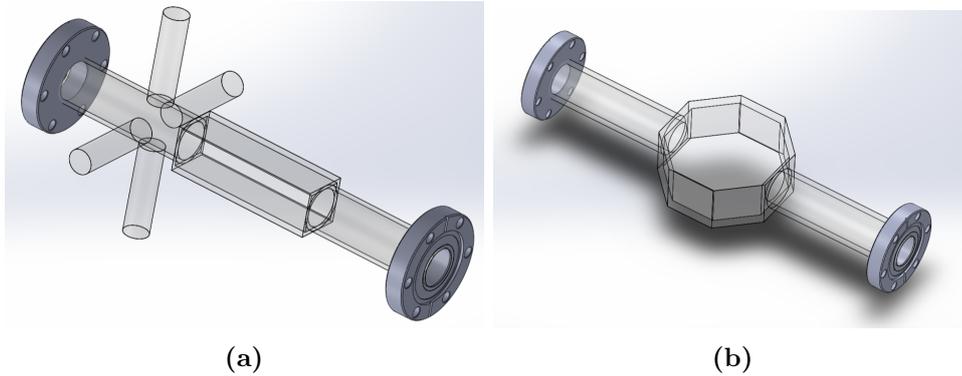


Figure 4.3: Glass cells. a) A Solidworks representation of the source cell, made from a rectangular cuvette with 4 glass tubes attached to one end in which the getters are installed. b) A Solidworks representation of the octagonal glass cell used as the science cell.

4.1.5 Overall system design

Throughout this thesis, I shall define the axes of the experimental apparatus as follows; the X axis of the experiment is coincident with the axis of line of sight along the differential pumping tube, from the source chamber to the science chamber, whilst the Z axis of the experiment is aligned with gravity and the Y axis orthogonal to these two directions. Figure 4.4a shows the a cut through of the resultant final vacuum system design when looking directly along the Source chamber. The system is centred around a 6" spherical octagon¹¹ which has two 6" faces, and eight 2 3/4" ports evenly spaced around its edge. Four of the eight 2 3/4" ports have 780 nm AR coated viewports attached, to allow optical access if the flux from the 2D source is to be quantified, while the other four ports attach the various vacuum gauges and pumps.

The Nextorr D200 is seen on the bottom right hand corner of the chamber, while the TSP is directly opposite it. This design was chosen as the TSP has a metal disk that prevents sublimation from the TSP being directed along the axis of the filaments, hence preventing direct line of sight from the filaments to the Nextorr NEG cartridge which protrudes into the main chamber. This ensures the NEG cartridge does not get coated in titanium, preventing it from pumping. The ion gauge is seen on the bottom left corner of the system contained within a 4 1/2" to 2 3/4" conical reducer, this was chosen to help prevent the filaments of the ion gauge shorting to the walls of the vacuum chamber¹², preventing it functioning. Finally, the RGA is mounted

¹¹Kimball Physics, MCF600-SphOct-F2C8

¹²A design flaw of a previous system.

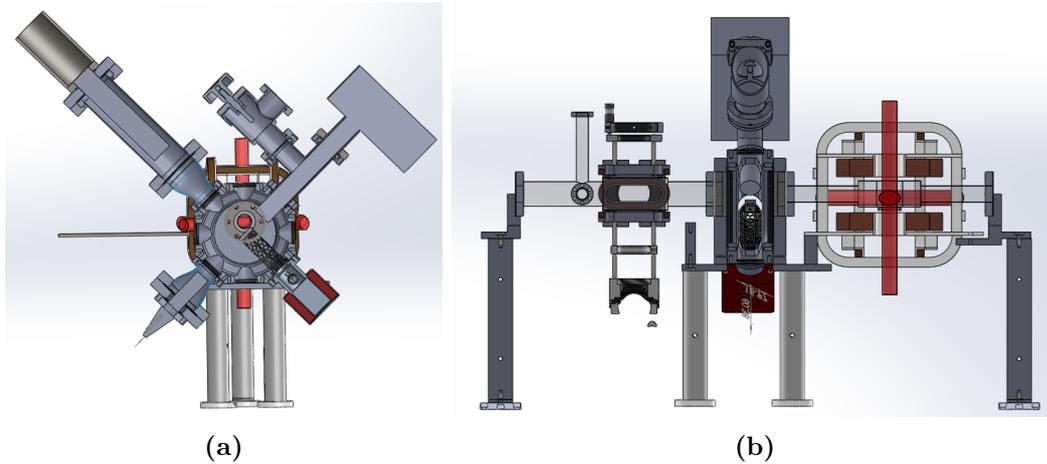


Figure 4.4: Vacuum system perspectives. **a)** A Solidworks model of the final vacuum system looking along the length of the system from the source end, when cut through the centre of the spherical octagon, the centrepiece of the system. **b)** shows the same model but viewed from the side instead and cut along the length vertically to show the path the atoms follow from the 2D MOT formed in the Source chamber on the left hand side, through to the 3D MOT in the octagonal chamber on the right hand side.

in a T piece on the top right corner, with the T being connected to a right angle valve for allowing the initial pump down of the system.

Figure 4.4b shows a cut through of the system when viewed from the side. The left hand side of the vacuum system is the Source chamber, connected to the spherical octagon via the custom 6" to 2 3/4" zero length reducer flange. The differential pumping tube can be seen protruding into the Source chamber, allowing the 2D MOT to be aligned as close to the DPT as practical to ensure the cold atomic beam is as small as possible when first entering the DPT. On the right hand side we can see the Science chamber again attached via the custom reducer flange, and the bellows and gate valve allowing a future extension of the system.

The vacuum system in its entirety in its final design is shown in figure 4.5. The design philosophy for the vacuum system was as simple as possible while maintaining all necessary features. This has been achieved through the use of getters as opposed to a chunk of rubidium in an oven allowing the footprint of the vacuum system to be greatly minimised, while the use of the spherical octagon in the centre allows all necessary components for the vacuum system to be compactly arranged, whilst maintaining good conductance through to the Science chamber ensuring good pumping.

One problem that typically arises with the design of a BEC apparatus that can

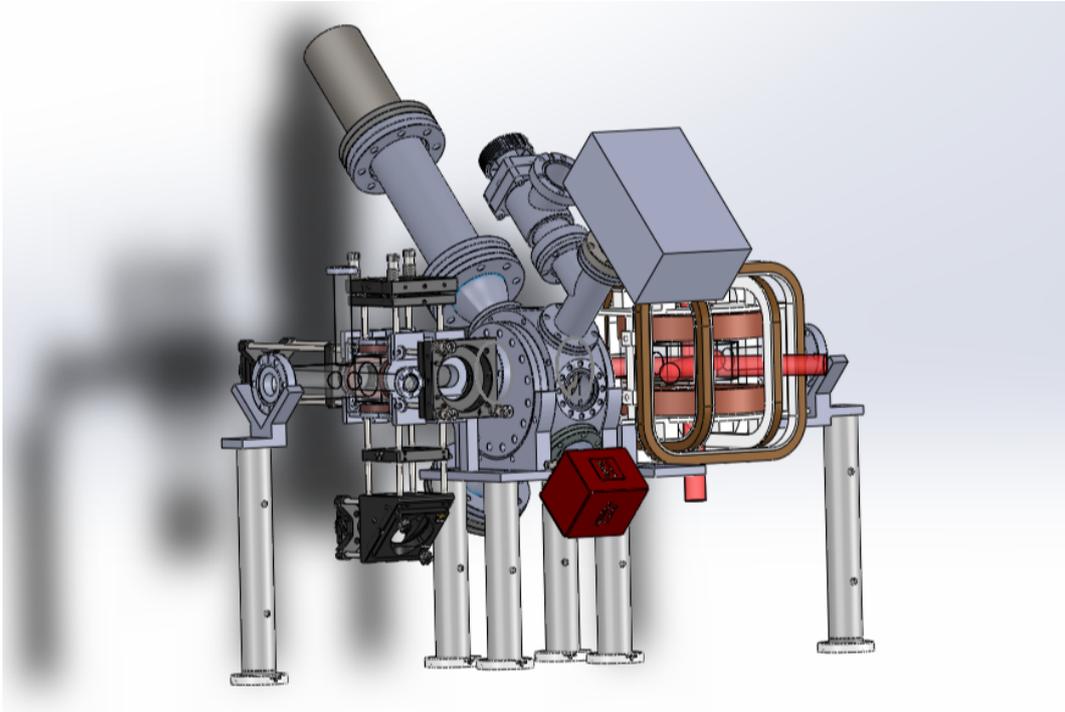


Figure 4.5: Finalised vacuum system. The final vacuum system design with mounts and some optics included. The source end with the 2D MOT optics as well as the coil formers is shown in the forefront, whilst the science side is in the background with the octagonal cell mostly obscured by the quadrupole coil formers.

both produce large BECs, as well as maintain high resolution imaging of the final result, is how to maintain optical access for the MOT beams whilst also getting an objective lens as close as reasonably possible to the atoms to maximise the possible numerical aperture of the imaging system. This problem arises as the objective lens cannot typically be placed on one of the axes that the MOT beams follows, as the MOT beam would then either have to pass through the objective, or the objective be actively moved throughout the experiment.

If the MOT beam were to pass through the objective, then the output beam from the objective would have to be a large diameter collimated beam, as the objective would be the last optical element before the atoms. This is difficult to achieve as large numerical aperture objective lenses typically have short focal lengths, thus if used in any telescopic arrangement the telescope will have to have a similarly short focal length for the input lens so as to not reduce the beam size too much. This telescopic arrangement is typically difficult to implement due to the limited access about the cell, furthermore, desirable MOT beams sizes are typically larger than the

diameter of the first lens in most high numerical aperture objectives, thus significant clipping will occur for large MOT beams. For similar reasons, this also removes the possibility of retroreflecting the MOT beam through the objective in a cateye type configuration, as access to place a mirror at the back focal plane of the lens is lacking.

Alternatively, the objective lens could be moved during the experiment, however, as the field of view and depth of focus of the objective lens are typically very small, whatever translates the objective lens must reliably reproduce the same position of the objective lens to within a few microns each time. By translating the objective lens, vibrations are also likely to be introduced via the translation mechanism, which can lead to imperfect cancellation of interference fringes during the image subtraction used in absorption imaging, reducing the quality of the resultant image.

My solution to this problem was to use a long working distance objective lens (≈ 40 mm), and to slide a very thin (2 mm) custom designed optic, a quarter wave plate at 780 nm with a high reflection coating on the back side, between the top of the objective lens and the bottom of the glass cell during the MOT stage, and then remove it when using the imaging. By using this solution, the MOT beam whose axis is coincident with the imaging axis can be retroreflected by the custom optic without having to pass through the objective lens.

This solution requires a small gap between the bottom of the cell and the top of the objective lens, I used an 8 mm gap, which places the focus of the objective lens 8 mm below the centre of the glass cell. As the focus of the objective lens is below the centre of the glass cell, where the MOT was found to load best, this requires the location of the atoms to be translated.

For my apparatus I designed the quadrupole coils to be centred on the imaging location of the atoms, not the centre of the glass cell. This allows only relatively small bias fields to be required to shift the field zero up to the centre of the glass cell during the MOT loading stage and during the transfer into the uncompressed magnetic trap, and then the translation into the final location is performed by reducing the bias field. I also considered the scenario where the quadrupole coils are instead centred on the glass cell, and a bias field used to translate the compressed magnetic trap to its final location. This would require large currents to be driven through the bias coils for long times, the duration of both RF evaporation in the compressed magnetic trap as well as during the loading of the hybrid trap. It would also be more difficult to maintain the location of the hybrid trap during its loading, as the bias field would have to be ramped in proportion to the changing quadrupole

field as it is decompressed, which may reduce the efficiency of the load.

The translation stage for the retroreflecting optic was made from 4mm PVC, thick enough to allow a 1 mm thick lip at the bottom of the aperture in which the optic sits, 2 mm for the optic itself, and a further 1 mm for a retaining ring to be glued into place to secure the optic. The aperture in which the optic sat was 30mm in diameter, whilst there was a second aperture, used for imaging through, 40 mm in diameter, placed 50 mm further along the translation stage. The separation of these two apertures, centre to centre, was chosen to be the same as the distance the actuator translated (50mm), so that the two positions the translating mount had to be in were at the two extremes of the actuator's range.

4.2 Achieving Ultra-High Vacuum

Achieving ultra-high vacuum is no mean feat, as at such low pressures the smallest impurity in the system can create a gas load large enough to limit the ultimate achievable pressure to above UHV. To achieve the desired base pressure for the vacuum system careful preparation of all components must be undertaken before they can be fixed together. Once the system is constructed, a vacuum system then also typically undergoes a stage known as 'bake out' where the entire system is brought to a few hundred degrees celsius to help outgas any impurities on the surface or within the walls of the vacuum system, typically lowering the lowest achievable pressure by 1-2 orders of magnitude.

4.2.1 Component preparation and construction

In preparation for the construction of a UHV vacuum chamber all components must carefully be cleaned. This is simplest with the aid of an ultrasonic bath and a series of different solvents, each used in turn to clean the previous solvent, which may leave residue. Each vacuum system component¹³ was first cleaned in a bath of trichloroethylene for 30 minutes, used to remove any grease from the manufacturing of the product. This was then followed with an ultrasonic bath of acetone and then methanol for similar amounts of time. Once this was completed, all metal components were then baked in a vacuum oven at low pressure and 250°C for 2 weeks

¹³We did not clean the ion gauge nor glass cells in this manner, as the manufacturers specified they were already UHV ready.

to help remove any additional residue, which may be clinging to the component, and begin removing hydrogen from the components.

Following this initial preparation of the vacuum components, the vacuum system was then constructed. All components of the vacuum system are CF flanged, with high purity oxygen free copper gaskets used to form the seal between each component. Before joining any two components, each knife edge seal must be checked for any damage as this would prevent proper sealing, and all knives edges as well as the copper gasket were wiped clean with methanol. Once the vacuum system was constructed and mounted in place, the entire system was pumped down using a turbo pump being backed by a scroll pump. At this point catastrophe struck with a small implosion from the science cell, with the fused seal between one of the glass to metal seals and the main body of the octagon failing. This required the cell to be removed from the vacuum system, and to be repaired by a local glassblower.

Once the cell was returned and reattached, the vacuum system was then pumped down again, this time to a pressure of 2×10^{-8} torr. This pressure is still far from our ultimate desirable pressure, however, at this point none of the vacuum pumps had been activated, and the system was still dominated by water pressure. The next stage to achieving UHV requires another bake of the entire system in situ. During this baking process, the pressure inside the vacuum chamber quickly rises as many different chemicals and elements outgas from the vacuum chamber walls. This increase in pressure then allows the turbo pump to actively remove these impurities from the system before the system will be closed off.

4.2.2 Bakeout

To bake the vacuum system, an oven must first be constructed. To begin, a number of K-type thermistors were placed across the vacuum system to monitor the temperature. These thermistors were placed in the areas most prone to leaking such as the interfaces of glass to metal seal, to ensure acceptable thermal gradients¹⁴. The structure of the oven was then formed from Thorlabs post's, post holders and metal rack unit stripping, around which a stiff anodized aluminium foil was wrapped, which was then wrapped with heater tapes. This frame ensures that no heater tapes came into direct contact with either of the glass cells, as the glass cells, along with the glass to metal seal to which they are attached, are typically the most susceptible

¹⁴If thermal gradients become too large across these interfaces, the differing thermal expansion of the materials can lead to vacuum system leaks opening up.

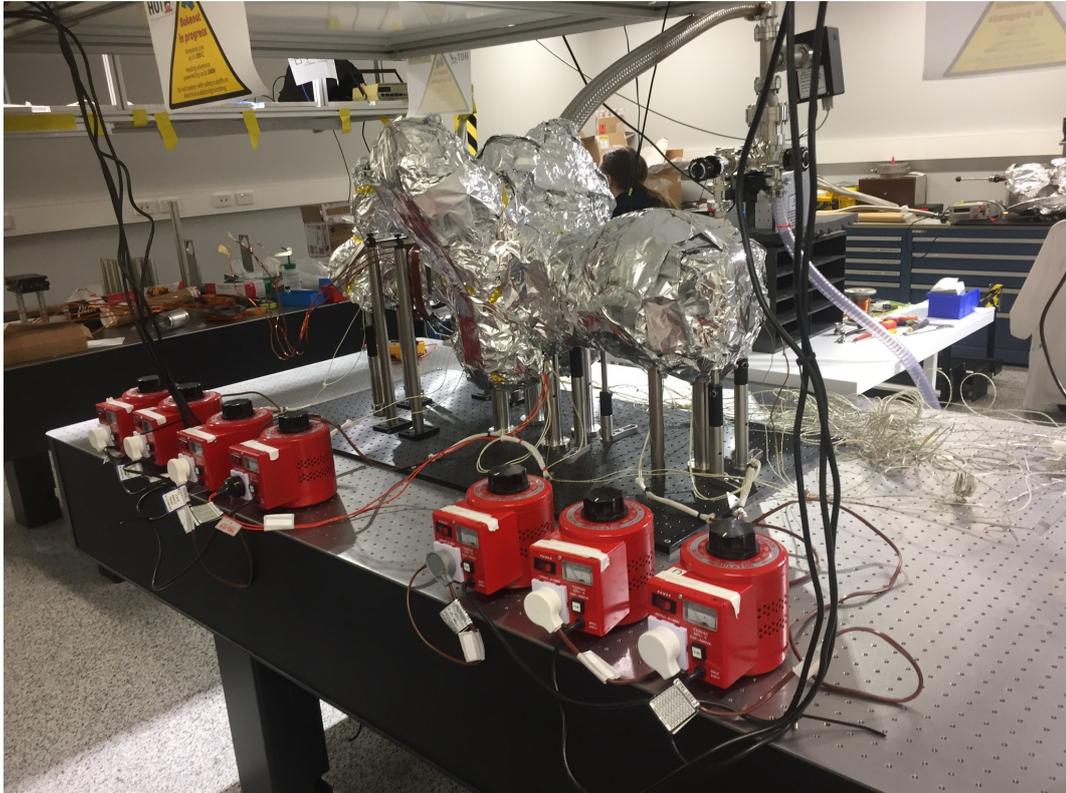


Figure 4.6: The bakeout oven. To achieve ultrahigh vacuum the entire vacuum system must be ‘baked out’ to help remove any water present and reduce the outgassing of hydrogen from the steel. Pictured is the vacuum system once an oven had been constructed about it and the entire system insulated with aluminium foil. On the table are 7 red variacs, that allow independent control of the seven heater tapes used to heat the vacuum system, ensuring the temperature of the vacuum system can be uniformly increased thereby preventing any leaks forming from differential expansion of the materials.

component of the system to thermal stresses leading to leaks. Next the outside of the heater tapes were then wrapped with multiple layers of standard aluminium foil to create a thick layer of insulation to keep the heat in the vacuum system, allowing a greater final temperature.

Seven heater tapes were used across the vacuum system each localised to a particular area to give greater control of thermal gradients, allowing the varying thermal mass of the system to be accounted for. Each heater tape was controlled by a variac, which determines the voltage across the heater tape, and thus power delivered by that heater tape. Figure 4.6 shows the vacuum system once the oven had finished being constructed, with the red variacs in the foreground of the image. Next, the

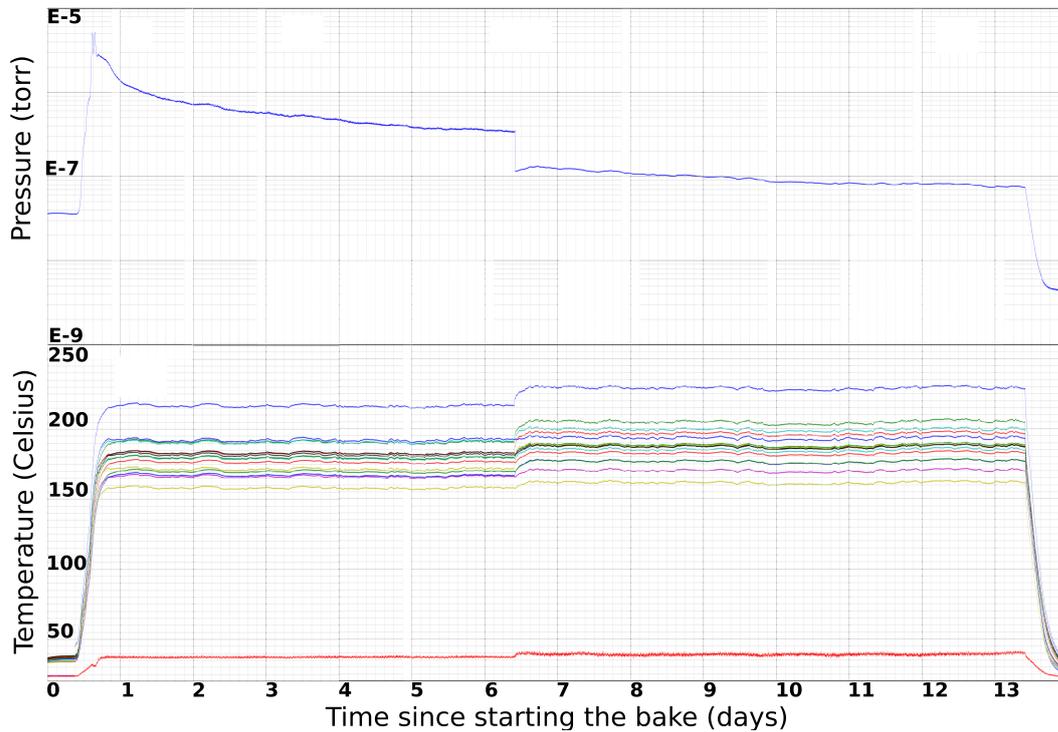


Figure 4.7: Bakeout temperature and pressure. During bakeout the temperature and pressure of the vacuum system was measured. The bottom graph displays the temperature measured by the 15 different temperature sensors on a linear scale, the lowest temperature reading (red) is of a thermistor placed upon a set of magnetic field coils, which had to be in-situ during bakeout, visible on the left of figure. 4.6 The top graph displays the pressure read by the RGA during the bakeout on a logarithmic scale. The pressure is seen to increase by 2 orders of magnitude as the temperature is increased, and falls approximately exponentially by over an order of magnitude over the two weeks. On the day we returned from Easter we slightly increased the temperature further, leading to a small increase in pressure, but there was also a significant drop in the pressure when we started backing the turbo pump with a scroll pump.

power of the heater tapes was slowly ramped up¹⁵, with each of the variac's being independently controlled to ensure the thermal gradients across the system were limited. To ensure the temperature of the entire system did not rise too quickly, we aimed for increasing the temperature by approximately 1°C per minute¹⁶.

Figure 4.7 shows the measured temperatures (bottom) and vacuum system pressure (top) as measured by the thermistors and RGA respectively, during the bakeout period. Over the course of the first day the temperature of the vacuum system

¹⁵As the heater tapes are constant resistance, the power varies with the square of the voltage, and thus the voltage was increased as a square root function in time for a constant heating rate.

¹⁶This is the typical recommended heating speed for glass to metal seals on vacuum components.

rose to 180°C, while the pressure rose by 2 orders of magnitude, note the scale for the pressure is logarithmic, whereas the temperature scale is linear. At this point the system was then left for six days to continue baking, before the temperature was further brought up to 200°C and a scroll pump added to back the turbo pump allowing faster pumping, and left to bake for another 6 days before being brought down to room temperature. During this time it can be seen that the pressures fell approximately exponentially, as the graph shows an approximate linear decrease on the logarithmic scale, initially starting at a pressure of 2×10^{-6} torr, and ending at a pressure of 8×10^{-8} torr. When the vacuum systems temperature was brought down, the final pressure read by the RGA was 4×10^{-9} , an order of magnitude below the initial pressure, and now limited by nitrogen and hydrogen, not water.

Figure 4.8 shows two mass spectrometer scans performed with a residual gas analyser attached to the turbo pump that was used to pump out the vacuum system, one before bakeout and one after. Clearly before bakeout the dominant source of gas in the system was due to water, at a mass of 18, whereas this is mostly removed during the bake. The dominant source of pressure in the vacuum system after bakeout was nitrogen, at a mass of 28, however, this was due to backflow through turbo pump, thus once the vacuum system was sealed this load was pumped out by the NexTorr. From comparison of the two scans there is also a clear decrease in the pressure, with the scale of the pre to post bake scans changing by an order of magnitude.

Once the bakeout was complete the ion gauge, titanium sublimation pump and neg cartridge on the Nextorr were degassed, and the ion pump on the Nextorr activated before the angle valve was closed to seal off the vacuum system and let it pump down to its ultimate pressure, reaching below 5×10^{-12} torr after one day, the lowest pressure the ion gauge can read. At this point we then leak checked the entire system by spraying Helium gas around the leak ports of all flanges, and near the glass cells. During this check we found that two leaks had opened up around the electrical feedthroughs for the getters built into the glass source cell, these were readily repaired through the repeated application of Vacseal over the course of a day. The getters were then degassed, however, this created a large pressure in the Source chamber, which would have required weeks to pump out, so at this point it was decided to perform a second 'mini-bake' of the source cell to 100 degrees to speed up the process¹⁷.

¹⁷I would recommend degassing the getters whilst still pumping on the system, however this was against the manufacturer's recommendations.

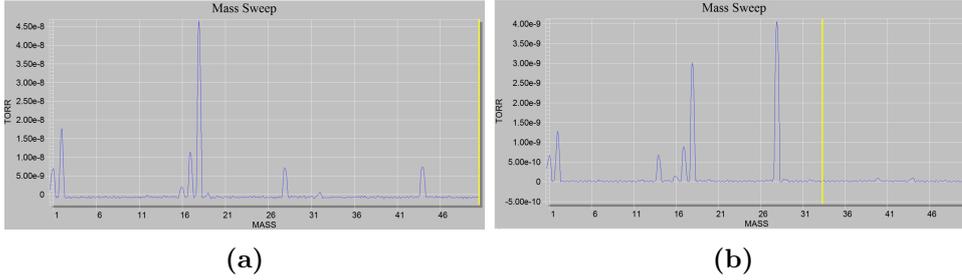


Figure 4.8: RGA scans pre and post bake. **a)** and **b)** show measurements of the partial pressure of gases remaining in the vacuum system made with a residual gas analyser on the turbo pump that pumped out the vacuum system, pre and post bake respectively. Clearly after the bakeout the partial pressures of all the gases have reduced considerably, notably the largest value on the scale of the post bake out scan is almost an order of magnitude less than on the pre bakeout scan.

4.3 Laser cooling system

With the vacuum system built and baked out it was then time to focus on the assembly of the laser system required for the laser cooling of the atoms. The light used for cooling and imaging of our atoms is generated from a single external cavity diode laser (ECDL) and tapered amplifier (TA) system¹⁸, which at the time of writing outputs 1 W of power¹⁹, while the repump light comes from a single cateye ECDL²⁰ providing 80 mW of light. These two lasers provide all the light required for laser cooling atoms down to 30 μ K, as well as the ability to image the result, with all the desired transitions reached by adjusting the frequency of the laser light with acousto-optic modulators..

Figure 4.9 shows a schematic of the hyperfine levels of the $5^2S_{1/2}$ to $5^2P_{3/2}$ transition which has all relevant levels to this thesis included. The figure indicates the lock points and their detunings below the relevant resonance as well as the detunings of the main beam lines.

¹⁸Toptica TA Pro

¹⁹The TA pro is rated to output 1.5 W but over time has drifted down to 1 W, however, as the apparatus still functions well it was not deemed necessary to optimise the power output as this would likely require realignment of most of the laser table.

²⁰Moglabs Cateye diode laser CEL

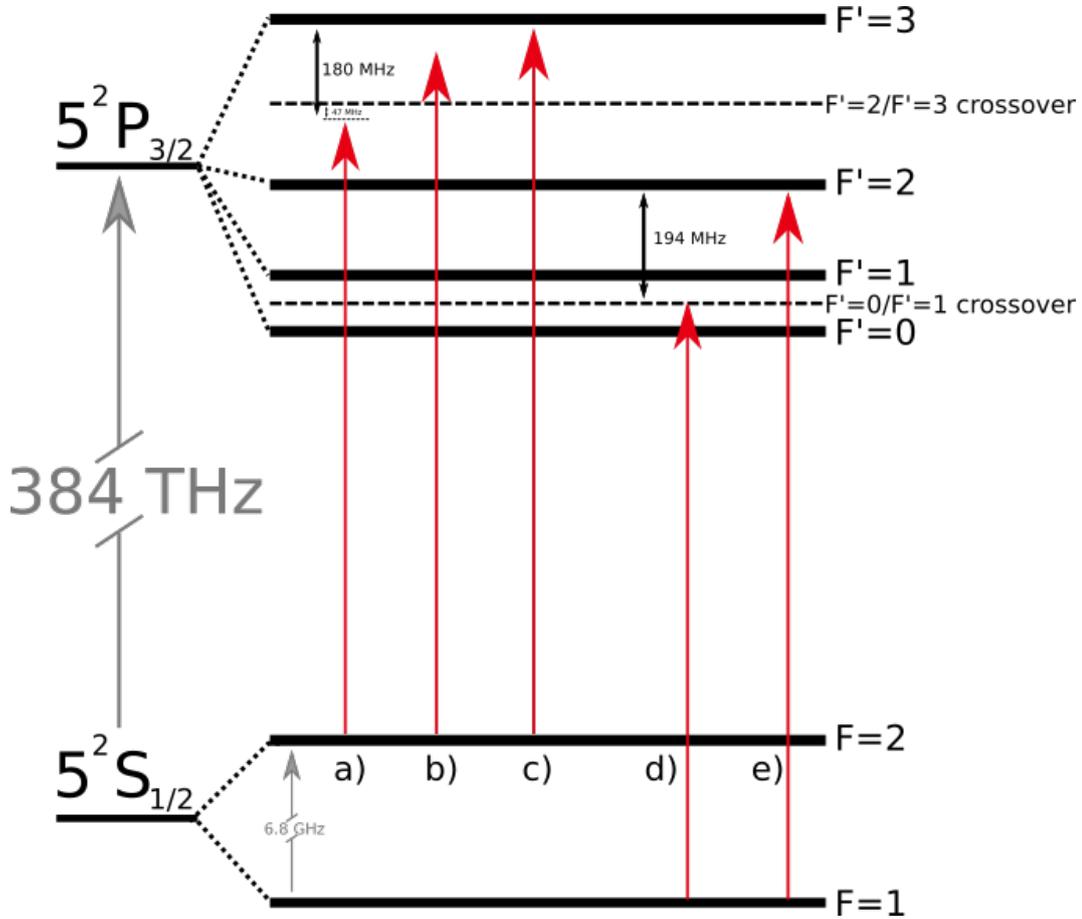


Figure 4.9: Hyperfine levels of $5^2S_{1/2}$ to $5^2P_{3/2}$ transition. The hyperfine levels of the $5^2S_{1/2}$ to $5^2P_{3/2}$ transition are shown, as well as the features that are used in this thesis. Arrow a) indicates the lock point for the cooling laser, 47 MHz below the $F=2$ to $F'=2/F'=3$ crossover, which places the lock point 180 MHz below the $F=2$ to $F'=3$ cycling transition. Arrows b) and c) represent the approximate detunings used for the MOT cooling light (in both the source and Science MOT, and close to the value used for the push beam) and the imaging light. Arrow d) represents the lock point of the repump laser locked directly to the $F=1$ to $F'=0/F'=1$ crossover, placing the lock 194 MHz below the repump transition, $F=1$ to $F'=2$, where the repump light for the Source and Science MOT (arrow e)) are fixed to.

4.3.1 Cooling laser

The first important element of laser cooling is a frequency stabilised laser. To achieve this, approximately 20 mW of the light from the laser diode of the cooling laser is split off before the TA to lock the laser. This light passes through an acousto-optic modulator (AOM) which shifts the frequency of the light upwards by 47 MHz before passing through a telescope to expand the beam, and then into a standard saturation absorption setup. The saturated absorption setup begins with a half

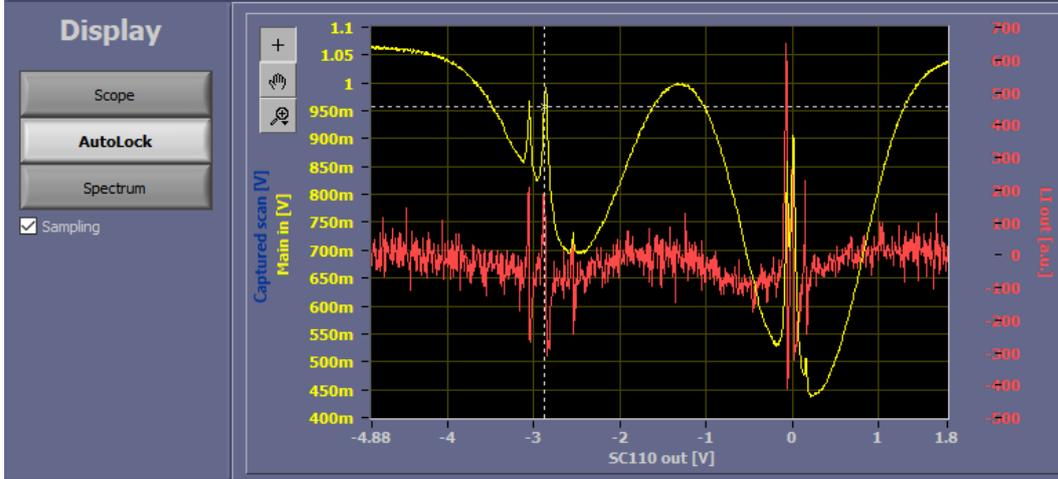


Figure 4.10: Saturated absorption signal for the cooling transition. The Digilock display of the saturated absorption signal (yellow) and resulting error signal (red) used for locking the cooling laser to the $F' = 2/F' = 3$ crossover, indicated by the dotted white lines.

wave plate to rotate the polarisation of the light so that it entirely passes through a polarisation beam splitter (PBS) cube, which passes the P polarisation of the light. The light then passes through a quarter wave plate producing circularly polarised light before entering the vapor cell containing rubidium and being retro-reflected by a mirror normal to the beam. Once reflected the beam passes back along its original path, thus acting as both the probe and pump beam in this setup, while the second pass through the quarter waveplate combined with the change in handedness of the polarisation due to the reflection off the mirror, rotates the polarisation to become S polarised, thus being reflected by the PBS rather than transmitted, after which the light is focussed with a 35 mm focal length lens onto a photodiode.

This setup is useful in its simplicity, requiring only a single beam to produce the saturated absorption signal, however, this comes at the cost of being able to independently chose the power and polarisations of the pump and probe beams. We find this not to be an issue though, with the saturated absorption signal showing well defined hyperfine peaks. Using this saturated absorption setup, we then lock this signal to the $F' = 2/F' = 3$ crossover, with the error signal being generated by a 200 kHz modulation on the current of the laser, generated by the Toptica Digilock 110 module. Figure 4.10 shows a typical measured saturated absorption peaks of the rubidium-87 D2, $F = 2 \rightarrow F' = 1 - 3, 2 - 3$ and 3 transitions.

The expanded beam is critical in our locking scheme, due to the increased error signal it provides. This is as an expanded beam with more power can be used (up to

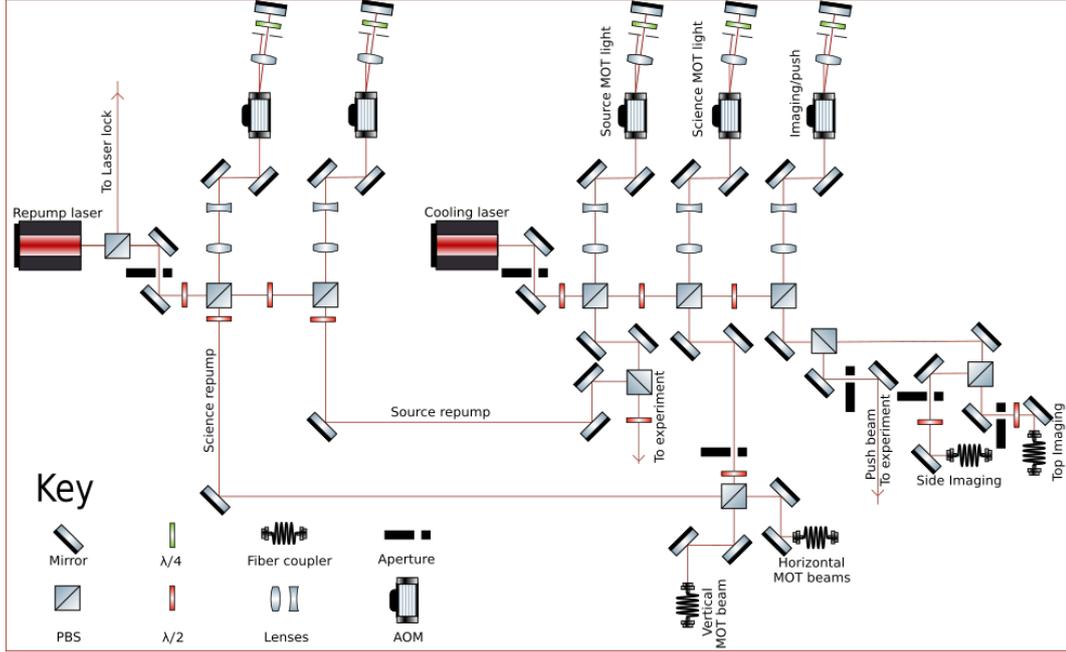


Figure 4.11: Schematic of the cooling and repump laser setup. A schematic of the cooling and repump laser system at the time of writing. Not shown are the saturated absorption setups for locking the lasers. The bottom lefthand side of the image indicates the key, where PBS is a polarising beam splitter cube typically used combined with a half wave plate, $\lambda/2$ to control the amount of power in a given beam line or to combine two beams, $\lambda/4$ is a quarter wave plate, which combines with a lens and a mirror to form a catyeye retroreflector for the double passed acousto-optic modulators (AOMs). The thick apertures are optical shutters, used to extinguish light leakage through the AOMs, whilst the thin apertures are irises used to remove the undiffracted mode in the double passed AOM configurations. Fiber couplers consist of a half wave plate used to align the axis of polarisation of the input light to the fast axis of the fiber, an aspheric lens to focus the light onto the fiber and an x-y mount used to align the fiber mount with the focussed spot.

the point that the photodiode saturates) without saturating the atomic transition²¹. If the original sized beam is used for locking (3 mm) a large current modulation is required to generate a large enough error signal to lock to, which increases the sidebands on the laser mode. Without the expanded beam, modulations 4 times larger are required to produce a usable error signal and when the linewidth was measured with a Fabry-Perot interferometer, the signal displayed significant sidebands. This locking scheme locks the output of the TA to 47 MHz below the $F' = 2/F' = 3$ crossover, placing it 180 MHz below the $F = 2$ to $F' = 3$ cooling transition, well within the reach of a double passed 80 MHz AOM.

²¹i.e. expanding the beam reduces the intensity of the beam (for a fixed power), so the power in the beam can be increased whilst maintaining the local intensity below the saturation intensity of the transition.

The output of the TA is used to generate three separate beam lines, one to create light for the Source MOT, one for the Science MOT, and the other for both imaging and a push beam. The TA output first passes through a $2\times$ cylindrical telescope to correct for the asymmetric mode generated by the tapered amplifier chip, producing a 3 mm $1/e^2$ diameter Gaussian beam with a Rayleigh range of ≈ 9 m[84]. This beam size is chosen so that the Rayleigh range is long enough such that the collimation of the beam does not change significantly over the propagation distances required for all three beam lines²². The circularised beam then passes along the main beam line of the system, which consists of 3 half wave plates and 3 PBS used to split the beam into the three separate beam lines.

Each beam line is identical in configuration until it passes back through the PBS, with the beam line first consisting of a $5\times$ telescope that demagnifies the beam to 0.6 mm for maximum diffraction efficiency through the AOM, two mirrors to steer the beam through the AOM for highest diffraction efficiency, and a cateye configuration after the AOM. The cateye configuration is used to retroreflect the laser beam back through the AOM from a range of incident angles, which allows a greater bandwidth to be achieved as compared to a conventional single pass AOM. This is as the cateye configuration ensures the position and angle of the beam that exits the beamline is independent of the frequency the AOM is driven at, and thus the angle the AOM deflects at. Our cateye configurations consist of a singlet 150 mm focal length lens placed 150 mm from the AOM, with a mirror placed 150 mm further from the lens²³. A quarter wave plate is also incorporated in this setup to ensure the polarisation of the light that exits the cateye is rotated by 90° , as well as an iris to select out the first order mode of the AOM. Upon exiting the cateye configuration, the beam then traces its incident path back, again being deflected by the AOM, and passing back through $5\times$ telescope to be magnified back to its original 3 mm diameter. Once the beam encounters the PBS, however, the beam is now transmitted rather than reflected, as the cateye configuration rotates the polarisation by 90° .

Typically we find the efficiency of a double pass configuration such as ours, using Gooch and Housego 3080-122 AOM's that have a typical maximal diffraction efficiency of 80%, is on the order of 64%, indicating the highest efficiency obtainable. We find that when using the reflection off the PBS to generate the light for the beam line that the highest power beamlines should come first. This is because the purity

²²The beam size cannot be too much larger either though, or not all the light will pass through the laser shutters.

²³We find the accurate placement of the lens with respect to both the AOM and mirror is critical in achieving a large bandwidth for the double-pass setup.

of the polarisation is worse for a weak relative reflection off the cube, as there is always a minimum percentage of light of the wrong polarisation that will be reflected. Hence if only a small amount of power is to be used the purity of the polarisation will be worse. This polarisation purity becomes an issue as AOMs have a preferred axis of polarisation, and thus can exhibit thermal birefringence, meaning that the overall efficiency of the beamline becomes dependent upon the temperature of the AOM. Alternatively, this can be overcome by using only the transmitted beam of the PBS for each beam line, however, this creates more complexity in the geometry of the layout.

Once each beam line has been double passed, the three beam lines then differ. The Source MOT light gets coupled into the 2D Source MOT with free space optics. We find this to be adequate as the 2D MOT is relatively insensitive to small deviations in the alignment and spatial purity of the mode. This allows us to have $\approx 2\times$ more power in the 2D MOT, resulting in a total of 225 mW in the 2D MOT, coming from 313 mW injected into the Source MOT beamline initially. The Science MOT light first passes through a laser shutter and half wave plate, before being split into two beamlines on a PBS. One path leads to a fiber coupler for the light for the vertical Science MOT beam, and the other leads to a fiber coupler for the two horizontal Science MOT beams. It is upon this PBS that the repump light for the Science MOT is also combined with the trapping/cooling light, however, it is only coupled into the horizontal MOT beams. The input power in the Science MOT beam line is 275 mW, while the output power is 190 mW. The splitting in power between the vertical MOT beam and the two horizontal MOT beams is roughly 1:2, and thus after fiber coupling the power in each of the MOT beams is approximately equal at 30 mW.

The final beam line is the imaging and push beam line, which first passes through a half wave plate and another PBS. This splits the beam into a push beam reflecting off the cube through a laser shutter to be coupled into the 2D MOT setup with free space optics, while the transmitted beam is then passed through another half wave plate and PBS to be split into a side imaging and top imaging beam. Both imaging beams pass through a laser shutter and are then fiber coupled to be transported to the Science MOT. The incident power to the imaging/push beam line is 63 mW resulting in 37 mW exiting the double pass configuration, with 14.5 mW going into the push beam, 14 mW into the side imaging beam and into the top imaging beam.

A schematic of the entire laser system for the MOT, imaging and repump lasers is found in figure 4.11. The schematic shows the approximate layout used to create the

2 beam lines of the repump laser light and the 3 beam lines of the MOT and imaging light, and how they are combined. In the bottom left hand side of the image is a key explaining what the symbols represent, in particular it should be noted that;

- PBS is a polarising beam splitter cube
- $\lambda/2$ is a half wave plate
- $\lambda/4$ is a quarter wave plate
- A fiber coupler consists of a half wave plate, aspheric lens in a z translation mount, and a fiber mount in an x-y translation mount.
- Lenses are used to either form telescopes, or in the cateye retroreflector configurations
- Apertures can either be electronic laser shutters or irises
- AOM are acousto-optic modulators

4.3.2 Repump laser

For creating a MOT with rubidium-87, a second laser frequency resonant with the $F = 1$ to $F' = 2$ transition is required. This is because atoms being cooled on the $5S_{1/2}F = 2$ to $5P_{3/2}F' = 3$ cooling transition have a small probability of being excited into the $5P_{3/2}F' = 2$ state, from which the atom can fall into the $5S_{1/2}F = 1$ 'dark state' which is not resonant with the MOT cooling/trapping light. Although the probability of this is rather small the atoms will almost definitely fall into this dark state, as the atoms typically have to undergo tens to hundreds of thousands of absorption/emission cycles on the cooling transition to be cooled down to the millikelvin temperatures achieved in a MOT. Thus to be able to recapture these atoms that have fallen into the dark state, a repumping beam is needed that optically pumps the atoms back up from the $F=1$ to the $F'=2$ excited state. Light at this frequency is typically either generated with a second ECDL, or alternatively an electro-optic modulator, as the cooling transition and repump transition are separated by 6.8 GHz. We choose to generate the light for our repumping beam from a second ECDL, a MOGLabs cateye laser, that is locked directly to the $F=1$ to $F'=0/F'=1$ crossover, as this is conveniently placed 194 MHz below the desired repump transition, well within the reach of a double passed 80 MHz AOM. The saturated absorption signal (yellow) and resulting error signal (blue) used for locking

is shown in figure 4.12, with the largest feature being the lock point.

Light from the output of the laser first passes through an optical isolator and half wave plate, before being split on a PBS for locking and the main beam line. The reflection from the PBS is used for locking the laser. This setup has a similar saturation absorption setup as the one used for locking the main laser, a half-waveplate and PBS is used to reflect the incident light along a path consisting of a beam expanding telescope, a quarter wave plate, a vapor cell containing rubidium, and a mirror, which retroreflects the incident beam. This saturated absorption setup behaves identically to that used to lock the main laser, apart from the placement of the beam expander. The beam expander used in the repump locking setup is placed within the retroreflected path rather than before, as the photodiode used to measure the saturated absorption signal is a MOGLabs photodiode, which comes with a lens preinstalled on the front, and can only focus a beam on the order of a few mm, thus the beam exiting the setup must be of similar size to that which enters. This locking setup uses ≈ 5 mW of power, as we find the error signal is greatest for the largest intensity passing through the cell that does not saturate the photodiode. The error signal for the locking scheme is similarly generated by the MOGLabs diode laser controller, with the controller adding an ≈ 250 kHz modulation onto the laser current, the demodulation of which provides the error signal.

The main beamline used for the experiment contains 75 mW and is split into two beamlines, one for the Source MOT, and one for the Science MOT. Each of these beamlines is identical to those described above for the main laser, the beams pass through a telescope demagnifying the beam, a pair of mirrors to steer the beam for maximum diffraction efficiency through the AOMs, and a cateye retroreflector setup. Approximately 35 mW of light passes along the Source repump beamline, which is free space coupled into the Source MOT setup, combined with the cooling light on a PBS, with 15 mW of light making it through the setup. The Science repump beamline is combined with the science cooling light on a PBS also, and fiber coupled into the horizontal MOT beam fibers. 40 mW of repump light passes along the Science repump beamline, with a total of 10 mW making it through the fiber coupled to the experiment. As previously mentioned, figure. 4.11 shows a schematic of the repumping laser system in its current state at the time of writing.

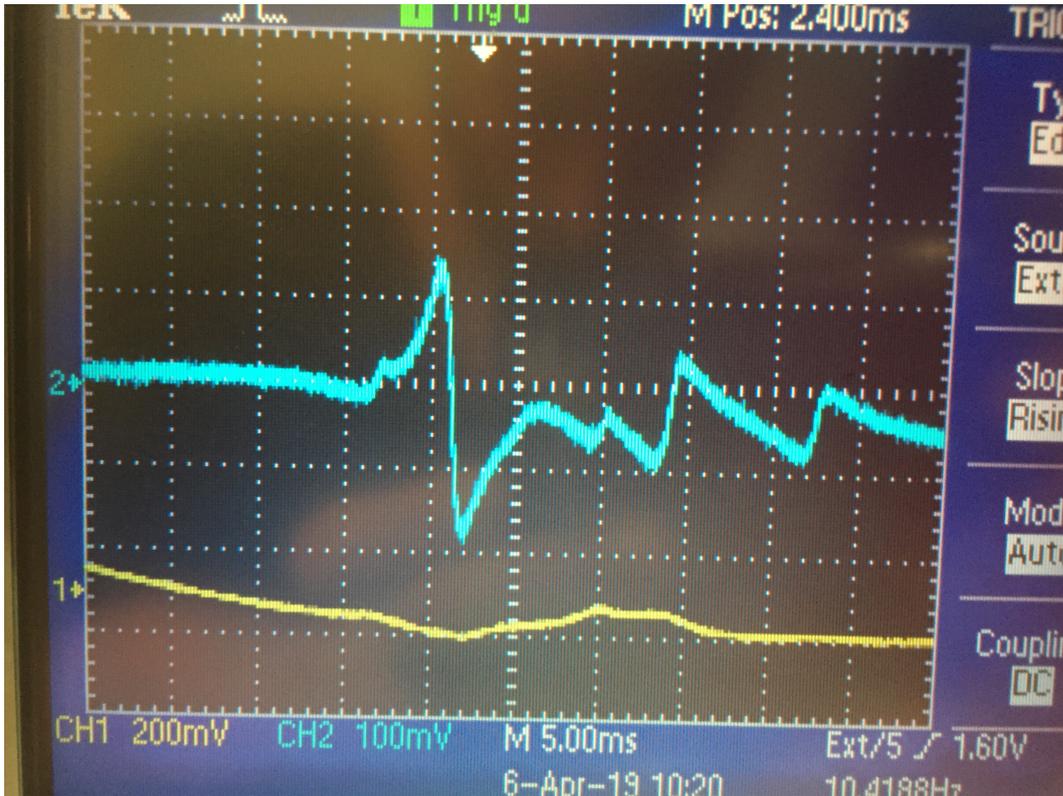


Figure 4.12: Error signal for the repump transition. An image of an oscilloscope displaying the error signal (blue) and saturated absorption signal (yellow) used to lock the repump laser to the $F = 1$ to $F' = 0/F' = 1$ crossover transition. The largest feature in the error signal is used as the lock point

4.4 Quadrupole coil design

There are a number of different types of magnetic traps that can be used to confine cold atoms, varying in size and complexity. For my apparatus, I chose to use a quadrupole magnetic field to magnetically trap the ultracold atoms, generated by a set of coils external to the vacuum system.

There are a large number of trade offs and requirements that must be met when designing quadrupole coils for cold atom trapping, with each factor that must be considered typically impacting upon the other factors, generally resulting in an iterative process until a final solution is achieved. The magnetic coils play a large role in a Bose-Einstein condensate apparatus as they determine the properties of the magnetic trap used to confine the atoms, which in turn plays a large role in the efficiency of the forced evaporative cooling stage necessary to achieve Bose-Einstein

condensation.

When designing coils for a magnetic trap the following properties must be considered:

- The geometric size of the coils
- The inductance and resistance of the coils
- The power dissipation of the coils
- The characteristics of the power supply driving the coils

Typically there are two solutions to designing coils that meet all the requirements for cold atom trapping; coils made of a large gauge hollow wire, which can be actively cooled, and coils made of much smaller gauge wire, which are passively cooled. The design considerations and how the two different solutions account for them are detailed in the following subsections.

4.4.1 Geometry of the coils

The geometry of the magnetic coils is dependent greatly upon the vacuum system design, as the coils must be centred on the the position in which the BEC will be prepared. The main consideration with regards to the geometry of the coils comes typically from the required optical access for imaging as well as the MOT beams, which will set the minimum inner diameter of the coils²⁴, while the chamber size will set the minimum coil separation. These two properties of the geometry, the inner coil diameter and the coil separation, both greatly affect the resulting magnetic field gradient from the coils for a given current density, and thus these geometric constraints will be considered first, and the number of turns of the coils²⁵ and the required currents then calculated based upon these constraints.

One further consideration is the orientation of the coils. The rotational symmetry axis of the coils typically shares the same axis as gravity, and as the magnetic field gradient along this axis is twice as strong as along the other axes and aligning these axes ensures lower currents²⁶ are required to support the trap against gravity.

²⁴Large enough to pass the MOT beams and/or for the desired numerical aperture for imaging.

²⁵Note this affects the geometric size of the coils, more turns increases either the radial extent of the coil or the number of layers, which may be limit optical access or be limited by other vacuum components.

²⁶Which allows faster switching speeds.

4.4.2 Switching speed of a coil

The inductance and resistance of the coils can greatly affect how much an atomic cloud will heat up when being loaded into a magnetic trap. This is because most BEC apparatus have a stage in which the magnetic field gradient is reduced to zero so as to cool the atoms to a lower temperature²⁷ before loading them into the magnetic trap. The magnetic field must then be quickly switched on to at least a large enough gradient to hold the atoms against gravity²⁸, with this typically known as 'the catch'. This generally must occur on a timescale of the order of 1 ms to efficiently capture the atoms, as during the catch the cooling light is switched off and the atoms are in freefall. Long switch on times thus lead to significant increase in velocity of the atoms as they fall under gravity as well as a change in position, leading to significant heating once the atoms are caught²⁹.

The speed at which you can switch a current on through a coil is limited by the inductance of the coil and the voltage trying to drive current through the coil. When current begins to flow through the coil a back emf, V_L , is induced across the coil opposing the change in magnetic flux through the coil. This back emf is related to the rate of change of current through the coil, $\frac{dI}{dt}$, and the inductance L by[92]

$$V_L = -L \frac{dI}{dt} \quad (4.6)$$

thus the rate of change of current through a coil is proportional to the driving voltage, and the inductance of the coils. The self-inductance of a single coil is given by[92]

$$L = \frac{\mu N^2 A}{l} \quad (4.7)$$

where N is the number of turns of the coil, A the cross sectional area of the coil, l the length of the coil and μ the magnetic permeability of the coil. As the driving voltage of the coils is fixed by the power supply, the maximum driving voltage of the power supply must be considered when calculating the maximum permissible inductance of the coils.

Throughout the following sections we will consider two coils, coil one, made from a

²⁷Through a stage of polarisation gradient cooling.

²⁸For rubidium-87 atoms in the $F = 1, m_f = -1$ state this is $\approx 30.5 \text{ Gcm}^{-1}$.

²⁹An increase in velocity will result in an increase in temperature, as the atoms initially oscillate back and forth in the magnetic trap once they are caught and as the atoms thermalise this motion damps out, but increases the overall temperature of the sample.

large diameter wire, and coil two made from a smaller diameter wire, say by a factor α , for comparison. We shall assume that the two coils have the same cross sectional area, as the cross sectional area is generally set by other considerations such as the required optical access³⁰. From equation 4.4.2 it is clear that as the second coil has more turns, by a factor of α , this coil shall also have a greater inductance, by a factor of α^2 . However, if we consider the current density in the two coils, as the current density is the key physical quantity in determining the magnetic field gradient³¹, then for the first coil to achieve the same current density as the second, a greater current, by a factor of α , must be run through the coil. Given this, via equation 4.4.2 it is clear that the generated back emf on the coil made from the larger diameter wire is only a factor α smaller than that generated on the first coil to switch the same overall current density.

4.4.3 Power dissipation of a coil

When maintaining large currents through magnetic coils, as required to generate large magnetic field gradients to tightly compress our atomic trap, a large amount of power is dissipated as ohmic heating. This resistive heating is readily calculable,

$$P_{Ohmic} = I^2 R \quad (4.8)$$

where I is the maximum current required for the tightest trap, and R is the resistance of the coils.

Considering our two coils from before, as the larger diameter wire has to carry a factor of α more current to achieve the same current density, naively we might assume that the ohmic heating in this coil would be a factor of α^2 larger. However, as the resistance of a wire is inversely proportional to its cross-sectional area³² and proportional to its length, the coil made from the smaller diameter wire will require a length that is a factor of α longer, and a cross sectional area that is a factor of α smaller, leading to a resistance that is α^2 larger. The overall power dissipated by the two coils should therefore be approximately equal. This analysis ignores the possibility of having a water cooled core for the larger diameter wire which would

³⁰This is akin to assuming that the cross sectional area of the first coil's wire is greater than the second coil's wire by a factor of α , and hence in the same cross sectional area the smaller diameter wire can have α more turns, assuming a 100% packing fraction.

³¹Given the cross sectional area is fixed.

³²As we are considering DC current only here, the skin effect, where the current density throughout the cross section of a wire is not uniform in alternating currents, need not be considered.

decrease the wires effective cross sectional area, and ignores any possible differences in the packing factor between the two cases.

The power dissipation of the coils is critical to the experiment as if the cooling, whether passive or active, is not enough to remove the heat generated by the coils and maintain a constant temperature, then runaway heating can occur. Runaway heating occurs due to a positive feedback cycle arising from the dependence of the resistance of the wire upon its temperature. It is well known that copper wire has a positive temperature coefficient^[93], and thus the resistance of the wire will increase with increasing temperature. Thus for a fixed current driven through the coils this will also increase the power dissipated, which in turn heats the coil further, leading to a positive feedback cycle. This can lead to the eventual melting of the enamel about the wire, creating an electrical short, or alternatively becoming hot enough to damage surrounding objects³³, thus it is critical that this regime not be reached during an experiment.

4.4.4 Power supply considerations

The power supply is the easiest element of the designs to change, however, is also the most expensive. The key limitations the power supply places upon the coil design are the maximum voltage the power supply can generate, and the maximum current it can supply. As shown in section 4.4.2, coils made from smaller diameter wire will typically have a greater resistance and inductance, and will thus generate a greater back emf when switching them on for the catch. This means that a higher voltage power supply will be needed to provide enough voltage to counteract this back emf and provide a reasonable switching speed. If using a simple transistor circuit to switch these coils, this large voltage requirement for the catch stage may lead to high power dissipation in the transistor when only low currents are being run through the transistor for the MOT stage, and thus active changing of the power supply voltage may be required. Our coils are driven by a Sorensen SGA power supply, which can generate up to 150 Amps at 100 Volts and has the ability to have the voltage changed by either an analog voltage fed into the back of the device, or a resistor placed between two pins on the device, making switching the voltage across the coils during an experiment possible.

³³Coil formers, etc.

4.4.5 Final quadrupole coil design

Based upon the considerations above, a larger diameter wire was ultimately chosen as this would reduce the inductance allowing a faster switching speed. The final coil design settled upon used wire with a square cross-sectional area, 4.6 mm external width and a 2.5 mm diameter hollow core allowing the core to be actively cooled with circulating water cooled to 19°Celsius³⁴. The design resulted in two coils separated by 56 mm, with each coil consisting of 6 turns in the vertical direction, and 8 turns radially, for a total of 48 turns per coil. The separation between the two coils is much larger than the 40 mm we were limited to by the glass cell, however, this increased separation allows the coils to be placed asymmetrically about the glass cell so that the quadrupole field zero is 8 mm below the centre of the glass cell. This asymmetry allows the imaging objective to be placed 8 mm below the bottom of the glass cell, far enough away to allow a quarter wave plate with a dielectric reflective coating on the back side to be placed between the objective and the cell. The purpose of this quarter waveplate with a high reflection coating is to retroreflect the vertical MOT beam, and thus prevent it having to pass through the objective lens³⁵, while also ensuring the objective lens need not move within the experiment to image the condensate.

With this coil design the quadrupole coils can achieve a magnetic field gradient of $1.75 \text{ Gcm}^{-1}\text{A}^{-1}$, allowing the desired final field gradient of 200 Gcm^{-1} to be achieved at 114 A, and the coil pair had a theoretical inductance of $312 \mu\text{H}$. These values require 17 A to catch the atoms at 30 Gcm^{-1} , and thus the back emf generated by these coils when switching over 1 ms is 5.5 V, well within the range our power supply can handle. However, during the experiment we typically catch at 80 Gcm^{-1} , requiring 46 A of current that is switched over 2 ms, this generates a back emf of 7 V. All simulations of these coils were performed with Radia³⁶, as these simulations take into account the true geometry of the coils by simulating the coils as a set of independent conductors.

³⁴This hollow core reduces the effective cross sectional area of the wire, but allows higher currents to be sustained due to the active cooling preventing runaway heating.

³⁵From the previous apparatus it was known that getting a large diameter MOT beam through the objective lens that was not distorted was very difficult due to the high numerical aperture of the objective lens.

³⁶Radia is a magneto-statics computation package created by ESRF for solving technical problems for synchrotron light sources.

4.4.6 Bias coils and formers

Bias coils are typically configured in Helmholtz configuration to generate a uniform magnetic field between the coil pair for the cancelling of stray magnetic fields, and the ability to move the zero of a quadrupole magnetic field around. The requirements for bias coils are generally much less strict, as much smaller current densities are required meaning both switching speeds and power dissipation do not become technical issues.

In the current design, a large bias field is needed along the Z axis, as this bias field pushes the quadrupole field zero from 7 mm below the centre of the glass cell to the centre for loading the MOT and the magnetic catch. For the 70 Gcm^{-1} field gradient used when catching the atoms, the Z bias coils must be able to produce and maintain a field of at least 49 G to maintain the field zero position during the catch. To achieve this the Z bias coils were wound directly on the outside of the quadrupole coils using the same square, hollow wire, with 6 turns in the vertical direction and 4 turns radially for a total of 24 turns per coil. This configuration allows the coils to generate 4.5 GA^{-1} , thus requiring 11 A to achieve the desired field bias, well within the capabilities of most power supplies. The Z bias coils were made out of the same wire as the quadrupole coils to allow them to also be actively cooled, enabling future experiments that might require large bias fields, with bias fields of up to 700 G being readily achievable with these coils.

The X bias coil pair and Y bias coil pair were both designed using rectangular formers, as this allows the coils to be placed much closer together, without being prohibitively large (as they have to fit around the quadrupole coils). This change from typical round coils to rectangular coils reduces the uniformity of the generated bias field, however, as these coil pairs are large in size³⁷, the effect on the uniformity of the field over the region occupied by the atoms is negligible. Square coils, however, are not easy to wind as copper wire does not like relatively sharp bends. Thus to make the coils easier to wind and reduce the chance of damaging the wire and enamel, rounded edges of 37mm radius have been used on the formers. The X bias pair were wound onto 170 mm square formers with 5 layers of 12 turns (60 turns total), with the pair separated by 85 mm, achieving 6.25 GA^{-1} . The Y bias pair were wound onto 200 mm square formers with 4 layers of 4 turns (16 turns total), and separated by 100 mm producing a field of 1.28 GA^{-1} .

To position the coils and align them geometrically to the glass cell, a set of formers

³⁷Relative to the size of the atom cloud.

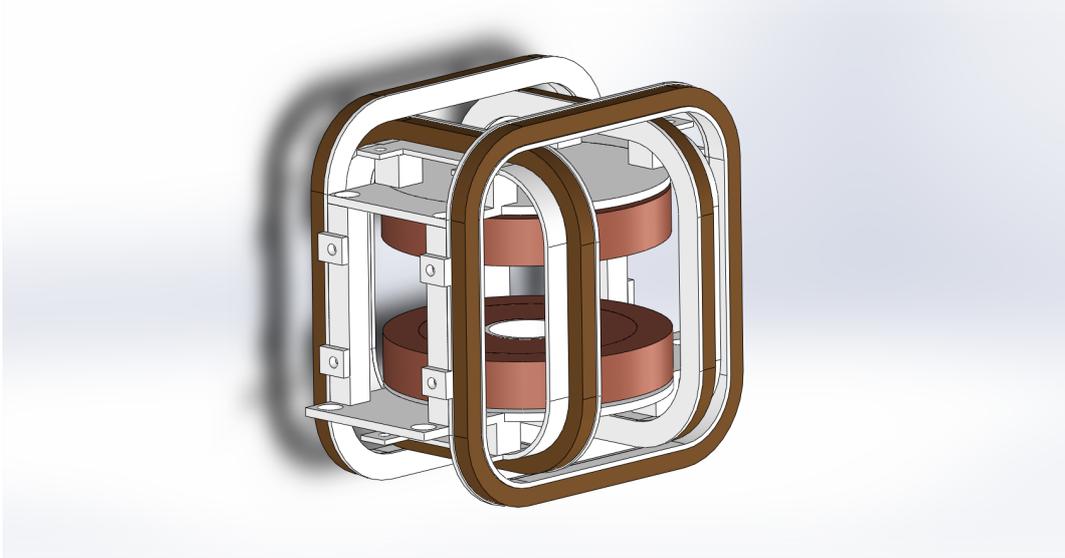


Figure 4.13: Coil formers. The interlocking coil former design, made from 3 mm PVC to prevent eddy currents persisting after the coils are switched off. The coil former design ensures the alignment and relative position of the coils with respect to one another, and allows for easy construction about the science cell.

were designed that fix them to one another. The formers were made from 3 mm PVC, ensuring enough rigidity of the former to have the coils wrapped under tension onto them and made of a non-conductive material to prevent the possibility of any eddy currents. The resultant complete coil former set is shown in figure 4.13.

4.5 Magneato coil drivers

Once the coils have been designed and made, a device that drives the desired amount of current through the coils must be procured. There are a number of 'off the shelf' solutions to this problem, power supplies that can output a current proportional to an analog line fed into the device, however, these are typically at least twice as expensive as a power supply that can simply handle the voltage and current requirements, and can generally only drive a single output. The current apparatus has 8 coils, each of which require the ability to be independently driven;

- 4 Source coils that, when independently driven, allow the position of the field zero and the magnetic field gradient of the 2D MOT to be controlled.
- The quadrupole coil pair, in which control of the coil allows the magnetic field

gradient of the magnetic trap and 3D MOT to be controlled.

- The three pairs of bias coils, which allow the position of the magnetic field zero generated by the quadrupole coil to be moved, and can be used to cancel stray magnetic fields.

To use an independent power supply for each of these coils would be prohibitively expensive, especially as most power supplies are typically not configured to drive such low resistance inductive loads with high stability. One solution to this problem is to build a coil driver for each element that needs to be independently driven. We call our design a Magneato coil driver. These coil drivers must be optimized for the resistance, inductance and current requirements of the specific coil they are driving, but this is readily achieved through changing the value of a few key resistors and capacitors.

4.5.1 Design of the Magneato

The principle behind the Magneato driver is relatively simple, the driver measures the current running through the coils and generates a voltage corresponding to this current, compares this voltage to the control voltage provided by the lab user, and then either allows more or less current through the coil depending on how the two voltages compare. A schematic of the Magneato circuit is shown in figure 4.14, with the three main sections of the circuit highlighted; the sensing circuit (green) which measures the current passing through the coil, the control circuit (red) which buffers and filters the control voltage, and the comparison circuit (blue) which compares the two signals from the control and sensing circuit and outputs a control voltage to an insulated gate bipolar transistor (IGBT) which acts as a variable resistor, allowing more or less current to flow through the coils in an attempt to balance the two inputs.

To measure the current through the coil the sensing circuit uses an LA 150-P³⁸ current transducer. This device outputs a current 2000 times smaller than the current passing through the a rectangular aperture in the device, which is then converted to a voltage through a 100 Ω sense resistor to ground³⁹. This resistor is chosen so as to produce a voltage in the range of 0-7.5 V for currents between 0-150 A for the quadrupole driver, however its value can be changed to select the range of

³⁸Manufactured by LEM USA inc.

³⁹The transconductance calculation for the Magneato can be found in appendix. C.

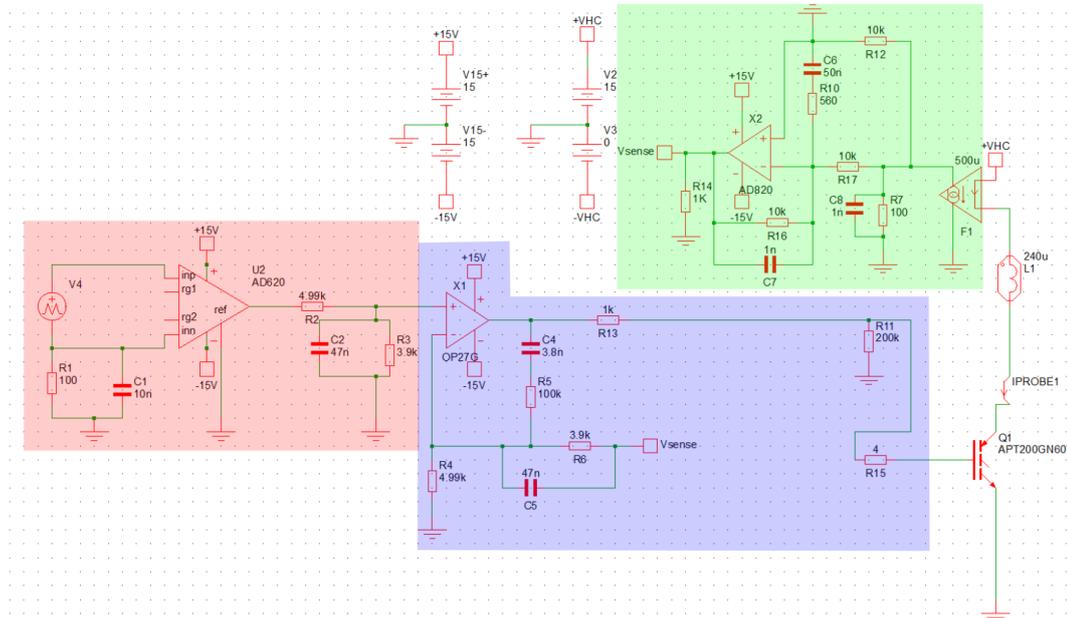


Figure 4.14: Unipolar Magneato schematic. A schematic of a unipolar Magneato circuit used to control the current through the quadrupole coils. Shaded in green in the sensing circuit, used to sense how much current is passing through the coil and convert it to a voltage. Shaded in red is the control circuit, used to buffer the input control signal, preventing feedback into the NI card. Shaded in blue is the comparison circuit used to compare the signals from the control and sense circuits, and provide a control signal to the IGBT used to modulate the current through the quadrupole coil.

currents the Magneato can measure. For example, if the Magneato is only to output currents between 0-10 A, then a 2 k Ω resistor would produce a voltage range of 0-10 V for the desired current range.⁴⁰

Typically, we would wind the coil through the current transducer multiple times to ensure the full current output range of the current transducer can be used so as to increase the stability of the driver, as this minimizes the effect of fluctuations in the output of the current transducer. In the above example, using a 100 Ω resistor and winding the coil through the transducer 15 times would produce the same result as using the 2 k Ω resistor. This voltage is then passed through an operational amplifier (op-amp), an AD820, configured as a voltage follower to act as a buffer for the signal⁴¹ which is then passed onto another op-amp, an OP27 for comparison with the control voltage. Meanwhile, the control voltage from the computer passes through an instrumentation amplifier, the AD620, which is also configured as a

⁴⁰The voltage generated at the high side of the resistor is given by $V_{Resistor} = \frac{I_{Coil}}{2000} 100$.

⁴¹This provides isolation of the measured signal from the feedback loop.

voltage follower to act as a buffer for the input signal, before also being passed onto the operational amplifier for comparison.

In the control circuit, there is a pair of resistors whose ratio sets the transconductance of the circuit by forming a voltage divider on the input signal, these are resistors R2 and R3 in the schematic. Resistor R2 also limits the current flow from the control voltage, which allows the response time of the operational amplifier to be tuned to reduce ringing in the driving circuit. Also on the control signal path is a capacitor C2, which offers a low resistance path to ground for high frequency signals, generally set to remove all frequencies greater than the fastest response time desired for the Magneato, which for our experiments is typically 1 kHz, as the catch is the fastest we need to change the current in the coils. The combination of capacitor C2 and resistor R2 form a low pass filter on the control input, which limits the maximal frequency response of the Magneato to

$$f_{max} \approx \frac{1}{2\pi R_2 C_2} \quad (4.9)$$

The comparison circuit then looks at the inputs from the sensing and control circuit and compares them with an operational amplifier. The OP27 is used as a closed loop amplifier, with a resistor, R5, used to provide a negative feedback path from the output to the negative input, which is also combined with the sensed voltage. The value of resistor R5 determines the closed loop gain of the op-amp, which is used to limit ringing in the output of the op-amp when quickly changing the output current of the Magneato, for example during the catch, however, it also limits how quickly the current in the coils will change when the op-amp is presented with a difference between the control and sense voltages, as reducing the ringing reduces the gain of the op-amp. The negative feedback path also contains a capacitor, C4, which offers a lower resistance path for higher frequency signals to pass back to ground without feeding into the IGBT. Additionally, there is a pair of resistors, R6 and R4, upon the path to the input of the negative terminal of OP27 (for the sensed voltage) which similarly act as a voltage divider, and a capacitor C5, which acts similarly to C2, these are typically set to the same values as R2, R3 and C2 respectively.

To actuate the output of the Magneato coil driver, in the case of the quadrupole driver⁴², an IGBT is placed between the coil and negative terminals of the power supply. The IGBT can be viewed as a variable resistor, with the resistance of the device determined by the voltage applied to the gate of the IGBT, thus the Magneato

⁴²In the case of the other drivers, lower power MOSFET's are suitable.

allows the desired current to pass through the coil by varying the gate voltage upon this device. An IGBT is chosen rather than a power MOSFET, as they are typically designed to handle larger currents and power dissipation, making them more robust for driving the large currents required for the magnetic trap. For coil drivers in which the actuator need not dissipate large amounts of power, or drive large amounts of current, a power MOSFET is suitable. In the comparison circuit, just before the gate voltage is applied to the gate of the IGBT, a final resistor, R13, is placed that limits the rate at which current can accumulate at the gate of the IGBT, limiting the rate at which it responds. This final resistor can also be tuned to optimize the response of ringing in the circuit.

4.5.2 IGBT power dissipation

One precaution that must be taken when using an IGBT is to ensure that the power dissipated by the IGBT remains below its maximum power dissipation at all times, otherwise the IGBT will fail, typically catching alight and shorting the circuit, which if not caught in time can lead to a smoke/water filled lab.

To calculate the power dissipated by the IGBT, the voltage across the IGBT and the current it is driving must be known, then the power dissipated via Ohmic heating is simply;

$$P_{IGBT} = V_{IGBT}I. \quad (4.10)$$

However, as the voltage drop across the IGBT is dependent upon how much voltage is dropped across the coils, this can be rewritten as

$$P_{IGBT} = (V_{Total} - IR_{Coils})I \quad (4.11)$$

$$\therefore P_{IGBT} = V_{Total}I - I^2R_{Coils}. \quad (4.12)$$

From equation. 4.12, it is plain to see that the power dissipated by the IGBT is quadratic in the current, with the peak power dissipated given by

$$P_{peak} = \frac{V_{Total}^2}{4R_{Coils}} \quad (4.13)$$

occurring when $I = \frac{V_{Total}}{2R_{Coils}}$ ⁴³. Typically, large voltages on the power supply are

⁴³This is easily found by taking the derivative of equation 4.12 and solving for the stationary point.

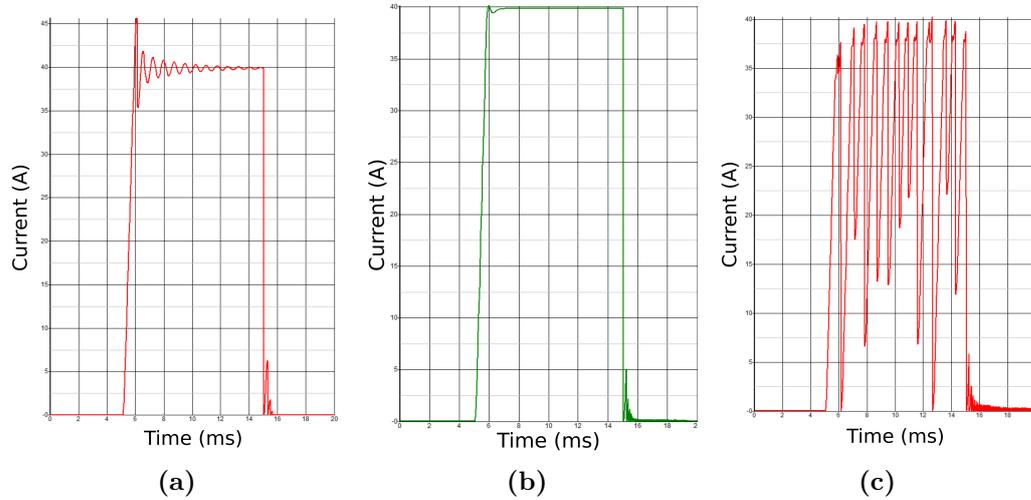


Figure 4.15: Simetrix simulations of the current switching for the catch a), b) and c) show the effect of changing the resistor, R_5 , in the Magneato driver on the resultant catch current. For values of R_5 that are too low, **a)**, some ringing is present as the system is underdamped, for values that are too large, **c)**, the driver can go into oscillation, whilst for a small range of values, **b)**, the system will be close to critically damped and will perform optimally. All simulations show some ringing when suddenly switched off, however, these oscillations are damped by a diode placed across the coil, not included in the simulation.

helpful when needing to drive large amounts of current quickly as this voltage decreases the switching time, however, care must then be taken to ensure the IGBT doesn't dissipate too much power. To avoid the IGBT dissipating large amounts of power, we actively switch between three different voltages on the power supply⁴⁴.

4.5.3 Magneato performance

To be able to tune the Magneato driver, the circuit must first be simulated. All simulations I performed of the coil driver were done in Simetrix. Tuning of the circuit's transconductance was first performed by changing the values of the ratio of resistors R_2 and R_3 until the desired current range for the coil was achieved. Simulations of the current through the coil as a function of the control voltage when varied from 0-10 Volts were used to measure this response, as shown in figure 4.16. Once this was completed the circuit was then tuned for the current switching for the catch. The control voltage that corresponded to 40 A of current through the coils

⁴⁴Actuated by two digital control lines, which power two separate relays that can switch additional resistors into a circuit between two pins on the back of the power supply, whose resistance determines the output voltage of the supply.

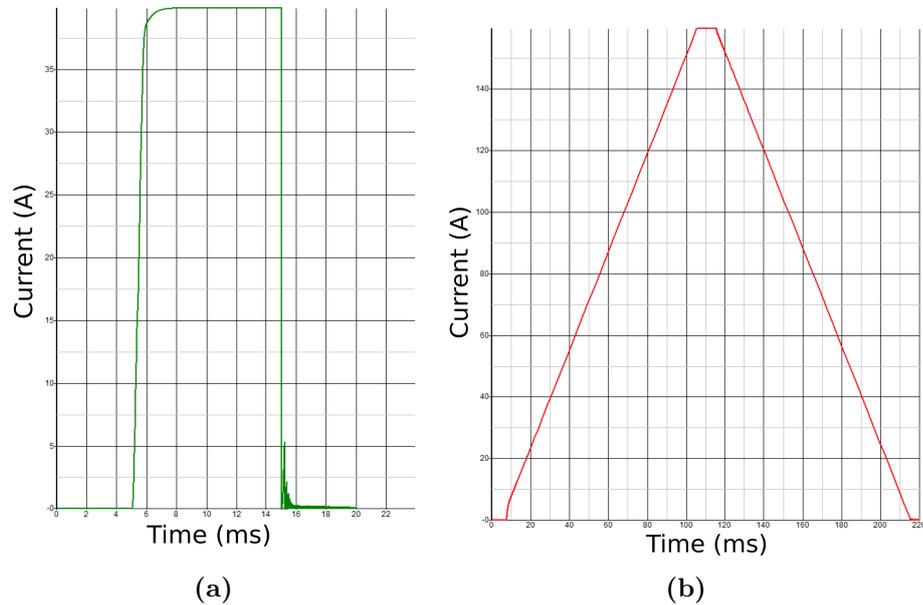


Figure 4.16: Quadrupole Magneato simulations **a)** A Simetrix simulation of the current through the quadrupole coil when suddenly switched on after 5 ms to the catch current value (40 A). **b)** A Simetrix simulation of the transconductance of the Magneato optimised for the Quadrupole coils. The control voltage is linearly ramped from 0 to 10 V over 100 ms, before being held constant for 10 ms and finally being linearly ramped back down to 0 V. The current through the quadrupole coil clearly follows the voltage linearly, however, there is a slight lag initially, this is not observed experimentally.

was found, and a 10 ms pulse with a $1\ \mu\text{s}$ rise and fall time was used to simulate the switching on and off of the coil to the desired current. Resistors R13 and R5, as well as capacitor C4 were then adjusted to reduce the ringing of the current in the coil while maintaining a rise time of $<1\ \text{ms}$. Figure 4.15 shows how increasing the value of R5, which reduces the closed loop gain of the op-amp can be used to decrease the ringing in the coils. In the case that ringing was still present on long time scales, resistor R13, which limits the response rate of the IGBT could then also be tuned, as well as capacitors C2 and C5. Figure 4.16 shows the resulting transconductance of the magneato driver, as well as the switch on characteristics for the catch.

4.5.4 Future improvements to the Magneato design

There are a number of improvements that could be made to the Magneato board layout, for both a more robust design and a more practical design. The main feature that would be useful to include in a future version, would be to replace key resistors,

which tune the response of the driver (resistors R13 and R5) with molex connectors on the board, so that a variable resistor could be mounted to the front panel of the box to allow active tuning of the driver to optimise its response, rather than relying upon the optimal configuration found through simulations.

Another improvement to the design would be to have an additional voltage dividing circuit that could provide the control voltage for the coils generated by the Magneato power supply. Whenever dynamic control of the coils are required this feature is redundant, however, for a number of coils, for example the four Source coils, once a particular optimal value has been found dynamic control is no longer necessary. The addition of the control voltage being generated by the Magneato power supply would thus allow a number of analog control lines from the control system to be freed for other purposes. In addition to the manually controlled voltage, a TTL powered relay could also be included so as to allow the coils to be switched on and off when desired through the use of a single digital line.

One further improvement would be to also include a set of high current relays, actuated by a TTL signal, into the Magneato box, that would allow for reversing the polarity of the coils. This would then negate the requirement for a bipolar driver, halving the number of MOSFETs per channel needed.

4.6 Coil cooling system interlock

One drawback of using large gauge wire for our coils is the need to actively cool the coils, as the consequences of thermal runaway or the possibility of destroying the enamel coating between the windings would be disastrous; leading to either the coils needing to be rewound, which would most likely require the entire rebuild of all optics about the chamber, or worse, damage to the vacuum system requiring a rebuild and rebake of the entire system. Another possible failure of the cooling system, with perhaps lesser consequences, would be a leak somewhere in the water cooling circuit leading to flooding of the lab and perhaps damage to electronics. A necessary measure to avoid these circumstances is the use of an interlock system that can monitor the temperature of the coils to ensure they are not overheating and measure the flow rate to ensure there is no leaking, and in either case take measures to prevent further consequences.

4.6.1 The interlock system

To build such an interlock system, three components are required;

- A sensing component, which can observe the variable of interest such as temperature or flow rate
- A decision making component, which can decide whether the variables of interest are within the expected ranges, or decide whether action need be taken
- An action component, which the system can use to prevent consequences from occurring if something does go wrong

these three components are all carried out with the use of a Galil RIO-47102 PLC⁴⁵. The Galil is a programmable logic controller with 8 analog inputs and outputs, 16 digital inputs, and 16 digital outputs, which is the brain of our interlock system. By using the analog inputs of the Galil, temperature measurements of the coils can be made, as well as water flow measurements through the water cooling system, and then by comparing the measured values to the limit values preprogrammed into the controller, the PLC can decide upon what type of error mode to go into, or whether to continue in a state of ‘all ok’.

4.6.2 Temperature measurement

To measure the temperature of the coils, a negative temperature coefficient (NTC) thermistor is used in a voltage dividing circuit. The voltage dividing circuit is formed with 5 V supplied by a power supply for the interlock, a 10 k Ω resistor to ground, and the NTC thermistor in series with the analog input of the Galil. The NTC thermistor is nominally 10 k Ω at 25 °celsius, with a β value of 3970. The β value of a thermistor allows the resistance of the thermistor to be calculated at other temperatures, through the equation^[94]:

$$\beta = \frac{\ln\left(\frac{R_{T1}}{R_{T2}}\right)}{1/T1 - 1/T2} \quad (4.14)$$

where R_{T1} and R_{T2} are the resistance of the thermistor at temperatures $T1$ and $T2$.

Using equation 4.14, and the formula for the voltage measured at the output of a

⁴⁵Galilmc

voltage dividing circuit^[95]:

$$V_{out} = V_{in} \frac{R_2}{R_1 + R_2} \quad (4.15)$$

where R_1 is the resistor in series with the voltage reading V_{out} , and R_2 is the resistor to ground, one can calculate the temperature through rearranging equation 4.15 for the resistance of the NTC, R_1 , and rearranging equation 4.14 for the temperature and then combining the equations, resulting in

$$T_2 = \frac{1}{\frac{1}{T_1} + \frac{\ln(\frac{V_{in}}{V_{out}} - 1)}{\beta}} \quad (4.16)$$

Unfortunately, the Galil does not have the capability to perform logarithm's, thus a Taylor expansion of $\ln(\frac{V_{in}}{V_{out}} - 1)$ must be used. As R_2 lies between 3500 Ω and 16000 Ω for the temperature range $[15^\circ, 50^\circ]$, the argument of the logarithm lies between $[0.3, 1.65]$, thus a Taylor expansion of $\ln(x)$ about $x=1$ is relatively accurate.

Four of the analog input channels of the Galil are dedicated to temperature monitoring. The Galil is programmed to compute the temperature using the acquired voltages, and if it is within a suitable range then the Galil powers a relay that connects two pins on the power supply for the coils, enabling the power supply. If instead the temperature falls outside of the allowed region, then the Galil stops powering the relay, breaking the contact between the two pins on the power supply and thus disengaging it, preventing further heating of the coils.

4.6.3 Water flow measurement

The flow of water through the coils is measured utilising an RS PRO radial flow turbine flow meter⁴⁶. This flow sensor measures the flow of water through a spinning turbine, a photodiode and an LED, and outputs a signal whose frequency corresponds to the flow rate of the water through the device. This frequency is then converted into a voltage through the use of a Texas instruments LM2907 Frequency to voltage converter. The Galil then compares this input voltage to predefined minimal and maximal voltages, and if it is within these limits the Galil then powers a relay which connects 18 V across a pair of Asco solenoid valves, engaging them and allowing water to flow through the coils⁴⁷.

⁴⁶part number: RS 257-149

⁴⁷Obviously when the interlock is not engaged no water flows through the flow meter, thus when engaging the interlock the solenoids are powered for a short time so as to be able to measure the

4.7 RF generation and shutter control

There are a number of other electronic devices that are required for the successful creation of a BEC, these include devices that can generate large amounts of radio-frequency electromagnetic radiation (RF) for driving acousto-optic modulators, which then control the laser frequency and amplitude, as well as to drive an RF antenna for evaporative cooling in a magnetic trap. Although acousto-optic modulators allow the control of the amplitude of light, light leakage even when the AOM is switched off is typically enough to significantly reduce the lifetime of the atoms. As such mechanical optical shutters are also required to provide slow⁴⁸, but absolute shuttering of the light. Off the shelf solutions for these are available, but typically lab made devices are cheaper and as effective.

4.7.1 RF generation

To produce the RF for our experiments we use a Novatech DDS9M signal generator module, capable of generating up to four independent RF outputs of up to 170 MHz in 0.1 Hz increments. Of these four channels, two are static channels whose values⁴⁹, must be set and unchanged throughout an experiment, while the other two channels are dynamic and can step through a range of values⁵⁰. These signals, however, are quite low power, and thus are passed through an amplifier⁵¹ to reach the required powers for the AOM's⁵². Before reaching the amplifier, however, a Mini-circuits RF switch⁵³ is used to allow TTL control of the RF signal.

The Novatech module is incorporated into a device we call a SuperNova, which integrates all components into the one box, along with thermal switches to ensure the amplifiers don't overheat and burn out, and relays to engage/disengage the amplifier power at will, disengaging if the thermal switches are triggered.

flow rate.

⁴⁸Typically ms switch on times, as opposed to μs with an AOM.

⁴⁹amplitude, phase and frequency

⁵⁰32,768 values can be stored

⁵¹Delta RF technologies LA2-1-525-30

⁵²Maximum 2 W

⁵³Mini-circuits ZX80-DR230-S+

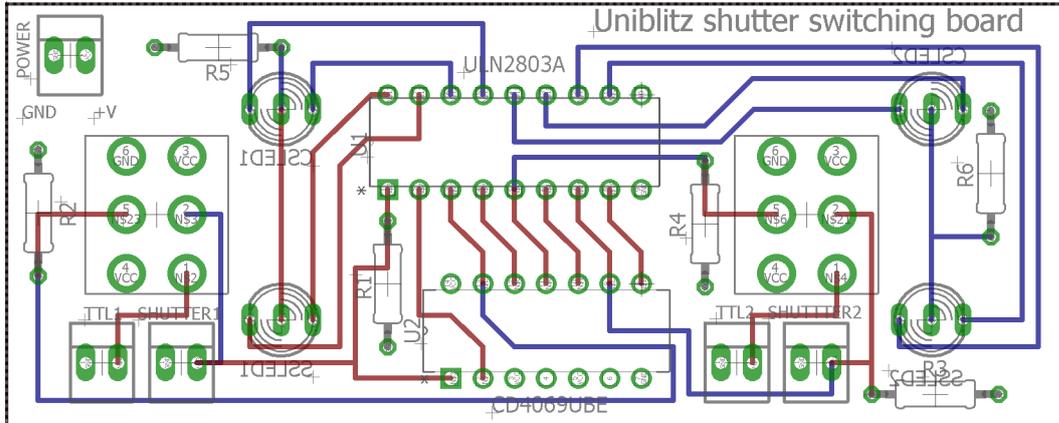


Figure 4.17: Shutter control circuit. The designed PCB circuit board used to switch between digital (TTL) control of the shutters, and manual on/off mode for diagnosis of problems. This board was designed to be mounted directly to the front panel of a 2RU box via the DPDT switches, and also switches the state of two LEDs to allow the user to know whether the shutters are under manual or TTL control, and whether they are on or off.

4.7.2 Laser shutter control

For the ability to completely extinguish laser light from entering the experiment eight laser shutters are used on the laser side of the experiment, before the optical fiber couples the light to where it is needed for the experiment. We chose to use Uniblitz TS6B laser shutters that have a 6 mm aperture, large enough for the largest laser beam on the laser side of the experiment, are bistable, thus preventing them from overheating in cases where they are held open for large amounts of time, and have a reported 3 ms closing and opening time. To drive these shutters, the Uniblitz ED12DSS bistable shutter driver is used. This shutter driver can power a single shutter, and can take a TTL signal as input to actuate the shutter.

The TTL input to these drivers was connected through the designed PCB shown in figure 4.17. This circuit allows the user to select between three modes, manual on, manual off and TTL through the position of a switch to be mounted to the front panel of a 2RU box, which also powers two bi-colour LEDs that indicate whether the shutter is being controlled manually or through the TTL signal, and whether the shutter is open or closed. The response of these shutters have been measured by placing a photodiode after the shutter and measuring the delay and closing/opening time of the shutter. Typical opening times of the shutters are less than a millisecond, whereas delays can be on the order of 3 ms, these values are taken into account in the labscript suite so that when the shutter is programmed to open or close at a

given time the control signal will be sent taking into account the delay.

Laser cooling of atoms

Having discussed the construction of the vacuum system, along with all the necessary electronics and optics required for a Bose-Einstein condensate apparatus, we now turn to the theory required for understanding how a Bose-Einstein condensate can be produced. In Appendix A I cover some of the basic theory behind Bose-Einstein condensates, however, I do not discuss how the state physically comes to be realised. The route to Bose-Einstein condensation can typically be broken into two separate stages, laser cooling and evaporative cooling. During the laser cooling stage large numbers of atoms are caught from a background thermal gas of atoms at 300 K, and cooled to a few mK, with temperatures of a few tens of μK achieved with sub-Doppler cooling mechanisms. The next stage, evaporative cooling, requires the atoms to be caught in a purely magnetic trap¹, where their temperature and lifetime are not limited by the resonant scattering of photons, and they then undergo forced evaporative cooling, where atom number is traded for a reduction in temperature by the selective removal of atoms with energies greater than the average thermal energy of the gas. In this chapter we shall discuss the theory behind the first of these two stages, the laser cooling and confining of atoms.

At the turn of the 19th century, the mechanical effects that light had upon matter were first starting to be demonstrated in a laboratory setting by Lebedev[96] in 1901 and Nichols and Hull[97] in 1903. However, these effects were relatively small, as the effect any individual photon had upon a massive object was greatly hindered by the small momentum a single photon carries, and as such intense sources of light were required to produce large effects upon matter. With the invention of the laser in the

¹Or alternatively an optical dipole trap.

60's though, access to a highly monochromatic intense source of light was possible, and the ability to trap and manipulate particles with laser light was put to the test by Arthur Ashkin in 1969[10]².

The idea of using a laser beam tuned to an atomic resonance to cool atoms was introduced approximately 5 years later, in 1975 by Hansch and Schalow[98], and independently by Wineland and Dehmelt[99]. The first laser cooling experiments were performed only three years after by Wineland, Drullinger and Walls[100], in which a cloud of magnesium ions were cooled in Penning trap. It is with these observations that the beginning of a new field of physics, the laser cooling of atoms, began.

5.1 Optical forces and molasses

To be able to discuss the mechanism by which we can cool an atomic sample, we must first understand the forces that arise from the interaction of a photon with an atom. There are two forces typically considered in laser cooling of atoms when discussing the interaction of the atoms with a photon; the scattering force and the dipole force, however, the dipole force plays only a small role in laser cooling³, and as such will not be discussed here.

The scattering force arises from the absorption of a photon by an atom, requiring the light to be tuned to an atomic resonance, before the photon is spontaneously emitted by the atom in a random direction. As the photon carries momentum⁴ and the photon changes direction through this interaction, the photon necessarily imparts some momentum to the atom that scattered the photon. The momentum change of the atom from the interaction, divided by the duration of the interaction provides an estimate of the force upon the atom known as the scattering force, F_s . To consider the scale of this force, a rubidium 87 atom has a mass of 1.4×10^{-25} kg and absorbs light at 780 nm for the D2 cooling transition. Given the lifetime of the atom is 26 ns in this excited state, the atoms can thus absorb $\approx 1.9 \times 10^7$ photons per second⁵. Each photon provides a momentum kick to the atom of $\hbar k = 8.5 \times 10^{-28}$ kg m s⁻¹

²Whose work was worthy of a Nobel prize

³However, it plays an important role in forced evaporation in an optical dipole trap.

⁴The momentum of a photon is given by $p = \frac{h}{\lambda}$, where p is the momentum, h is the Planck constant and λ the wavelength of the photon.

⁵This assumes a photon scattered every 52 ns, if one tries to scatter photons faster then stimulated emission starts to compete with spontaneous emission.

and thus the total acceleration applied to the atom is $\approx 10^5 \text{ m s}^{-2}$.

The scattering force described above, however, does not cool atoms in itself, as it needs to be coupled with some anisotropy in the system that makes it more likely for the atoms to absorb a photon from a laser whose k-vector opposes the direction of the atomic motion, rather than whose k-vector is aligned⁶. Hansch and Schawlow[98] realised that the Doppler shift that arises due to the atoms own velocity, coupled with detuning the incident light field below resonance could provide the required anisotropy in the system⁷. An atom moving towards a laser beam will act to blue-shift the laser light with respect to the atom, i.e. the atom will see a light field that is oscillating at a greater frequency than in the lab frame. By choosing the right detuning below resonance of a two level atomic transition the optical field can then be made to be more likely absorbed by atoms whose velocity vector opposes the k-vector of the optical field. This effect has been demonstrated to significantly reduce the velocity of atoms, slowing them to near zero.

The realisation of Hansch and Schawlow that allowed the laser cooling of atoms in counterpropagating laser beams was seen to rapidly cool the atoms over a few milliseconds to a fraction of a degree Celsius above absolute zero, nearly halting the atoms in their path as they move such a small distance over this time. It was later realised that this configuration also realises some level of confinement as the atoms start to exhibit diffusive motion, demonstrating random walks within the beams, and only slowly drifting out of the confines of the beams. Due to the ‘viscous’ nature of the force this resulting geometry became known as ‘optical molasses’.

To begin, we will consider the force an atom feels from a single beam in this arrangement. We will assume a laser of intensity I and frequency ω , and an optical transition of saturation intensity I_{sat} where the saturation intensity is defined as $I_{sat} = \frac{\hbar\omega_{atom}\Gamma}{2\sigma_0}$, where Γ is the natural linewidth of the transition, ω_{atom} the resonance frequency and σ_0 is the on resonance scattering cross section.⁸ Next, we assume that the splitting between the ground and excited state energies of our two level atom is ω_0 , and then we can define the difference in frequency between the optical field and the atomic transition as $\delta = \omega - \omega_0$. As our atoms are moving

⁶This is akin to ensuring that the scattering force always opposes the atom’s velocity making it a damping force, since if the k-vector is aligned with the atom’s motion the atom will speed up rather than slow down, whereas if it is opposed it will slow down and thus ‘cool’ the atom.

⁷Note that this anisotropy is in momentum space, not position space.

⁸The saturation intensity is the light intensity at which the on absorption cross section will reduce by a factor of 2 from its weakly pumped value, thus the amount of light the atom will scatter as a fraction of the incident light is also reduced by a factor of two.

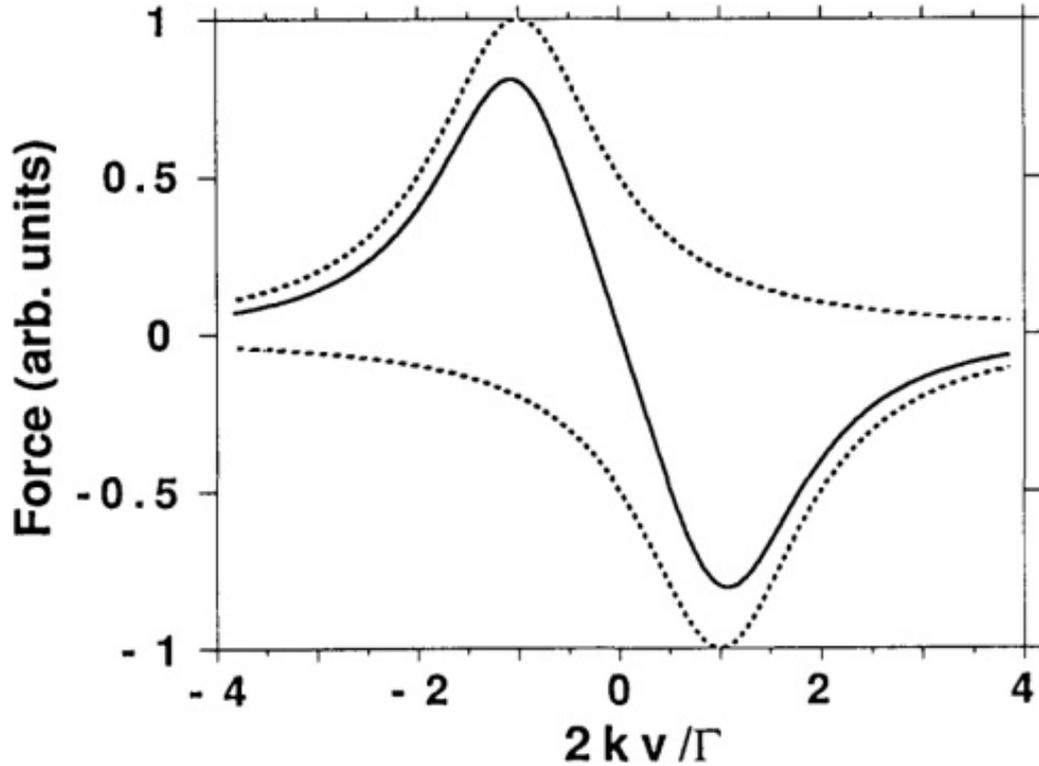


Figure 5.1: Force in optical molasses The force upon an atom due to a pair of low intensity counter propagating plane waves detuned half a linewidth below resonance. The dashed lines indicate the force arising from the two plane waves independently, whilst the solid line indicates the resultant force. There is a clear linear region near $v = 0$, which is utilised for laser cooling. Reproduced from ref. [5].

through this molasses, we need to consider the effect of the Doppler shift the atom experiences. An atom moving with a velocity, v , towards one of the lasers will see the laser frequency Doppler shifted down by kv , where k is the magnitude of the wavevector of the optical field. The combined Doppler shift of the light and detuning of the optical field provides a total detuning of $\delta \pm kv$ that the atoms see, where the plus sign is taken if the k vector of the optical field points in the opposite direction to the velocity of the atom, and the minus sign is taken if they are aligned. Finally, the rate at which the atom can scatter photons from this light field will be determined by how long the atom stays in the excited state, which is given by the inverse of the linewidth, Γ , and thus the excited state population will radiatively decay to the ground state at a rate of Γ .

The average force that the atom will feel due to the light field is the product of the momentum that a single photon imparts to the atoms, $\hbar k$, the fraction of incident

photons that the atom scatters, and the number of incident photons, yielding^[5]

$$F_{scat} = \mp \hbar k \frac{\Gamma}{2} \frac{I/I_{sat}}{1 + I/I_{sat} + (2(\delta \pm kv)/\Gamma)^2} \quad (5.1)$$

where in the case that the atoms velocity is antialigned with the k-vector of the light field we take the upper signs, resulting in a force that opposes the atoms motion, whereas if the velocity is aligned with the k vector we take the lower sign and the force is in the direction of the atomic motion. From equation 5.1 we can see that in the limit of $I \gg I_{sat}$ the force approaches its maximal value of $\frac{\hbar k \Gamma}{2}$ ⁹.

Note that this is the force that an atom travelling in the presence of a single optical field will feel, to extend this we now need to consider the effect of a second counterpropagating beam. To do this, in the simplest case we can assume the forces are independent and thus act independently on the atoms and simply take the sum of the positive and negative versions of equation 5.1 as they represent the force felt by an atom from a plane wave in the positive and negative direction. If we then assume that the Doppler shift of the optical field is much smaller than both the linewidth of the transition and the detuning of the field from resonance, the force upon the atom can be written^[5]

$$F_{scat} = 4\hbar k \frac{I}{I_{sat}} \frac{kv(2\delta/\Gamma)}{(1 + (2\delta/\Gamma)^2)^2} \quad (5.2)$$

importantly, in the case where the detuning is red, i.e. $\delta < 0$, the sign of this force is always negative and opposes the velocity of the atom, thus we have a damping force upon the atoms. This resultant force, as well as the force arising from the two plane waves independently, is shown in figure 5.1.

5.1.1 Limitations of optical molasses

Two issues quickly arise with optical molasses, however, to elucidate these issues we will consider the scattering force upon a Rubidium-87 atom.

The first of these issues arises from Rubidium-87 not being a true two level atomic system, rather there are two ground hyperfine states and four excited hyperfine states, as described in section. 4.3. Typically when laser cooling Rubidium-87, the $F = 2$ to $F' = 3$ cycling transition is used to cool the atoms. In a small proportion

⁹As the term on the left hand side approaches unity for $I \gg I_{sat}$.

of cases, when the atom is in the excited state, it may spontaneously emit a photon and fall down into the $F = 1$ ground state, which is separated from the $F = 2$ ground state by 6.8 GHz. Due to the narrow linewidth of the lasers used, typically ≤ 10 MHz, and the transition linewidth, 6 MHz[101], this means that the atoms will no longer be able to absorb the laser light, and thus the atom will no longer be cooled. The standard method of correcting this loss of atoms from the $F = 2$ ground state is to add an additional laser, known as a repumping laser, to optically pump the atoms from the $F = 1$ ground state back to the $F = 2$ excited state, where it is likely to fall back into the $F = 2$ ground state and continue being cooled on that cycling transition.

The second issue that arises from the laser cooling of atoms also arises from the narrow linewidth of the laser and atomic transition, as well as the Doppler shift of the light an atom experiences due to its motion. As already discussed, an atom moving with a velocity, v , towards a photon will see the frequency of the photon shifted by an amount $\frac{2\pi v}{\lambda}$. As we are cooling the atoms, the velocity of the atoms necessarily change with time, implying the Doppler shift of the photons seen by the atoms change with time, eventually shifting the photons out of resonance with the atoms cycling transition. For a linewidth of 6 MHz, after only slowing an atom initially moving at 300 m/s, to 292 m/s, the change in the experienced Doppler shift will already be greater than the natural linewidth of the transition, meaning the atom will no longer be slowed by the laser light. This causes atoms of only a particular velocity, for the given laser detuning from the atomic resonance, to be slowed and only by a small amount.

There are two solutions to this problem, one in which the cooling mechanism is temporally changed, i.e. the detuning of the laser is changed with time to continue cooling the atoms, known as ‘chirp cooling’, while the other utilises a spatial change in the cooling mechanism, known as ‘Zeeman cooling’. We shall focus on the latter of these two methods as that is what is employed in a MOT. Success has also been shown with the former of these methods¹⁰, however, this results in bunches of cooled atoms rather than a continuous beam.

Zeeman cooling is known as such as it utilises the Zeeman effect, the change in energy between the ground and excited state of an atoms due to the interaction between its magnetic moment and an external magnetic field, and is used to shift the energy levels of the atoms, keeping them in resonance as they are cooled. Initial

¹⁰First theoretically proposed in 1976 by Letokhov, Minogin and Pavlik[102] being clearly demonstrated by Phillips and Prodan in 1983[103] and Ertmer, Blatt and Hall in 1984[104]

work on this method of cooling an atomic beam resulted in a Zeeman slower, a long tube along which a thermal atomic beam counter propagates against a cooling laser beam in the presence of a spatially varying magnetic bias field that keeps the cooled atomic beam in resonance with the laser. This was first demonstrated in 1982 by W.D. Phillips and H. Metcalf[105], which managed to efficiently cool atoms to ≈ 40 m/s. The first design which both simultaneously cooled and trapped atoms from a background vapor was the magneto optical trap developed in 1987 by Raab et al.[106].

5.1.2 Magneto-optical traps

The creation of the first MOT was a major step towards producing BECs, as it allowed a large number of atoms to be cooled and trapped before being transferred into a purely magnetic trap in which forced evaporation could be performed to further cool the cloud. A MOT typically utilises a weak quadrupole field generated by a pair of coaxial coils in an anti-Helmholtz configuration. This magnetic field has a field zero at the geometric centre of the pair of coils, where the contribution from each coil cancels out, and the magnitude of the magnetic field increases in magnitude approximately linearly for small displacements from the field zero. It is this magnetic field that produces a spatially dependent Zeeman shift which in turn generates a spatially varying force, providing the trapping in the MOT. The optical component of a MOT is formed by pairs of counter-propagating laser beams intersecting at the field zero region of the magnetic field, detuned below the cycling transition upon which cooling will occur. The number of pairs of beams determine the number of axes along which the atomic gas is to be cooled. The polarisation of each laser beam of any given pair is typically of orthogonal circular polarisation with respect to a space fixed axis, ensuring that the laser beams will only excite magnetic substates that are such that the magnetisation of the substate is aligned with the magnetic field on the side of the field zero the laser beam is incident from. This resonance condition is most readily understood considering a one dimensional example, however, it is easily extended to two and three dimensions.

Consider an atom moving with velocity v_x , initially starting at position $-x_0$, in the presence of a laser detuned by $\Delta\omega$ below resonance at zero velocity. Now, as the atom is moving at velocity v_x , the atom sees the frequency of the laser light as shifted by only $\Delta\omega - \frac{2\pi v_x}{\lambda}$ below resonance, due to the Doppler shift of the light induced by the velocity of the atom. Finally, as the atom passes the position x' such that

the magnetic field, $B(x) = B'x$, where B' is the magnetic field gradient, satisfies $m_f \mu_B g_j B(x) = \hbar(\Delta\omega - \frac{2\pi v_x}{\lambda})$, where m_f is the magnetic moment of the state, μ_B is the Bohr magneton and g_j is the g factor of the atom, the atom will become resonant with the laser light and absorb a photon. It is this position dependent absorption that generates the trapping nature of the MOT, with hotter atoms meeting the modified resonance condition closest to the field zero.

With the development of the MOT physicists' now had the ability to both confine and cool atoms to millikelvin temperatures with a relatively simple, compact setup, creating a momentous step towards achieving Bose Einstein condensation, with MOTs becoming the standard method for trapping and confining atoms for cold atom experiments.

5.1.3 Experimentally relevant quantities of a MOT

Having discussed the basic theory of a MOT, we can now discuss some of the experimentally relevant quantities of a MOT. For a MOT that is being used as a source of cold atoms to begin the path towards Bose-Einstein condensation, the most experimentally relevant quantities are;

- The capture velocity of the MOT, v_c
- The lifetime of atoms in the MOT, τ
- The total number of atoms trapped in the MOT, N
- The temperature of atoms in the trap, T
- The density of atoms in the MOT, n

These parameters are critical to a BEC experiment, as the capture velocity and lifetime of a MOT influence the number of atoms that can be loaded into the MOT, while the number, temperature and density of the atoms all determine the initial phase space density of the MOT, the sole parameter to be considered when producing a condensate.

MOT number

When loading a MOT, the number of atoms the MOT can hold in its steady state will be determined by a balance between the rate at which the MOT loads, and

the rate at which the MOT loses atoms. The main loss mechanism from a MOT is typically due to background collisions with hot atoms that have not been captured by the MOT, which present a loss rate that is proportional to the number of atoms in the MOT, defining a characteristic time scale, τ , over which the MOT will load. At high densities, however, a second loss mechanism comes into play, where two body collisions between atoms in the trap can act as a loss mechanism. To load the largest number of atoms into a MOT, it is therefore desirable to have the largest load rate possible, which in turn depends greatly on the capture velocity.

It has been shown that the total number of atoms loaded into a MOT for a given background pressure is given by [107][108]¹¹

$$N = R\tau \approx \frac{1}{5} \left(\frac{v_c}{\bar{v}} \right)^4 \frac{A}{\sigma} \quad (5.3)$$

where R is the load rate of the MOT, \bar{v} is the average speed of an atom in the background vapor, A is the surface area of the trapping region, τ is the lifetime of the atom in the MOT and σ is the cross section of collision for an atom in the MOT with a background vapor atom. It can be seen from equation 5.3 that the equilibrium number of atoms in the MOT is independent of the background vapour density¹², a simple explanation for this is that larger background vapor pressures lead to faster load rates, but also higher loss rates.

Equation 5.3 indicates that the total number of atoms loaded into the MOT will rise rapidly with the capture velocity of the MOT, thus optimising this parameter is critical for obtaining large numbers of atoms in a MOT quickly. The capture velocity of a MOT can be estimated as the largest velocity that will be stopped by half the maximum scattering force [107], over a distance equal to the radius of the trapping beams, this yields

$$v_c = \sqrt{2ar} \quad (5.4)$$

where r is the radius of one of the trapping beams, and

$$a = \frac{h\Gamma}{4m\lambda} \quad (5.5)$$

where Γ is the natural linewidth of the transition, and m the mass of the atom.

¹¹This is shown by considering the equilibrium state achieved when the load and loss rate are balanced, and assumes loading from a background gas.

¹²This assumes that the vapor pressure is not so low that a different atomic species dominates the background gas.

Given the dependence of the critical velocity upon the radius of the trapping beams, it can then be seen that the total number of atoms loaded into the MOT is proportional to area of the trapping beams¹³. The above holds true when the intensity of the laser beams is above or equal to the saturation intensity of the atoms, thus maximising the beam size for given amount of power available, ensuring it remains above the saturation intensity, is critical for large atom numbers in a MOT.

It should be noted, however, that in most typical MOT setups used for achieving Bose–Einstein condensation utilize a two chamber apparatus, where the loading of the MOT is typically not from a background vapor, but rather from another MOT or Zeeman slower. This allows large loading rates from a high background pressure for the first MOT (or Zeeman slower), which can then be efficiently transferred into the second MOT, without compromising the quality of the vacuum at the second MOT. In these schemes it is typically desirable to have the fastest load rate for the first MOT possible as this maximises the load rate of the second MOT, assuming the efficiency of the transfer is independent of the rate at which the atoms are loaded, rather than maximising the number of atoms that can be loaded into the first MOT. This is why greater vapor pressures in the first chamber are generally better as they lead to faster load rates of the second MOT, even though the size of the MOT is independent of the pressure¹⁴. However, this leads to a caveat that some forget, as the source of atoms for the second MOT is already significantly pre-cooled, the capture velocity of the MOT need not be anywhere near as large to efficiently transfer the atoms from the cold beam into the MOT. The limitation in this case typically comes from the divergence of the source, which can lead to a relatively large atomic beam by the time it reaches the MOT, and for efficient transfer the size of the MOT beams must at least be large enough to interact with this atomic beam.

MOT temperature

The next question that naturally arises then, is what is the temperature of the atoms in a MOT? It is not obvious whether or not this question is a reasonable one to ask, as when defining the temperature of an object we generally consider the system under consideration, the MOT, to be in thermal contact with a reservoir¹⁵. In the case of a MOT we have a system that is thermally isolated from its environment, and thus what acts as the reservoir is unclear. However, Reif shows[109] that for a system of

¹³As $v_c \propto \sqrt{r}$ and $N \propto v_c^4$, therefore $N \propto r^2$.

¹⁴However, the load rate is not.

¹⁵Some theories of laser cooling consider the ‘bath’ of photons to be a reservoir.

particles with a frictional damping force proportional to velocity, as is the case with optical molasses, and with random noise independent of velocity, arising from the re-emission of photons in random directions, the solution is a Maxwell-Boltzmann distribution, hence assigning a temperature to the system seems reasonable.

As a MOT traps and cools atoms to $\langle v \rangle \approx 0$, what causes the heating of these atoms that maintains them at a finite temperature? When considering the forces in a MOT, we typically consider the average force on the atoms summed over many thousands of absorption and emission cycles. Assuming that the spontaneous emission from an atom is isotropic, these emissions will not contribute to the average force, however, they will cause fluctuations upon the force that results in a finite temperature. To elucidate this, consider the cooling in a MOT as akin to a drunkards walk, with a friend attempting to guide him home (to zero velocity) from the pub (his initial velocity).

The drunkards walk is used to model many processes in physics, a notable example being A. Einstein's modelling of Brownian motion with a random walk[110], and the experimental verification leading to further proof that atoms and molecules exist. In the case of Brownian motion that Einstein considered however, there is no friend to help guide the drunkard, rather he just stumbles randomly about¹⁶. In our example, think of the friend as the MOT beams, the friend initially helps the drunkard take a step in the right direction (reducing his velocity), but the drunkard stumbles after another step in a random direction (equally likely increasing or decreasing his velocity) before the friend can help him again. Even when the friend has guided the drunkard all the way home and said goodnight, the drunkard can still stumble about in his house and it is how far the drunkard is stumbling about in his house that is akin to how the atoms in a MOT maintain a finite temperature. Alternatively, we can say that it is the fluctuations of the frictional damping force due to the spontaneous emission of photons from the atoms that creates a heating mechanism for atoms in the MOT, and the equilibrium temperature will be determined by the balancing of the cooling and heating rate.

Figure 5.2 graphically represents the cooling process in a MOT. The left hand side of the figure represents what occurs in a single absorption/emission cycle, with the atom initially absorbing a photon from the MOT laser (red headed arrow) before undergoing spontaneous emission and re-emitting the photon (blue headed arrow)

¹⁶The process that more closely resembles our MOT is Brownian motion in a harmonic potential[111], but we will simply consider Brownian motion here as it is only for illustrative purposes.

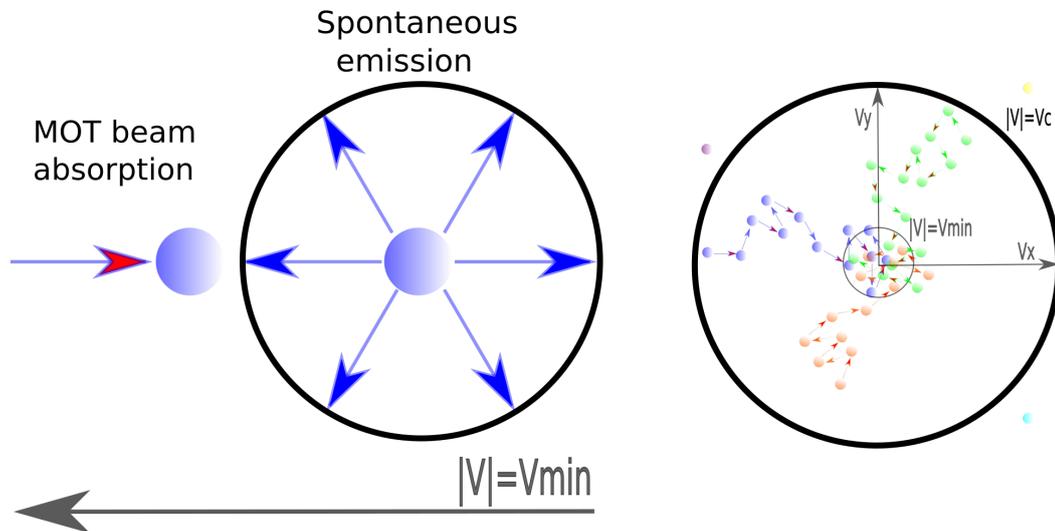


Figure 5.2: Representation of the drunkards walk A representation of the drunkards walk and how it can lead to cooling in a MOT. On the left hand side we see the effect of a single absorption/spontaneous emission event. The atom initially absorbs a photon from the MOT laser (red headed arrow) which reduces the atoms velocity, and then undergoes spontaneous emission and re-emits the photon in a random direction (blue headed arrow). The right hand side shows how undergoing this process multiple times reduces the atoms initial velocity to some finite velocity. The two concentric circles are representative of the capture velocity (larger circle), v_c , determined by the MOT parameters, and the minimum velocity (small circle) of the atom, whilst the coloured circles represent atoms undergoing multiple absorption re-emission events as they are cooled. Note that once the atoms are within the smaller circle they can still undergo absorption re-emission cycles, these determine the minimum temperature of the MOT.

in a random direction. For atoms with a velocity greater than the recoil velocity¹⁷, the absorption from the MOT beam always reduces the atoms velocity and thus the net effect of both the absorption and emission is at worst for the atoms to remain at the same velocity, or at best to reduce its momentum by $2\hbar k$, as shown by the range of positions the blue-headed arrow can result in.

The right hand side of figure 5.2 shows how an atom undergoing many absorption and re-emission cycles can be cooled from some finite initial velocity to the minimum temperature achievable in the MOT. Shown in the figure is a number of different random walk trajectories that the atoms can follow, represented by different coloured circles. The larger of the two concentric circles is representative of the capture velocity of the MOT, v_c , whilst the smaller circle is representative of the minimum

¹⁷When the velocity of the atom is below the recoil velocity, the absorption of a photon from a MOT beam can lead to an increase in the atoms velocity, this is discussed in more detail later.

velocity an atom in the MOT will attain. An atom initially within the capture velocity of the MOT will undergo many absorption/re-emission cycles, with each absorption kicking the atom towards $|v^2| = 0$, whilst the emission kicks the atom in a random direction. After many of these cycles the atom finally approaches $|v| = 0$ but it will still undergo absorption/emission cycles, and these momentum kicks lead to the finite temperature of the MOT.

As the frictional forces in a MOT arise from the optical molasses caused by the counter propagating laser beams we shall assume the temperature in a MOT is comparable to that of the optical molasses that would result with no magnetic field. The temperature of optical molasses, when balancing the heating rate with the cooling rate of the molasses has been calculated by Lett *et al.*[5] to be

$$k_B T = \frac{\hbar\Gamma}{12} \frac{1 + 6I/I_0 + (2\Delta/\Gamma)^2}{2|\Delta|/\Gamma} \quad (5.6)$$

where T is the temperature of the atoms in the MOT, Γ the natural linewidth of the transition and Δ the detuning below resonance of the lasers, and we have assumed a MOT in three dimensions. From equation 5.6 it is clear that as the intensity increases, so too does the temperature, thus to calculate the lowest temperature achievable in a MOT, let $I/I_0 \ll 1$. By then taking the derivative with respect to the detuning, and solving the equation when the derivative is set to zero¹⁸, we have the Doppler temperature limit for a MOT, which is $\approx 145\mu K$ for ^{87}Rb when cooling on the D_2 transition.

One additional interesting property worth noting here is the dependence of the temperature of the MOT on the atom number. In a simple description of a MOT the atoms do not interact with one another and the heating rate of the MOT solely determines the temperature of the MOT. However, this simple picture doesn't take into account the possibility of photons that have initially been absorbed and then re-emitted by one atom then being absorbed by another atom. Cooper *et al.* describe this effect in reference [112] and measure the resultant effect of the number of atoms in the MOT on the temperature of the atoms, finding that it scales with the third root of the number, with MOTs on the order 5×10^8 atoms being at least twice as hot as the theoretical limit for the MOT temperature.

¹⁸This occurs when $\frac{2\Delta}{\Gamma} = -1$, indicating the lowest achievable temperatures occur when the detuning is on the order of the linewidth of the transition.

MOT density

Typically when loading a MOT, the size of the MOT is initially determined by the temperature of the atoms, which according to the equipartition theorem will force the atoms to spread to an rms radius, and have a rms velocity given by[113][107]

$$k_B T = \kappa_i \langle r_i \rangle^2 = m \langle v \rangle^2 \quad (5.7)$$

where κ_i is the spring constant and $\langle r_i \rangle$ is the root mean square radius in the i^{th} direction, and $\langle v \rangle$ is the root mean square velocity. If the atoms in the MOT are assumed to behave as simple damped harmonic oscillators, the solution of the Fokker-Planck equation suggests this distribution should be Gaussian. For the Doppler limit calculated above, along with the spring constant of a MOT, which is typically on the order of 100's of kHz[113], this would indicate the size of the MOT would be on the order of a fraction of a millimetre, which generally agrees for MOTs smaller than $\approx \times 10^5$ atoms.

When the MOT is initially loading, only the density of the MOT increases, with the shape and size of the cloud remaining constant, as more atoms get loaded into the trap at the same temperature. However, once the MOT becomes dense enough an additional force comes into play due to the scattering of light off the atoms, which produces a repulsive force between the atoms.[107] This force arises when the likelihood of a photon to be absorbed by one atom that has been emitted by another atom becomes large. The photon that is transmitted between the two particles then acts as a force mediator, much like the photon is the mediator of the electromagnetic force, however, in this case it is not the emission and absorption of a virtual particle, but rather a real particle. The limiting density in the MOT will thus be determined by the interplay between the repulsive forces between the atoms and the spring constant of the trap.

By a somewhat similar mechanism, there is an additional effect that acts to compress the MOT further. As atoms at the edge of the cloud can scatter the laser light before it reaches the centre, they cast a shadow near the centre of the MOT. This creates an intensity imbalance across the cloud that should act to further compress the cloud. This effect has been shown[114] to be weaker than the repulsive force arising from the scattering of light between the atoms in a typical MOT, however, thus the density achievable is limited by the repulsive effect.

The effect of the repulsive force is to limit the maximal achievable density, and

produce a cloud that is shaped such that the repulsive scattering force between the atoms perfectly balances the trapping force. Steane *et al.* and others[107][115] have calculated this density distribution for optical molasses and found that it results in near uniform density given by

$$n = \frac{kc}{I\sigma_L^2(\sigma_R/\sigma_L - 1)} \quad (5.8)$$

where I is the intensity of the trapping light, k the spring constant of the trap, c the speed of light, σ_L the cross section for absorbing a photon from the MOT beam and σ_R the cross section for an atom to absorb a re-radiated photon. From equation 5.8 it is clear that the density decreases with increasing laser power, as the repulsive interactions between the atoms becomes stronger, however, this can be somewhat mitigated through detuning the light further, which results in a lower scattering rate of light, leading to a reduced repulsive interaction. In reality, measurements of atoms in a MOT in the multiple scattering regime show that rather than attaining a uniform density, the MOT still shows a somewhat Gaussian distribution, however, the density does get limited to a particular value[116] and the width of the Gaussian increases after this value is reached. The signature of being in the multiple scattering regime is thus not determined by the uniformity of the distribution of atoms, but rather a limited density with increasing atom number, or alternatively the mean radius of the atoms in the MOT is larger than the temperature limited mean radius.

From equation 5.3 it is clear that for a large MOT loaded from a vapor, large MOT beams are a necessity as the load rate is proportional to the area of the beams, as long as the beams remain above saturation intensity. We also saw that as a MOT loads it initially attains a Gaussian profile, the size of which is determined by the temperature of the MOT through equation 5.6, and as it grows the density simply increases while the cloud size does not. As the density of the MOT starts to become appreciable, however, the MOT's shape starts to deform, becoming approximately uniformly dense as the repulsive force between the atoms starts to balance the trapping force provided by the MOT, with the ultimate density given by equation 5.8, rather than being determined by the temperature. It is this regime that is to be reached, as typically a few by 10^9 atoms in a MOT gives a reasonable number of atoms to produce decent sized BECs within 10s of seconds.

5.1.4 Optimal parameters for two and three dimensional MOTs

The above discussion of MOTs will now be applied to two different scenarios; the use of a two-dimensional MOT as a source of cold atoms collected from a background vapor, and a three dimensional MOT used to capture atoms from the cold atomic beam generated by the 2D MOT, and to transfer into a purely magnetic trap at high phase space density. The requirements of the two cases differ, as in the 3D MOT high atom numbers, high densities and low temperatures are all desirable, whereas in the 2D MOTs case the most significant factor is the load rate¹⁹, as the 2D MOT is simply being used as a source of cold atoms, not a trap for cold atoms.

2D MOTs

The first two-dimensional MOT as a source of cold atoms was demonstrated by Dieckmann *et al.* in 1998[90], who measured a flux of almost 10^{10} atoms per second, with a mean velocity of 8 ms^{-1} and a beam divergence of 43 mrad. A standard two-dimensional MOT consists of two pairs of rectangular ‘racetrack’ coils in anti-Helmholtz configuration generating a cylindrical quadrupole field with a line of field zero along the symmetry axis. Pairs of MOT beams are then incident in two dimensions on the line of field zero, perpendicular to each other with $\sigma^+ - \sigma^-$ polarisation. The beam of atoms is then coincident with the line of field zero. The above describes what is typically known as a pure two-dimensional MOT, however, a common adaptation is to add a pair of counter propagating cooling beams along the line of field zero to help provide Doppler cooling along the longitudinal axis allowing a greater flux to be achieved as atoms with greater longitudinal velocities get cooled as well, known as a 2D⁺ MOT.

A schematic of the configuration required to create a 2D MOT is shown in figure 5.3, adjusted from reference [6]. Shown is a long rectangular cuvette, optimal²⁰ for a 2D MOT due to its length along the longitudinal dimension of the MOT, surrounded by four racetrack coils that are carrying current such that the current in wires adjacent to one another flows in the same direction. The MOT beams are also shown, with the required handedness of the polarisation indicated, however, what is not included in this schematic is how the MOT beams are generated. I chose to retroreflect the MOT beams, recycling the power in the MOT and simplifying the optics required. At the

¹⁹The temperature of the atoms is also important, however, this temperature need simply be low enough that the 3D MOT can recapture the atoms.

²⁰Compared to a square cuvette.

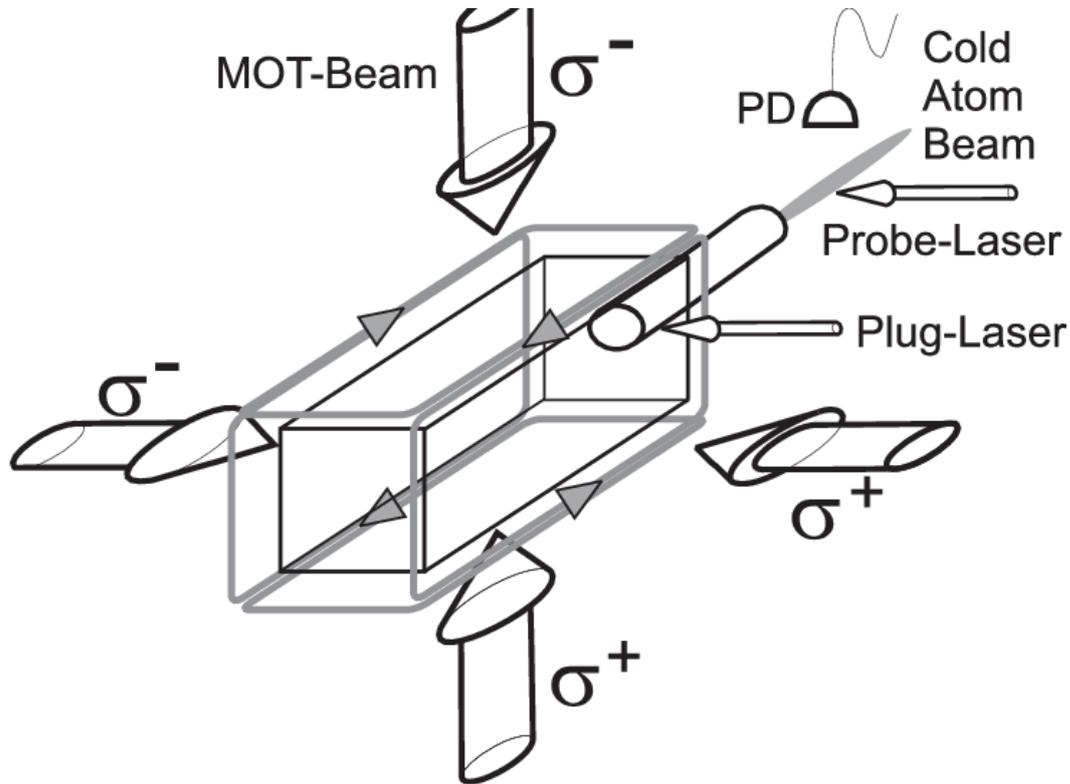


Figure 5.3: Schematic of a pure 2D MOT A schematic of a pure 2D MOT including the rectangular cuvette, the four racetrack MOT coils required to produce the line of field zero in the centre of the chamber, the MOT beams with their respective polarisations as well as the differential pumping tube is shown. The probe-laser, plug-laser and photodiode (PD) were not utilised. Reproduced from reference [6].

far end of the cuvette (along the z direction), a cylinder is shown to represent the differential pumping tube, which provides differential pumping between the Source and Science cell, as well as being a velocity selective element for the MOT. If a pair of counterpropagating beams were added along the length of the cuvette (the z direction), this pure 2D MOT would become a $2D^+$ MOT.

In both the 2D and $2D^+$ MOT, the atomic beam then passes through a small aperture, typically acting as a differential pumping tube, which selects out a smaller range of longitudinal velocities, v_z , as only atoms that spend enough time within the MOT volume will be cooled and driven close enough to the aperture to make it through. In the case of Dieckmann using a $2D^+$ MOT, the measured resultant longitudinal velocity of 8 m s^{-1} is 3% of the mean longitudinal velocity at room temperature, 270 m s^{-1} . One of the other benefits the 2D MOT provides as a source

of cold atoms is that as the MOT is continually emptying²¹, interatomic collisions thus only play a minor role, reducing the collisional loss from the 2D MOT and thus increasing the load rate.

In 2002 J. Schoser[6] further investigated the properties of a pure 2D MOT as a source of cold atoms, in particular the dependence of the capture velocity and flux upon the intensity and length of the cooling beams[6]. Figures 5.4a, 5.4c and 5.4b, all reproduced from ref [6], show the measured dependencies upon the MOT beam power and length.

From figure 5.4a it is clear that there is only a relatively small dependence of the capture velocity upon the power per beam, with the capture velocity saturating to 38 m s^{-1} for intensities on the order of 17 mW cm^{-2} . Interestingly the capture velocity only increases by $\approx 30\%$ even though there is an increase in the intensity of the MOT beams by a factor of ≈ 20 , however, it should be noted that as the MOT flux depends strongly on the capture velocity²² this small change still leads to a significant change in flux.

Figures 5.4c and 5.4b show the dependence of the velocity distribution and flux on the length of the MOT beams, respectively. In figure 5.4c there is a clear increase in both the peak and mean longitudinal velocity of the emerging cold atomic beam distribution with increasing MOT length. This result is not unexpected as increasing the MOT length will increase the range of longitudinal velocities that spend enough time within the MOT volume to be cooled enough such that they make it through the differential pumping tube aperture. Figure 5.4b shows that the flux from the 2D MOT scales linearly with the length of the MOT beam, as a consequence of the increased longitudinal capture velocity.

Schoser saw a greater flux than Dieckmann, with a flux as high as 6×10^{10} atoms per second, and a similarly small divergence of 32 mrad , however, their measured longitudinal velocity was much higher, with a peak velocity on the order of 50 m s^{-1} , and was quite broad with a width of 75 m s^{-1} . The striking differences in mean velocities between the measurement of Schoser and Dieckmann arises due to the increased width of the MOT beams as well as a larger aperture. Schoser, using beams twice as wide, created a much greater volume and increased length over which cooling occurs, allowing atoms of greater longitudinal velocities to be cooled and make it through the aperture, which is also 6 times larger.

²¹As opposed to a pulsed 3D MOT.

²²In section 5.1.2 we saw that the flux was proportional to the fourth power of v_c .

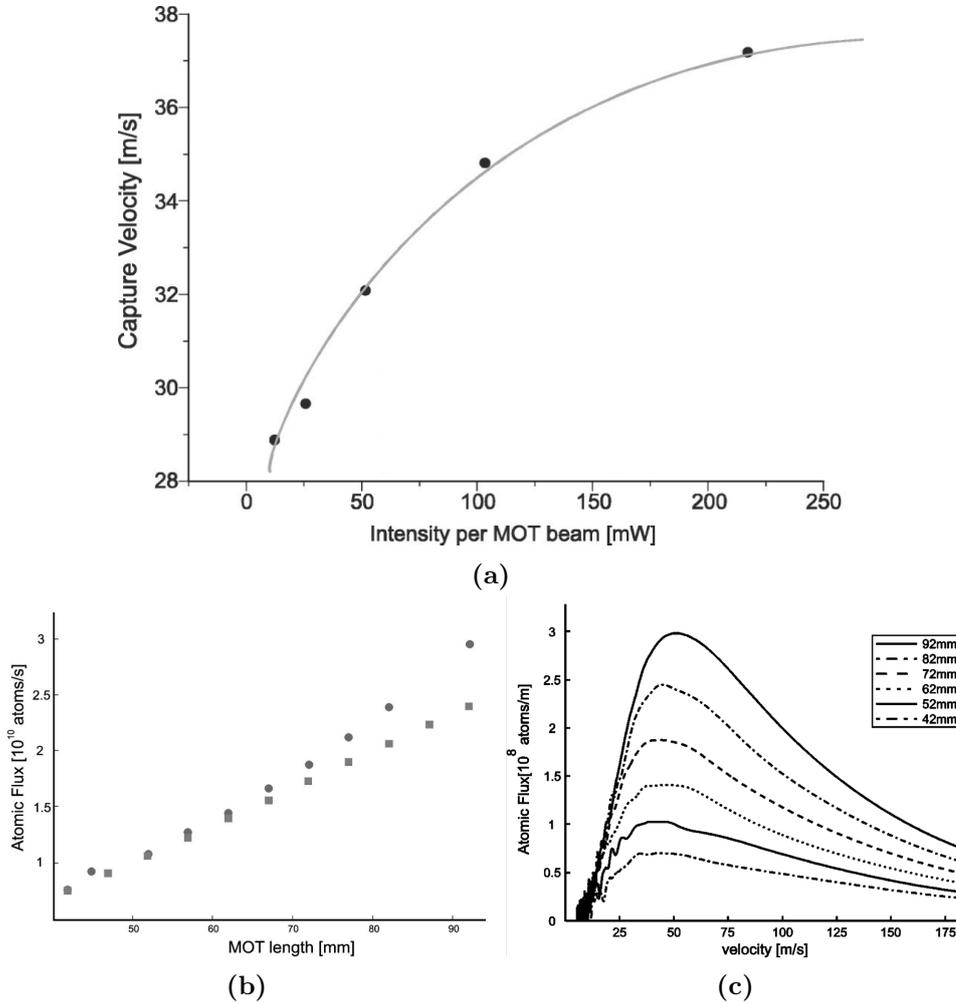


Figure 5.4: 2D MOT dependencies. a) **Capture velocity of a 2D MOT.** The dependence of the capture velocity of the pure 2D MOT upon the intensity per MOT beam is shown, Schoser *et al.*[6] saw a saturation of the capture velocity, and thus a saturation of the load rate, for intensities on the order of 17 mW cm^{-2} . b) **Velocity dependence on MOT beam length.** Shown above is the dependence of the velocity distribution of the cold atomic beam emerging from the 2D MOT on the MOT beam length. With longer MOT beams the cooling volume is extended and thus atoms with larger axial velocities can be cooled, leading to an increase in the peak axial velocity of the distribution. Circles indicate measured data whilst squares are the results of a model. c) **Flux dependence on MOT beam length.** The dependence of the flux of the 2D MOT on the length of the MOT beams, there is a clear almost linear dependence of the flux on the length of the MOT beams. The squares show the theoretical values whilst the circles show the measured data. All figures reproduced from ref. [6].

Chaudhuri, Roy and Unnikrishnan[87] investigated similar parameters as well as the dependence upon the detuning of the MOT beams, but directly comparing the case of a $2D^+$ MOT and a pure 2D MOT. For the $2D^+$ MOT they found fluxes as high as 2×10^{10} atoms/s, with a mean longitudinal velocity of 15 m/s and divergence of 26 mrad. Their results showed similar dependences of the flux upon the length of the transverse beams, however, they saw a much greater increase in the flux as a function of the power of the MOT beams for the $2D^+$ MOT as compared to the 2D MOT, and saw fluxes of up to 10 times greater in the $2D^+$ configuration as opposed to the pure 2D MOT.

Given the experiments described above, it seems that a $2D^+$ MOT is the most suitable choice as a source of cold atoms, as even though the pure 2D MOT yielded similar, if not greater fluxes, the greatly reduced longitudinal velocity of the $2D^+$ MOT reduces the required capture velocity of the 3D MOT and increases its collection efficiency. This reduction in longitudinal velocity occurs due to the presence of the longitudinal cooling in the $2D^+$ MOT which is absent in the pure 2D MOT. From the work of refs. [87] and [6], it is also clear that longer MOT beams leads to higher total fluxes, however, in the pure 2D MOT case this is also accompanied by an increase in the mean longitudinal velocity of the atoms.

3D MOTs

When considering the required parameters for a 3D MOT, the source of cold atoms for the 3D MOT must be considered. In equation 5.3 we saw the dependence of the atom number upon the capture velocity, however, this was for the case of a 3D MOT loaded from a background vapor. Here we are considering the case of a 3D MOT loaded from a cold atom source, instead. For simplicity lets work with Chaudhuri's parameters[87] and thus consider a source of 30 mrad divergence, and mean longitudinal velocity of 15 m s^{-1} .

The first parameter we shall consider is the required size of the 3D MOT beams, this will either be determined by the size of the cold atom beam by the time it reaches the 3D MOT²³, or alternatively the required capture velocity. For a beam of divergence 30 mrad, over the 50 cm ²⁴ distance between the aperture of the differential pumping

²³Note that the atoms loaded into the 3D MOT need to be directly caught from the cold atomic beam, as if they bounce off the walls of the chamber etc. before entering the MOT region, they may stick to the walls and heat up from the collision.

²⁴This distance is obviously highly dependent upon the vacuum system design, but 50 cm is a reasonable distance of separation.

tube and the center of the 3D MOT the beam will diverge to a full width, half maximum (FWHM) size of 2 cm, and assuming a Gaussian distribution²⁵ of atoms, then $\approx 95\%$ of atoms will be contained within beam of diameter 4 cm. From this calculation, it seems as though the MOT beams should be as large as possible to capture the entire flux as the windows of our science cell are only 4 cm wide, and we shall require on the order of ≈ 60 mW to ensure we are above the saturation intensity.

From equation 5.4, substituting in the MOT beam radius of 2 cm, we calculate a capture velocity of ≈ 45 m s⁻¹ well above that required for capturing the atoms from Chaudhuri's source and even large enough to capture a significant fraction of atoms from Schoser's pure 2D MOT source. Thus it is clear that the minimal beam size is determined by the divergence of the source as opposed to the required capture velocity. Thus to minimise the power required for the 3D MOT either a smaller divergence source, or (perhaps more readily achievable) a smaller distance between the 2D and 3D MOT is desirable.

Given we know the dependence of the load rate of our 3D MOT on the beam size, the next important parameters to consider would be the density and the temperature, as these would determine the phase space density (D) of the atoms in the MOT, and thus the initial D in a magnetic trap. From above we saw that the density of a MOT increases until the repulsive force from the re-radiation of photons within the cloud acts to limit the maximal achievable density. In equation 5.8 we saw that the density of the cloud was inversely proportional to the intensity of the trapping light, as a lower intensity would lead to a less repulsive force between the atoms. Similarly, in equation 5.6 we saw that the temperature of the MOT was also dependent upon the intensity, increasing with increasing intensity.

From the above two considerations alone it would seem that less intense light would be more optimal for cold, dense MOTs, however, less intense light also leads to a reduction in the scattering force²⁶, which leads to a reduction in the capture velocity of the MOT, as well as the spring constant of the MOT. In the case of loading atoms from a 2D⁺ MOT one would expect the reduction in power to not lead to a significant reduction in load rate as the capture velocity of the MOT is much greater than the mean velocity of the atoms, however, in the case of a pure 2D MOT one would

²⁵This is a reasonable assumption, as the atoms behave effectively as if they are in free fall during this time, thus the initial Gaussian distribution from the MOT evolves into a wider Gaussian distribution dependent upon the temperature of the atoms.

²⁶As less photons are scattered per second.

expect a decrease. Thus the optimal power and detuning²⁷ of the MOT beams becomes a balancing act between loaded number of atoms against temperature and density, highly dependent upon the velocity characteristics of the source of cold atoms. The common approach to solving this issue of having the optimal laser power both be small (for temperature and density) and large (for atom number) is to temporally adjust the MOTs parameters, using the optimal parameters for loading a large number of atoms into the MOT initially, then changing those parameters to optimise the density and temperature of the cloud. Thus for loading the MOT the laser intensity should be kept at its maximum value.

In the next section we shall outline how the density and temperature can be maximised and minimised, respectively, by changing the MOT parameters by undergoing a compression of the MOT followed by a period of optical molasses in which sub-Doppler temperatures can be achieved, before loading the atoms into the magnetic trap.

5.2 Sub-Doppler cooling

Once the theoretical lowest temperature achievable by optical molasses had been established it was time for experimentalists to test this theory. Temperature measurements of cold atoms, however, can be quite difficult, as any traditional method²⁸ fails. In order to measure the temperature of these cold atoms, early experimentalists²⁹ had to devise ingenious methods and, the standard method at the time was the release and recapture method.^[117]

In the release and recapture method the optical molasses is switched off for a finite amount of time and the atoms allowed to ballistically expand before the molasses is switched back on and the atoms are recaptured. The fraction of atoms that then remain in the molasses then yields information of how quickly the atoms were expanding and thus the temperature.³⁰ However, this method can have large uncer-

²⁷As the detuning of the MOT beams also determines the scattering rate and hence the repulsive force between the atoms, as well as the lowest achievable temperature.

²⁸Such as placing the atoms in thermal contact with a thermometer, not only is this impractical given the atoms are confined to ultrahigh vacuum, as there are such few atoms the thermometer will not equilibrate to the atoms temperature, but rather stay at its own temperature.

²⁹Such as Paul Lett, Rich Watts, Chris Westbrook, Phil Gould, Hal Metcalf and Bill Phillips.^[117]

³⁰Note this method requires the thermal expansion of the atoms to be much greater than the change in position due to gravity, or alternatively the effect of gravity need be accounted for or otherwise the temperature will be overestimated.

tainties associated with it, for example the measurement depends strongly upon the beam size and their exact overlap, and these properties can be difficult to measure experimentally³¹. Gould, Lett and Phillips in 1987 instead decided to measure the dependence of the lifetime of atoms in optical molasses as a function of detuning, but were greatly surprised when they found their results disagreeing strongly with the Doppler theory of optical molasses, finding the lifetime to be greatest for much larger detunings than expected[5]. These measurements led the experimentalists to devise new methods for measuring the temperature of the atoms in order to gain more data via different methods to corroborate their discrepancy with theory, these included;

- The time of flight method, in which the atoms are allowed to free-fall under the influence of gravity as well as ballistically expand. The time dependent fluorescence from a probe beam placed below the cloud then yields information about the temperature of the atoms, as the arrival time of the atoms will be determined by a combination of their velocity and the acceleration due to gravity.³²
- The fountain method, where a small probe beam is placed above the cold cloud, and the fluorescence as a function of probe beam position as it is displaced vertically yields information of how far the atoms can travel before their initial kinetic energy is all converted into gravitational potential energy. This method ultimately did not produce good results, due to unexpectedly low temperatures[5].

Using these new methods Lett *et al.* went on to measure the temperature of their atoms in the optical molasses finding temperatures well below that of the Doppler limit[5][118], up to an order of magnitude lower. These measurements spurred on new work on the laser cooling of atoms to try to explain how sub-Doppler cooling mechanisms could occur, with most of the work centering on the multi-level nature of the atoms, as well as the Zeeman sublevels once it was discovered that these temperatures depended strongly upon external magnetic fields as well as the polarisation of the incident light. The outcome of these results was a new theory of laser cooling, proposed by Jean Dalibard and Claude Cohen-Tannoudji, that the sub-Doppler cooling mechanisms were driven by a combination of polarisation gra-

³¹Especially the width of a Gaussian, as a true Gaussian beam is infinite in spatial extent, thus where does one draw the cutoff?

³²Nowadays a variation on the time of flight method is the standard tool of measuring the temperature of cold atoms, with the ability to spatially resolve the atoms with high resolution cameras, uncertainties can be reduced to $\leq 10\%$ quite easily with as few as 5 measurements at various times.

dients, the differential light shifts these induce for different Zeeman sublevels, and the optical pumping between these sublevels[119]. At the same time a theory was proposed by Ungar and Chu[120], which described the sub-Doppler cooling mechanism very similarly, again arising from the Zeeman structure of the atomic sublevels, and polarisation gradients in the optical molasses.

5.2.1 Polarisation gradient cooling

The works of Dalibard and Cohen-Tannoudji[119] as well as those of Ungar and Chu[120] explain in detail the sub-Doppler cooling mechanisms at work within the optical molasses, and as such won't be described in detail here, but we will touch upon the main points relevant to our purposes. To consider how the mechanism arises, we will consider the case of two counterpropagating beams of orthogonal linear polarisation. The interference of these two beams creates a standing wave pattern in which the polarisation of the resultant field changes, and given the beams are counterpropagating, this pattern changes on a subwavelength scale³³. In the case of orthogonal linear polarisation, at a position in which the polarisation of the two beams are in phase with each other, the resultant polarisation is linear, whereas at a position where the polarisation is out of phase by 90° , the resultant polarisation will be circular, and in between the polarisation will smoothly vary from linear through elliptical to circular. Now, if we start to consider the different Zeeman sublevels of our atoms, because the Clebsch-Gordan coefficients³⁴ are not constant between the differing Zeeman levels of the ground and excited states, the light shifts between these states are different, and change with polarisation³⁵, and therefore the light shifts also change with position of the atom.

Figure 5.5 shows a graphical representation of this mechanism. An atom initially in say the $m_f = 1/2$ (blue) state moving along the x axis climbs a potential hill (blue curve) that arises as the resultant polarisation of the light changes spatially. Once the atoms nears the top of the hill it is absorbs a photon from the laser and is excited to the excited state before it is preferentially optically pumped into the $m_f = 1/2$ state (red) which is at a lower potential, given the position of the atom. The atom then begins to climb the next potential hill (red curve) which again arises

³³As they are counterpropagating, the electric field will be periodic every $\lambda/2$

³⁴Clebsch-Gordan coefficients describe the coupling between different angular momentum states of an atom.

³⁵As the polarisation of the optical field determines which transitions are driven.

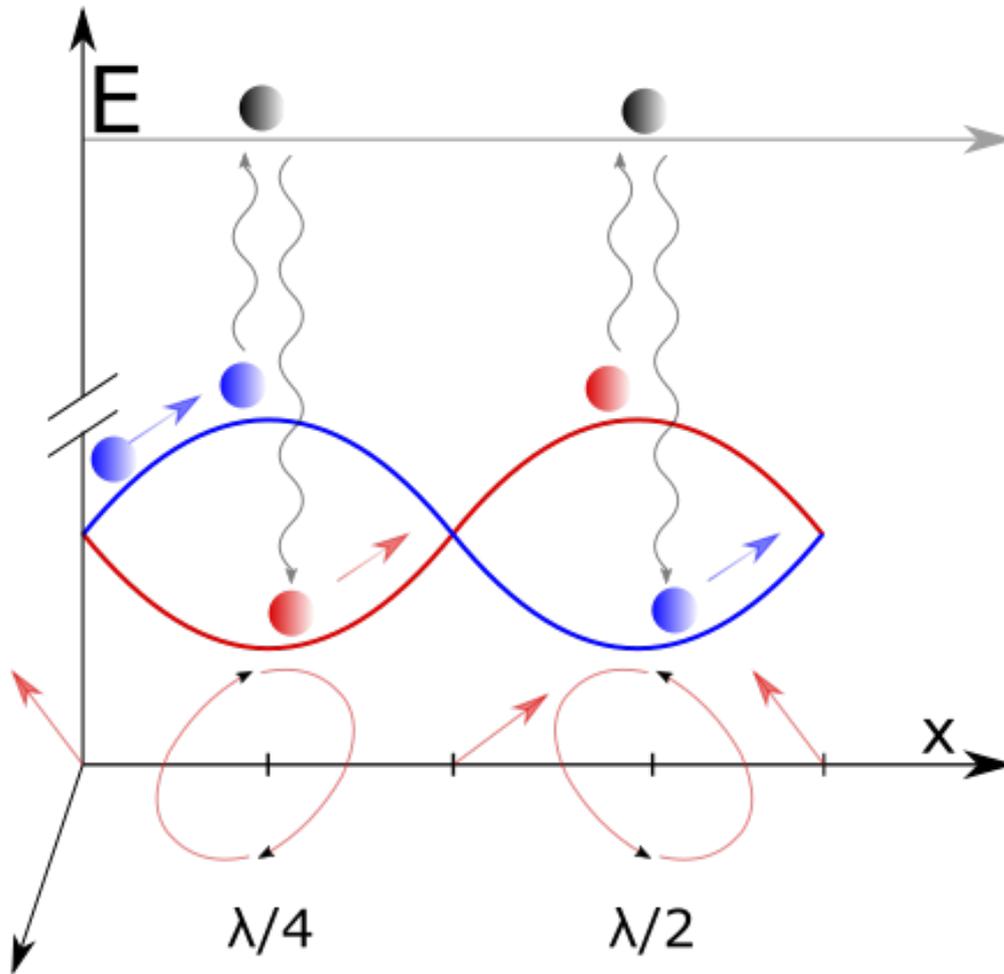


Figure 5.5: Graphical representation of polarisation gradient cooling A graphical representation of how polarisation gradient cooling occurs. Along the x-axis the polarisation of light is shown to vary over the course of half a wavelength, represented by the red arrows, from initially linearly polarised, through to circularly polarised, back to linearly polarised and finally circularly polarised but of the opposite handedness. The atom is represented by the coloured circle, with the colour indicating the magnetic substate the atom is in, either $m_f = 1/2$ (blue) or $m_f = -1/2$ (red), and the red and blue curves depict how the light shifts of the different substates varies as a function of the polarisation. An atom initially in $m_f = 1/2$ moving along the x-axis climbs the potential hill that arises due to the spatially varying polarisation before being optically pumped into the $m_f = -1/2$ and the process repeating until the atom no longer has enough kinetic energy to climb the potential hill.

from the spatially changing polarisation. This process continues to occur until the atom no longer has enough kinetic energy to further climb the potential hills.

From what we have described above we have a situation in which the difference

between the ground and excited state of the atom depends upon the Zeeman sublevel of the ground state that the atom is in, as well as the position, and these changes in energy will be periodic on the scale of half the wavelength of the cooling light. Polarisation gradient cooling (PGC) arises from these shifts in the energy of the ground state³⁶, and the variation in the Clebsch-Gordan coefficients meaning atoms will preferentially get pumped into a particular Zeeman sublevel of the ground state. This variation in the Clebsch-Gordan coefficients is such that the optical pumping of the atoms is always such that atoms tend to populate the lowest energy state. This optical pumping means that as atoms move through the standing wave potential the atoms will initially be at a potential minima, but will then have to start climbing the ‘potential hill’ as the polarisation changes. Once the atoms are at the top of this potential hill, however, they will then be optically pumped into the other Zeeman sublevel, which will then be at a minima not a maxima. This process continues as the atoms move, forcing them to continually climb potential hills, every time reaching the top being pumped back to the bottom, which is how this process came to be known as Sisyphus cooling.

A question that quickly arises is, as the atoms are still absorbing and reradiating photons, why does the conventional Doppler limit not apply? In the case of conventional Doppler cooling, atoms absorb light from a laser beam and then reradiate it, however, the light absorbed is of lower energy than that reradiated and hence energy is carried away from the atom. This excess energy arises from the Doppler shift the atom experiences, and ultimately determines the size of the energy fluctuations the cooling force experiences. However, in the case of polarisation gradient cooling the excess energy that the light carries away is due to the difference in energy of the Zeeman sublevels, much smaller than the difference in energy caused by the Doppler shift of the atoms. This causes the fluctuations in this cooling force to be much smaller, ultimately allowing much lower temperatures to be reached.

5.2.2 Experimentally relevant qualities of sub-Doppler cooling

Having established the method by which the sub-Doppler cooling mechanism works, it is now time to turn to the experimentally relevant parameters, in particular we shall focus upon the the temperature limit. The mechanisms and resultant parameters of polarisation gradient cooling have been extensively researched[121][122][120][119],

³⁶We ignore the shift in the excited state, as the atom typically spends little time in the excited state compared to the ground state.

we shall focus upon the results presented by Dalibard and Cohen-Tannoudji for polarisation gradient cooling in $\sigma^+\sigma^-$ light, as this is most relevant for atoms in a MOT.

One of the most surprising results that was found when the first theories of sub-Doppler cooling arose was that for small intensities the damping force provided by PGC is independent of the laser intensity, in contrast to what we calculated earlier for the damping force due to classical optical molasses, equation 5.2. The damping coefficient calculated by Dalibard and Cohen-Tannoudji was found to be[119]

$$\alpha = -\frac{3}{2}\hbar k^2 \frac{2\delta}{\Gamma} \quad (5.9)$$

where α is the damping coefficient and all other parameters carry their usual meaning. What is perhaps most striking about this relationship is the dependence upon detuning, which finds the damping coefficient increasing with detuning, as opposed to decreasing as is the case with classical molasses. This can be understood as the damping coefficient in classical molasses arises from the rate at which photons are scattered, which decreases rapidly with further detuning, whereas the damping in PGC arises due to the relative light shifts of the different magnetic sublevels, which increases with larger detuning.

As we saw that the heating rate was proportional to the intensity, this implies that the lowest achievable temperature due to PGC is proportional to the intensity also, as opposed to independent of it as was the case for optical molasses. Dalibard and Cohen-Tannoudji calculated the minimum temperature for $\sigma^+\sigma^-$ cooling and found the temperature was given by

$$k_B T = \frac{\hbar\Gamma I/I_{sat}}{2|\delta|/\Gamma} \left[\frac{29}{300} + \frac{254}{75} \frac{1}{1 + (2\delta/\Gamma)^2} \right]. \quad (5.10)$$

This implies that by simply reducing the intensity of the cooling light the temperature of the atoms could be arbitrarily reduced. However, this simple analysis is not accurate, and is limited by another feature of PGC that is absent in classic optical molasses, a critical cut-off velocity after which the force reduces dramatically.

This cut-off velocity arises due to the requirement that the atomic populations of the sublevels as the atoms move through the polarisation gradient cannot adiabatically follow the polarisation gradients that lead to the effect of Sisyphus cooling. At velocities above the critical velocity, the atoms will feel very little damping from PGC and will mainly be acted upon by the Doppler cooling mechanism. For the

force due to PGC to be linear in velocity, we require that the Doppler shift be much smaller than the light shift for the atoms, given by

$$\Delta \propto \frac{I}{I_{sat}} \frac{\Gamma^2}{\delta} \quad (5.11)$$

from this we can see that the cut-off velocity will scale proportionally to the intensity and inversely to the detuning. Equation 5.11 thus implies that as we try to cool the atoms further by either detuning the lasers further or reducing their intensity further, the cut off velocity will start to reduce significantly, limiting the fraction of atoms that we can cool. If we substitute the requirement that $kv \ll \Delta$ into equation 5.10 we find that $Mv_{rms} \gg \hbar k$, and thus to maintain a linear damping force the rms velocity of the atomic motion must be greater than the recoil velocity.³⁷ If we chose to reduce our intensity further, or equivalently the detuning was further increased, we would expect a great reduction in the damping coefficient as PGC becomes ineffective, and perhaps even a resultant non thermal distribution of atoms. Thus we can see that with PGC lower temperatures can be reached by reducing the intensity of the light, and that this reduction does not come at the cost of a reduction in the restoring force, however, there is an ultimate lowest temperature achievable by PGC, which is approximately limited by the recoil velocity.

The final experimentally relevant parameter that we shall consider is the effect of stray magnetic fields upon polarisation gradient cooling. As PGC inherently relies upon the differing light shifts between different Zeeman sublevels of the atomic ground state, it is expected that magnetic fields will also play a large role in the efficiency of the cooling, as the magnetic fields themselves will shift the energies of the sublevels. In the first paper demonstrating PGC, Lett *et al.*[118] noted that applying a magnetic field of up to 2 G increased the temperature of the atoms by up to a factor of almost 20, increasing from 30 μK up to almost 500 μK . Figure 5.6 shows their results, which show a clear increase in temperature as a function of magnetic field strength.

This phenomenon was later investigated experimentally in more detail and it was found that the effect of the magnetic field was to shift the temperature to which the atoms were cooled, with the temperature being linearly proportional to the magnetic field. This was explained via velocity selective resonances[123] whereby the atoms undergo coherent two photon transitions when the difference in frequency

³⁷The recoil velocity is the velocity at which an atom initially at rest would recoil to when absorbing a resonant photon of wavenumber k , which results in a velocity of $\frac{\hbar k}{m}$, where m is the mass of the atom.

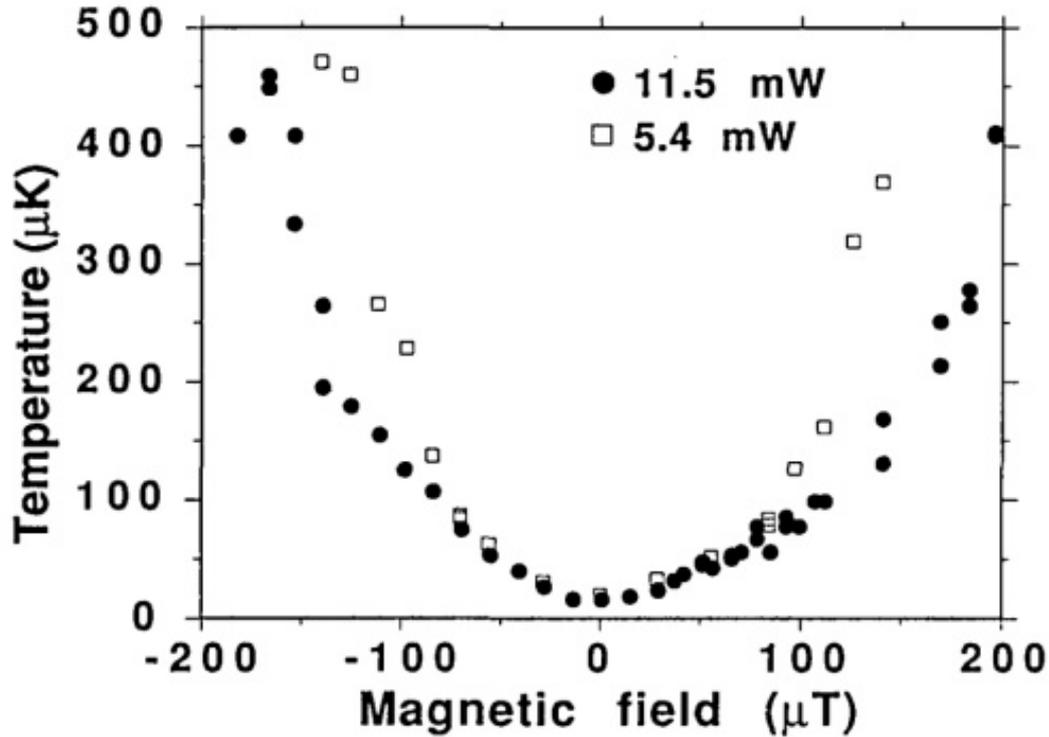


Figure 5.6: PGC temperature dependence on magnetic fields The experimental results of Lett *et al.*, in which they demonstrate how the minimum temperature achieved by PGC shows a strong dependence upon external magnetic fields. Two different sets of measurements at different powers are shown. Reproduced from ref. [5]

of two light beams seen by a moving atom is equal to the splitting of the ground state levels the atom transitions between. Phillips *et al.* went on to investigate and formalise these effects further[124], finding the same dependence of temperature upon applied magnetic field, as well as the inhibition of PGC forces in large magnetic fields. It should be noted, however, that the increase in temperature of the atoms when applying a magnetic field is not due to an increased heating rate nor decreased cooling rate, but rather it is a translation of the the velocity at which the PGC force is found, thus atoms can still be sub-Doppler cooled, however, they are being cooled to a finite velocity.

5.2.3 Sub-Doppler cooling in a MOT

From the discussion of PGC above and the dependence of the cooling upon applied magnetic fields, which can perturb the energies of the Zeeman sublevels to a greater

extent than is achieved through the radiation field, it would appear as though a caveat of cooling atoms in a MOT is that you do not gain the benefits of sub-Doppler cooling. However, in 1990 Steane and Foot[122] measured temperatures well below the Doppler limit in a MOT of caesium atoms and showed that the new sub-Doppler theory was required for describing the cooling mechanisms in a MOT, since when the beams in the MOT are well balanced the resulting MOT can be confined to a region where the perturbation due to the magnetic field is small enough to allow PGC to still provide a reasonable cooling force. This work was later extended to include a derivation of the region over which the cooling due to PGC in a MOT can extend, and its dependencies upon the MOT parameters[116].

In reference [116] the properties of a MOT in the temperature limited, multiple scattering and the two component regime are discussed. So far we have only described the temperature limited regime, where the number of atoms trapped by the MOT is so small that inter-atom repulsion is negligible, and the multiple-scattering regime, where the repulsive forces between the atoms start to dominate, and the MOT tends to a uniform density. The two component regime arises due to PGC occurring within the MOT, and occurs when the MOT becomes large enough such that the cloud is larger than the volume in which PGC occurs. In this case the MOT is decomposed into two components, the central component, which contains a dense cold cloud where position dependent forces due to PGC dominate, and a diffuse component in which standard Doppler cooling takes place.

In references [125] and [121] it is shown through semiclassical and quantum mechanical calculations, respectively, that in a one-dimensional MOT the radius at which PGC stops dominating as a cooling mechanism, r_l , occurs when the Zeeman shift due to the magnetic field gradient is on the same order as the light shift of the ground state. Townsend *et al.* assume this dependence in three dimensions and find the radius is given by

$$\mu_B b r_l = C_l \frac{\hbar \Omega^2}{\delta} \quad (5.12)$$

where μ_B is the Bohr magneton, b is the magnetic field gradient, C_l is a constant proportionality factor, whose value we will not concern ourselves with here, Ω is the Rabi frequency, given by $\Omega = \Gamma \sqrt{I/2I_s}$, and all other parameters are as previously defined. Importantly, if we substitute in the dependence of the Rabi frequency upon the intensity of the light we see that the radius at which PGC occurs³⁸ is proportional to the intensity of the light, and inversely proportional to the detuning

³⁸Or alternatively the radius at which the two component regime is reached.

and magnetic field gradient.

5.3 Compression of atoms in a MOT (CMOT)

The other key quantity to achieving a high initial phase space density, and thus good initial conditions for reaching Bose-Einstein condensation is the initial spatial density of the atoms. As we saw in section 5.1.3 the density of a large MOT is limited by the repulsive interactions between the atoms in the MOT caused by the emission and reabsorption of photons from within the MOT. We also saw that this density is inversely proportional to the light intensity, as seen through equation 5.8, indicating that reducing the light intensity will further increase the density of the MOT. Somewhat equivalently the density can also be increased by further detuning the laser, resulting in lower scattering in the centre of the MOT and thus further compression of the MOT. To date, further detuning the laser beams and reducing the intensity of the MOT beams is the most common method utilised to compress a MOT.

In 1994 Petrich, Cornell *et al.* investigated the possibility of compressing a MOT by temporally increasing the magnetic field gradient and found over an order of magnitude increase in the peak spatial density[126]. The authors note the conclusion that we have already made, that large detunings and large magnetic field gradients do not lead to large numbers of atoms as they are not optimal parameters for the load rate, but are rather only optimal for the density. However, by transiently changing the parameters of the MOT, leading to what the authors call a compressed MOT (CMOT), they were able to see over an order of magnitude increase in the spatial density of the MOT. As we saw in equation 5.7 and 5.8, the density of the MOT in both the temperature limited regime and the multiple scattering regime is dependent on the spring constant of the MOT. In this paper the authors studied how the size and shape of the cloud varied as they changed the spring constant of the MOT, actuated by varying the gradient of the magnetic field and the detuning of the cooling laser, for a range of different numbers of atoms in the MOT.

The detuning of the cooling light is thought to increase the spatial density of the MOT as it reduces the force between atoms due to re-radiation of photons, which is proportional to δ^{-4} while the trapping force in the central region of the MOT³⁹

³⁹This is the low magnetic field region where polarisation gradient cooling is the dominant cooling mechanism.

scales with δ^{-1} and in the outer region scales with δ^{-3} . Thus detuning further from resonance allows the cooling force to grow stronger compared to the repulsive interaction, further increasing the density. The authors studied this phenomena by jumping the detuning to three different detunings while maintaining the gradient at their MOT load value. At the collection gradient, they observed that the MOT remained an irregular shape, but they saw a 50% reduction in the horizontal width of their cloud, with no reduction in the vertical size, they attributed this to entering the temperature limited regime.

Next they studied the dependence upon the field gradient, initially jumping the detuning from -9 MHz to -44 MHz, and then ramping the field gradient from their MOT load value of 11 Gcm $^{-1}$ to a range of values up to 228 Gcm $^{-1}$ in 5 steps over 5 ms before allowing the cloud to equilibrate for 20 ms. For intermediate values up to 60 Gcm $^{-1}$, the atoms compress into a Gaussian shaped peak, while for larger field gradients the authors observed a diffuse cloud surrounding the narrow central peak, with the shape and position of this cloud depending strongly upon the alignment of the cooling lasers, whereas the narrow peak was insensitive to these parameters. This two component nature has been predicted[107] and, as described above in section 5.2.3, is attributed to the central region being dominated by polarisation gradient cooling, while the outer cloud undergoes only Doppler cooling. This description also explains why the central region is relatively insensitive to the alignment of the beams, as the volume and shape of that region is determined by the magnetic field parameters, and polarisation gradient cooling has been found in some cases to be insensitive to intensity imbalances⁴⁰.

Interestingly the authors found that the size of the cloud scaled with the square root of the field gradient, which is the expected behaviour for a cloud that maintains constant temperature, indicating that even as the atoms are compressed they do not heat up.⁴¹ Noting that the spring constants for MOTs are typically on the order of 100's of kHz, the atoms should adiabatically follow the temporal changes of the MOT. Thus a reduction in the volume should lead to an increase in temperature of the MOT. However, as polarisation gradient cooling is the dominant cooling mechanism in the central region of the MOT, and can cool atoms in milliseconds to its limit, this indicates that as the atoms are being compressed they are also being significantly cooled, which leads to greater phase space density. Some of the other

⁴⁰See the further remarks of reference [107].

⁴¹As PGC can bring atoms to the sub-Doppler cooling limit within ms, this is not surprising, since the atoms have enough time to cool as they compress.

important results for our application include the authors noting negligible change in atom number with the compression, overall seeing a 25 times increase in the peak spatial density and obtaining a relatively pure Gaussian profile for the density of the atoms. The final point may not seem particularly relevant, however, when catching the atoms in a purely magnetic trap from a MOT, the heating of the atoms during this process is dependent upon the overlap of spatial profiles of the MOT and magnetic trap. Given the distribution of the atoms in the magnetic trap is Gaussian, this would indicate the atoms from the CMOT will heat less as they get caught in the magnetic trap.

More recently some further studies with larger atom numbers have been performed. For example, DePue *et al.* studied a similar phenomenon using Cs atoms in a MOT, and reached peak densities of 10^{12} atoms per cm^3 [127]. In their case, the MOT was originally loaded with more atoms than there are when the peak density occurs, as they have significant losses while the atoms are compressing, but the total density during the compression still increases as the atoms that are falling in from the edges of the MOT replenish the atoms lost from the increased losses. DePue's results corroborate the measurements made by Petrich *et al.*, also seeing the two component nature of the MOT, and confirming the dependence of the central Gaussian peak upon the laser parameters, originally determined by Townsend *et al.* in reference. [116], as described in section. 5.2.3 and equation. 5.12.

Given the above works it seems that a temporal compression of atoms in a MOT can lead to significant enhancements of the initial phase space density, as well as improving the spatial overlap of the MOT and atoms at an equivalent temperature in a magnetic trap. The works do not mention the resultant final temperature of the atoms, however, but as the central component of the CMOT is determined by the polarisation gradient cooling forces at play, the temperature is expected to be determined by this mechanism. The above two sections describing the sub-Doppler cooling mechanisms and the temporal compression of the atoms in a MOT provide a way forward that should lead to very high initial phase space densities in a MOT for loading a magnetic trap with minimal heating; implement a CMOT stage where the atoms are forced into a very high density Gaussian distribution, and then apply further PGC by rapidly switching off the quadrupole gradient and nulling stray magnetic fields with the bias fields to achieve as low temperature as possible. As PGC typically only takes milliseconds to cool atoms to their final temperature, and the viscous damping force it applies is quite large, the spatial distribution of the atoms during this stage should change only minimally, leading to a very cold, very

dense cloud of atoms to be loaded into a purely magnetic trap.

5.4 Experimental implementation of laser cooling

In section 4.3 the laser setup used to produce the light for laser cooling in the 2D and 3D MOT was described. Here I shall specifically outline how the light is then used to laser cool the atoms, and the optimisation process that was involved to achieve optimal loading and cooling. My BEC apparatus utilises the dual MOT scheme, where a cold atomic beam is formed by a two-dimensional MOT in a region of relatively high background pressure, which is then captured in a three-dimensional MOT in a region of ultra-high vacuum. This scheme allows the three dimensional MOT to be loaded with a large number of atoms within seconds, without compromising the ultra-high vacuum required for the production of a BEC.

5.4.1 Two-dimensional Source MOT

The 2D Source MOT produces a pre-cooled, low divergence beam of atoms for loading the 3D Science MOT. The 2D Source MOT is formed in the source chamber, described in section 4.1.4, by placing 4 racetrack coils about the chamber, each coil parallel to one of the faces of the glass chamber.

Each racetrack coils consists of 4 layers of 5 turns (20 turns total) of 2 mm insulated copper wire, which has been wound onto a coil former, with a layer of epoxy being applied after each layer of the coil had been wound. The coil formers were designed in Solidworks, and printed by a 3D printer using ABS. These formers are designed to lock together with one another, guaranteeing the geometric alignment of the coils with one another, and have 60 mm cage rod holes for compatibility with the Thorlabs 60 mm cage system, guaranteeing the geometric alignment of the Source MOT optics with the coils. The inner profile onto which the coils are wound consists of two 25 mm inner diameter circular ends, separated by 50 mm, providing an almost entirely unobstructed view of the glass cell. When running the Source MOT, on average 7.5 A is driven through each coil, however each coil has its own Magneato driver allowing the coils to each be independently driven, which is particularly useful for aligning the MOT with the differential pumping tube. To geometrically align the racetrack coil formers to the cell, the formers were mounted on adjustable mounts

in both the x and y direction, and a set of callipers were used to ensure the distance of each of the formers from each face of the cell was the same.

For the MOT light, both the cooling and repump light used for the 2D MOT, once double passed through their respective AOMs as described in section 4.3, are combined upon a PBS. Both the repumping and cooling beams have an ≈ 3 mm Gaussian profile at this point with approximately 15 mW in the repump beam and 250 mW in the cooling beam. The light then passes through a half wave plate used to rotate the initially linear polarisation of the two beams, so that when the beam is split into two beams on a PBS further down the path, the two beams have approximately the same power. The beam then passes through a $2\times$ cylindrical telescope, before being further expanded by a $2\times$ spherical telescope to produce a 6 mm by 12 mm beam. At this point the light then passes through the aforementioned PBS and is split into the vertical and horizontal 2D MOT beams. Each of these MOT beams then passes through a quarter waveplate producing circularly polarised light⁴², before being further expanded by a $4\times$ spherical telescope, producing an approximately 24 mm \times 48 mm beam, as large as reasonable when using two inch optics and a 25 mm cell, before entering the source cell. After passing through the glass cell, the beams are then retroreflected by a custom⁴³ two inch diameter quarter waveplate and mirror, ensuring the handedness of the polarisation remains unchanged when the beam propagates back through the MOT.

To align the MOT beams to the glass cell an iterative two step process was followed utilising either of two mirrors. First the MOT beam being aligned was apertured down to 1 mm and the fluorescence caused by this beam, when viewed with a camera along the length of the 2D MOT, was aligned to the DPT using the first mirror⁴⁴. Next the reflection off the front and back surfaces of the glass cell was aligned upon the incident beam using the second mirror. Iterating this process until the two conditions are satisfied, that the beam is both centred on the DPT, and that it is propagating perpendicular to the glass cell's surfaces, guarantees both the position and angle of the beam are theoretically optimal for the MOT. This process was repeated on both the horizontal and vertical MOT beams.

Once the coils and cooling light have been aligned to the cell, the push beam was then aligned. The push beam propagates along the length of the 2D MOT, centred

⁴²Here one needs to ensure the polarisation is of the correct handedness with respect to the magnetic field.

⁴³This is a custom optic made by Union optic.

⁴⁴Furthest from the cell.

on the MOT's field zero, and acts to push atoms out of the 2D MOT towards the 3D MOT. To align the push beam, the beam is first centred on the DPT (observed by a camera) with the first mirror, while the second mirror is used to ensure the push beam exits the opposite end of the apparatus centred on the final window, this ensures the push beam traces out the trajectory the atoms ideally will follow.

Pressurising the source cell

So far I have not discussed how the pressure of rubidium in the source cell is varied. As previously mentioned in section 4.1.4, the source chamber has several SAES rubidium and potassium getters built into it with electrical feedthroughs extending through the glass chamber so as to allow current to be passed through the getters. These getters are used as the source of rubidium in the vacuum system, and are actuated by passing a current through the getter that heats it, catalysing a chemical reaction between the salts contained within the getter, creating a gas of rubidium. Typically I actuate the getters once a day by passing ≈ 7 A through the getters until the pressure in the main chamber reaches 2×10^{-9} Torr, before turning them off. This creates a large enough vapor pressure in the source chamber to run experiments for most of the day.

However, by the end of the day the pressure in the source chamber has dropped significantly enough that the load rate and thus saturated size of the MOT has decreased by approximately a factor of two. To counteract this, I also installed a pair of ≈ 1 W UV LEDs, which are directed onto the differential pumping tube, allowing us to perform light induced atomic desorption (LIAD). LIAD uses non-resonant light to help desorb atoms off the walls of the chamber, which appear to preferentially stick to metallic surfaces, hence we observe the largest pressure increase when the LEDs are directed at the differential pumping tube. The current through the LEDs is modulated with a Mornsun KCH24H LED driver, which drives ≈ 1 A through the LEDs in series and allows the LEDs to be modulated with a TTL signal.

An image of the source chamber when illuminated with UV light is shown in figure 5.7. During the morning the UV LEDs have little to no effect upon the load rate of the 3D MOT, as the pressure in the source chamber is saturated. However, as the day progresses the pressure in the source chamber drops as it is slowly pumped out through the differential pumping tube, and LIAD then helps maintain the pressure in the source chamber. During an experiment the LEDs are only switched on dur-

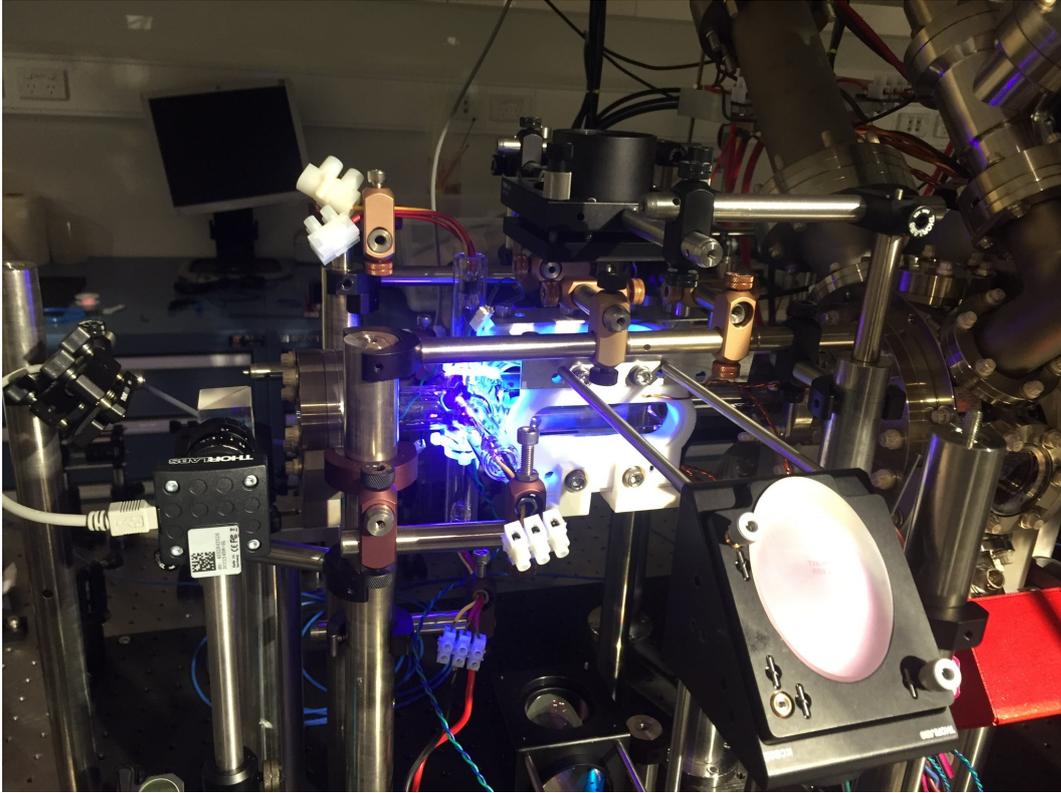


Figure 5.7: Light induced atomic desorption An image of the source chamber whilst illuminated by UV LEDs to perform light induced atomic desorption. When the pressure in the source chamber is low enough LIAD can provide up to a factor of 20 increase in the load rate of the 3D MOT.

ing the loading of the MOT, as the high pressure in the source chamber need not be maintained during the rest of the experiment, however, I observe no significant change in the lifetime of the magnetic trap even with the LEDs on, testament to the pressure difference the differential pumping tube can maintain.

It should be noted that when the LEDs are desorbing atoms off the walls of the chamber there is initially a larger increase in the load rate of the 3D MOT than in steady state⁴⁵. It is not known whether this increase is due to either initially having a slightly larger pressure in the source chamber when desorption starts, before it reaches an equilibrium pressure, or whether the pressure in the chamber increases too much and high density collisions start to deplete the 2D MOT. My thought is

⁴⁵Because of this effect if one runs three sequential identical experiments the loaded number of atoms into the MOT will decrease with each experiment, this effect has to be considered during optimisations or a false maxima might be chosen. Normally after 3 experiments the load rate settles to a constant value.

that this arises due to the first of these suggestions as when varying the current through the LEDs the largest load rate in steady state was achieved with the largest current, corresponding to the largest optical power, thus suggesting we have not yet reached this high density regime that would deplete the 2D MOT.

Optimising the Source MOT

Once the Source MOT has been roughly aligned it can then be optimised, however, to optimise the Source MOT, the Science MOT must be used, as the load rate of the Science MOT is the key parameter to optimise for the 2D MOT. Luckily, the parameters for the Source MOT to create a fast load rate for the 3D MOT are typically independent of the 3D MOT parameters. This is as the capture velocity of the 3D MOT is typically much greater than the mean velocity of the cold atomic beam from the Source MOT, so optimising the load rate of the 3D MOT amounts to simply optimising the flux from the Source MOT without having to consider changes to the mean velocity. This allows the optimisation of the Source MOT parameters and then the 3D MOT parameters, without having to iterate back and forth.

To optimise the load rate of the Source and Science MOT, a 0.2 ms fluorescence image of the 3D MOT after 3 seconds of loading and a 0.5 ms drop was taken, and the sum of the value of the pixels was used as the metric for atom number. To ensure that there were no multiple scattering events occurring within the cloud during imaging, and to ensure the fluorescence light reached the atoms, the imaging light was detuned 40 MHz below resonance. If the imaging light is instead on resonance, only the outer layer of atoms is likely to see the light as the cloud would typically be optically thick, and this metric would then optimise for large diffuse clouds as opposed to dense clouds.

To begin, the position of the field zero has to be aligned with the DPT. This can initially be performed by eye using a camera, but scanning the current through two coils on orthogonal axes will lead to a more precise optimum value. Figures 5.8a and 5.8b shows the fluorescence count rate as a function of the current through the bottom source coil and the source coil mounted on the south side of the experiment, respectively. From this data there is a clear optimal value for each. Also, it is clear from the figures that the functional form of the load rate is relatively symmetric about the optimal value, this indicates that at the optimal value the push beam is well aligned coaxially with the 2D MOT, as if this was not the case the push beam would break the symmetry of the 2D MOT, and the load rates would be

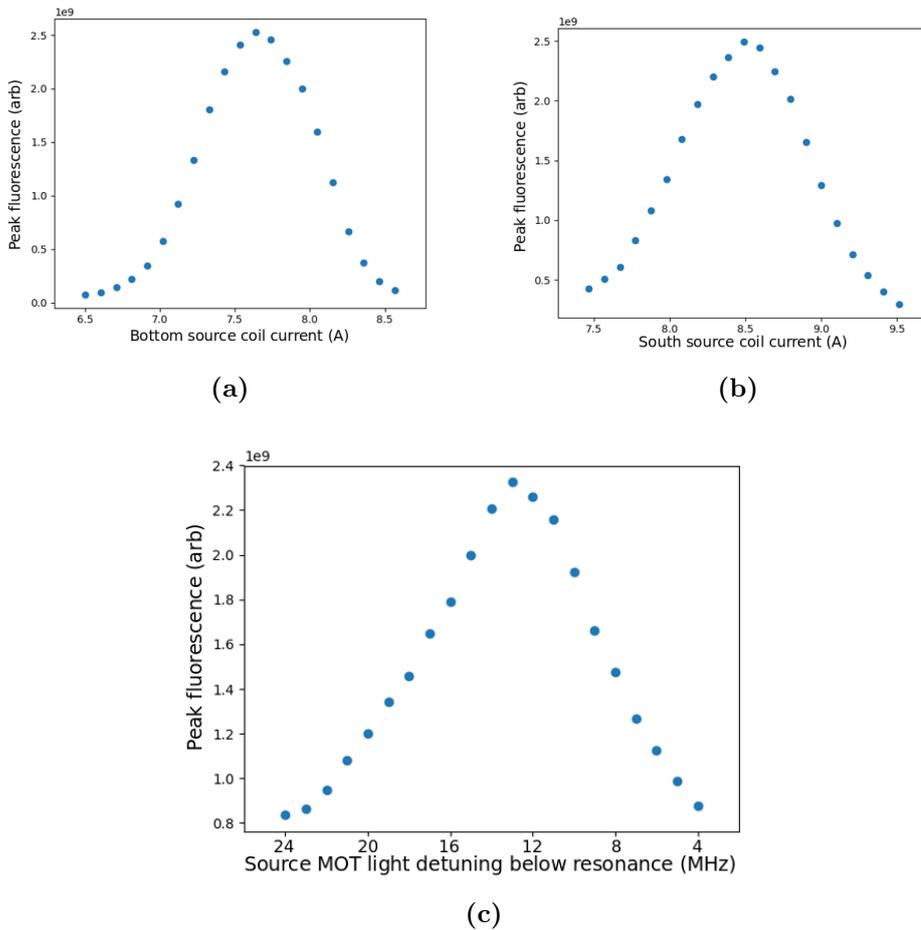


Figure 5.8: 2D MOT optimisation The measured peak fluorescence count of the Science MOT after 3 seconds of loading from the Source MOT while scanning three different parameters; **a)** the current in the bottom source coil, **b)** the current in the source coil on the south side of the chamber and **c)** the detuning of the 2D MOT below resonance. Clearly there is an optimum value for the current through each coil that, when viewed with a camera along the length of the Source MOT, coincides with the 2D MOT aligning with the hole in the differential pumping tube, whilst there is also a clear optimal detuning for the load rate, which is ≈ 13 MHz below resonance.

asymmetric⁴⁶.

Once the optimal position with regards to the DPT has been found the collimation of the final telescopes before the MOT can be tuned to maximise the peak density of the 2D MOT. To do this, the 2D MOT is observed on the camera while adjusting the collimation of the telescope, and simply maximising the peak fluorescence, which appears to also minimise the width of the cloud. The current in the source coils can then be optimised again, however, I found that adjusting the collimation of the 2D MOT beams has little effect on this parameter. Once the Source MOT has been optimised in this fashion it has a full width half maximum size of 0.12 mm in the vertical and horizontal directions, respectively. Figure 5.9 shows an image of the resulting 2D MOT when looking along its length, as well as a vertical and horizontal lineout through the centre of the MOT.

Next, the optimal source coil gradient and detuning must be found. To optimise the gradient of the source coils for a particular detuning, the current in each of the source coils in a given pair must be increased proportionally, to ensure the position of the field zero does not move away from the DPT. Once this has been optimised for a particular detuning, the current through the coils can then be scanned along with the detuning of the cooling light to maximise the flux into the 3D MOT⁴⁷. Figure 5.8c shows a scan of the 3D MOT fluorescence as a function of the 2D MOT detuning, which was found to have a maximal value at 18 MHz detuned below resonance. After the optimal detuning for the cooling light was found the amplitude of the cooling light as well as the amplitude and detuning of the repump light was also scanned. The load rate monotonically increased with increasing cooling laser power up to the maximal available (250 mW total) as did the load rate as a function of the repump power (15 mW total), and the optimal repump frequency was found to be on resonance.

With the 2D MOT parameters optimised the final element to optimise is the push beam. The load rate of the 3D MOT is a function of both the amplitude and detuning of the push beam, however this dependency is not separable in these two variables, thus a 2-dimensional scan of both power and detuning is required. Figure 5.10 shows the dependence of the load rate upon the detuning and amplitude of the push beam when both blue and red detuned. For this data, an approximately 1 mm Gaussian $1/e^2$ diameter beam with a maximum power of 3.8 mW was used. Clearly from the

⁴⁶Use of a ‘push beam’ to increase the flux from a MOT was first demonstrated by Lu *et al.*[128].

⁴⁷Note that the same rule applies regarding proportionally scanning the current to ensure the field zero remains fixed in position.

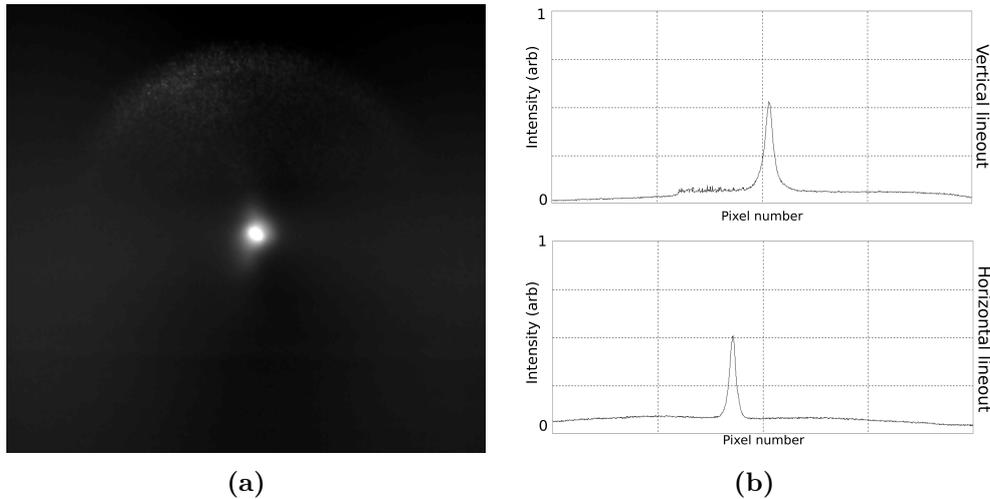


Figure 5.9: 2D MOT image. **a)** A 2.37 ms exposure taken of the Source MOT, looking along its length. Above the fluorescence of the 2D MOT, scattered light off the differential pumping tube from the vertical MOT beam can be seen. **b)** A vertical (top) and horizontal (bottom) lineout taken through the centre of the 2D MOT.

data, the load rate of the 3D MOT is largest when the push beam is red-detuned rather than blue detuned, with the largest integrated fluorescence in the red detuned scan being approximately 3 times larger than when it was blue detuned. I attribute this to the fact that in a 2D MOT, any atom that is already travelling in the direction towards the 3D MOT will continue to do so as there is no cooling along that axis. Thus a blue detuned beam would only serve to speed up those atoms that are already moving in that direction, thus narrowing the cold atom beams divergence, slightly increasing the load rate of the 3D MOT. With a red detuned push beam, however, atoms that are moving towards the push beam will instead be pushed in the opposite direction, increasing the flux of atoms out of the 2D MOT in the direction of the 3D MOT. I also found that the polarisation of the push beam has little effect upon the load rate.

Once all the above parameters have been optimised, it is then time to turn to the Science MOT.

5.4.2 Three-dimensional Science MOT

The 3D Science MOT is the initial trap used upon the path towards Bose–Einstein condensation. The coils used for forming the quadrupole field, as well as providing bias fields for shifting the position of the quadrupole field are described in section 4.4.

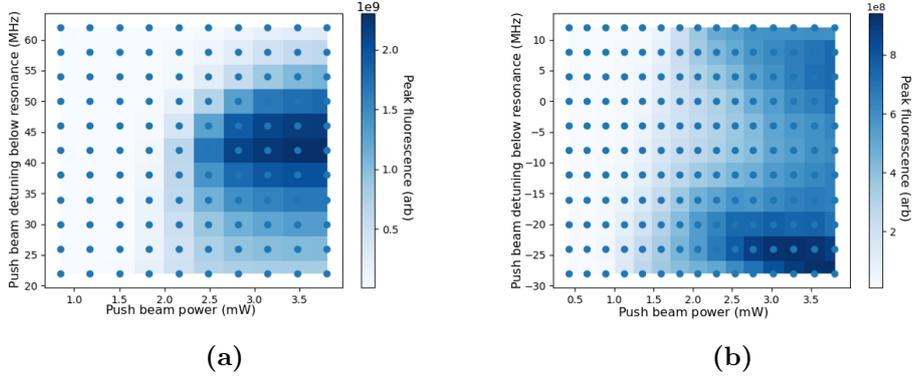


Figure 5.10: Push beam optimisation. a) and b) show the peak fluorescence from an image of the MOT, used as a metric for the load rate, as a function of the detuning below and above resonance (positive numbers indicate below) of the push beam, as well as a function of beam power. The colour coding of the 2D graphs are shown to the right of each, respectively. Note that the integrated fluorescence for the red detuned cases is approximately 3 times larger than for the blue detuned case.

The geometric centre of all the bias coil pairs as well as the quadrupole coils are coincident, however, the centre of the coils is not coincident with the centre of the octagonal glass cell described in section 4.1.4, but is rather displaced 7 mm below the centre. This feature allows a 14 mm clearance between the bottom of the glass cell and the top of the quadrupole coils, a large enough gap for a translation stage with a mounted optic and hole to be placed, so that the translation stage can move between two positions, one in which the optic is located directly below the centre of the cell, which retroreflects the MOT beam and ensures the handedness of the polarisation does not change, and the other in which a clear aperture is present for imaging through.

Both the repump and cooling light for the Science MOT is fiber coupled from the laser into two fibers, one which contains cooling light only and forms the vertical MOT beam, while the other contains both the repumping and cooling light and forms the two horizontal MOT beams. This light is transmitted along polarisation maintaining, FP/APC fibers which reduces back reflections from the fiber tips, minimising optical feedback into the laser. The light from the fibers is then recollimated with 50 mm focal length lenses, producing a Gaussian $1/e^2$ diameter of 6 mm⁴⁸. For the horizontal MOT beams, the collimated beam passes through a halfwave plate and is split into two equal power beams on a PBS, whereas the vertical MOT beam

⁴⁸In the rest of this chapter, when referring to the size of a Gaussian beam, the diameter I refer to is the $1/e^2$ diameter.

also passes through a waveplate and PBS, but for polarisation stability of the beam it is the fully transmitted light through the PBS that is utilised. After the cube each beam then passes through a $4\times$ spherical telescope⁴⁹, enlarging the beam to a 28 mm diameter before entering the cell. As $\approx 95\%$ of the power in a Gaussian beam is transmitted through an aperture that is $1.2\times$ the diameter of the beam, this ensures that we can have the largest beams possible entering our cell, with only a very small loss in power. Each of the three MOT beams is then retroreflected on an optic that is a quarter waveplate with a high reflectivity coating for 780 nm on the back side, ensuring the handedness of the polarisation remains unchanged upon reflection. In the case of the horizontal MOT beams, these optics also act as dichroics for passing and dumping the 1064 nm light used to form the dipole trap described in section 6.5.5.

Initial alignment of the 3D MOT is performed in a similar manner to the 2D MOT, however, as there is little background rubidium vapor in the Science cell, one cannot use the fluorescence of the MOT beams to align them. Instead the MOT beams are aligned to the cell itself. To perform this alignment, the horizontal MOT beams were aligned first, with the two mirrors after the final telescope of each beam that allow two point steering of the beam. A small piece of card that was within 1 mm of the size of the window, with a 1 mm hole cut into the centre of it, was used to align the beam to the centre of each window via the first mirror, and the second mirror was then used to align the reflection off the surface of the cell to the input beam. This alignment similarly guarantees, to within the manufacturer of the glass cells tolerances, the geometric alignment of the two horizontal beams.

To align the vertical MOT beam a similar process was used with regards to ensuring the reflection off the cell counter-propagated with respect to the original beam, but it was not possible to create an alignment tool that would accurately align the vertical beam to the centre of the two horizontal beams. Instead, the two horizontal beams were apertured down to 5 mm, and then the first mirror on the vertical beam would be used to steer the position of the beam such that a MOT was visible, and the second mirror could then guarantee the reflection off the cell counter propagated to the original beam.

⁴⁹A spherical telescope is a what one thinks of as a typical telescope, made from two spherical lenses that acts similarly upon both axes.

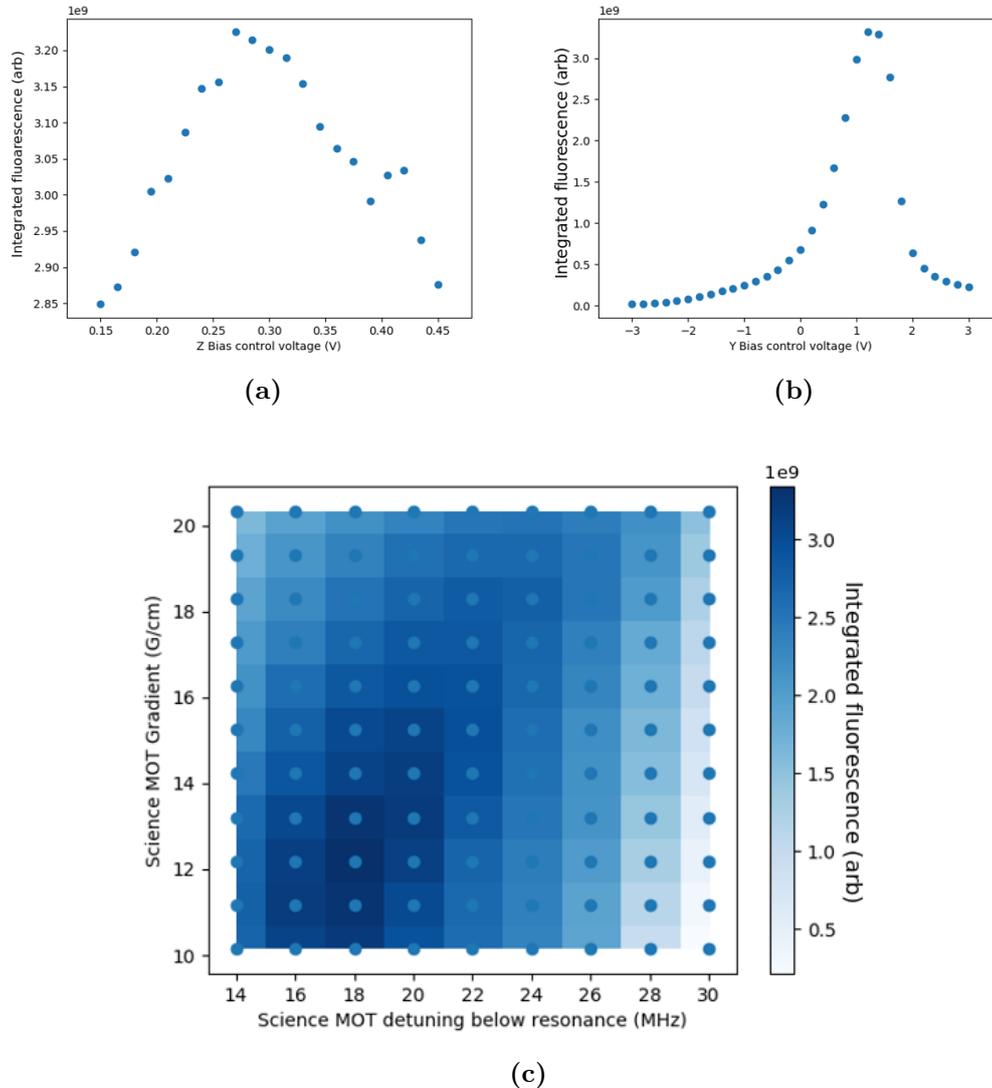


Figure 5.11: 3D MOT optimisation. a) and b) show the Science MOT fluorescence as a function of the control voltage applied to the bias coils Magneato in the Z (along gravity) and Y (along the optical axis for side imaging) directions respectively, which determine the location of the quadrupole field zero, and hence the centre of the MOT. Each has a clear optimal value for the load rate. c) The MOT fluorescence measured for different detunings of the cooling light below resonance, and values of the quadrupole gradient.

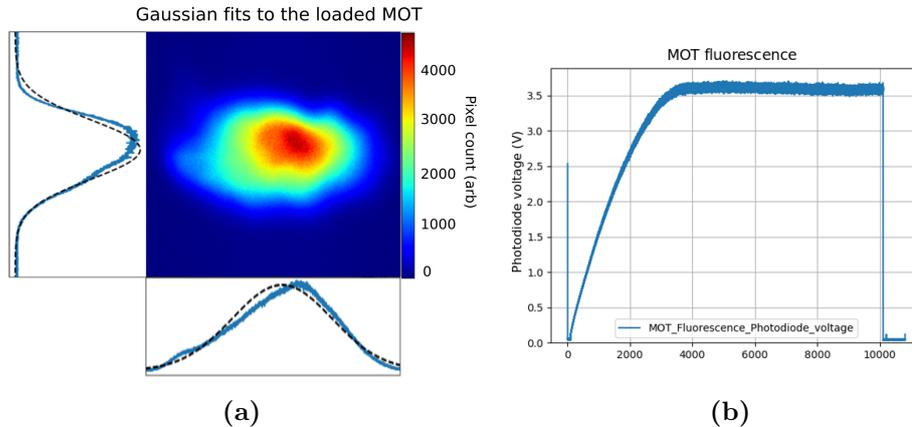


Figure 5.12: Loading of, and the final 3D MOT. **a)** A fluorescence image of the MOT when fully loaded, taken with the optimal parameters found in the optimisation procedure. It also shows two attempted Gaussian fits to the data, resulting in widths of 2.6 mm and 1.35 mm in the horizontal and vertical directions, however, the MOT clearly has an irregular shape thus these are only rough indicators of size. **b)** The measured fluorescence signal from the MOT while the MOT loads, the signal saturates to its maximal value in just under 4 seconds.

Optimising the Science MOT

With the MOT beams roughly aligned, and the polarisation made roughly correct by retroreflecting the beams through the setup and ensuring the beam is transmitted (reflected) off the PBS it was originally reflected (transmitted) through, the MOT can then start to be optimised. A similar procedure used to optimise the parameters of the Source MOT is performed for optimising the 3D MOT, first with the load rate dependence upon the position of the field zero, shown in figure 5.11a and figure 5.11b, which show clear optimal values for each of the bias coils that determine the loading position of the MOT.

Next, the quadrupole field gradient and cooling detuning from resonance must be scanned together as there will be an optimal detuning for a given gradient. Figure 5.11c shows the result of this 2D scan, the peak load rate was found to occur at a gradient of 12 Gcm^{-1} , and a detuning of 18 MHz below resonance. These values are comparable to what others find for maximising the load rate of the MOT. Note that the detuning is greater than the optimal theoretical value for the lowest temperature of the MOT. The data also shows a clear trend between the optimal load rate for a given field gradient, with a greater detuning required for an increased field gradient, as this would maximise the capture volume for the particular gradient.⁵⁰

⁵⁰In a MOT with a larger gradient atoms at a fixed distance from the field zero will experience

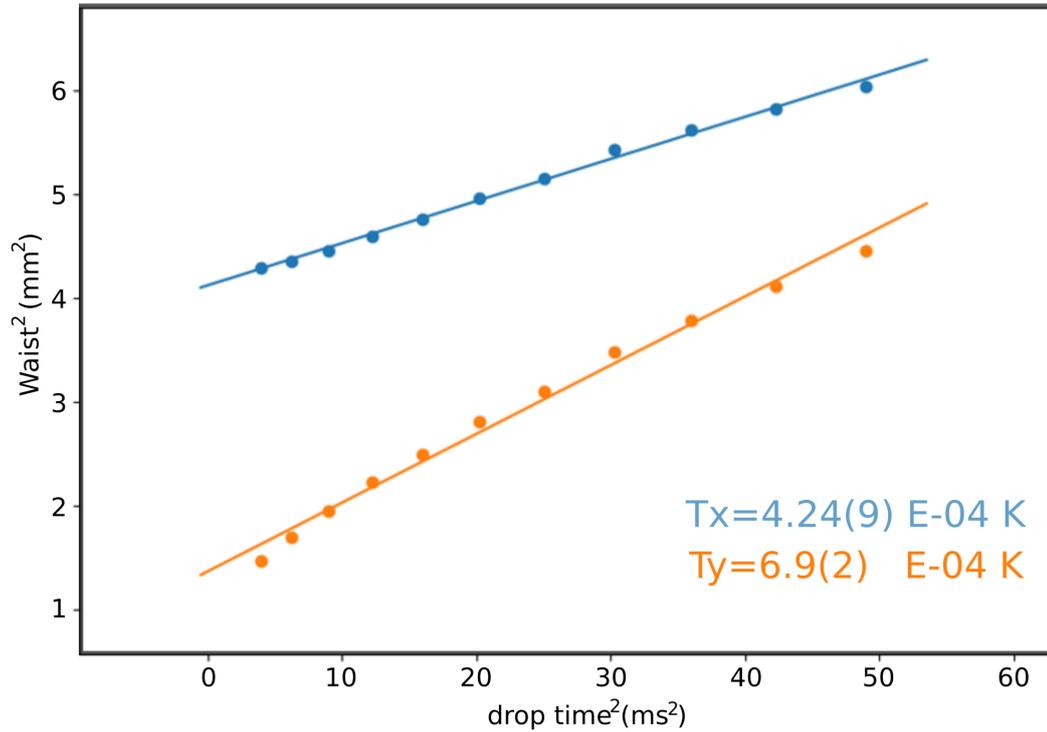


Figure 5.13: MOT Thermometry. Thermometry of the MOT once fully loaded, with data taken over the first 0.5-7 ms. The vertical axis shows the width of the cloud squared, whilst the horizontal axis shows the time squared, with the resultant temperature of 700 μ K found from the gradient of the fitted lines (shown).

Finally, one can turn back to the optical components, making small tweaks to optimise the size of the MOT, typically varying the collimation of the telescopes first, and then adjusting the alignment of the beams by small amounts. Interestingly, by adjusting the collimation of the top MOT beam the number of atoms in the MOT in steady state can be increased. This was seen as an increase in the volume of the MOT but also coincided with a reduction in density and load rate, thus the collimation was adjusted to maximise load rate and density, as number of atoms in the MOT was already ample to achieve large BECs.

Once this optimisation has been carried out, the Science MOT loads around 4×10^9 atoms in 4 seconds. Figure 5.12b shows the loading curve of the MOT, measured by reimaging the fluorescence of the MOT onto a photodiode approximately 150 mm from the atoms. Figure 5.12a shows the fluorescence of the MOT after a 0.5 ms

a greater Zeeman shift than in a MOT with a lower gradient, thus to increase the volume within which the atoms are likely to absorb a photon from the cooling beams, the detuning of the light must be further increased, however, this comes at the cost of reducing the scattering rate in the MOT and thus reducing its spring constant.

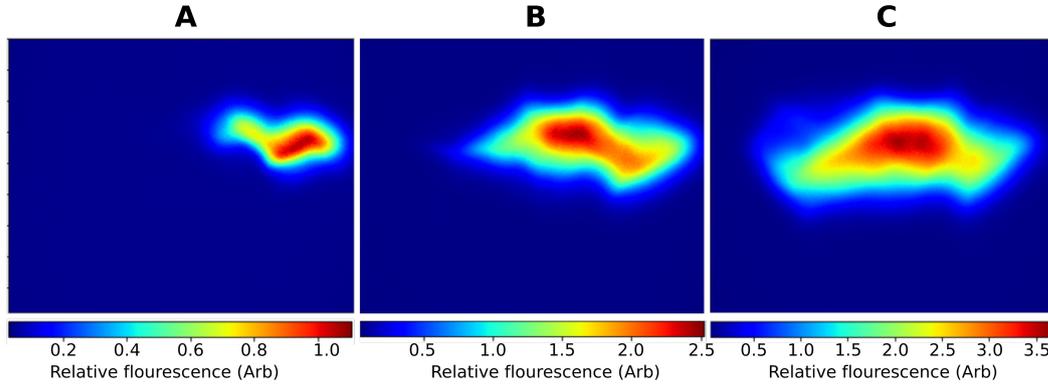


Figure 5.14: 3D MOT loading movement a), b) and c) show the 3D MOT at three sequential times during its loading. As the MOT loads the beam imbalance caused by the atoms casting a shadow in the MOT beams pushes the MOT back towards the centre of the cell.

drop⁵¹, with the cloud having a Gaussian width of approximately 2.5 mm and 1.1 mm in the horizontal and vertical directions, respectively. Clearly the MOT is quite irregular, with the large number of atoms in the MOT placing it in the multiple scattering regime and thus the shape being determined mostly by imbalances in the optical field and the resulting repulsive force that arises between the atoms as described in section 5.1.2. This is clear throughout the loading process, shown in the series of images in figure 5.14, where the MOT initially starts to load in the right side of the frame, but then moves to the centre as the optical forces become more balanced. As the MOT is made from retroreflected beams that are slightly converging, initially before the MOT becomes optically thick the retroreflected beams are of greater intensity, and thus unbalance the MOT pushing it towards the side of the cell that they arose from. As the MOT loads and the atoms start to create a significant shadow in the MOT beams, the retroreflected MOT beams intensity drops and the cloud gets pushed back towards the field zero, which is approximately centred on the screen, resulting in the final steady state MOT shown in figure 5.12a.

At this stage the temperature of the atoms is approximately 700 μK as is shown in figure 5.13, which shows how the square of the width of the cloud varies as a function of the square of the drop time over the first 7 ms of expansion. This temperature is in rough agreement with that predicted for a large atom number MOT in which the temperature of the MOT scales with the third root of the atom number[112]. For a MOT on the order of 5×10^9 atoms the predicted temperature of the MOT should

⁵¹Essentially in-situ.

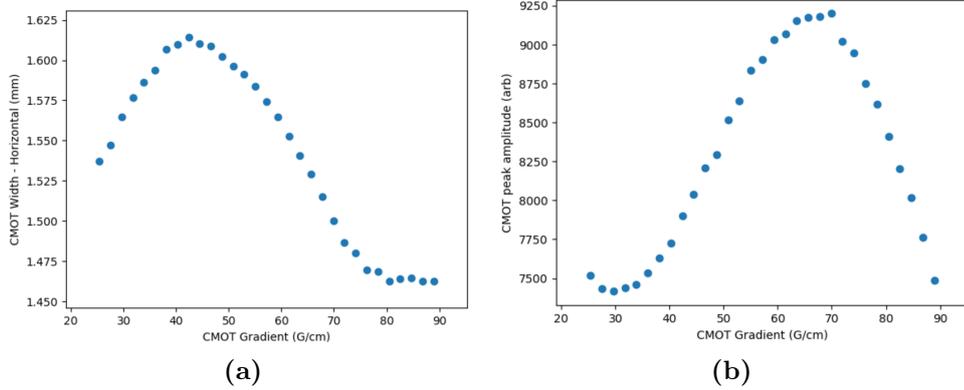


Figure 5.15: CMOT gradient optimisation. **a)** The dependence of the width of the resulting compressed MOT on the magnetic field gradient. A clear minimum in width occurs for gradients of 80-90 Gcm^{-1} . This data was taken after compressing the gradient over 200 ms. **b)** The integrated fluorescence from the MOT for the same set of experiments as in **a)** is shown, used as a metric for the atom number. This data clearly shows that whilst the width is decreasing the integrated fluorescence is increasing, indicating that the spatial density is increasing and is not being limited by other loss mechanisms.

increase by approximately a factor of ≈ 8 , and as we are not at the optimal detuning for temperature, but rather for load rate, the resulting temperature of 700 μK is not surprising.

5.4.3 Compressed MOT optimisation

After the Science MOT has been loaded to saturation the Source MOT is turned off, the push beam extinguished and the next stage of laser cooling, the compressed MOT stage described in section 5.3, where the spring constant of the MOT is increased by increasing the magnetic field gradient and the detuning of the MOT beams to spatially compress the MOT, is realised. This stage is relatively easy to optimise, as there are few parameters, namely the duration over which the gradient of the quadrupole trap as well as the detuning of the optical fields are ramped, the final magnetic field gradient and the final detuning of the cooling and repumping light.

To optimise this stage the dependence of the widths and peak fluorescence of the MOT was investigated for a range of different gradients and detunings. Figure 5.15 shows the dependence of the compressed MOT's width and peak fluorescence upon the gradient the quadrupole field is ramped. This data was acquired while also ramping the cooling light detuning to approximately 40 MHz below resonance, as large losses were observed if only the quadrupole gradient was ramped. In this data

there is a clear initial increase in the width of the cloud up to a gradient of 40 Gcm^{-1} , but this is then followed by a decrease in the cloud width, reaching a minimum at around 80 Gcm^{-1} . The initial increase in the width of the cloud is attributed to a mismatch of the gradient to the detuning, allowing atoms to move further from the MOT centre before they come into resonance with the cooling light. When observing the peak fluorescence count from the CMOT there is a clear increase until a gradient of approximately 70 Gcm^{-1} , after which the peak fluorescence begins to drop. The initial increase in the peak fluorescence is expected if the spatial volume of the cloud decreases whilst the loss rate is not significantly increased over the compression time. For larger gradients the decrease in the peak fluorescence is attributed to an increased light assisted collisional loss rate. From these scans we chose to compress our quadrupole field to a final value of 70 Gcm^{-1} .

Next the detuning of the cooling and repumping light from resonance was varied and the effect on the final width of the cloud was measured, shown in figure 5.16. This dependence of the peak fluorescence on these parameters was also investigated with this data set, and the data qualitatively showed the expected form, with the amplitude being maximised when the width was minimised. From these data sets an optimal detuning of 32 MHz and 54 MHz below resonance for the cooling and repumping light, respectively, was found to maximise the density of the cloud. These values are not surprising, and are close to the values found by Petrich *et al.*[126]. The amplitude of the cooling and repumping light was also scanned, but optimal conditions were found to occur with the maximal intensity available.

Finally, the duration of the ramping time for the changing parameters was then varied and the time at which the peak amplitude occurred was chosen, as this will coincide with the time at which the MOT is at its densest, as the atoms have had long enough to respond to the changes in the magnetic field, but not long enough to start significantly losing atoms. This was found to occur with a 40 ms ramp time.

The resultant CMOT after the compression stage with a 0.5 ms drop time is shown in figure 5.17, along with Gaussian fits in the horizontal and vertical directions. Clearly the CMOT now has a more Gaussian form than the initially loaded MOT and is also denser, with waists of 1.3 mm and 0.68 mm in the horizontal and vertical direction respectively, corresponding to a decrease in MOT volume by a factor of approximately 8 and thus an increase in the density of the MOT by a factor of 8 as there is no observed atom loss. This corresponds to an average density of approximately 10^{12} atoms per cm^3 . However, there is no clear sign of the two component MOT described in reference [126], but by performing thermometry on the

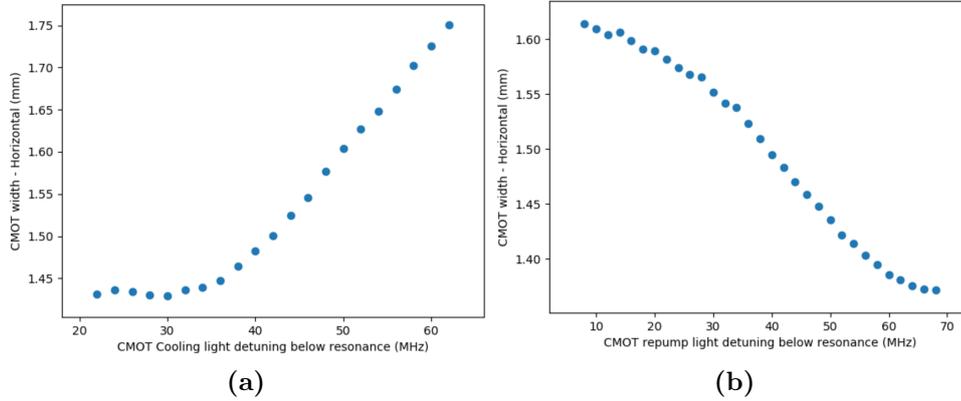


Figure 5.16: CMOT detuning optimisation. a) and b) show the dependence of the width of final compressed MOT on the detuning of the cooling and repumping light, respectively. Each dataset was taken with the quadrupole gradient set to 85 Gcm^{-1} . Each has clear optimal values for minimising the width of the cloud, and the peak fluorescence for the two data sets shows the expected inverse relationship, being near a peak when the width is minimised. A 2D scan of the detuning of both the repump and cooling light was also performed which showed that each parameter showed little dependency upon the other.

cloud, as shown in figure 5.18a, we can see two clear regions, one in which the CMOT undergoes a rapid expansion by a factor of 2 during the first 7 ms, followed by a slower expansion over the remaining 23 ms in which it undergoes a further expansion by a factor of 4. This unexpected behaviour in the ballistic expansion of the cloud is attributed to some fraction of the atoms being in the second component of the MOT, and are thus at a much higher temperature, accounting for the initial rapid expansion, while the temperature of the dense colder core is what is measured in the second regime. Figure 5.18b shows a fit to the shorter expansion time indicating a temperature on the order of $600 \mu\text{K}$, quite similar to the MOT temperature, as would be expected when only Doppler cooling is being employed. Figure 5.18c is the temperature fit to the longer drop times and shows that the rest of the cloud is expanding much slower, with an approximate temperature of $50 \mu\text{K}$, a reasonable temperature for a sub-Doppler cooled cloud⁵²

Note that for simplicity I chose to keep the ramp time for the optical and magnetic field parameters the same, however, this need not be the case. Also, the optical fields were chosen to be ramped from their MOT values, rather than leaving this as a free parameter. Further optimisation of the compression stage by having these as independent parameters could be performed, but was deemed not necessary as the magnetic trap catch after PGC was found to already be highly efficient.

⁵²Note that this is below the Doppler limit for rubidium.

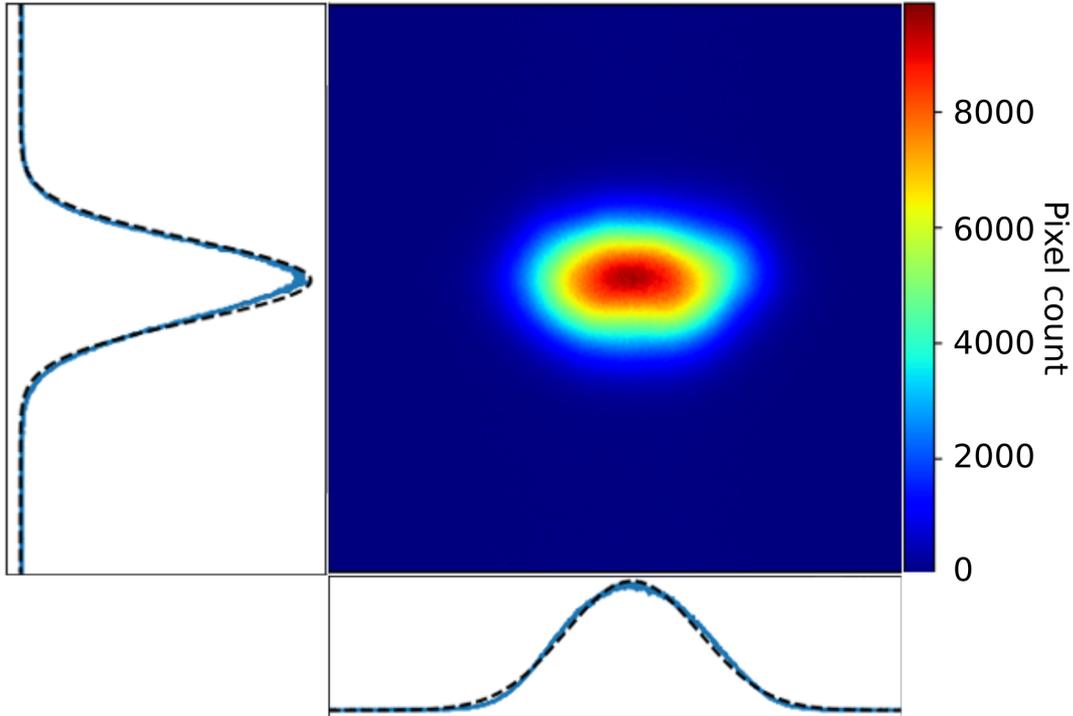


Figure 5.17: In-situ fluorescence image of the resulting CMOT. The compressed MOT in situ after the compression sequence. Clearly the MOT now has a more regular Gaussian form with the displayed fits to the line outs fitting the data quite well, indicating waists of 1.3 mm and 0.68 mm in the horizontal and vertical directions respectively.

5.4.4 PGC optimisation

With the compressed trap stage optimised, the final stage of laser cooling is polarisation-gradient cooling, used to cool the entire cloud, including the ‘second’ component of the CMOT, to as low temperature as possible whilst maintaining atom number, to load the largest number of atoms and highest phase space density into the magnetic trap. During PGC the detuning and amplitude of the cooling and repumping light are typically varied, and the magnetic fields are quickly changed to a value that produces approximately no net magnetic field at the atoms as this will minimise the temperature the atoms can be cooled to, and extends the volume over which sub-Doppler cooling can occur⁵³. PGC typically lasts on the order of 5 ms-15 ms, and during this time there is a strong damping force that is not position dependent, so the size and shape of the cloud should not differ considerably from its CMOT size and shape, but the entire cloud should now be reduced to the lowest temperature

⁵³Which as discussed earlier, is limited to the range over which the light shift is larger than the Zeeman shift.

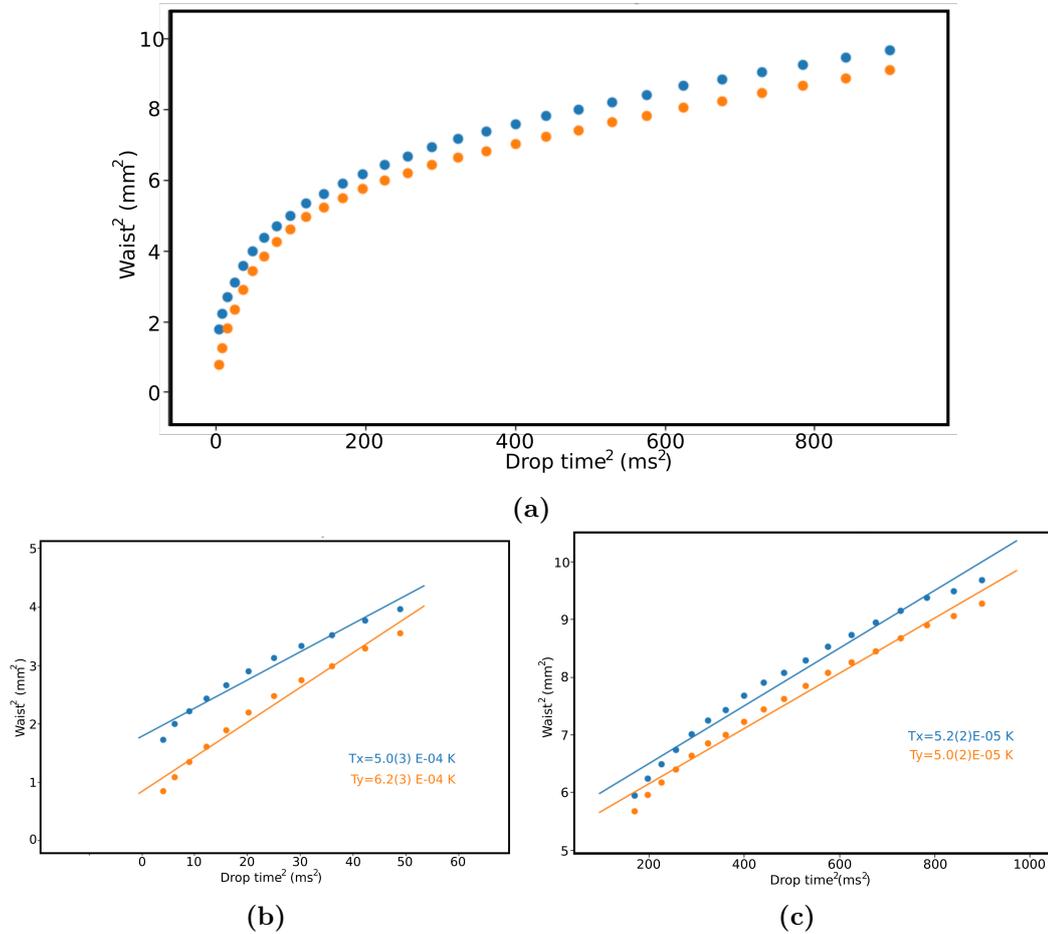


Figure 5.18: CMOT thermometry. **a)** Thermometry performed on the CMOT from 0.5ms to 30ms of expansion. The data is clearly not linear, **b)** and **c)** Temperature fits to the short (0.5-7ms) and long (7-30ms) drop time data. The short drop time data has a temperature found to be quite close to the MOT temperature, whereas the long drop time data indicates that sub-Doppler cooling is occurring in the centre of the CMOT.

achievable by our PGC, not just the inner component, which is already undergoing PGC.

PGC appears to be most sensitive to stray magnetic fields, so the key parameter to scan to ensure PGC is correctly working is the bias fields, while the quadrupole field should be quickly switched off. For my apparatus I found an initial quick reduction of the quadrupole and Z bias coils to a finite value over 4 ms before switching them to their final value was much smoother for the atoms, and greatly improved the load into the magnetic trap. To optimise the bias fields for PGC I typically find the bias fields that will ensure the cloud remains in position after 500 ms of PGC, however, to narrow in on the appropriate bias fields it may be easier initially to use shorter

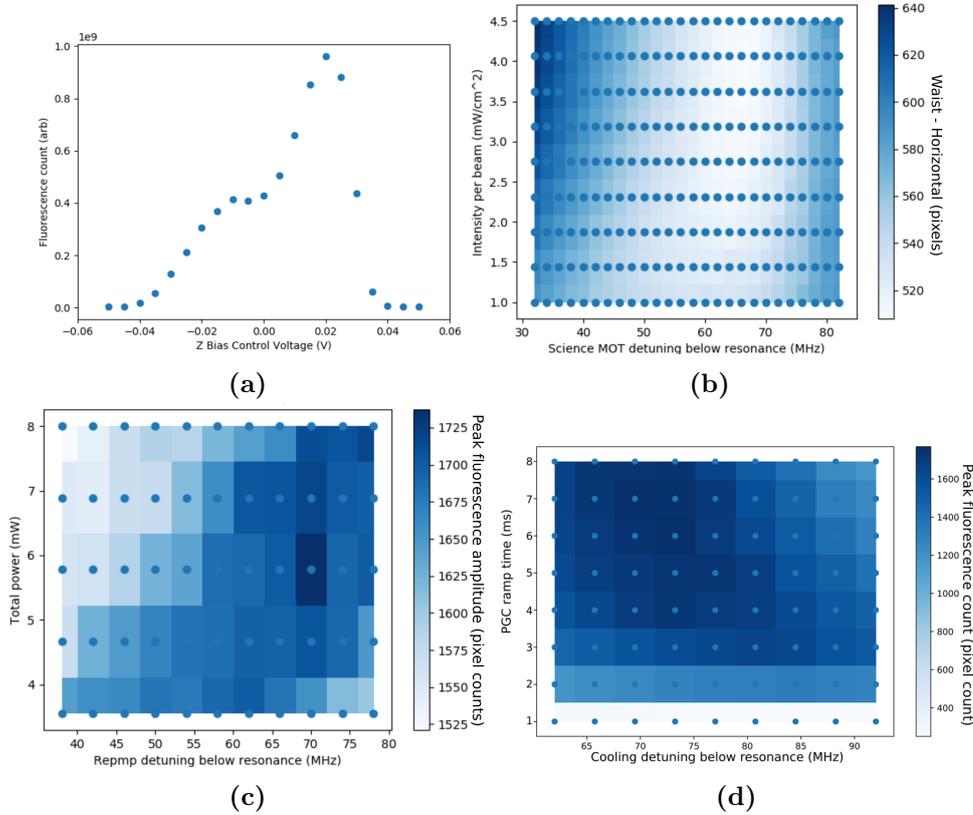


Figure 5.19: Optimisation of polarisation gradient cooling. **a)** Fluorescence as a function of Z bias control voltage. Note the very small range of control voltages (≈ 10 mV), outside this range no atoms were observed after 500 ms of PGC. **b)** Cloud width after a 20 ms drop as a function of detuning and intensity. **c)** Peak fluorescence after PGC as a function of repump intensity and detuning. **d)** Peak fluorescence as a function of PGC time and detuning

PGC times, such as 20 ms-50 ms, as on these time scales it will be easier to see if the atoms are being forced in a particular direction. Figure 5.19a shows the effect of the value of the Z bias field upon the final atom number after 500 ms. As is clearly seen there is a very strong dependence upon the bias field. Note the bias field control voltage is only tuned over 10 mV, as outside this range there were no atoms left at the end of the PGC.

Once the appropriate bias fields have been found to minimise the movement of the cloud during PGC, the next stage is to optimise the temperature of the cloud. As the cloud size should not change significantly during PGC, the width of the cloud after the longest drop time observable with the imaging system is a simple, one experimental shot metric that can be used to optimise the final temperature, opti-

minimising for the smallest width as the width will be proportional to the temperature⁵⁴. This is a rather large parameter space and requires a number of scans, as there is the duration of PGC, the detuning and amplitude of the cooling light, as well as the detuning and amplitude of the cooling light. Figures 5.19b, 5.19c and 5.19d show some of the optimisation in these spaces.

In figure 5.19b one can see the effect of the detuning and amplitude of the cooling light upon the final cloud width after a 20 ms drop. The data has a relatively broad minimum, with little change of the final width over an approximate 20 MHz range in the final detuning value, but it does show that the width is also minimised with the maximal available power when in this range, whilst larger powers lead to an increase in the width of the cloud for smaller detunings. This decrease in width with a greater power indicates that in our configuration a larger laser power is required to maintain the critical cut off velocity for PGC above the temperature of the MOT to ensure all the atoms are cooled, and that this effect dominates the heating caused from the laser.

Figure 5.19c shows a similar scan for the repumping detuning and amplitude. Here we find that there is an approximately linear relationship between the optimal detuning and intensity of the repumping light to maximise the density of the cloud after PGC, clear from the approximate straight line that can be drawn between the peak values in each row of data. Reducing the repumping light during PGC acts to ‘hide’ atoms in a dark state in which they will not be scattering light from the molasses, leading to a smaller repulsive force between the atoms and thus higher overall densities. However, with no repumping light present, once the atoms have fallen into the dark state they can simply drift out of the PGC region or fall under gravity as there is no molasses force upon them. By detuning the repumping light we can make the repump velocity selective, with a greater detuning selectively repumping atoms of a larger velocity due to the Doppler effect. The relationship we see between the repumping power and its detuning thus describes the optimal tradeoff between further detuning the repumping light from resonance so that we are mostly repumping atoms that are likely to escape the MOT volume during the duration of PGC, as opposed to repumping atoms that are near their lowest achievable velocity near the centre of the cloud, which would then just cause a repulsive interaction, and having enough repumping light to ensure that losses during PGC are minimised.

⁵⁴Minus an offset, which is the clouds initial size.

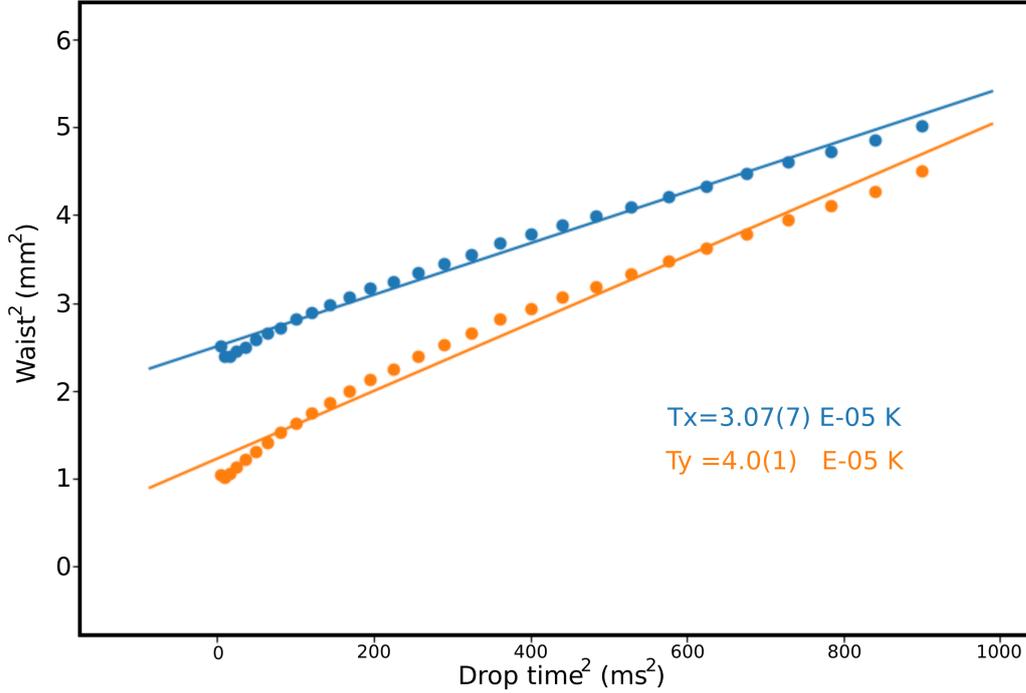


Figure 5.20: PGC Thermometry. Thermometry for the PGC cooled cloud for drop times in the range of 0-30ms. The data is much closer to a linear trend than the thermometry performed on the CMOT, indicating that nearly all the atoms are cooled to this final temperature and thus the two component nature of the cloud has been removed.

Finally, figure 5.19d shows the relationship between the duration over which the optical fields are ramped,⁵⁵ and the peak density, as well as rescanning the dependence upon the detuning of the cooling light. Here we see a clear minimum for a short PGC time of 1 ms, during which one expects not much cooling to occur, and then a gradual increase in the density up to a peak value occurring at around 8 ms of ramp time.

Optimisation of PGC resulted in a 4ms ramp of the coils from their compressed MOT values to their final PGC values before the cooling light was detuned by a further 38 MHz from the compressed MOT value to a final value of ≈ 70 MHz below resonance, while the repumping light was similarly further detuned by a further 24 MHz to a final value of 78 MHz below resonance over 8 ms. During this detuning the amplitude of the beams was not ramped, however, due to the double passed AOM efficiency varying as a function of driving frequency, the amplitude was reduced from the CMOT values of 4.5 mW cm^{-2} to an intensity of approximately 1.5 mW cm^{-2} ⁵⁶.

⁵⁵Note this occurs once the magnetic fields have been ramped off.

⁵⁶Note that these intensities are per MOT beam.

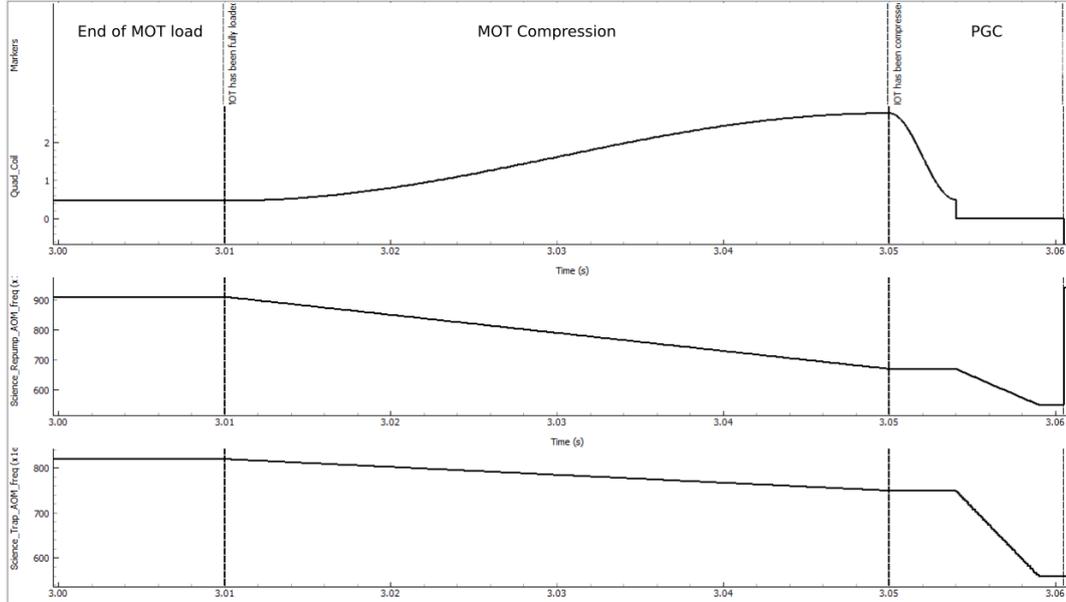


Figure 5.21: Runviewer trace of the laser cooling process A trace of relevant parameters that change during the laser cooling process is shown. The included parameters, from top to bottom, are the control voltage applied to the quadrupole coils Magneato driver, the RF frequency applied to the science repump AOM and the RF frequency applied to the Science MOT cooling AOM.

At the end of this ramping, the cooling light was held on for a further 1.5 ms, while the repumping light was shut off 0.5 ms before the end of the ramp to allow passive optical pumping of the atoms into the $(1, -1)$ state.

Figure 5.20 shows 30 ms of thermometry on the resultant cloud. Clearly the evidence of a bimodal distribution in the thermometry has been greatly reduced, with the effect only seen at very short drop times. The final temperature of the bulk of the atoms is approximately $30 \mu\text{K}$ - $40 \mu\text{K}$, a factor of three below the Doppler temperature for rubidium. This results in a final cloud of $4 - 5 \times 10^9$ atoms, with a horizontal and vertical width of 1.3 mm and 0.68 mm, corresponding to a trap volume of approximately 4.3 mm^3 , an average density of 1×10^{12} atoms per cm^3 . This is the final stage of the laser cooling of our atoms before they will be confined in a purely magnetic trap, in which forced evaporative cooling of the atoms will take place to further increase the PSD at the cost of a reduction in atom number, before finally being loaded into a hybrid magnetic and optical trap for the evaporation to BEC.

Finally, figure 5.21 shows a trace of the control voltage of the quadrupole coils Magneato, the RF frequency applied to the Science repump AOM and the RF frequency

applied to the Science trap AOM (top to bottom), during the entire laser cooling process, generated by Runviewer⁵⁷. The three stages shown correspond to the end of the MOT load⁵⁸, the compression of the MOT and polarisation gradient cooling.

⁵⁷Runviewer is software package part of the Labscript suite, that shows the user how various properties change throughout an experiment, such as analog control lines, RF frequencies, amplitudes etc.

⁵⁸Note that no parameters change during the MOT load, so it was not deemed necessary to show this entire period.

Evaporative cooling of atoms

To increase the PSD of an atomic cloud to a large enough value for Bose–Einstein condensation to occur, laser cooling of atoms is not enough as this typically only achieves PSD’s on the order of 10^{-5} , five orders of magnitude less than that required for BEC. This limitation in PSD occurs due to the limited temperatures that can be achieved through laser cooling, and the limit on the maximal density that can be achieved due to the repulsive force between atoms that occurs once the optical density of the cloud becomes large. The second caveat that occurs with laser cooling of atoms is a loss mechanism induced by the presence of atoms in the excited state, which leads to inelastic scattering between the atoms in the excited and ground state, causing heating of the sample.

To further increase the PSD of the atomic cloud, one must thus avoid using resonant laser light to avoid the presence of atoms in the excited state, limiting this loss mechanism and limiting the repulsive interactions between the atoms. There are two traps that are typically used to avoid this; either pure magnetic traps, where a spatially varying magnetic field with a field minimum is used to trap low field seeking atomic states, or far off resonant dipole traps, where the atoms are trapped by the dipole force and far detuned light is used to minimise heating in the trap.

Once the atoms are trapped by one of these methods, forced evaporative cooling can be used to greatly enhance the PSD at the cost of losing atoms. Forced evaporative cooling is the process of deliberately lowering the depth of a trap such that atoms with an energy greater than the mean energy of the sample can escape, thus taking away a portion of energy greater than their ‘fair share’ and lowering the overall

average energy of the sample. Evaporative cooling occurs in a wide range of everyday situations, the most obvious being the evaporative cooling of coffee in your mug as you blow on it, or the evaporation of water on a surface or sweat from your body. In the following sections I will describe the basic theory behind evaporative cooling, the rates at which cooling occurs and how its efficiency can be maximised, and then consider two examples of forced evaporation, in a magnetic trap as well as an optical dipole trap.

6.1 Forced evaporative cooling

There are a number of models that can be used to analyse the effects of forced evaporative cooling on a trapped atomic sample, however, we shall only consider the simplest method in which the trap parameters are changed in a single step, and the thermodynamic properties of the trap afterwards are then calculated, following the treatment by Ketterle and Van Druten[7]. Although evaporative cooling is typically a continuous process, whereby the trap depth would be continually lowered and the rethermalisation process constantly occurs, this simple model will provide enough insight into the working of forced evaporative cooling to optimise the procedure for achieving condensation.

There are a number of assumptions that are critical to this model, and which are typically upheld in experiments, but they are important to keep in mind;

1. The distribution of atoms in phase space depends solely upon the energy and nature of the trap
2. The cloud is far from the BEC phase transition point
3. The rate at which rethermalisation in the trap occurs is greater than the cooling rate
4. Atoms that escape the trap do so without colliding or exchanging energy with atoms within the trap
5. The ratio of inelastic collisions to elastic collisions in the trap is small

The first and second assumption are required so that we can compute the thermodynamic properties of the trapped atoms with knowledge solely of the trap properties and the energy of the thermal gas, using classical mechanics. The third assumption is to ensure that the sample remains thermalised, and thus the properties of

the trapped atoms are those of a thermal gas, this is typically the case as long as the rate at which the depth of the trap is reduced is not too great. Finally, the fourth and fifth assumptions are required to ensure there are no additional heating mechanisms within the trap which must be considered.

Using these assumptions, we will consider the general case of a thermal gas of atoms trapped in a power law potential given by;

$$U(x, y, z) = \epsilon_1 \left| \frac{x}{a_1} \right|^{s_1} + \epsilon_2 \left| \frac{y}{a_2} \right|^{s_2} + \epsilon_3 \left| \frac{z}{a_3} \right|^{s_3} \quad (6.1)$$

where ϵ_i , a_i and s_i are the relative amplitude, characteristic length and power of the trap in the i^{th} direction. From this potential, it can be shown[113] that the volume of a thermal gas scales with the temperature as

$$V \propto T^\xi \quad (6.2)$$

where ξ is given by

$$\xi = \frac{1}{s_1} + \frac{1}{s_2} + \frac{1}{s_3} \quad (6.3)$$

In the case of a crossed dipole trap where the potential is harmonic, $s_i = 2$ and thus $V \propto T^{\frac{3}{2}}$, whereas in the case of a quadrupole trap where the potential is linear, $s_i = 1$ and thus $V \propto T^3$.

The simple model of forced evaporative cooling for a thermal gas trapped in the above power law potential will follow three steps:

1. We assume a trap containing N atoms, thermalised at a temperature T , in an infinitely deep potential
2. The trap depth is reduced to $\eta k_B T$ and the sample is allowed to thermalise through collisions
3. The properties of the system, such as PSD, T , V and N are then computed

With this model we can characterise the efficiency of the forced evaporation in cooling the sample. Here we have introduced what is typically known as the truncation parameter, η , which is the fraction of the mean energy of the thermal gas that the trap is lowered to.¹

¹Note that η is always greater than 1 for forced evaporation to lower the resultant temperature of the sample, as if $\eta \leq 1$ then the mean energy of the atoms removed from the trap is lower than the average energy of atoms in the trap.

Using this model there are two parameters that are useful to identify and that can be used to determine all other thermodynamic properties of the evaporated system, these are

$$\nu = \frac{N'}{N} \quad (6.4)$$

the fraction of atoms remaining after the system has thermalised, where N' is the final number of atoms, and

$$\alpha = \frac{\ln(\frac{T'}{T})}{\ln(\frac{N'}{N})} \quad (6.5)$$

where α is a parameter characterising the efficiency of the evaporative cooling process² and where T' is the new temperature after thermalisation. Note that the efficiency of the process, α , and the fraction of atoms that remain after truncating the distribution, ν , both depend critically on how far the distribution is being truncated, η , the strength of the potential in which the atoms are confined as well as the ratio of good to bad collisions.

It should be noted that there is also a more intuitive understanding of the meaning of α . From the definition of α , it is clear that it is dimensionless and in reference [7] Van druten and Ketterle show that the parameter characterises how much more than the average energy of the distribution a particle evaporates, thus larger values of α correspond to more energy being removed from the system with smaller particle loss.

Using these parameters it can be seen that the temperature of the gas scales as

$$T' = T\nu^\alpha. \quad (6.6)$$

The fractional reduction in temperature thus has a power law dependence on α caused by the loss of evaporated atoms, thus maximising the value of α increases the efficiency of the evaporation. Similarly we can then calculate how other thermodynamic properties of the gas change after this truncation, such as how the volume of the trap scales by combining equation 6.6 with equation 6.2 to yield

$$V' = V\nu^{\xi\alpha} \quad (6.7)$$

²Note that α is defined in terms of the logarithmic ratios of the change in T and N , as evaporative cooling occurs on an exponential time scale.

Thermodynamic quantity	\propto
Number, N'	$N\nu$
Temperature, T'	$T\nu^\alpha$
Volume, V'	$V\nu^{\alpha\xi}$
Number density, n'	$n\nu^{1-\alpha\xi}$
Phase space density, ρ'	$\rho\nu^{1-\alpha(\xi+\frac{3}{2})}$
Collision rate, k'	$k\nu^{1-\alpha(\xi-\frac{1}{2})}$

Table 6.1: Dependencies of different thermodynamic quantities on the efficiency parameter α .

the trap density changes, using the relation $n = \frac{N}{V}$ and equation. 6.4 and equation 6.7

$$n' = \frac{N'}{V'} = \frac{N\nu}{V\nu^{\xi\alpha}} = n\nu^{1-\xi\alpha}, \quad (6.8)$$

changes to the phase space density, given by $\rho \propto nT^{-3/2}$, using equations 6.6 and 6.8

$$\rho' = n'T'^{-3/2} = n\nu^{1-\xi\alpha}T^{-3/2}\nu^{-3\alpha/2} = \rho\nu^{1-\alpha(\xi+3/2)} \quad (6.9)$$

and finally the elastic collision rate, given by $k \propto nT^{1/2}$, using equations 6.6 and 6.8

$$k' = n'T'^{1/2} = n\nu^{1-\xi\alpha}T\nu^{\alpha/2} = k\nu^{1-\alpha(\xi-1/2)}. \quad (6.10)$$

The change in the thermodynamic properties of the gas above is given in table 6.1[113], where the left hand column indicates the new value of the quantity, X' , and the right hand column expresses this quantity in terms of its value before evaporation, X , and the quantities ν and α , as well as the strength of the potential, ξ . From this table it is clear that for a given truncation parameter³ for a given potential, the scaling of all quantities is dependent solely upon α making it the single critical parameter for the efficiency of the process.

Given table 6.1 we can see how the thermodynamic properties of a trapped gas change during a single step of forced evaporation if the trap potential and thus ξ is known, as well as the fraction of atoms removed during the process, ν , and the efficiency of the process, α . It is now left to calculate how α and ν depend upon the truncation parameter and the trap characteristics.

To calculate the effect that a given truncation parameter, η , will have on the number

³Which determines the change in temperature, and thus determines both ν and α .

and temperature of the atoms, one performs the following steps:

1. Calculate the density of states of the distribution of atoms in the presence of the confining potential
2. Calculate the fraction of atoms remaining in this trap after reducing the trap depth, by integrating the density of states multiplied by the Boltzmann distribution from 0 up to the truncated trap depth
3. Calculate the average energy of the thermal gas after reducing the trap depth by integrating the density of states multiplied by the Boltzmann distribution and the thermal energy of the initial sample from 0 up to the truncated depth, $\eta k_B T$

Clearly step two gives the ratio of atoms left after the truncation, whilst step 3 yields the final temperature of the atoms. This model, however, has assumed no dependence upon time, i.e. we have ignored any ‘bad collisions’ which limit the time over which evaporation can occur. To consider the how evaporation in an experiment proceeds, we must introduce a dependence on time. In the case of evaporation the relevant quantity is the ratio of good to bad collisions, R , where good collisions are elastic collisions, which rethermalise the sample, whereas bad collisions are collisions that cause trap loss or heating, such as collisions with background gas particles etc.

6.1.1 Cooling rates during forced evaporation

There is no theoretical limit to the cooling efficiency of evaporative cooling. If one considers the case in which the truncation parameter was infinite, and the system was allowed to equilibrate for an infinite amount of time, then eventually a single atom would carry away the entire energy of the system, cooling the remaining atoms to absolute zero. This is obviously not practical in a laboratory setting as the time over which forced evaporative cooling occurs is an important parameter, typically determined by the characteristic time scales of loss mechanisms. If we consider the case in which the truncation parameter is rapidly reduced the gas won’t have time to rethermalise invalidating our third assumption and reducing the efficiency of the evaporation as time progresses, whereas in the case when it is reduced too slowly, the loss of atoms from inelastic collisions will reduce the efficiency of the evaporation. Thus the rate at which evaporative cooling can occur whilst maintaining efficient evaporation is set by the rate at which rethermalisation of the gas occurs.

It has been theoretically[129] and experimentally[130] shown that ≈ 2.7 elastic collisions between the atoms are required to rethermalise the gas once the distribution has been truncated, however, theoretical modelling based upon the Boltzmann equation[131] indicates that the restoration of the high energy tail of the distribution, the tail from which we remove the hottest atoms, requires many more collisions, on the order of ≈ 60 . From the above the rate of rethermalisation can be estimated as simply $\frac{1}{3\tau_{elastic}}$, where $\tau_{elastic}$ is the elastic collision time.

A more thorough analysis utilises the principle of detailed balance[7] which states that in the high energy tail of a thermalised distribution, the rate at which elastic collisions removes particles from the high energy tail is equal to the rate at which elastic collisions between particles in the low energy tail produce particles in the high energy tail⁴. In a thermalised distribution, almost every elastic collision that involves an atom in the high energy tail removes it from the high energy tail, thus the rate at which atoms are produced with an energy greater than $\eta k_B T$ is simply given by the number of atoms with energy greater than $\eta k_B T$ divided by the elastic collision time,

$$\frac{dN}{dt} = -N f(\epsilon > \eta) k_{el} = n\sigma\bar{v}\eta e^{-\eta} N = \frac{-N}{\tau_{evap}}, \quad (6.11)$$

i.e. the rate of change in the number of atoms, $\frac{dN}{dt}$, is given by the product of the number of atoms and the fraction of atoms with energy above the truncation energy, multiplied by the elastic scattering collision rate, as this is simply the number of atoms that are produced with energies greater than $\eta k_B T$ and thus is the number of atoms evaporated per second⁵. Here all parameters have their usual definition, with σ being the elastic scattering cross section and \bar{v} the average velocity of an atom in the gas, and we have introduced the evaporation time constant, $\tau_{evap} = (n\sigma\bar{v}\eta e^{-\eta})^{-1}$, which is simply the decay constant of the number of atoms when loss is due solely to evaporation and bad collisions are not considered. From the above we can see that with a constant truncation parameter, η , the number of atoms in the trap will decay exponentially with a time constant given by τ_{evap} .

If we compare the above rate required to produce an atom to be evaporated, compared to the elastic collision time, we find[113]

$$\frac{\tau_{evap}}{\tau_{elastic}} = \frac{\sqrt{2}e^\eta}{\eta} \quad (6.12)$$

⁴This is clear as the distribution would not remain thermalised otherwise.

⁵Note that we have used the average velocity of an atom with thermal energy $\eta k_B T$ to calculate \bar{v} , not the average velocity of an atom in the distribution, as to remove a particle from the high energy tail at least one of the particles must have energy greater than $\eta k_B T$.

From the above we can see that this ratio depends exponentially on the truncation parameter, which is greater than one for cooling and thus indicates that the ratio of the evaporation time to the elastic collision time decreases almost exponentially as η decreases. This is relatively intuitive to understand, as when η decreases, a larger fraction of atoms are removed from the distribution, and thus there are a greater number of atoms with energies close to the truncation energy, thus the rate at which atoms are ‘evaporated’ increases as atoms only need to gain a small fraction of energy before they will have an energy greater than the truncation energy. However, this poses the question of how fast do we evaporate to maximise the PSD, as ideally we wish to minimise the ratio of the evaporation time and the elastic collision time as this will allow the sample to thermalise faster, but in doing so we require a smaller truncation parameter which in turn cuts away more of the distribution and leads to less efficient evaporation.

6.1.2 Maximising PSD

The aim of any experimental apparatus designed to produce BECs is to reach ‘the biggest’ BEC, the condensate with the largest number of atoms. As this is the aim, it follows that maximising the quantity γ_{tot} is the key aim of any experiment, where

$$\gamma_{tot} = -\frac{\ln(\rho_f/\rho_i)}{\ln(N_f/N_i)}, \quad (6.13)$$

and all parameters carry their usual meaning. This global parameter characterises the cost of increasing the PSD towards achieving condensation, in terms of the loss of atoms. This parameter is a more useful parameter for general considerations of forced evaporation than the loss of atoms for a reduction in temperature⁶, as it also accounts for how the trap volume depends upon the strength of the trap, and thus can be used to compare forced evaporation in different potentials.

In reference [7] Ketterle and Van Druten show that if the evaporation process is split into a number of smaller steps, then γ_{tot} is maximised if γ_i ⁷ is maximised at each step of the evaporation. The derivation is not particularly instructive or insightful and so is not repeated here, however, the arguments relies on the fact that maximisation of γ_i relies on maximisation of the initial ratio of good to bad collisions, R , which

⁶The parameter α , however, α is a simpler parameter to compute when the efficiency of evaporation is compared within the same potential.

⁷where γ_i is the efficiency of the i^{th} step.

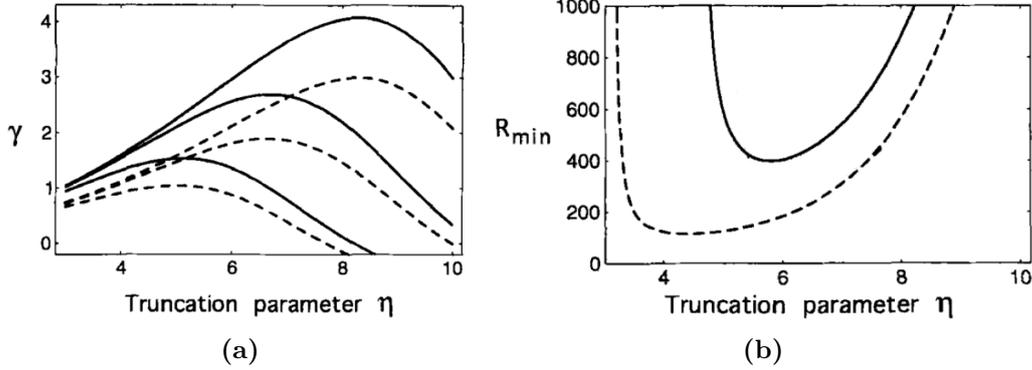


Figure 6.1: Forced evaporation efficiency and runaway evaporation dependence on the potential and ratio of good to bad collisions. **a)** shows the dependence of the efficiency of the evaporation, characterised by γ , upon the depth at which the distribution is truncated, $\eta k_B T$ for a 3D harmonic potential (solid lines) and a 3D linear potential (dashed lines) for 3 different ratios of good to bad collision, R , 5000, 1000 and 200 respectively from the upper to lower sets of lines. From comparison of the curves, at a given value of R , the peak efficiency in the parabolic trap is always greater than in the linear trap, and peak efficiency increases with increasing R , as expected. **b)** shows how the minimum ratio of good to bad collisions, R , required to achieve runaway evaporation varies as a function of the truncation parameter for a parabolic potential (solid lines) and a linear potential (dotted lines). **a)** and **b)** were both altered from Reference. [7].

in turn is maximised by having the largest number of atoms, N . Thus if we split our trajectory into a number of small changes in phase space density, $\Delta\rho$, then if we maximise γ_i for the first step, this will maximise the number of atoms at the given PSD after the first step, which in turn maximises R , and thus creates the optimal initial conditions for γ_{i+1} .

Experimentally this is very convenient as it provides a single metric to be optimised for each step of the forced evaporation, with a guarantee that this will also optimise the entire trajectory. In particular, it implies that to optimise γ_{tot} over a particular trajectory, first that trajectory should be broken into N steps, then the evaporation conditions for the first step should be changed to maximise γ_2 , and then the evaporation conditions for the second step should be changed to maximise γ_3 and so on, always optimising γ_{N+1} by changing the evaporation conditions of the N^{th} step.

Figure 6.1a shows how the evaporation efficiency in two different potentials, a 3D parabolic potential (solid lines) and a 3D linear potential (dashed lines) depends upon the truncation parameter for three different ratio's of good to bad collisions, 5000, 1000 and 200 respectively for the uppermost line to the lowest, respectively. From comparison within a given potential, it is clear that increasing the ratio, R , significantly increases the peak efficiency, with the peak efficiency increasing by ≈ 1

with each factor of 5 increase in the ratio R . It is also clear that as the ratio, R , increases, the truncation parameter at which the peak efficiency occurs increases. This is expected, as at larger values of R the distribution will rethermalise quicker and thus restore the high energy tail of the distribution faster thus increasing the evaporation rate.

By comparing the curves for the parabolic potential with the linear potential for a given R , it is clear also that the parabolic potential consistently outperforms the linear potential. However, it should be noted that as the collision rate typically decreases in a parabolic potential as evaporation proceeds, the efficiency will typically decrease as evaporation proceeds, whereas this is not the case in a linear potential due to the increased compression effect from the potential.

It should also be noted that as γ is the exponent of the fractional change in phase space density for a fractional change in number of atoms, small changes in γ lead to large changes in the experiment. For example, consider a change of γ from 3 to 2 over the entire trajectory from magnetic trap to BEC, requiring an increase in PSD of 10^6 . Now, as we know $\gamma = -\frac{\log(D'/D)}{\log(N'/N)}$, we can thus rearrange to yield

$$\frac{N'}{N} = \left(\frac{D'}{D}\right)^{-1/\gamma}. \quad (6.14)$$

If we then substitute in $D'/D = 10^6$ and $\gamma = 2$ and 3 , respectively, we find N'/N equal to 0.001 and 0.01, showing that an increase in γ by 1 leads to an increase in the number of atoms in the condensate by an order of magnitude.

6.1.3 Runaway evaporation

The most efficient way of reaching the largest increase in phase space density is to have the trapped atoms undergo runaway evaporation, where the elastic scattering rate of the sample either remains constant or increases in time. It has been shown that for a given truncation parameter, there is a minimum ratio of good to bad collisions, R_{min} , required to achieve runaway evaporation, and that this value depends upon the potential.[7]

Figure 6.1b shows how the dependence of R_{min} varies as a function of the truncation parameter for a parabolic potential (solid line) and a linear potential (dashed line). It is clear that runaway evaporation in a linear potential is much easier to achieve, with the minimum ratio being much lower than in the case of a parabolic potential

for all values of the truncation parameter considered. This is due to the increased compression of the trapped gas in a linear potential as opposed to quadratic, which further increases the elastic collision rate.⁸ This dependence of R_{min} on the trap potential is the reason why runaway evaporation is difficult to achieve in a harmonic potential, as larger values of R are required.

It should be noted that evaporation can still proceed in a potential even if runaway evaporation does not proceed, however, this will cause R to decrease during evaporation, which will further reduce the efficiency of evaporation until eventually no further increase in PSD can be achieved and evaporation stagnates.

6.2 Forced evaporation in a magnetic trap

We will now consider the case of forced evaporation in a magnetic trap, the first method used for achieving Bose–Einstein condensation[132]. To accurately describe this process we must first discuss how atoms are trapped in a spatially varying magnetic field, and the mechanism we can use to force evaporation.

6.2.1 Trapping in a magnetic field

Although magnetic trapping of ions was first demonstrated as early as 1953[133], magnetic trapping of neutral atoms was only first achieved in 1985[134]. The limitation preventing neutral atoms being trapped was due to the relatively small trap depths of magnetic traps for neutral atoms, due to the much smaller magnetic dipole moment compared to the electric dipole moment of ions, requiring mK temperatures to be achieved before the atoms could be trapped. The force, \vec{F}_B , on a neutral atom in a magnetic field arises from the interaction between the inhomogeneity of the magnetic field, \vec{B} , and the atomic magnetic moment, $\vec{\mu}$ given by

$$\vec{F}_B = \nabla(-\vec{\mu} \cdot \vec{B}) \quad (6.15)$$

In the case where the velocity of the atoms is low enough that the atomic magnetic moment can follow the changes in the magnetic field direction, this can be reduced to

$$F_B = -\mu_B g_f m_f \nabla B \quad (6.16)$$

⁸Thus it is easier to achieve runaway evaporation in a linear potential than a quadratic potential.

where μ_B is the Bohr magneton and g_f and m_f are the g-factor and projection of the magnetic dipole moment of the atom upon the quantisation axis, respectively. As Earnshaw's theorem applied to Maxwell's equations forbid the existence of a local maxima in a static magnetic field, only magnetic field minima are thus allowed, and thus the product g_fm_f must be positive for the states to be trapped.

A common magnetic field used for trapping neutral atoms, originally suggested by W. Paul[133], is a quadrupole field, formed by two identical coaxial coils ideally separated by $1.25\times$ their radius, carrying current in opposite directions. The magnitude of the magnetic field is given by[135]

$$\vec{B} = B_x\vec{x} + B_y\vec{y} + B_z\vec{z} \quad (6.17)$$

where $\vec{x}, \vec{y}, \vec{z}$ are unit vectors pointing along the usual three Cartesian coordinates, and B_x, B_y, B_z are the magnetic field gradients in the x, y and z directions given by $-\frac{a}{2}, -\frac{a}{2}$ and a respectively, where a is the field gradient at $(x, y, z) = (0, 0, 0)$, which is also clearly a magnetic field minimum. From the form of equation 6.17 it is clear that the confining potential for atoms in the trap is not harmonic but rather is linear and is also not central, due to the additional factor of 2 in the gradient in the z direction, and thus does also not preserve angular momentum.

6.2.2 Majorana losses

An important consideration in the magnetic trapping of atoms is the ability of the magnetic moment of the atoms to remain aligned with the magnetic field, as if this is not the case the atom will most likely be ejected from the trap due to a process known as a Majorana spin-flip loss. Majorana spin-flip losses are very detrimental to cold samples of atoms trapped in quadrupole field, and were one of the key issues that had to be navigated to achieve Bose-Einstein condensation. These spin flips are most likely to occur near the field zero due to the rapid change in direction of the magnetic field along any line that crosses through the zero, and as such are most likely to occur for atoms that spend the most time near the field zero, the coldest atoms. As the coldest atoms are preferentially lost from the trap, Majorana losses thus not only present a loss mechanism for the atoms, but a heating mechanism also.

The rate of Majorana losses from a quadrupole trap can be estimated as[8]

$$\Gamma_{Maj} = 1.85 \frac{\hbar}{m} \left(\frac{m_f g_f \mu_f B'_q}{k_B T} \right)^2 \quad (6.18)$$

where m is the mass of the atom, B'_q is the magnetic field gradient, g_f is the g factor of the state and m_f is the quantum magnetic number of the state. From the above it is clear that the Majorana loss rate is largest for low temperatures and large magnetic field gradients, with the lifetime of the magnetic trap on the order of a second for a 20 μ K cloud in a 140 Gcm⁻¹ magnetic field[8]. Given the dependence of equation 6.18 on the magnetic field gradient and temperature, the Majorana loss rate can be somewhat decreased through adiabatic expansion of the trap as the temperature scales with $B_q^{2/3}$ and thus overall B'_q/T decreases as B'_q decreases, however, this comes at the cost of a reduction in the elastic collision rate.

There are two key methods that have been used to reduce Majorana losses in a quadrupole trap, either plugging the ‘hole’ with a optical beam, or alternatively time varying the magnetic field potential so as to effectively remove the field zero. A time varying magnetic field was the first solution to alleviate this loss, creating what is known as a time orbiting potential (TOP) trap[136], and used to create the first BEC[132]. In the TOP trap, a strong bias fields in the x and y direction which vary sinusoidally in time and out of phase with one another, are used to make the magnetic field zero orbit about its original position at kilohertz rates. As the trapping frequency of a magnetic trap is typically on the order of 10 Hz to 100 Hz, the atoms do not see this time varying potential, but rather the time average of the potential which creates a finite magnetic field at the original position of the field zero, removing it. A similar approach to achieve condensation in a purely magnetic trap is the QUIC trap[137], which combines the standard quadrupole field with an additional Ioffe coil and traps atoms at a finite bias field. In the case of an optical plug used to reduce Majorana losses from a quadrupole trap, either a focused blue detuned beam can be used to force atoms away from the field zero, or alternatively a red detuned beam can be placed approximately one waist below the field zero to trap the coldest atoms before they undergo a spin flip.

6.2.3 Radiofrequency evaporation in a magnetic trap

With the basics of both forced evaporation and magnetic trapping discussed, we can now consider forced evaporation in a quadrupole magnetic field trap. As already

Thermodynamic quantity	\propto
Number, N'	$N\nu$
Temperature, T'	$T\nu^\alpha$
T	$B'^{2/3}$
Volume, V'	$V\nu^{3\alpha}$
V	B'^{-1}
Number density, n'	$n\nu^{1-3\alpha}$
n	B'
Phase space density, ρ'	$\rho\nu^{1-\alpha(\frac{9}{2})}$
Collision rate, k'	$k\nu^{1-\alpha(\frac{5}{2})}$
k	$B'^{\frac{4}{3}}$

Table 6.2: Dependencies of different thermodynamic quantities on the efficiency parameter α during forced evaporation in a quadrupole trap, as well as the dependency of these parameters on the field gradient, B'

mentioned the potential of a quadrupole magnetic field is linear in position, thus if we consider the thermodynamic properties dependence upon the quantities γ and ν we result in table 6.2. Here I have included the dependence of the same quantities on the magnetic field gradient when it is changed adiabatically, as this is an additional parameter we can control. Note that the phase space density remains constant during adiabatic changes of the trap.

Forced evaporation requires the trap depth to be lowered to a point, $\eta k_B T$, such that our Boltzmann distribution gets truncated at that energy and the hottest atoms escape. This can be achieved in a magnetic trap via two methods, lowering the gradient of the magnetic trap, which in turn lowers the trap potential, or by introducing a radio frequency field that is resonant with the energy level splitting between a trapped and untrapped Zeeman state at a particular magnetic field value.

From the table above, the first method of reducing the magnetic field gradient to force evaporation is undesirable as is clear from the dependence of the elastic scattering collision rate on the magnetic field gradient. If the trap depth were to be lowered the rate of rethermalisation would also lower, stagnating the forced evaporation process and reducing the ratio of good to bad collision making runaway evaporation difficult, and reducing the efficiency of evaporation.

The second method relies upon the energy shift of the Zeeman sublevels of the atom due to the presence of the external magnetic field. To elucidate this we will consider an atom moving in one dimension in the presence of a quadrupole field of gradient

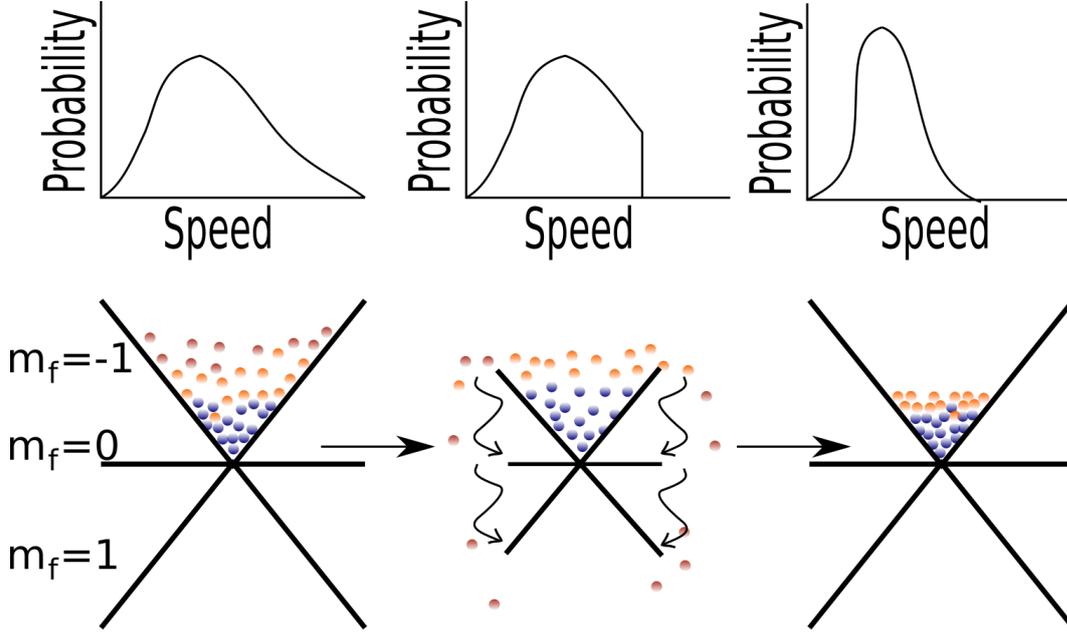


Figure 6.2: Graphical representation of RF evaporation. A graphical representation of how RF evaporation occurs in a quadrupole potential. The top part of the figure shows how the velocity distribution changes through the three times considered, before evaporation, during evaporation and after thermalisation, whilst below a graphical interpretation of the process is shown. In the representation atoms are shown as coloured circles with the colour indicating the kinetic energy the atom has, and the potential is drawn as a black curve, with the potential seen by the $(F = 1, m_f = -1)$, $(F = 1, m_f = 0)$ and $(F = 1, m_f = 1)$ indicated.

B'_q , in the $m_f = -1$ state. At a position, x , the internal energy of the $m_f = -1$ state relative to the $m_f = 0$ state will be shifted by

$$\Delta E(x) = g_f \mu_B B'_q x \quad (6.19)$$

thus, if an RF field is applied such that $\hbar\omega_{RF} = \Delta E(x)$, the radiofrequency field will selectively drive an RF transition between the Zeeman sublevels into an untrapped state, thus causing the atom to be removed from the trap.

Thus the use of the RF knife allows the selective removal of atoms that cross through the position, x , at which the energy of an RF photon is equal to the Zeeman splitting of the atomic substates. In three dimensions this idea is easily extended from a single point to the equipotential surface of constant magnetic field strength, which in a quadrupole field is an ellipsoid centred about the field zero. If the strength of the coupling of the RF field to the atoms is large enough, then any atoms that have a large enough kinetic energy to pierce this surface will undergo a spin flip and be

ejected from the trap.

A representation of this process is shown in figure 6.2, this figure shows graphically what is occurring to the atoms in the potential during evaporation (bottom), as well as how the velocity distribution of the trapped clouds changes during this process (top). In the graphical representation of what occurs to the atoms during evaporation, the black curve in which the atoms are confined represents the potential the atoms in the ($F = 1, m_f = -1$) see, whereas the black curve directly below is indicative of the potential seen by atoms in the anti-trapped ($F = 1, m_f = 1$) state. On the left hand side of the figure the initial distribution of atoms in the magnetic trap is shown, with the colour of the atoms representing the kinetic energy the atom has, and thus how far up the potential the atom can climb. The next part of the figure(center) shows what occurs once the RF coupling is turned on. This effectively truncates the height of the potential as it couples atoms in the ($F = 1, m_f = -1$) state to the ($F = 1, m_f = 1$) state if the atoms have a large enough velocity to reach the position at which the difference in the potential energy of the two states is equal to twice the energy of an RF photon. During this stage the velocity distribution shows a clear truncation, with the probability of an atom having a velocity greater than some particular value, set by the RF frequency, being zero. In the final part of the figure, the trap is left to thermalise over some time, and the resulting velocity distribution is now peaked at a lower velocity, hence cooling of the sample has occurred.

We can equate the RF knife with a truncation temperature by using the equipartition theorem. By equating the thermal energy of the atom with the potential energy due to the magnetic field⁹, which is equal to the energy of an RF photon and we thus have

$$\frac{3}{2}k_B T = \hbar\omega_{RF}. \quad (6.20)$$

Rearranging equation 6.20 and evaluating the constants leads to

$$\omega_{RF} \approx \frac{T}{5} \quad (6.21)$$

$$\therefore f_{RF} \approx \frac{T}{30} \quad (6.22)$$

where f_{RF} , the frequency of the RF field, is expressed in MHz, and T in μK . This relation is useful in determining evaporation parameters such as the truncation energy

⁹Here we have assumed that the atom has spent all its kinetic energy ‘climbing’ the potential hill.

for a given applied RF field.

6.2.4 Adiabatic compression of a magnetic trap

Adiabatic compression of atoms confined to a potential, a process where heat is not transferred between the atoms and their external environment, leads to changes in temperature and volume of the gas without changing the phase space density. In the case of a magnetic trap, the gradient of the magnetic field determines the volume of the trap, which scales inversely to the gradient. As such, increasing the magnetic field gradient over an appropriate time will decrease the volume of the trap, thus increasing the density and the temperature. As both the temperature and density of the gas increases, the elastic scattering rate also increases, and as such adiabatic compression will increase the ratio of good to bad collisions, R .

The only time this is not the case is when three body recombination loss, which is dependent upon the density of the gas, starts to become appreciable, or alternatively when the rate at which atoms are evaporated from the trap approaches the rate at which rethermalisation occurs, thus limiting the rate of evaporation. Thus evaporative cooling in a magnetic trap should occur in the tightest possible confinement before either three body combination sets in, when the rate of evaporation approaches the rate of rethermalisation or when Majorana losses become a significant loss mechanism.

6.3 Forced evaporation in an optical dipole trap

Quadrupole magnetic fields have a severe limitation when attempting to achieve BEC due to the significant Majorana spin-flip losses that occur in low temperature clouds. One method of reducing this problem is the use of a far off resonant red detuned optical dipole beam that forms a deep optical dipole trap located one waist below the field zero, moving the potential minimum to below the magnetic field zero. The atoms are then loaded into the hybrid magnetic/dipole potential by reducing the field gradient to a value below that which compensates the force of gravity, ‘dropping’ the atoms into the new potential. Alternatively, a pure optical dipole trap can be loaded in this same fashion, or directly from a MOT, but in this case the atoms are typically confined in the potential minima formed at the crossing of two dipole traps, providing trapping in all three dimensions. In this section we will

outline how the dipole force can be used to strongly confine atoms, and how the trap parameters can be used to force evaporation.

6.3.1 The optical dipole force

Neutral atoms have permanent magnetic dipole moments that allow them to interact with an inhomogeneous static magnetic field, experiencing a force. Neutral atoms do not, however, have a permanent electric dipole moment, as this is forbidden by the inversion symmetry of the atomic wavefunction. Thus to interact with an electric field an electric dipole moment must first be induced. Although inducing electric dipole moments can be achieved through the use of an electrostatic field, this is typically not the method used for alkali atoms, but rather somewhat near resonant optical fields are typically used.

When an atom is placed in an oscillating electric field, such as the case when the atom is exposed to an optical field oscillating at frequency, ω_0 , the oscillating electric field induces an oscillating dipole moment, \vec{p} , proportional to the electric field strength and oscillating at the same frequency. The magnitude of the dipole moment is given by

$$p = \alpha E \quad (6.23)$$

where α now describes the complex polarizability of the atom¹⁰, which is also dependent upon the driving frequency.

The interaction of this induced dipole moment with the external electric field creates an induced potential, known as the dipole potential. The dipole potential has been shown to arise from the dispersive interaction of this induced dipole moment with the intensity gradient of the light field[13] and, as it results in a conservative force, it can be written as the gradient of a light induced potential,

$$U_{dip} = -\frac{1}{2\epsilon_0 c} \text{Re}(\alpha) I, \quad (6.24)$$

where U_{dip} is the dipole potential, I is the intensity of the light field and $\text{Re}(\alpha)$ is the real component of the atomic polarizability.¹¹

Calculations of the atomic polarizability have been performed elsewhere and will

¹⁰ α is chosen to keep this thesis in line with standard notation in the literature, even though it introduces an ambiguity. In this section α shall only refer to the complex polarizability.

¹¹This is arrived at by time-averaging the product of the electric field and the induced dipole moment.

not be repeated here, but the resulting potential and scattering rate calculated by Grimm[13] are

$$U_{dip}(r) = \frac{3\pi c^2}{2\omega_0^3} \frac{\Gamma}{\Delta} I(r) \quad (6.25)$$

and

$$\Gamma_{sc}(r) = \frac{3\pi c^2}{2\hbar\omega_0^3} \left(\frac{\Gamma}{\Delta}\right)^2 I(r) \quad (6.26)$$

where Γ is the on-resonance damping rate given by $\Gamma = \frac{e^2\omega_0^2}{6\pi\epsilon_0 m_e c^3}$, Δ is the detuning of the driving field from resonance and Γ_{sc} is the scattering rate.

There are two important points to note from equation 6.25 and equation 6.26, the first is that as the dipole potential depends upon the sign of the detuning, the dipole force can be made to be attractive or repulsive. In the case of a red detuned optical field¹², the sign of the detuning is negative, and thus atoms will be attracted to regions of high intensity. In the case of a blue detuned optical potential¹³ the reverse is true, the resulting potential is positive and thus the light acts as a barrier to the atoms, with atoms being repulsed away from regions of high intensity.

The second point of interest is the dependence of the potential and the scattering rate upon the intensity and the detuning. As the potential is proportional to I/Δ , while the scattering rate is proportional to I/Δ^2 , the scattering rate can be made quite low while maintaining a given depth by increasing the intensity whilst detuning the light further from resonance.¹⁴

The simplest trap that can be imagined that utilises the optical dipole force is a single focussed Gaussian laser beam detuned by Δ below resonance, whose intensity varies as a function of position at the waist, w_0 of the beam as

$$I(r) = I_0 e^{-2\frac{r^2}{w_0^2}} \quad (6.27)$$

where $I_0 = \frac{2P}{\pi w_0^2}$ is the peak intensity of a laser beam of power P . Unless the laser beam is very tightly focussed, the intensity typically varies slowly along the direction in which it focuses¹⁵, intensity gradients along this axis are thus typically

¹²An optical field detuned below atomic resonance.

¹³with a frequency greater than the optical resonance.

¹⁴e.g. if the intensity of the light field was to be double as was the detuning, the potential would remain constant whilst the scattering rate would reduce by a factor of 2.

¹⁵The Rayleigh range can be used to approximate a Gaussian width along this dimension and is given by $z_R = \frac{\pi w_0^2}{\lambda}$, thus the ratio of the Rayleigh range to the waist is of order $\frac{\pi w_0}{\lambda}$ and thus only approaches unity for $w_0 \approx \lambda/2$.

small compared to the transverse dimension.

Given the detuning below resonance of this optical field, ground state atoms will thus experience a force attracting them towards the region of highest intensity in this optical field. In the case where the temperature of the atoms is much less than the depth of the trap, the Gaussian can instead be approximated as a Taylor series, truncated to second order¹⁶ yielding

$$I(r) \approx \frac{2P}{\pi w_0^2} \left(1 - \frac{2}{w_0^2} x^2\right). \quad (6.28)$$

In this limit the force provided by the dipole trap will be approximately harmonic.

Substituting equation 6.28 into equation 6.25 yields the potential resulting from a focussed Gaussian beam in the limit that the depth is much greater than the temperature of the atoms, as

$$U_{dip}(r) = \frac{3c^2}{\omega_0^3} \frac{P\Gamma}{w_0^2 \Delta} \left(1 - \frac{2}{w_0^2} x^2\right). \quad (6.29)$$

Note that equation 6.29 can be written in the form $U_0 - \frac{2U_0}{w_0^2} x^2$, where $U_0 = \frac{3c^2}{\omega_0^3} \frac{P\Gamma}{w_0^2 \Delta}$ is the potential depth of the dipole trap.

Now, as the dipole force is conservative we can calculate the force upon the atoms by taking the gradient of the potential by combining equation 6.28, our approximation for the Gaussian intensity profile, with equation 6.25 to yield

$$F_{dip} \approx -\frac{12c^2}{\omega_0^3} \frac{P\Gamma}{w_0^4 \Delta} x. \quad (6.30)$$

From equation 6.30 we can see that this force is proportional to the position but with opposite sign and thus acts as a restoring force, as expected for a harmonic confining potential.

Utilising the form for the confining force of a harmonic potential, $F = -kx$ we can see the spring constant of the trap is given by

$$k_{dip} = \frac{12c^2}{\omega_0^3} \frac{P\Gamma}{w_0^4 \Delta}, \quad (6.31)$$

and utilising the relation between the trap frequency and the spring constant for a

¹⁶Note the first term of the Taylor series vanishes due to the symmetry of the Gaussian about zero.

simple harmonic oscillator we see that

$$\omega_{trap} = \beta \sqrt{\frac{P}{w_0^4}}, \quad (6.32)$$

where we have incorporated all constants into a single constant, $\beta = \sqrt{\frac{12c^2\Gamma}{\omega^3\Delta m}}$ to make explicit the dependence upon the power and waist of the Gaussian. Through equations 6.29 and 6.32 we can see that although large volume traps are possible through the use of large waisted beams, the power requirements increase very rapidly, and for a similar depth potential the trapping frequency reduces dramatically as a function of size, limiting the rate of rethermalisation.

6.3.2 Forced evaporation in an optical trap

The potential of a crossed optical dipole trap is close to harmonic in all directions, and as such $\xi = \frac{3}{2}$. This is not ideal for runaway evaporation as the ratio of good to bad collisions has to be 3 times larger than in a linear potential, indicating runaway evaporation is difficult, if not impossible, to achieve. However, the benefit of the transition to an optical dipole trap is the removal of the possibility of Majorana spin flip losses, which present a considerable heating at low temperatures, as well as typically larger trapping densities and trapping frequencies, leading to more efficient evaporation.

For completeness, the dependence of the different thermodynamic properties of the trapped gas as a function of ν and γ for the harmonic potential are included in table 6.3.

Forced evaporation in an optical dipole trap typically occurs through lowering the height of the trapping potential, from equation 6.29 we can see a linear dependence of the depth of the potential on the power, for a fixed waist. However, by reducing the power in the laser by a factor x , the trapping frequency is also reduced by a factor \sqrt{x} , as can be seen through equation 6.32 which lowers the elastic collision rate. Thus in the simplest case the reduction in the trap depth is inextricably linked with a reduction in the ratio of good to bad collisions. This reduction then limits the rate of rethermalisation and thus evaporation, which further leads to a reduction in the efficiency of the evaporation, which can eventually lead to stagnation of the evaporation, where further lowering of the potential only results in loss of atoms and not an increase in PSD. As such, table 6.3 also shows the dependence of the

Thermodynamic quantity	\propto
Number, N'	$N\nu$
Temperature, T'	$T\nu^\alpha$
T ,	P
Volume, V'	$V\nu^{\frac{3\alpha}{2}}$
V ,	$P^{3/4}$
Number density, n'	$n\nu^{1-\frac{3\alpha}{2}}$
n ,	$P^{-3/4}$
Phase space density, ρ'	$\rho\nu^{1-3\alpha}$
Collision rate, k'	$k\nu^{1-\alpha}$
k ,	$P^{-1/4}$

Table 6.3: Dependencies of different thermodynamic quantities on the efficiency parameter α during forced evaporation in a crossed dipole trap, as well as the dependency of these parameters on the power in the trap, P

thermodynamic quantities upon the power of the laser.

A number of techniques have been employed to mitigate this reduction of the confinement strength during evaporation, for example an optical dimple trap has been employed[138] where a more tightly focussed dipole trap is centred on the larger volume trap allowing the loading of the smaller trap from the larger volume, and a corresponding increase in confinement. Another method that has been successfully employed is the use of a zoom lens to dynamically change the size of the dipole trap, creating a compressible dipole trap[139][140], allowing the size of the dipole trap to be changed during evaporation, which can result in an increase of the trapping frequency. One final example is the use of a dimple trap inside a larger volume dipole trap, however, in this case the tighter dipole trap is not centred on the larger volume trap, and the larger volume trap is displaced during evaporation[141]. This displacement generates a potential difference that acts to reduce the depth of the tighter trap, resulting in evaporation at approximately constant confinement. All the methods above require additional complexity, however, as they require either additional dipole traps or the ability to dynamically change the dipole trap.

6.4 Experimental considerations of forced evaporation in combined optical and magnetic potentials

So far we have considered how forced evaporation can take place in two different strength potentials, a magnetic quadrupole trap, which is linear in its confinement, and a dipole trap, which is approximately harmonic. In this section we will consider the experimental implications of evaporation in the two differing potentials, their key benefits and their drawbacks.

6.4.1 Magnetic traps

Magnetic traps are particularly well suited to capture atoms from a laser cooled cloud of atoms. This is because magnetic traps are typically on the order of a few mK deep for traps made with large coils¹⁷, and effectively extend all the way to the edge of the coil, making them very large volume traps as well. These two key features mean that close to 100% of atoms in magnetically trappable states from a MOT can be transferred into the magnetic trap.

Forced evaporation in a magnetic trap then typically occurs through the introduction of an RF 'knife', which will remove atoms that are hot enough to reach the position at which the RF knife is resonant with the Zeeman level splitting for the given magnetic field. Thus the frequency of the RF knife determines the truncation parameter for forced evaporation. Due to the linear nature of the trap, $V \propto T^3$, and thus as the temperature of the sample decreases the volume of the trap also decreases rapidly helping maintain, or even increase, the elastic collision rate, making runaway evaporation achievable. This increase in the elastic collision rate of the trapped gas allows for faster evaporative cooling as the gas cools, implying the rate at which the RF knife is increased should increase with increasing time. We saw in figure 6.1b that this can occur because the minimum ratio of good to bad collisions for runaway evaporation to occur in a linear potential is low relative to a harmonic potential, and as such is easier to achieve, meaning forced evaporation in a magnetic trap will typically increase in efficiency as the gas cools, rather than decrease, absent any loss mechanism dependent upon the temperature or density of the gas, which increases the number of bad collisions.

¹⁷As the MOT beams typically have to go through the coils, this typically limits the inner radius of the coils to a few cm, making them 'large'

In figure 6.1a, however, we saw that the peak efficiency of evaporation for a given R was lower in the linear trap than in the parabolic trap. Thus although magnetic traps are good for capturing large numbers of atoms from a MOT, and increasing the scattering rate of the sample during evaporation, they are not as efficient at forced evaporation as a dipole trap for equivalent conditions, and thus typically need an exceptional vacuum to allow long evaporation times to achieve condensation.

The second drawback of the standard magnetic trap employed, a quadrupole trap, is that as the atoms become colder Majorana losses begin to dominate the bad collisions, leading to a decrease in evaporation efficiency and an increase in the heating and loss rate of the trap. This must be mitigated, as condensation is not possible in a trap without having this loss mechanism removed, and this requires additional hardware, either coils or a dipole trap.

6.4.2 Dipole traps

On the other hand, whilst dipole traps are particularly efficient at forced evaporation they are particularly ill suited to capture atoms out of a MOT and can lead to evaporative stagnation. As dipole traps are usually far detuned, commonly detuned 100's of nm away from resonance to reduce the scattering rate and thus increase the lifetime of the trap, they thus require large amounts of power to create only relatively shallow traps¹⁸ This creates a tradeoff between the depth of the trap and its trapping volume for a given power, both of which are typically much smaller than those of the magnetic trap. A second problem that occurs with dipole traps is that for the same trap depth, a smaller waist dipole trap will have a greater trapping frequency¹⁹. As can be seen in appendix A the threshold temperature for condensation is proportional to the trapping frequency, thus for a smaller waisted trap evaporation need not proceed as far to reach condensation. A higher trapping frequency will also lead to a greater elastic scattering collision rate, and thus will improve the ratio of good to bad collisions, making evaporation more efficient.

We are thus left with somewhat of a conundrum as to how to operate the dipole trap, making a dipole trap smaller will make the trap more efficient at cooling and deeper, allowing it to capture a larger thermal fraction of atoms out of the MOT, whilst making it larger will make the trap volume larger and increase the spatial

¹⁸For example, an 6 Watt beam focused to 90 micron Gaussian waist diameter at 1064 nm creates an approximately $100\mu K$ deep potential for rubidium-87.

¹⁹Xie *et al.*[138] explore this dependence for waists ranging from $30\mu m$ to $200\mu m$.

fraction of atoms from the MOT that the dipole trap can capture, but at the cost of reducing the depth of the trap and thus the fraction of thermal atoms the dipole trap can capture. We are also still left with the question of which is a better method for producing BEC, magnetic traps whilst not as efficient in their forced evaporation appear more reliable as they can begin with much larger atom number²⁰ and thus have many more atoms to spare during evaporation, and do not reach evaporative stagnation as runaway evaporation is possible, however, the timescale for evaporation to BEC is typically on the order of 100 seconds. Evaporation in a dipole trap on the other hand is much more efficient and quicker, with typical timescales being on the order of 5-15 seconds, improving the condensate producing duty cycle by an order of magnitude, however, condensates are usually smaller, and unless high initial PSD are reached, a condensate may not be produced at all. In 2009 Lin *et al.* demonstrated a solution to this problem[8]²¹, by combining the large capture volume of a magnetic trap with the evaporation efficiency of a dipole trap in what has become known as the hybrid trap.

6.4.3 The hybrid trap

The hybrid trap of Lin *et al.* seems to be the next logical step of the evolution of BEC experiments, combining a large volume trap to load a smaller volume trap. This is achieved by capturing atoms in the magnetic trap first then undergoing forced RF evaporation to cool the sample until it approaches the volume of the dipole trap before adiabatically ramping down the magnetic field gradient to just below gravity so the atoms sag in the magnetic trap and are pulled into the dipole trap, which sits just below the magnetic field zero.

In some sense, the hybrid trapping scheme of using one trap to capture and compress atoms to load another is what has already been done, consider the loading of a magnetic trap from a MOT; the spatial volume of these two traps is of similar size, as the magnetic fields of the coils used²² will extend all the way to the edge of the cell, whilst the volume formed by the intersection of the MOT beams will only be a small factor smaller. However, the MOT has a much greater capture range in

²⁰2 to 3 orders of magnitude larger.

²¹There technique was originally noted by Ketterle and Van Druten in [7], however, only in the most basic of forms, and later numerically investigated by Comparat *et al.*[142].

²²Typically the same coils are used for both stages, however, the volume of the MOT will be limited by the size of the MOT beams, thus the magnetic trap is typically larger by perhaps a factor of two.

velocity space, from the calculation in section 5.1.4 we found this to be on the order of 50 m s^{-1} , whereas for a magnetic trap that is a few mK deep, this is of the order of 1 m s^{-1} ²³. Thus the MOT can capture from a much greater velocity range, and the effect of laser cooling is to compress the distribution in velocity space and reduce the mean velocity to below the capture range of the magnetic trap, allowing efficient loading of the magnetic trap.

Similarly, the hybrid approach makes use of the large capture volume of a magnetic trap in both geometric space as well as velocity space, to efficiently load a large number of atoms into the magnetic trap. Compression in both geometry and velocity is then achieved using RF induced forced evaporation to further reduce the temperature of the atoms, and the linear nature of the trapping potential to then further compress the volume of the atoms. This is continued until the sample is within the spatial and thermal capture range of the optical dipole trap and can be efficiently loaded. The major difference between these two cases is that in the case of the laser cooling of atoms in a MOT, the number of atoms is increasing as the MOT loads as opposed to the hybrid approach where the number of atoms decreases due to forced evaporation.

The hybrid approach has allowed for an increase by one to two orders of magnitude of the number of atoms in the final condensate as opposed to conventional dipole traps loaded from a MOT[143] without using more complicated techniques to improve evaporation. This increase arises from three effects;

- A greatly enhanced number of atoms loaded into the dipole trap, an increase of typically an order of magnitude or more,
- The reduced size of the dipole trap, which improves the trapping frequency at a given depth allowing fast evaporation and a reduction in the required laser power,
- The reduced temperature of the atoms loaded into the trap

The paper by Lin *et al.*, ‘*Rapid production of ^{87}Rb Bose–Einstein condensates in a combined magnetic and optical potential*’ describes the key features of the technique very well, however, there are a few key features of the hybrid evaporation trapping scheme that are worth discussing here as they are counter-intuitive to the conventional understanding of loading a dipole trap.

²³Where we have simply calculated the average thermal velocity for a ^{87}Rb atom at 2 mK, as an order of magnitude estimate.

The hybrid approach relies upon the spatial mode matching of the optical dipole trap to the magnetic quadrupole trap for efficient transfer. As such, the approach taken by Lin *et al.* is to calculate the trap volume in both the high and low temperature limit, where the trapping is dominated by the quadrupole field and the dipole trap respectively, and calculate the crossover temperature, the temperature at which these two volumes coincide, and thus the temperature at which efficient adiabatic transfer from the quadrupole trap into the dipole trap can occur.

The overall potential that the atoms see can be written as

$$U(r) = \mu B' \sqrt{\frac{x^2}{4} + \frac{y^2}{4} + z^2} - U_0 \exp(-2[x^2 + (z - z_0)^2]/w_0^2) + mgz + E_0, \quad (6.33)$$

where $U(r)$ is the potential as a function of position $r = (x, y, z)$, μ is the atoms magnetic moment, m is its mass and g the acceleration due to gravity, B' is the quadrupole field gradient in the tight direction, U_0 the dipole trap depth, w_0 is it's waist, z_0 the location of the waist with respect to the field zero,²⁴ and E_0 is the difference in energy between the zero field location with no dipole trap present, and the trap minimum with the dipole trap present, ensuring the minimum of the potential is zero.²⁵

Now, in the high temperature limit we can ignore the effect of the dipole trap and consider only the quadrupole trap and the effect of gravity. In this case we simply ignore the second term in equation 6.33, and using the definition for the volume in section A.0.1, the trap volume is found to be

$$V(T) \approx \frac{32\pi e^{-E_0/k_B T}}{[1 - (mg/\mu B')^2]^2} \left(\frac{k_B T}{\mu B'}\right)^3, \quad (6.34)$$

where we have assumed the trap depth is much greater than the thermal energy of the atoms, typically upheld for a laser cooled cloud confined in a quadrupole trap. Note that in the case that $\mu B'$ approaches mg , the trap volume blows up to infinity, this arises as in this case the trap is no longer strong enough to support the atoms against gravity, and thus the atoms begin to fall out of the trap, this approximation will begin to break down at this point.

As discussed in section 6.3.1, when the atoms are much colder than the depth of the

²⁴Assumed to be located at $z=0$.

²⁵Note in this approximation the focussing of the dipole trap is not taken into account, however, typically the dipole trap is large enough such that the Rayleigh range is on the order of mm, and the effect of focussing is negligible.

dipole potential the potential is approximately harmonic and can be described by

$$U(r) \approx \frac{1}{2} \left[\frac{\omega_x^2}{m} x^2 + \frac{\omega_y^2}{m} y^2 + \frac{\omega_z^2}{m} (z - z_{min})^2 \right], \quad (6.35)$$

where z_{min} is the location of the potential minima in the combined trap and ω_i is the trapping frequency in the i^{th} direction.²⁶²⁷ Similarly calculating the volume of this trap yields

$$V(T) \approx \frac{(2\pi k_B T)^{3/2}}{\omega_x \omega_y \omega_z m^{3/2}}. \quad (6.36)$$

In the approach by Lin *et al.*, they numerically integrated equations 6.36 and 6.34 with their given dipole parameters, a 49 μK deep potential formed by a Gaussian beam of waists 65 μm and 78 μm along the x and z axes respectively, for a range of different quadrupole field gradients. Along with simulating the change in trap volume, they also calculated the dependence of the entropy per particle during these changes in volume and trap potential shape. Figure 6.3, reproduced from reference [8] shows the result of their calculations. Importantly they note two distinct regions, the high temperature and low temperature regions, where the dependence of the entropy on the temperature dramatically changes due to the changing trap shape from linear to quadratic, marked on the bottom graph by horizontal lines. This change occurs at what they define as the crossover temperature, T_x , which occurs when equations 6.36 and 6.34 are approximately equal, which is also the temperature at which the atoms will be transferred into the dipole trap.

What is also clear is that with decreasing field gradient, the region over which the change in potential occurs becomes much steeper, with a maximal slope occurring for $\mu B' \approx mg$, and that in all cases the final entropy per particle remains approximately the same. Thus to make the greatest use of the reduction in entropy per particle due to the changing potential shape, the transfer should thus occur at the lowest possible field gradient. From these results they also conclude that adiabatic expansion from a wide range of initial temperatures will lead to a roughly constant final temperature, approximately the crossover temperature, equal to a tenth of the depth of the dipole trap, $T_x \approx 0.1U_0/k_B$.

The graphs in figure. 6.3 have been calculated for a particular dipole potential, changing this potential will obviously change the dependence of the volume on the

²⁶Note that as we are considering a single beam trap, the curvature in one of the dimensions will be determined by the magnetic field gradient.

²⁷Given by equation 6.32 for the dimensions dominated by the dipole trap.

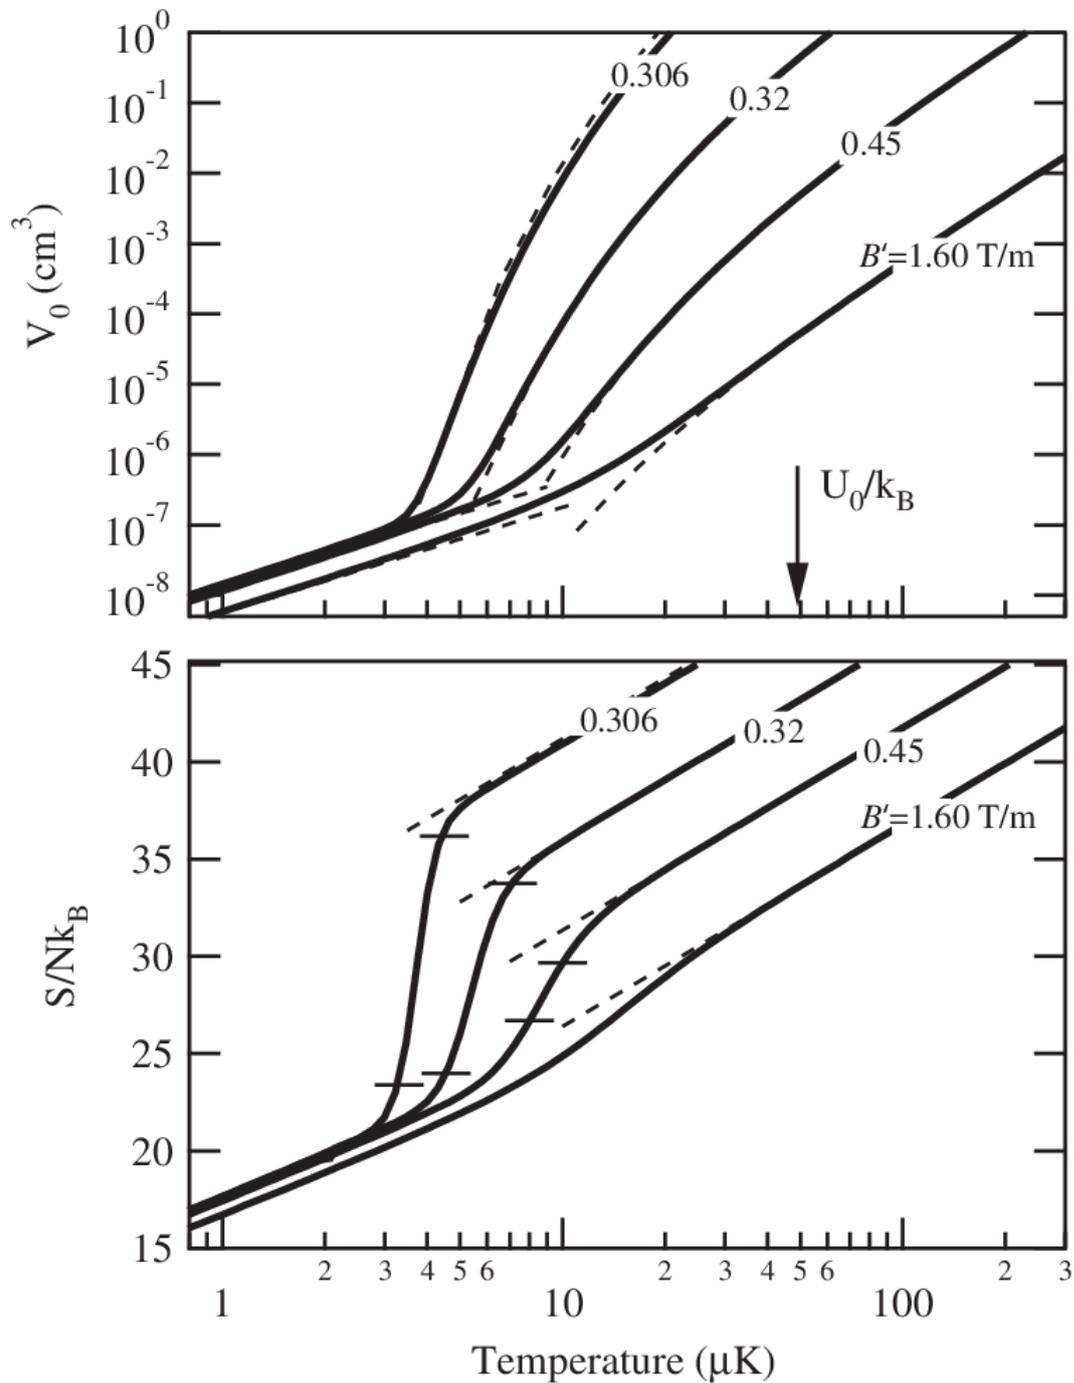


Figure 6.3: Mode-matching a hybrid trap. **Top** shown above are the calculated trap volumes in the hybrid trap as a function of cloud temperature for a range of different field gradients, given in Tm^{-1} , the dashed lines indicate the trap volumes in the dipole and quadrupole traps alone. **Bottom** The bottom figure shows how the entropy per particle changes with temperature in the hybrid trap (solid lines) and quadrupole only trap (dashed lines). Reproduced from reference [8]

temperature, thus the obvious question is what properties of the dipole trap are optimal for this process? To answer this question we must first look at what temperatures are achievable in a quadrupole trap.

The limiting factor to the lowest achievable temperature in a pure quadrupole trap is the heating induced by Majorana losses. In section 6.2.2 we estimated that for a cold atomic cloud at $\approx 20 \mu\text{K}$ in a quadrupole trap with a gradient of 140 Gcm^{-1} , Majorana losses limits the lifetime to approximately one second. In the case of a hybrid trap one would expect the lifetime to be greater, as the dipole potential acts to shift the potential minima further from the field zero, and thus reduce the density of atoms near the field zero, however, this effect won't be particularly large, thus we shall use $20 \mu\text{K}$ in a 140 Gcm^{-1} trap as our estimation of a reasonable final temperature achieved in the quadrupole trap before Majorana losses become too severe.

As the process of loading the dipole trap from the quadrupole trap requires reducing the quadrupole field gradient to the value which just holds the atoms against gravity, $B_0 = \frac{mg}{\mu}$, allowing the atoms to 'drop' into the dipole trap, there will be a reduction in temperature if the change in potential is performed adiabatically. For ^{87}Rb in (1,-1), this occurs at approximately 30.5 Gcm^{-1} , thus the temperature will also reduce by a factor of ≈ 2.8 , reducing to $\approx 7 \mu\text{K}$. This, however, will also increase the lifetime due to Majorana losses, as discussed in section 6.2.2, due to the weaker dependence of the temperature on the magnetic field gradient. Thus if we were to maintain the lifetime at one second we would thus allow forced evaporation during decompression to reduce the temperature of the atoms by a further factor of $(\frac{3}{14})^{1/3} \approx 0.6$, resulting in a cloud at a temperature of approximately $4 \mu\text{K}$. Now, as we saw in figure 6.3, the crossover region is steepest when the magnetic field is lowered to just hold the atoms against gravity. Thus over this crossover region the temperature of the atoms will change negligibly and we can set $4 \mu\text{K}$ as our crossover temperature. As the crossover temperature is approximately a tenth of the depth of the dipole trap, this requires our dipole trap to be approximately $40 \mu\text{K}$ deep.

Given the depth of the dipole trap is now known, the next question is what volume is best? Intuitively one might think that the optimal size would be the largest possible given the available power, however, this is not necessarily so.

The volume of the trap obviously plays a large role in the loading dynamics during decompression, as the volume of the trap will determine the flux of atoms crossing into the dipole trap. However, in equation 6.29 we saw that the trap depth, U_0 scaled

proportionally to the power, and inversely proportional to the waist squared. Thus to maintain a constant dipole trap depth the power must be increased in proportion to the square of the waist, i.e. $\frac{P}{w_0^2} = \alpha = \text{constant}$. If we then substitute this requirement on the relation between power and waist into equation 6.32 we find

$$\omega_{\text{trap}} = \frac{\beta\sqrt{\alpha}}{w_0} \tag{6.37}$$

for constant power, i.e. to maintain a particular potential depth for increasing waist, the trap frequency must decrease in proportion to the waist. This increase in volume of the dipole trap necessitates a decrease in the resulting trapping frequency, hampering both the loading of the dipole trap as well as reducing its forced evaporation efficiency as the scattering rate in the dipole trap will reduce, decreasing the ratio of good to bad collisions.

One may then wonder whether we could relax the requirement that the depth of the dipole potential be constant. As the depth of the dipole potential determines the final temperature of the atoms in the trap, increasing the dipole trap depth will necessarily correspond to an increase in the final temperature. This increase in temperature leads to a reduction in the phase space density of the atoms loaded into the trap, reducing the initial elastic collision rate in the dipole trap, a key parameter to the efficiency of evaporation in the trap.

The final question then becomes, how small can the dipole trap be made before the loading of the dipole trap due to its small volume is inhibited? If the loading of the dipole trap did not depend upon the volume of the trap, then the volume of the dipole trap could be reduced indefinitely, resulting in an infinitely tight dipole trap with infinitely large trapping frequencies and requiring no power. This is obviously unphysical²⁸, however, the question of how small the dipole trap can be made before its loading capability is significantly affected, to my knowledge, has not yet been well studied. One caveat of smaller volume dipole traps with higher trapping frequencies is that they necessarily create very high number densities, which can quickly lead to a high three body loss rate, especially once condensation is achieved.

In the case demonstrated by Lin *et al.*, they chose to match the volume of the dipole trap to the quadrupole trap, thus only requiring changes in the distribution of the atoms during loading, not the overall density and allowing a higher efficiency of the loading. However, it is not clear that this is necessarily the best choice, as there could

²⁸An infinitely small dipole trap in itself is limited by the diffraction limit.

be increased gains in PSD by loading into a smaller volume trap, which would also lead to an increase in the efficiency of evaporation, however, this would also require further evaporation during decompression. Lin *et al.* implemented their hybrid approach with a 49 μK deep red detuned potential, which was formed by 3.5 W of power at 1550 nm, focussed to a 65 μm and 78 μm waist along x and z , respectively, and achieved a 25% transfer efficiency from their compressed magnetic trap into the dipole trap, with a reduction in temperature by a factor of ≈ 6 , indicating some forced evaporation occurred during the decompression.

6.5 Experimental implementation of forced evaporation

Having discussed the theory behind both forced evaporation and the benefits of the hybrid trap, in the following section I shall discuss how this is implemented; from the loading of the atoms into the magnetic trap from the laser cooled cloud, through the forced evaporation in the magnetic trap and the subsequent loading of the hybrid trap, followed by the final evaporation to BEC.

6.5.1 The catch

One of the most critical stages of any BEC apparatus that makes use of a purely magnetic trap is the catch. The catch is the point at which laser cooling ends, with the cooling lasers being rapidly switched off, and the quadrupole field is quickly turned on to a value large enough to hold the atoms against gravity, and thus ‘catch’ the atoms in the magnetic trap. This stage is particularly critical, as if the catch is not smooth and quick, then it will result in significant heating, reducing the initial PSD in the magnetic trap, and thus reducing the efficiency of forced evaporation in the magnetic trap.

Lets consider some numbers here for a concrete example of a catch to find which requirements are most stringent, beginning with a PGC cooled cloud of ^{87}Rb atoms at 30 μK . Typical switch on times for coils to generate large enough field gradients to hold the atoms against gravity are on the order of ms. For simplicity we will assume that the quadrupole field is initially completely off until it is instantaneously switched on to its final value after say 3 ms, as this will make the estimations simpler. The first effects we will consider are those that arise from the atoms falling under gravity, taking $g \approx 10 \text{ m s}^{-2}$. During this time the atoms will accelerate in the

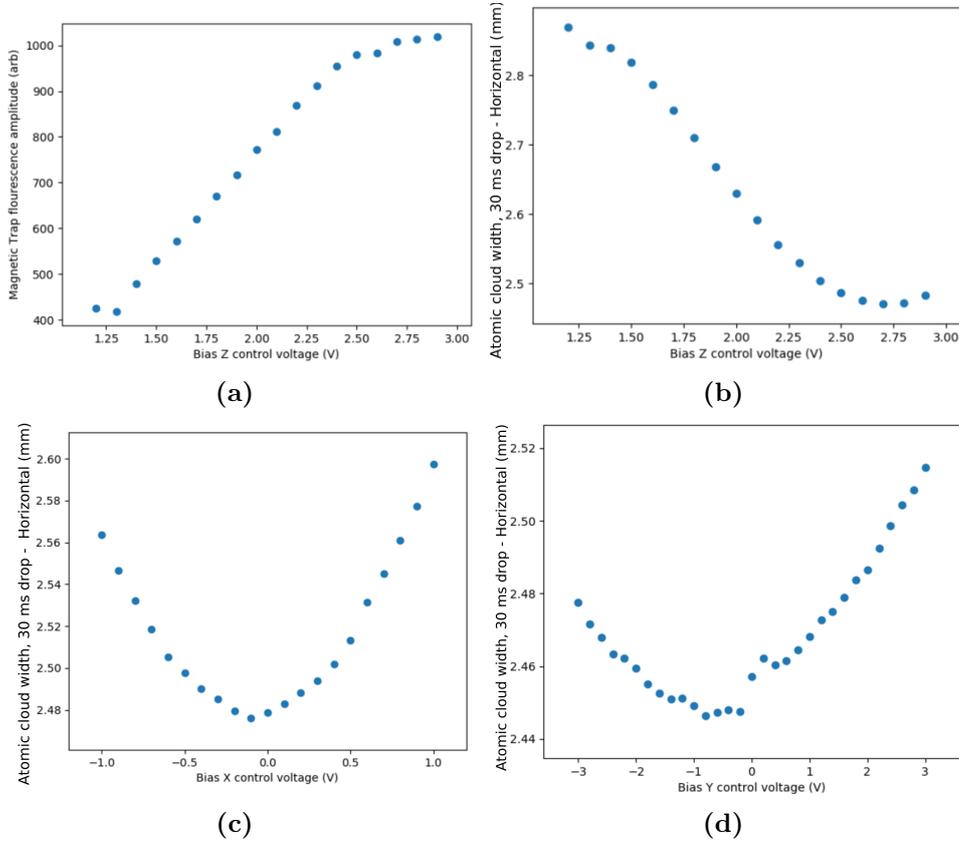


Figure 6.4: Optimisation of the catch. **a)** Peak fluorescence from the magnetic trap as a function of the Z bias control voltage, clearly more atoms are caught as the bias increases until it reaches a peak at approximately 3 V. Interestingly this corresponds to a magnetic trap position above the initial cloud. **b), c), d)** Cloud width measured after a 30 ms drop once the atoms have been allowed to thermalise in the magnetic trap as a function of the bias coils control voltages. This is the metric used to optimise the magnetic trap catch. There is a clear minimum for each bias coil, however, over the ranges investigated, there is only minimal heating seen, at most a 20% increase in the width occurred for the Z bias coil scan.

vertical direction to a final velocity of $v = at = (10)(3 \times 10^{-3}) \text{ m s}^{-1} = 3 \times 10^{-2} \text{ m s}^{-1}$, although it should be noted that this velocity is only in the vertical direction. Using the equipartition theorem we can equate this to a temperature, which comes out to a temperature on the order of $6 \mu\text{K}$ ²⁹, thus in this time the atoms will already heat by an appreciable fraction of their initial temperature.

Next we consider the movement of the atoms, as the atoms as an ensemble have

²⁹Where in using the equipartition theorem I have considered the velocity of the atoms in one dimension only, but have considered the temperature in three dimensions, i.e., I have used $\frac{3}{2}k_B T = \frac{1}{2}mv^2$, assuming the kinetic energy gained from falling under gravity will eventually thermalise in all dimensions in the final trapping potential.

no net velocity at this point we can use $s = ut + \frac{1}{2}at^2$ to calculate the distance moved, s , with the initial velocity set to zero. Substituting in the numbers leads to an estimation of $s \approx 5 \times 10^{-5}\text{m}$, thus the position of the atoms effectively does not change during the drop, with only a small motion of 10's of microns of movement. Finally, we can consider the ballistic expansion of the atoms during this time, and thus the change in size of the trap, however at such a low temperature this effect is negligible, for example at $30\ \mu\text{K}$ the average velocity of an atom in the ensemble will be $\approx 5 \times 10^{-2}\text{m s}^{-1}$, and thus the average atom will only move by $\approx 10^{-4}\text{m}$, or $100\ \mu\text{m}$. As the initial size of the cloud is on the order of a few mm, this change in position will thus be less than a tenth of the cloud's initial width.

From the above three considerations we can see that the largest change to the atoms during the catch stage is to their mean velocity in the vertical direction, with appreciable increases in the effective temperature of the atoms occurring for a few ms of drop. The increase in the temperature of the cloud can also be seen to scale with the square of the drop time, thus minimising the catch time will minimise the heating of the cloud. However, it should be noted that here we have only considered the effect of the atoms falling and the ballistic expansion of the atoms during the drop. In the case where the quadrupole field is quickly switched on but is not coincident with the centre of the atoms, one could expect further heating of the cloud. For example, if the quadrupole field is switched on at a position below the cloud then the additional magnetic force on the atoms would serve to accelerate their velocity faster than if they were just in free fall, creating a larger increase in their final temperature. Similarly, if the field zero was displaced from the centre of the atoms perpendicular to gravity, one would expect the atoms to heat up as they accelerate towards the field zero and overshoot its position.

I like to think of the optimisation of the catch from the perspective of mode matching the PGC cooled cloud to the magnetic trap³⁰. Perfect mode matching of the distribution of the atoms in both position and velocity space is not possible, however, as in the PGC cloud there should be little spatial variation in the average velocity of the atoms, whereas this is not the case in the magnetic trap.³¹ To mode match the PGC cooled cloud to the magnetic trap there are two elements to consider. First is the location of the centre of the magnetic trap with respect to the cloud, and second is the gradient of the magnetic trap, which will match the spatial distribution of the

³⁰As one would mode match focussed beam to an optical fiber, having to both match the position as well as the size and divergence of the mode to the fiber.

³¹A fact we make use of to employ forced evaporation at a fixed field gradient in the magnetic trap.

atoms in the PGC cooled cloud to that of the atoms in the magnetic trap. These two properties are physically controlled by the quadrupole field gradient, which determines the spatial distribution of the atoms, and the bias fields generated by the bias coils, which can be used to move around the location of the field zero.

Optimisation of the catch is somewhat easier than optimisation of laser cooling, as there are no hardware parameters to optimise as there is with laser cooling, such as beam alignment, collimation etc. but rather everything is controlled by software. To implement the catch, once the atoms have undergone PGC and the optical shutters have been closed to prevent any near resonant light reaching the atoms, the quadrupole and bias coils are all rapidly switched to some finite value to approximately center the quadrupole field upon the location of the atoms. To match the spatial distribution of the atoms in the PGC cloud to that of the atoms in the magnetic trap, I choose to switch the quadrupole gradient on to the same value as is used in the CMOT stage, as in this stage the cloud shape is more strongly determined by the magnetic fields as opposed to the optical distribution, and thus will most closely match the magnetic trap distribution. If the field gradient is instead increased to a value greater than the CMOT stage a greater final temperature after compression to the final field gradient is seen, this is attributed to the atoms effectively ‘imploding’ as they rush inwards due to the initial cloud size being larger than the final cloud size in the magnetic trap. If the field gradient is instead decreased below the CMOT value heating is also observed, this occurs due to a non-adiabatic expansion of the cloud volume into the magnetic trap.

With the quadrupole gradient set, it is then time to optimise the location of the trap. For my apparatus the Z bias coils play a much larger role than in most apparatus, as the geometric centre of the coils, which should approximately be the location of the quadrupole field zero in the absence of any stray magnetic fields, is not coincident with the centre of the octagonal glass chamber, which is approximately where the MOT is loaded. Thus the Z bias coils must provide a large bias field, on the order of 50 G, to move the location of the field zero of the uncompressed magnetic trap approximately 8mm above the field zero’s normal position to the optimal location for loading the magnetic trap. As described earlier, this offset of the field zero from the glass chamber centre allows our objective lens to be placed far enough away from the cell such that it can image the atoms *in-situ* in their final location, as well as have enough clearance between the cell and the objective so that a thin quarter waveplate with a high reflection coating for 780nm on the backside can slide in and out of position above the objective for either retroreflecting the MOT beam, or

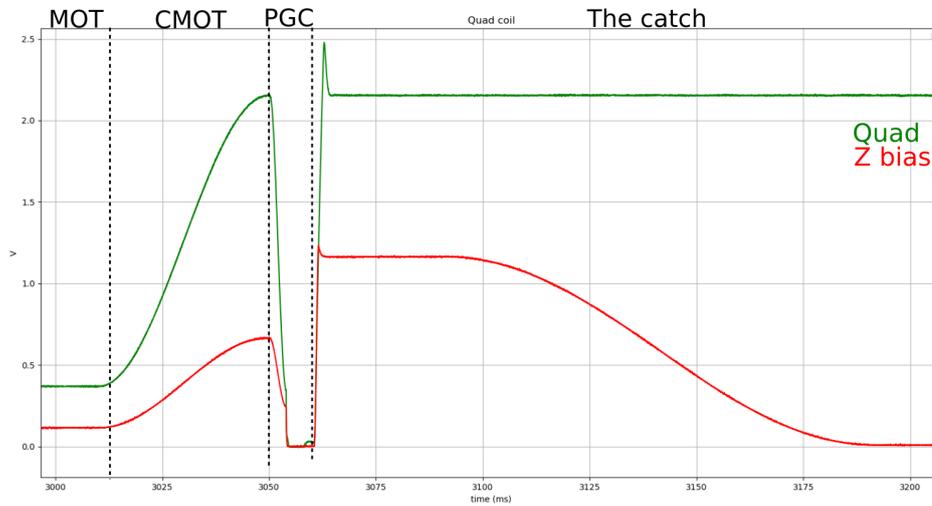
imaging the atoms respectively.

To optimise the location of the field zero of the uncompressed trap, the atoms are initially caught and then are left to thermalise in the trap for 5 seconds before they are allowed to expand in time of flight for a further 30ms, before they are imaged. The width of the cloud after this time of flight will be approximately proportional to the temperature, and thus this is used as my metric for optimal phase space density, which in a magnetic trap will also correspond with a minimum in temperature for a fixed number of atoms, whereas the *in-situ* peak fluorescence is used as a metric for the total number of atoms caught.

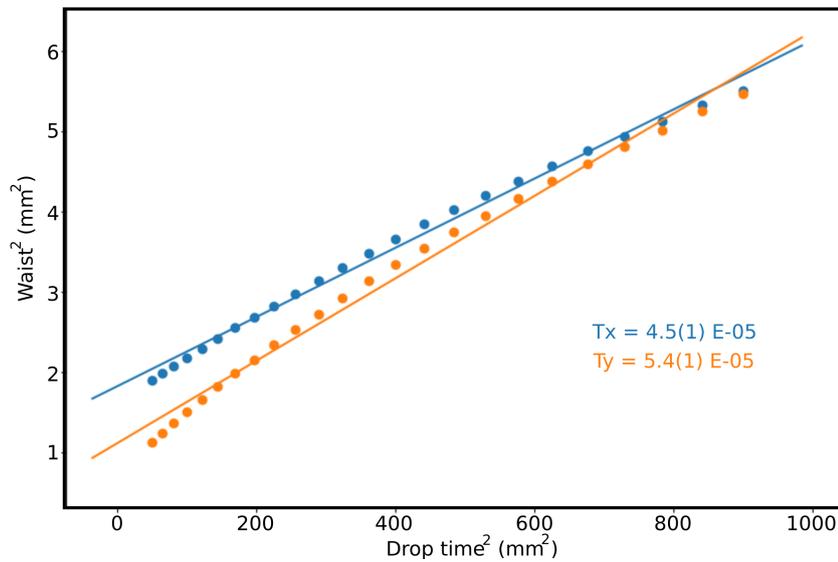
Figures 6.4a and 6.4b show the peak fluorescence and width of the cloud after dropping the atoms respectively as a function of the control voltage for the Z bias coil. Clearly the number of atoms caught in the trap depends strongly upon the location, with a reduction of the peak fluorescence by up to a factor of 2 when the field zero is moved away from the optimal position. Similarly, there is a clear optimal control voltage for the Z bias coil to minimise the heating of the atoms as they are transferred into the trap, occurring at roughly the optimal value for catching the largest number of atoms. The heating effect is smaller than the reduction in atom number, however, with only a 20% increase in the width of the cloud at its worst. Interestingly, the optimal location for the field zero for these two conditions is above the location of the PGC cooled cloud as opposed to centred on the cloud. This is understandable through considering the effect of gravity on the atoms, as during the catch stage the atoms are initially in free fall, and thus can quickly gain velocity, and hence increase their temperature. However, if the field zero is located above the initial location of the atoms then instead the atoms will feel a force opposing gravity, reducing the acceleration of the atoms during the catch, minimising the heating and thus maximising the fraction of atoms that can be caught in the trap.

Once the control voltage for the Z bias coil had been optimised, the X and Y bias coils control voltages were then also optimised, as shown in figures 6.4c and 6.4d. This data also shows a clear minima in the width of the ballistically expanded cloud as a function of control voltage, but the net effect of these bias fields upon the final width is much smaller, with the maximal change in width over the range of control voltages investigated being on the order of 5%.

With the bias fields for the catch optimised, another parameter that can be adjusted is the time over which the coils are ramped on. One might assume that ramping the coils on as fast as possible would result in the highest initial PSD in the magnetic



(a)



(b)

Figure 6.5: Uncompressed magnetic trap thermometry and measured coil response a) shows the measured response of the quadrupole coil and Z bias coil during the CMOT stage, PGC and the catch. There is clearly an overshoot of about 15% as the quadrupole coils are suddenly ramped on, but I find this has little effect upon the resulting temperature of the atoms. b) Thermometry performed on the uncompressed magnetic trap, 5 seconds after catching the atoms to allow for adequate thermalisation.

trap, however, this is not always the case as if the control voltage is suddenly switched to some finite value the Magneato drivers can overshoot the desired current and then oscillate about the desired value for some time, which can heat the atoms,

reducing the PSD in the trap. The time over which the coils are ramped on thus also plays a role in the efficiency of the catch, and thus is another parameter that can be optimised. Figure 6.5a shows the monitor line for the current through the quadrupole coil and the Z bias coil during the compression of the MOT, polarisation gradient cooling and the catch. There is clearly an overshoot of the current in the quadrupole coil as it is switched on of about 15%, but I find this has little effect upon the temperature of the atoms.

With all parameters for the catch optimised and a 3 second MOT load time, the magnetic trap catches approximately 1×10^9 atoms at $\approx 50 \mu\text{K}$ in a quadrupole gradient of 69 Gcm^{-1} , corresponding to an initial phase space density of 8.3×10^{-6} . Figure 6.5b shows thermometry of the uncompressed magnetic trap after a 5 second hold period. Clearly the temperature of the atoms has increased somewhat compared to their final PGC temperature, however, this is not unexpected due to the slightly irregular shape of the laser cooled cloud, as well as the change in the dependence of the velocity upon the position of the atoms between the PGC cloud and the magnetically trapped atoms.

6.5.2 Moving and compressing the trap

Once the atoms have been caught in the magnetic trap, the field zero is then moved to its final location by ramping the bias fields whilst maintaining the quadrupole field at its catch current. The location of the field zero, or equivalently the magnitude of the bias fields, is chosen such that during compression and decompression of the magnetic trap, the field zero does not move location. To achieve this the location of the magnetic trap for a range of quadrupole gradients and bias fields is measured, and the bias fields that minimise the movement of the cloud are chosen.³² Figure 6.6a shows how the position of the atoms in the quadrupole trap moves as a function of the Z-bias control voltage as the trap is decompressed to different gradients, clearly a control voltage close to 0 minimises this movement. This condition on the bias fields is chosen so that during the decompression of the magnetic trap to load the atoms into the hybrid trap, the field zero does not significantly move, as if it does it will reduce the efficiency of the transfer.

³²Note that the location of the cloud can be more precisely found if a colder, and thus smaller, cloud is used, and typically the gradients used for observing the motion are only reduced to approximately 90 Gcm^{-1} , as below this the cloud starts to sag under gravity, reducing the ability to accurately locate the field zero.

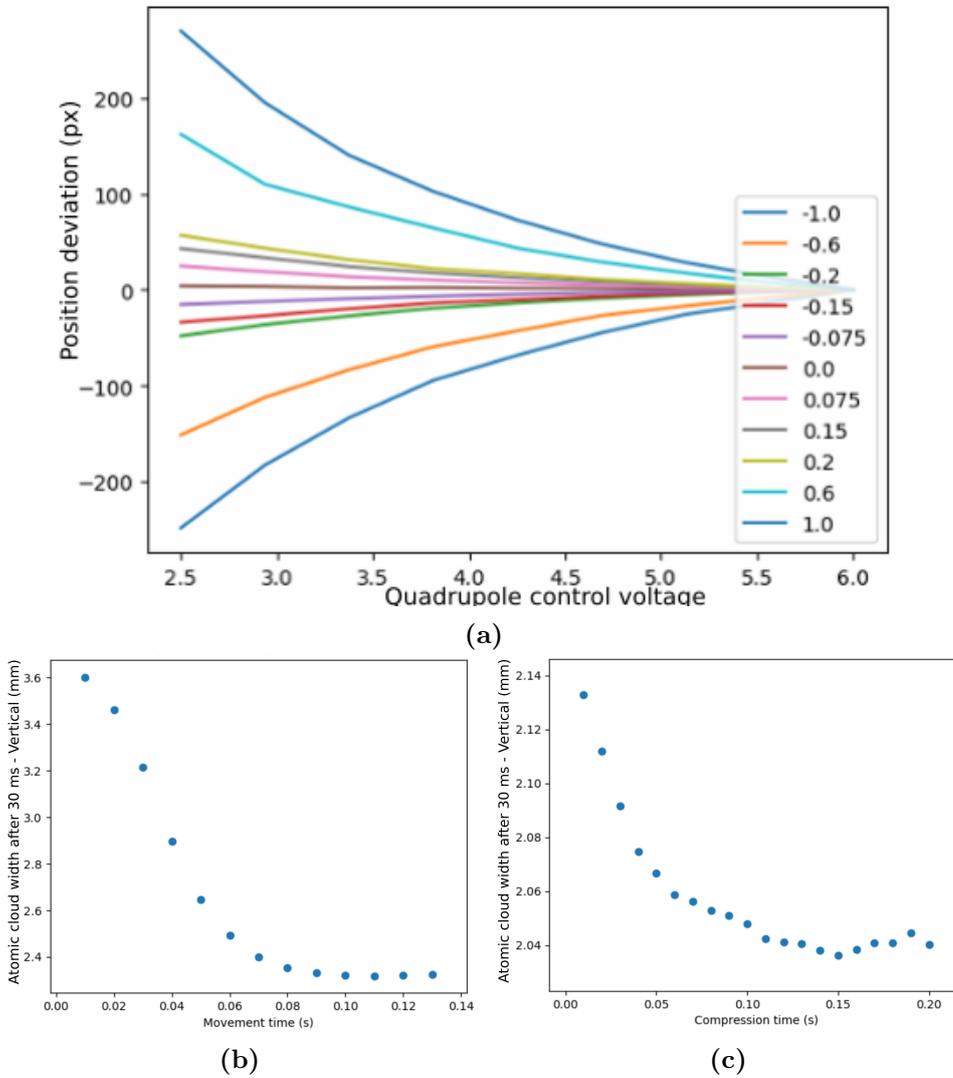


Figure 6.6: Adiabatic movement and compression of the magnetic trap. **a)** shows how the position of the atoms in the quadrupole field moves as the trap is decompressed for a range of different control voltages for the Z-bias coil. This data is used to minimise the movement of the quadrupole trap whilst loading the hybrid trap. **b)** and **c)** show the dependence of the width of the ballistically expanded cloud on the time over which the magnetic trap is moved and compressed, respectively. Clearly there is a stronger dependence upon the time over which the cloud is moved as opposed to the time over which it is compressed, with the data showing a maximal 50% and 5% increase, respectively. This data was used to determine the rate at which the cloud could be moved and compressed that ensured the cloud did not significantly heat.

Once the atoms have moved to their final location and settled, the quadrupole gradient is then increased to its final value of 190 Gcm^{-1} , to further compress the atoms and thus increase the elastic scattering rate, hence improving the efficiency of forced evaporation. Figures 6.6b and 6.6c show the dependence of the width of the ballistically expanded cloud on the time over which the atoms are moved and compressed respectively. Ideally the changes to the magnetic trap will occur adiabatically, and thus maintain PSD. From figure. 6.6b there is clearly a strong dependence upon the time scale over which the atoms are moved with the width of the cloud increasing by up to 50% if the cloud is moved over 10 ms, whilst there is no discernible increase in the width of the cloud after expansion if it is moved over a time greater than 80 ms. Figure 6.6c shows less of a dependence upon the time over which the atoms are compressed, with a maximal increase in the width of the cloud of only 5% if it is rapidly compressed over 10 ms. This data indicates that heating of the trap is more likely if the location of the trap is shifted as opposed to if the gradient is increased. It should also be noted that the quadrupole gradient is maintained during the movement of the field zero as heating was observed if it was not held constant.

Once the atoms have thermalised in the compressed magnetic trap, the atoms have a final temperature of $140 \mu\text{K}$ corresponding to an increase in temperature by a factor of ≈ 2.3 , which agrees closely with the expected increase in temperature due to adiabatically compressing the trap from 69 Gcm^{-1} to 190 Gcm^{-1} .

6.5.3 Characterising the magnetic trap

Once the atoms have been caught, settled and compressed in the magnetic trap, it is time to characterise the trap to ensure it is behaving as expected. A pure quadrupole trap is a relatively simple trap to characterise, especially when the width of the trap in-situ is much smaller than the radius of the coils used to produce the trap, as this will ensure the trap is relatively smooth and thus free from aberrations. There are three key quantities that characterise the magnetic trap, the quadrupole field gradient, the magnetic trap lifetime and the heating rate in the trap.

As the atoms are within the vacuum chamber, the field gradient cannot be directly measured with a magnetometer. One method of measuring the field gradient is to first measure the current through the quadrupole coils required to levitate the atoms against gravity, and then by knowing the gradient is proportional to the current,

the gradient at all other currents can be known. This measurement is most accurate with a very small cold cloud, as this will allow levitation times on the order of 100's of ms whilst still having a small enough cloud such that uncertainties in fitting the centre of the cloud are small, and still having a large enough OD to have an appreciable signal. This calibration implies the quadrupole coils generate a magnetic field gradient of $1.7 \text{ Gcm}^{-1}\text{A}^{-1}$ agreeing well with the theoretical simulation of the coils, which indicated a field gradient of $1.75 \text{ Gcm}^{-1}\text{A}^{-1}$.

To measure the lifetime of the atoms all one must do is simply hold the atoms in the trap for variable lengths of time and measure the decay of the number of atoms. The lifetime of the atoms is most likely going to be limited by background collisions, as 2-body collision for ^{87}Rb in $(1, -1)$ are relatively small and the density of the atoms in the trap is not large enough for 3-body collisions to play a dominant role. To get the most accurate vacuum limited lifetime, the temperature of the atoms should be large enough, or the trap weak enough, such that the Majorana loss rate is much larger than the lifetime. Figure 6.7b shows the decay of the number of atoms in the magnetic trap as a function of the hold time, an exponential decay fit to the data indicates a lifetime on the order of 70 seconds³³, indicating that the pressure in the Science chamber is approximately as expected.

It should be noted though that there is another source of atom loss that can occur in a magnetic trap, the scattering of photons. The scattering of photons from atoms in the magnetic trap can cause a severe limitation to the magnetic trap lifetime, for example, consider the case of ^{87}Rb atoms trapped in $(1, -1)$. Now if one of these atoms was to scatter a photon from the repumping light, then as the photon can only cause a change of 1 unit of angular momentum, the atom will either end up in $(2,0)$ or $(2, -2)$ neither of which are magnetically trappable states. This effect is similar in the hybrid trap, however, in this case the atoms which absorb a photon from the repump laser may stay within the dipole beam if the atom is pumped into $(2,0)$ as this state has no magnetic dipole moment, but would spread out along the length of the dipole beam as there will be no radial confinement along the beam. In the case of a pure dipole trap, such as a crossed trap, this effect is unlikely to be noticed as the atoms will likely remain in the dipole trap as the trap is not state selective, however, the trapped sample will end up in a mixture of different trapped states which may present a higher loss rate, nevertheless this effect could be observed

³³To accurately measure the lifetime of the atoms, the dependence of the atom number on time should be measured over multiple lifetimes. This was not possible in our apparatus, as running the coils at such high current for such a long time created undesired heating, tripping the safety interlock.

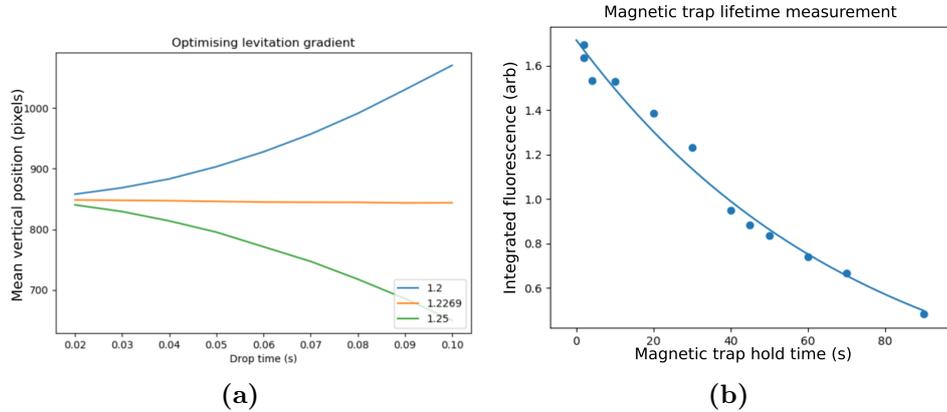


Figure 6.7: Characterising the quadrupole trap gradient and lifetime. **a)** The measurement used to calibrate the field gradient as a function of current for the quadrupole coils. The data shows the mean fitted position of a cold cloud of atoms for a range of drop times from 10-100 ms, for three different position control voltages. The largest variation in position is only 1 pixel $4.5\ \mu\text{m}$ over the range of drop times investigated, suggesting a cancellation of gravity on the order of 1 part in 10^4 . **b)** The fluorescence from the magnetic trap, measured after a 0.5 ms drop using light detuned 20 MHz below resonance, decays as a function of time in the magnetic trap, a reasonable fit to the data (solid line) indicates a lifetime on the order of 73 seconds, this is indicative of a close to vacuum limited lifetime.

through a Stern-Gerlach type imaging scheme.

Lets consider quantitatively how large this effect can be, for example we shall consider a cloud of 10^9 atoms, and an apparent decay time of 5 seconds. As the *in-situ* optical depth of a quadrupole trap with this number of atoms is typically greater than 10 we will assume that any photon that enters the cloud will be absorbed, thus as we expect the number of atoms to decay to approximately one third of its initial value over 5 seconds. This corresponds to approximately 10^8 photons entering the cloud each second. As the energy of a photon is given by $\frac{hc}{\lambda}$, where λ is the wavelength of the photon, multiplying this by the number of photons gives as little power as 10^{-11}W being enough to cause this loss rate in the cloud. If we estimate the Gaussian width of the atoms in the magnetic trap as approximately 1 mm, then the area that cloud subtends is on the order of 10^{-2}cm^2 , and thus the intensity of the atoms need only be 10^{-9}W cm^{-2} . This is a surprisingly small value, corresponding to an intensity that can only just be measured with a typical power meter, and most likely not visible with an IR viewer either. This intensity is easily greater than the light leakage through an AOM, indicating that mechanical shutters must be used after AOMs to prevent stray light reaching the atoms, and is also easily within the realm of the intensity of stray light scattered off objects on the experiment table. In the case of unexpectedly short lifetimes stray photons should thus be considered

the first possibility, which can be easily tested by blocking the direct output of the lasers with a black object.

I found that the lifetime of the atoms in the magnetic trap and dipole trap was severely hampered by this effect, with my initial measurements of the lifetime in the quadrupole trap indicating a lifetime on the order of 4 seconds. After installing an additional shutter directly at the main output of the repump laser just after the PBS cube that split off light for the repump lock, the lifetime was increased to 8 seconds. To achieve the lifetime of 70 seconds observed in figure 6.7b, a box of anodized aluminium foil had to be constructed around the repump lock³⁴ to prevent any light escaping.

The final property of the quadrupole trap that should be investigated is the heating rate of the atoms in the magnetic trap. This was investigated by measuring the temperature of the atoms 100 ms after compressing the trap and then again at 5 and 10 seconds later. These measurements show no clear evidence of heating over the three measured times, with the measured temperatures being within the measurement uncertainty of each other. This indicates that there is no significant heating of the atoms in our pure quadrupole trap even though a small ($< 1\%$) approximate 1 kHz modulation is observed on the monitor line of the quadrupole coil.

6.5.4 Forced evaporation in the compressed magnetic trap

With the magnetic trap loaded, compressed and characterised, it is time to perform forced RF evaporation to reduce the temperature and increase the phase space density to improve the load into the hybrid trap. Forced evaporation in the magnetic trap is instigated by applying RF to the atoms through a 2 turn RF antenna approximately centred on the atoms and placed 20 mm below the cloud, sitting atop the bottom quadrupole coil. The RF antenna is made by winding 2 mm enamelled copper wire around a 65 mm diameter circular former and then taping the wire together to keep the antenna fixed in shape, the resultant antenna is shown in figure 6.8a.

As a wide range of frequencies are used for evaporation it is difficult to impedance match the coil for the entire range. For my RF coil, some resonances are observed in the range of frequencies used, 2-40 MHz, these can be measured through the use

³⁴A shutter could obviously not be placed before the repump lock, as otherwise the laser would lose its lock point every experiment!

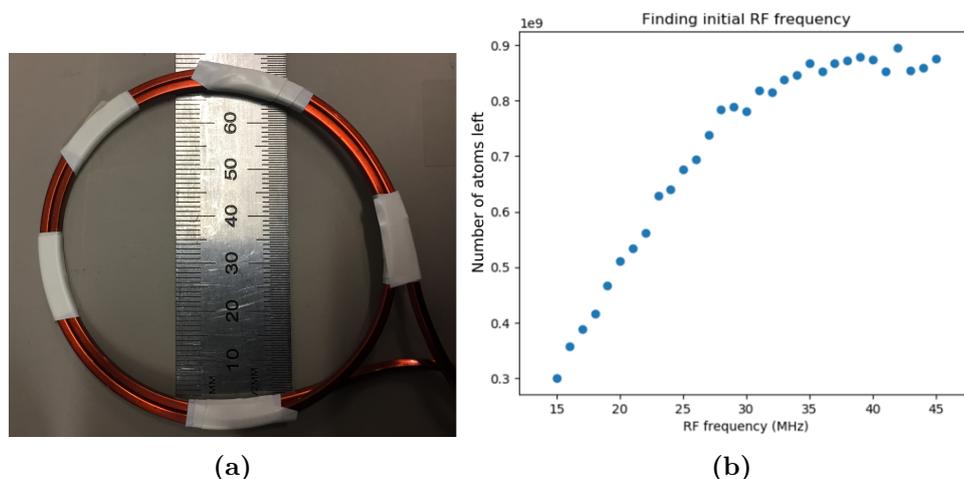


Figure 6.8: RF Antenna and forced evaporation a) The RF antenna, formed by two loops of round 2 mm enamelled copper wire wound about a 65 mm circular former. b) To find the initial RF frequency used for evaporation the RF is held on at a constant frequency for 2 seconds and the atom number is measured.

of a small pickup coil placed near the RF antenna or alternatively I see them upon the monitor lines of the MOT fluorescence photodiode or the coil current monitors. These resonances do not appear to pose a problem with my evaporation. I believe this is since the atoms are at such a high temperature the power broadening of the RF knife does not present an issue as it does not approach the temperature of the atoms. Others have observed this^[144], with the efficiency of the first stage of evaporation being quite insensitive to evaporation parameters including RF power, before optimisation of power broadening at lower temperatures became necessary. If the resonances were an issue, it would be possible to measure the power coupled out of the RF antenna as a function of frequency with a network analyser, and adjust the input power accordingly to maintain constant RF power.

To ensure the RF from the coil is successfully coupling to the atoms, and to find the frequency at which evaporation should begin the RF is switched on to a constant frequency and held for 2 seconds before the atom number is measured. Figure 6.8b shows how the atom number as a function of frequency changes, for the maximum RF power available applied to the RF coil. Clearly there is a large decrease in atom number for frequencies below about 35 MHz, whilst for larger frequencies there seems to be minimal effect. This agrees with what we would expect, an RF frequency of 40 MHz corresponds to a temperature of approximately $1200 \mu\text{K}$ ³⁵, 8 times the temperature of the atoms in the trap, thus evaporation would approximately halt

³⁵via equation 6.20.

at this large depth.

With the RF known to influence the atoms, forced evaporation can begin. Forced evaporation in the magnetic trap is performed by linearly ramping the RF frequency from 30 MHz down to 7 MHz in 5 seconds. The final RF frequency used is chosen to set the final temperature of the atoms at approximately 30 μK , based upon the considerations in section 6.4.3, and the initial frequency and duration of the ramp are chosen to maximise the number of atoms at the end of evaporation. If one was to perform forced evaporation in the magnetic trap to BEC, the forced evaporation trajectory should be optimised by optimising γ , as defined in equation 6.13, at each point along the trajectory, however, this is unnecessary for my apparatus, with the number of atoms loaded into the hybrid trap being relatively insensitive to these evaporation parameters³⁶

At the beginning of evaporation we began with 1×10^9 atoms at 150 μK in a quadrupole trap with a field gradient of 190 Gcm^{-1} , corresponding to an initial phase space density of 8.3×10^{-6} . At the end of evaporation at constant gradient we are left with 2×10^8 atoms at 28 μK , resulting in a final phase space density 3.2×10^{-3} , corresponding to an increase by a factor of ≈ 380 . Once BEC had been achieved in the hybrid trap, however, the parameters of the evaporation ramp were adjusted to see how the condensate number was affected, and it was found that reducing the duration of the ramp from 5 seconds to 2 second showed no reduction in the final condensate number thus a 2 second ramp is now used.

6.5.5 Loading the hybrid trap

The hybrid trap is formed by the combination of a far off-resonant optical dipole trap, which provides confinement in the radial directions along the optical axis of the beam, whilst the remaining quadrupole gradient provides confinement along the axial direction of the beam. The light for the optical dipole trap is provided by a Keopsys 1064 nm fiber laser, providing a maximum output of 20 W, however, typically only 6 W is used in the hybrid trap.

The optical setup used to form the dipole trap in which forced evaporation to BEC occurs, as well as an additional light sheet trap used to confine the atoms for later experiments, is shown in figure 6.9 The output mode of the Keopsys is a 2.6 mm

³⁶I did perform a three step optimisation by splitting the evaporation ramp into three linear ramps and optimising γ at the end of each ramp, but this did not lead to any significant increase in atom number, thus the simplicity of a single ramp was used instead.

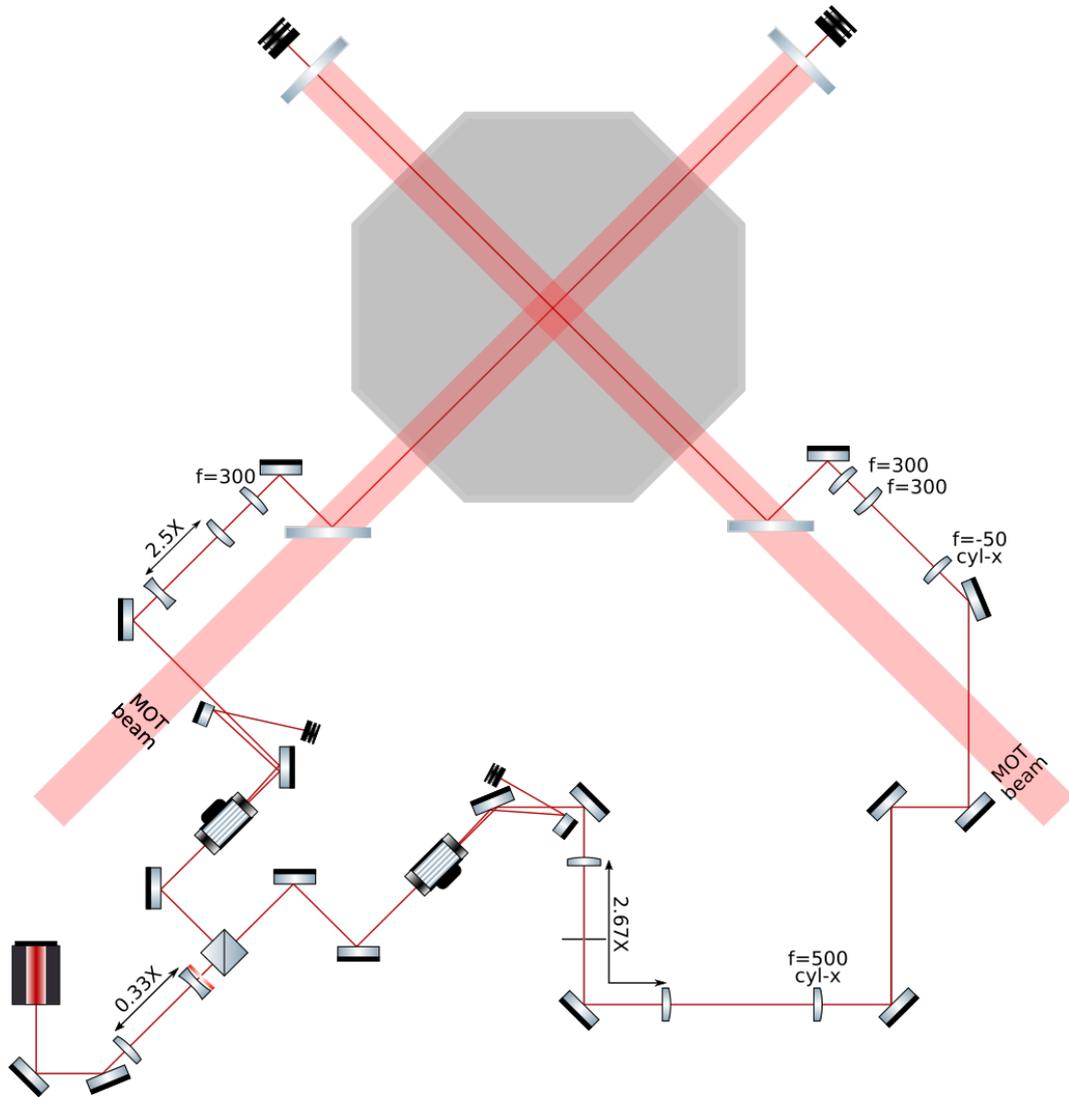


Figure 6.9: Dipole setup schematic A schematic of the optical setup used to generate the optical dipole beam used for hybrid trapping and evaporation (left hand side) and the sheet trap (right hand side) used for axial confinement of the atoms for later experiments. All optical elements are as in the cooling and repump laser schematics.

$1/e^2$ Gaussian beam, which first passes through a $0.33\times$ telescope, producing a beam of a Gaussian diameter of 0.87 mm ³⁷ before passing through a halfwave plate and polarising beam splitter cube to form the two beamlines for the dipole traps. The beamline that forms the dipole trap in which forced evaporation occurs, which I will refer to as the hybrid beam, reflects off the PBS and a mirror before passing through

³⁷This size is chosen as it maximises the diffraction efficiency through the AOM used, a Gooch and Housego 3110-197 AOM.

the AOM, whilst the other beam, which I will refer to as the sheet trap, reflects off two mirrors before passing through its AOM. The angle of the beam through the AOM for the hybrid beam is adjusted so that the power in the diffracted order that deflects towards the input RF is maximised, whilst the power in the diffracted order that deflects away from the RF is maximised for the sheet trap, this ensures the two beams have a large frequency difference and thus any interference effects will be on a timescale of 100's of MHz, imperceptible to the atoms. D-shaped pick off mirrors are placed after the AOMs to pick off the undiffracted (zero) order beams, and dump them safely.

Next, the hybrid beam passes through a $2.5\times$ telescope formed by a $f=-50$ mm lens and a $f=125$ mm lens to increase the beam size to ≈ 2.2 mm before being focussed by a $f=300$ mm lens onto the atoms, which should theoretically create a beam waist of ≈ 90 μm . The output of the $2.5\times$ telescope is made as collimated as possible, as this should minimise aberrations in the system and as the focus of the trap will be adjusted by translating the final focussing lens, this will make the position of the trap move by the same amount the final lens is translated³⁸. The telescope and final focussing lens are all mounted in a Thorlabs 30mm cage system to align the lenses optical axes. The final focussing lens is not mounted in a Z-translation stage as it was decided the deflection of the beam caused by the tilt of the lens in the translation mount as it is actuated was not worth the increased precision in the ability to focus the trap the translation stage provides as the Rayleigh range of the beam is approximately 2 cm, thus accurate focus of the beam to within a Rayleigh range is readily achievable with manual positioning of the lens by hand. After the final lens a mirror and a dichroic mirror that reflects 1064 nm light and transmits 780 nm light are used to combine the dipole beam with the MOT light, whilst on the other side of the experiment a custom air spaced dichroic is used to reflect the 780 nm light and transmit the 1064 nm light to be safely dumped. Details for the sheet trap can be found in section 8.2.1.

To load the hybrid trap, first the dipole beam needs to be located. This is a non-trivial task as the dipole beam is typically only 10's to a few 100 microns wide, the Rayleigh range is at most only a few millimetres, and the depth is typically only 10's to 100's of μK deep, thus the beam has to be well placed in all 3 dimensions for any signal to be observed. The easiest method I have found to align the dipole beam is to first align the beam to an existing beam that is known to pass through the atoms, in my experimental configuration this is relatively simple as the dipole

³⁸This will not be the case if the input beam is not collimated.

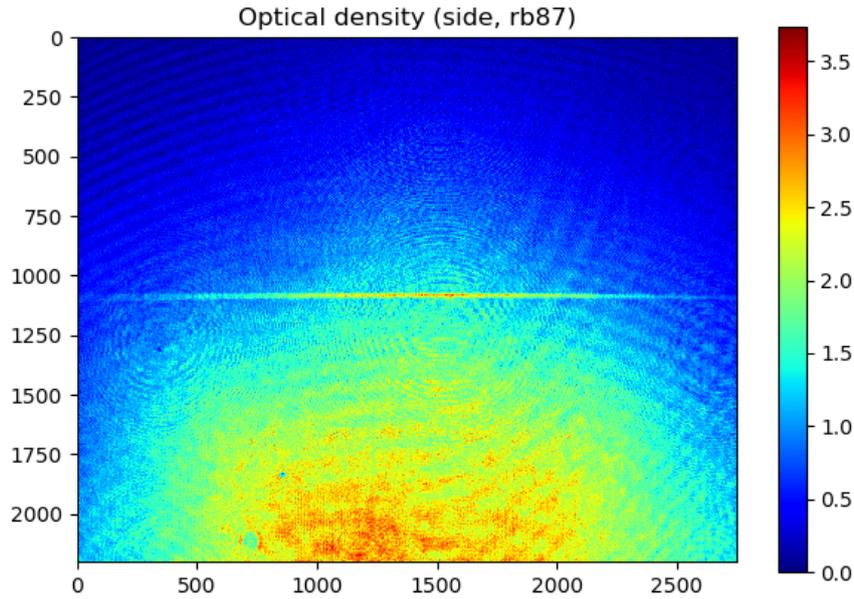


Figure 6.10: Finding the dipole beam An absorption image showing how the dipole beam is found through allowing ballistic expansion of the compressed magnetic trap to expand the volume of the cloud whilst keeping the dipole beam on to capture the atoms that are cold enough showing the location of the beam.

beam co propagates with one of the MOT beams. The alignment thus begins by aperturing down the MOT beam to the smallest size possible whilst still observing a MOT³⁹ and then aligning the dipole beam to this MOT apertured MOT beam at either side of the cell. As the dipole beam is on the order of a millimetre outside the cell and the apertured MOT beam is a few millimetres across, this gets the dipole beam to within a few millimetres of the atoms.

Next, the dipole beam is found by loading the magnetic trap and switching on the dipole beam, then releasing the atoms from the magnetic trap and after some expansion time switching off the dipole beam and imaging the atoms almost immediately after. This method works by capturing the atoms that are cold enough to be trapped in the dipole beam from the ballistically expanding cloud, which allows the dipole beam to be seen as a dense line of atoms. Typically expansion times of 20-30 ms are used, however, there is a trade off as longer expansion times will allow the cloud to expand further filling a bigger volume and thus creating a larger target for the

³⁹For my apparatus where the location of the field zero for the MOT load and the final compressed magnetic trap differ the MOT was moved to the compressed magnetic trap for this alignment.

dipole beam to hit but this comes at the cost of a reduced density of atoms, giving a smaller signal. Alternatively the atoms from an RF evaporated compressed magnetic trap or decompressed magnetic trap can be released and the same procedure followed, this is useful for finding shallow dipole traps as the atoms will be of lower temperature, but they will also expand slower creating a smaller target. Figure 6.10 shows a resulting absorption image when the dipole beam is within the volume of the ballistically expanded cloud, the dipole beam seen as the dense line of atoms within the cloud. Once the dipole beam has been found it is important to level the trap with respect to gravity, this ensures the trap is symmetric. If the dipole trap is not symmetric there will be a potential gradient across the trap, and once forced evaporation begins atoms will only evaporate along one direction, reducing the efficiency of the evaporation.

With the dipole beam found the hybrid trap can then be loaded. To load the hybrid trap the magnetic field gradient is ramped down from its compressed magnetic trap value to a value just below that which compensates for gravity, I typically use 27 Gcm^{-1} , over 3 seconds. During this time the RF frequency is also swept, beginning at the final value occurring at the end of evaporation in the compressed magnetic trap, and in my case ending at 3 MHz. The final frequency and duration of this ramp are both scanned to optimise the number of atoms in the trap, but these scans are performed once the position of the dipole trap has been optimised.

Optimisation of the position of the dipole trap begins with optimising the position in the plane transverse to gravity which begins by aligning the dipole trap vertically with the center of the atoms in the compressed magnetic trap. Next the position of the dipole beam is scanned in the horizontal plane, the compressed magnetic trap is decompressed to load the trap and the final parameters are held constant for 1-2 seconds before the atom number is measured, the results of this scan are shown in figure 6.11a. There is a clear dip in the center of the scan, this dip arises from close alignment of the dipole trap to the field zero, which increases the Majorana loss rate depleting the trap, however, this feature is typically only noticeable if the atoms are held in this position for a few seconds. Although the atom number is lowest when the dipole beam is aligned to the field zero, the dipole beam is positioned here as this loss effect is reduced by displacing the beam vertically, the red dot indicates the position the dipole beam was left at. Next the position of the beam is scanned vertically, the results of which are shown in figure 6.11b. The horizontal axis indicates the mean position of the atoms in the vertical direction in pixels, where an increase in pixel value corresponds to a displacement downwards. Clearly in this data there is

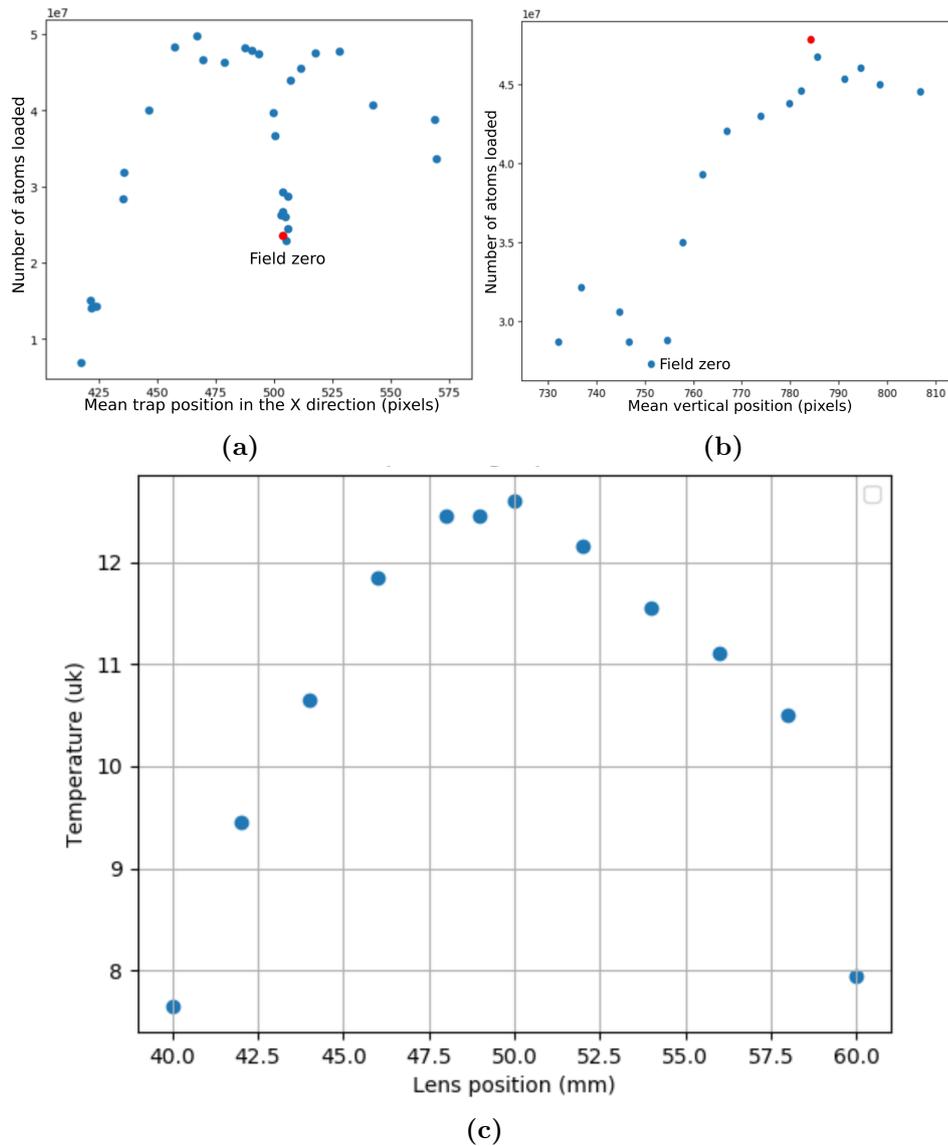


Figure 6.11: Optimisation of dipole location a) and b) show the dependence of the number of atoms loaded into the hybrid trap upon the position of the dipole beam. In both a) and b) there is a clear dip in the number of loaded atoms that arises from precise alignment of the dipole trap with the field zero, greatly increasing the Majorana loss rate. The red dots indicate the position the dipole beam was left at. In a), the horizontal position is scanned whilst the vertical position is fixed at the location of the field zero, and similarly for b) the vertical position is scanned whilst the horizontal position is fixed at the location of the field zero. c) shows the measured temperature of the atoms in the hybrid trap. As the temperature in the hybrid trap is set by the depth of the dipole trap, translating the lens to maximise the temperature thus maximises the depth at the field zero and thus places the focus at the atoms.

still the presence of the field zero, and the position that loads the largest number of atoms is located below the field zero by approximately 30 pixels corresponding to a distance of $\approx 135 \mu\text{m}$, larger than the distance suggested by Lin *et al.*[8] where they placed the beam one waist below the field zero, however, I find that evaporation at this position is still effective. A possible explanation for this is that I load a larger number of atoms into the trap than Lin *et al.*, and thus the increased distance may be required to relax the axial confinement and decrease the density in the dipole trap, reducing 3-body losses.

With the position of the beam optimised in the x-y plane, the next step is to optimise the focus of the beam by translating the lens that focusses the beam onto the atoms. One method of measuring how well focussed the dipole beam is near the field zero is to measure the temperature of the atoms, as when the trap is the most tightly focussed the trap will be the deepest and as we saw in section 6.4.3 the depth of the trap determines the temperature of the atoms once loaded into the hybrid trap. Optimisation is thus performed by translating the lens that focuses the trap, measuring the temperature of the atoms in the trap and repeating. Figure 6.11c shows the result of such a measurement, with the horizontal axis showing the displacement of the lens from an arbitrary reference position and the vertical axis showing the measured temperature. The data shows a reduction in temperature of approximately a factor of almost two over a range of 1 cm, as a factor of two reduction in the trap depth indicates an increase in the waist by a factor of $\sqrt{2}$ ⁴⁰, this thus corresponds to approximately one Rayleigh range, roughly in agreement with the expected size of the dipole beam. Once the dipole beam has been well focussed, it is worthwhile checking to see if it is still at the optimal position as translation of the lens can affect the position if the dipole beam is not exactly centered on the lens and aligned with its optical axis.

With the position in all 3 dimensions of the dipole beam optimised for the hybrid trap load the decompression parameters can then be scanned. As discussed these mainly require scanning the final RF frequency and duration of the decompression, although adding a period of constant RF frequency at the end of the evaporation and changing the functional form of the RF sweep can improve the loaded number also. The optimal parameters for loading the largest number of atoms into the hybrid trap results in the quadrupole coils being decompressed from their initial gradient of 190 Gcm^{-1} to 26.5 Gcm^{-1} over 1.5 seconds. During this time the RF frequency is swept exponentially from 7 MHz to 3.5 MHz in the first 0.5 seconds with a 1.25 second

⁴⁰Implied by equation 6.29.

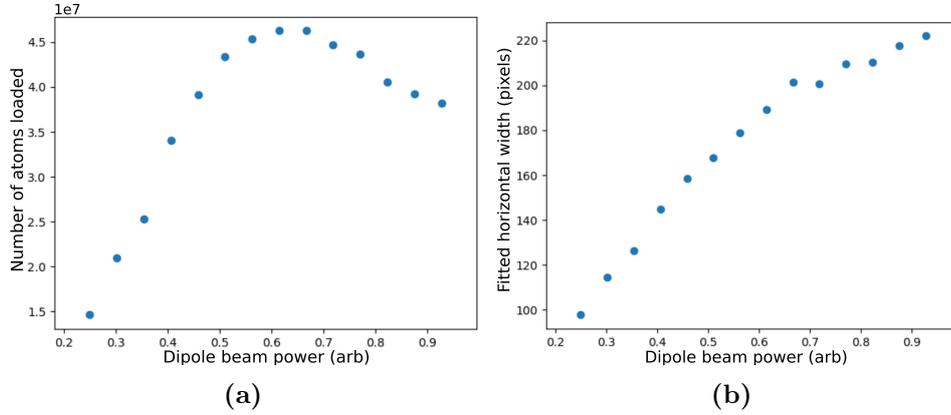


Figure 6.12: Hybrid load dependence upon dipole power **a)** and **b)** show how the number of atoms loaded into the hybrid trap and the waist of the cloud after a 25 ms drop time, used as a proxy for temperature, depends upon the power in the dipole beam. **a)** shows a clear increase in number as a function of power until it reaches a value of 0.62, after which it decreases, the decrease is thought to be attributable to an increase in 3 body losses caused by the increased density. **b)** shows a monotonic increase with laser power as expected.

time constant, and then held constant for the remainder of the decompression.

With this optimisation completed, the power in the dipole beam should also be scanned, figure 6.12a and 6.12b show the results of this scan. Interestingly, figure 6.12a shows that the loaded atom number decreases with increasing power after a value of 0.62⁴¹, corresponding to 4 W of power at the atoms, this is attributed to an increased 3 body loss rate arising from the increased density. Figure 6.12b shows how the waist of the cloud of atoms after being allowed to ballistically expand for 25ms varies as a function of laser power, this is used as a proxy for temperature. Unsurprisingly the waist increases with increasing laser power, again agreeing with the theory discussed in section 6.4.3, that the depth of the trap determines the final temperature of the atoms in the trap. These two data sets make an important point, however, that using the maximum available laser power does not necessarily lead to the largest atom number nor the largest initial phase space density.

With all parameters of the hybrid trap load optimised, we typically load 5×10^7 atoms at $\approx 8 \mu\text{K}$ into the hybrid trap, corresponding to an initial phase space density of 0.056. For comparison, a measure of the temperature and number of atoms in the magnetic trap when decompressed to 45 Gcm^{-1} with no dipole beam turned on was also made. This measurement indicated a temperature and atom number of

⁴¹The horizontal axis shows in arbitrary units the RF power applied to the AOM controlling the dipole beam, the power in the dipole beam is monotonic in this variable, but not linear.

8 μ K and 6×10^7 atoms respectively, showing that greater than 80% of the atoms are transferred into the hybrid trap, and indicates a phase space density in the quadrupole trap of 3.6×10^{-3} showing this large stage of evaporation does not produce a great increase in phase space density for the loss of atoms, most likely because the efficiency of the evaporation is greatly limited by the Majorana loss rate. This also shows the impact of changing the shape of the potential, as almost the same number of atoms at the same temperature are loaded into the dipole trap, but the phase space density increases by almost a factor of 16.

6.5.6 Characterising the hybrid trap

With the atoms finally loaded into the hybrid quadrupole/dipole trap, it is then time to characterise this trap. In the hybrid trap the potential along the length of the dipole beam is dominated by the quadrupole trap, whereas the dipole potential dominates the two transverse dimensions. As the dipole trap can be approximated as a harmonic trap, we can thus associate a trapping frequency with the harmonically trapped directions, these trapping frequencies, as well as knowing the heating rate and lifetime in the hybrid trap will fully characterise our trap.

As we saw in section 6.3.1 the trapping frequency depends upon both the power of the laser as well as the waist of the beam, thus both these properties must be known to know the trapping frequency. To measure the power in the dipole beam at the location of the atoms, the power before and after the cell can be measured. These measurements can then be used to estimate the power loss from the reflections off the windows of the cell, and the power at the atoms can then be estimated. The initial power in the dipole trap is estimated to be 5 W at the atoms, based on the power measured with a power meter before the cell, and then accounting for the losses a lower power beam experiences upon transmission from where the high power measurement was made⁴² and losses from reflection off the glass cell.

Measurement of the waist of the dipole beam at the atoms is more difficult, however, as the beam can only be profiled outside of the cell, and thus the effect the thick glass windows have upon the focussing dipole beam, as well as the exact alignment of the dipole beam onto the atoms, cannot be accounted for. Instead we can measure the trapping frequency directly from the atoms by creating a sloshing motion in the cloud. This is performed by using a bias field or increasing the quadrupole

⁴²A lower power beam is used for reasons of practicality, the power meter capable of measuring 5 W is too large to fit into most places.

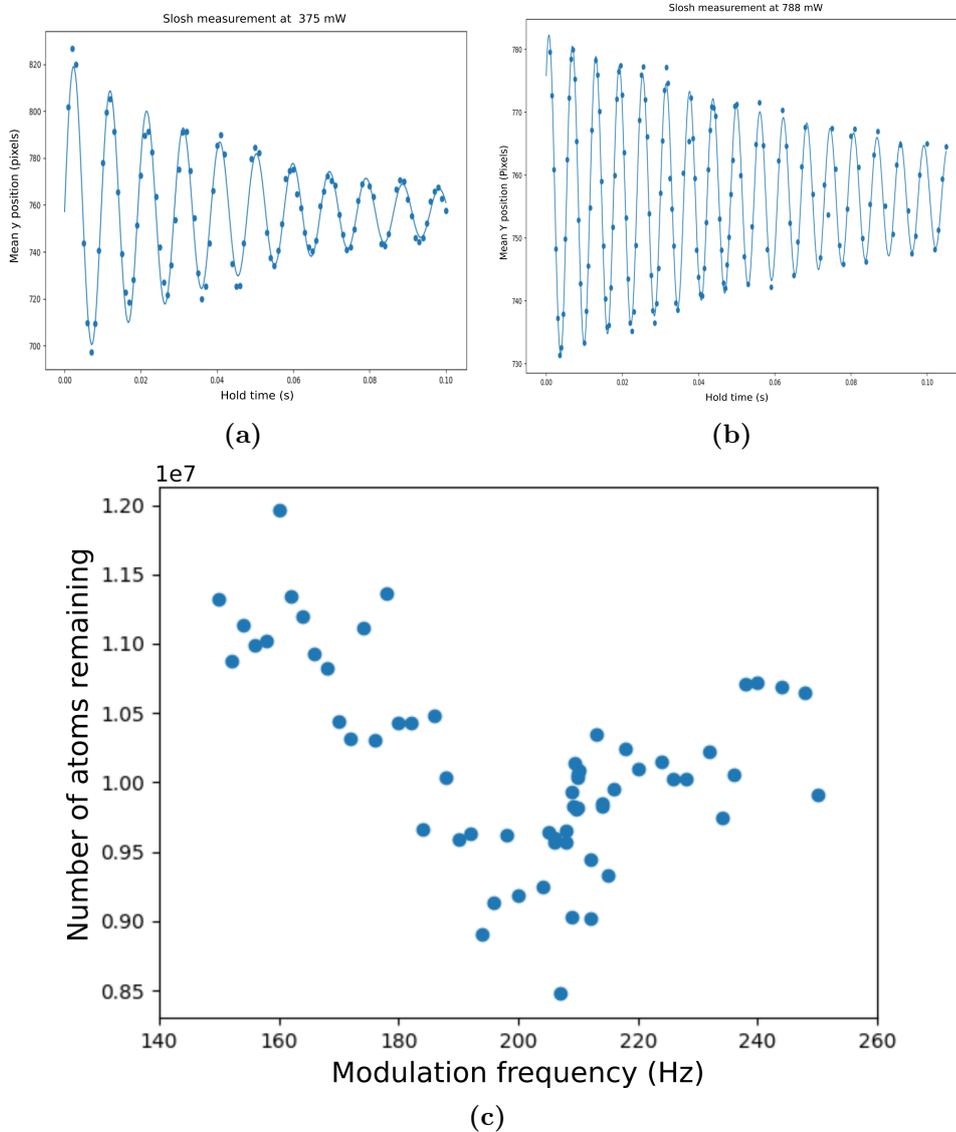


Figure 6.13: Hybrid trap frequency measurements a) and b) shows for 2 different dipole powers the measured (dots) mean position in the vertical direction of a cloud of cold atoms in the hybrid trap after displacing the atoms a small amount from the trap centre, waiting a variable amount of time indicated by the horizontal axis, and then allowing them to ballistically expand for 25 ms. There is a clear sinusoidal sloshing motion of the atoms, the frequency of which is equal to the trap frequency. Solid curves are decaying sinusoidal fits to the data. c) shows the measurement of the trapping frequencies of the hybrid trap using parametric resonance, there is a clear decrease in atom number when the frequency of modulation is twice that measured using the sloshing motion, but the decrease is very wide in frequency space.

gradient slightly to displace the atoms a small amount from the trap centre, and then releasing them and measuring the mean position of the cloud after a range of different dropping times. Figure 6.13a and 6.13b show the resulting oscillations for two different dipole trap powers. Plotted is the vertical position of the centre of a Gaussian fit to the atom as a function of hold time in the trap since the displacement after 25 ms of ballistic expansion. The data shows a clear sloshing motion, and the decaying sinusoidal fits to the data (solid lines) indicate trapping frequencies of 162.5 Hz and 104 Hz for powers of 788 mW and 375 mW in the dipole beam respectively.

By comparing the ratio of these trapping frequencies to the ratio of the power in the beams we also see the square root dependence of the trapping frequency on the power as calculated in equation 6.32. Also, by using the data obtained for a dipole power 788 mW and a trapping frequency of 162.5 Hz and equation 6.32, we can calculate the expected waist of the focussed trap. The calculated waist was found to be 60 μm , in good agreement with the measured waist via reflecting the beam with a mirror before it enters the cell, profiling the beam with a camera and fitting a Gaussian beam to the resultant image, which measured a waist of 55 μm and 65 μm in the vertical and horizontal directions respectively⁴³.

Another method that can be used to measure the trapping frequencies is via parametric heating. Parametric heating in a harmonic oscillator occurs when a property of the trap is modulated at twice the natural resonance frequency of the trap, which results in considerable energy being injected into the system. An everyday example of a parametric resonance is when one is in a swing and uses their legs to increase their swinging height, by swinging their legs twice as fast as the swing is swinging they will gain amplitude much faster.

To perform parametric heating, the power in the dipole trap is modulated by 8%⁴⁴ in a sinusoidal fashion about its mean power. This modulation occurs for 1 s and the cloud is then released from the trap and allowed to ballistically expand for 35 ms before being imaged. The frequency of the modulation is then changed, and the experiment repeated. Figure 6.13c shows the result of the parametric heating upon the atom number for a dipole power of 375 mW. There is a clear decrease in the number of atoms remaining in the trap after it has been parametrically excited for a wide range of frequencies about ≈ 210 Hz, where we would expect the resonance,

⁴³Note that when calculating the in-situ waist size via this method, the largest amount of power available should be used as this will minimise the effect gravity has upon reducing the trapping frequency.

⁴⁴ $\pm 4\%$

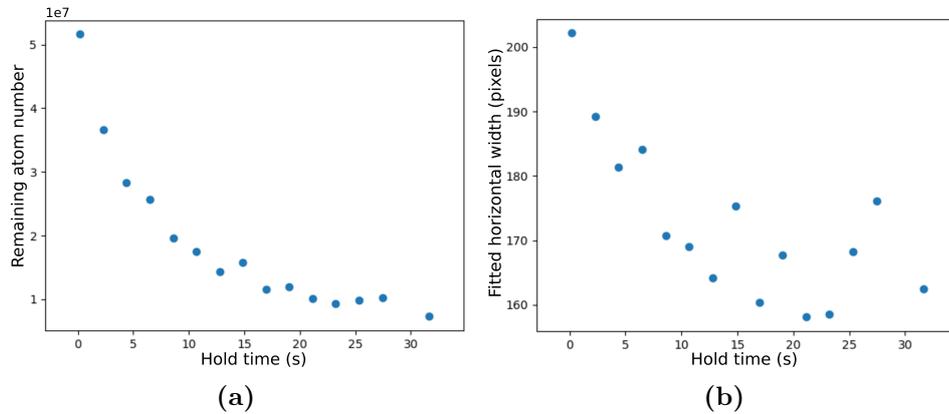


Figure 6.14: Hybrid trap lifetime a) shows the atom number decay as a function of time, there is a faster than exponential initial decay. The data in b) shows the variation in the waist of the cloud for the same experiments as in a), and show that the waist of the cloud decreases as a function of time indicating some ‘plain evaporation’ occurs. Thermometry after a 7 second hold agrees with this.

however, the width of this resonance is much larger than others have seen.

In an ideal parametric resonance measurement, the trapping frequency alone would be modulated, however, in a hybrid trap when the dipole potential is modulated the mean trap position will also move due to the influence of gravity. It is thought that this coupling of the movement in the position of the trap to the dipole power broadens the resonance.

The other key property that characterises the hybrid trap is the lifetime of the atoms in the trap, which will no longer likely be limited by background collisions, as off-resonant scattering will also remove atoms from the trap. Figure 6.14a shows a measurement of the number of atoms left in the trap as a function of time. There is clearly a fast initial reduction in the number of atoms, followed by a slower decay, as characteristic of an exponential decay in time, however when plotting the width of the cloud as a function of hold time, as shown in figure. 6.14b there is also a reduction in width during this time. This indicates that some ‘plain evaporation’ is occurring over the first 7 seconds in the dipole trap, a measurement of the temperature after 7 seconds shows it has reduced to $\approx 7 \mu\text{K}$. Fitting the data from 7 seconds onwards indicates a lifetime on the order of 15 seconds.

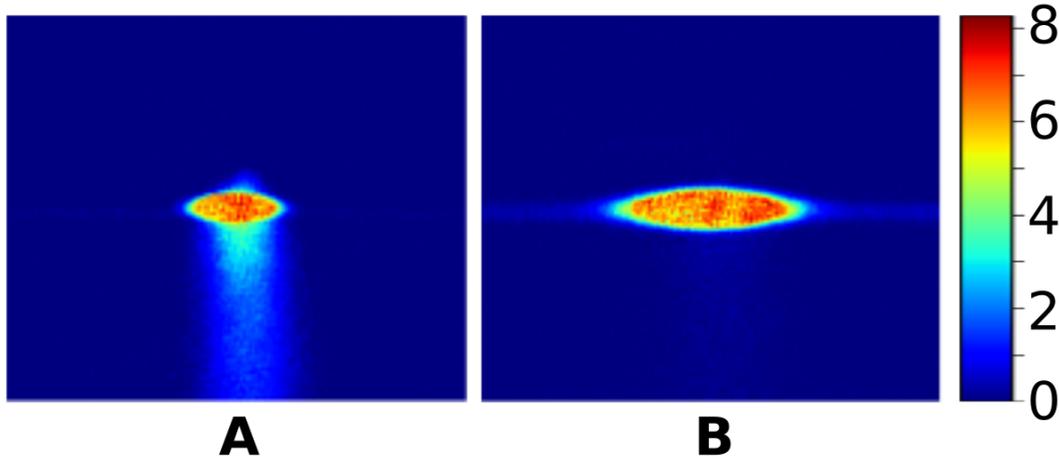


Figure 6.15: Optimal evaporation in the hybrid trap. a) and b) show how evaporation in the hybrid trap occurs. In a) the evaporation ramp is faster than the optimal rate, leading to evaporation occurring only in one dimension, along gravity, whereas in b) evaporation is occurring at the optimal rate, and atoms are seen to evaporate in 3-directions, along the length of the dipole beam as well as along gravity.

6.5.7 Forced evaporation in the hybrid trap

Forced evaporation in the hybrid trap proceeds by reduction of the dipole power. As the hybrid trap is the trap in which BEC will be made, optimisation of forced evaporation for condensates is performed as described in section 6.1 by optimising the parameter γ as defined in equation 6.13, or alternatively by optimising α as defined in equation 6.5⁴⁵, which as discussed is optimised by optimising α at each point during the evaporation.

To perform this optimisation, first the initial temperature and number of atoms in the hybrid trap is measured. Next a fixed time interval is chosen, this time interval is chosen to be long enough such that the optimal evaporation trajectory results in a reduction of the temperature of the atoms by a factor of 2-3⁴⁶. This reduction in temperature by a factor of 2-3 is chosen so that when scanning the evaporation parameters there is a clear maximum, as otherwise the uncertainty in α , arising from the uncertainty in the measured temperature and number, will obscure the

⁴⁵ α is typically an easier parameter to optimise experimentally, as it depends solely upon the change in number and temperature, and does not require detailed knowledge of the trapping potential.

⁴⁶As the optimal trajectory is not yet known there is obviously some guess work involved here, but in a hybrid trap with good initial conditions and lifetime, an evaporation time on the order of seconds is reasonable, with the time increasing as evaporation proceeds and the rethermalisation rate slows.

maximum⁴⁷. With an evaporation time chosen, the next step is to linearly⁴⁸ ramp the power in the dipole beam to a range of different values, measuring the final number and temperature of the atoms in the trap for each power, and calculating α for each evaporation ramp.

Figure 6.15 shows how evaporation proceeds in the hybrid trap, with the evaporating atoms been clearly seen as lines of atoms falling out of the trap. In figure 6.15 two different rates of evaporation are seen, in **a**) the evaporation rate is faster than required for the optimal efficiency, and it is clear that evaporation is mostly occurring in one direction, along gravity. In **b**) evaporation is occurring at the optimal rate, and evaporation is seen to occur in three directions, along the lengths of the dipole beam as well as along gravity. This increase in the area of the surface through which evaporation can occur leads to the increased efficiency, and is why levelling the dipole beam with respect to gravity is important, as otherwise evaporation would only occur along one direction of the dipole beam.

Evaporation ramp	Duration (s)	P_f (W)	N_f (10^7)	T_f (μk)	α
1	0.5	2.17	4	5	2.11
2	0.5	0.788	3	2.2	2.85
3	1	0.245	1.45	0.59	1.81

Table 6.4: Properties of the optimised evaporation ramps

Table 6.4 shows the results of this optimisation resulting in three linear ramps of the dipole power before the first signs of condensation were observed, as shown in figure 6.16. These ramps take the initial dipole power from 4 W down to 2.17 W then 788 mW and finally to 245 mW over durations of 0.5 seconds, 0.5 seconds and 1 second, respectively, and each ramp results in a temperature reduction of a factor greater than 2. Once the first signs of condensation occur this optimisation procedure is no longer used as accurate fitting of the temperature and number of atoms becomes difficult, and when considering atom losses one has to not only account for evaporated atoms, but has to take into account atoms that have condensed as well.

Forced evaporation in the hybrid trap began with a cloud of 5×10^7 atoms at $\approx 8 \mu\text{K}$. After evaporation, just on the edge of condensation, 1.45×10^7 atoms remained at

⁴⁷A larger change in temperature will make changes in α more apparent, but reduces the number of ramps used to optimise the evaporation.

⁴⁸Obviously there are other functional forms of the ramp that could be used, however, they all involve additional parameters that would need to be varied to find the optimal trajectory, creating a much larger space in which the optimisation is performed, I achieve good efficiency without the need for this additional complexity.

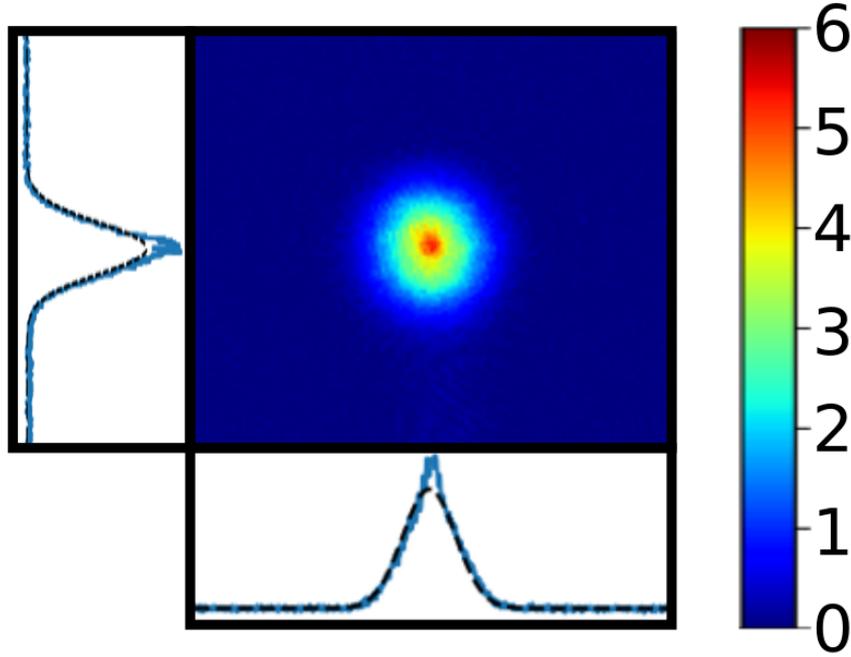


Figure 6.16: First sign of condensation. Shown are the very first signs of condensation occurring at the end of the 3rd evaporation ramp, indicated by the peaked centre of the cloud, the first signs of a bimodal distribution.

a temperature of ≈ 590 nK, a volume of $6.7 \times 10^{-6} \text{ cm}^3$ and thus average density of $1.85 \times 10^{14} \text{ atoms/cm}^3$ and a peak phase space density of 2.5^{49} , corresponding to an increase in PSD by a factor of 44, and thus a resulting evaporation efficiency, γ , of 3.06.

6.5.8 Bose-Einstein condensation

With the first signs of condensation observed we have almost achieved our main goal for the apparatus. As discussed, optimisation using γ as defined in equation 6.13 is difficult as one has to accurately count the number of atoms in the condensate as well as the thermal cloud, account for the increase in PSD due to the increase in condensed atoms, and still accurately measure the temperature of the thermal

⁴⁹Note that in this increase in PSD, one also has to account for the reduction in the trapping frequency due to the reduced power, one cannot simply compare the ratio of the temperature and number to the initial conditions, appendix. D shows the calculation used.

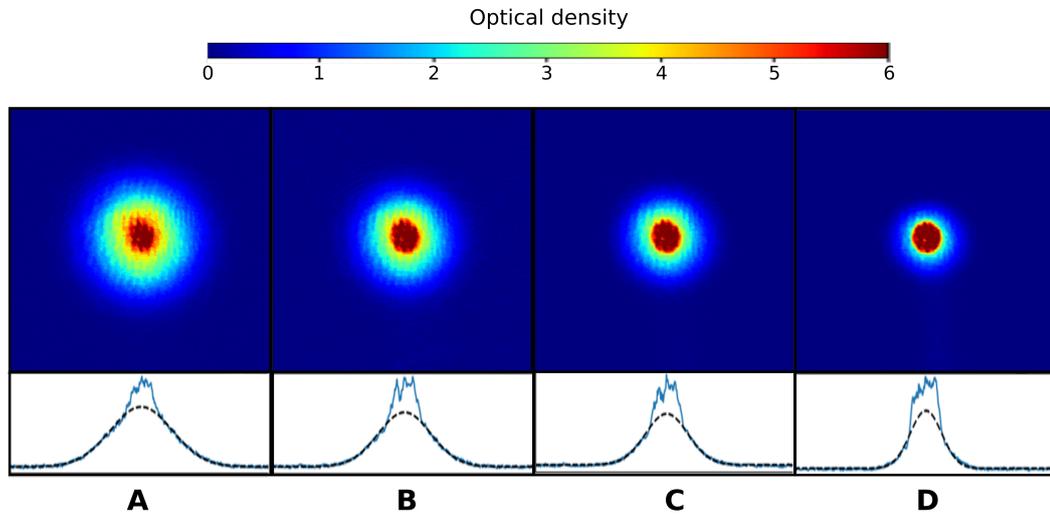


Figure 6.17: Emergence of bi-modality a), b), c) and d) show the emergence of a bimodal distribution in time of flight as the final dipole power is lowered. As the power is lowered the fraction of thermal atoms reduces whilst the condensate fraction increases, although this is difficult to see at this relatively short drop time (35 ms) as the OD becomes very large.

cloud. Instead I chose to optimise the large stage of evaporation in a single step, aiming for a ‘pure’ BEC⁵⁰, and maximising the number of atoms in the condensate.

Figure 6.17 shows the effect of reducing the laser power on the fraction of atoms in the condensate, with the power being reduced from figure a) through to d). Clearly as the laser power is reduced the fraction of atoms in the thermal component is reduced whilst the fraction of atoms in the condensate increases, however, at the short drop times used for these images (35 ms), the optical density of the condensate is too large and the measurement saturates, typically leading to underestimating the optical density.

Obviously condensation can occur at any power given a large enough number of atoms that are cold enough, and thus as evaporation becomes more efficient the power at which a ‘pure’ condensate will form is increased, thus ideally one wants to maximise the power at which the condensate forms. Optimisation of the final evaporation ramp is thus performed by scanning both the final power as well as the final evaporation ramp time, and maximising the number of atoms in the condensate, or alternatively by maximising the power at which a pure condensate forms, as this should also maximise the number of atoms in the condensate. As the metric for

⁵⁰Here I consider a condensate fraction greater than 80% to be pure.

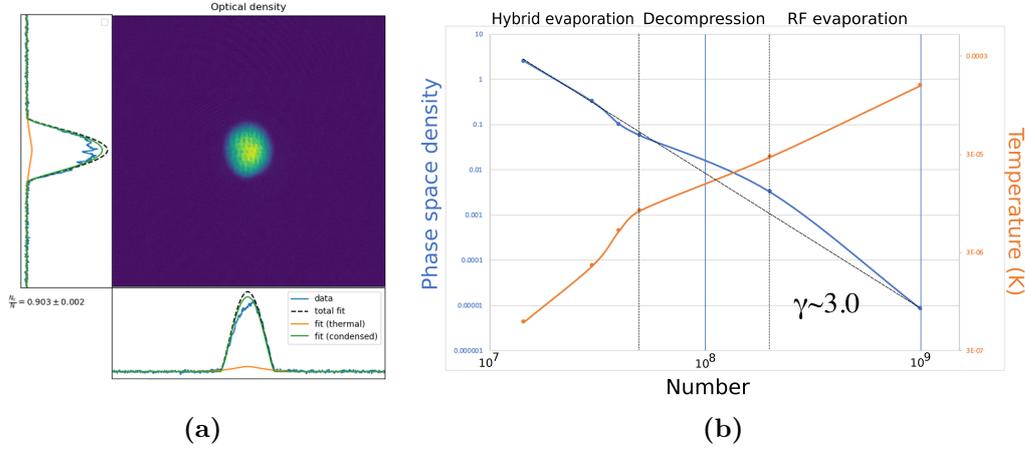


Figure 6.18: Final BEC and PSD versus N **a)** shows the final condensate after 100 ms of ballistic expansion. A Thomas-Fermi and Gaussian fit to the data indicate a 90% pure condensate with $\approx 2.2 \times 10^6$ atoms, produced in just under 10 seconds. **b)** shows how the phase space density and temperature varies as a function of atom number during forced evaporation in the magnetic trap and the hybrid trap, the different stages being indicated by the dotted lines. All axes are logarithmic, with the horizontal axis indicating the atom number, the left hand side vertical axis the phase space density, and the right hand vertical axis the temperature. The data shows an almost linear trend throughout the entire process, with an average evaporation efficiency of $\gamma = 3$, indicated by the dashed line.

this optimisation, a drop time of 100 ms is used to reduce the optical density of the condensate to ≈ 4 , so the number can be accurately measured, and a Gaussian fit to the thermal cloud as well as a Thomas-Fermi fit to the condensate was performed to measure the condensate fraction as well as the atom number.

Figure 6.18a shows the final condensate once the optimisation ramp was optimised, resulting in a 1 second ramp from 245 mW down to 45 mW, a further reduction in power by a factor of 5.45, followed by 0.2 seconds of free evaporation. This evaporation results in a condensate of 2.2×10^6 atoms, and given equation A.5 and a condensate fraction of 90% suggests $\frac{T}{T_C} = 0.45$.

Figure 6.18b shows how the phase space density and temperature of the cloud varies with atom number throughout the entire trajectory of the experiment up to the first sign of condensation. All axes display a logarithmic scale, with the left hand side vertical axis indicating the phase space density, the right hand vertical axis the temperature and the horizontal axis the atom number. Dotted lines are used to separate the three different stages, forced evaporation in the compressed magnetic trap, decompression to load the hybrid trap and forced evaporation in the hybrid trap, beginning from the right hand side. A dashed line also shows the average

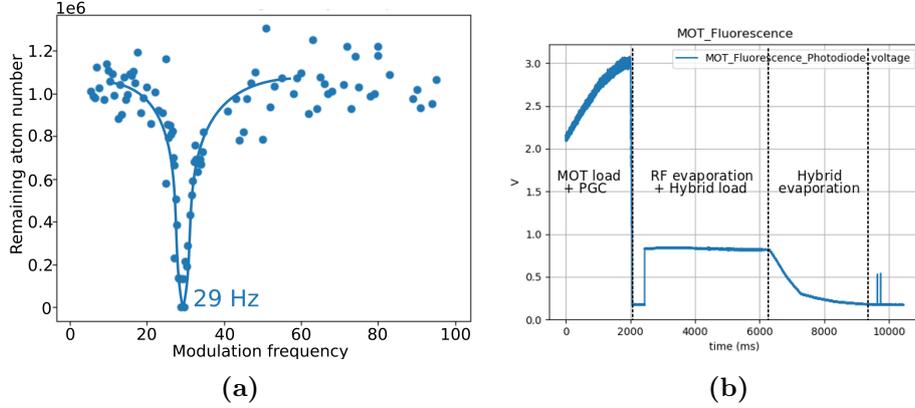


Figure 6.19: Final BEC characterisation **a)** shows the result of modulating the dipole beam potential for 1 second with a 2% modulation at various frequencies. Only one clear resonance was observed at 29 Hz, which coincides with the trap frequency as measured by a dipole mode oscillation measurement. **b)** shows the voltage measured by a photodiode used to monitor the MOT fluorescence during an experiment. The first 2 seconds show an increase in fluorescence as the MOT loads before suddenly switching off during the PGC stage. After a few hundred milli-seconds the Hybrid dipole beam is switched on and the scatter from the beam is picked up by the photodiode. From ≈ 6.5 s to 9.5 s the power in the dipole beam is seen to ramp down as hybrid evaporation occurs, followed by two pulses that correspond to the imaging beam switching on for absorption imaging.

evaporation efficiency, $\gamma = 3.0$, across the entire trajectory. Clearly evaporation in the magnetic trap is initially more efficient, however, as we saw earlier this efficiency drops remarkably once Majorana losses become large, resulting in a decreased efficiency during the decompression and loading of the hybrid trap. Evaporation in the hybrid trap then begins with a slightly greater efficiency than the average, before decreasing during the last stage of evaporation.

To characterise the BEC, we need to measure both the trapping frequency of the potential, and the Thomas-Fermi radius of the condensate, and then through use of equation A.11 we can calculate the chemical potential. To measure the Thomas-Fermi radius of the condensate an image from both the side and bottom are taken of the condensate, effectively *in-situ* (after 2 ms drop). The measured Thomas-Fermi radii from these images were $81\ \mu\text{m}$, $91\ \mu\text{m}$ and $44\ \mu\text{m}$ along the Z, Y and X axes, respectively, where the Z axis is along gravity, the Y axes along the optical axis of the dipole beam and the X axis is transverse to the optical axis of the dipole beam and perpendicular to gravity. This indicates an approximate 2:2:1 ratio of the condensate in the Z, Y and X dimensions, respectively. One may assume that the Thomas-Fermi radius of the condensate in the X and Z directions would be the same given the dipole beam dominates the potential in these directions, however, as

gravitational sag forces the trap potential and trapping frequency to reduce to zero at a finite power, this is not the case. These Thomas-Fermi radii indicate the relative trapping frequencies along the different axes, indicating the trapping frequency in the vertical direction should be approximately half of that along the other two axes.

To measure the trap frequency we can again perform a dipole mode oscillation measurement. This was performed along the X-direction, and agrees well with the expected trap frequency of ≈ 30 Hz, as calculated via equation 6.32 and the measured trap frequencies at higher powers. A parametric resonance type measurement was also performed by modulating the Hybrid beam power by 2% for 1 second, which resulted in the data in figure 6.19a. This data shows substantial atom loss at 29 Hz in a narrow peak, but no other features where we might expect them.

This peak could be attributed to the expected 15 Hz trapping frequency along the vertical direction, this trapping frequency being confirmed by a dipole mode oscillation measurement. However, I expect that this resonance is more likely due to coupling with the dipole oscillation mode in the X direction causing loss, as reduction in the laser power will exaggerate the gravitational sag in the potential, coupling the modulation of the trap frequency to a modulation of the position of the trap minimum in the vertical direction, causing atoms to roll down along the length of the dipole beam to the new trap minimum. This may also explain why no parametric resonance was observed at 60Hz, twice the 30 Hz trap frequency in the X direction, as the coupling of the modulation of the trapping frequency to the position of the atoms prevents a parametric resonance measurement being made.

To produce a BEC of 2×10^6 atoms an experiment consists of 2 seconds of MOT load followed by 2 seconds of evaporation in the compressed magnetic trap, a further 1.5 seconds of decompression to load the hybrid trap and finally 3 seconds of evaporation in the hybrid trap, resulting in an entire experiment runtime of about 9 seconds. This entire sequence can be ‘seen’ from a photodiode signal, shown in figure 6.19b, which was initially used to monitor the fluorescence of the MOT. The photodiode signal is seen to increase in the first two seconds, corresponding to an increasing number of atoms in the MOT as it loads, after which the CMOT stage and PGC occur before the MOT beams are finally switched off. Next at $t=2.3$ s we see the scatter of the Hybrid dipole beam from the glass cell as we suddenly switch it on, followed by it being held on at constant power for ≈ 3.5 s, before being ramped down over 3 seconds to its final value to force evaporation in the hybrid trap. Finally, just before $t=10$ s we can see two pulses corresponding to scatter from the imaging beam as we image our result.

My BEC production time is almost 40% faster than Lin *et al.*, which is attributed to my much faster MOT load rate, as I typically load 3 times as many atoms in the MOT loading stage than they do in the same amount of time, allowing faster evaporation in the compressed magnetic trap leading to a larger phase space density for faster loading of the hybrid trap, and finally a faster evaporation in the hybrid trap.

Imaging cold atoms

One of the key questions we ask when probing cold atoms is “how cold?” As cold atom samples are so small in number compared to most macroscopic objects, even in the case of the smallest probes available the atoms cannot be considered a thermal reservoir, and thus measurement through contact probes would quickly find the atoms equilibrating to the probe’s temperature, not vice-versa, and thus information about the temperature of the atoms cannot be obtained this way. We saw in section 5.2 a number of other ways one could measure the temperature of the atoms. These techniques utilised a single probe beam causing fluorescence in the atoms which was detected, and then probe beam parameters such as detuning or position were changed and the dependencies of the fluorescence on these parameters were measured. From these measurements the temperature could be estimated. Nowadays the most common technique is to use time of flight expansion, measuring how the atomic cloud ballistically expands during time of flight, however, this requires spatially resolved imaging over a wide field of view.

The other key question one asks when trying to achieve Bose-Einstein condensation is whether you have achieved condensation or not? One of the first characteristic signs of condensation is a bimodality seen in time of flight measurements, distinguishing the condensate from the thermal atoms. Another key characteristic is the distinctive Thomas-Fermi profile arising from the atoms adopting the shape of the potential as opposed to a Gaussian distribution, or alternatively in an anisotropic potential seeing an aspect ratio inversion as the tighter dimension expands more rapidly. All these signs require some form of spatially resolved imaging also.

Optically probing cold atoms is a versatile technique through which most of our measurements of cold atoms have been made, typically in which an image is taken of the atoms that reflects the density distribution of the atoms. To understand how we can use optical probing to gain quantitative information from the atoms, we must first understand how atoms interact with light.

7.1 Optical interactions

There are three processes we must consider when discussing the interaction of atoms with light; the absorption of photons, the spontaneous emission of photons and the effect of the atoms on the phase of light that passes through them, these three processes are depicted in figure 7.1. All of these interactions can be encapsulated into a single parameter, the complex index of refraction, n_{atom} , given by

$$n_{atom} = \sqrt{1 + 4\pi n\alpha}, \quad (7.1)$$

where α is the complex atomic polarizability of the atoms, as we saw in section 6.3.1, and n is the number density of the atoms. In section 6.3.1 we were mostly concerned with the real component of the atomic polarizability as this describes the dispersive interaction of the atoms with the light field, however, here we shall be more concerned with the imaginary component describing the attenuation. In the case where $n_{atom} - 1 \ll 1$, the index of refraction can be written as

$$n_{atom} = 1 + \frac{\sigma_0 n \lambda}{4\pi} \left(\frac{i}{1 + \delta^2 \frac{I}{I_{sat}}} - \frac{\delta \sqrt{\frac{I}{I_{sat}}}}{1 + \delta^2 \frac{I}{I_{sat}}} \right), \quad (7.2)$$

where all symbols carry their usual meaning, σ_0 being the on resonance cross section, λ the wavelength, δ the detuning from resonance expressed in half linewidths¹ and I and I_{sat} are the light intensity and saturation intensity, respectively.

The form of equation 7.2 is useful as it separates the contributions of the real and imaginary part of the refractive index. The imaginary part of the refractive index is clearly the first term in the brackets, and is responsible for the absorption and thus attenuation of the optical field, whereas the second term accounts for the dispersive nature of the atoms when interacting with an optical field that is off resonance. Note

¹i.e. $\delta = \frac{\omega - \omega_0}{\Gamma/2}$, where ω and ω_0 are the frequency of the imaging light and the transition frequency, respectively, and Γ is the linewidth of the transition.

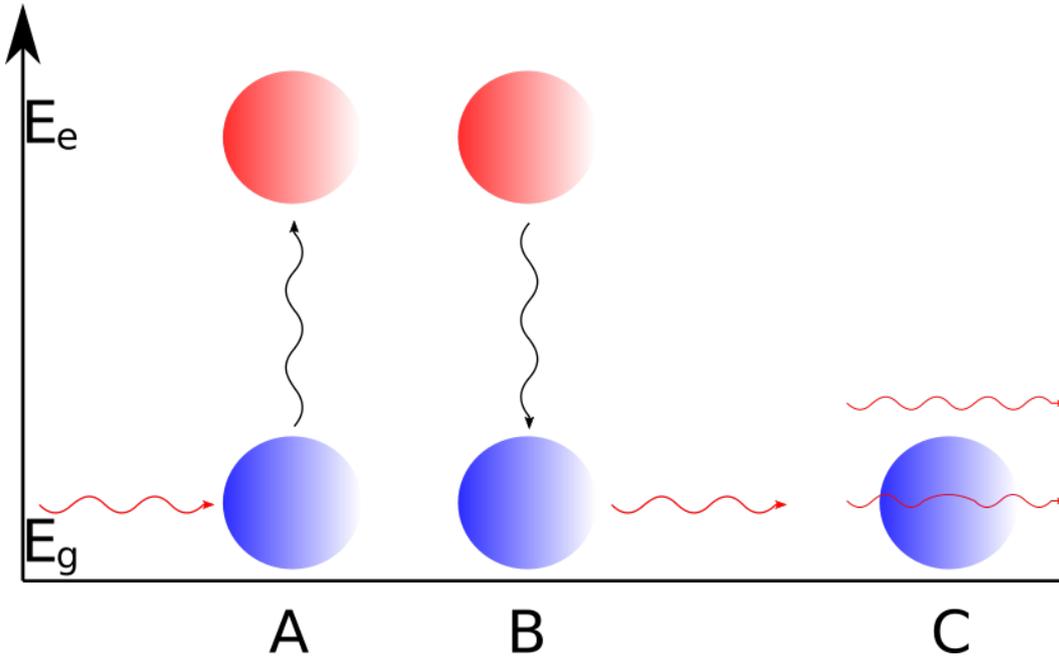


Figure 7.1: Interaction of photons with atoms **A**, **B** and **C** show the three types of interactions photons have with an atom. In **A**, the absorption of a photon by an atom is shown, exciting the electron in the atom into a higher orbital. **B** shows the process of spontaneous emission of photons as an atom in an excited state decays to its ground state. **C** shows how the phase of a photon as it passes through an atom is shifted with respect to a photon passing through vacuum.

that as the dispersive term is proportional to the detuning, it clearly vanishes for zero detuning, whereas the absorptive term is largest for zero detuning.

The simplest approximation we can make for the interaction of the optical field with the atoms is the equivalent of the thin-lens approximation in optics, that the light that enters the atoms at coordinate (x,y) will also exit at the same coordinate, and only be attenuated and phase shifted. In this case the atoms acts as a non-transparent phase object. If the incident optical field can be described by an electric field $E_0(x, y)$, then the light exiting the atoms can be described by [145]

$$E_T(x, y) = tE_0(x, y) \exp(i\phi) \quad (7.3)$$

where we have introduced the transmission coefficient, t and the phase shift, ϕ .

The transmission coefficient describes the attenuation of the optical field by the atoms, and is simply given by the imaginary component of $\frac{4\pi n_{atom}}{\lambda}$. In the thin-lens approximation this can be calculated by integrating the atomic density along the

optical axis and replacing the atomic density, $n(x, y, z)$, with the integrated column density, $\int \sigma_0 n \cdot dz$ ² thus the transmission coefficient is given by

$$t = \exp\left(-\frac{\int \sigma_0 n \cdot dz}{2} \frac{1}{1 + \delta^2 \frac{I}{I_{sat}}}\right). \quad (7.4)$$

A common quantity that is used experimentally is the off resonance optical density, OD , given by

$$OD = \frac{\int \sigma_0 n \cdot dz}{1 + \delta^2 \frac{I}{I_{sat}}}, \quad (7.5)$$

this is particularly useful, as the transmission coefficient can then be expressed as $t = \exp(-OD/2)$, yielding the simple interpretation of the OD as twice the log of the fractional reduction in the electric field amplitude through the sample. i.e. $OD = -2\log(t)$, or alternatively, the log of the fractional reduction in intensity. The phase shift can be similarly expressed in terms of the off resonance optical density, or alternatively through the dispersive component of the complex refractive index of the atoms,

$$\phi = -\frac{\delta OD}{2} = \left(\frac{\int \sigma_0 n \cdot dz}{2} \frac{\delta \sqrt{\frac{I}{I_{sat}}}}{1 + \delta^2 \frac{I}{I_{sat}}}\right). \quad (7.6)$$

Now as the integrated density plays a role in both the phase shift and the transmission coefficient, spatially resolved measurement of either of these two quantities allows one to infer the atomic density distribution integrated along the optical axis, through measurement of the optical density.

Measurement of the phase shift induced by the atoms typically requires a more complex setup than standard absorption imaging, as cameras are not sensitive to the phase of the optical field and thus must have the phase converted into an amplitude modulation to allow it to be measured. There have been several methods which have been demonstrated that achieve this, one such method is dark ground imaging[146]. In dark ground imaging, the atoms are illuminated by a collimated laser beam, and then reimaged onto a camera, for simplicity we will consider the case of a pair of lenses used to reimage the atoms in a 4f configuration. In the Fourier plane of the first lens, the unscattered light will come to a focus where it can be blocked, whereas the scattered light will be collimated, the second lens then reimages the atoms onto the camera.

In a similar manner, phase contrast imaging of cold atoms has also been demonstrated[147]

²Here we assume the optical axis is aligned with the z-axis.

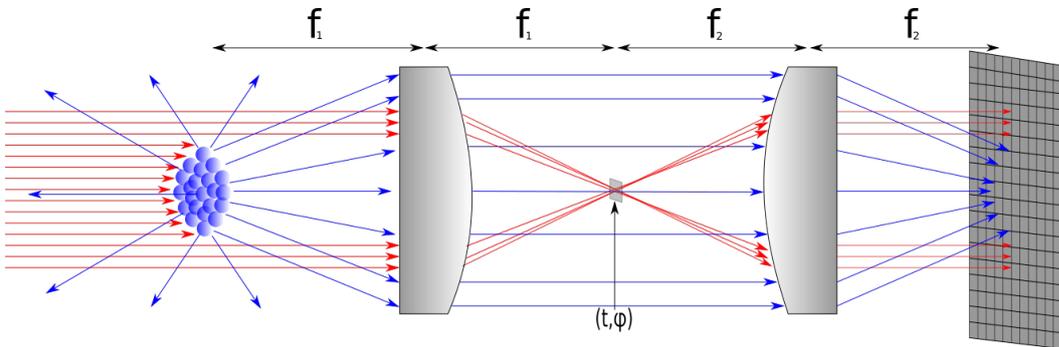


Figure 7.2: Generalised imaging scheme A generalised imaging scheme is shown above. A probe beam is shown to be incident on the atoms from the left hand side (red arrows), which scatter some of the light (blue arrows). These photons are then captured with a lens and reimaged onto the camera using a 4f configuration. In the Fourier plane of the first lens, a grey box represents either an amplitude (t) or phase (ϕ) modulating element. In the case of standard absorption imaging, the element is such that $(t, \phi) = (1, 0)$, in the case of darkground imaging $(t, \phi) = (0, 0)$ and in the case of phase imaging typically $(t, \phi) = (1, \pi/2)$.

which utilises modulation of the unscattered radiation in the Fourier plane of the first lens, however, this time by placing a well characterised phase object, which typically shifts the phase of the light by $\pm\pi/2$. The modulated unscattered light then interferes with the scattered light on the camera, allowing the phase induced by the atoms to be measured through the interference pattern. These methods require additional complexity to the optical setup and are most useful if one is wanting to use a non-destructive imaging technique³, however, in this work this was not of concern. Thus in this thesis I have solely relied upon either measurement of the transmission coefficient, or alternatively the fluorescence of an atomic sample.

A generalised setup in which absorption, dark ground imaging or phase contrast imaging can be performed is shown in figure 7.2. Shown on the left hand side of the figure is the incident probe beam indicated by the red arrows. As the light interacts with the atoms it is absorbed and re-emitted (blue arrows), or phase shifted by the atoms, for clarity the phase shifted light is not shown. Of the light that is fluoresced by the atoms some of it is captured by the imaging system, in this case shown in a 4f configuration, and is reimaged onto the camera. In a standard absorption imaging setup no other optical components are required. In the case of phase contrast imaging a phase object is placed at the focus of the probe beam, in the focal plane of the first lens, which phase shifts the transmitted light, typically

³As the phase shift is proportional to the detuning, the detuning is typically many linewidths, sometimes even up to hundreds of linewidths[148], which greatly reduces the absorptive component of the refractive index, minimising heating of the sample.

by $\pi/2$. In the case of dark-ground imaging a small opaque object is used to block this light instead.

7.2 Absorption imaging

On-resonance absorption imaging is a particularly quick, simple method of obtaining relatively accurate quantitative atom numbers, as the signal to noise is maximised for zero detuning⁴ and as there is no phase component, one does not have to worry about lensing effects that arise from the varying phase shift imparted to the optical field. To calculate the optical density of the cold atom cloud through an absorption image, one only need invert equation 7.5, to find

$$OD(x, y) = -\ln\left(\frac{I_{out}(x, y)}{I_{in}(x, y)}\right) = \sigma\tilde{n}(x, y), \quad (7.7)$$

where I_{out} and I_{in} are the incident and transmitted light intensity respectively, and $\tilde{n}(x, y)$ is the integrated column density, and we have assumed on-resonant light. Using the second equality in the above equation allows the integrated column density to be inferred from the optical density.

The main criticism of on-resonant absorption imaging is its dynamic range, which is typically limited by the bit-depth of the camera used, for example consider our case of a camera with a bit depth of 14. The counts of this camera can thus range from 0-16383, and thus to maximise the dynamic range of the camera, the exposure of the atoms by the laser should be adjusted to reach the maximal counts recorded by the camera. Now, for my camera the typical count rate for an image with no light is ≈ 1000 counts, with a deviation of ≈ 100 counts across the chip. Thus this allows a maximal attenuation of the laser light from ≈ 15000 counts to ≈ 100 counts before the change in the laser light will be lost due to the noise on the camera. This change in laser intensity corresponds to an OD of $\approx 2^5$, indicating this is the maximal OD we can accurately detect. However, in a range of situations, such as *in-situ* absorption imaging, or even short-drop time (< 20 ms) ballistic imaging, OD's can range be in the range of 10-300, making accurate measurement of the OD somewhat difficult. This issue can be circumvented by allowing the cloud to ballistically expand for longer enough such that the peak OD drops to ≈ 2 , but this

⁴As the most photons are scattered at zero detuning.

⁵Note that the absorption measured by the camera is given by t^2 , as defined in equation 7.4, not t , as the camera measures $|E|^2$ not E .

requires a field of view wide enough to encompass all of the atoms, which typically prevents accurate imaging of small objects⁶.

If one were to instead want to maximise the dynamic range of the optical density the camera can detect, one could image off resonance and choose a particular value of detuning such that the resulting off-resonance OD is approximately 1 and yields $\approx 63\%$ absorption.⁷ When the OD of the cloud is much greater than unity, the optimal detuning for maximal contrast is given by $\delta = \sqrt{OD}$. However, in this case the atoms will refract the light as we are no longer on-resonance, since the real component of the atomic refractive index is non-zero. The angle at which the cloud refracts the light can be estimated as $\frac{2\lambda\phi}{\pi d}$, where d is the size of the cloud and ϕ the maximal phase shift imparted by the cloud. Now, due to the finite size of the atoms, the atoms will also diffract the light at an angle given by $\frac{\lambda}{d}$, thus if the phase shift is kept smaller than $\pi/2$ then the refraction angle will remain smaller than the diffraction angle, and the spatial resolution of the imaging system will be limited to its diffraction limit. However, if the phase shift is larger than $\pi/2$, then the spatial resolution of the imaging system will be degraded and not all of the refracted rays will be collected by the imaging system, leading to a loss in signal apparent as an increase in the absorptive signal. This increase in the absorptive signal will vary as a function of the phase shift, and thus will vary across the cloud, affecting both relative and absolute density measurements.

An alternative method to image with on-resonant light is to use high-intensity imaging, where the light intensity is greater than the saturation intensity of the transition, and thus saturation effects cannot be neglected. In the case of high intensity imaging, equation 7.7 is modified to become

$$OD(x, y) = -\ln\left(\frac{I_{out}(x, y)}{I_{in}(x, y)}\right) + \frac{I_{in}(x, y) - I_{out}(x, y)}{I_{sat}(x, y)} = \sigma\tilde{n}(x, y). \quad (7.8)$$

There is now an additional term which is linear in the difference of the input and transmitted intensity, and depends on the saturation intensity. In the case of low intensities compared to the saturation intensity the linear term vanishes and we arrive back at equation 7.7.

We can now utilise this linear term to improve our imaging, as by increasing the

⁶As to have such a large field of view, either a camera with an extremely high number of pixels is required, or alternatively the magnification of the optical system will be low such that each pixel is quite large, making it difficult to resolve small features.

⁷As this will maximise the contrast for the desired range of OD's.

intensity of the light to be greater than the saturation intensity, we can effectively make the atoms more transparent, as the rate of incident light becomes greater than the rate at which the atoms can scatter the photons. Let's elucidate this with an example. Consider the case of a dense cloud that has a large optical density and let's calculate the light intensity required, in multiples of the saturation parameter, to have an absorption of only 50% of the light, i.e. $\frac{I_{out}}{I_{in}} = \frac{1}{2}$, and thus $I_{in} - I_{out} = \frac{I_{in}}{2}$. Substituting the above relations into equation 7.8 yields

$$OD(x, y) = \ln(2) + \frac{I_{in}}{2I_{sat}} \quad (7.9)$$

and rearranging the above, and noting that we have assumed that $\ln(2) \ll OD$ yields

$$I_{in} = 2OD \times I_{sat}, \quad (7.10)$$

i.e. to have approximately half of the probe light absorbed by the atoms, the intensity of the imaging beam should be larger than the saturation intensity by approximately twice the optical density of the atoms. Note that this relation holds for on-resonance imaging when the optical density of the atoms is much greater than 2.

7.2.1 Measuring saturation intensity

High intensity absorption imaging has been the workhorse imaging technique for quantitative atom numbers in my apparatus, however, to utilise this technique, accurate knowledge of the absolute intensity of the probe light is required, not simply the ratio of the incident to transmitted intensity as is the case with low intensity imaging of dilute clouds. To measure the intensity of the imaging light a few different methods have been developed.

The simplest but typically least accurate is to measure the incident probe power with a power meter, account for any transmission loss through the optical system, and then infer the pixel count on the camera that corresponds to the saturation intensity. This method, however, introduces a number of uncertainties, such as the uncertainty associated with the power meter, your measurements of transmission loss through the optical system as well as accurate knowledge of the beam profile.

An alternative method suggested recently[149] is to instead use the imaging beam as effectively a kicking beam to impart momentum to the atoms during free-fall then infer the momentum imparted by measuring the displacement of the cloud after some

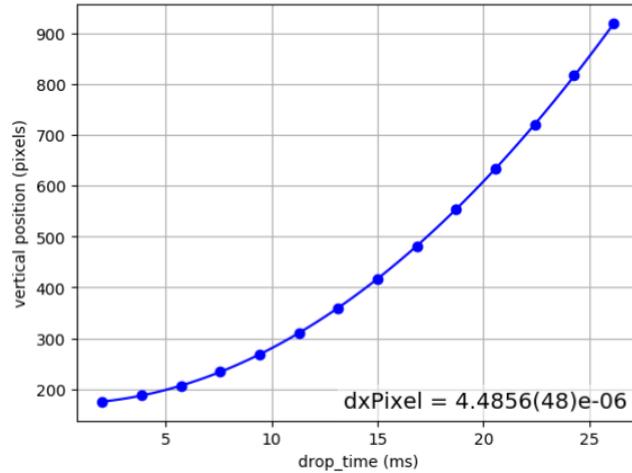


Figure 7.3: Side imaging calibration A calibration of the effective side imaging camera pixel size is shown above, performed by dropping a cold cloud for various different drop times, and fitting the displacement along the vertical axis.

fixed time as a function of beam power. The momentum imparted to the cloud will depend upon the polarisation, detuning and intensity of the probe light. Hueck *et al.* showed that this method could be used to calibrate all three parameters, finding the optimal polarisation, detuning, and the pixel count rate that corresponds to the saturation intensity. One drawback of this method is the complexity of the setup required to perform the calibrations, as one requires independent control of the probe beam one is attempting to calibrate, as well as an imaging system that is transverse to that probe beam used to measure the displacement of the cloud, as such I have not performed this measurement.

Another method first suggested by Reinaudi *et al.*[150], measured the saturation parameter by performing absorption imaging at a range of different probe intensities, but fixed duration, with a fixed number of atoms in identical conditions. As the number of atoms and their conditions are fixed, the optical density is thus also fixed, and thus each image should result in the same calculated optical density if the correct pixel value for the saturation intensity is used. Reinaudi *et al.* then found this value by performing a least squares fit to their data, minimising the differences in calculated optical density between shots. The benefit of this method of calibrating the imaging saturation parameter is that it is not as prone to the uncertainties of the first method suggested, however, the drawback is that it requires consistent conditions for the atoms. This is the method I employ to measure the saturation intensity in terms of pixels.

7.2.2 Measuring imaging system magnification

Clearly from equation 7.8 the physical size of the cloud must be known, not simply the size in pixels. As such, an accurate calibration of the magnification of the optical system must be known.

To calibrate the imaging system when the optical axis is perpendicular to gravity, measurement of the mean position of the cloud as a function of drop time can be used. This is as once all trapping potentials are switched off the cloud will accelerate due to gravity and fall along the vertical axis. By fitting the mean position of the cloud as a function of drop time, the mean position should obey

$$s = \frac{1}{2}gt^2 \quad (7.11)$$

where s is the displacement of the mean position from its *in-situ* position, g is the acceleration due to gravity, and t is the drop time. Now, we can rewrite the displacement in position as the product of the displacement in pixels, dy , and the effective size of a pixel⁸, px , i.e. $s = dy \times px$. Using this, along with the knowledge that $g = 9.8 \text{ m s}^{-2}$ ⁹, we can then fit the mean position as a function of drop time using equation 7.11, with the only free parameter being px , and thus obtain the effective pixel size.

Figure 7.3 shows the result of this calibration for the side imaging, which indicates a pixel size of $4.48 \mu\text{m}$ corresponding to a magnification of 1.01 ± 0.01 , agreeing with the expected value. Clearly the data agrees with the theory quite well, however, it should be noted that if one uses the displacement of the cloud along a particular axis, say along the y-axis (i.e. $s^2 = dy^2$), as opposed to the total displacement (i.e. $s^2 = dx^2 + dy^2$), then one may overestimate the pixel size as they will not account for the velocity in the orthogonal direction. The same reasoning applies if the camera is tilted with respect to gravity. As an example, consider the extreme case of when the optical axis is aligned with gravity. In this case the mean position of the cloud will not change x-y coordinates with time, and thus the pixel size will be calculated to be infinitely large, which is obviously incorrect. Alignment of the optical axis to be perpendicular to gravity can be readily performed by placing a spirit level on

⁸We say the effective size of the pixel as the optical system will provide some level of magnification, and thus a given pixel will represent an effective size in physical space given by the pixel size divided by the magnification.

⁹Note that g varies geographically, but the variation in g across Australia is less than 0.02%, whilst across the globe the maximum variation is only 0.07%^[151] and thus it is likely the uncertainty in fitting will be greater than this uncertainty.

the camera, this typically gets the alignment to within a few degrees. Alignment of the gravitational axis with either the x or y axis of the camera pixels can then be performed by rotating the camera until the initial and final coordinate along one of the two axes is the same after some drop time. For my side imaging this alignment has been measured to be within 0.15° of parallel with gravity¹⁰

Clearly from the above example, a camera whose optical axis is aligned with gravity, such as a top or bottom imaging system, cannot be calibrated this way. Instead, one can use a calibrated imaging system whose optical axis is perpendicular to the top or bottom imaging, and then by displacing the cloud along an axis which both imaging systems can see, the magnification of the top/bottom imaging can be measured by comparing the displacement in the two systems. This was performed to calibrate the bottom imaging system, and the measured magnification was found to be 39.6 ± 0.1 , very close to the expected value of 40.

7.2.3 Taking an absorption image

To take an absorption image of our cold atomic cloud we image on the $F = 2$ to $F' = 3$ cycling transition so that the atoms can scatter multiple photons within a single shot. This requires us to first optically pump our atoms from $F = 1, m_f = -1$ into the $F' = 2$ manifold, where they will undergo spontaneous emission and fall back into the $F = 2$ manifold from which they may undergo the cycling transition. This optical pumping is typically performed $100 \mu\text{s}$ before imaging, utilising the MOT repumping light with a $20 \mu\text{s}$ pulse typically being long enough to pump all the atoms into the $F = 2$ manifold.

The atoms are then exposed to on-resonance light¹¹, with a typical exposure time of $200 \mu\text{s}$, and the light is then imaged onto the camera producing an image, I_{atoms} , via the imaging systems that are described in detail in section 7.5. This first exposure is what we typically call the atom frame as it contains the shadow cast by the atoms, a typical image is shown in figure 7.4a. In this frame there is a clear shadow where the atoms are located, whilst in the rest of the frame interference fringes are clearly present. During the imaging process it should be noted that the bias coils are used to generate a bias field along the imaging axis to provide a quantisation axis for the

¹⁰This is measured by the displacement of the cloud along the x direction for the longest drop time viewable.

¹¹The resonance frequency is found by simply scanning the imaging frequency and choosing the frequency that maximises the absorption of the probe beam for a dilute cloud.

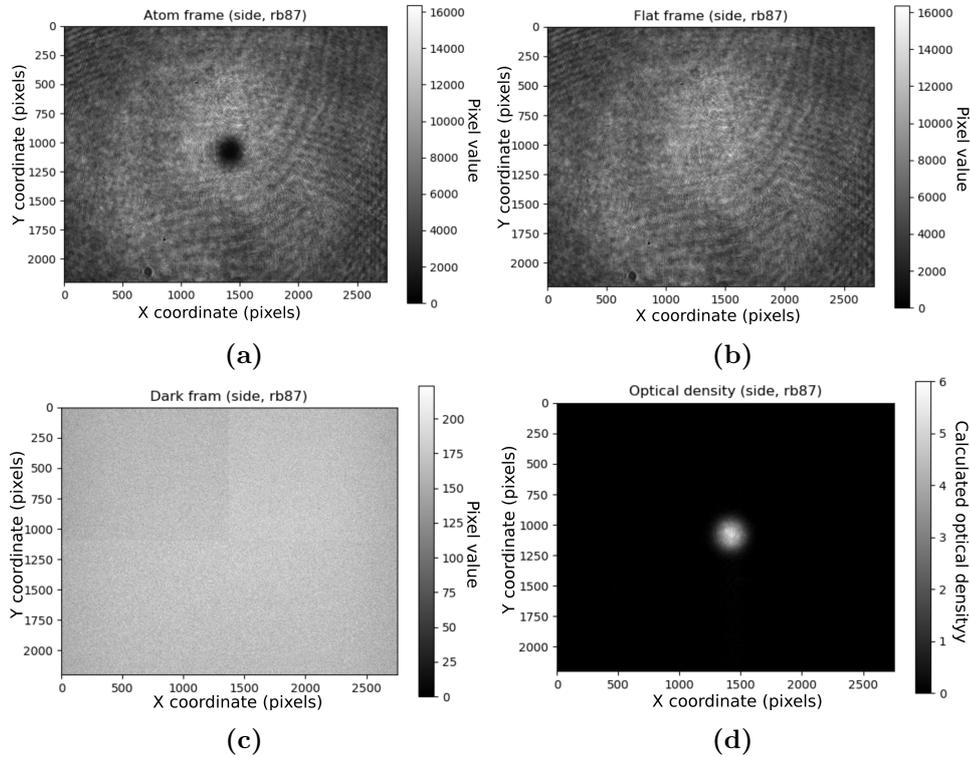


Figure 7.4: Taking an absorption image **a)**, **b)** and **c)** show the images captured and used to extract the optical density of a cloud via absorption imaging, with **d)** showing the resulting optical density of the cloud. **a)** shows the first frame taken, the ‘atom’ frame, in which the atoms cast a shadow in the laser. **b)** shows the flat frame, the exact same conditions of the exposure as in the ‘atom’ frame are used, but the atoms have had enough time to disperse. **c)** shows the dark frame, an exposure taken with no laser sources switched on, to capture the background noise, whilst **d)** shows the resultant calculated optical density using frames **a)**, **b)** and **c)**.

atoms, and circularly polarised light is used to drive the transition.

Next we take a flat frame, I_{flat} , shown in figure. 7.4b. This frame is to account for any spatial variations in the probe beam, such as the interference fringes arising from coherent laser light scattering off particles in the imaging beam path, or from multiple reflections within the cell and gives us a reference to measure how much absorption of the probe beam has occurred. To perform this image we typically wait 100ms so that any atoms that were in the trap have had time to fall out of the field of view, and then perform the exact same imaging procedure, including the repumping of atoms in case this adds to the background light that could be scattered onto the camera. Finally, we wait another 100ms before taking a dark frame, I_{dark} with no imaging beams turned on, so as to account for the background noise, shown

in figure 7.4c.

In order to extract the optical density from the images, equation 7.8 is used, but with the replacements of $I_{in} \rightarrow I_{flat} - I_{dark}$ and $I_{out} \rightarrow I_{atom} - I_{dark}$ resulting in

$$OD(x, y) = -\ln\left(\frac{I_{atom}(x, y) - I_{dark}(x, y)}{I_{flat}(x, y) - I_{dark}(x, y)}\right) + \frac{I_{flat}(x, y) - I_{atom}(x, y)}{I_{sat}(x, y)} = \sigma\tilde{n}(x, y) \quad (7.12)$$

Where all quantities can now be expressed in camera pixel counts. From this equation we can thus infer the optical density, and then infer the atom number via

$$N = \int \tilde{n}(x, y) dx dy = \int \frac{OD(x, y)}{\sigma} dx dy \quad (7.13)$$

which is typically performed as a summation over the pixels of the camera. Figure 7.4d shows an image of the calculated optical density, using equation 7.12, the frames in figures 7.4a, 7.4b and 7.4c, and the measured saturation intensity value.

7.2.4 Interference fringes

One problem that does frequently arise with the use of absorption images is the imperfect cancellation of interference fringes between the atom frame and the flat frame. As we are using coherent monochromatic light, the presence of interference fringes on the resulting image is unavoidable, since any particle or imperfection in the optical path, a piece of dust, a scratch on a lens etc. will act as a point source, scattering light from the probe beam that will then interfere, producing interference fringes.

If these interference fringes were constant they would present less of a problem, as subtracting the flat frame should remove them from the calculated image. In reality though, small variations in the path length between the scatterer and the camera can occur between the atom frame and flat frame, causing the phase of the fringes to change in the observed image, leading to imperfect cancellation. It is useful to remember that to cause a phase shift of π a path length difference of only $\lambda/2$ is required, thus very small vibrations of components can lead to significant changes in the resultant image.

To remove the effects of the fringes, the first step is to ensure that every element in the optical setup is as mechanically secure as possible, as mechanical vibrations are the leading cause of the fringe variations. Next, if all elements are secured, one

can then try to find the offending particle/scratch, and attempt to clean/replace it to remove the fringes. Finally, the last step is to minimise the time between the atoms frame and the flat frame, as if this time can be minimised the range of vibration frequencies that the imaging system is sensitive to will be decreased. For example, if we say that a fringe shift of up to $\lambda/10$ is acceptable in our imaging system, and if the frequency that the fringes are moving at is given by $f_{vibration}$, then by having an interframe time of $\frac{1}{10f_{vibration}}$ the fringes will only move by $\lambda/10$ during the interframe time¹². Note that the vibrational frequency we define here is not necessarily the frequency at which the offending element is vibrating, but rather the frequency that the phase of the fringes is oscillating at, as if the amplitude of the vibrations is quite large the fringes could undergo many 2π phase shifts within one oscillation of the vibrating element. Alternatively the above argument can be rephrased, with the same condition upon the allowed fringe variation, as the lowest frequency to which the imaging system is sensitive is given by $f_{min} = \frac{1}{10t_{interframe}}$, thus by minimising $t_{interframe}$ we can make sure $f_{min} > f_{vibration}$.

As the timing between the atom frame and flat frame can be critical to the performance of absorption imaging, this is where the particular type of camera used is important, as there are a few specific cameras that are optimal for what we call ‘back to back’¹³ imaging. The Proscillica GT2750 is one such type of camera, an interline CCD camera¹⁴. Typically the limitation on the frame rate of a camera is placed by the readout time of the ADC, not the rate at which charge can be moved about the CCD chip. An interline CCD camera is effectively a CCD array where every second column of pixels is covered from external light, which allows an image to be taken, the charge in each pixel moved horizontally into a pixel that is not exposed to light and where it can be stored during readout, and then another image is taken by exposing those same original pixels. Figure 7.5 shows a representation of how this process works. Initially the CCD chip is exposed to light and the photosensitive pixels accumulate charge. Next, the charge that has accumulated during the exposure gets transferred into the non-photosensitive ‘pixels’ to be stored whilst the next exposure occurs. The second exposure then occurs, with the photosensitive pixels again accumulating charge, and the entire camera chip, photosensitive

¹²Where we have assumed that the exposure time is much less than the interframe time.

¹³Where the images are taken with a very short interframe time

¹⁴It should be noted that this is not the only type of camera that can do fast back to back imaging, other camera’s that are not interline may have similar capabilities, such as frame transfer, where the entire image is shifted from a set of exposed pixels to a set of covered pixels that are not interline. These cameras have the benefit of not wasting the light that would otherwise fall upon the interline pixels.

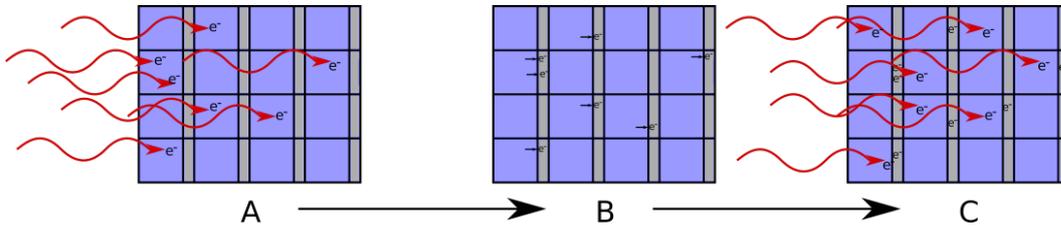


Figure 7.5: Interline CCD mechanism The process that allows an interline CCD camera to take very fast back to back images is shown. **A** shows the initial accumulation of charge in the photosensitive pixels of the chip during the exposure. **B** shows how this accumulated charge is then transferred to a non-photosensitive area of the chip to be stored whilst the next exposure, shown in **C** occurs, before the entire chip is read out.

and non-photosensitive pixels are all read out. With interline CCD cameras, frame rates can be improved by orders of magnitude. For example, imaging with my side-imaging camera and not utilising the interline nature of the CCD, interframe times are on the order of 100 ms. Whereas if I used the interline CCD they could be on the order of 20 μ s, which is over three orders of magnitude improvement in the frame rate and hence minimum frequency to which the imaging system is susceptible¹⁵.

One final method of reducing the effect of fringes upon absorption imaging is to reduce their relative intensity in the image plane. To achieve this, one needs to reduce the intensity of the light at the scatterer compared to the intensity of the light in the atom plane. This can be achieved through focussing of the probe beam, which does not affect the resultant image, but will increase the intensity in the atoms plane, and thus reduce the ratio of the intensity scattered to the intensity in the atom plane.

7.3 Fluorescence imaging

So far we have only discussed the role of absorption imaging in our experiments, however, another frequently used method to gain quantitative information about the cold atom cloud is to use fluorescence imaging. Fluorescence imaging is very similar to absorption imaging in that the same photons are being collected by the imaging system, however, in the case of fluorescence imaging the photons are ‘counted’ photons, as opposed to ‘missing’ photons. Fluorescence images have the benefits of having little background light, allowing smaller signals to be detected than in the

¹⁵Although our side imaging camera has this capability, we have not had to employ it as our fringes have not been too severe.

case of absorption imaging, where the signal may be lost due to fluctuations in the imaging probe. However, the signal in a fluorescence image is also typically smaller than that of a corresponding absorption image by a factor of ≈ 100 [145], thus making the signal typically very small. It is only the greatly reduced ‘noise’ in fluorescence imaging that allows these low light level signals to be detected.

To perform fluorescence imaging all one need do is illuminate the atoms with near resonant light¹⁶ and image the result. Fluorescence imaging is thus very simple to perform, however, it can be very difficult to accurately calibrate as the absolute intensity of the light at the atoms, since the collection efficiency of the optical system and the quantum efficiency of the camera must all be known to give an accurate number. As such, I do not use fluorescence imaging for quantitative number measurements of cold atoms, however, I do use it for temperature measurements.

To use the fluorescence signal for temperature measurements one need only see the spatial distribution of the fluorescence signal and have an accurately known effective pixel size (i.e. magnification and known camera pixel size) to yield accurate temperature measurements. There are a few caveats here though, the first is that the exposure time must be kept relatively short, as otherwise imbalances in the beams used to cause the atoms to fluoresce can impart momentum to the atoms. The second arises with optically thick clouds, as if the cloud is too thick, only the atoms on the edge of the cloud are likely to scatter photons, and thus one observes only the outer shell of atoms. To remove this effect one can image off resonance as this will reduce the optical density, and as we are imaging only those photons that are re-emitted by the atoms there is no phase-shift of the resulting light to be concerned with.

The benefit of fluorescence imaging as opposed to absorption imaging occurs when the atomic cloud has a very large optical density, as the resulting absorption images show a ‘flat top’ due to the limited OD the imaging system can accurately detect for the given imaging parameters, and thus Gaussian fits to the resulting profiles can be inaccurate. This makes the absorption imaging of a cold cloud at short drop times fairly inaccurate. As described in section 7.2, this effect can be mitigated by using high intensity imaging, but this can become prohibitive in the amount of power required for imaging very dense clouds with a large field of view.

Throughout this thesis a mixture of fluorescence and absorption images are used in

¹⁶Note the light used to make the atoms fluoresce should not be coincident with the imaging axis, or the resulting signal will be lost.

the quantitative analysis of our results, however, atom number measurements are solely performed with absorption imaging, whilst temperature measurements before the atoms are loaded into the magnetic trap come solely from fluorescence signals.

7.4 Measuring the temperature of a cold gas

So far we have described two techniques for imaging ultracold atoms, and in the case of absorption imaging have described how this can be used to extract the number of atoms from a single absorption image. In the quest to achieve BEC, however, atom number is not the single most important parameter, the phase space density is. As the phase space density depends upon the temperature of the atoms, some method has to be used to ascertain the temperature of the sample. A number of methods has been described in section 5.2, but the method we use is a variant on the time of flight method.

In the time of flight method a probe beam is placed above or below the cold atom sample, and during time of flight the cloud expands and falls under gravity encountering the probe beam and causing fluorescence. By measuring the time dependence of the fluorescence, the temperature of the atoms can be calculated by fitting the resultant data with a model such as presented in reference [152], which takes into account the size and shape of the probe beam. The variant on this method that we employ is to instead obtain a spatially resolved image of the atoms, and repeat this for a number of different drop times, and by applying a similar model we can measure the temperature of the atoms by how rapidly the cloud expands in time of flight.

In reference [152] Brzozowski shows that for a Gaussian distribution of initial waist σ_0 and temperature T , the Gaussian waist at a later time, t is given by

$$\sigma(t) = \sqrt{\sigma_0^2 + \frac{k_B T}{m} t^2}. \quad (7.14)$$

By squaring both sides we yield a linear relationship between the Gaussian waist squared, $\sigma(t)^2$, and t^2 given by

$$\sigma(t)^2 = \sigma_0^2 + \frac{k_B T}{m} t^2. \quad (7.15)$$

From equation 7.15, measurement of the temperature is relatively simple, measure

the waist of the cloud in free fall for a range of different drop times and then perform a linear fit to the relationship between the square of the waist and the square of the drop time, and the gradient will be equal to $\frac{k_B T}{m}$, from which the temperature can be inferred, and the intercept will be given by σ_0^2 . Typically time of flight thermometry provides very good agreement between theory and experiment, there are multiple examples throughout this thesis that demonstrate this, such as in figure 6.5b, but as also demonstrated in this thesis there are some exceptions in cases where the atomic cloud has not thermalised¹⁷. Some caveats also apply to time of flight imaging. The first is that we have assumed a collisionless expansion of the cloud, which may not necessarily be the case in particularly dense clouds. The second is that the isotropic expansion of the cloud is only guaranteed for drop times $t \gg \omega^{-1}$, where ω is the frequency of the trap in which the atoms were confined before the expansion, thus isotropic expansion will only occur once the cloud has expanded to much larger than its initial size.

7.5 Optical imaging systems

In this final section on imaging cold atoms, I will detail the optical systems used for side and bottom imaging. These two systems are quite different, as they are used for two different purposes.

Side imaging is the workhorse diagnostic and optimisation tool of the apparatus, as such it must be able to accurately measure both the number and temperature of the cold atoms cloud. To meet these requirements a large field of view is required so that hot thermal clouds can expand for many ms allowing accurate measurement of the temperature. The resolution of the side imaging system, however, need not be particularly good, as since it is only used as a diagnostic tool for the optimisation of temperature and number, the smallest feature it has to observe is the waist of the cloud, which will typically be on the order of 10s of microns.

The bottom imaging system has quite different requirements as it has been designed for high resolution, *in-situ* imaging. As such, the focus of this imaging system is upon the resolution of the imaging system and the quality of the resulting *in-situ* image. For the optimal magnification, typically one diffraction limited spot should be spread across 2-3 pixels, for our camera which has 1024 by 1024 pixels, this results

¹⁷See the thermometry of the CMOT, figure 5.18a, where the data is not singularly linear, but rather shows two linear regions.

in a field of view of only 340 micron, much too small for thermometry of hot clouds, but ideal for imaging *in-situ* condensates.

7.5.1 Side imaging

For side imaging we image along a horizontal axis at 45° to the MOT beams with a laser beam of 2 cm $1/e^2$ diameter, containing ≈ 7 mW of power, placing the peak intensity at $\approx 3I_{sat}$ and giving a relatively uniform distribution of power across typical cloud sizes.

The optical imaging system for side imaging consists of a 75 mm focal length 1" achromatic doublet placed ≈ 75 mm from the atoms, and is followed by another 75 mm focal length achromatic doublet¹⁸ placed 75 mm from the camera, a Prosillica GT2750. The lenses are spaced approximately 150 mm apart, but this separation is not crucial to the performance of the imaging system. The 75 mm focal length lens is chosen as this is as close to the atoms from the side as we can reasonably get, and should provide a diffraction limited resolution of $\approx 3 \mu\text{m}$, which is below the magnified pixel size of the camera. All lenses are mounted directly off the camera using Thorlabs 1" lens tube, ensuring the optical axis of all elements are coaxial. Figure 7.6a shows a schematic of the side imaging optics, note the orientation of the achromats changes to minimise spherical aberration.

The Prosillica GT2750 is an interline CCD camera, which consist of 2200 by 2750 pixels that are $4.5 \mu\text{m}$ squares. With this camera, and the one to one imaging described above, we have an approximate 10 mm by 12 mm field of view allowing long drop times of up to 30 ms to be observed, useful for accurate time of flight thermometry, whilst the pixels are still small enough to observe the qualitative changes during condensation.

To focus the side imaging system, first the final lens that refocuses the image onto the camera is placed one focal length from the CCD chip. To achieve this a large collimated beam¹⁹ is passed through the lens onto the camera, and the lens translated along the optical axis until the smallest spot-size is observed on the camera chip²⁰. Next, the first lens is spaced so that the entire lens configuration is in an ap-

¹⁸It should be noted that achromats are chosen not for their chromatic performance, but rather their reduced spherical aberration compared to singlets when oriented correctly.

¹⁹For beams larger than a few mm, a shear interferometer is useful for checking the collimation of a beam.

²⁰This method should place the lens quite precisely, to within the Rayleigh range of the focussed

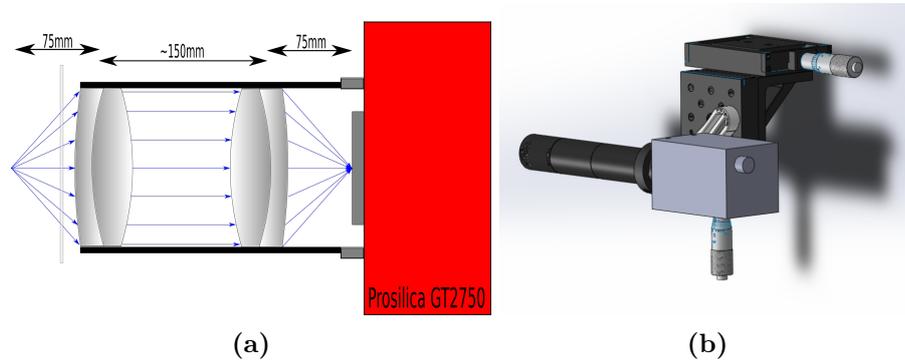


Figure 7.6: Side imaging **a)** A schematic of the side imaging optical system, shown are the two achromatic lenses of 75 mm focal length placed approximately in a 4f imaging configuration. Note the orientation of the achromats changes to minimise spherical aberration and the rectangle in front of the first lens represents the glass of the Science cell. **b)** A Solidworks model of the entire side imaging system, including camera, lenses mounted directly off the camera in 1" lens tube, and the two translation stages used for positioning the camera.

proximate 4f configuration, although this spacing is not critical to the performance of the imaging system. Finally, the lens and camera system is mounted on a pair of translation stages allowing translation in both the vertical direction as well as along the optical axis, and the entire system is then moved to focus the fluorescence image of the atoms²¹ onto the camera. The entire side-imaging system is shown in figure 7.6b as a Solidworks model.

7.5.2 Bottom imaging

Bottom imaging makes use of a custom made objective, designed to have a diffraction limited resolution of less than a micron at 780 nm, able to image through 4 mm of glass, has an effective focal length of 30 mm and a working distance of 20 mm. This objective is followed by an $f=300$ mm achromat, placed approximately 300 mm from the back of the objective which produces an image with $10\times$ magnification at its focus. Next, an $f=100$ mm achromat and an $f=400$ mm achromat configured in 4f configuration reimage the image formed by the $f=300$ mm achromat onto the camera with a further $4\times$ magnification, forming an overall $40\times$ magnification. A schematic of this imaging system is shown in figure 7.7, note that all pairs of lenses are placed approximately in 4f configuration, i.e. the objective lens and the first achromat, the

²¹typically I would focus the camera on the magnetic trap in-situ.

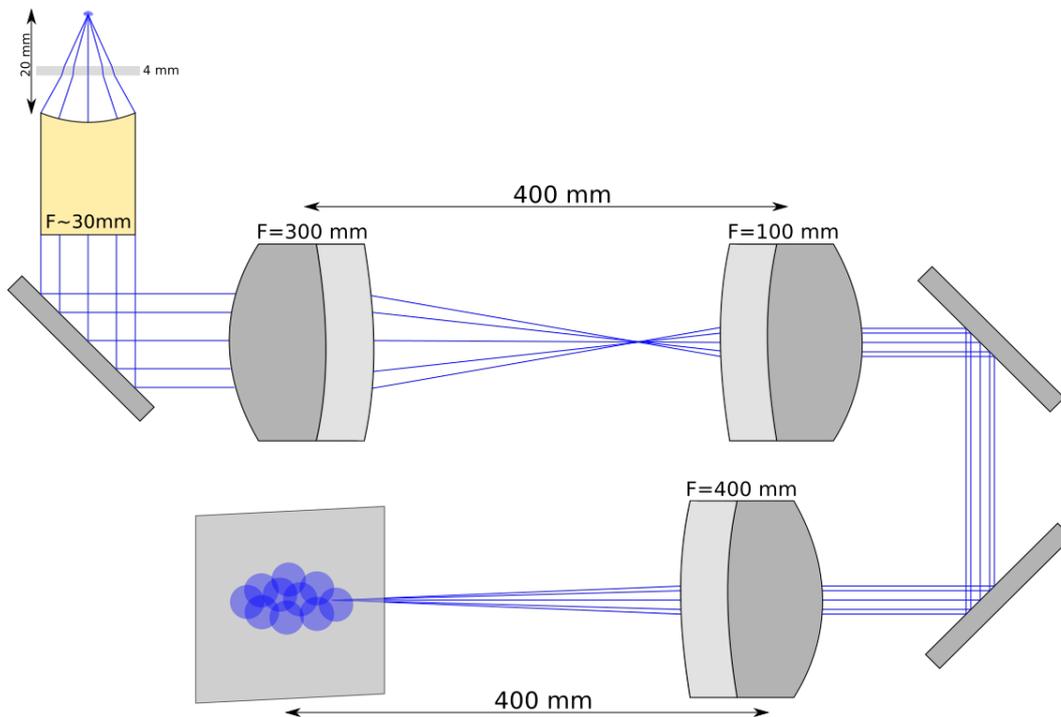


Figure 7.7: Bottom imaging a) A schematic of the bottom imaging system, shown are the custom objective lens as well as the three achromats, which overall provide a magnification of $40\times$. Note all pairs of lenses are placed in a $4f$ configuration, however, this is only particularly critical to the performance for the first pair of achromats.

first and the second achromat as well as the second and the third achromat. This $4f$ configuration of all pairs of lenses is not critical to achieving diffraction limited performance of the imaging system, as the image is in infinity space between the objective lens and the first achromat, and similarly between the second and the third achromat, however, this configuration will help minimise and aberrations for a collimated beam entering the objective lens.

The camera used for bottom imaging is an Andor iKon M, which has a CCD containing 1064 by 1064 pixels that are $13\mu\text{m}$ by $13\mu\text{m}$. With this camera and the optical imaging system described above, the resultant effective pixel size is $0.325\mu\text{m}$ by $0.325\mu\text{m}$, has a field of view of $\approx 350\mu\text{m}$ and results in approximately 3 pixels per diffraction limited spot.

To align the bottom imaging system, first the $f=300\text{mm}$ achromat was placed one focal length away from the Andor camera. The position of this lens was optimised by the same procedure as described in the side imaging section above, by focussing a collimated beam onto the camera and minimising the spot size. Next, a 1mm

aperture was placed upon the top and bottom of the objective lens *in-situ* and the LED light from my phone was used to illuminate the objective. For the LED light to make it through both the apertures and the objective lens it has to both be centred on the objective lens as well as coaxial with the optical axis, thus making it useful for tracing out the imaging path. The light is then centred on the mirror underneath the objective and directed through the 300 mm lens, coaxial with the optical axis and centred on the lens. This procedure aligns the centre of the objective to the centre of the camera, and ensures that the camera is placed at the focus of the achromat.

Next the fluorescence from a cold cloud of atoms in the magnetic trap is used to centre the objective lens upon the atoms by physically translating the objective lens in the plane transverse to the optical axis, and the focus of the objective moved to the atoms by translating the objective lens along the optical axis. During this process the peak fluorescence count on the camera is used as a metric, as even if the atoms are not within the field of view some fluorescence is likely to be observed, thus one can walk towards the atoms without necessarily seeing them. Similarly the peak fluorescence count is used as a metric for the focus as the objective lens should capture the largest fluorescence when placed one focal length from the atoms. This procedure obviously ‘misaligns’ the optical system, and if large displacements of the objective lens are required I would suggest repeating the initial alignment procedure that utilise the apertures on the objective, however, for my imaging system only very small corrections to the position of the objective lens had to be made to get the image of the atoms centred on the camera, thus it was not deemed necessary to repeat the initial aperture alignment step. As the camera is placed at the focus of the achromat, only the objective lens should be used to focus the image onto the camera as otherwise one would likely be attempting to correct for the atoms not being at the focus of the objective lens, which would prevent diffraction limited performance. After aligning the fluorescence signal onto the camera, the probe beam to be used for absorption imaging was then aligned through the system, ensuring it was coaxial and centred on the achromat.

With the atoms imaged onto the camera with $10\times$ magnification, it is then time to add in the rest of the optical system that provides the additional $4\times$ magnification²². To align the next two lenses the probe beam used for absorption imaging was used to trace out the imaging path, ensure the imaging path was centred and coaxial

²²It is not wise to go straight to $40\times$ magnification, as this can make it very difficult to initially find the image of the atoms.

with each lens. To find the correct position of the second achromat with respect to the first, i.e. ensure they are spaced by the sum of their focal lengths, an additional collimated probe beam was passed through the two lenses and the position of the second achromat translated until the beam exiting the lens was again collimated. Again, a similar procedure was used to ensure the final achromat was placed one focal length away from the camera, with a collimated beam being transmitted through the final achromat and the position of the achromat adjusted to focus the beam onto the camera. With this alignment procedure completed, the $40\times$ magnified image of the atoms was found immediately upon the camera and only small touch ups of the final mirror before the camera²³ were required to get the image exactly centred on the camera.

This entire optical system was bench tested before being employed to image the atoms, utilising absorption imaging of a USAF resolution target with a small sample of the same glass as the science cell placed in front of it. The smallest pair of lines that can be resolved in the image formed were the group 9, element 2 lines, which corresponds to a resolution of $0.87\ \mu\text{m}$, approximately as expected from the design specifications of the objective lens.

High intensity imaging is used for bottom imaging, with approximately 7 mW of on-resonant light in a 1 mm Gaussian $1/e^2$ diameter beam, producing a peak intensity of $\approx 445\ \text{mW cm}^{-2}$, corresponding to $\approx 100I_{Sat}$, allowing imaging of OD's up to ≈ 50 , as per the considerations in section. 7.2.

²³Not shown in the schematic.

Holographic potentials for Bose-Einstein condensates

Having finally achieved Bose-Einstein condensation it is time to consider what potential geometries we wish to confine the atoms to, and how to achieve these potentials. As described in section 1.1.2 one of the most interesting potentials to study cold atoms in is a uniform potential, as this removes problems associated with harmonic potentials when one wants to study uniform quantum systems, and as simulated by Groszek *et al.*[20], could be used to observe the inverse energy cascade. As such, this chapter focusses on the implementation of very sharp walled uniform potentials. This chapter thus outlines the considerations of creating this confining potential, how it is experimentally implemented and how atoms are loaded into the potential.

8.1 Creating uniform potentials

Having established the desire for a uniform potential, the question then arises as to how one can create a uniform potential for cold atoms using structured light? The first question that needs to be answered to tackle this problem is whether red or blue detuned light is to be used, i.e. do we use an attractive or repulsive potential? To answer this question we also need to define our metric for the uniformity of the

flat bottom circular potential¹.

In defining the metric for the uniformity of the flat bottom circular potential two considerations come into play. The first of these considerations is of the global structure of the trap, how uniform is the trap on a global scale? To this end we can define a metric based upon the size of the overall trap, and the corresponding ‘width’ of the trap walls, i.e. the ratio of the diameter of the trap to the region over which the restoring force will come into play, the wall thickness. Another metric that could be used here instead to characterise the trap globally would be to fit the potential with a power law function, and the order of the function would be our metric. Figure 8.1 shows an example of this metric for different power law potentials, plotted for the range $[-1, 1]$. Clearly as the power increases the potential begins to look more and more uniform. Alternatively if we consider the distance over which these potentials go from say 10% to 100% of its maximum value for the different potentials, e.g. R1 and R2 in figure 8.1, this distance obviously decreases with increasing power, and thus the ratio of the size of the potential to the ‘width’ of the walls of the potential decreases.

The second consideration that comes into play is the local variations of the potential, ideally there should be none as the trap is to be uniform, but this is never the reality. Small scale local variations in the potential will typically arise from two different mechanisms, the first of these mechanisms is simply imperfect optics and dust resulting in scattering of the laser light creating local variations in the resulting potential. The second mechanism that typically comes into play arises from reflections off the optical components in the imaging system, as well as the glass cells of the wall, which can lead to interference with the light creating the uniform potential. These local variations in the potential can be characterised by the average root-mean square deviations in the intensity of the optical field across the potential.

We can now consider the two options available to us, a blue detuned and a red detuned potential. In the case of a red detuned potential we would have a uniformly bright circle, whereas in the case of a blue-detuned potential we would have a ring of light, with the atoms confined to the centre. In both cases the metric for the global uniformity of the potential is mostly determined by the resolution of the optical system that images the potential onto the atoms, set by the numerical aperture of the imaging system. As the numerical aperture is set by the custom objective lens, the resolution is thus set by the wavelength used and as 532 nm light would be used

¹Note you need not have a circular potential for it to be flat bottomed, but as that is the shape of the potential investigated in this thesis, thus is the potential we will consider throughout.

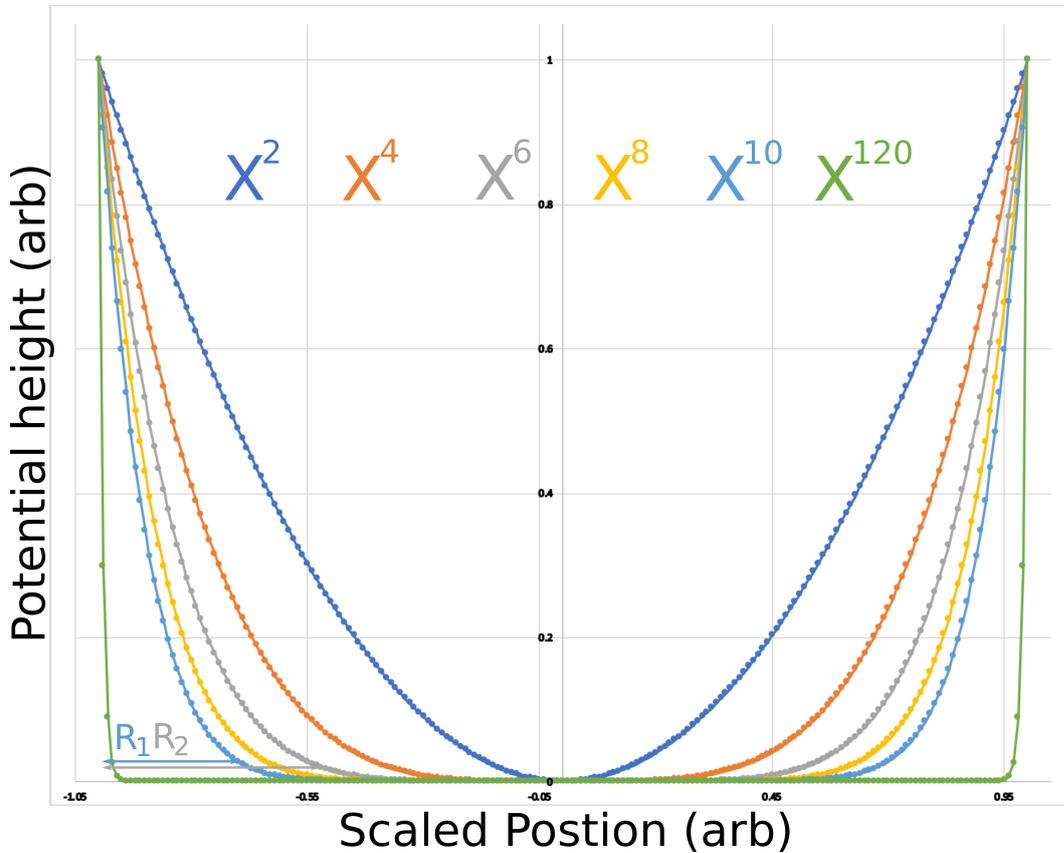


Figure 8.1: Metrics for a uniform potential Shown above is how the shape of a potential, approximated by a power law, varies as a function of the power. Clearly as the power increases the potential becomes more uniform. Indicated in the bottom left hand side of the figure are the widths, R_1 and R_2 , over which the potential increases from 10% to 100% for the X^{10} and X^6 potentials respectively, clearly as the power of the potential increases this width decreases.

for a blue-detuned potential whilst 1064 nm light would be used for a red-detuned potential, the smallest feature created with blue detuned light should be twice as small then that created by red-detuned light, assuming similar aberration limited performance at the two wavelengths . This means that the global metric for blue detuned light will be twice as good as that for red detuned light.

Another consideration with regards to the global metric for the uniformity of the trap is the power required to create the potential. The depth of the potential will be set by the local intensity of the light, I , which for a red-detuned potential must be constant across the entire potential and thus the power required will be proportional to the area of the uniform trap, i.e. $P_{red} = I\pi R^2$, where R is the radius of the trap. In the case of a blue-detuned potential, however, the power will only be proportional

to the area of the ring which will be proportional to the circumference of the ring-shaped wall and its thickness. Thus the power required for a blue-detuned ring will be given by $P_{blue} = I2\pi R \times 2w$, where w is the thickness of the ring-shaped wall. Comparing the required powers, we can see that the ratio of the power required for the red-detuned potential to the blue-detuned potential is given by $\frac{P_{red}}{P_{blue}} = \frac{R}{4w}$. Now as our global metric can be phrased in terms of the ratio of the width of the wall, w to the size of the trap, R , we want to minimise the parameter $\frac{w}{R}$, or alternatively maximise $\frac{R}{w}$, which will necessarily maximise the ratio $\frac{P_{red}}{P_{blue}}$. For the trap achieved in this potential, $\frac{R}{w} \approx 100$, thus $\frac{P_{red}}{P_{blue}} \approx 25$, so for an equivalently deep potential we would require 25X the power.

This consideration is not simply of the power required to form the trap, however, but it also factors into the heating rate of the trap. Grimm and Weidemuller show in reference [13] that the ratio of the heating rate in a blue-detuned potential compared to a red-detuned potential of the same steepness scales inversely proportional to the steepness of the trap, and thus a blue detuned potential offers a substantially lower heating rate. This is intuitively understandable, as the atoms in the blue detuned potential are confined to the region where there is no light, only being forced back to the centre of the trap when they encounter one of the steep walls, whereas in the red-detuned potential the atoms are confined to the regions of high intensity. The heating in a dipole trap results from scattering of photons, and is proportional to the light intensity, thus the scattering rate in the red-detuned potential will be much larger.

Next we can consider the difference between a blue-detuned and red-detuned uniform potential on the local variations in the trap. Now for both potentials the percentage of light scattered off imperfections in the optical setup and from reflections off the surfaces of lenses and the glass cell will be the same, however, the effect they have on the resulting potential will be different.

As we are interested in the local variations of the trap potential due to the scattered light, we will consider the effect upon the bulk of the potential, the interior, not the edges. For the case of a red-detuned potential the scattered light will interfere with the trap light, resulting in a modulation of the local intensity of the light that is greater than the intensity of the scattered light itself, whilst for a blue-detuned potential there is no light in the interior of the trap to interfere with and thus this modulation is smaller. To make this more concrete lets consider an example, considering a point in the interior of the potential where the electric field resulting from the trap light alone is given by \vec{E}_0 for the red-detuned potential and 0 for the

blue detuned potential, and the electric field resulting from scattered light is given by $\alpha\vec{E}_0$, where α is now the fraction of light scattered from surrounding objects that makes it back to the trapping potential, in both cases. Here I have set the electric field intensity at the peak of the ring to be E_0 for the blue detuned potential. Now for the case of the blue detuned potential the resulting intensity in the centre of the trap is given by

$$\delta I_{blue} = |\alpha\vec{E}_0|^2, \quad (8.1)$$

and normalising this with respect to the peak intensity in the trap gives a fractional fluctuation in the intensity of

$$\frac{\delta I_{blue}}{I_{peak}} = \alpha^2. \quad (8.2)$$

In the case of the red-detuned potential, however, the resulting intensity in the centre of the trap is given by

$$I = |\alpha\vec{E}_0 + \vec{E}_0|^2 = |\vec{E}_0|^2 + \alpha^2|\vec{E}_0|^2 + 2\alpha|\vec{E}_0|^2. \quad (8.3)$$

here we have considered a position where the relative phase of the scattered light and trap light is such that constructive interference occurs, obviously this will vary throughout the trap, but this will give us an upper estimate. Subtracting off the average intensity ($|\vec{E}_0|^2$) throughout the trap, and normalising the resultant fluctuation with respect to the peak intensity gives a fractional fluctuation of

$$\frac{\delta I_{red}}{I_{peak}} = \alpha^2 + 2\alpha. \quad (8.4)$$

Thus the ratio of the fluctuations in the red and blue detuned trap will be given by

$$\frac{\delta I_{red}}{\delta I_{blue}} = 1 + \frac{2}{\alpha} \quad (8.5)$$

and given that for a reflection α is typically between 0.2 and 0.3, the fluctuations in the red-detuned trap will be larger by a factor of 6-10.

Given the considerations above a blue detuned potential requires less power, offers less heating and will most likely be more uniform, as such a blue-detuned potential is my potential of choice. We now need to look at how to generate a flat bottom circular potential.

8.1.1 Blue-detuned uniform potentials

We now need to turn to how to create a uniform blue-detuned potential, i.e. a bright ring with a dark centre. The most obvious choice for such a potential comes from the Laguerre-Gaussian modes, which are the eigenmodes of a cylindrically symmetric cavity. These modes are characterised by two integers, $\text{LG}_{(l,p)}$, with l being known as the azimuthal index, which can be any positive or negative integer, whilst p is known as the radial index and can be any positive integer. The electric field of a Laguerre-Gaussian mode in the plane $z = 0$ can be described by[84]

$$E(r, \phi, z = 0) = \sqrt{\frac{2p!}{\pi(p+|l|)!}} \left(\frac{r\sqrt{2}}{w_0}\right)^{|l|} \exp\left(-\frac{r^2}{w_0^2}\right) L_p^{|l|}\left(\frac{2r^2}{w_0^2}\right) \exp(il\phi) \quad (8.6)$$

where r and ϕ are the radial and azimuthal coordinates, w_0 is the waist of the beam and $L_p^{|l|}$ is the generalised Laguerre polynomial.

The azimuthal index is clearly involved in both the phase of the optical field as well as the intensity. In the case of the phase it is clear that in one rotation about the beam profile the phase of the beam increases (or is retarded) by $2\pi l$. This increase in the phase of the optical field is an example of an optical vortex, and is thus tied to the orbital angular momentum of the mode with each photon carrying $l\hbar$ of orbital angular momentum with the sign of the index indicating the handedness of the angular momentum, The radial index on the other hand determines how the amplitude of the optical field varies radially, with the index p determining the number of concentric rings of zero intensity that are centred on the ring, each radial ring also results in a π phase step as the ring is crossed.

The modes that are of most interest to us are the $\text{LG}_{(l,0)}$ modes, as they contain a single central minimum in which the atoms will be confined. In this case the generalised Laguerre polynomial simplifies to

$$L_0^{|l|}\left(\frac{2r^2}{w_0^2}\right) = 1 \quad (8.7)$$

thus we are left with

$$E(r, \phi, z = 0) = \sqrt{\frac{2}{\pi|l|!}} \left(\frac{r\sqrt{2}}{w_0}\right)^{|l|} \exp\left(-\frac{r^2}{w_0^2}\right) \exp(il\phi) \quad (8.8)$$

here we can see that the radial profile is formed from the product of a power law

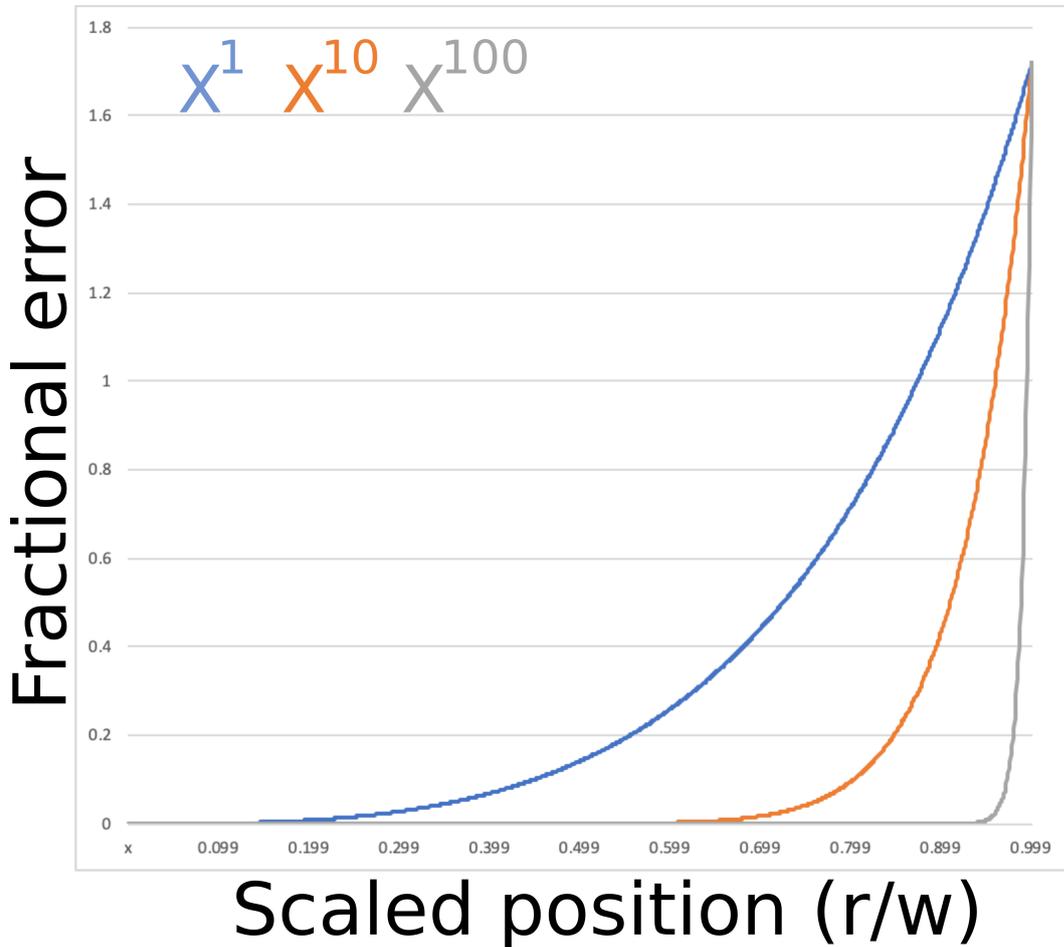


Figure 8.2: Approximating an LG mode Shown above the fractional error in approximating an $\text{LG}_{(l,0)}$ mode by an $x^{|l|}$ mode is shown for $l = 1, 10$ and 100 , for $x \in [0, 1]$ where x represents the scaled variable $\frac{r}{w}$. Clearly, as the azimuthal index increases the Laguerre-Gaussian mode more clearly resembles a power law potential.

potential, $r^{|l|}$ and a Gaussian $\exp(-r^2)$, thus by increasing the azimuthal index of the mode the Laguerre-Gaussian mode becomes closer approximated by a power law potential. This is easily seen by taking the difference of the potential and a power law potential of exponent l , and plotting the fractional difference as a function of a scaled position variable, $\frac{r}{w}$ for a range of different azimuthal indices. Figure 8.2 shows the results of this calculation for three different indices, $l = 1, 10$ and 100 . Clearly as the index increases the fractional error is further reduced over the central uniform region.

A high order $\text{LG}_{(l,0)}$ mode is thus our potential of choice, with the largest possible reasonable value of l to generate the most uniform trap. The limitations placed upon

how large l can be are set by the numerical aperture of our imaging system and the largest sized ring considered reasonable. For this, a number of LG modes generated by an SLM were bench tested with the objective lens that would eventually focus the ring onto the atoms, and it was determined that the largest value of l that could be used whilst maintaining a ring of diameter less than $100\ \mu\text{m}$ was $l \approx 120$.

8.1.2 A dynamic uniform potential

So far we have found an optical field that can be used to approximate a uniform potential. In this section we will add one small addition to this potential that allows the trap to become dynamic, and possibly may be used to impart angular momentum to a cold cloud of atoms.

If we consider the resulting potential from an $\text{LG}_{(l,0)}$ and a $\text{LG}_{(-l,0)}$ in the $z=0$ plane we have

$$E(r, \phi) = \sqrt{\frac{2}{\pi|l|!}} \left(\frac{r\sqrt{2}}{w_0}\right)^{|l|} \exp\left(-\frac{r^2}{w_0^2}\right) \left[\exp(il\phi) + \exp(-il\phi)\right]. \quad (8.9)$$

Using the complex exponential identities, we can rewrite $[\exp(il\phi) + \exp(-il\phi)]$ as $2\cos(l\phi)$, thus the electric field reduces to

$$E(r, \phi) = 2\sqrt{\frac{2}{\pi|l|!}} \left(\frac{r\sqrt{2}}{w_0}\right)^{|l|} \exp\left(-\frac{r^2}{w_0^2}\right) \cos(l\phi) \quad (8.10)$$

which has the same radial form as equation 8.8 but now also has a spatial variation in the amplitude azimuthally due to the $\cos(l\phi)$, which upon one revolution of the optical field has $2l$ minima. This resulting potential is depicted in figure 8.3, which shows the amplitude and phase of an $\text{LG}_{(40,0)}$ and $\text{LG}_{(-40,0)}$ mode, and the amplitude that results from their interference, clearly showing the corrugated pattern. We can thus think of this resulting potential as the same as the uniform potential, but with an additional corrugation about the edge of the ring, note that if the intensities of the two modes are not balanced, then this will simply results in a corrugation that only modulates the potential by a smaller fraction.

Now, as the position of the fringes is set by the relative phase difference between the modes, ϕ , if we change the frequency of the second mode by δf with respect to the frequency of the first mode this relative phase also changes causing the corrugations to begin to rotate. As the relative phase of the modes will now evolve according to

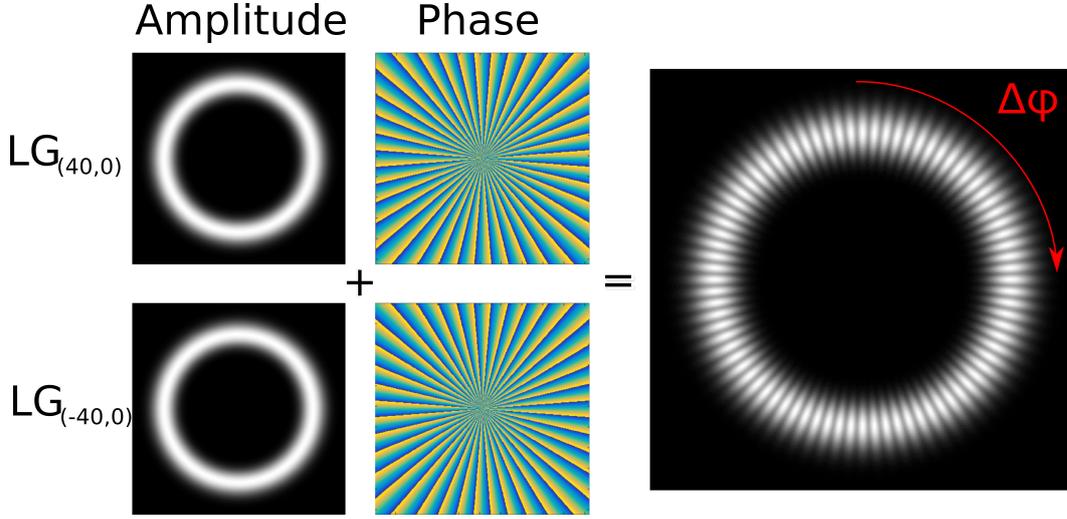


Figure 8.3: A dynamic uniform potential shown above are the amplitude and phase profiles of two opposite handedness Laguerre-Gaussian modes of $|l| = 40$, and the resulting amplitude profile that arises from their interference. The interference of the two modes produces a ring with corrugations, whose position is determined by the relative phase of the modes. By changing this relative phase, ϕ , the corrugations can be made to rotate about the edge of the trap.

$\delta\phi(t) = \frac{2\pi}{\delta f}$, the position of one minima will move to the position of the next minima every $1/\delta f$ seconds. Given there are $2l$ minima around the circumference of the trap the rotation rate of the trap will be $\delta f/2l$, depicted in figure 8.3 by the red arrow.

A single $\text{LG}_{(120,0)}$ mode can thus act as a uniform potential for our atoms, but then by switching on a second mode, an $\text{LG}_{(-120,0)}$, we can modulate the depth of the walls by adding an azimuthal modulation to the ring which will rotate at a frequency given by $\delta f/2l$. This thus gives us the ability to dynamically control the optical field, creating a uniform rotational movement in the trap that could perhaps impart angular momentum to the atoms.

If we consider a ring of diameter $100\ \mu\text{m}$ generated by an $\text{LG}_{(120,0)}$ and $\text{LG}_{(-120,0)}$ mode, which would result in 240 minima azimuthally about the ring, spaced by approximately $1.3\ \mu\text{m}$, then for the ring trap to impart j units of angular momentum to the BEC, then the rotation rate of the ring must satisfy,

$$j\hbar = mvr \quad (8.11)$$

where m is the mass of an atom, v is its velocity and r is the radius of the trap. Now, if we assume in equilibrium that the corrugations rotate at the same velocity

as the atoms we have

$$v = 2\pi r f_{ring} \quad (8.12)$$

where we have calculated the velocity by simply multiplying the circumference of the ring, $2\pi r$ by the number of revolutions per second, f_{ring} of the ring. If we then rearrange to solve for the frequency at which the ring must rotate to impart a given amount of angular momentum we have

$$f_{ring} = \frac{j\hbar}{2\pi m r^2}. \quad (8.13)$$

Here we can see that the frequency at which the ring needs to rotate to impart angular momentum to the atoms scales linearly with the angular momentum and inversely with the radius of the trap squared. Using $r = 50 \mu\text{m}$ and the mass of a rubidium-87 atom, this yields

$$f_{ring} = 0.047 \times j \text{ Hz} \quad (8.14)$$

If we next substitute in the relationship between the rate at which the ring rotates and the frequency difference of the modes, using an $\text{LG}_{(120,0)}$ mode as an example yields

$$\delta f = 11.5 \times j \text{ Hz} \quad (8.15)$$

thus a frequency difference between the beams on the order of 10 Hz is required to impart one unit of angular momentum to the atoms.

One method of generating such a small frequency difference between the beams is to pass the light for both beams through two separate AOMs driven by slightly differing frequencies of RF. In the next section we will describe how exactly this is implemented. Note though that either a single device should be used to generate the RF for the two AOMs, or alternatively two separate devices that use the same reference signal to frequency stabilise the output as otherwise the accuracy of the devices can easily deviate by 10s of Hz, given the carrier signals are typically on the order of 100s of MHz. This thus seems to be a possible method of imparting angular momentum to a condensate through modulations that occur on a scale comparable to the healing length of the condensate.

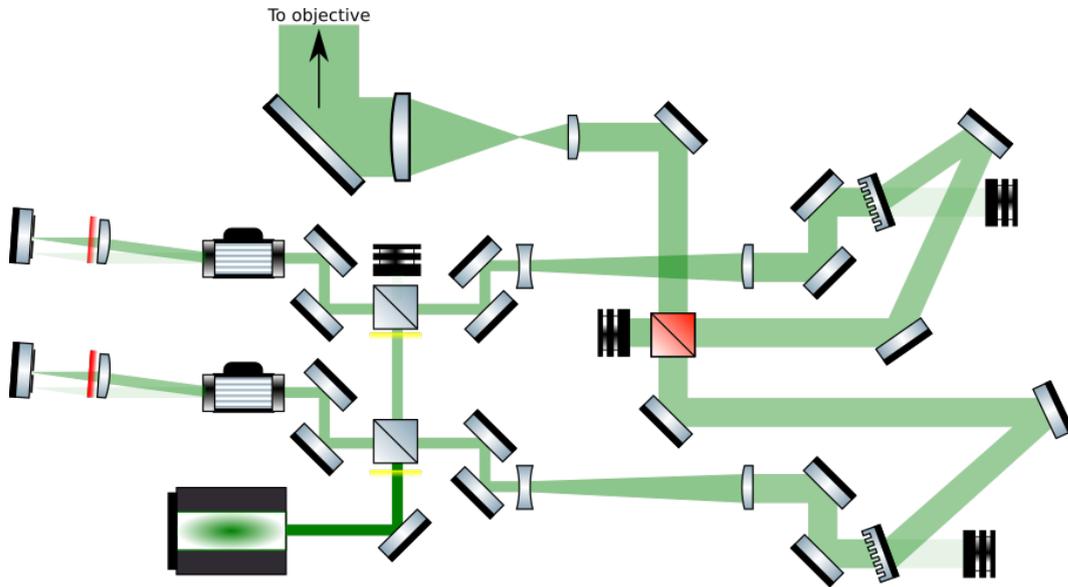


Figure 8.4: Rotating trap schematic. A schematic of the optical setup used to create a sharp walled potential for the atoms, with the ability to also rotate corrugations upon the wall. The output of a 1 W, 532 nm laser is split into two beamlines using two polarising beam splitter cubes and half-wave plates that are then double passed through separate AOMs. The resultant diffracted modes are then magnified before diffracting off two separate holograms to generate two $LG_{(120,0)}$ modes that are then combined upon a non-polarising beam splitter cube before being magnified by a final telescope.

8.2 Implementing a rotating potential

To implement the rotating potential described in section 8.1 we need two high order Laguerre-Gaussian modes of opposite handedness, whose frequency and power can also be controlled. This requires both beamlines to have their own acousto-optic modulators, and to ensure that the beamlines remain aligned when the frequency of one of the AOMs is changed they should also be double passed. The entire optical setup used to generate these beams is shown in figure 8.4.

To achieve this a 1 W, single longitudinal mode, 532 nm laser is used as the light source, which has a 0.7 mm $1/e^2$ Gaussian diameter output. This beam then passes through a pair of polarising beam splitters and half wave plates, the reflection off of the PBS being used to generate the two beam lines for the pair of $LG_{(120,0)}$. From when the two beams are split the two beamlines are made as identical as possible, matching the path lengths as close as possible and minimising the path length where possible, as such only a single beamline is described from here on. After the PBS the beam then passes through an acousto-optic modulator in a double pass configuration

with a $f=150$ mm lens configured in a cat-eye retro reflector arrangement, similar to those described in section 4.3, with a quarter waveplate used to rotate the output polarisation by 90deg with respect to the input polarisation and an iris to remove the undiffracted mode. The double passed mode then is transmitted through the PBS from which it was originally reflected, and a pair of mirrors steers the beam through a $12\times$ telescope to magnify the beam to the size of the recorded hologram, and two mirrors are then used to steer the beam onto the hologram.

The $LG_{(120,0)}$ hologram is generated by the methods described in section 3.4, however, to maximise the efficiency of the hologram a uniformly bright circle with the phase of an $LG_{(60,0)}$ mode is generated by the SLM and interfered with itself² upon the photopolymer film rather than a true $LG_{(60,0)}$ mode as the spatial overlap with a Gaussian is limited. The resultant hologram is shown in figure 8.5a. To maximise the diffraction efficiency off the hologram the input beam is first aligned level with the table and then apertured down to a 1 mm beam by placing a 30 mm cage alignment tool upon the final lens in the telescope. The 1 mm beam is then centred on the hologram by physically moving the hologram mount, this alignment is relatively simple as at the very centre of the hologram is a small hole visible by eye and clearly seen in figure 8.5a, and it is this hole in the hologram that the beam is aligned to.

With the beam centred on the hologram, the target is removed and the hologram is then rotated to maximise the diffraction efficiency. This will necessarily move the position of the beam with respect to the hologram, thus the last two steps need to be performed iteratively to get the beam both centred on the hologram and at the right angle with respect to the holograms surface. Once I was satisfied with the beam being well aligned to the hologram, the resultant mode was imaged onto a camera approximately 600 mm away and the intensity around the edge of the ring was fine tuned to create a uniform intensity by using the final mirrors before the hologram. Figure 8.5b shows the resultant mode imaged onto the camera 600 mm from the hologram once the alignment was complete, whilst figure 8.5c shows the resulting intensity pattern when the $LG_{(120,0)}$ and $LG_{(-120,0)}$ mode are both on at roughly equal intensities and at the same frequency.

Note that for the two beamlines the orientation of the holograms have to be the same, as they were recorded under identical conditions, to ensure the diffracted mode in both cases is an $LG_{(120,0)}$ mode, not an $LG_{(120,0)}$ and a $LG_{(-120,0)}$ mode. Although ultimately we wish to combine an $LG_{(120,0)}$ with an $LG_{(-120,0)}$ mode, it was decided

²but of opposite handedness of course.

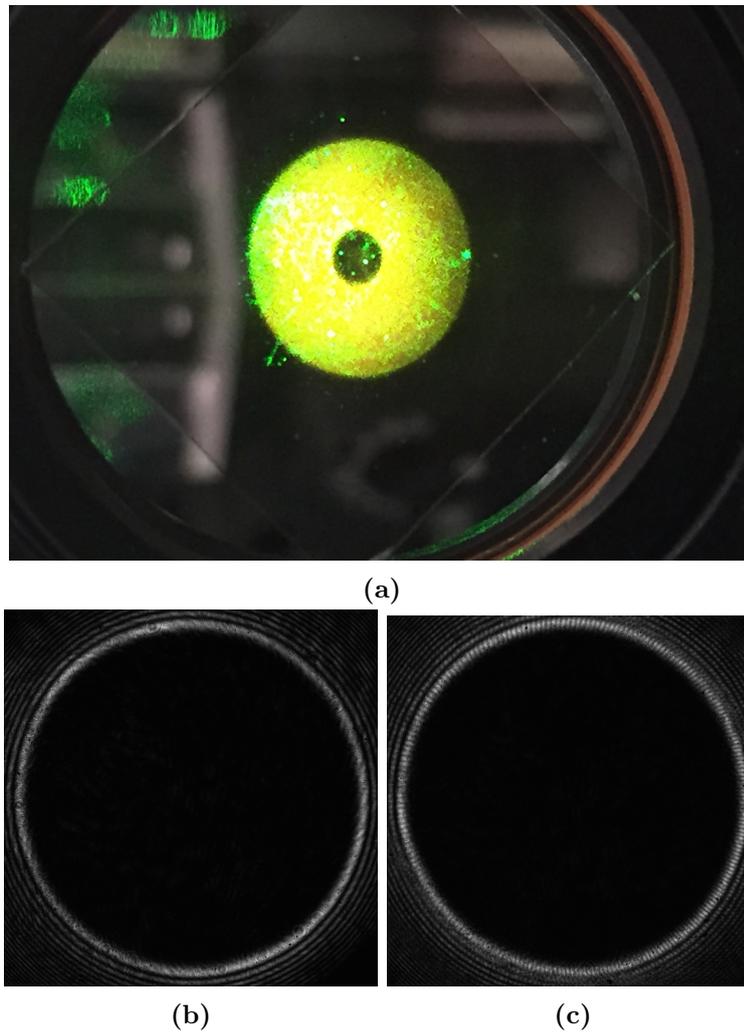


Figure 8.5: LG hologram and resulting profile **a)** The recorded hologram used to generate an $LG_{(120,0)}$ mode from a Gaussian mode, illuminated by the beams used to record it. The photo shows a square (the photopolymer film) upon which the hologram is recorded, whilst the hologram shows a small hole in the centre that arises due to the dark centre of the $LG_{(120,0)}$ mode, and alignment of the two beams is performed by ensuring the hole in each beam overlaps. **b)** and **c)** show the resulting intensity pattern produced by the $LG_{(120,0)}$ hologram, and the intensity resulting from the interference of the $LG_{(120,0)}$ and the $LG_{(-120,0)}$ mode.

that by recording the holograms identically and playing back the holograms in the same manner, the resultant modes will share the greatest similarity. Instead, to get the desired interference of an $LG_{(120,0)}$ with an $LG_{(-120,0)}$ the number of reflections each beam undergoes differs by 1, ensuring the handedness of the two beams is opposite. To this end after the hologram on either beamline two mirrors are used between the hologram and the non-polarising beam splitter cube used to combine

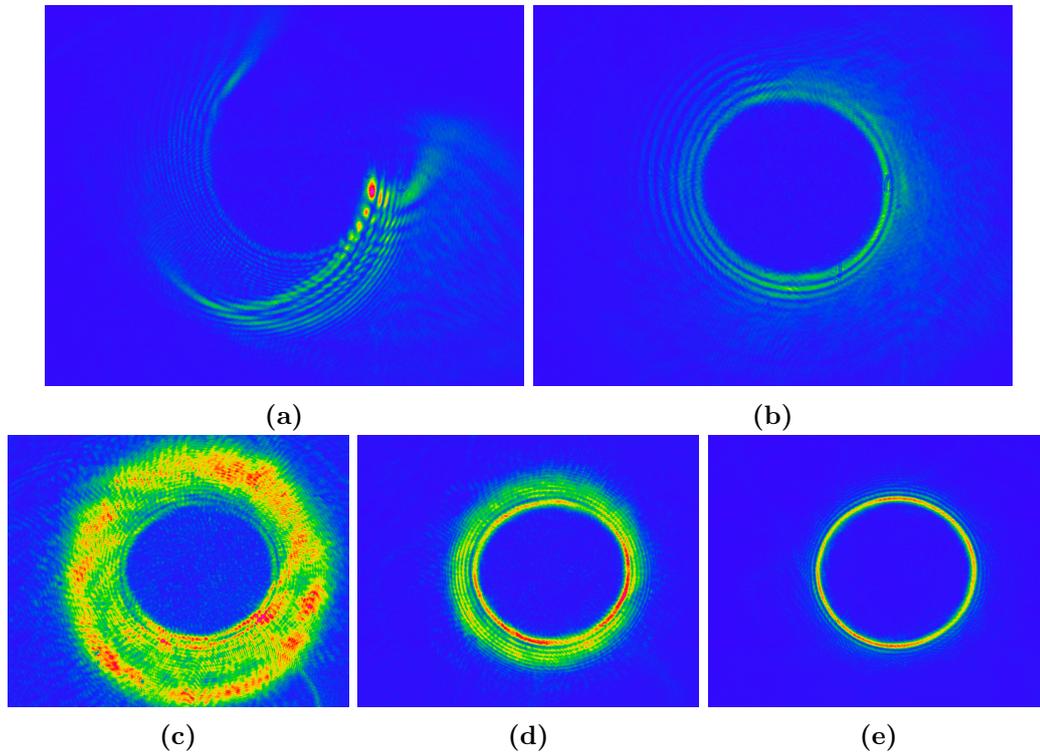


Figure 8.6: Aligning the ring trap a) and b) show the effect of passing through the first telescope off axis, showing the ring clearly deteriorating when aberrations are introduced, whilst c), d) and e) shows the ring at varying stages of defocus, with e) showing the final resultant focussed ring. Note that these images were taken with auto-exposure on, so the exposure of each image is scaled to prevent saturation.

the modes, and thus the beam that is transmitted by the NPBS will undergo two reflections from hologram to combination (the two mirrors) whilst the beam that is reflected off the NPBS will undergo three (two mirrors plus the NPBS). Alignment of the beams to one another after the cube was performed initially by setting a large frequency difference between the beams so that no interference effects were obvious by eye, and the beam profiles were overlapped in the near and far field. With the beams closely aligned final touchups were performed by setting the beams to the same frequency and observing the radial ‘spokes’ pattern, which should be symmetric when the beams are accurately aligned.

Once the beams were combined they then pass through a $4\times$ telescope, expanding the beams to approximately 32 mm. The beams then reflect off a 2” mirror, and a dichroic mirror that reflects 532 nm light and transmits 780 nm light combining the ring trap path with the imaging beam path. This allows the tight focussing of the $LG_{(120,0)}$ mode, producing a ring that is approximately 100 μm in diameter, with

the width of the ring in the radial direction being approximated by a $1 \mu\text{m}$ $1/e^2$ Gaussian diameter.

Alignment of the ring trap through the first telescope is critical to the resulting mode, as small aberrations such as astigmatism seem to cause the mode to deteriorate. To see the effects of misalignment through this telescope the mode was imaged onto a camera placed approximately 2 m after the first telescope, with the telescope slightly defocussed to produce a relatively focussed ring upon the camera. Figure 8.6a shows the resulting effect of the $\text{LG}_{(120,0)}$ mode passing through the telescope off centre, the mode is seen to deteriorate, the ring splitting up into multiple rings, whereas by tweaking the input to the telescope so that the beam is aligned with the optical axis the ring returns, as seen in figure 8.6b.

Figures 8.6c, 8.6d and 8.6e shows the uniform trap at varying degrees of focus, actuated by defocussing the telescope. For a true $\text{LG}_{(120,0)}$ mode, the mode should simply appear as a ring whose size and thickness change as the beam is focussed. However, as the hologram acts more so as a uniform diffraction grating that imparts angular momentum as opposed to a grating that also shapes the beam into a ring, the resulting mode is actually a superposition of $\text{LG}_{(120,l)}$ modes, where $l \in [0, \infty]$. From comparison of figures 8.6c, 8.6d and 8.6e we can see that when the ring is not focussed, the central dark region does not change in size a great deal, but rather as the ring focusses the width of the ring simply decreases, increasing the local intensity of light. At the focus of the ring, shown in figure 8.6e, the ring is almost perfectly circular with a dark core, with the presence of some outer rings arising from being a superposition of $\text{LG}_{(120,l)}$ modes, and has a ring thickness to diameter ratio of 50.

8.2.1 Axial confinement- The sheet trap

The uniform potential generated by the $\text{LG}_{(100,0)}$ mode can be used to confine the atoms radially, however, it will not confine them axially, instead another trap is needed to provide this confinement. For the experiments that would match the theory of Groszek *et al.*[20] or alternatively to study any atomically thin condensed matter systems, the condensate should be quasi-two dimensional, and thus should be trapped much more tightly in the axial direction. One option would be to use the original dipole beam in which forced evaporation to condensation occurs as this provided axial confinement during evaporation, however, this beam would also provide harmonic radial confinement, undoing our efforts to create a uniform trap.

Another option would be to use two blue detuned light sheets and confine the atoms between the sheets, or alternatively use a blue detuned Hermite-Gaussian mode, appearing qualitatively similar to two light sheets, but having a $\pi/2$ phase step between the two sheets, resulting in true darkness at the centre of the mode. A blue detuned option requires much more power than was readily available, however, although we have demonstrated a blue-detuned Hermite-Gaussian mode used to confine the atoms axially[4]. Instead, the option I chose to use is a cylindrically focussed 1064 nm light sheet that is particularly tight in one dimension, and weak in the others. To achieve this I used a sheet of light with Gaussian waists of 0.5 mm in the long direction and 10 μm in the tight direction.

The light sheet is generated from the same laser used to form the Hybrid trap, and the schematic of figure 6.9 shows the optical layout used. Once diffracted through the AOM, the light for the sheet trap first passes through a spatial filter formed by a 75 mm lens focussing the beam onto a 100 μm aperture, which is then re-collimated with an $f = 200$ mm lens, creating a $2.67\times$ telescope magnifying the beam to a Gaussian diameter of 2.32 mm. This telescope is then followed by a pair of cylindrical lenses that are orthogonally aligned, and a pair of $f = 300$ mm lenses, which together act as an approximate $f = 150$ mm lens. This arrangement is an interesting design and allows one axis of the beam, aligned along gravity, to be tightly focussed at the position of the atoms, whilst the other axis exits the final pair of lenses collimated. This design was chosen as if both of the axes are focussed it is difficult to achieve a large aspect ratio of the beam and almost impossible to prevent back reflections off the exit window of the glass cell. Whereas for the arrangement where one axis is collimated and the other focussed it is easy to achieve an aspect ratio on the order of 200, and the beam entering the cell only has to be tilted by a few degrees to prevent back reflections from interfering with the trap.

The first of the cylindrical lenses is an $f = 500$ mm lens and the axis along which the lens focuses is oriented perpendicular to gravity. This lens is placed approximately 650 mm from the final lens forming a $0.3\times$ telescope, demagnifying the beam along the horizontal direction to approximately 1 mm. The second cylindrical lens on the other hand is oriented parallel to gravity, and is a $f = -50$ mm lens. This lens creates a virtual focus 50 mm behind the lens, which is placed approximately 300 mm from the final lens in the system, which reimages this focus in an approximate $2f$ imaging configuration onto the atoms. Figure 8.7 shows a ray trace of the dipole beam through the lenses. The vertical view shows a ray trace if viewed from above the lenses, whilst the horizontal view shows a ray trace if the lenses were viewed from

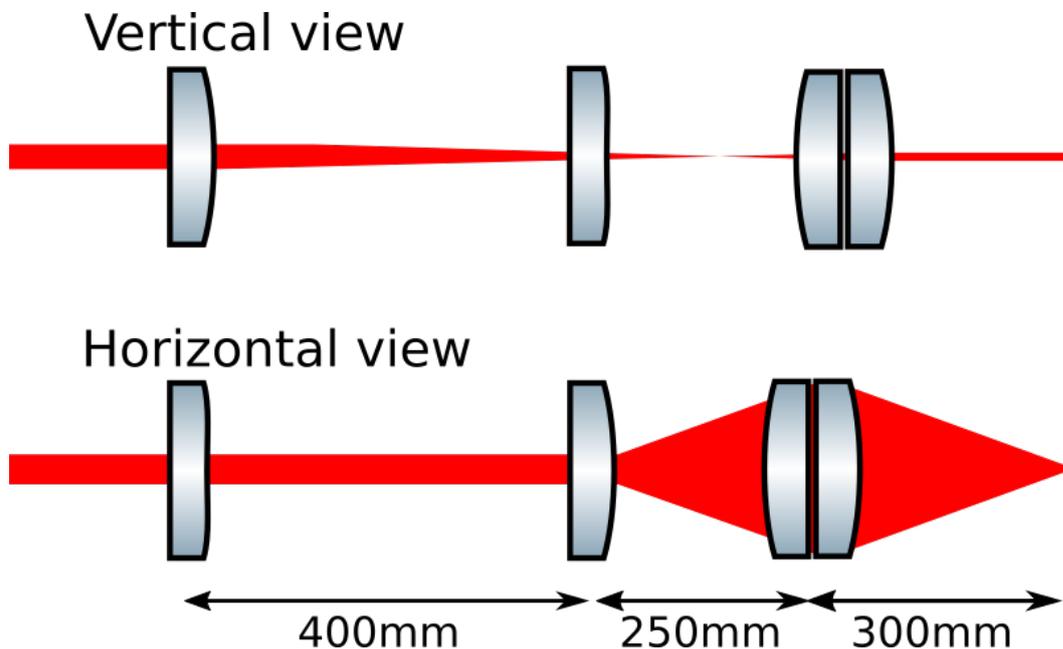


Figure 8.7: Sheet trap ray trace A ray trace showing how a pair of cylindrical lenses and a pair of spherical lenses are used to tightly focus the beam in the vertical direction whilst collimating the beam along the horizontal axis.

the side. The rays enter the lens system from the left hand side. From the vertical view the rays are seen to be focused from the first lens, pass through the second lens un-refracted and are collimated by the final two lenses, whilst from the horizontal view the rays are seen to pass through the first lens unaffected, diverge rapidly after the second lens, and the virtual focus is reimaged by the final two lenses onto the atoms.

The light for the sheet trap is combined with the second horizontal MOT beam, as seen in figure 6.9, with a 2" dichroic mirror placed after the final pair of lenses used to image the focus onto the atoms.

8.2.2 Aligning the uniform trap

The sheet trap used to confine the atoms axially is formed from the same laser as used for the hybrid trap and thus there is up to 20 W of power available to find the dipole beam, producing a trap that is 10s of μK deep. Finding the light sheet trap is thus very similar to the method employed to find the hybrid beam in section 6.5.5, the beam is first aligned using a MOT beam as reference and then found by keeping

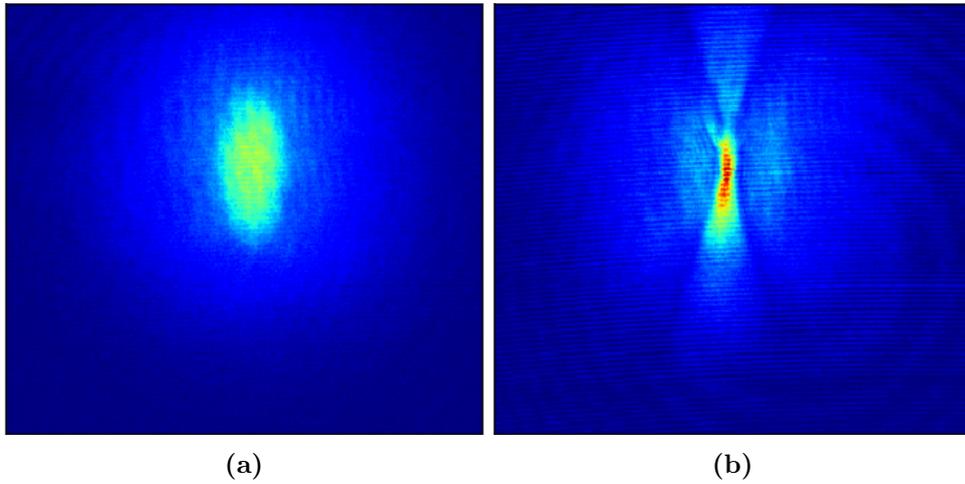


Figure 8.8: Finding the ring trap **a)** and **b)** show time of flight absorption images used to find the uniform ring trap, **a)** shows a condensate that has been loaded into the sheet trap and imaged after 100 ms of ballistic expansion without the ring on whilst **b)** shows the effect of switching the ring trap on, confining the atoms radially whilst they can still expand axially, allowing the shape of the trap to be visualised.

the dipole beam on whilst the cold cloud released from a compressed magnetic trap is allowed to ballistically expand.

Finding the blue-detuned ring is much more difficult as although the walls are tightly focussed, there is much less available laser power and the beam path has much higher losses. These losses are due to the custom objective, which consists of four lenses, none of which are AR coated for 532 nm light. Overall the transmission through the objective and the glass cell results in approximately 25% transmission, thus the overall intensity is much smaller. As the intensity is much smaller, the trap has to be much better aligned, and the atoms much colder, to see the effect upon the atoms.

Initial alignment of the ring trap to the atoms is performed by using the bottom imaging beam. As the dichroic does not have 100% transmission, there is some reflection off the top surface of the dichroic. It is this reflection that is used to initially align the ring trap to the atoms, by first centering the dichroic upon the imaging beam, and then using the dichroic and final mirror to steer the imaging beam reflection through the final telescope that the ring trap comes through. An alternative method that can be used is to align the multiple back reflections that come from the uncoated lenses in the custom objective atop one another, as using these reflections one can get the ring both centred and aligned with the optical axis.

With the position and angle of the ring trap roughly aligned, the other property that needs to be adjusted is the collimation. As the objective lens has been placed the working distance from the atoms, using the alignment described in section 7.5.2, if a collimated 780 nm beam passed through the objective it should be focussed at the atoms. However, as the ring trap is made from 532 nm light there will be some chromatic shift of the focus, but this is not of concern for the initial alignment.

To observe the effect of the ring trap upon the atoms a condensate is first formed and then allowed to expand for 100 ms, being levitated by a magnetic field gradient, whilst the ring trap is switched on. This allows the atoms to expand into the walls of the ring trap, being repelled from the regions of high intensity, and allows the position and shape of the ring trap to be seen. Figure 8.8a shows an absorption image of a condensate after 100 ms of expansion, whilst figure 8.8b shows an absorption image with the ring trap switched on. Once the ring trap has been found, the position is optimised by loading the trap by slowly switching off the hybrid beam and switching on the ring and sheet trap, and then looking at the position of the atoms in the loaded trap compared to the position in which the condensate is formed and ensuring they coincide.

8.3 Characterising the uniform trap

With the uniform trap found, the atoms can finally be loaded into the trap. The loading of the trap is performed by first ramping up the light sheet power whilst switching off the hybrid beam over 1 second, transferring the atoms from one hybrid potential to another, with the quadrupole potential still providing radial confinement to the atoms. This transfer results in approximately 8×10^5 condensed atoms being transferred to the sheet hybrid trap, with the loss being mostly attributed to the increased density of the atoms and thus increased 3-body loss rate. Although the radial area of the condensate increases by roughly a factor of 2 as the atoms expand into a spherical shape, determined by the quadrupole potential, the axial extent of the condensate becomes much denser as the atoms become confined by the sheet rather than the hybrid beam, increasing the overall density by a factor of ≈ 10 .

Next the ring trap is switched on to full power suddenly³, the quadrupole field gradient is increased to 30.5 Gcm^{-1} whilst the Z-bias field is increased to ≈ 100

³This doesn't appear to have any negative effect on the atoms as the ring is larger than the condensate size in-situ.

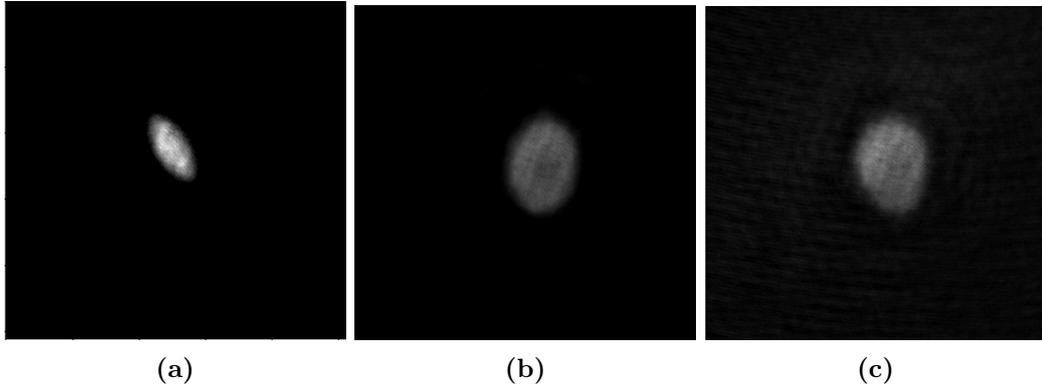


Figure 8.9: Loading the uniform trap **a)** shows an image of the condensate *in-situ*, taken from below before it is loaded into the ring trap whilst **b)** and **c)** shows the condensate imaged *in-situ* when loaded into the ring trap formed by each of the LG-modes separately.

Gcm^{-1} to move the position of the field zero ≈ 3 cm above the atoms while relaxing the radial confinement of the quadrupole field, providing solely a levitation field for the atoms. As the radial confinement of the atoms is released, the atoms expand into the ring, loading $\approx 3 \times 10^5$ atoms into the ring trap. Figure 8.9a shows the condensate, as imaged from below, *in-situ* in the hybrid trap in which the condensate was formed, whereas figures 8.9b and 8.9c show the condensate confined to the uniform ring formed from the two separate $\text{LG}_{(120,0)}$ modes. Clearly there is a small difference between the confined atoms in figures 8.9b and 8.9c, indicating that the confining potentials are not identical. The difference between the resulting confining potentials is attributed to an astigmatism introduced by the glass cell being tilted slightly with respect to the final objective lens. This hypothesis was bench tested by focussing the uniform trap with an achromatic lens and inserting a 4 mm thick piece of glass in-between the final lens and the focus. By tilting the glass with respect to the optical axis of the beam, the circular profile of the ring trap was seen to distort in a similar manner to the distortion observed in the shape of the BEC.

Figure 8.10a shows the condensate in the uniform trap when viewed from the side, clearly a large fraction of atoms remain in the light sheet but are not confined to the uniform trap, as is evidenced by both the image and the reduction in atom number in the condensate when loading the ring from the hybrid sheet trap. It is thought that this low efficiency in loading the ring trap is due to the small potential depth of the ring, arising from the very low overall efficiency of converting the light into the desired mode and directing it onto the atoms. As the light from the laser output must first be double-passed through AOMs, there is initially a 50% reduction in the power from this process, followed by a further 50% reduction due to the amount of

light diffracted off the hologram⁴ and a further 50% reduction in the laser power from combining the two modes upon a cube, these three losses combined reduce the efficiency to 12.5%. After the beams have been combined, however, there is a further reduction in laser power due to reflections off the lenses in the final objective as discussed earlier, resulting in only 25% of the light that enters the objective reaching the atoms. These losses combined result in an efficiency of only $\approx 3\%$, resulting in a measly 30 mW of power reaching the atoms from the initial 1 W available.

To test this hypothesis, two tests were performed. First, with a single ring confining the atoms, the intensity of the ring was reduced and the number of atoms remaining in the condensate was measured as a function of power. This test saw a reduction in the number of atoms as a function of power, but before long the ring was not strong enough to confine any atoms at all and fitting the difference became difficult. A second method employed to test this hypothesis was performed by increasing the power of the second LG mode whilst maintaining the first ring at full power, and keeping the frequency of the two beams the same. This increase in power of the second LG mode results in the static modulation described in section 8.1, reducing the overall depth of the potential. Figure 8.10b shows the result of this test, performed by ramping on the second beam to a variable power. As expected, there was a reduction in atom number as the second beams power was increased.

Now, from equation A.9 we can see that once we have a condensate, the atoms effectively ‘fill up’ the potential, much like a water fills up a glass, until the confining potential is equal to the chemical potential. For the case of our uniform potential we can approximate the potential as $U(r) = 0$ if $|x| < R$ else $U(r) = U_0$. Within the uniform potential, the density is constant at μ/g whereas outside of the potential if $U_0 > \mu$, $n(r) = 0$. On the other hand, if $U_0 < \mu$ then some finite condensate fraction will exist outside of the trap. This explains what we see in figure 8.10a if the chemical potential is larger than the potential depth of our ring trap, as the atoms fill up the ring they begin to overflow as the traps potential is not deeper than the chemical potential of the atoms, and thus cannot contain them all.

Now to reduce the effect of atoms spilling out of the ring we have three options, the first is to increase the power in the ring trap as the potential will scale proportional to the power. However, as no more power is readily available to us, this is not a practical solution.

⁴The overall diffraction efficiency of the hologram is $\approx 80\%$, but due to the spatial overlap of the ring mode with a Gaussian, only 50% of the incident light gets converted into the final mode.

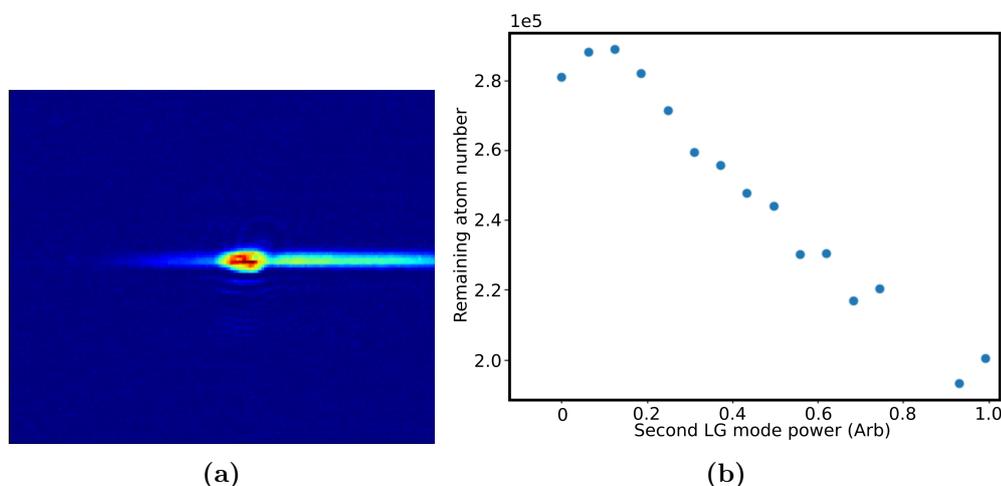


Figure 8.10: Atom number in the uniform trap **a)** shows an image of the condensate *in-situ* in the combined ring and sheet trap when viewed from the side, clearly a large fraction of atoms remain in the sheet but not in the ring. **b)** shows the atom number as a function of the power in the $\text{LG}_{(-120,0)}$ beam, as the power in the second LG mode is increased, the corrugations caused by the interference of the two modes allow atoms to leak out, as the effective potential height of the trap is reduced.

The next solution we have is to change the size of the uniform trap. If we decrease the size of the uniform trap by a factor α by changing the order of the LG mode used to form the trap⁵, then the width of the ring will stay constant, but the circumference will decrease by a factor α increasing the peak intensity and thus potential by a factor α . This thus allows a greater density of atoms to be confined within the trap, however, this will also decrease the area of the trap by a factor of α^{26} , thus overall the number of atoms that can be confined will decrease by a factor of α^7 . Decreasing the size of the ring thus does not help us but increasing it does, however, the ring trap is already as large as we would like, thus this is not an option.

The final possibility is to reduce the chemical potential. The chemical potential is equivalent to the energy of the ground state wavefunction, thus to reduce this energy we need to reduce the energy of the ground state wavefunction. As the energy of the ground state in the Thomas-Fermi limit is mostly determined by particle interactions, we thus want to reduce the condensate density whilst maintaining the condensate number. To do this without increasing the size of the ring, we can instead increase the size of the condensate along the axial direction by either reducing the

⁵Increasing the size of the LG mode that is focussed so the full numerical aperture of the objective lens is used.

⁶Up until the density in the trap is limited due to collisional losses.

⁷The product of the increased density with the decreased area.

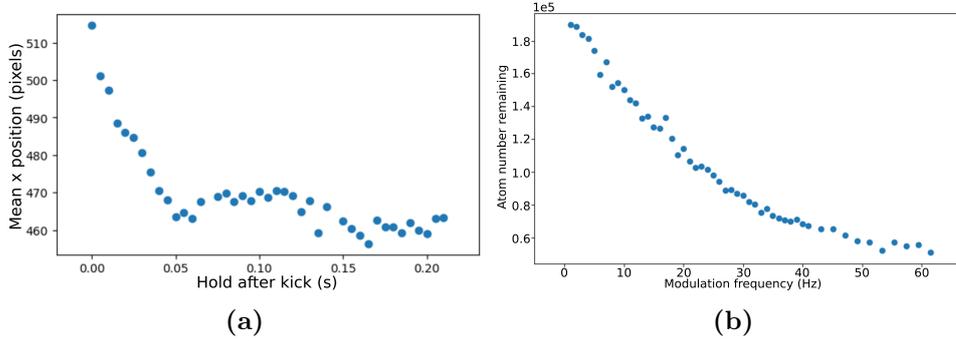


Figure 8.11: Perturbing the uniform trap **a)** shows the mean position of the cloud, as measured using the bottom imaging system, when displaced from the centre of the uniform trap by applying a bias field. Clearly the atoms move back to their central position and remain there without sloshing as would be the case in a harmonic trap. **b)** shows how the atom number in the uniform trap decays as a function of the frequency at which the power in the ring is modulated.

power in the confining sheet trap, or alternatively focussing the trap less tightly. This essentially allows us to increase the integrated density along the axial direction of the trap, without increasing the integrated density transverse to the optical axis of the uniform trap.

As the atoms are condensed in the trap, the density of the atoms reflects the underlying potential, rearranging equation A.8 we have

$$U(r) = \mu - n(r)g, \quad (8.16)$$

and noting that the peak density is given by $n_0 = \mu/g$ and substituting in we yield

$$U(r) = g(n_0 - n(r)) \quad (8.17)$$

thus from the density we can measure the normalised potential.

A fit to the normalised density profiles in figures 8.9b and 8.9c yields a potential that is most accurately approximated by a $U(r) = r^{18}$ potential. This is smaller than expected from the order of the LG mode, but agrees with the simulations of A. Gaunt[153] where the steepness of the potential is simulated as a function of a truncation parameter, determining how far up the walls of the potential the atoms explore. They found that the steepness of the potential varies dramatically with the truncation parameter, saturating at large truncation parameters, but being significantly reduced at small truncation parameters, which is the regime in which our potential is operating due to the limited laser power.

Unfortunately the intensity noise on the imaging beam is large enough that extracting the uniformity of the trap from these images is difficult, as since the optical density is given by $D(x, y) = \ln(T(x, y))$, where $T(x, y)$ is the normalised transmission of the image, then $\delta D(x, y) = -\delta T/T$, and thus fluctuations in the transmission⁸ cause significant fluctuations in the measured optical density. However, beam profiling of the ring trap outside of the glass cell show that the noise within the uniform trap is less than 1% of the peak intensity.

To further characterise the uniform trap, I measured the response of the uniform trap to different perturbations. In figure 8.11a we can see the mean position of the cloud when displaced from the trap centre, this was performed by applying a bias field to move the position of the field zero, also displacing the cloud, holding the displaced cloud for 5 ms, and then setting the bias field back to its original value. The atoms are seen to initially move back to their central position at $x = 460$, taking approximately 50 ms, before undergoing a slight bounce and then resettling at the centre of the trap. This motion does not show the expected sloshing motion one would expect in a harmonic trap, but rather is more consistent with a hard walled potential and shows that the imparted momentum is very quickly damped by the uniform potential.

The next perturbation I applied to the uniform trap was a modulation of the ring power at different frequencies. To perform this perturbation the power in the ring trap was modulated sinusoidally by 5%, much like one would to measure a parametric resonance, and the number of atoms left in the trap after 2 seconds of modulation was measured. Figure 8.11b shows the result of this measurement, which shows a clear increase in the atom loss rate with an increase in frequency up to a modulation rate of 60 Hz, after which the remaining atom number became too few to image faithfully. This data shows that the loss rate increases monotonically with the increased rate of modulation, again, in the case of a harmonic potential this is unusual, as one would expect on large loss rates near parametric resonances, however, as the atoms are confined to a uniform trap, modulation at a range of frequencies will cause atom loss.

⁸Caused for example by the varying imaging beam intensity arising from fringes moving during an image.

8.3.1 Rotating the uniform trap

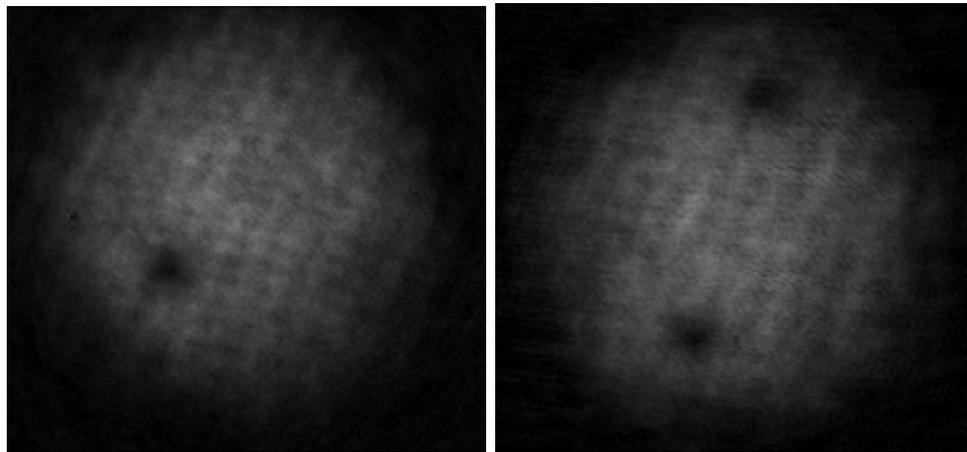
With the atoms loaded into the uniform ring trap, I next attempted to impart angular momentum to the atoms in the trap. The signature of imparting angular momentum to the condensate is the presence of vortices in the condensate.

The size of a vortex is set by the healing length, defined via equation A.14, which for densities on the order of 2×10^{14} atoms cm^{-3} is on the order of 200 nm, thus too small for us to observe *in-situ*. Instead, we must ‘blow up’ the vortex so that we can image it, this is typically achieved by allowing the condensate some time of flight expansion, which decreases the condensate density, increasing the vortex size. To image vortices in the uniform trap I typically levitate the atoms for 50 ms before imaging, figure 8.12a shows an image of a vortex in the condensate. This vortex, however, was not generated by rotating the trap, instead it was generated purely by loading the trap. Many attempts were made to remove these vortices, such as increasing the time over which the trap is loaded and ensuring that all optical traps were accurately aligned, however, it was not possible to remove these vortices in the end.

With the ability to image vortices, I next set out to rotate the optical trap to see if I could impart angular momentum to the atoms. To rotate the uniform trap the second LG mode was ramped on over 1 second at the same frequency as the first mode, to 10% of the power in the first LG mode, creating a modulation of the depth of the potential of about 50%⁹. Next, the frequency of the second mode was suddenly changed to a variable frequency, and held constant for 2 seconds before the result was imaged.

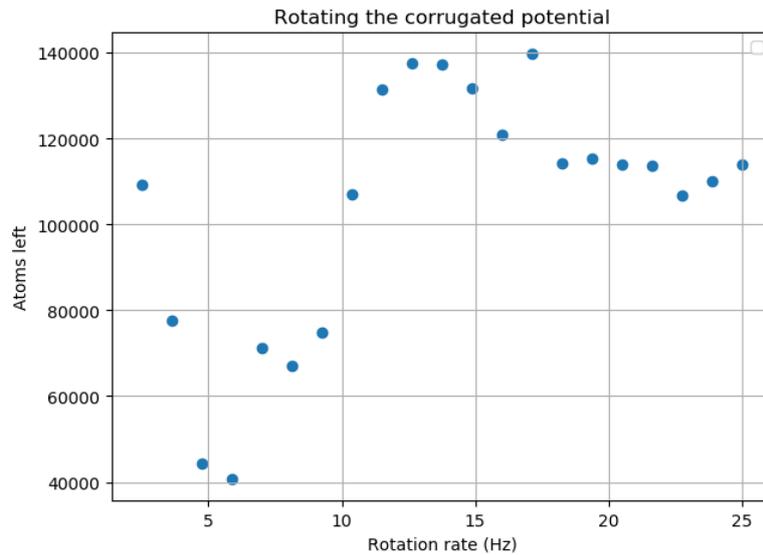
Although vortices were clearly observed in a number of images, figure 8.12b shows an image taken after rotating the corrugations at ≈ 0.1 Hz for 2 seconds, which resulted in two vortices, it was not possible to separate out the vortices that may have been generated by the rotation of the uniform trap and those that were generated solely by the loading of the uniform trap. This value does closely agree with the ‘back of the envelope’ calculation we performed earlier however, predicting a rotation rate of 0.098 Hz for the measured trap size. Atom loss was also clearly observed for a range of different rotation frequencies, as evidenced by figure 8.12c. In this figure we can see that with increasing rotation rate the atom loss rate increases up to ≈ 5

⁹The modulation depth caused by the interference of two beams of intensity I_1 and $I_2 = \beta I_1$ is given by $\frac{2\sqrt{\beta}}{1+\beta}$, this can be calculated from considering the size of the interference term when two coherent plane waves are added together.



(a)

(b)



(c)

Figure 8.12: Rotating a condensate a) and b) shows images of the condensate released from the uniform potential after 50 ms of ballistic expansion, imaged from below, which clearly show the presence of one and two vortices respectively, whilst c) shows the effect of rotating the corrugations generated by the interference of the two $LG_{(120,0)}$ modes at different frequencies.

Hz, beyond which the atom loss reduces until ≈ 12.5 Hz, after which negligible atom loss was observed. This shows that the atoms confined to the uniform potential are interacting with the corrugations.

8.4 Future outlook

In this chapter I have demonstrated the ability to trap a condensate via an optical potential generated by one of the holographic optical elements, which would not have been possible using an SLM, given there would be a further reduction in efficiency by a factor of ≈ 10 , thus the power in the dipole trap would only be on the order of a few mW. However, it would be interesting to study the characteristics of the condensate in the same shaped potential but with much greater power, such that the potential depth is much greater than the chemical potential, preventing the condensate from spilling over the edges of the trap, and exploring as far up the walls of the potential. In a much deeper trap it would also be interesting to explore the effect of rotating the potential, as although some evidence of the generation of vortices via rotating the potential was observed, this effect may be more apparent when the depth of the potential is much larger, preventing the condensate wavefunction extending outside the confining potential.

This demonstration, however, does show that HOEs can fulfil the same role as a spatial light modulator in generating interesting potentials for cold atoms, but at a much reduced cost, furthering the possibility of the commercialisation of BEC apparatus used for precision sensing, for example as a magnetic field sensor, where the atoms must be confined optically.

Appendix **A**

Bose-Einstein condensation

Bose-Einstein condensation was originally predicted by Albert Einstein as early as 1925 as a consequence of the newly developed Bose-Einstein statistics, which describes the thermodynamic behaviour of integer spin particles. The prediction that was made was that in a gas of non-interacting atoms, if the atoms were cooled to low enough temperature or alternatively a large enough density, there would be a macroscopic occupation of the ground state of the system, and the atoms would condense into the ground state of the system. It was later realised that this macroscopic occupation of the ground state would lead to a new phase of matter, in which every atom in the gas occupies the same single particle state set by the external potential in which they are confined. This has allowed the study of a quantum system which displays macroscopic quantum phenomena, which can be manipulated through magnetic and optical potentials.

The equation key to Bose-Einstein condensation is the Bose-Einstein distribution function, which gives the occupation of the i^{th} energy state of a system of indistinguishable non-interacting integer spin (bosons) particles as a function of temperature[154],

$$\langle n_i \rangle = \frac{1}{e^{(\epsilon_i - \mu)/k_B T} - 1} \quad (\text{A.1})$$

where $\langle n_i \rangle$ is the number of particles occupying the i^{th} energy state of energy ϵ_i , μ is the chemical potential of the atoms in the given potential, which can be determined by the requirement that the sum of the occupation of all states is equal to the total number of particles N , i.e. $N = \sum_i \langle n_i \rangle$, and T is the temperature of the gas. The Bose-Einstein distribution reduces to the more familiar Maxwell-Boltzmann

distribution in the case of low densities or high temperatures, however given high enough densities or low enough temperatures $\langle n_0 \rangle$ can become macroscopic, with the number of particles in the ground state approaching the number of particles in the gas, i.e. with $\langle n_0 \rangle \rightarrow N$ as $T \rightarrow 0$.

It is instructive to consider the case of Bose-Einstein condensation in a harmonic potential as this is the potential in which condensation will be achieved in this thesis. In the case of a harmonic potential the critical temperature, the temperature at which the ground state occupation becomes macroscopic, is given by [154]

$$k_B T_C = \hbar \omega \left(\frac{N}{\zeta(3)} \right)^{1/3} \quad (\text{A.2})$$

where ω is the geometric mean of the oscillation frequencies of the trap, and $\zeta(3)$ is the Riemann-zeta function. A more useful form that this condition can be expressed in, is in terms of the phase space density, D or PSD throughout, [154]

$$D = n \lambda^3 = \zeta(3/2) \gtrsim 2.6 \quad (\text{A.3})$$

where n is the number density of the gas, $n = \frac{N}{V}$, and λ is the deBroglie wavelength, given by [154]

$$\lambda = \sqrt{\frac{2\pi\hbar^2}{mk_B T}} \quad (\text{A.4})$$

where m is the mass of a single particle.

Expressing the condition for Bose-Einstein condensation in terms of the phase space density is experimentally useful, as the phase space density is a quantity that is conserved throughout a number of different thermodynamic processes, such as adiabatic compression, and expresses the condition in terms of readily observable trap parameters, such as the temperature, number and volume of the gas. The phase space density can also intuitively be understood as the probability of a particle being found within a box the size of the thermal deBroglie wavelength. The condition stated in equation. A.3 can thus interpreted as the point at which the atomic wavefunctions, of volume $\approx \lambda^3$, have begun to overlap.

Once the atoms in a harmonic potential have reached the critical temperature the atom do not immediately condense, instead the condensed fraction, $\frac{N}{N_C}$, in a harmonic potential at some temperature T below T_C can be described by

$$\frac{N}{N_C} = 1 - \left(\frac{T}{T_C} \right)^\alpha \quad (\text{A.5})$$

where α is a parameter dependent upon the trap geometry, for a harmonic oscillator $\alpha = 3$ whilst in a uniform potential $\alpha = 3/2$ [154]. Clearly at the critical temperature, $T = T_C$ and thus $\frac{N}{N_C} = 0$ and no atoms are condensed, whereas it is also clear from equation A.5 that for $\frac{N}{N_C} = 1$ then $T=0$, which is unphysical, thus it is expected that some thermal fraction will always remain¹. The final useful fraction to consider is what I shall define as a pure BEC, somewhat arbitrary given the consideration above, however I shall define it as a condensate fraction of 90%, from equation A.5 it is simple to calculate that this will occur when $\frac{T}{T_C} \approx 0.464$.

One of the most intriguing properties of a Bose-Einstein condensate is the fact that the condensate can be described by a single macroscopic wavefunction, $\psi(r, t)$ even though it consists of typically 10^4 to 10^7 atoms. If there was truly no interaction between the particles in a condensate then the condensate wavefunction would simply obey the time-dependent Schrodinger equation. In reality, there is enough interaction between the atoms that their interaction cannot be neglected, instead it is usual to employ the mean field approximation, where we include a potential term proportional to the density of the atoms, $|\Psi(r, t)|^2$, to account for the interaction. This results in what is known as the Gross-Pitaevskii equation,

$$i\hbar \frac{d\Psi}{dt} = \left[-\frac{\hbar^2}{2m} \nabla^2 + U(r, t) + g|\Psi(r, t)|^2 \right] \Psi(r, t), \quad (\text{A.6})$$

where we have introduced the potential in which the atoms are confined, $U(r, t)$, the mass of the atom, m , and $g = 4\pi\hbar^2 a/m$ is a term setting the strength of the interaction between the particles, determined by the s-wave scattering length, a .

Now, the ground state solution of equation A.6 will be a stationary state solution, thus we can separate the time and space variables in the condensate wavefunction allowing us to write it as $\Psi(r, t) = e^{-i\mu t/\hbar} \psi(r)$. Substituting this into equation A.6 yields a single second order ordinary differential equation,

$$\mu\psi(r) = \left[-\frac{\hbar^2}{2m} \nabla^2 + U(r) + g|\psi(r)|^2 \right] \psi(r), \quad (\text{A.7})$$

where all variables have the same meaning, but we have assumed a time independent potential.

Solving equation A.7 typically required numerical integration, however if we make the assumption that the mean field energy dominates the ground state kinetic energy, typically true in the case of large atom number or alternatively strong repulsive

¹Note that this is only an approximation however.

interactions, then we can ignore the kinetic energy term and we are left with what is known as the Thomas-Fermi approximation

$$\mu\psi(r) = \left[U(r) + g|\psi(r)|^2 \right] \psi(r), \quad (\text{A.8})$$

which can be trivially rearranged to give

$$n(r) = |\psi(r)|^2 = \frac{\mu - U(r)}{g}, \quad (\text{A.9})$$

where we note that as $|\psi(r)|^2$ is strictly positive, $n(r) = 0$ if $\mu - U(r) < 0$.

For the case of a harmonic potential the density will vary quadratically, as the potential does. For example, consider a harmonic potential described by $U(r) = \frac{m\omega^2}{2}x^2$, where we have used the relation $\omega = \sqrt{\frac{k}{m}}$. Now, substituting this potential into equation A.9 gives

$$n(r) = \frac{\mu}{g} \left[1 - \sum_{i=1,2,3} \frac{m\omega_i^2 x_i^2}{2\mu} \right], \quad (\text{A.10})$$

where we assumed a harmonic potential in three dimensions and summed over the orthogonal directions. From this equation we can define what is known as the Thomas-Fermi radius, R , the distance over which the density falls to zero, or equivalently the distance over which the harmonic potential increases to the chemical potential, by setting $n(r) = 0$ and solving, yielding

$$R_i = \sqrt{\frac{2\mu}{m\omega_i^2}}, \quad (\text{A.11})$$

allowing the condensate density to be written as

$$n(r) = \frac{\mu}{g} \left[1 - \sum_{i=1,2,3} \frac{x_i^2}{R_i^2} \right], \quad (\text{A.12})$$

which can be rewritten in terms of the peak condensate density, $n_0 = \frac{\mu}{g}$ as

$$n(r) = n_0 \left[1 - \sum_{i=1,2,3} \frac{x_i^2}{R_i^2} \right]. \quad (\text{A.13})$$

The final quantity that we shall define that will be relevant to this thesis is the healing length, ξ . The healing length defines the shortest distance over which the

condensate wavefunction can grow from 0 to its peak density, it is known as such as it is the distance over which the wavefunction can ‘heal’. As vortices require the condensate density to drop to zero at the core², the healing length thus sets the typical size of a vortex, alternatively it can be thought of as the smallest size a perturbation can exist in the wavefunction. The healing length is given by[155]

$$\xi = \sqrt{\frac{\hbar^2}{2mgn_0}}, \quad (\text{A.14})$$

where all variables are as defined in this appendix.

A.0.1 Thermodynamics of trapped gases

Although BECs display macroscopic quantum phenomenon, achieving BEC is most readily understood through a classical thermodynamic perspective. The thermodynamics of a gas of particles trapped in a potential can be fully described through the use of a single function, the partition function[8],

$$\zeta = \frac{V_0}{\lambda^3} \quad (\text{A.15})$$

where V_0 is the effective trap volume, given by

$$V_0 = \int \exp[U(r)/k_B T] d^3r \quad (\text{A.16})$$

where $U(r)$ is the confining potential, defined such that its minimum is zero and hence V_0 is strictly positive.

Once the partition function is known, all other thermodynamic quantities for the gas can then be calculated, in particular the peak phase space density is given by

$$D = \frac{N}{\zeta}, \quad (\text{A.17})$$

and the density can be calculated through

$$n(r) = \frac{N}{V_0} \exp(-U(r)/k_B T). \quad (\text{A.18})$$

Most relevant parameters have so far been described as cold atom systems are rela-

²discussed more in section. B

tively simple, and can mostly be characterised by the number of atoms, their temperature and thus the volume they occupy in a given potential. From these parameters the most important parameter experimentally when attempting to achieve BEC, the phase space density³, D , can also be calculated.

One parameter however which has not been described is the elastic scattering rate, k_{el} , this is given by

$$k_{el} = n\sigma v, \quad (\text{A.19})$$

where n is the density of atoms, v is the average velocity of the atoms and σ is the elastic scattering cross section. The elastic scattering rate is an important parameter in all BEC experiments to date, as BEC is yet to be achieved without forced evaporation. Forced evaporation relies upon the rethermalisation of a system once the high energy tail of the distribution has been truncated, removing atoms with greater than the average thermal energy, and overall reducing the average thermal energy once thermalised. As thermalisation of a gas occurs through the atoms in the gas scattering off one another, the elastic scattering rate thus determines the rate at which the gas will rethermalise thus determining the rate at which forced evaporation can occur and ultimately its efficiency.

³As the point at which condensation occurs is specified in terms of the phase space density, experimentally this is our most important parameter.

Appendix B

Vortices

Vortices arise in an enormous range of physical systems, both classical and quantum, from macroscopic such as a hurricane, through to microscopic, such as in a superfluid Helium, wherever turbulent flow is present, vortices will follow. One of the final great unsolved problems in classical physics is that of turbulence, as turbulence arises in chaotic systems that are too complex to model, however vortices are prolific in any system with turbulent flow and by understanding the behaviour of the vortices one might be able to understand the chaotic behaviour of the flow.

In this thesis quantum vortices¹, also known as singularities, present themselves in two separate systems, a spatially coherent optical field and a spatially coherent atomic wavefunction, and arise as a consequence of the ability to describe both systems with a single macroscopic wavefunction, $\psi(r)$. The interested reader is referred to *'Twisted photons - Applications of light with orbital angular momentum'* [79] which provides a detailed introduction to vortices in light, as well as their use to generate vortices in atomic wavefunctions, as the theory will not be detailed here. However, the pertinent features of a vortex are such;

- Quantised vortices arise in systems that can be described by a macroscopic wavefunction
- They arise due to the macroscopic phase coherence of the system, which ensures the phase is singly defined at all points
- Around any closed loop, the phase must thus change only by $2n\pi$, where n is

¹As opposed to classical vortices, in which the circulation about the vortex is not quantised.

an integer, as otherwise the phase will be multiply defined.

- The integer n determines the total circulation that is enclosed within the loop
- At the centre of a loop containing a single quantised vortex, the phase is undefined, to ensure the wavefunction remains defined at this point the wavefunction must also vanish at this point.

Thus we can see that vortices are mostly characterised by windings in the phase of the wavefunction not the density, although they do require the density vanish at the centre of a vortex. In the case of optical fields these vortices carry orbital angular momentum, with an optical field containing N vortices carrying $N\hbar$ units of angular momentum per photon. In the case of a macroscopic atomic wavefunction, i.e. a condensate, these vortices also carry angular momentum, and for a single vortex centred in a condensate, each atom will carry \hbar of angular momentum per atom. As vortices carry angular momentum, they also have a handedness dependent upon the direction in which the phase winds about the vortex and also are topological features in the wavefunction, thus cannot spontaneously decay, but rather have to annihilate with another vortex of opposite handedness.

Appendix C

Transconductance of the Magneato driver

To calculate the transconductance of the Magneato driver, we begin by looking at the sense voltage, V_s , defined as the voltage input to the sensing op-amp, AD820. This sense voltage is generated by the flow of current from the LA-150P current transducer, which induces a voltage across the sense resistor as the current flows to ground. The sense current I_s is 2000 times smaller than the coil current, I_c , thus

$$I_s = \frac{I_c}{2000}. \quad (\text{C.1})$$

As this current flows through the sense resistor of resistance R_s , a voltage difference is generated lifting the input of the op amp to

$$V_s = I_s R_s \quad (\text{C.2})$$

above ground. This op-amp is configured as a voltage follower thus the input is equal to the output voltage. The output of the op amp then passes through the voltage divider formed by R_4 and R_6 , producing a voltage to the negative input of the OP27, V_- of

$$V_- = V_s \frac{R_4}{R_4 + R_6}. \quad (\text{C.3})$$

Meanwhile the control voltage, V_c passes through the op-amp AD620 which is also configured as a voltage follower, and then through the voltage divider formed by R_2

and R_3 , yielding an input voltage to the positive side of the op amp, V_+ of

$$V_+ = V_c \frac{R_3}{R_3 + R_2}. \quad (\text{C.4})$$

Now, noting that in equilibrium the positive and negative inputs to OP27 will be balanced, this yields

$$V_s \frac{R_4}{R_4 + R_6} = V_c \frac{R_3}{R_3 + R_2} \quad (\text{C.5})$$

and further noting that $R_2 = R_4$ and $R_3 = R_6$, the above equation can be rearranged, and through the substitution of eq:C.2 we arrive at

$$I_c = V_c \frac{2000 R_3}{R_s R_4} \quad (\text{C.6})$$

yielding the transconductance of the Magneato driver.

Phase space density calculations

It is useful to be able to calculate the phase space density of the atoms in a harmonic potential to be able to know how far one is from condensation. To be able to calculate the phase space density, the number of atoms, the trap volume as well as the temperature of the atoms all must be known.

As the in-situ optical density of the atoms is typically on the order of 10-100, accurate imaging of the size of the cloud is difficult, and would need to be performed along two orthogonal axes. Instead, we can calculate the trap volume knowing the temperature of the atoms as well as the trap frequency via equation. 6.35. We can rewrite equation. 6.35 in terms of the measured parameter, the trap frequency, $2\pi f = \omega$ yielding

$$V(T) = \left(\frac{k_B}{2\pi m}\right)^{3/2} \frac{T^{3/2}}{f_x f_y f_z} \quad (\text{D.1})$$

where T is the temperature of the atoms and f_i is the trapping frequency in the i^{th} direction. Next, knowing the definition of the peak phase space density, equation. A.17, we have

$$D = \frac{N\lambda^3}{V} \quad (\text{D.2})$$

and combining with equation. A.4 and equation. D.1 we yield

$$D = N \left(\frac{2\pi\hbar}{k_B}\right)^3 \frac{f_x f_y f_z}{T^3} \approx 1.1 \times 10^{-31} N \frac{f_x f_y f_z}{T^3} \quad (\text{D.3})$$

which allows the phase space density to be calculated for a harmonic potential with knowledge solely of the atom number, temperature and trapping frequencies.

Bibliography

- [1] V. Arrizon, U. Ruiz, R. Carrada, and L. Gonzalez, “Pixelated phase only computer holograms for the accurate encoding of scalar complex fields,” *Journal of the Optical Society of America*, Vol.24(11), pp.3500-7, 2007.
- [2] S. J. Tempone-Wiltshire, S. P. Johnstone, and K. Helmerson, “Optical vortex knots one photon at a time,” *Scientific Reports*, vol. 6, p. 24463, Apr. 2016.
- [3] H. Berneth, F.-K. Bruder, T. Fäcke, D. Jurbergs, R. Hagen, D. Hönel, T. Rölle, and G. Walze, “Bayfol HX photopolymer for full-color transmission volume Bragg gratings,,” *Proceedings of SPIE*, vol. 9006, p. 900602, 2014.
- [4] S. Tempone-Wiltshire, S. Johnstone, and K. Helmerson, “High efficiency, low cost holographic optical elements for ultracold atom trapping,” *Optics Express*, vol. 25, pp. 296–304, 2017.
- [5] P. D. Lett, W. D. Phillips, S. L. Rolston, C. E. Tanner, R. N. Watts, and C. I. Westbrook, “Optical molasses,” *J. Opt. Soc. Am. B*, vol. 6, pp. 2084–2107, Nov 1989.
- [6] J. Schoser, A. Batär, R. Löw, V. Schweikhard, A. Grabowski, Y. B. Ovchinnikov, and T. Pfau, “Intense source of cold Rb atoms from a pure two-dimensional magneto-optical trap,” *Phys. Rev. A*, vol. 66, p. 023410, Aug 2002.
- [7] W. Ketterle and N. V. Druten, “Evaporative cooling of trapped atoms,” vol. 37 of *Advances In Atomic, Molecular, and Optical Physics*, pp. 181 – 236, Academic Press, 1996.

-
- [8] 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,” *Phys. Rev. A*, vol. 79, p. 063631, Jun 2009.
- [9] T. Kovachy, P. Asenbaum, C. Overstreet, C. A. Donnelly, S. M. Dickerson, A. Sugarbaker, J. M. Hogan, and M. A. Kasevich, “Quantum superposition at the half-metre scale,” *Nature*, vol. 528, pp. 530 EP –, 12 2015.
- [10] A. Ashkin, “Acceleration and trapping of particles by radiation pressure,” *Phys. Rev. Lett.*, vol. 24, pp. 156–159, Jan 1970.
- [11] A. Ashkin and J. Dziedzic, “Optical trapping and manipulation of viruses and bacteria,” *Science*, vol. 235, no. 4795, pp. 1517–1520, 1987.
- [12] S. Chu, J. E. Bjorkholm, A. Ashkin, and A. Cable, “Experimental observation of optically trapped atoms,” *Phys. Rev. Lett.*, vol. 57, pp. 314–317, Jul 1986.
- [13] R. Grimm, M. Weidemller, and Y. B. Ovchinnikov, “Optical dipole traps for neutral atoms,” vol. 42 of *Advances In Atomic, Molecular, and Optical Physics*, pp. 95 – 170, Academic Press, 2000.
- [14] J. Kirk and A. Jones, “Phase only complex-valued spatial filter,” *Journal of the Optical Society of America Vol. 61 No. 8*, pp. 1023–1028, 1971.
- [15] M. Pasienski and B. DeMarco, “A high-accuracy algorithm for designing arbitrary holographic atom traps,” *Opt. Express*, vol. 16, pp. 2176–2190, Feb 2008.
- [16] A. L. Gaunt and Z. Hadzibabic, “Robust digital holography for ultracold atom trapping,” *Scientific Reports*, vol. 2, p. 721, 10 2012.
- [17] E. R. Shanblatt and D. G. Grier, “Extended and knotted optical traps in three dimensions,” *Opt. Express*, vol. 19, pp. 5833–5838, Mar 2011.
- [18] J. A. Rodrigo, T. Alieva, E. Abramochkin, and I. Castro, “Shaping of light beams along curves in three dimensions,” *Opt. Express*, vol. 21, pp. 20544–20555, Sep 2013.
- [19] E. Noether, “Invariante variationsprobleme,” *Math-phys. Klasse*, pp. 235–257, 1918.

-
- [20] A. J. Groszek, T. P. Simula, D. M. Paganin, and K. Helmerson, “Onsager vortex formation in Bose-Einstein condensates in two-dimensional power-law traps,” *Phys. Rev. A*, vol. 93, p. 043614, Apr 2016.
- [21] G. Boffetta, A. Celani, and M. Vergassola, “Inverse energy cascade in two-dimensional turbulence: Deviations from Gaussian behavior,” *Phys. Rev. E*, vol. 61, pp. R29–R32, Jan 2000.
- [22] L. Biferale, S. Musacchio, and F. Toschi, “Inverse energy cascade in three-dimensional isotropic turbulence,” *Phys. Rev. Lett.*, vol. 108, p. 164501, Apr 2012.
- [23] M. T. Reeves, T. P. Billam, B. P. Anderson, and A. S. Bradley, “Inverse energy cascade in forced two-dimensional quantum turbulence,” *Phys. Rev. Lett.*, vol. 110, p. 104501, Mar 2013.
- [24] A. N. Ganshin, V. B. Efimov, G. V. Kolmakov, L. P. Mezhov-Deglin, and P. V. E. McClintock, “Observation of an inverse energy cascade in developed acoustic turbulence in superfluid helium,” *Phys. Rev. Lett.*, vol. 101, p. 065303, Aug 2008.
- [25] M. A. Rutgers, “Forced 2d turbulence: Experimental evidence of simultaneous inverse energy and forward enstrophy cascades,” *Phys. Rev. Lett.*, vol. 81, pp. 2244–2247, Sep 1998.
- [26] R. M. B. Young and P. L. Read, “Forward and inverse kinetic energy cascades in Jupiter’s turbulent weather layer,” *Nature Physics*, vol. 13, pp. 1135 EP–, 08 2017.
- [27] T.-L. Ho and Q. Zhou, “Obtaining the phase diagram and thermodynamic quantities of bulk systems from the densities of trapped gases,” *Nature Physics*, vol. 6, p. 131, Dec. 2009.
- [28] Y. Sagi, T. E. Drake, R. Paudel, and D. S. Jin, “Measurement of the homogeneous contact of a unitary Fermi gas,” *Phys. Rev. Lett.*, vol. 109, p. 220402, Nov 2012.
- [29] A. L. Gaunt, T. F. Schmidutz, I. Gotlibovych, R. P. Smith, and Z. Hadzibabic, “Bose-Einstein condensation of atoms in a uniform potential,” *Phys. Rev. Lett.*, vol. 110, p. 200406, May 2013.

- [30] B. Mukherjee, Z. Yan, P. B. Patel, Z. Hadzibabic, T. Yefsah, J. Struck, and M. W. Zwierlein, “Homogeneous atomic Fermi gases,” *Phys. Rev. Lett.*, vol. 118, p. 123401, Mar 2017.
- [31] I. Bloch, J. Dalibard, and W. Zwerger, “Many-body physics with ultracold gases,” *Rev. Mod. Phys.*, vol. 80, pp. 885–964, Jul 2008.
- [32] Z. Hadzibabic, P. Krüger, M. Cheneau, B. Battelier, and J. Dalibard, “Berezinskii–Kosterlitz–Thouless crossover in a trapped atomic gas,” *Nature*, vol. 441, no. 7097, pp. 1118–1121, 2006.
- [33] J. Struck, C. Ölschläger, R. Le Targat, P. Soltan-Panahi, A. Eckardt, M. Lewenstein, P. Windpassinger, and K. Sengstock, “Quantum simulation of frustrated classical magnetism in triangular optical lattices,” *Science*, vol. 333, no. 6045, pp. 996–999, 2011.
- [34] 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, no. 6867, pp. 39–44, 2002.
- [35] B. T. Seaman, M. Krämer, D. Z. Anderson, and M. J. Holland, “Atomtronics: Ultracold-atom analogs of electronic devices,” *Phys. Rev. A*, vol. 75, p. 023615, Feb 2007.
- [36] 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,” *Phys. Rev. Lett.*, vol. 106, p. 130401, Mar 2011.
- [37] S. Moulder, S. Beattie, R. P. Smith, N. Tammuz, and Z. Hadzibabic, “Quantized supercurrent decay in an annular Bose-Einstein condensate,” *Phys. Rev. A*, vol. 86, p. 013629, Jul 2012.
- [38] Y. Shin, M. Saba, T. A. Pasquini, W. Ketterle, D. E. Pritchard, and A. E. Leanhardt, “Atom interferometry with Bose-Einstein condensates in a double-well potential,” *Phys. Rev. Lett.*, vol. 92, p. 050405, Feb 2004.
- [39] 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, 03 2012.

- [40] D. Barredo, V. Lienhard, S. de Léséleuc, T. Lahaye, and A. Browaeys, “Synthetic three-dimensional atomic structures assembled atom by atom,” *Nature*, vol. 561, no. 7721, pp. 79–82, 2018.
- [41] H. Rubinsztein-Dunlop, A. Forbes, M. V. Berry, M. R. Dennis, D. L. Andrews, M. Mansuripur, C. Denz, C. Alpmann, P. Banzer, T. Bauer, E. Karimi, L. Marrucci, M. Padgett, M. Ritsch-Marte, N. M. Litchinitser, N. P. Bigelow, C. Rosales-Guzmán, A. Belmonte, J. P. Torres, T. W. Neely, M. Baker, R. Gordon, A. B. Stilgoe, J. Romero, A. G. White, R. Fickler, A. E. Willner, G. Xie, B. McMorran, and A. M. Weiner, “Roadmap on structured light,” *Journal of Optics*, vol. 19, p. 013001, nov 2016.
- [42] D. Jurbergs, F.-K. Bruder, F. Deuber, T. F äcke, R. Hagen, D. Hönel, T. Rölle, M.-S. Weiser, and A. Volkov, “New recording materials for the holographic industry,” *Proceedings of SPIE*, vol. 7233, p. 72330K, 2009.
- [43] S. Abend, M. Gebbe, M. Gersemann, H. Ahlers, H. Müntinga, E. Giese, N. Gaaloul, C. Schubert, C. Lämmerzahl, W. Ertmer, W. P. Schleich, and E. M. Rasel, “Atom-chip fountain gravimeter,” *Phys. Rev. Lett.*, vol. 117, p. 203003, Nov 2016.
- [44] S.-Y. Lan, P.-C. Kuan, B. Estey, P. Haslinger, and H. Müller, “Influence of the coriolis force in atom interferometry,” *Phys. Rev. Lett.*, vol. 108, p. 090402, Feb 2012.
- [45] I. Dutta, D. Savoie, B. Fang, B. Venon, C. L. Garrido Alzar, R. Geiger, and A. Landragin, “Continuous cold-atom inertial sensor with 1 nrad/sec rotation stability,” *Phys. Rev. Lett.*, vol. 116, p. 183003, May 2016.
- [46] T. Miller, M. Gilowski, M. Zaiser, and et al., “A compact dual atom interferometer gyroscope based on laser-cooled rubidium,” *Eur. Phys. J. D*, vol. 53, pp. 273–281, 2009.
- [47] B. Battelier, B. Barrett, L. Fouch, L. Chichet, L. Antoni-Micollier, H. Porte, F. Napolitano, J. Lautier, A. Landragin, and P. Bouyer, “Development of compact cold-atom sensors for inertial navigation,” p. 990004, 04 2016.
- [48] V. Mnoret, R. Geiger, G. Stern, P. Cheinet, B. Battelier, N. Zahzam, F. P. D. Santos, A. Bresson, A. Landragin, and P. Bouyer, “A transportable cold atom inertial sensor for space applications,” in *International Conference on Space Optics ICSSO 2010* (E. Armandillo, B. Cugny, and N. Karafolas, eds.),

- vol. 10565, pp. 698 – 704, International Society for Optics and Photonics, SPIE, 2017.
- [49] M. Vengalattore, J. M. Higbie, S. R. Leslie, J. Guzman, L. E. Sadler, and D. M. Stamper-Kurn, “High-resolution magnetometry with a spinor Bose-Einstein condensate,” *Phys. Rev. Lett.*, vol. 98, p. 200801, May 2007.
- [50] J. Scherschligt, J. A. Fedchak, D. S. Barker, S. Eckel, N. Klimov, C. Makrides, and E. Tiesinga, “Development of a new UHV/XHV pressure standard (cold atom vacuum standard),” *Metrologia*, vol. 54, pp. S125–S132, nov 2017.
- [51] J. P. McGilligan, P. F. Griffin, R. Elvin, S. J. Ingleby, E. Riis, and A. S. Arnold, “Grating chips for quantum technologies,” *Scientific Reports*, vol. 7, no. 1, p. 384, 2017.
- [52] M. Vangeleyn, P. F. Griffin, E. Riis, and A. S. Arnold, “Single-laser, one beam, tetrahedral magneto-optical trap,” *Opt. Express*, vol. 17, pp. 13601–13608, Aug 2009.
- [53] G. P. Lewis, Z. Moktadir, C. O. Gollasch, M. Kraft, S. Pollock, F. Ramirez-Martinez, J. Ashmore, A. Laliotis, M. Trupke, and E. A. Hinds, “Fabrication of magneto-optical atom traps on a chip,” *Journal of Microelectromechanical Systems*, vol. 18, pp. 347–353, 2009.
- [54] J. Arlt, O. Marag, S. Webster, S. Hopkins, and C. Foot, “A pyramidal magneto-optical trap as a source of slow atoms,” *Optics Communications*, vol. 157, no. 1, pp. 303 – 309, 1998.
- [55] C. C. Nshii, M. Vangeleyn, J. P. Cotter, P. F. Griffin, E. A. Hinds, C. N. Ironside, P. See, A. G. Sinclair, E. Riis, and A. S. Arnold, “A surface-patterned chip as a strong source of ultracold atoms for quantum technologies,” *Nature Nanotechnology*, vol. 8, pp. 321 EP –, 04 2013.
- [56] J. P. McGilligan, P. F. Griffin, E. Riis, and A. S. Arnold, “Diffraction-grating characterization for cold-atom experiments,” *J. Opt. Soc. Am. B*, vol. 33, pp. 1271–1277, Jun 2016.
- [57] M. B. Crookston, P. M. Baker, and M. P. Robinson, “A microchip ring trap for cold atoms,” *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 38, pp. 3289–3298, sep 2005.

- [58] J. A. Sauer, M. D. Barrett, and M. S. Chapman, “Storage ring for neutral atoms,” *Phys. Rev. Lett.*, vol. 87, p. 270401, Dec 2001.
- [59] F. K. Fatemi and M. Bashkansky, “Spatially resolved magnetometry using cold atoms in dark optical tweezers,” *Opt. Express*, vol. 18, pp. 2190–2196, Feb 2010.
- [60] S. Kulin, S. Aubin, S. Christe, B. Peker, S. L. Rolston, and L. A. Orozco, “A single hollow-beam optical trap for cold atoms,” *Journal of Optics B: Quantum and Semiclassical Optics*, vol. 3, pp. 353–357, oct 2001.
- [61] N. Davidson, “Dark optical traps for ultra-cold atoms,” in *19th Congress of the International Commission for Optics: Optics for the Quality of Life* (G. C. Righini and A. Consortini, eds.), vol. 4829, pp. 481 – 482, International Society for Optics and Photonics, SPIE, 2003.
- [62] G. Ferrari, N. Poli, F. Sorrentino, and G. M. Tino, “Long-lived bloch oscillations with bosonic Sr atoms and application to gravity measurement at the micrometer scale,” *Phys. Rev. Lett.*, vol. 97, p. 060402, Aug 2006.
- [63] D. Engstrom, M. Persson, J. Bengtsson, and M. Goksor, “Calibration of spatial light modulators suffering from spatially varying phase response,” *Optics Express: Vol.21 No.13*, 2013.
- [64] A. Jesacher, C. Maurer, A. Schwaighofer, S. Bernet, and M. Ritsch-Marte, “Full phase and amplitude control of holographic optical tweezers,” *Optics Express*, 2008.
- [65] J. Leach, M. R. Dennis, J. Courtial, and M. J. Padgett, “Vortex knots in light,” *New Journal of Physics*, Vol. 7, pp. 55, 2005.
- [66] S. Tempone-Wiltshire, *Towards the generation of knotted vortices in a Bose-Einstein Condensate*. PhD thesis, Monash University School of Physics and Astronomy, 2014.
- [67] R. Feynman, M. Sands, and R. Leighton, *The Feynman Lectures on Physics*. Ingram publisher services US, January 2011.
- [68] S. Parker, “A single-photon double-slit interference experiment.,” *American Journal of Physics*, vol. 39, pp. 420–424, 1971.
- [69] W. Rueckner and P. Titcomb, “A lecture demonstration of single photon interference,” *American Journal of Physics*, vol. 64, pp. 184–188, 1996.

- [70] S. Kocsis, B. B. S. Ravets, M. J. Stevens, R. P. Mirin, L. K. Shalm, and A. M. Steinberg, “Observing the average trajectories of single photons in a two-slit interferometer,” *Science*, vol. 332, pp. 1170–1173, 2011.
- [71] P. G. Meril, G. F. Missiroli, and G. Pozzi, “On the statistical aspect of electron interference phenomena,” *American Journal of Physics*, vol. 44, pp. 306–307, 1976.
- [72] A. Tonomura, J. Endo, T. Matsuda, T. Kawasaki, and H. Ezawa, “Demonstration of single-electron buildup of an interference pattern,” *American Journal of Physics*, vol. 57, pp. 117–120, 1989.
- [73] R. Bach, D. Pope, S. H. Liou, and H. Batelaan, “Controlled double-slit electron diffraction,” *New Journal of Physics*, vol. 15, 2013.
- [74] O. Carnal and J. Mlynek, “Young’s double-slit experiment with atoms: A simple atom interferometer,” *Phys. Rev. Letters*, vol. 66, pp. 2689–2692, 1991.
- [75] M. Arndt, O. Nairz, J. Vos-Andreae, C. Keller, G. van der Zouw, and A. Zeilinger, “Wave-particle duality of C₆₀ molecules,” *Nature*, vol. 401, pp. 680–682, 1999.
- [76] S.-W. Lee and J. Kim, “Single-photon characteristics of superposed weak coherent states,” *Phys. Rev. A*, vol. 99, p. 013847, Jan 2019.
- [77] G. Roumpos and S. T. Cundiff, “Photon number distributions from a diode laser,” *Optics Letters*, vol. 38, pp. 139–141, 2013.
- [78] M. Razeghi, *Technology of quantum devices*, vol. 1. Springer Science and Business Media, 2009.
- [79] J. P. Torres and L. Torner, *Twisted photons - applications of light with orbital angular momentum*. Wiley-VCH, 1 ed., 2011.
- [80] M. Reicherter, J. Liesener, T. Haist, and H. Tiziani, “Advantages of holographic optical tweezers,” *Proceedings of SPIE-OSA Biomedical Optics*, vol. 5143, 2003.
- [81] J. García-Márquez, V. López, A. González-Vega, and E. Noé, “Flicker in a twisted nematic spatial light modulator,” *Optics and Lasers in Engineering*, vol. 51, p. 741, 2013.

- [82] J. García-Márquez, V. López, A. González-Vega, and E. Noé, “Flicker minimization in an lcos spatial light modulator,” *Optics Express*, vol. 20, p. 8431, 2012.
- [83] H. Berneth, F. K. Bruder, T. Fäcke, R. Hagen, D. Hönel, D. Jurbergs, T. Rölle, and M.-S. Weiser, “Holographic recording aspects of high-resolution bayfol HX photopolymer,” *Proceedings of SPIE*, vol. 7957, p. 79570H, 2011.
- [84] A. E. Siegman, *Lasers*. University Science Books, 1986.
- [85] K. S. Hardman, *A BEC Based Precision Gravimeter and Magnetic Gradiometer: Design and Implementation*. PhD thesis, Research school of Physics and Engineering, Australian National University, 2016.
- [86] D. E. Fagnan, J. Wang, C. Zhu, P. Djuricanin, B. G. Klappauf, J. L. Booth, and K. W. Madison, “Observation of quantum diffractive collisions using shallow atomic traps,” *Phys. Rev. A*, vol. 80, p. 022712, Aug 2009.
- [87] S. Chaudhuri, S. Roy, and C. S. Unnikrishnan, “Realization of an intense cold Rb atomic beam based on a two-dimensional magneto-optical trap: Experiments and comparison with simulations,” *Phys. Rev. A*, vol. 74, p. 023406, Aug 2006.
- [88] SAES getters, “Nextorr d 200-5.” Datasheet, Accessed 2018.
- [89] Gamma Vacuum, “Titanium sublimation pumping (tsp).” Brochure, Accessed 2018.
- [90] K. Dieckmann, R. J. C. Spreeuw, M. Weidemüller, and J. T. M. Walraven, “Two-dimensional magneto-optical trap as a source of slow atoms,” *Phys. Rev. A*, vol. 58, pp. 3891–3895, Nov 1998.
- [91] J. H. Moore, C. C. Davis, and M. A. Coplam, *Building Scientific Apparatus*. Perseus Books, 2002.
- [92] L. S. Lerner, *Physics for Scientists and Engineers, Vol. 2*. Jones and Bartlett Learning, 1997.
- [93] J. H. Dellinger, “The temperature coefficient of resistance of copper,” *Bulletin of the Bureau of Standards*, vol. 7, pp. 71–97, 1911.
- [94] Ametherm, “NTC thermistor calculating the temperature coefficient of a thermistor.” Webpage, Accessed 2017.

- [95] T. R. Kuphaldt, “Voltage divider circuits.” Online textbook, Accessed 2017.
- [96] P. Lebedev, “Untersuchungen ber die druckkrfte des lichtetes,” *Annalen der Physik*, vol. 4, 1901.
- [97] E. F. Nichols and G. F. Hull, “The pressure due to radiation,” *The Astrophysical Journal*, vol. 17, pp. 315–351, 1903.
- [98] T. W. Hansch and A. L. Schawlow, “Cooling of gases by laser radiation,” *Optics communications*, vol. 13, pp. 68–69, 1975.
- [99] D. Wineland and H. Dehmelt, “Proposed $10^{14}\Delta\nu < \nu$ laser fluorescence spectroscopy on TI^+ mono-ion oscillator III,” *Bull. Am. Phys. Soc.*, vol. 20, p. 637, 1975.
- [100] D. J. Wineland, R. E. Drullinger, and F. L. Walls, “Radiation-pressure cooling of bound resonant absorbers,” *Phys. Rev. Lett.*, vol. 40, pp. 1639–1642, Jun 1978.
- [101] D. A. Steck, “Rubidium 87 d line data.” Online reference material, Accessed 2019.
- [102] Letokhov V., V. Minogin, and B. Pavlik, “Cooling and trapping atoms and molecules by a resonant laser field,” *Opt. Comms.*, vol. 19, p. 72, 1976.
- [103] Phillips, W. D., and J. V. Prodan, “*Chirping the Light-Fantastic?*” in *Laser-Cooled and Trapped Atoms*. Natl. Bur. Stand., Washington, DC, Spec. Publ. 653, 1983.
- [104] W. Ertmer, R. Blatt, J. L. Hall, and M. Zhu, “Laser manipulation of atomic beam velocities: Demonstration of stopped atoms and velocity reversal,” *Phys. Rev. Lett.*, vol. 54, pp. 996–999, Mar 1985.
- [105] W. D. Phillips and H. Metcalf, “Laser deceleration of an atomic beam,” *Phys. Rev. Lett.*, vol. 48, pp. 596–599, Mar 1982.
- [106] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard, “Trapping of neutral sodium atoms with radiation pressure,” *Phys. Rev. Lett.*, vol. 59, pp. 2631–2634, Dec 1987.
- [107] A. M. Steane, M. Chowdhury, and C. J. Foot, “Radiation force in the magneto-optical trap,” *J. Opt. Soc. Am. B*, vol. 9, pp. 2142–2158, Dec 1992.

- [108] C. Monroe, W. Swann, H. Robinson, and C. Wieman, “Very cold trapped atoms in a vapor cell,” *Phys. Rev. Lett.*, vol. 65, pp. 1571–1574, Sep 1990.
- [109] F. Reif, *Fundamentals of Statistical and Thermal Physics*. McGraw-Hill, New York, 1965.
- [110] A. Einstein, “On the motion of small particles suspended in liquids at rest required by the molecular-kinetic theory of heat,” *Annalen der Physik*, vol. 17, pp. 549–560, 1905.
- [111] W. T. Coffey and Y. P. Kalmykov, *The Langevin Equation*. WORLD SCIENTIFIC, 3rd ed., 2012.
- [112] C. J. Cooper, G. Hillenbrand, J. Rink, C. G. Townsend, K. Zetie, and C. J. Foot, “The temperature of atoms in a magneto-optical trap,” *Europhysics Letters (EPL)*, vol. 28, pp. 397–402, nov 1994.
- [113] H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping*. Springer-Verlag New York, 1999.
- [114] T. Walker, D. Sesko, and C. Wieman, “Collective behavior of optically trapped neutral atoms,” *Phys. Rev. Lett.*, vol. 64, pp. 408–411, Jan 1990.
- [115] D. W. Sesko, T. G. Walker, and C. E. Wieman, “Behavior of neutral atoms in a spontaneous force trap,” *J. Opt. Soc. Am. B*, vol. 8, pp. 946–958, May 1991.
- [116] C. G. Townsend, N. H. Edwards, C. J. Cooper, K. P. Zetie, C. J. Foot, A. M. Steane, P. Szriftgiser, H. Perrin, and J. Dalibard, “Phase-space density in the magneto-optical trap,” *Phys. Rev. A*, vol. 52, pp. 1423–1440, Aug 1995.
- [117] W. D. Phillips, “Laser cooling and trapping of neutral atoms,” *Reviews of Modern Physics*, vol. 70, pp. 721–741, 1998.
- [118] P. D. Lett, R. N. Watts, C. I. Westbrook, W. D. Phillips, P. L. Gould, and H. J. Metcalf, “Observation of atoms laser cooled below the [doppler] limit,” *Phys. Rev. Lett.*, vol. 61, pp. 169–172, Jul 1988.
- [119] J. Dalibard and C. Cohen-Tannoudji, “Laser cooling below the Doppler limit by polarization gradients: simple theoretical models,” *J. Opt. Soc. Am. B*, vol. 6, pp. 2023–2045, Nov 1989.
- [120] P. J. Ungar, D. S. Weiss, E. Riis, and S. Chu, “Optical molasses and multilevel atoms: theory,” *J. Opt. Soc. Am. B*, vol. 6, pp. 2058–2071, Nov 1989.

- [121] J. Werner and H. Wallis, “Laser cooling by σ^+ - σ^- circularly polarized beams of unequal intensities,” *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 26, pp. 3063–3080, sep 1993.
- [122] A. M. Steane and C. J. Foot, “Laser cooling below the Doppler limit in a magneto-optical trap,” *Europhysics Letters (EPL)*, vol. 14, pp. 231–236, feb 1991.
- [123] S.-Q. Shang, B. Sheehy, H. Metcalf, P. van der Straten, and G. Nienhuis, “Velocity-selective resonances and sub-Doppler laser cooling,” *Phys. Rev. Lett.*, vol. 67, pp. 1094–1097, Aug 1991.
- [124] M. Walhout, J. Dalibard, S. L. Rolston, and W. D. Phillips, “ σ^+ - σ^- optical molasses in a longitudinal magnetic field,” *J. Opt. Soc. Am. B*, vol. 9, pp. 1997–2007, Nov 1992.
- [125] A. M. Steane, G. Hillenbrand, and C. J. Foot, “Polarization gradient cooling in a one-dimensional σ^+ - σ^- configuration for any atomic transition,” *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 25, pp. 4721–4743, nov 1992.
- [126] W. Petrich, M. H. Anderson, J. R. Ensher, and E. A. Cornell, “Behavior of atoms in a compressed magneto-optical trap,” *J. Opt. Soc. Am. B*, vol. 11, pp. 1332–1335, Aug 1994.
- [127] M. T. DePue, S. L. Winoto, D. Han, and D. S. Weiss, “Transient compression of a MOT and high intensity fluorescent imaging of optically thick clouds of atoms,” *Optics Communications*, vol. 180, no. 1, pp. 73 – 79, 2000.
- [128] Z. T. Lu, K. L. Corwin, M. J. Renn, M. H. Anderson, E. A. Cornell, and C. E. Wieman, “Low-velocity intense source of atoms from a magneto-optical trap,” *Phys. Rev. Lett.*, vol. 77, pp. 3331–3334, Oct 1996.
- [129] H. Wu and C. J. Foot, “Direct simulation of evaporative cooling,” *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 29, pp. L321–L328, apr 1996.
- [130] C. R. Monroe, E. A. Cornell, C. A. Sackett, C. J. Myatt, and C. E. Wieman, “Measurement of Cs-Cs elastic scattering at $t=30 \mu\text{k}$,” *Phys. Rev. Lett.*, vol. 70, pp. 414–417, Jan 1993.

- [131] O. J. Luiten, M. W. Reynolds, and J. T. M. Walraven, “Kinetic theory of the evaporative cooling of a trapped gas,” *Phys. Rev. A*, vol. 53, pp. 381–389, Jan 1996.
- [132] 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, no. 5221, pp. 198–201, 1995.
- [133] W. Paul and H. Steinwedel, “Ein neues massenspektrometer ohne magnetfeld,” *RZeitschrift fr Naturforschung*, vol. 8, pp. 448–450, 1953.
- [134] A. L. Migdall, J. V. Prodan, W. D. Phillips, T. H. Bergeman, and H. J. Metcalf, “First observation of magnetically trapped neutral atoms,” *Phys. Rev. Lett.*, vol. 54, pp. 2596–2599, Jun 1985.
- [135] V. G. Minogin, J. A. Richmond, and G. I. Opat, “Time-orbiting-potential quadrupole magnetic trap for cold atoms,” *Phys. Rev. A*, vol. 58, pp. 3138–3145, Oct 1998.
- [136] W. Petrich, M. H. Anderson, J. R. Ensher, and E. A. Cornell, “Stable, tightly confining magnetic trap for evaporative cooling of neutral atoms,” *Phys. Rev. Lett.*, vol. 74, pp. 3352–3355, Apr 1995.
- [137] T. Esslinger, I. Bloch, and T. W. Hänsch, “Bose-Einstein condensation in a quadrupole-IOFFE-configuration trap,” *Phys. Rev. A*, vol. 58, pp. R2664–R2667, Oct 1998.
- [138] D. Xie, D. Wang, W. Gou, W. Bu, and B. Yan, “Fast production of rubidium Bose-Einstein condensate in a dimple trap,” *J. Opt. Soc. Am. B*, vol. 35, pp. 500–503, Mar 2018.
- [139] T. Kinoshita, T. Wenger, and D. S. Weiss, “All-optical Bose-Einstein condensation using a compressible crossed dipole trap,” *Phys. Rev. A*, vol. 71, p. 011602, Jan 2005.
- [140] K. Yamashita, K. Hanasaki, A. Ando, M. Takahama, and T. Kinoshita, “All-optical production of a large Bose-Einstein condensate in a double compressible crossed dipole trap,” *Phys. Rev. A*, vol. 95, p. 013609, Jan 2017.
- [141] J.-F. Clément, J.-P. Brantut, M. Robert-de Saint-Vincent, R. A. Nyman, A. Aspect, T. Bourdel, and P. Bouyer, “All-optical runaway evaporation to Bose-Einstein condensation,” *Phys. Rev. A*, vol. 79, p. 061406, Jun 2009.

- [142] D. Comparat, A. Fioretti, G. Stern, E. Dimova, B. L. Tolra, and P. Pillet, “Optimized production of large Bose-Einstein condensates,” *Phys. Rev. A*, vol. 73, p. 043410, Apr 2006.
- [143] K. Arnold and M. Barrett, “All-optical Bose-Einstein condensation in a 1.06m dipole trap,” *Optics Communications*, vol. 284, no. 13, pp. 3288 – 3291, 2011.
- [144] H. Lewandowski, D. Harber, D. Whitaker, and et al., “Simplified system for creating a Bose-Einstein condensate,” *Journal of Low Temperature Physics*, p. 132: 309. <https://doi.org/10.1023/A:1024800600621>, 2003.
- [145] W. Ketterle, D. Durfee, and D. Stamper-Kurn, “Making, probing and understanding Bose-Einstein condensates,” *arXiv:cond-mat/9904034*, 1999.
- [146] K. E. Wilson, Z. L. Newman, J. D. Lowney, and B. P. Anderson, “In situ imaging of vortices in Bose-Einstein condensates,” *Phys. Rev. A*, vol. 91, p. 023621, Feb 2015.
- [147] R. Meppelink, R. A. Rozendaal, S. B. Koller, J. M. Vogels, and P. van der Straten, “Thermodynamics of Bose-Einstein-condensed clouds using phase-contrast imaging,” *Phys. Rev. A*, vol. 81, p. 053632, May 2010.
- [148] P. B. Wigley, P. J. Everitt, K. S. Hardman, M. R. Hush, C. H. Wei, M. A. Sooriyabandara, P. Manju, J. D. Close, N. P. Robins, and C. C. N. Kuhn, “Non-destructive shadowgraph imaging of ultra-cold atoms,” *Opt. Lett.*, vol. 41, pp. 4795–4798, Oct 2016.
- [149] K. Hueck, N. Luick, L. Sobirey, J. Siegl, T. Lompe, H. Moritz, L. W. Clark, and C. Chin, “Calibrating high intensity absorption imaging of ultracold atoms,” *Opt. Express*, vol. 25, pp. 8670–8679, Apr 2017.
- [150] G. Reinaudi, T. Lahaye, Z. Wang, and D. Guéry-Odelin, “Strong saturation absorption imaging of dense clouds of ultracold atoms,” *Opt. Lett.*, vol. 32, pp. 3143–3145, Nov 2007.
- [151] C. Hirt, S. Claessens, T. Fecher, M. Kuhn, R. Pail, and M. Rexer, “New ultrahigh-resolution picture of Earth’s gravity field,” *Geophysical Research Letters*, vol. 40, no. 16, pp. 4279–4283, 2013.
- [152] T. M. Brzozowski, M. Maczynska, M. Zawada, J. Zachorowski, and W. Gawlik, “Time-of-flight measurement of the temperature of cold atoms for short

- trap-probe beam distances,” *Journal of Optics B: Quantum and Semiclassical Optics*, vol. 4, pp. 62–66, jan 2002.
- [153] A. L. Gaunt, *Degenerate Bose Gases: Tuning interactions and geometry*. PhD thesis, Cavendish laboratory, University of Cambridge, 2014.
- [154] C. J. Pethick and H. Smith, *BoseEinstein Condensation in Dilute Gases*. Cambridge University Press, 2 ed., 2008.
- [155] F. Dalfovo, S. Giorgini, L. P. Pitaevskii, and S. Stringari, “Theory of Bose-Einstein condensation in trapped gases,” *Rev. Mod. Phys.*, vol. 71, pp. 463–512, Apr 1999.