


A New Apparatus for Experiments with Caesium Bose-Einstein Condensates



Craig David Colquhoun

Experimental Quantum Optics and Photonics Group
Department of Physics and SUPA

University of Strathclyde

A thesis presented in the fulfilment of the requirements for
the
degree of
Doctor of Philosophy

2019

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

The copyright of this thesis belongs to the author under the terms of the United Kingdom Copyright Acts as qualified by University of Strathclyde Regulation 3.50. Due acknowledgement must always be made of the use of any material contained in, or derived from, this thesis.

Signed:

Date:

Abstract

This thesis reports on the design and construction of an experimental apparatus capable of generating Bose-Einstein condensates (BECs) of spin-polarised caesium-133 atoms. Caesium condensates offer excellent control over interatomic interactions due to the rich landscape of low-field magnetic Feshbach resonances, which enables the study of quantum gases in attractive and repulsive interaction regimes.

An ultra-high vacuum system, laser systems and magnetic field coils were assembled to trap and cool the atoms from room temperature to temperatures on the order of 1 nK. Absorption imaging was implemented as a means to detect the number and density distribution of the atoms.

Laser cooling and trapping methods are introduced, and the effects of each cooling stage on the gas are demonstrated. The final cooling stage, evaporative cooling, is presented by way of examining the gas after it has undergone each evaporation phase so we observe the onset of Bose-Einstein condensation.

Evaporative cooling produces BECs containing on the order of 2×10^5 atoms, with a condensate fraction of 0.48. We demonstrate that the atom number can be fine-tuned by removal of the most thermal atoms in the trap. We also exhibit our ability to observe expanding condensates in a guiding beam for durations of up to 1 s, and our ability to cause the condensate to implode, by tuning interactions. Measuring the temperature of the gas in sub-nK regimes is currently a challenge when using traditional time-of-flight thermometry. Some modifications to the apparatus have been described, that would permit thermometry using dilute caesium atoms in a different spin state to probe the main gas.

Supervisors: Dr Elmar Haller
Prof. Stefan Kuhr

*Dedicated to the loving memory of my father,
David Colquhoun
(1968 - 2017)*

Acknowledgements

I genuinely can't believe it's been almost four years since I set foot in the lab. They say time flies when you're having fun, and I have to say they're not kidding! Okay, enough, let me thank my supervisors and some of the people that kept me sane throughout my PhD.

First, I must thank Elmar for the opportunity to work on this extremely interesting, stimulating and fruitful experiment. I am grateful for your endless patience, support and benevolence, and for the frequent freedom of choice you afforded us whenever there were forks in the road.

I also need to thank my second supervisor, Stefan Kuhr, for always taking time out of your busy schedule to stop by and check that things were running smoothly. You welcomed me into the Experimental Quantum Optics and Photonics group at Strathclyde, and I could not be more glad that you did.

I would like to thank my lab partner Andrea Di Carli, your perseverance and indomitable spirit accelerated progress in the lab significantly. You were the undeniable workhorse of our lab, and I feel extremely fortunate to have spent my time in the lab with you. Thank you for sharing your wealth of physics knowledge with me where I was lacking, for sharing parts of Italian and Roman culture with me, and for improving my knowledge of Italian profanities by an order of magnitude. You can keep your musical taste to yourself, though. I wish you, Pamela and Riccardo all the very best for your future together.

Gail McConnell, you helped with my application for this PhD in the first place. As if that wasn't enough, you agreed to be my PhD counsellor and kept in regular contact throughout my time here. I am indebted to you for providing so much moral support, advice and assistance.

Griff, I appreciate your open door policy, your dignity, attentiveness, good humour and good sense of humour. You seem to have a sixth sense for knowing just the right thing to say at the right time, please do all you can to keep this. I think this is what kept me going at times.

It was an utter joy to share an office with Rudy, Matt, Manuel and Rémy. Thank you all for your friendship, thank you Rudy, Matt and Manuel for allowing your quirks to constructively interfere with my own for maximum tomfoolery, and thank you Rémy for tolerating us daily.

Thank you to everyone who looked at this thesis at various stages and provided

so much useful feedback, you each helped me improve the document significantly, anyone who reads it will thank you too. In no particular order...

Ilian, I've enjoyed our chats about anything and everything, our games of poker where we both lost far too much money, our joint procrastination attempts, etc. Any future labs you will work in, much like the lab you work in now, will be very lucky to have you.

Carolyn, you are incredibly resilient against my efforts to heckle you. I am continuously impressed with how inquisitive you are about all subjects, it's refreshing to see. You will do great things (once SCAPA turn off their stuff).

Andrew, I suspect I owe you for a lot more than you'd ever admit to. I hope to be a virtuous man like you when I grow up.

Vicki, I don't think I ever had a conversation with you where I didn't end up grinning from ear to ear. Maybe we can share a couple of steins in Germany once all this Brexit shite blows over.

Oliver, you shared *a lot* of your delicious homebrew beers with me. I believe I owe you a keg or two.

Stuart, I always enjoyed how lively you made lunchtime chats. You're wrong about putting brown sauce on a chippy, but your determination to continue doing so is admirable. Bless.

Rachel O, remember how we started off wrestling for lab space, and after I left you started winning awards? You're welcome. You seem to possess a well of positivity, thank you for sharing some with me.

Bruno, the group simply has not been the same since you moved back to France. I thank you for the coffee you left behind, it certainly aided my productivity. CNRS will undoubtedly benefit from your wisdom and experience, as did our group as a whole.

Largest thanks of all go to my Mum and Dad. Without your unconditional love, care and tolerance throughout my tumultuous years, without you nurturing my thirst for knowledge and always guiding me towards better things, I would not have gotten here. I will spend the rest of my days trying to repay this debt with everything I do. I will try to make you both proud.

Alan, you have been a constant reminder that there are more things to life than work. You carefully navigated my work-related grumpiness over the last few years in addition to my intrinsic angst. For bearing with me through thick and thin, you have my sincerest gratitude. Now let's go on a wee holiday or something, eh?

Contents

Abstract	ii
Acknowledgements	iv
Contents	vi
List of Figures	ix
1 Introduction	1
1.1 History of Caesium Cooling	2
1.2 Why Condense Caesium?	4
1.3 Thermometry of Bose-Einstein Condensates	5
1.4 Thesis Overview	6
1.5 Publications Arising From This Work	7
1.6 Author Contributions	7
2 Cooling to a Bose-Einstein Condensate	9
2.1 Doppler Cooling	10
2.2 Magneto-Optical Trap	12
2.3 Degenerate Raman Sideband Cooling	18
2.4 Magnetic Levitation	20
2.5 Dipole Trap	23
2.5.1 Dipole Force	23
2.5.2 Dimple Trap	26
2.5.3 Evaporative Cooling	28
3 Bose-Einstein Condensation	32
3.1 Bose-Einstein Statistics	33
3.2 Phase Transition	34
3.3 Mean Field Approach	36
3.4 Thomas-Fermi Approximation	38
3.4.1 Thermal Density Distribution	39
3.4.2 Bimodal Distribution	40

4	Experimental Setup	43
4.1	Vacuum Setup	44
4.1.1	Caesium Source	44
4.1.2	Glass Cells	44
4.1.3	Pumping	46
4.1.4	Magnetic Shielding	47
4.1.5	Preparation and Assembly	47
4.2	Laser Configuration	48
4.2.1	Saturated Absorption Spectroscopy	51
4.2.2	Polarisation Spectroscopy	53
4.3	Magnetic Field Generation	56
4.4	Magnetic Field Coils	58
4.5	Computer Control	61
4.6	Microcontrollers	63
4.6.1	Optical Shutters	65
4.6.2	DDS controller	70
5	A Bose-Einstein Condensate of Caesium	75
5.1	Diagnostic Procedures	76
5.1.1	Absorption Imaging	76
5.1.2	Temperature Measurement	77
5.1.3	Bimodal Fitting	79
5.1.4	Microwave Spectroscopy	82
5.1.5	Optical Trap Frequency	84
5.1.6	Levitation Gradient Curvature	85
5.2	Route to a Bose-Einstein condensate	87
5.2.1	2D+ MOT	87
5.2.2	3D MOT	89
5.2.3	Degenerate Raman Sideband Cooling	91
5.2.4	Magnetic Levitation	95
5.2.5	Dipole Trap	96
5.2.6	Evaporative Cooling	99
5.3	Bose-Einstein Condensate Characterisation	107
5.3.1	BEC Benchmarks	107
5.3.2	Thermal Atom Removal	110
5.3.3	BEC Collapse	111
5.3.4	Long expansion durations	111
6	Outlook and Conclusions	115
6.1	Future Work	116

CONTENTS

A Publication: A simple laser shutter with protective shielding for beam powers up to 1 W	120
B Publication: Interferometric measurement of micro-g acceleration with levitated atoms	124
C Publication: Excitation modes of bright matter wave solitons	136
D DDS Driver Code	146
Bibliography	160

List of Figures

List of Figures

1.1	Feshbach spectrum for the $F=3$, $m_F=3$ state of caesium	4
2.1	Doppler cooling diagram	11
2.2	Magneto-Optical Trap operation diagram	13
2.3	Illustration of the $\sigma^+ - \sigma^-$ polarisation gradient cooling mechanism	16
2.4	Illustration of Sisyphus cooling mechanisms	17
2.5	Raman sideband cooling energy level scheme for caesium	19
2.6	Levitation gradient horizontal curvature	22
2.7	Plot of real and imaginary parts of atomic polarisability vs detuning	25
2.8	Reservoir beam and dimple beam potentials	27
3.1	Illustration of the Bose-Einstein condensate phase transition	35
3.2	Theory plot of condensate fraction vs temperature	36
3.3	Slice through the centre of a simulated quantum gas	41
3.4	Simulated quantum gas integrated along ‘imaging’ axis	42
4.1	Images of the vacuum chamber	45
4.2	Caesium energy level diagram with useful atomic transitions	49
4.3	Overview of lasers and optics used to prepare light for the experiment	50
4.4	Saturated absorption spectra for caesium cooling and repump lasers	52
4.5	Diagram of the apparatus used for polarisation spectroscopy	54
4.6	Polarisation spectra for caesium cooling and repump lasers	55
4.7	Results of magnetic field simulations for our setup	57
4.8	Construction of the main magnetic field coils	60
4.9	Control computer data flowchart	61
4.10	Plots produced by our lab monitoring setup as they appear online	64
4.11	Photos of our shutter and shutter mount design	66
4.12	Circuit diagram for the shutter driver	67
4.13	Transmission data from tests on a shutter	68
4.14	Histograms generated from tests on a shutter	69
4.15	Pinout diagram of the Arduino and DDS chip setup	72

LIST OF FIGURES

5.1	Demonstration of our absorption imaging process	77
5.2	Example time-of-flight temperature measurement	78
5.3	1D fitting procedure for bimodal density profiles	80
5.4	2D fit of bimodal density profiles	81
5.5	Microwave spectroscopy energy level scheme and example result .	83
5.6	Measurement of trap frequency	85
5.7	Measurement of horizontal curvature due to magnetic levitation .	86
5.8	Computer render of the 2D+ MOT	88
5.9	Plots indicating 3D MOT loading rate and temperature	90
5.10	Diagram of degenerate Raman sideband cooling beam geometry .	92
5.11	Plots demonstrating optimisation of degenerate Raman sideband cooling	93
5.12	Atom number and temperature after Raman sideband cooling . .	94
5.13	Stern-Gerlach separation caused by magnetic levitation	96
5.14	Diagram of beam geometry of reservoir and dimple traps.	97
5.15	Lifetime and temperature measurements of atoms in dipole trap .	98
5.16	Thermalisation in dimple beams and lifetime measurements . . .	100
5.17	The timing diagram for the evaporative cooling sequence	101
5.18	Measurements performed after each stage of evaporative cooling .	103
5.19	Atoms imaged after each evaporative cooling stage, onset of Bose- Einstein condensation	104
5.20	Examination of phase space density at all cooling stages	106
5.21	Demonstration of temperature measurements using a bimodal fit- ting procedure	107
5.22	Lifetime of a Bose-Einstein condensate trapped in dimple beams .	108
5.23	Plots indicating thermalisation of a BEC in different trap depths .	109
5.24	The study of our thermal atom reduction method	110
5.25	Demonstration of a BEC collapse	112
5.26	Demonstration of long BEC expansion durations in a guiding beam	113
6.1	Illustration of a proposed non-destructive thermometry method .	118

Chapter 1

Introduction

When laser cooling was experimentally realised in 1985 by Steven Chu and co-workers [1], it marked the beginning of a new field of physics based around cold atoms and triggered a race to new low temperature regimes, which stimulated the advancement of laser and vacuum technology. Atomic vapours were soon routinely being cooled to ultracold regimes below 10 μK using newly developed laser cooling methods [2]. The importance of these milestones was apparent, as it was widely speculated that laser cooling could serve as a gateway to Bose-Einstein condensation.

After evaporative cooling was developed, rubidium-87 was evaporatively cooled to quantum degeneracy by the group of C. E. Wieman and E. A. Cornell at the National Institute of Standards and Technology in Boulder, Colorado, forming the first Bose-Einstein Condensate (BEC) [3]. This was quickly followed by sodium [4] and lithium [5] BECs. As an indication of the substantial importance of their discoveries, Nobel prizes were awarded to William Phillips [6], Steven Chu [7] and Claude Cohen-Tannoudji [8] in 1997 for the development of laser cooling methods; and Carl Wieman, Wolfgang Ketterle [9] and Eric Cornell [10] in 2001 for the experimental realisation of BECs.

While work on other atomic species was ongoing, a new field was developing

as this interesting state of matter finally allowed the observation of some effects usually described as wave behaviour. Atom lasers [11], solitons [12, 13], vortices [14, 15] and matter-wave interference [16] are some examples of the behaviours that were observed in BECs.

Some years passed by before caesium was condensed, but it was not from lack of trying. All prior BECs were formed in magnetic traps, however the states of caesium that would permit magnetic trapping always led to catastrophic losses during evaporative cooling. Caesium was an elusive BEC goal until late 2002, when Rudolf Grimm's group in Innsbruck evaporatively cooled a spin-polarised sample of atoms in a state that required some unusual techniques to be used [17, 18].

This thesis describes the development and construction of a new apparatus designed to produce caesium BECs and to study quantum systems in optical lattice potentials. The results of each cooling stage that leads to our Bose-Einstein condensation are presented herein, and a preliminary experiment to measure ultra-low temperatures in small systems is discussed.

1.1 History of Caesium Cooling

Caesium has been used as the frequency and time standard since 1967, based on improvements made to the first atomic clock made from caesium in 1955 by L. Essen at the National Physics Laboratory in the UK [19]. Continual improvements to clock stability were made over the years and in 1989, after the advent of laser cooling, M. A. Kasevich and co-workers designed an atomic fountain clock [20]. The idea of atomic fountains was that atoms could be laser cooled and launched into a ballistic trajectory using a light pulse, allowing for a low velocity spread and longer interrogation times than previously possible, further improving the accuracy of atomic clocks.

Research was focused on reaching lower temperature regimes with caesium because it is the heaviest stable alkali metal, so it was thought to be easier to slow caesium atoms than lighter alkali atoms, and because of the expected improvement to the performance of atomic fountain clocks. Due to the fact fountain clock precision was being limited by atom collisions [21, 22], it became a priority to create a Bose-Einstein condensate (BEC) of caesium atoms [23].

The first attempts to create a caesium BEC used states that could be cooled in magnetic traps, and initial low temperature results for atoms in the $|F, m_F\rangle = |4, 4\rangle$ state, seemed promising [24]. After the first BEC was created using rubidium by E. Cornell and co-workers in 1995 [3], a similar evaporative cooling scheme was attempted with caesium. For the $|4, 4\rangle$ state, however, it was found that spin relaxation rates were too high for evaporative cooling to succeed [25–27]. The $|3, -3\rangle$ state, which was previously considered as a viable option [23], became the new focus [28, 29]. The $|3, -3\rangle$ state also resulted in high levels of inelastic collision loss - the best phase space density achieved with this state was two orders of magnitude too low for Bose-Einstein condensation to occur [30, 31], even with suppression of the dipolar relaxation that caused the loss [32].

During this time, many studies were performed on collisions in caesium [27, 33–38], with some suggesting that the $|3, 3\rangle$ ground state could yield more favourable results, however the energetic ground state is always a ‘high-field seeking’ state, and non-magnetic traps were necessary [36, 37]. Success was eventually found with this state in Innsbruck in 2002, when T. Weber and co-workers used an optical trap with magnetic levitation and a magnetic field bias to tune interatomic interaction strengths [39, 40].

1.2 Why Condense Caesium?

With its large mass and large hyperfine splitting, caesium is well suited to optical cooling methods [24, 41]. The separation of the D1 and D2 lines is also ideal for studies with spin-dependent optical lattices, which was demonstrated with rubidium-87 [42, 43], and while the work in this thesis was ongoing, caesium-133 [44].




Figure 1.1: Calculated Feshbach spectrum for the $|3, 3\rangle$ state of caesium. The solid red line shows the s-wave scattering length in Bohr radii with respect to magnetic field, the dashed red lines indicate Feshbach resonances, where the scattering length goes to infinity, and the dotted black line indicates 0 G. Figure obtained from [45].

Another reason for the use of caesium is the rich landscape of Feshbach resonances at low external magnetic fields, as is shown in Fig. 1.1. In particular, a very wide magnetic Feshbach resonance at -11.7 G results in a very large negative scattering length at zero magnetic field with a zero crossing at 17 G. Using this Feshbach resonance, the scattering length of caesium can be reliably tuned from being extremely attractive at external magnetic fields of 0 G, to being repulsive at 17.1 G, with stability provided by the width of the resonance. For a similar

amount of control over interactions in rubidium-87 one would need to resort to magnetic fields of 1007 G [46], or 907 G for sodium-23 [47]. While this tunability is crucial for efficient evaporative cooling in caesium, it also allows for the observation of interesting interaction regimes without the requirement for bulky magnetic field coils and cooling systems.

One more reason to use caesium is simply that there is only a handful of caesium Bose-Einstein condensates worldwide. At the time of writing, there are known to be caesium BECs in Innsbruck [39], Chicago [48], Durham [49], Ljubljana [50] and now Glasgow [51].

1.3 Thermometry of Bose-Einstein Condensates

Temperature measurements of cold atoms are usually performed by measuring the time-of-flight expansion of atom clouds, because their rates of expansion are determined by the temperature of the sample. Once a Bose-Einstein condensate is produced, it gathers in the centre of the cloud, forming a high-density core which does not expand at the same rate as the thermal component due to interactions. Usually a ‘bimodal’ fitting procedure which deconstructs the cloud into thermal and condensed components is used, but works at its best when these components are of roughly equal proportions [52]. For highly pure Bose-Einstein condensates, it is difficult to measure the temperature of the sample because the size and expansion energy of the thermal component are comparable to those of the condensate. It has been reported that accurate time-of-flight thermometry of a Bose gas is limited to temperatures of more than 30% of the critical temperature for Bose-Einstein condensation [53]. Because of this limitation, and the fact that time-of-flight thermometry is a destructive technique, the development of alternative thermometry methods is underway.

Thermometry schemes that use small amounts of ‘probe’ atoms to thermalise

with and measure the temperature of a larger reservoir of atoms have been implemented, both with two species of atom [54], and with two spin states of the same species [53]. Similar non-destructive schemes have been proposed, either by measuring the phase change of quantum dot probes as a result of the measurement [55], or by measuring the position and momentum spread of probes after coming into contact with the system being measured over many measurements [56], but these have not yet been experimentally realised.

The work in this thesis is geared towards performing such temperature measurements for Bose gases with very low thermal fractions. Preliminary results are obtained using time-of-flight measurements, however this method is less accurate for sub-nK temperatures which tend to have very high condensate fractions [57], so we consider some improvements for future work.

1.4 Thesis Overview

The majority of the work detailed in this thesis was performed during the building of our experimental apparatus. The Chapters that follow indicate the theory used during this process, as well as some preliminary results for cooling and experiments with a caesium Bose-Einstein condensate, with the following layout:

- Chapter 2 describes the theory used when cooling caesium atoms from room temperature to a Bose-Einstein Condensate.
- Chapter 3 outlines the theory of Bose-Einstein Condensates in various interaction regimes.
- Chapter 4 introduces the vacuum setup, laser systems and computer control solutions which ensure the smooth running of our experiments (a more in-depth description of our laser arrangement can be found in the thesis of Andrea Di Carli [58]).

- Chapter 5 contains the results obtained throughout the course of cooling and early benchmarks for our Bose-Einstein Condensate.
- Finally, chapter 6 details our next steps for realising quantum temperature probes using our setup.

1.5 Publications Arising From This Work

- C. D. Colquhoun, A. Di Carli, S. Kuhr, and E. Haller *Note: A simple laser shutter with protective shielding for beam powers up to 1 W* [Review of Scientific Instruments](#), **89**, 126102 (2018) - Editor's Pick
- A. Di Carli, C. D. Colquhoun, S. Kuhr, and E. Haller *Interferometric measurement of micro-g acceleration with levitated atoms* [New Journal of Physics](#), **21** 053028 (2019)
- A. Di Carli, C. D. Colquhoun, G. Henderson, S. Flannigan, G. L. Oppo, A. J. Daley, S. Kuhr, and E. Haller *Excitation modes of bright matter wave solitons* [Submitted to Physical Review Letters](#) (2019)

1.6 Author Contributions

The caesium Bose-Einstein condensate apparatus was constructed in collaboration with one other PhD student, Andrea Di Carli, whose project ran concurrently with the author's. As such, many of the aspects, tasks and procedures described in this thesis were shared equally, such as interfacing hardware with our computer control system, constructing the vacuum chamber, assembling and calibrating laser systems, arranging and aligning laser optics, spectroscopies, designing and making magnetic field coils, as well as the day-to-day running of the lab. The author of this thesis contributed more strongly towards the simulation, imaging and

1.6. AUTHOR CONTRIBUTIONS

fitting of results, as exemplified in Sections 3.4 and 5.1.3; towards the design and fabrication of electronic devices, and programming of microcontrollers as shown in Section 4.6; and towards the programming of the control computer and USB and GPIB communications for devices that require it. Andrea focused on placing and aligning most of the optics surrounding the vacuum chamber, optimising the experiment, programming much of the experimental sequence and providing theoretical insight for the experiments we performed. All activities were performed under the supervision and guidance of Dr Elmar Haller.

Chapter 2

Cooling to a Bose-Einstein Condensate

Laser cooling of atoms has become commonplace in cold atom labs around the world and has shaped modern physics research since its first implementation over 30 years ago. The usefulness and importance of laser cooling cannot be overstated, with several new fields of physics being realised as a direct result of its development. This Chapter provides the theoretical background for each stage of the process we execute to cool atoms to degeneracy.

Caesium atoms at room temperature move with a velocity distribution centred around 190 m/s which makes the atoms difficult to trap, manipulate or even observe for more than a few milliseconds. This can be a problem for precision experiments where results are obtained by observing the atoms over long durations. Our aim is to cool the atoms to nanoKelvin temperatures, with velocities of 10^{-3} m/s, so that we may observe quantum physics on a macroscopic scale. We reach our desired temperature using a sequence of cooling stages, utilising magneto-optical traps (MOTs), degenerate Raman sideband cooling [59], and a combination of magnetic levitation and small and large dipole traps to perform all-optical evaporative cooling [39, 40]. The theories behind these methods will

be discussed in this Chapter.

Laser cooling uses the radiation pressure of light to slow down the atoms. The concept of radiation pressure was theorised in the 17th century by Johannes Kepler as the reason a comet's tail is always directed away from the Sun [60]. At the time this theory was criticised by the likes of Sir Isaac Newton, but it was later predicted by James Clerk Maxwell in his theory of light in 1873 [61]. It wasn't until 1933 that radiation pressure was experimentally realised by Frisch [62], using a sodium lamp to deflect a beam of sodium atoms. Upon the invention of the laser, Ashkin demonstrated that it was possible to accelerate macroscopic pellets suspended in water using laser radiation, and indicated that atoms and molecules could be accelerated in the same way [63]. It wasn't long before this effect was proposed by Hänsch and Schawlow to cool neutral atoms [64], and Wineland and co-workers to cool ions [65]. The slowing of an atomic beam using radiation pressure was demonstrated by W. D. Phillips and H. Metcalf in 1982 [66], and 3D cooling and optical molasses were realised in 1985 by S. Chu and co-workers [1], securing radiation pressure as the main effect responsible for the initial cooling of atoms in most ultracold physics labs.

2.1 Doppler Cooling

Radiation pressure can transfer momentum to atoms in an atomic vapour if the light has a frequency close to an atomic transition, which causes the atoms to scatter photons more frequently. The transition rate depends strongly on the frequency of the light with relation to the transition being excited. Due to the optical Doppler effect, the frequency of the light experienced by the atoms differs when they move relative to the light source. If they propagate towards (away from) a resonant light source, the light frequency will appear higher (lower) than the source is producing, resulting in an apparent 'blue-detuning' ('red-detuning')

of the light.

The method of Doppler cooling is based on this frequency shift, and the corresponding change of absorption rates [64]. Each time an atom absorbs a photon, the momentum from that photon is transferred to the atom in the direction the photon was travelling. The atom becomes excited for a short time, before emitting another photon in some random direction, pushing the atom in the opposite direction. The momentum vector from that photon is given by $\mathbf{p} = \hbar\mathbf{k}$ where \hbar is the reduced Planck constant, $\mathbf{k} = 2\pi/\lambda$ is the wave vector of the light and λ is its wavelength.




Figure 2.1: Diagrams of Doppler cooling mechanisms. a) The atoms absorb more photons from the light source towards which they propagate. b) The transfer of momentum that occurs when a photon is absorbed, then emitted.

Consider an atom moving in one dimension towards a monochromatic laser, which has been slightly red-detuned from the atom's resonance. In the reference frame of the atom, the light is closer to its resonance frequency than in the lab reference frame, so it absorbs a photon and receives a momentum 'kick' of $\hbar\mathbf{k}$ from the light source with a higher probability. Although it emits another photon in a random direction and receives a recoil 'kick' in the opposite direction, over several million absorption and emission events, the net effect is that the atom is eventually pushed away from the light source. If we add a second red-detuned laser to counter-propagate towards the first as illustrated in Fig. 2.1, the atom will preferentially absorb photons from the light source towards which it propagates, slowing it in that dimension. Extending this idea to three dimensions so that

we have three orthogonal pairs of counter-propagating beams, we can slow the expansion of a cloud of billions of thermal atoms, cooling it.

The motion of the atom never stops completely - the random directionality of spontaneous emission causes a recoil heating effect comparable to the rate of cooling at lower temperatures. These competing effects eventually reach an equilibrium, resulting in a lower temperature limit for Doppler cooling [67]

$$T_D = \frac{\hbar\Gamma}{2k_B}, \quad (2.1)$$

where Γ is the atomic natural linewidth, and k_B is the Boltzmann constant. The natural linewidth for the $6S_{1/2} \rightarrow 6P_{3/2}$ transition of caesium is $2\pi \times 5.22$ MHz, which corresponds to a Doppler temperature of $125 \mu\text{K}$ [68].

As the atoms near the Doppler temperature, the laser beams provide a viscous damping force in the opposite direction to which the atoms propagate, forming an ‘optical molasses’. Optical molasses can sometimes cool the atoms to below the Doppler temperature, and although the atoms are confined for finite periods of time by the molasses, they are not trapped - trapping is achieved by combining optical molasses with the Zeeman effect caused by the presence of external magnetic fields.

2.2 Magneto-Optical Trap

In 1982, William Phillips and Harold Metcalf devised a method of slowing and bunching beams of atoms with a position dependent force using laser beams and a magnetic field gradient [66]. This process is based on the idea that as the atoms lose their velocity and their Doppler shift is reduced, their transition energies are lowered by virtue of the Zeeman effect, so that they can still be excited by the red-detuned laser. E. L. Raab and co-workers extended this principle to three dimensions, enabling the simultaneous cooling and trapping of sodium

2.2. MAGNETO-OPTICAL TRAP

atoms [69]. Where Doppler cooling exerts a velocity-dependent force on the atoms, the combination of lasers and magnetic field gradients produces a position-dependent force, resulting in a magneto-optical trap (MOT).




Figure 2.2: Schematic of a magneto-optical trap in 1D with only three excited m_F states shown, for simplicity. The magnetic field varies with position, so do the $m_F = \pm 1$ levels. The $m_F = -1$ level is lowered to become resonant with the circularly polarised beams, which cause σ^- transitions on the sides they enter. Because the direction of the magnetic field changes in the centre position of the MOT, these same beams will cause σ^+ transitions after crossing the 0 position point, but because the $m_F = 1$ level has been shifted further from resonance with the light, these transitions are far less likely.

A MOT works in two or three dimensions with the correct configuration of quadrupole magnetic field and circular polarisations of red-detuned laser light, as shown in Fig. 2.2. In our case, the cooling and trapping lasers are red-detuned from the $F = 4 \rightarrow F' = 5$ transition of caesium, because if the atoms are excited to the $F' = 5$ state, they can only decay back down to the $F = 4$ state, allowing for continuous cycling of this transition. In the case of an off-resonant $F = 4 \rightarrow$

2.2. MAGNETO-OPTICAL TRAP

$F' = 4$ transition occurring, the atoms could decay into the $F = 3$ ground state, which because it is 9.2 GHz detuned from the $F = 4$ state, is dark to the cooling transition. Repump light is necessary to excite these atoms from $F = 3 \rightarrow F' = 4$, so these atoms can decay back to the $F = 4$ ground state, rejoining the cooling cycle.

At the centre of the trap, the atoms experience no magnetic field, so they absorb an equal number of photons from each beam. As they drift from the centre position and experience magnetic fields, the energies of the m_F levels are Zeeman shifted by

$$\Delta E = \hbar\Delta\omega = m_F g_F \mu_B \frac{\partial B}{\partial z} z, \quad (2.2)$$

where μ_B is the Bohr magneton, $\partial B/\partial z$ is the magnetic field gradient in the direction z and g_F is the Landé g-factor associated with the F state. The direction of the magnetic field defines whether σ^+ or σ^- transitions are caused by circularly polarised light of the same handedness. In the 1D case in Fig. 2.2, the magnetic field is always positive, however the direction of the magnetic field, and hence the quantisation axis of the atoms, changes. Once the atoms have travelled far from the centre position, towards positions of positive magnetic field, their excited negative m_F states are shifted to lower energies, so the atoms will preferentially absorb light that will induce σ^- transitions. The polarisation of the light is selected so that it will induce σ^- transitions on the side that it enters, so atoms that absorb it will be sent to the centre position. After it passes the centre position, this same light will cause σ^+ transitions because of the change in direction of the quantisation axis of the atoms but, because the $m_F = 1$ levels have been shifted further away from resonance, this is unlikely to happen.

This trapping method is often used in 2D, creating a cigar-shaped trap that can be used as a cold atom source for an experimental chamber [70]; or in 3D as a method to cool and trap atoms so that they can be cooled further for loading weaker traps.

An atom cloud can also be compressed in the MOT by increasing the magnetic field gradient and the detuning of the cooling beam, this usually results in a higher atom density which can be convenient for performing the next step in the cooling sequence with minimal atom losses, at the expense of heating the sample.

With the complex magnetic field and laser polarisation arrangement that results in a 3D MOT, two sub-Doppler cooling mechanisms often occur and are difficult to suppress - polarisation gradient cooling and Sisyphus cooling. In the following Sections, these mechanisms are considered in their simplest case, formed in 1D by two overlapping counter-propagating laser beams in either circular polarisations of opposite handedness, or orthogonal linear polarisations. In a six-beam MOT, a complicated polarisation arrangement is formed in the beam overlap volume, which causes both cooling mechanisms to occur simultaneously.

Polarisation Gradient Cooling

In the 1D picture, if we keep the same beam configuration as in Fig. 2.2 and remove the magnetic field gradient, the two circularly polarised beams interfere to form a standing wave of linearly polarised light, which precesses in a spiral along the direction of light propagation [2]. An atom that propagates towards one of the light sources observes its light blue-shifted towards resonance, and the other light source red-shifted away from resonance. As light with a linear polarisation can be said to consist of equal amounts of right-handed and left-handed circular polarisations, the atom then sees more of one circular polarisation than another, causing an asymmetry in the likelihood of absorbing light from the source towards which it propagates over the other. Figure 2.3 illustrates this effect, with the relative transition strengths for an atom in stationary and moving cases. If the atom moves towards the σ^+ beam, its energy levels are shifted further apart with increasing m_F states. The amount of energy lost from each absorption event is equal to the shift in energy from the m_F to the m_F+1 state, which allows for

2.2. MAGNETO-OPTICAL TRAP

lower temperatures than the Doppler limit calculated from Eq. 2.1.




Figure 2.3: Illustration of polarisation gradient cooling in the $\sigma^+ - \sigma^-$ configuration. a) The combination of counter-propagating right handed and left handed circularly polarised beams form a rotating linear polarisation. b) A stationary caesium atom in this beam configuration experiences equal amounts of both circular polarisations, which are now denoted as σ^+ and σ^- for the transitions they induce. The AC stark shift for the different Zeeman states is proportional to the transition strengths indicated, which have been normalised to integer values. The atom is most probably in the lowest-lying $m_F = 0$ state. c) The atom is now moving towards the σ^+ polarised light source, which is Doppler shifted towards resonance whereas the σ^- is shifted further from resonance. As a result, the transition strengths are altered to favour the higher $m_F \rightarrow m_F'$ transitions, and so the light shift affects these states more. This difference in energy from m_F to $m_F + 1$ is the amount of kinetic energy in the atom that is converted to potential energy as the atom is pumped towards $m_F = 4$ (such as the decay indicated by the dotted line). b) and c) have been adapted from a figure in [71].

Sisyphus Cooling

If we now consider the MOT beams in 3D without a magnetic field gradient, the configuration of polarisations in the beam overlap volume is not so trivial. As well as the aforementioned $\sigma^+ - \sigma^-$ configuration of polarisation gradient cooling, the atoms are also subjected to the lin- \perp -lin, or ‘Sisyphus cooling’, configuration. Sisyphus cooling usually occurs when orthogonal linearly-polarised beams are overlapped with one another. The beams interfere to form a pattern of linear and circular polarisations with varying degrees of ellipticity, which shifts the ground

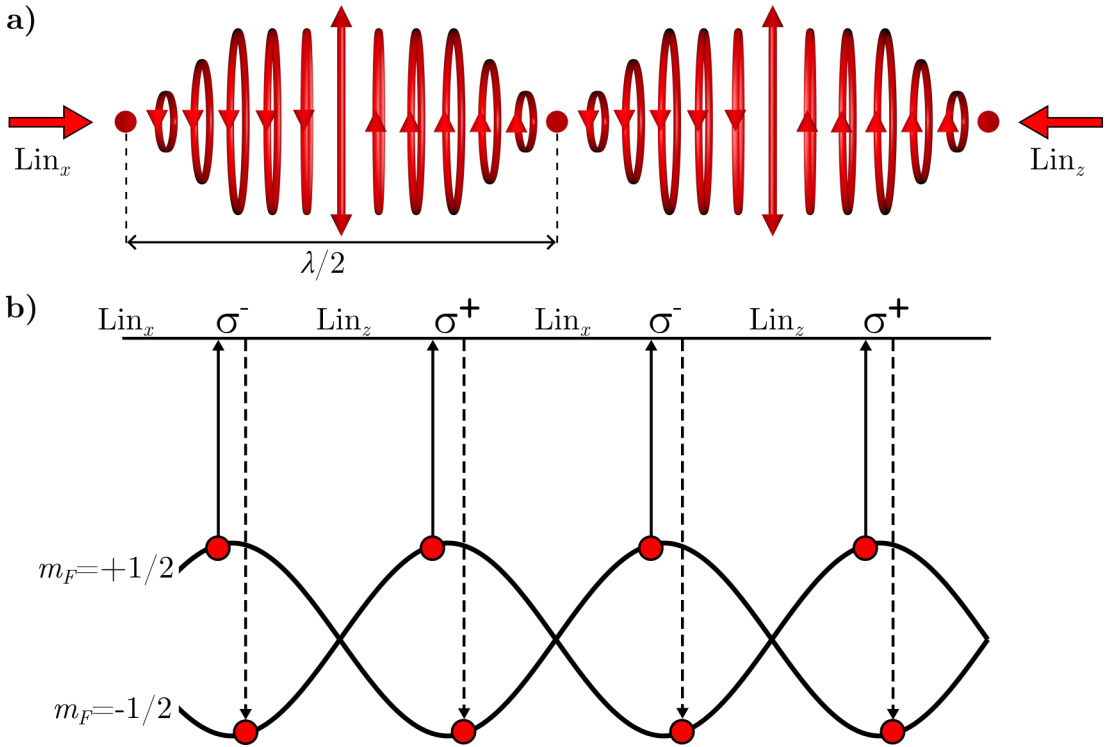


Figure 2.4: Illustration of polarisation gradient cooling in the lin- \perp -lin configuration. a) The polarisation configuration produced when two counter-propagating, orthogonally linearly polarised beams are overlapped. The polarisation of the light deforms through space, with Lin_x , right-handed circular, Lin_z and left-handed circular polarisations all occurring periodically, with varying degrees of elliptical polarisations occurring intermediately. b) The potential shift of the ground m_F states of an atom as it propagates through this polarisation configuration, and the absorption-emission cycle that results in Sisyphus cooling.

state m_F levels sinusoidally in position space, as illustrated in Fig. 2.4. An atom experiencing a positive potential shift in an m_F state it occupies can be said to be climbing a ‘potential hill’. As it climbs the hill, it gains potential energy bringing it closer to resonance with the light, making it more likely to absorb a photon. When the atom absorbs a photon, it spontaneously emits, either decaying back down into its original m_F state, losing no energy, or decaying into the other lower m_F state, losing an energy of $E = \hbar(\omega_{\text{abs}} - \omega_{\text{sp}})$, where ω_{abs} and ω_{sp} refer to the frequency of the light absorbed by the atom and the frequency of the light emitted by the atom, respectively. If the latter occurs, the atom is ready to climb the potential hill once more.

2.3 Degenerate Raman Sideband Cooling

After the atoms have been cooled and trapped in the MOT, they are in a mixture of hyperfine and Zeeman states. This is problematic for magnetically levitating them because different Zeeman states are affected differently for the same magnetic field gradient. We use degenerate Raman sideband cooling as the next cooling stage because it cools the atoms to their lowest motional level, with the added advantage of polarising them in a high-field seeking state [18] so there is no need for an additional optical pumping stage. In addition, Raman sideband cooling is well suited for the all-optical BEC production because compared to other methods, it cools atoms faster, to similar temperatures [59, 72] with a lower level of loss [73]. A diagram of the degenerate Raman sideband cooling process for caesium is shown in Fig. 2.5. In the case of caesium, an optical lattice is formed using lasers that are resonant with the $F = 4 \rightarrow F' = 4$ transition, which pumps atoms to the energetically lower $F = 3$ state, so the lattice beams are then 9.2 GHz red-detuned with respect to all the possible transitions. The atoms are then attracted to the points of highest intensity in the lattice, where if they




Figure 2.5: A diagram of the Raman sideband cooling process. The green double-ended arrows indicate Raman transitions between Zeeman states, either increasing or decreasing the vibrational state. The red arrows show the transition to an AC-Stark shifted excited state, and the subsequent decay back to the ground vibrational level in the lowest energy m_F state, which is a dark state.

have low energies, they are said to be in the Lamb-Dicke regime. In this regime, the atoms are so tightly bound that photon recoils are suppressed, allowing the cooling of atoms to their motional ground state. The atoms experience harmonic potentials that conserve their vibrational levels ν , and they are mostly unaffected by recoils due to absorption and emission. The lattice light induces two-photon Raman transitions, which couple the degenerate vibrational levels in different Zeeman states. Using a magnetic field to lift the degeneracy of the vibrational levels by $\hbar\omega$, the level ν in state m_F becomes degenerate with the $\nu \pm 1$ levels in the $m_F \pm 1$ states.

A polarising beam is blue-detuned from the closed transition $F=3 \rightarrow F'=2$, and is mostly σ^+ polarised, with a small π component. When the atoms become excited and decay back to this ground state, they increase in m_F and conserve vibrational states because they are in the Lamb-Dicke regime. Eventually, we end up in the case of the atoms alternating between the $|3, 3, \nu=1\rangle$ and $|3, 2, \nu=0\rangle$ levels, the former being a dark state to the polarising beam, and the latter being dark to only the σ^+ component of the light. The role of the π component of the beam is to finally pump the atoms into the dark $|3, 3, \nu=0\rangle$ state. The atoms are now in the lowest motional level, spin polarised in the high-field seeking state we desire so they can be magnetically levitated against gravity. The theoretical recoil cooling limit is given by

$$T_R = \frac{\hbar^2 k^2}{mk_B}, \quad (2.3)$$

which is 198 nK for caesium having absorbed and emitted a single 852 nm photon [68]. The first attempts to perform Raman sideband cooling on caesium yielded temperatures of 280 nK [72] and 250 ± 30 nK [59], approximately 1.5 recoil temperatures. The lowest temperature limit in Raman sideband cooling is set by the rate of photon reabsorption, which occurs in dense clouds.

2.4 Magnetic Levitation

Once the atoms have been cooled so that they can be contained in much weaker traps than before, their weight becomes a limiting factor. Caesium is the heaviest stable alkali and any trap containing caesium atoms needs to counteract the effects of gravity. In optical dipole traps, this would require the use of an impractical amount of laser power. Instead, we magnetically levitate the atoms against gravity before introducing the dipole trap so that it takes less optical power to trap the atoms. This additionally eliminates the effects of gravitational

sag during the evaporative cooling process, and ensures that all the atoms are of the same spin [39, 40, 74]. Magnetic levitation is also advantageous for observing atoms over longer durations than time-of-flight measurements would allow, and makes it convenient for a single region of interest to be selected for several images across multiple measurements. If our cloud is in a mixture of states and we apply a magnetic field gradient in the vertical direction (z axis) before releasing them from a trap, each state will experience a different gravitational acceleration, and the states separate spatially, because each m_F level experiences a different potential. If we levitate one of these states perfectly, this will result in a large loss of atoms as we lose all other states, because each m_F state experiences a different potential U_{mag} due to the magnetic field in the z axis, B_z , according to

$$U_{\text{mag}} = \mu_B g_F m_F B_z, \quad (2.4)$$

where μ_B is the Bohr magneton. We need the gradient of this potential to counteract the gravitational potential energy $U_{\text{grav}} = mgz$, so we obtain the equation for magnetic field gradient

$$\frac{\partial B_z}{\partial z} = \frac{mg}{\mu_B g_F m_F}, \quad (2.5)$$

which for the $|3, 3\rangle$ state of caesium is 31.14 G/cm. However, this perfect levitation in the vertical direction z introduces an ‘anti-trapping’ potential in the horizontal xy plane (see Fig. 2.6). In free-space, this causes an expansion of the levitated cloud in the horizontal plane, and in a dipole trap this reduces the horizontal trap depth [39, 40]. This anti-trapping potential takes the form of a parabola, with an outward radial force originating from its turning point at $(x, y) = (0, 0)$

$$F = m\alpha^2 r, \quad (2.6)$$

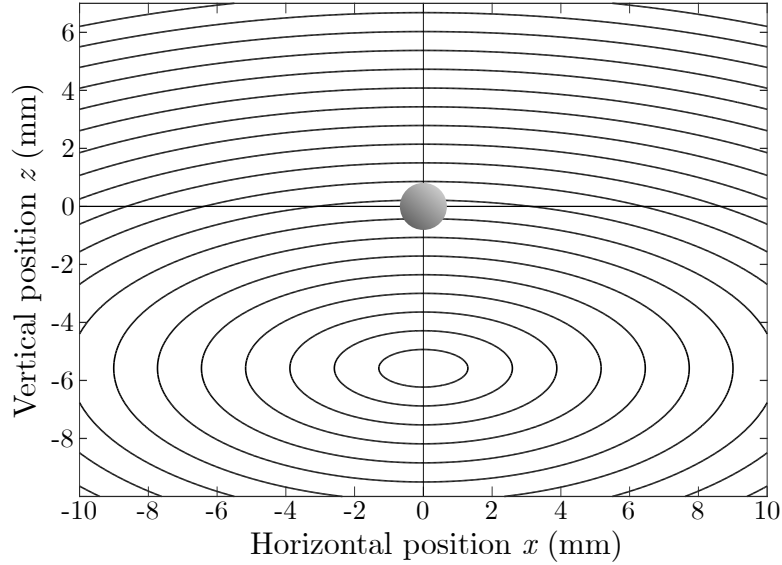


Figure 2.6: The curvature of the levitation field flattens off at higher magnetic field values, the atoms are located at the bias field value $B_0 = 17$ G. The equipotential surfaces all centre on zero magnetic field, each subsequent surface indicating a difference of 2 G.

where $r = \sqrt{x^2 + y^2}$ is the distance from $(0, 0)$, B_0 is the magnetic field at $(0, 0)$ and α is a parameter defining the curvature of the parabola

$$\alpha = g \sqrt{\frac{m}{3\mu_B B_0}}. \quad (2.7)$$

This α parameter causes a position-dependent acceleration, which results in the following equations of motion with time-dependent position $r(t)$ and velocity $v_r(t)$ [75].

$$r(t) = r(0) \cosh(\alpha t) + \alpha^{-1} v_r(0) \sinh(\alpha t) \quad (2.8)$$

$$v_r(t) = v_r(0) \cosh(\alpha t) + \alpha r(0) \sinh(\alpha t) \quad (2.9)$$

$$z(t) = v_z(0)t + z(0). \quad (2.10)$$

As the cloud spreads out, the magnetic field gradient along the z axis is reduced, causing the atoms to fall. According to Eq. 2.7, there is a $1/\sqrt{B_0}$ dependence on

the curvature of the levitation potential, so it becomes practical to introduce a homogeneous bias field to prevent atoms sliding downwards over short levitation durations.

2.5 Dipole Trap

After Raman cooling, the atoms are in a temperature regime where an additional dissipative process would only serve to heat them, so we now implement a method of trapping the atoms using a conservative force. By using far-detuned, off-resonant laser beams, we suppress absorption and emission events and a different trapping force begins to dominate. Unlike the radiation pressure force which is parallel to beam direction, the dipole force is strongest where the intensity gradient is steepest, which is usually in the transverse axis of the beam. Because the dipole force is conservative, a dipole trap can be used to trap atoms, but additional processes are required to cool them to degeneracy.

2.5.1 Dipole Force

We can approximately describe this force as a result of the induced electric dipole of an atom oscillating in an electromagnetic field of frequency ω [76]. In the case of a neutral atom in a oscillating electric field $\mathbf{E}(t) = \mathbf{E} \cos(\omega t)$, the oscillation of the dipole moment $\mathbf{p}(t)$ is approximated to be a damped harmonic oscillator

$$\frac{d^2 p(t)}{dt^2} + \Gamma_\omega \frac{dp(t)}{dt} + \omega_0^2 p(t) = \frac{e}{m_e} E_0 \cos(\omega t) \quad (2.11)$$

where e and m_e are the charge and mass of an electron, and we have made scalar approximations $E(t) = |\mathbf{E}(t)|$ and $p(t) = |\mathbf{p}(t)|$. The rate of radiative energy loss,

2.5. DIPOLE TRAP

Γ_ω , serves as the damping rate of this oscillator, and is given by the equation

$$\Gamma_\omega = \frac{e^2\omega^2}{6\pi\epsilon_0 m_e c^3}, \quad (2.12)$$

where ϵ_0 is the permittivity of free-space and c is the speed of light. The induced dipole moment, p_0 , is related to the \mathbf{E} field amplitude by the complex polarisability of the electron,

$$\alpha_p = \frac{e^2}{m_e} \frac{1}{\omega_0^2 - \omega^2 - i\omega\Gamma_\omega}, \quad (2.13)$$

where ω_0 is the resonance frequency of the oscillator. The real part of the polarisability is used to calculate the interaction potential of the dipole moment U_{dip} ,

$$U_{\text{dip}} = -\frac{1}{2\epsilon_0 c} \text{Re}(\alpha_p) I. \quad (2.14)$$

Since I is the intensity of the driving field, and the intensity of a Gaussian laser beam is inhomogeneous in position-space, the dipole force is obtained by taking the gradient of this potential

$$\mathbf{F}_{\text{dip}}(\mathbf{r}) = \frac{1}{2\epsilon_0 c} \text{Re}(\alpha_p) \nabla I(\mathbf{r}). \quad (2.15)$$

The nature of the real part of the polarisability is illustrated in Fig. 2.7, where it can be seen that $\text{Re}(\alpha_p)$ has a dispersive response centered around ω_0 ($\Delta = \omega - \omega_0 = 0$). For a blue-detuned driving field ($\Delta > 0$), $\text{Re}(\alpha_p)$ is positive, and for a red-detuned field it is negative. This indicates that for a red-detuned field the dipole potential is negative and therefore the dipole force attracts the oscillator towards intensity maxima; and similarly that a blue-detuned field repels the oscillator from the intensity maxima. The radiative scattering rate which arises from the imaginary part of the polarisability is calculated by considering the power




Figure 2.7: A graph of the real and imaginary parts of the polarisability vs detuning in atomic linewidths near an atomic resonance at frequency ω_0 , illustrating the nature of the dissipative and conservative forces experienced by atoms in an electromagnetic field with respect to the detuning of the field.

absorbed by the oscillator at ω , as a stream of photons with energy $\hbar\omega$

$$P_{\text{abs}} = \frac{\omega}{\epsilon_0 c} \text{Im}(\alpha_p) I, \quad (2.16)$$

$$\Gamma_{\text{sc}}(\mathbf{r}) = \frac{P_{\text{abs}}}{\hbar\omega} = \frac{1}{\hbar\epsilon_0 c} \text{Im}(\alpha_p) I(\mathbf{r}). \quad (2.17)$$

By substituting $\text{Re}(\alpha_p)$ and $\text{Im}(\alpha_p)$ into Eqs 2.14 and 2.17 respectively, we get the equations

$$U_{\text{dip}}(\mathbf{r}) = -\frac{3\pi c^2}{2\omega_0^3} \left(\frac{\Gamma}{\omega_0 - \omega} + \frac{\Gamma}{\omega_0 + \omega} \right) I(\mathbf{r}), \quad (2.18)$$

$$\Gamma_{\text{sc}}(\mathbf{r}) = \frac{3\pi c^2}{2\hbar\omega_0^3} \left(\frac{\omega}{\omega_0} \right)^3 \left(\frac{\Gamma}{\omega_0 - \omega} + \frac{\Gamma}{\omega_0 + \omega} \right)^2 I(\mathbf{r}) \quad (2.19)$$

which are valid for any value of ω . Often, by making the rotating wave approximation and the assumption that $\omega/\omega_0 \approx 1$, these equations can be simplified to

$$U_{\text{dip}}(\mathbf{r}) = -\frac{3\pi c^2}{2\omega_0^3} \frac{\Gamma}{\Delta} I(\mathbf{r}) \quad (2.20)$$

$$\Gamma_{\text{sc}}(\mathbf{r}) = \frac{3\pi c^2}{2\hbar\omega_0^3} \left(\frac{\Gamma}{\Delta}\right)^2 I(\mathbf{r}). \quad (2.21)$$

This simplification emphasises an outcome that is crucial for the possibility of dipole traps. We can see that the dipole potential scales as I/Δ , and the scattering rate scales as I/Δ^2 , which indicates that for large values of detuning the atoms can be trapped with a low probability of being heated due to scattering events. We can calculate the radial and axial trap frequencies, ω_{rad} and ω_{ax} , of a single dipole trap beam using the harmonic approximation [77]

$$\omega_{\text{rad}} = \sqrt{\frac{4U_{\text{dip}}}{mr^2}} = \sqrt{\frac{12Pc^2}{mr^4\omega_0^3} \left(\frac{\Gamma}{\Delta}\right)} \quad (2.22)$$

$$\omega_{\text{ax}} = \sqrt{\frac{2U_{\text{dip}}}{mz_r^2}} = \sqrt{\frac{24Pc^4}{mr^6\omega_0^5} \left(\frac{\Gamma}{\Delta}\right)} \quad (2.23)$$

where r is the beam waist radius, $z_r = \pi r^2/\lambda$ is the Rayleigh length and P is the optical power. Using this notation to identify the trap frequencies, the geometric mean trap frequency $\bar{\omega}$ is calculated by $\bar{\omega} = \sqrt[3]{\omega_{\text{rad}}^2\omega_{\text{ax}}}$.

2.5.2 Dimple Trap

There are many different dipole trap configurations [76] - single beam traps are used to trap atoms in the focus of a laser beam, which only offers strong confinement in the radial direction. Confinement in the axial direction is much weaker, as the Rayleigh length of a Gaussian beam is typically much larger than its waist. Crossed dipole traps, where the foci of two beams overlap, provide confinement




Figure 2.8: Graphs of the potential experienced by an atom in a) a Gaussian beam and b) the same Gaussian beam with a more tightly-focused Gaussian beam in its centre. The shape of this potential is what gives the ‘dimple trap’ its name.

in three dimensions with higher atom densities. This is because the increase in intensity in the beam overlap region provides a deeper trapping potential, and tighter axial confinement in each beam. Note that simply increasing the trap depth would not improve confinement substantially (with the $\sqrt{U_{\text{dip}}}$ dependency seen in Eqs 2.22 and 2.23), but it is the deformation of the potential provided by the second beam that increases the capacity for cooling.

This effect was investigated by Pinkse and co-workers [78], who showed that by introducing a much tighter beam in the centre of the dipole trap, the phase space density of the trapped gas can be increased without affecting its temperature. This trap configuration was dubbed a ‘dimple’ trap, because of the shape of its potential (a comparison between dipole and dimple potentials can be seen in Fig. 2.8). Atoms collide in the original, wider part of the trap and occasionally find themselves in the dimple section, where they are confined more tightly. In this way, the wider beam behaves as a reservoir that loads the dimple through elastic collisions, increasing the density of the gas. This results in an increase of phase space density, with minimal effect on the temperature of the gas since the dimple is in thermal contact with the reservoir. By increasing the depth of the dimple potential, Stamper-Kurn and co-workers [79] demonstrated the ability to

form a BEC reversibly without evaporative cooling, measuring a 50-fold increase in the phase space density of their sample. Because a high phase space density is required to form a BEC, this dimple potential is critical for the beginning of an optical evaporation process.

2.5.3 Evaporative Cooling

Because a dipole trap is a conservative potential, additional methods need to be implemented for cooling atoms to degeneracy while they are in the trap. These ‘evaporative cooling’ methods result in the release of the most energetic atoms from the trap, cutting off the high-end tail of the Maxwellian velocity distribution, and cooling the cloud as a whole. The most common method of evaporation involves using rf pulses on atoms contained in a magnetic trap to transfer them from a magnetically trapped state to a non-trapped state, which is not suitable for cooling caesium [18, 39]. Instead, we take an optical approach to evaporative cooling which takes advantage of the dimple beams and magnetic levitation.

By slowly ramping down the intensities of the ‘reservoir’ beams and keeping the ‘dimple’ beam intensities constant, we deform the potential in a way that allows the hottest atoms to escape the trap. These atoms carry more than the average kinetic energy of the gas with them, cooling it. In order for thermalisation to occur, and for the dimple trap to be loaded with atoms, the atoms need to undergo elastic collisions in the reservoir.

The efficiency of the evaporative cooling process depends on the ratio between elastic and inelastic collisions, the latter of which causes losses that do not contribute to the desired cooling effect. Three different effects constitute inelastic collisions: collisions of the atoms with the background gas, and two-body and three-body losses. Background gas collisions are unlikely to cause significant loss in an ultra-high vacuum environment, and because the atoms are spin-polarised into the lowest energy state, spin-exchange collisions which result in two-body

losses are also negligible. The only effect which can cause significant loss in this system is three-body recombination, caused when two atoms form a molecule while interacting with a third. Each three-body recombination event can result in the loss of two or three atoms, and recombination heating [80].

The rate of inelastic collisions is given by the differential equation

$$\dot{n}(\mathbf{r}, t) = \frac{-n(\mathbf{r}, t)}{\tau} - L_2 n^2(\mathbf{r}, t) - L_3 n^3(\mathbf{r}, t), \quad (2.24)$$

where $n(\mathbf{r}, t)$ is the atom density in the trap, τ is the $1/e$ trap lifetime, and L_2 and L_3 are the coefficients of two-body and three-body losses, respectively. The exponent of the density in each loss term scales with the number of atoms participating in the collision, for instance two-body losses require two atoms to converge on the same position, so the probability of this occurring scales with n^2 . Integrating Eq. 2.24 over the atom cloud gives

$$\Gamma_{\text{in}} = \frac{dN}{dt} = -\frac{N}{\tau} - L_2 \langle n \rangle N - L_3 \langle n^2 \rangle N, \quad (2.25)$$

where N is the atom number, $\langle n \rangle$ is the mean atom density in the trap given by

$$\langle n \rangle = \frac{1}{N} \int n^2(\mathbf{r}) d^3\mathbf{r}, \quad (2.26)$$

and $\langle n^2 \rangle$ is given by

$$\langle n^2 \rangle = \frac{1}{N} \int n^3(\mathbf{r}) d^3\mathbf{r}, \quad (2.27)$$

which for a Gaussian distribution can be written as $\langle n^2 \rangle = 8/\sqrt{27} \langle n \rangle^2$ [81]. The three-body loss rate coefficient is given by

$$L_3 = n_i C \frac{\hbar}{m} a^4, \quad (2.28)$$

where n_i is the number of atoms lost per three-body recombination event and

2.5. DIPOLE TRAP

C is a dimensionless parameter. $L_3 \langle n^2 \rangle$ can be expressed in terms of N and T , giving a three-body loss rate of

$$\frac{dN}{dt} = -\frac{L_3}{\sqrt{27}} \left(\frac{m\bar{\omega}^2 N}{2\pi k_B T} \right)^3, \quad (2.29)$$

so the only direct way three-body loss can be reduced is by keeping the scattering length low.

The rate of elastic collisions is calculated by

$$\Gamma_{\text{el}} = \langle n \rangle \sigma \langle v_{\text{rel}} \rangle, \quad (2.30)$$

where $\langle v_{\text{rel}} \rangle$ is the average relative velocity of the atoms in the trap

$$\langle v_{\text{rel}} \rangle = \sqrt{\frac{16k_B T}{\pi m}} \quad (2.31)$$

and σ is the collision cross-section given by

$$\sigma = \frac{8\pi a^2}{1 + k^2 a^2}, \quad (2.32)$$

where k is the de Broglie wave vector given by

$$k = \sqrt{\frac{16mk_B T}{\pi \hbar^2}}. \quad (2.33)$$

Finally the evaporation efficiency, γ , which describes the fractional gain of phase space density for a fractional atom loss, is defined as [82]

$$\gamma = \frac{\ln(PSD'/PSD)}{\ln(N'/N)} \quad (2.34)$$

where N and PSD are the initial atom number and phase space density, and N' and PSD' are the final atom number and phase space density for evaporative

2.5. DIPOLE TRAP

cooling. The phase space density gives a measure of how close the gas is to Bose-Einstein condensation, hence γ is an important parameter during evaporative cooling - it measures what ratio of the 'lost' atoms contribute to the desired cooling effect.

Chapter 3

Bose-Einstein Condensation

Bose-Einstein Condensates (BECs) are a state of matter of bosonic particles which occurs at temperatures close to absolute zero. Due to the quantum statistics of bosons, a macroscopic number of particles accumulates in the energetically lowest quantum state at low temperatures. The quantum statistical treatment for bosons was first presented for photons by Bose in 1924 [83], which Einstein extended to bosonic matter particles, predicting Bose-Einstein Condensation in an ideal gas of bosons in 1925 [84]. The first BEC wasn't experimentally realised until 70 years after Bose and Einstein's predictions, when C. Wieman and co-workers produced a BEC made of 2×10^4 rubidium-87 atoms in July 1995 [3].

The properties of this new state of matter opened up a new world of research possibilities. The phase coherence and wave-like nature of BECs allows for experiments in atom interferometry [16] and atom lasers [11]; collective excitations allow for the observations of solitons [12], vortices [14] and the speed of sound in the condensate [52]; and the tunability of interactions in BECs enables studies in condensate collapse [85, 86], as well as molecule formation [75, 87] and loss mechanisms [80, 88].

In addition, work on creating BECs from other species of alkali atoms continued, and before the end of 1995 there were already lithium [5] and sodium [4]

BECs. However, the creation of caesium BECs remained elusive due to its complex scattering properties, and it was not until 2003 that caesium was successfully condensed [39], but it required some deviation from the standard methods [17].

3.1 Bose-Einstein Statistics

The Bose distribution function [52],

$$N_\nu(T) = \frac{1}{e^{(\epsilon_\nu - \mu)/(k_B T)} - 1} \quad (3.1)$$

gives the number of atoms N_ν in any given state ν , where ϵ_ν is the energy of that state, μ is the chemical potential, k_B is the Boltzmann constant and T is the temperature of the gas. At high temperatures, the total atom number is

$$N = \sum_\nu N_\nu(T), \quad (3.2)$$

where the atoms are distributed across many states, and the occupation of each state $N_\nu \ll 1$. The occupation of any given state is limited such that $\mu < \epsilon_\nu$, however this breaks down for low temperatures which causes occupation numbers to diverge. Instead, we split Eq. 3.2 into two parts, allowing us to distinguish between atom numbers into these two temperature regimes so that

$$N = N_0 + N_T = N_0 + \sum_\nu N_\nu(T) \quad (3.3)$$

where N_0 is the number of atoms that have condensed, and N_T is the number of thermal atoms.

3.2 Phase Transition

The transition from a cloud of thermal atoms to a BEC depends on the phase space density of the atoms. The phase space density is given by

$$PSD = \hat{n}\lambda_{\text{dB}}^3, \quad (3.4)$$

where \hat{n} is the peak atom density, and λ_{dB} is the thermal de Broglie wavelength given by [52]

$$\lambda_{\text{dB}} = \sqrt{\frac{2\pi\hbar^2}{mk_B T}}, \quad (3.5)$$

where \hbar is the reduced Planck constant, and m is the mass of an atom. The BEC transition occurs when the thermal de Broglie wavelength becomes comparable to the spacing between the atoms, so that the phase space density is on the order of unity. This increase in phase space density is achieved through reducing the temperature and increasing the peak atom density of the cloud as illustrated in Fig. 3.1.

In a three dimensional (3D) harmonic trap with trapping frequencies $\omega_{x,y,z}$ and potential

$$U(\mathbf{r}) = \frac{m}{2}(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2), \quad (3.6)$$

the peak density \hat{n} is given by the equation

$$\hat{n} = N\bar{\omega}^3 \left(\frac{m}{2\pi k_B T} \right)^{3/2}, \quad (3.7)$$

where $\bar{\omega} = \sqrt[3]{\omega_x\omega_y\omega_z}$ is the geometric mean of trap frequencies. The phase space density of a thermal gas can be determined by inserting Eqs 3.5 and 3.7 into Eq. 3.4, however if the parameters of the trap are well known, for example in the case of a harmonic trap,

$$PSD = N \left(\frac{\hbar\bar{\omega}}{k_B T} \right)^3, \quad (3.8)$$




Figure 3.1: A simple illustration depicting the formation of a Bose-Einstein Condensate, adapted from Ref. [52]. a) The atoms are well separated, and can be treated as point particles at high temperatures. b) When reducing the temperature of the gas, the size of the atomic wave packets increases, until c) the wave packets overlap and form a condensate. d) At $T = 0$, all particles accumulate in the ground state and form a BEC.

can give a more accurate measure of the phase space density in optically dense samples. An upper limit on the temperature of a trapped BEC is defined by the critical temperature T_C

$$k_B T_C = \hbar \bar{\omega} \left(\frac{N}{\zeta(3)} \right)^{1/3}, \quad (3.9)$$

where $\zeta(x) = \sum_{n=1}^{\infty} 1/n^x$ is the Riemann-zeta function, with $\zeta(3) \approx 1.202$ [89]. Using this definition of T_C in Eqs 3.5 and 3.7, we find that the phase space density at the critical temperature is close to unity, with $PSD = \zeta(3) \approx 1.202$.

Below T_C , the fraction of atoms in the ground state, the condensate fraction, increases for lower temperatures, with

$$\frac{N_0}{N} = 1 - \left(\frac{T}{T_C} \right)^3, \quad (3.10)$$

which is plotted in Fig. 3.2. The sharp increase of the condensate fraction at temperatures just below T_C is indicative of a phase transition.




Figure 3.2: Theoretical calculation of condensate fraction vs the temperature of the gas relative to the critical temperature, plotted from Eq. 3.10.

3.3 Mean Field Approach

For a time-dependent potential, the equation describing the ground state is

$$i\hbar \frac{\partial}{\partial t} \Phi(\mathbf{r}, t) = \left(-\frac{\hbar^2 \nabla^2}{2m} + U(\mathbf{r}, t) + g|\Phi(\mathbf{r}, t)|^2 \right) \Phi(\mathbf{r}, t), \quad (3.11)$$

where g is an interaction coefficient related to the inter-atom s-wave scattering length a by

$$g = \frac{4\pi\hbar^2 a}{m}. \quad (3.12)$$

For a static potential, the ground state is described by

$$\left(-\frac{\hbar^2 \nabla^2}{2m} + U(\mathbf{r}) + g|\phi(\mathbf{r})|^2 \right) \phi(\mathbf{r}) = \mu\phi(\mathbf{r}), \quad (3.13)$$

where $\Phi(\mathbf{r}, t)$ has the form $\Phi(\mathbf{r}, t) = \phi(\mathbf{r})e^{-i\mu t/\hbar}$. These Gross-Pitaevskii equations (derived independently by Gross [90] and Pitaevskii [91] in 1960) take the form of non-linear Schrödinger equations because of the non-linear $|\phi|^2$ and $|\Phi(\mathbf{r}, t)|^2$ terms, which correspond to the density $n(\mathbf{r})$. The ground-state energy of the

3.3. MEAN FIELD APPROACH

system is a function of the density, giving

$$E[n] = \int d\mathbf{r} \left[\frac{\hbar^2}{2m} |\nabla \sqrt{n}|^2 + nU_{\text{ext}}(\mathbf{r}) + \frac{gn^2}{2} \right] = E_{\text{kin}} + E_{\text{ho}} + E_{\text{int}}, \quad (3.14)$$

where E_{kin} , E_{ho} and E_{int} correspond to the kinetic ‘quantum pressure’ energy only present in inhomogeneous systems, the harmonic oscillator and mean-field interaction energies, respectively. Taking the ratio of $E_{\text{int}} \propto N^2|a|/a_{\text{ho}}^3$ to $E_{\text{kin}} \propto N/a_{\text{ho}}^2$ leads to a parameter which emphasises the importance of interactions in a BEC

$$\frac{E_{\text{int}}}{E_{\text{kin}}} \propto \frac{N|a|}{a_{\text{ho}}}, \quad (3.15)$$

where $a_{\text{ho}} = \sqrt{\frac{\hbar}{m\omega}}$ is the harmonic oscillator length of the trap containing the BEC. The ratio of interaction to kinetic energies gives a useful gauge of stability in an attractive ($a < 0$) BEC. When the scattering length is negative, the condensate contracts to minimise its overall energy and so the ‘quantum pressure’ increases. When the interactions become strongly attractive ($a \ll 0$), the kinetic energy is outweighed by the interaction energy and the condensate becomes very dense, causing collisional losses. Because E_{kin} and E_{int} scale differently with atom number, a critical atom number can be calculated using [92, 93]

$$N_C = 0.575 \frac{a_{\text{ho}}}{|a|}, \quad (3.16)$$

where the prefactor is determined by approximating a spherical trap. If the condensate exceeds this atom number, the kinetic energy is no longer large enough to maintain stability in the condensate. A BEC collapse known as a ‘Bose-nova’ can be caused if a large BEC is produced under conditions of large positive scattering length before the interactions are adjusted to be strongly attractive [85, 86]. Equation 3.15 also leads to an important approximation for a highly repulsive BEC ($a \gg 0$) - the Thomas-Fermi approximation.

3.4 Thomas-Fermi Approximation

In the case of a BEC with a high atom number, the ratio $N|a|/a_{\text{ho}}$ (Eq. 3.15) becomes very large and the kinetic energy of the BEC is dwarfed by its interaction energy. In this case, the kinetic energy can be neglected leading to a simple solution for the density distribution of the BEC [89]

$$n_0(\mathbf{r}) = |\phi(\mathbf{r})|^2 = \max\left(\frac{\mu - U(\mathbf{r})}{g}, 0\right), \quad (3.17)$$

which takes the form of an inverted parabola with zero density cut-off points at $\mu = U(\mathbf{r})$. In this approximation, the size of the BEC, the Thomas-Fermi radius, is determined by the chemical potential and the trap potential. Assuming a spherical potential so that $\mu = U(|\mathbf{r}| = R) = m\omega^2 R^2/2$, the Thomas-Fermi radius is calculated to be

$$R_{\text{TF}} = \sqrt{\frac{2\mu}{m\omega^2}} = a_{\text{ho}} \frac{\bar{\omega}}{\omega} \left(\frac{15Na}{a_{\text{ho}}}\right)^{1/5}. \quad (3.18)$$

By normalising the density distribution $N = \int d\mathbf{r} n(\mathbf{r})$, an equation relating the chemical potential to the atom number [89]

$$\mu = \frac{\hbar\omega_{\text{ho}}}{2} \left(\frac{15Na}{a_{\text{ho}}}\right)^{2/5}, \quad (3.19)$$

and an equation for the peak density

$$\hat{n} = \frac{\mu}{g} \quad (3.20)$$

are obtained.

When the condensate is released from the trap and allowed to expand, its shape evolves as a rescaling of its original parabolic form [94–96]. Now considering a cigar shaped trap, the condensate will expand faster in the axis within which

it was more tightly confined. The expansion rate of the Thomas-Fermi radii for the radial and axial directions of the condensate are given by [52]

$$R_r(t) = R_r(0)\sqrt{1 + \tau^2} \quad (3.21)$$

$$R_{\text{ax}}(t) = R_r(0)\epsilon^{-1} \left(1 + \epsilon^2 \left[\tau \arctan \tau - \ln \sqrt{1 + \tau^2} \right] \right) \quad (3.22)$$

where $R_r(0)$ and $R_{\text{ax}}(0)$ are the original Thomas-Fermi radii in the radial and axial directions which can be calculated from Eq. 3.19, $\epsilon = \omega_{\text{ax}}/\omega_r$ is the aspect ratio of the trap and $\tau = \omega_r t$. This allows us to determine the density of the BEC during expansion.

3.4.1 Thermal Density Distribution

For temperatures approaching the BEC critical temperature T_C , the gas is described by a semi-classical regime, where the influence of quantum statistics becomes significant and results in a divergence from Maxwell-Boltzmann statistics. In this regime, the density of the gas is given by [97]

$$n_T(\mathbf{r}) = \frac{g_{3/2}(ze^{-U(\mathbf{r})/k_B T})}{\lambda_{\text{dB}}^3}, \quad T \geq T_C \quad (3.23)$$

where for a harmonic potential of the form $U(\mathbf{r}) = m\omega^2 r^2/2$ is the trap potential with frequency ω , $z = e^{\mu/k_B T} < 1$ is the fugacity of the gas, and the function $g_{3/2}$ is a polylogarithm of the form $g_a(z) = \sum_{n=1}^{\infty} z^n/n^a$, which indicates the role of quantum statistics on the density distribution of thermal atoms.

Integrating this over the imaging axis y , the column density,

$$\tilde{n}_T(x, z) = \sqrt{\frac{2\pi k_B T}{m\omega_y^2}} \frac{g_2(ze^{-U(x,z)/k_B T})}{\lambda_{\text{dB}}^3}, \quad (3.24)$$

is obtained, which can be fitted to images of the gas to derive its properties. Below T_C , the fugacity of the gas becomes 1 and we require a different fitting

procedure that includes both its thermal and condensed components.

3.4.2 Bimodal Distribution

The study of the density distribution of a quantum gas facilitates the measurement of many properties of a Bose gas over the temperature range $0 \leq T \leq T_C$. A combined study of density distributions of a quantum gas both in-situ and after free-expansion allows us to determine the temperature and interaction energy of the gas. At finite temperature, the density distribution of the quantum gas typically combines the features of a thermal gas and a BEC (Sections 3.4 and 3.4.1).

Given that the density distribution of the atoms depends on which regime they are in, finding a fitting procedure that can simultaneously determine the properties of both the thermal and condensed atoms can be challenging. In an interacting gas, the condensed and thermal components of an atom cloud impose forces on one another, resulting in the thermal atoms being pushed outwards by the condensed atoms from the centre of the cloud. This regime requires coupled equations for the different density profiles. In this case, the in-situ density distribution for the thermal component of the atom cloud for $T < T_C$ is

$$n_T(\mathbf{r}) = \frac{g_{3/2}(e^{-(U(\mathbf{r})+2gn_0(\mathbf{r}))/k_B T})}{\lambda_{\text{dB}}^3}, \quad (3.25)$$

where n_0 is the condensate density given by the Thomas-Fermi approximation in Eq. 3.17. By summing Eqs 3.17 and 3.25, we find the 3D density distribution of an interacting gas containing both thermal and condensed atoms - a bimodal distribution.

Figures 3.3 and 3.4 present our simulations of a quantum gas using this method. Figure 3.3 shows a slice through the centre of a condensate with and without the thermal component, which illustrates the interactions between the

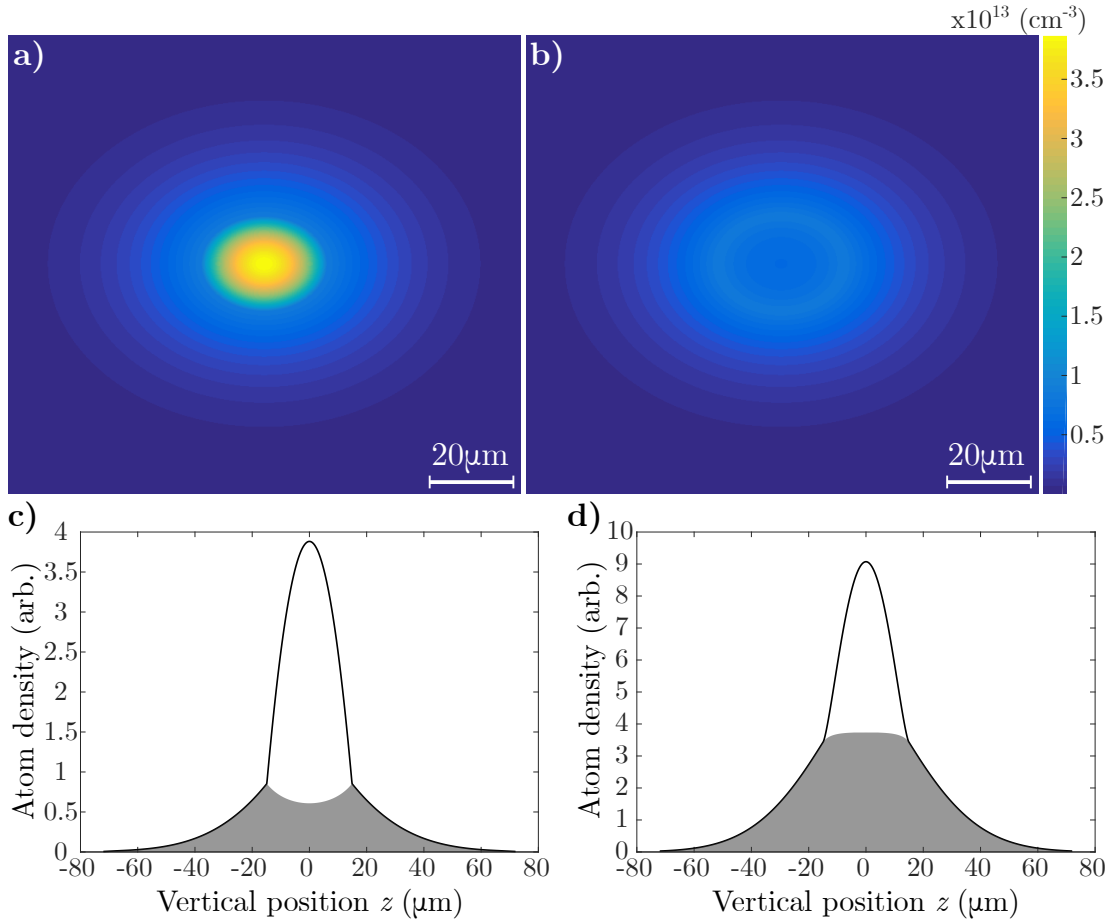


Figure 3.3: Density profiles of a slice from the centre of a quantum gas simulated using Eqs 3.17 and 3.25, with a temperature $T/T_C = 0.96$ and repulsive interactions. a) Density profile of the gas containing thermal and condensed components. b) Only the thermal component of the gas. Interactions with the condensed fraction displace the thermal component from the centre to the edge of the cloud. c) The line profile through the centre of the quantum gas with (white area) and without (grey shaded area) the condensate fraction. The effect of interactions between the components is even more apparent. d) The integrated density profile of this slice with and without condensate. Some detail from the interaction artefact is lost.

thermal and condensed components. The line profile also shows the effects of interactions in the densities of these components, but this information is lost in the integrated density profile. In Fig. 3.4, we can see the full condensate, integrated over the imaging axis to appear like absorption images we would produce experimentally. From looking at the simulated image with and without the condensed




Figure 3.4: A simulated quantum gas with the same parameters as in Fig. 3.3, integrated along the imaging axis. a) Density profile for the gas, with and b) without the condensed component. c) Integrated density profile of the gas with (white area) and without (grey-shaded area) the condensed component.

component, and the integrated density profile, we can see that all the information regarding interactions is now lost. This helps to illustrate the challenge of fitting bimodal density distributions to absorption images.

When imaging time-of-flight density distributions of a quantum gas, the effects of interactions between the thermal and condensed components on the expansion of the gas are not easily determined. As a result, we disregard the interaction term in the equation describing the density of the thermal component. The column density of the full bimodal distribution after a time-of-flight duration is then [52, 98]

$$\begin{aligned} \tilde{n}(x, z, t) = & \frac{1}{(\lambda_{\text{dB}} \bar{\omega} t)^3} \sqrt{\frac{2\pi k_B T}{m}} t^2 g_2\left(e^{-\frac{m(x^2+z^2)}{2k_B T t^2}}\right) \\ & + \tilde{n}_0(0) \max\left(1 - \frac{x^2}{R_r^2(t)} - \frac{z^2}{R_{\text{ax}}^2(t)}, 0\right)^{3/2}, \end{aligned} \quad (3.26)$$

where $\tilde{n}_0(0)$ is the peak column density and R_r and R_{ax} are the radial and axial Thomas-Fermi radii after the duration t . An example of this fitting method is shown in Section 5.1.3, and is used throughout Section 5.3 to determine the properties of the BEC.

Chapter 4

Experimental Setup

This Chapter describes the experimental apparatus we use to create and probe Bose-Einstein condensates. The vacuum setup, laser systems and computer control all must be synchronised with microsecond precision to perform experiments reliably and reproducibly. Here, I discuss the main parts of the vacuum chamber, lasers and electronics we use, how they were prepared and how they are interfaced with the computer system.

Section 4.1, describes the vacuum setup including the vacuum pumps, the source of caesium atoms and the glass cells within which experiments take place. The laser configuration, atomic transitions addressed by the lasers and their roles in the cooling process are outlined in Section 4.2. Sections 4.3 and 4.4 detail our method of generating magnetic fields, and the construction of the coils used to accomplish this. In Section 4.5, the computer control arrangement, and how the computers interface with the experiment and one another are described, and finally in Section 4.6 our uses for microcontrollers in the lab environment are presented.

4.1 Vacuum Setup

Ultracold atoms experiments require an ultra-high vacuum setup to minimise heating and loss due to collisions with background gases. Our vacuum setup was designed to be compact and versatile so there is plenty of space for the plethora of optics required to operate the experiment, and the experimental cells have excellent optical access. A diagram and a photo of the vacuum apparatus at an early stage of development are shown in Fig. 4.1. The setup is made primarily of stainless steel, with exceptions of the 2D+ MOT cell and main experimental cell, which are made of quartz glass. The cells are connected by a differential pumping tube, which enables a pressure differential of three orders of magnitude between them. The pressure differential is maintained by two ion pumps and a Ti sublimation pump, and we monitor the vacuum using pressure gauges.

4.1.1 Caesium Source

The setup contains two oven sections, implemented by small bellows close to the 2D+ MOT section. Currently one oven contains a glass ampoule with 5 g of caesium. The second bellow is currently empty, which permits the addition of a second atomic species at a later date, if desired. The caesium was initially released by bending the bellow after the bakeout procedure, cracking the ampoule inside. The caesium vaporises in the high vacuum environment and diffuses into the 2D+ MOT. We control the caesium flux with a manual valve, which separates the oven section from the glass cell.

4.1.2 Glass Cells

The use of glass cells in our vacuum rather than metal science chambers with glass windows increases the number and size of laser beams that can be used in our experiment, and reduces magnetisation and eddy currents. Both glass cells

4.1. VACUUM SETUP




Figure 4.1: Images of the vacuum chamber. a) A computer generated render of the vacuum apparatus. Atoms are pre-cooled in the ‘high-pressure’ 2D+ MOT in the glass cell on the left, accelerated through the hole in the gold mirror and through the differential pumping tube to the glass cell on the right. Here the rest of the cooling process and experiments are performed. b) A photo of the vacuum setup in an early development stage, just after bakeout and before the addition of optics or magnetic field coils. The red tint around all the edges of the 2D+ MOT cell is the epoxy used to bond the cell, which is not present on the glass cell on the right, the UHV region used as our main experimental cell, as this was bonded using an optical contact method. The bellow containing caesium is wrapped in aluminium foil.

(manufactured by Japan Cell) are made from quartz, and are of identical dimensions, both 40 x 40 x 140 mm internally. They differ, however, in the methods used to bond them. The glass plates forming the 2D+ MOT cell are held together using an epoxy which, because it would be likely to outgas in the UHV region, would not be suitable for the main experimental cell. The experimental cell is formed using optical contact bonding, which relies on intermolecular forces. An anti-reflective coating is used on both the inside and outside of the windows on the 2D+ MOT cell, but only on the outside of the main experimental cell, because the optical bonding method would not work if the inside were also coated. The differences in these glass cells can be clearly seen from the photo of the vacuum setup shown in Fig. 4.1. Furthermore, the neck of the main experimental cell is 45 mm longer than that of the 2D+ MOT cell (100 mm compared to 55 mm). This increased distance from the vacuum chamber reduces the effects of eddy currents circulating in the stainless steel body, which would produce undesired magnetic fields at the position of the atoms after the magnetic field coils were switched off. Above and below this cell are pieces of anodised aluminium breadboard, which hold magnetic field coils in place, and add more surfaces to which optics can be secured for three dimensional optical access.

The glass cell that houses the 2D+ MOT section contains a vacuum pressure of 10^{-8} mbar. At the neck of this cell, there is a gold plated aluminium mirror, rotated 45° from the longitudinal axis, to reflect the ‘plus’ beam. There is a hole in the centre of this mirror to allow atoms to propagate through the differential pumping tube towards the main experimental cell. The experimental cell contains a pressure of 10^{-11} mbar.

4.1.3 Pumping

The difference in pressure between the two glass cells is possible because of the differential pumping tube. The conductance of the differential pumping tube is

given by

$$C \approx 0.125 \times \frac{d^3}{L} \quad (4.1)$$

where $d=4$ mm is the diameter of the tube, and $L=150$ mm is its length, giving a conductance value of 0.053 l/s. The pressure difference is maintained by an 8 l/s ion pump (Agilent Vacion pump 8 l/s) in the 2D+ MOT section, and by a 55 l/s ion pump (Agilent Vacion plus 55 l/s) for the experimental cell, for which we also use a titanium sublimation pump (VacGen ST22). A hot filament pressure gauge (Agilent Varian MBA2-200T) used in combination with a controller (XGS 600), provides readings of the pressure in the UHV region of the vacuum chamber.

4.1.4 Magnetic Shielding

Because the ion pumps contain permanent magnets, their magnetic fields will interfere with magnetic fields produced during our experimental procedures if left bare. We have μ -metal shields that encase the ion pumps to reduce the magnetic field experienced by the atoms in the main experimental cell to less than 1 G. In addition, both glass cells are surrounded by shim coils, to minimise stray magnetic fields. The shim coil pairs are made from 0.5 mm wire and can generate homogeneous magnetic fields up to 5 G with the flow of co-propagating electric currents.

4.1.5 Preparation and Assembly

To prepare the vacuum chamber for our experiments, it was necessary to clean each part individually then heat the vacuum chamber as a whole. This process served to remove dirt and impurities, such as water, which reduced the potential rate of background collisions. All the metal parts we used were initially bathed in an ultrasonic bath containing acetone and then an ultrasonic bath of isopropanol, and wrapped in aluminium foil until they were ready to be used. We constructed

the vacuum chamber incrementally and checked for leaks at each stage using a turbo pump, helium and an atomic mass spectrometer. Once fully constructed, we wrapped the entire chamber in aluminium foil to evenly distribute heat before introducing heating tape and thermocouples for the bakeout process, and wrapping in foil once more for heat retention. Bakeout is necessary to minimise impurities, especially water, sticking to the inner walls of the chamber after the pumping process. We performed the bakeout by heating the chamber over two days to a maximum temperature of 150°C. Great care was taken to avoid introducing a temperature gradient across different regions, as this could damage the glass cells or windows in the chamber. While monitoring the temperature using thermocouples at several sections of the chamber between the layers of aluminium foil, we increased the temperature by 5°C every hour until we reached the maximum temperature, which we maintained for ten days before reducing the temperature equally as slowly. A turbo pump was run continuously throughout this process, with our ion and Ti sublimation pumps being switched on during the cooling stage.

4.2 Laser Configuration

This Section is intended only to give an overview of the lasers used in our experiment, and the techniques used to tune and lock them - the optical setup is described in more detail in the thesis of Andrea Di Carli [58].

We laser cool caesium-133 atoms using laser wavelengths in the region of 852nm to excite the $6S_{1/2} \rightarrow 6P_{3/2}$ transition. Different near-resonant laser frequencies are used for the MOT/imaging, repumper beams and Raman cooling beams as illustrated in Fig. 4.2. The light frequencies are all produced using extended cavity diode lasers (ECDLs) using 852nm laser diodes (Thorlabs L852P150).

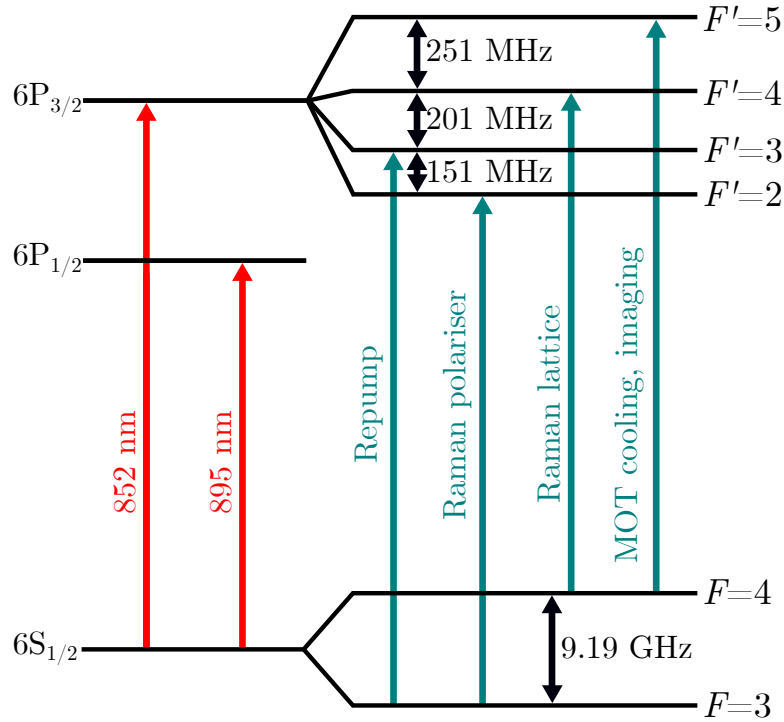


Figure 4.2: A caesium energy level diagram with the transitions used for various stages of the experiment.

When not used as part of ECDLs, these diodes typically have linewidths on the order of 100's of GHz, which is far too broad to excite a single atomic hyperfine state - a problem caused mainly by a short cavity length inside the diode. This problem is alleviated by building the diode into a longer cavity, with a diffraction grating for frequency-selective feedback, in the Littrow configuration [99].

Each ECDL is connected to a laser driver which adjusts the temperature, current and piezo voltage of the laser to select its frequency. The temperature of the extended cavity is controlled and stabilised using a Peltier element on the underside of the ECDL body. The electrical current used to drive the diode controls the optical output power generated by the diode, but also affects the temperature of the small cavity inside, which in turn adjusts the output frequency. A voltage is applied to a piezo which controls the position of the diffraction grating, so is used to fine-tune the frequency produced by the laser. Applying a

4.2. LASER CONFIGURATION

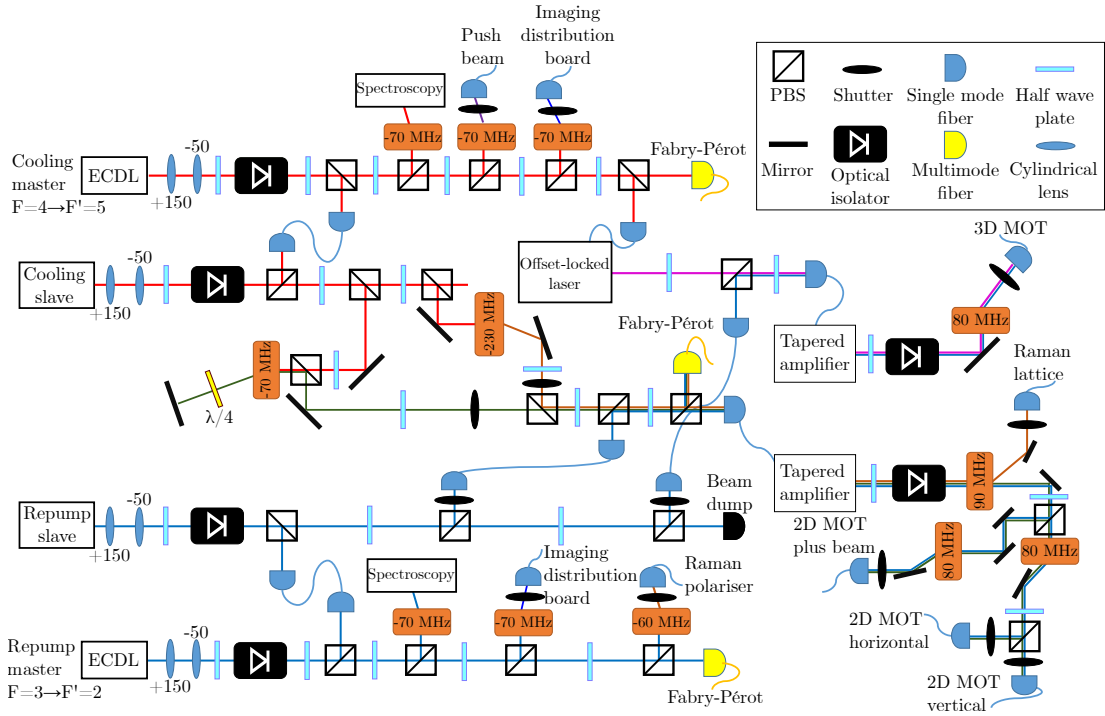


Figure 4.3: Diagram of the configuration of diode lasers and optics used. Two extended cavity diode lasers serve as ‘master’ lasers for the cooling and repump transitions. Injection seeded ‘slave’ diode lasers provide additional light power at the same frequencies as the masters. The frequency and intensity of each beam is controlled by acousto-optic modulators before the light is coupled into optical fibers, which lead to the experiment. Figure adapted from Ref. [58].

sawtooth wave to the piezo makes it possible to scan over a range of frequencies for the purpose of spectroscopy. When no frequency sweep is used, our ECDLs have linewidths of about 300 kHz and can be locked to hyperfine transitions using polarisation spectroscopy (see Section 4.2.2).

We use saturated absorption spectroscopy to tune the lasers into the desired frequency range, and the error signal obtained by polarisation spectroscopy to lock the lasers to specific transitions. Our master cooling laser is locked 70 MHz blue-detuned from the closed $F = 4 \rightarrow F' = 5$ transition, and the repumper master laser is locked 70 MHz blue-detuned from the $F = 3 \rightarrow F' = 2$ transition. The lasers operate in a single mode at powers of about 100 mW, but to run our experiment we need to generate significantly more power. We achieve this by

dedicating a small amount of power from the master lasers to injection locking of free-running laser diodes, forcing them to output at the same frequency, each with output powers of up to 150 mW. For our MOTs we need to further increase the available power by an order of magnitude, which we achieve using tapered amplifiers (TA, Thorlabs TPA850P10). To fine-tune the laser frequency before the light is sent to the experiment and for fast switching of powers, acousto-optic modulators (AOMs) are used. The configuration of these lasers and optics are shown in Fig. 4.3.

The crossed dipole trap ‘reservoir’ beams are generated by a 200 W, 1070 nm ytterbium fiber laser (IPG YLR-1070-WC). The reservoir beams have a relatively large 800 μm waist to catch as many atoms as possible after Raman cooling. In addition, we use light from a 1064 nm Nd:YAG laser (Coherent Mephisto) as the injection seed for a 1064 nm laser diode, which generates 60 mW of optical power. This light is then passed through a 50 W fiber amplifier (Nufern NUA-1064-PD-0050-D0), after which this light is used for crossed dimple beams, which have tighter waists of 50 μm horizontal and 90 μm vertical. In a typical experimental sequence, less than 500 mW of the amplified light is used for the dimple beams, so this light can also be used to implement an optical lattice. The cooling methods are covered in more detail in Section 5.2.

4.2.1 Saturated Absorption Spectroscopy

We observe and monitor the frequencies of our lasers by detecting the absorption spectrum of room temperature caesium atoms in vapour cells. A problem with this method of spectroscopy is that at room temperature, the Doppler broadening of the absorption spectrum masks the hyperfine transitions, so the observed ‘dip’ on the spectrum is 100s of MHz wide and locking to a specific atomic transition is not possible. Saturated absorption spectroscopy is a method which reveals these transitions within the Doppler-broadened absorption profile, which span tens of

4.2. LASER CONFIGURATION




Figure 4.4: The saturated absorption spectra of room temperature caesium atoms for the cooling (left) and repump (right) lasers. The cooling laser spectrum indicates the transitions from the $F = 4$ ground state to each allowed excited state, and each crossover between these transitions. The repump laser spectrum shows all allowed transitions from the $F = 3$ ground state.

MHz instead of hundreds [100, 101].

The principle behind saturated absorption spectroscopy is as follows: a laser that is scanned across the atomic transitions of caesium is split into two beams - a strong ‘pump’ beam, and a weak ‘probe’ beam [102] - which counter-propagate through a vapour cell, crossing one another. After leaving the cell, the pump beam is blocked and the probe beam is directed into a photodiode. The photodiode signals for the saturated absorption spectroscopy of caesium are shown in Fig. 4.4. There are three regimes for laser frequency in this scheme: on resonance, between resonances and non-resonant.

When the laser is resonant with a hyperfine transition, the atoms that are moving with zero longitudinal velocity with respect to the beam axis, $v = 0$, absorb the light. Since the pump beam is close to the saturated intensity of that transition, the majority of the $v = 0$ atoms in the beam path are pumped into the excited state. The counter-propagating probe beam is also resonant with the $v = 0$ atoms, which are now transparent to the probe light, so this beam is hardly absorbed. The transitions for this case are annotated with single F' values in

Fig. 4.4.

When the laser is between resonances, the pump beam and probe beam are both resonant with different velocity classes due to the Doppler shift, so the probe beam is absorbed by the same amount as if the pump beam were not present.

The third case is when the laser frequency is directly in the centre of two transitions - the pump beam is absorbed by atoms propagating towards the light source at a velocity v , which blue-shifts the laser frequency into the higher resonance, and by atoms of the velocity $-v$, which red-shifts the frequency into the lower resonance. As a result, most of the atoms of these same two velocity classes are transparent to the probe beam. These peaks in the absorption spectrum are known as ‘crossover’ peaks, and tend to be more prominent than the peaks from the resonance case since they result from two velocity classes rather than one. If we scan over a Doppler broadened transition using this technique, we can observe and identify the peak corresponding to each hyperfine transition and cross-over, which simplifies the action of locking our laser to a specific transition. The crossover peaks are indicated in Fig. 4.4 with two F values (e.g. $F' = 3, 4$).

4.2.2 Polarisation Spectroscopy

Once we find the desired hyperfine transition using saturated absorption spectrography, the laser must be locked to prevent the frequency from drifting. We use polarisation spectroscopy to provide a dispersive signal at the desired frequency for the laser driver to lock the laser to [103–105].

The setup for polarisation spectroscopy, seen in Fig. 4.5, is not dissimilar to the one for saturated absorption spectroscopy in that they both require pump and probe beams to be counter-propagating through a reference cell, and the spectroscopy can be performed by scanning over the Doppler broadened transmission spectrum of the atoms.




Figure 4.5: A simplified schematic of the optical arrangement that enables polarisation spectroscopy.

In polarisation spectroscopy, however, the polarisation of each beam becomes important as the mechanism depends on the birefringence of the medium that arises when certain closed atomic transitions are populated. The pump beam should be circularly polarised, repeatedly inducing σ^+ or σ^- transitions, to pump the atom towards the closed transition at $m_F = \pm F$. The probe beam is linearly polarised at 45° from horizontal so that with no pump beam present the beam is 50% transmitted and 50% reflected by the polarising beamsplitter. In this case, the beam can be said to consist of equal portions of left-handed and right-handed circular polarisations of light. In the presence of the pump light, the different m_F states are populated asymmetrically, so in the case that $m_F = +F$, we have that the line strength for the σ^+ transition is stronger than that of the σ^- transition, so light of one handedness is absorbed more than the other. In addition to the loss of light due to spontaneous emission, both circular polarisation components experience a different refractive index in the medium, so one polarisation is retarded in relation to the other. For the $m_F = +F$ case, this refractive index retards the light that causes σ^+ transitions more than the light that causes σ^- transitions, which results in a relative phase shift of the different circular polarisations, so a rotation of the linear polarisation. The light is no longer 50% transmitted/re-

flected by the beamsplitter, and when the two components are detected by the differential photodiode, the result is a dispersive signal.




Figure 4.6: The polarisation spectra for the cooling (left) and repump (right) lasers, measured using a difference photodiode. The cooling laser is locked to the negative slope of the $F=4 \rightarrow F'=5$ transition, and the repump, the positive slope of the $F=3 \rightarrow F'=2$ transition (both indicated). The locking circuit monitoring the difference photodiode signal locks to the sharp dispersions of these transitions, crossing the 0 V point.

The signal obtained from the differential photodiode is used by the laser controller as an error signal for locking, so it is important to maximise this signal for the desired locking transition. A small coil (not shown) is used to adjust the magnetic field around the vapour cell to cancel the Earth's magnetic field [105], and the polarisations of the beams are adjusted to strengthen the signal for any closed transition over the others. In Fig. 4.6, we can see the polarisation spectroscopy signal for caesium on the cooling and repump transitions. The difference of the two photodiode signals is used as an error signal which the laser driver uses for locking, so it is important to maximise the signal for the desired locking transition. The transitions used for locking are marked with arrows.

Polarisation spectroscopy has a large capture range, as observed in Fig. 4.6, where the width of the dispersive signals produced by the difference photodiode define the capture range. We lock to the negative slope of the signal for the

$F = 4 \rightarrow F' = 5$ cooling transition, and the positive slope of the signal for the $F = 3 \rightarrow F' = 2$ repump transition. The range of the slope for the locking signal of the cooling transition is ~ 10.5 MHz, and 22 MHz for the repump transition.

4.3 Magnetic Field Generation

Magnetic fields can easily be generated by passing electrical currents through copper wire. The strength of the magnetic field produced is calculated using the Biot-Savart law:

$$d\mathbf{B} = \frac{\mu_0 I d\mathbf{L} \times \mathbf{r}}{4\pi |\mathbf{r}|^3}, \quad (4.2)$$

where \mathbf{B} is the magnetic field at the distance \mathbf{r} away from the wire, $d\mathbf{L}$ is the wire element at the origin, and μ_0 is the permeability of free-space. Winding the wire to form a coil of multiple loops allows us to generate the desired magnetic field strengths using moderate currents. Multiple coils can be combined in different configurations to generate the desired magnetic field regime. Most coil configurations have that the coils are placed parallel to one another and have either co-propagating or counter-propagating currents passed through them, the former generating a uniform magnetic field and the latter a magnetic field gradient. In the process of building each phase of the experiment, we model the magnetic field at the position of the atoms using Eq. 4.2, to determine the currents required for a given number of windings in the coil. Our code takes into account the number of windings, the thickness and resistivity of the wire, the coil dimensions and their respective positions. The results of this model are presented in Fig. 4.7, featuring the visualisation of the coils, and the magnetic field and magnetic field gradients with respect to position, where the position $0 \mu\text{m}$ indicates the desired location of the atoms.




Figure 4.7: Magnetic field coil simulations. a) A visualisation of the magnetic field coils corresponding to the coils in our setup. Shown are the main experiment coils (blue), the vertical shim coils (green), and the horizontal shim coils that shift the magnetic field left/right (brown) or forward/backwards (red) on the imaging plane. b) The magnitude of the magnetic field vs vertical position, and c) the magnetic field gradient produced by simulations of coils for a 3D MOT. In this simulation, the main coils consist of 156 windings of copper wire each, driven at 4.3 A at z positions of 41 mm and -41 mm.

4.4 Magnetic Field Coils

The coils used for the 2D+ MOT and main experimental cells differ significantly due to the magnetic fields required from them. Fixed near all the long sides of the 2D+ MOT glass cell, there are four magnetic field coils for the required 2D quadrupole field. The coils are rectangular in shape with inner dimensions of 30×150 mm, and made from 0.8 mm diameter copper wire in 14 layers from the centre, with 10 windings per layer. These are mounted around the glass cell using a selective laser sintered, graphite reinforced plastic mount. These coils are operated in the counter-propagating current configuration, creating a magnetic field gradient in two dimensions. The 2D+ MOT requires magnetic field gradients of approximately 10 G/cm in two axes, which can be provided by driving the coils with 1 A of current.

Resting on the pieces of breadboard above and below the experimental cell, are two main sets of coils which are much larger. The wire used for these coils is rectangular, measuring 2×1 mm. Each coil is separated into seven sections, so different combinations of sections can be used in various configurations based on what is required for the experiment, as seen in Fig. 4.8. The red sections are used to generate magnetic field gradients of ~ 12 G/cm (4.3 A) for the 3D MOT, and 34.1 G/cm (11.85 A) for magnetic levitation. The blue section is used to create a magnetic field offset for tuning scattering length and increasing the curvature of the magnetic levitation field (see Fig. 2.6). The zero crossing for scattering length in the $|3, 3\rangle$ state of caesium occurs at 17 G, above which the scattering length is positive. Positive scattering lengths are required for Bose-Einstein condensation, so we tend to use magnetic fields above 17.12 G (9.39 A). As a result, two power supplies are used to drive the coils - connected in the co-propagating current configuration is the ‘offset’ supply (Delta Elektronika SM120-50), and the ‘gradient’ supply (Delta Elektronika SM7020-D) is used to drive coils in the counter-propagating current configuration. These power supplies would

4.4. MAGNETIC FIELD COILS

allow a maximum of 20 A to drive the gradient coils, generating a maximum gradient of 56.9 G/cm, and 50 A to drive the offset coils, generating a maximum offset of 113 G. In addition, a 450 V, 4700 μ F discharge capacitor is used to drive the green section of the coils in the counter-propagating current configuration, to provide a levitation gradient quickly. When magnetic levitation is performed, and all three sets of coils are ramped up simultaneously, the capacitor begins generating its maximum current after 50 μ s, as opposed to the 8.4 ms it takes the two Delta power supplies to level off on the desired currents.

During a routine cooling sequence, while performing levitation and tuning the scattering length to be positive, the main experimental coils dissipate \sim 120 W of power. Because this can generate a lot of heat, these coils are sealed inside mounts made of selective laser sintered nylon to allow cooling water to circulate while in contact with the wire. To aid heat transfer from the wire to the water, copper fins are used at regular intervals around the mount, serving as heat sinks (see Fig. 4.8). We cool the coils using a closed circuit of water, which is pumped by a water cooling pump through one partition of a heat exchanger. Mains water flows through the other partition in the opposite direction, promoting the transfer of heat from the closed circuit. Left unregulated, there is a danger that the closed circuit water could be cooled too much, causing water to condense on the outside of the coils on the optical bench - for this reason the flow of the mains water is controlled by a solenoid valve, operated by an Arduino microcontroller connected to a motor shield, using a digital temperature sensor for feedback. Also around this cell are six shim coils made using 0.5 mm wire, which compensate for stray magnetic fields to shift the position of the magnetic field minima to the centres of the cells in the case of magnetic field gradients.

4.4. MAGNETIC FIELD COILS




Figure 4.8: a) Main experimental coils in their enclosure before being sealed. The coil sections and the heat sink fins are clearly visible. b) The enclosure after being sealed. The hosepipes providing cooling water and drainage have been secured to the hose barbs protruding from the mount. c) An illustration of the different sections of a single coil. The centremost (green) section is powered by a discharge capacitor for fast magnetic field generation. The red sections are gradient coils used for levitation, and the outermost (blue) section is a coil used to generate a magnetic field offset for tuning interactions. The rest of the sections (grey) are currently unused. d) A cross-sectional render of the coil mount. The four grooves surrounding the coils serve as cooling water channels, facilitating the flow of cooling water around the coils throughout the mount.

4.5 Computer Control

For any experiment with a complex sequence such as ours, computers are essential to execute each stage of our sequence with microsecond precision, and in a consistent, reproducible manner. The computers in our lab carry out camera control, vacuum monitoring, image processing and analysis, and analog and digital control of devices during the experimental sequence. In addition, there are several microcontrollers used as regulation devices, and used to monitor the conditions surrounding our lasers and vacuum chamber.

The experimental control and data analysis is provided by several computers in synchronicity, which communicate with one another via a UDP (User Datagram Protocol) network protocol. The experimental control tasks and communications are split between the computers to promote fault tolerance in the system. The following paragraphs detail how the computers work synchronously with one another, a flowchart summarising the data flow is shown in Fig. 4.9.

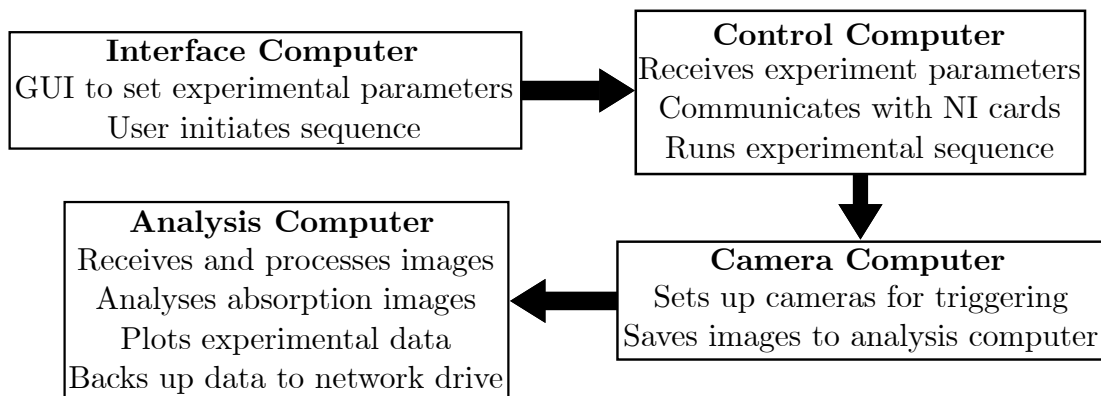


Figure 4.9: Flowchart outlining the role of each computer in the experimental control arrangement. The user begins the sequence from a GUI on the interface computer, after which the data is transferred as shown.

Interface computer This is the ‘front end’ computer, which is mainly used to set experimental parameters and execute the experimental sequence. A Graphical User Interface (GUI) written in Matlab allows for each parameter to be set, and

for the sequence to be executed. When we begin the sequence, this computer saves all the parameters to a file containing Matlab matrices with channel timing and output information. The file is transferred to the control computer via network file sharing. Because the interface computer does not control any hardware, we can change and reprogram sequences while the experimental setup is running.

Control computer A c++ program on the control computer processes incoming files and uses the information therein for hardware control. There are five National Instruments cards connected to this computer (2x NI-6534(digital), 2x NI-6733(analog), 1x NI-6723(analog/digital)), giving 40 analog and 64 digital outputs which generate signals for electronic devices around the lab with timing limited to either 5 μ s or 10 μ s. This computer also sends Standard Commands for Programmable Instruments (SCPI) via USB and GPIB connections to signal generators, to prepare them to be used in the sequence.

Camera computer Three cameras are connected to this computer: two CMOS cameras (MatrixVision BlueFox3-1012bG) and a CCD camera (Allied Vision Manta G-125B), which are each activated by hardware triggers. At the end of each cycle, the camera pictures are stored in a Matlab file and a trigger file is sent via the network to the analysis computer.

Analysis computer Once each set of three images with the experimental parameters has been saved, this computer processes them to generate an absorption image, which it displays in a Matlab GUI. The region of interest (ROI) can be selected by the user, and the application calculates the profile for that ROI, before performing a selected profile fit. From these fits, the application obtains properties (width of cloud, atom number etc.), which can be plotted against any of the experimental parameters.

4.6 Microcontrollers

Where we use PCs for detailed control of the experiment and analysis, we use microcontrollers (typically Arduino Unos or Arduino Mega 2560s) for continuous autonomous monitoring of the lab environment and the feedback-based regulation of some devices, and the communication between our control software with other devices. The advantages of using microcontrollers include customisability, the availability of attachments that are simple to implement, the multitude of software libraries that simplify programming, and their small form factors. The microcontrollers in our lab are used in several dedicated devices, for example our shutter drivers, motor drivers for rotation mounts, or the temperature controllers for our MOT coils and diode lasers. However, the bulk of our microcontrollers are used for monitoring magnetic fields, temperatures, vacuum pressures and the presence of water leaks. There are several factors that can affect the stability of the experiment - air temperature, vacuum pressure, table tilt and leak detection are some of the properties we monitor constantly. Keeping a close eye on these values allows both for us to react to drifts before they become critical, and to ensure the consistency of the results obtained from the experiment.


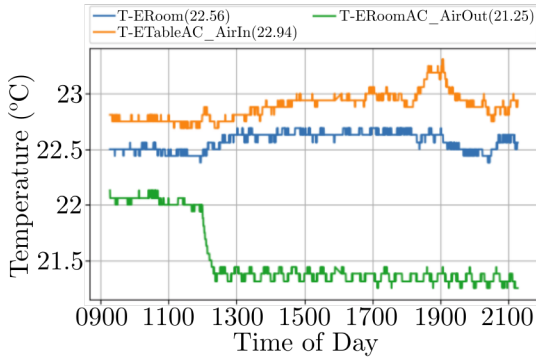
Each microcontroller that records these parameters and transfers them in UDP packets via a network switch to a Raspberry Pi which saves the data, compares the data to our set limits, graphs it, and uploads the graphs to a server so they can be viewed remotely. There is a warning system put in place for certain parameters, which alerts users by SMS and email in the case that they exceed safe limits.

Temperature Small changes in the air temperature surrounding the lasers can cause them to become unstable, unlock or mode-hop and bigger changes can cause the optics on the bench to expand or contract, misaligning the laser beam and decoupling optical fibres. We have several DS18B20 waterproof temperature

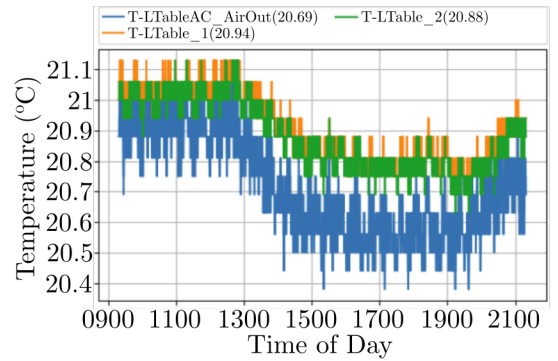
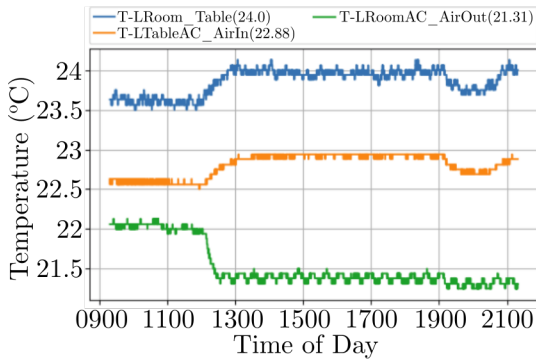
4.6. MICROCONTROLLERS

probes around the lab which monitor the water and air from the air conditioning unit, and the optical benches for any drifts in temperature. There is a warning system in the event that any of these temperatures drift out of our defined bounds.

Temperatures Exp Room



Temperatures Laser Room



Magnetic Fields

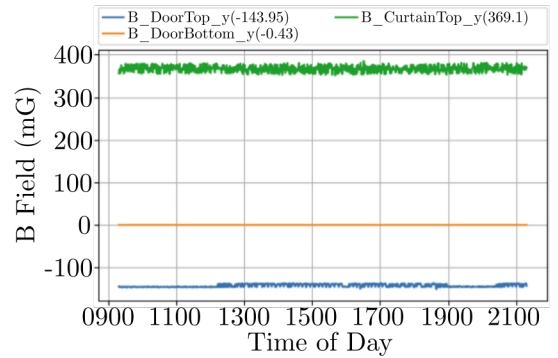
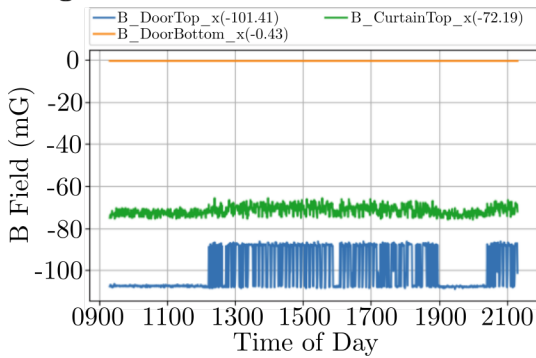


Figure 4.10: A series of plots as they appear on our lab monitoring webpage. The temperatures and magnetic fields in the lab are monitored, the time evolution is shown.

Table tilt Our vacuum chamber is on an optical bench that is suspended above the ground on top of an air-bed to isolate it from vibrations in the building. As a result, this bench could easily be tilted by a heavy weight on one end of the bench or a loss of pressure in the air-bed. This can affect the position stability of our atom cloud, so the tilt of the table is tracked using an MPU6050 accelerometer and gyroscope breakout board, which is secured to the surface of the optical bench.

Leak detection Our air conditioning units use water to cool the air flowing through them, and we use water cooling for our 3D MOT coils. In the event of a water leak, we need to be able to act quickly to mitigate any damage. For this reason, we have ‘rain sensor’ modules (YL-83) placed strategically around the lab. There is a warning system in place for the event of a leak being detected.

4.6.1 Optical Shutters

In most laser labs, the ability to independently extinguish laser beams in a time-dependent manner is a basic necessity - particularly when stray light can cause interference or significantly change the observed results. Mechanical optical shutter systems are often used for this purpose, but a driver unit and shutter units with desirable properties can cost upwards of £1000. Some groups design their own laser shutters using repurposed materials to save this cost [106–108], and we have done something similar modifying a mechanical diaphragm shutter that would otherwise be found in digital cameras. Our laser shutter design and benchmarks can be found in more detail in Ref. [109], which is featured in Appendix A.

The shutter consists of three solenoids, each of which controls a set of shutter blades and has its own pair of solder points, allowing for each set of blades to be controlled independently of the others. One of these blades only serves to provide a smaller aperture, which we have no use for, so we do not solder to




Figure 4.11: Photos of our shutter design. a) Front of the shutter, where the filter blade, and the manual override sliders can be seen. Red arrow 1 indicates the slider responsible for the motion of the scissor blades, arrow 2 indicates the slider for the filter blade, and arrow 3 indicates the slider for the (unused) smaller aperture blade. b) The back of the shutter with solder points clearly visible. Arrows 1 indicate the solder points for the scissor blades, arrows 2 the solder points for the filter blade, and arrows 3, the solder points for the (unused) smaller aperture blade. c) The shutter inside its 3D printed enclosure, with an anodised disc to prevent backscattered light interfering with the experiment. The foil that covers the filter blade can be seen through the hole in this disc.

the terminals corresponding to this blade. We refer to the other sets of blades as ‘scissor blades’, which close in from either side of the aperture and overlap in the centre, and the ‘filter blade’, which is a larger blade that reaches all the way across the aperture and is made of a lightly tinted plastic, so only slightly attenuates the light. We modify the latter by adhering a strip of aluminium foil to its surface for heat dissipation, which enables us to block laser powers of up to 1 W. After connecting both of the solenoids corresponding to these sets of blades in series using wire, we secure the shutter in a mount. Images of the shutter before and after modification are shown in Fig. 4.11.

We designed the mount to interface with optics posts so it can easily be secured to the optical bench. The mount is 3D printed using PLA plastic, which is used to dampen vibrations. We enhance this effect by placing the shutter between two rubber ‘O’ rings, and we place a black anodised aluminium disc in the front of the mount before screwing it closed. This disc has a 4 mm diameter hole in its centre, and is used to reduce the amount of light backscattered from the aluminium foil and prevent the shutter being damaged by a misaligned laser beam.




Figure 4.12: The circuit diagram for the shutter driver. The design is based around an ATmega2560 microcontroller connected to a motor driver shield. We added a digital input, a manual override switch and an LED for each shutter. For simplicity, this diagram only shows the components for one shutter, but the motor shield facilitates the use of four shutters simultaneously.

The shutter driver is built using a microcontroller with a motor shield, which has four independent outputs. These outputs are each connected to a BNC connector via a switch so the shutters can be controlled manually from the driver unit (see Fig. 4.12 for the connectors used for one shutter). To control the shutters using the control computer, we attach BNC connectors to four of the microcontroller’s digital pins so that it can accept TTL inputs as triggers to open or close the shutters. When a TTL input or manual switch changes its state, the driver supplies a 100 ms burst of $\sim \pm 200$ mA of current to the corresponding shutter as being constantly driven would likely damage the shutters. For the correct current, a supply voltage of 5-6 V is required to operate both of a shutter’s solenoids in series. For multiple shutters to operate simultaneously without failure, we use an external supply to power them and the microcontroller.

We began testing on a single shutter by exposing it to a prolonged duration of moderate power. Parts of diaphragm shutters such as the ones we use normally




Figure 4.13: Transmission signals from a simple experiment designed to test the blades of our shutter. Both of these plots show results for the shutter blades (red) and filter blade (blue) a) opening and b) closing. The time on the x axis refers to the time after the TTL signal has been sent to the shutter driver.

melt or bend when exposed to powers of 50 mW or more. However, because we adhere aluminium foil to its surface to dissipate and reflect heat, after one hour of a 1 W, 1064 nm laser beam being incident on the closed shutter, the blades still work normally.

Next, we performed a simple experiment to test the performance of a single shutter, which consisted of guiding a laser beam through the shutter aperture into a photodiode. The photodiode was connected to an oscilloscope so the detected light levels could be recorded. A signal generator was connected to the oscilloscope and the TTL input of the shutter driver so both could be triggered simultaneously. For the purpose of this experiment, only one shutter blade was operated at any given time so they could be measured separately. We recorded and analysed 500 consecutive results and allowed the experiment to continue for a further 5×10^4 cycles, after which no degradation was detected. The normalised transmission results of these 500 cycles for each blade are shown in Figure 4.13.

The results of the experiment show that the scissor blades open and close faster than the filter blade; with an average opening duration from 5-95% transmission of $790(10) \mu\text{s}$ and a closing duration from 95-5% transmission of $573(7) \mu\text{s}$, compared to averaged opening and closing durations of $1.51(3) \text{ms}$ and $1.46(2) \text{ms}$




Figure 4.14: Histograms generated using the data featured in Fig. 4.13. These histograms show the distribution of a) half-opening and b) half-closing times of the filter blade, and c) half-opening and d) half-closing times of the scissor blades.

for the filter blade. The lower opening and closing durations from the scissor blades are expected because they have half the distance to travel, and they may be lighter than the filter blade with the aluminium foil strip adhered to it. This is also reflected in the delay from when a change in the TTL signal has been received to when the blades have blocked or revealed 5% of the light power. Upon a change in the TTL signal, the scissor blades took 2.73(2) ms to block 5% of the light, but only 2.29(2) ms to reveal 5% of it; the filter blade is slower once again, taking 3.71(3) ms to reveal 5% of the light, and 2.71(3) ms to block 5% of it. The discrepancies between opening and closing are possible if the beam was not completely centred, the filter blade was working against gravity in one direction and assisted in the other, or the amount of friction to be overcome when opening is greater than that when closing. Knowing that both sets of blades take

different amounts of time to begin opening and closing brings the opportunity to connect them to separate outputs of the driver and program to account for this so they both close simultaneously if this is desired. However, we elect to connect them both in series to a single output to allow a shutter driver to address more shutters, as the relative timings of our shutter blades are not critical.

The jitter in a shutter system is determined by measuring the standard deviation in our data, which is usually obtained using histograms (seen in Fig. 4.14) of the time taken to reach 50% transmission after the TTL signal changes. Unfortunately, there was a skew to our data which would affect obtaining the width parameter of the histogram (see Fig. 4.14). Instead, we simply found the standard deviations directly from our data. The filter blade has a jitter of $60\ \mu\text{s}$ when opening and $40\ \mu\text{s}$ when closing, and the scissor blades have jitter values of $21\ \mu\text{s}$ opening and $24\ \mu\text{s}$ closing. This gives a measure of both the reliability and durability of the shutter blades.

4.6.2 DDS controller

An optical lattice has been introduced to the experimental setup (more details about our optical lattice setup can be found in the thesis of Andrea Di Carli [58]). Optical lattices are periodic trapping potentials for atoms, as opposed to harmonic potentials, formed by standing waves of two or more interfering laser beams [110]. Some optical lattice experiments may involve ‘shaking’ the lattice, for use in studies of atom tunnelling and band structure. There are different ways of changing the depth or phase of the lattice potential. Our method involves using a Direct Digital Synthesiser (DDS) as an external frequency input for our AOM drivers and modulating the frequency.

Experimental Requirements The DDS is required to have two independent output channels, be capable of generating sine waves at frequencies of up to

110 MHz, and be able to sweep ± 200 kHz from a central frequency, at a rate of 1 kHz. We use an Analog Devices AD9958 DDS chip because it has two channels that are independent of one another in frequency, phase and amplitude which facilitates modulation of one beam while keeping the other constant. This chip outputs at frequencies of up to 200 MHz, with 0.12 Hz resolution which is also the smallest step size for frequency sweeps [111]. Depending on the range of a frequency sweep, the amount of time spent on a single step for this sweep can be somewhere between 8 ns and 2.064 μ s. The breakout board amplifies the signal generated by the chip to an amplitude of 2 V peak-to-peak.

Hardware Because there are many breakout board, driver and software variants for this chip, a microcontroller was deemed the best option to control it for consistency. Doing so meant it could be implemented in the current computer control structure without additional software requirements. An Arduino Due is used to drive the AD9958, because it works on a 3.3 V digital logic, and all other Arduinos use 5 V logic which would damage the chip.

The Arduino is connected to an Ethernet port (Funduino W5100) and the data pins of the DDS breakout board, both of which use the SPI (Serial Peripheral Interface) protocol to communicate, for which the Arduino only has one set of pins. As a result, the breakout board and the Ethernet port share SPI clock and data pins, which makes them susceptible to crosstalk, and corruption of data. The chip select pins on the Ethernet port and the DDS are used to select the device that should listen for data, and the software is separated into sections where either the Ethernet port or the DDS chip is listening.

Software The Arduino has three phases of behaviour - the listening phase, when it receives UDP packets via the Ethernet port; the parsing phase, when it interprets the UDP packets; and the execution phase, when it passes instructions to the DDS chip. When the Arduino is first switched on, it instructs the DDS to

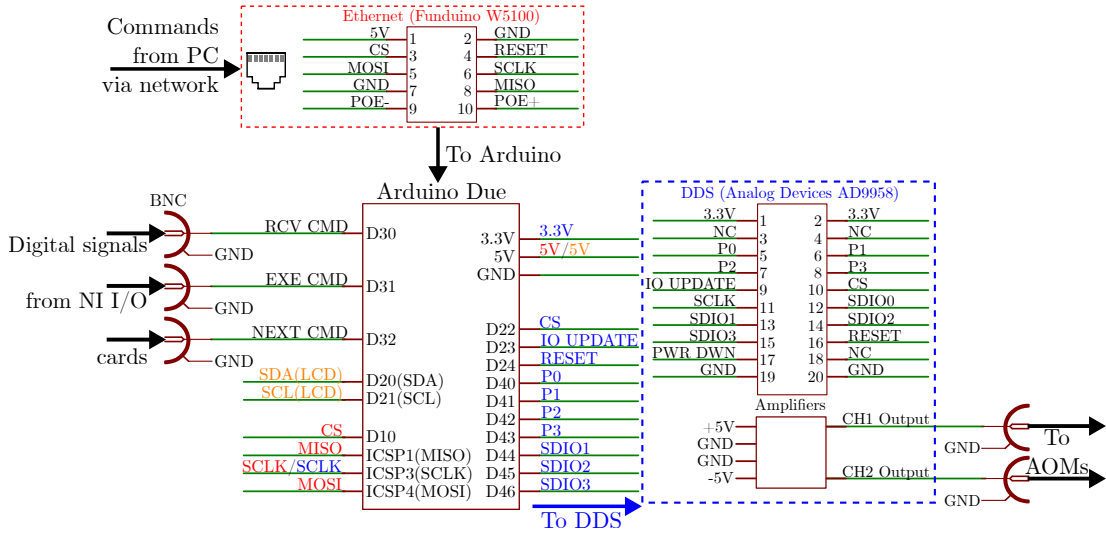


Figure 4.15: Pinout diagram of the Arduino Due, Ethernet port and DDS circuit board. The pins labels are colour coded, red for the Ethernet shield, blue for the DDS and orange for an LCD (not shown). The flow of data from the computer control system to the Arduino, to the DDS is also represented.

output at a ‘safe value’, so not to damage any connected AOM drivers.

When the RCV CMD pin (shown connected to a BNC in Fig. 4.15) is activated, the Arduino enters the listening phase. The Ethernet port is selected using its chip select pin, and the Arduino only listens for data sent to its IP address, usually consisting of text strings. When the RCV CMD pin is deactivated, the listening phase ends, the Arduino deselects the Ethernet port and interprets the UDP strings it has just received, separating each string down to its components. Each of the strings contains five parameters, each separated by a ‘#’ symbol, and ends with an ‘@’ symbol so the Arduino can interpret where each parameter begins and ends. The strings take the format

$$a\#b\#c.x\#d.y\#e.z\@ \quad (4.3)$$

where a is the frequency mode (constant or swept), b is the output channel, $c.x$ is the start frequency in MHz, $d.y$ is the end frequency in MHz and $e.z$ is the duration of the frequency sweep ($d.y$ and $e.z$ are always used in commands, but

are ignored during execution time if single frequency mode is selected). When the Arduino separates each parameter, it recasts them to the number format the respective value should take (integer for mode and channel, float for frequencies and durations), checks that they are within acceptable ranges, and then converts them to binary that the DDS can interpret. For instance, the frequency is converted to a ‘frequency word’ by

$$FW = \frac{2^{32} f}{f_{\text{clk}}} \quad (4.4)$$

where f is the desired frequency in Hz, f_{clk} is the clock rate (125 MHz by default), 2^{32} is the capacity of the DDS’s phase accumulator, and the frequency word has a value $0 \leq FW \leq 2^{31}$. All the converted values are stored in an array so that no time is spent processing parameters during the experimental sequence.

During the execution phase, the Arduino selects the DDS using its chip select pin and awaits digital triggers on the EXE CMD and NEXT CMD pins (shown connected to BNCs in Fig. 4.15). If the EXE CMD pin is activated, the command currently selected in the array is sent to the DDS and removed from the array, the next command being selected. If the NEXT CMD pin is activated, the selected command is removed from the array without being executed, and the next command is selected. In the event that a frequency sweep, rather than a constant frequency, has been executed the Arduino uses the profile pins of the DDS (marked P0, P1, P2 and P3 in Fig. 4.15) to determine the output channel on which the sweep occurs and in which direction. This can be repeated to ramp the frequency up and down, for shaking an optical lattice. If the last of the commands is skipped or executed, the DDS is instructed to output at its safe value, and the Arduino returns to the start of the program control sequence.

If the RCV CMD pin is activated at any point, the Arduino removes all the remaining commands from the list, instructs the DDS to output at its safe value, deselects the DDS and selects the Ethernet port to listen for more commands.

The code for the DDS driver is included in Appendix D.

Testing The Arduino was tested to confirm that it could receive commands, understand them and transfer them to the DDS chip successfully. Both output channels of the DDS were tested simultaneously to ensure that they output at the desired frequencies, that they were capable of performing the desired frequency sweeps, and that outputting a signal or performing a linear sweep on one output channel would not affect the other. These tests were usually performed at low frequencies for viewing on an oscilloscope screen, so that it could be seen whether the sweeps were phase coherent. Once we were confident that they were, we performed tests at higher frequencies using an rf spectrum analyser. We observed that when we selected a frequency sweep duration, that it took the requested amount of time to perform the sweep. This was also partially a test of the Arduino - there is an algorithm in the code that was written to calculate the smoothest sweep possible given the parameters, i.e. making the frequency steps as small as possible. As an example, the range of time steps is between 8 ns and 2.048 μ s, the range of frequency steps is 0.12 Hz to the full range of the sweep. If we request a 200 kHz sweep in 1 ms, the sweep can not consist of 0.12 Hz steps of 8 ns in duration, as the sweep would take 13.33 ms to complete. Instead, the step size must be increased to 1.6 Hz to accommodate the large sweep range. The Arduino was successful in calculating this need, and the DDS carried these sweeps out reliably. There were frequency sweeps that could not be tested, as the sweep could take less time than the refresh rate of the spectrum analyser, but there are no reasons to suspect that the DDS would not perform them reasonably well.

All the components of the control system described in this chapter, are assembled, configured and programmed to operate synchronously during each experimental sequence. This results in the reliable production of Bose-Einstein condensates as is shown in the following chapter.

Chapter 5

A Bose-Einstein Condensate of Caesium

In the previous Chapters, I discussed the theory of cooling atoms from room temperature to degeneracy, the properties of Bose-Einstein Condensates (BECs), and the physical apparatus which facilitates the methods discussed. This Chapter connects these concepts, presenting experimental results obtained using the apparatus, and using the theory to discover the properties of the gas we observe.

In Section 5.1, I describe some of the basic experiments we perform to measure the properties of the atoms and their environment. I detail our method of detecting the atom cloud and observing its spatial density distribution as well as methods to measure the conditions at the position of the atoms. In Section 5.2 I present results for each stage of the cooling process. These results reveal how the atoms behave at several checkpoints in the cooling sequence, and serve as optimisation benchmarks. Finally, in Section 5.3 I show initial results for our BEC, with emphasis on our current thermometry capabilities in regimes where both thermal and condensed atoms are present, and our ability to reproducibly control the atom number in our condensate.

5.1 Diagnostic Procedures

By imaging the atoms and determining their internal states, we can measure the magnetic field at their position, and the powers and widths of the beams that trap them. This Section details how we perform these measurements, and how we perform temperature measurements on the atoms in the thermal and partly condensed temperature regimes.

5.1.1 Absorption Imaging

In order to determine the properties of a gas of ultracold atoms either in-situ or after some time of flight, it is crucial to image the spatial density distribution of that gas. This density distribution can be observed by absorption imaging, which is performed by shining a resonant laser beam onto an atom cloud [52]. The atoms will absorb some of the light and scatter it elsewhere, casting a ‘shadow’ in the beam, which is directed into the sensor of a CCD camera. If the beam is propagating along the y axis, the amount of light that reaches the CCD sensor at each position on the x, z plane is given by the Beer-Lambert law:

$$I(x, z) = I_0(x, z)e^{-\tilde{n}(x, z)\sigma}, \quad (5.1)$$

where $\sigma = 3\lambda^2/2\pi$ is the absorption cross section, and $\tilde{n}(x, z) = \int n(x, y, z)dy$ is the column density of the sample at the position (x, z) . A map of the spatial density distribution of the atoms is obtained by imaging the laser beam both with $(I(x, z))$ and without $(I_0(x, z))$ the atoms present, and taking the ratio of these images. This technique requires that the intensity of the light is less than the saturation intensity of the atomic transition so the exponential decay of equation 5.1 applies. As a result, the sample must be imaged within a range of optical densities that allows for a good signal-to-noise ratio, i.e. the atoms are detected but not all of the light is absorbed. The effects of fringes and shot-to-




Figure 5.1: a) An image of the beam incident on the camera sensor. b) An image of the same beam partially blocked by a cloud of atoms. c) The result of dividing image b) by image a) and extrapolating for the column density \tilde{n} . Not shown is the image used as background, to subtract from images a and b before they are divided. All images are rotated so the z axis of the images perfectly matches the direction of gravity.

shot fluctuations observed in individual images (Fig. 5.1 a) and b)) are reduced in the calculation of the final absorption image (Fig. 5.1 c)). As well as giving us the density of the sample, by summing over all pixels on the image, we can also obtain the total atom number. Further properties of a thermal cloud can be determined by using a Gaussian fit function for the sum of the columns or rows for the x or y axes, respectively.

5.1.2 Temperature Measurement

Time-of-flight thermometry, first described in 1988 [67, 112], is the most widely used technique to perform temperature measurements on cold atom clouds. The temperature of an atom cloud can be measured by obtaining its width from fits to its 1D density profile over a range of expansion durations. Once a Gaussian of the form

$$f(x) = a \exp \left[- \left(\frac{x - b}{c} \right)^2 \right] \quad (5.2)$$

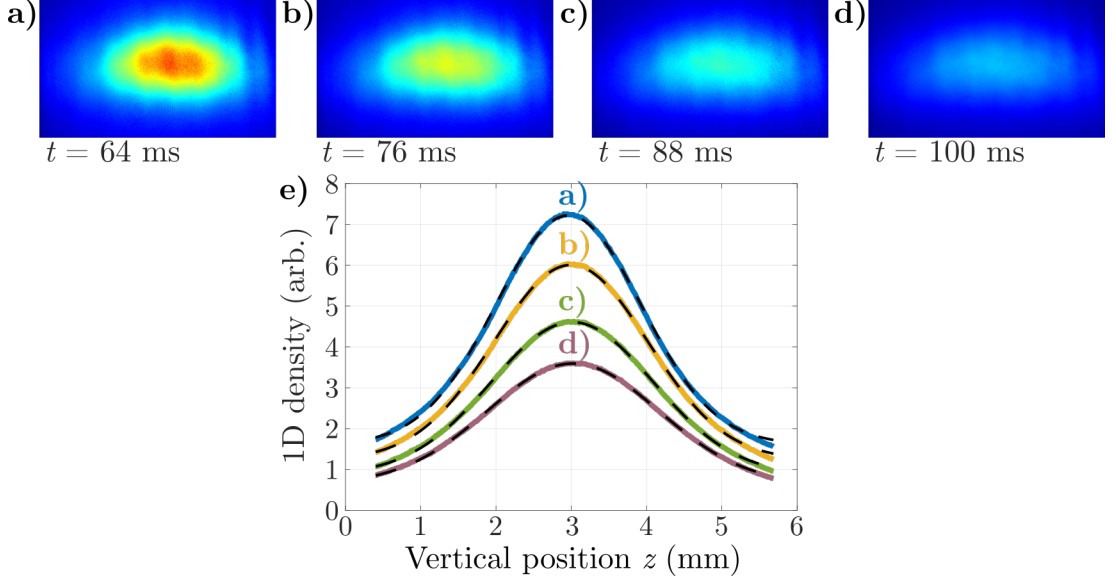


Figure 5.2: An example temperature measurement. Images a) - d) are images of an expanding cloud after 4 ms of Raman cooling after different expansion times. The plot shows the profiles of each of these clouds in the z -axis, the Gaussian fits represented by black dotted lines are used to obtain the $1/e$ cloud width. The widths are a) $1371(3) \mu\text{m}$, b) $1462(3) \mu\text{m}$, c) $1535(2) \mu\text{m}$ and d) $1616(2) \mu\text{m}$. These widths are used to measure the temperature of the sample in Fig. 5.12.

has been fit to the imaged cloud, its width, c , is extracted. By varying the time-of-flight or expansion time duration while keeping all other parameters constant, we can measure the rate of expansion of the cloud and extrapolate its temperature. Because we know the atoms spread out ballistically with a Maxwell-Boltzmann velocity distribution, the width of the cloud in the axis of expansion i at time t is given by

$$\sigma_i(t) = \sqrt{\sigma_i(0)^2 + \frac{k_B T_i}{m} t^2}, \quad (5.3)$$

where $\sigma_i(0)$ is the width of the cloud at release. Since the temperature depends on the rate of expansion, the $\sigma_i(0)$ term can be neglected. Rearranging equation 5.3 for T , we have

$$T_i = \frac{\sigma_i(t)^2}{t^2} \frac{m}{k_B}, \quad (5.4)$$

which makes it trivial to measure the temperature by plotting the width against

the expansion time and taking the gradient of the line. An example of one such temperature measurement is shown in Fig. 5.2.

5.1.3 Bimodal Fitting

As discussed in Section 3.4.2, once the cloud has been cooled to temperatures below the critical temperature for Bose-Einstein condensation, T_C , the density profile of the cloud changes. If we attempt to fit a Gaussian to the 1D density profile of a cloud below T_C , we find that it does not fit as closely as it did with thermal atoms because of a new, more pronounced peak (shown in Fig. 5.3). The column density profile is given by Eq. 3.26. Optimising this equation for fitting to time-of-flight images, we combine most of the variables to form free-fitting parameters, resulting in a simplified expression

$$\tilde{n}(x, z, t) = Ag_2 \left(e^{1 - \frac{x^2}{\sigma_x^2} - \frac{z^2}{\sigma_z^2}} \right) + B \max \left(1 - \frac{x^2}{R_r^2} - \frac{z^2}{R_{ax}^2}, 0 \right)^{3/2} \quad (5.5)$$

where $A = \frac{1}{\lambda_{dB}^3} \frac{1}{(\bar{\omega}t)^3} \sqrt{\frac{2\pi k_B T}{m}} t^2$, $\sigma_i = \frac{k_B T}{m} t^2$, $B = \tilde{n}_0(0)$, R_r and R_{ax} are all free fitting parameters. Because there are so many free parameters, we need to analyse the different components separately to obtain starting parameters and upper and lower bounds. The procedure for the 1D fit is as follows:

Fit the condensate Because the parabolic condensate portion of the density profile is responsible only for the peak, this makes a good starting point for our bimodal fit. A parabola is fit to the top 50% of the 1D profile, the width of this parabola where it intercepts the x axis is taken to be the initial guess for the Thomas-Fermi radius. Everything in this fit outside the Thomas-Fermi radius is set to zero as described by Eq. 3.17. This stage of the fitting procedure is shown in Fig. 5.3b).




Figure 5.3: The steps of the fitting procedure described in the main text using an absorption image. a) Absorption image of a BEC close to T_C . b) Initial fit of the parabola to the top 50% of the profile. c) Bose-enhanced Gaussian fit to the data outside the Thomas-Fermi radius determined by the parabola. d) Bimodal fit using the results of the previous fits.

Fit the thermal atoms As mentioned before, the thermal atoms take a Bose-enhanced Gaussian form which can be fitted by taking only the data outside the Thomas-Fermi radius into account. To reduce computation time, rather than using a polylogarithm summed to infinity, we sum only to 100. The results of this approximation are remarkably close to those of the infinite sum, and takes seconds to compute rather than minutes. From this fit (seen in Fig. 5.3c)), we can obtain parameters describing the thermal component of the cloud, allowing us to measure the temperature of the gas once more. Additionally, we can use the results of both the fits we have performed so far to determine the condensate fraction in the cloud. We take the parameters of both the fit to the condensate and the fit to the thermal fraction as starting parameters for the bimodal fit encompassing both components.

Fit the bimodal Now we have all the starting parameters we need to apply the full bimodal fit using Eq. 5.5. The fitting procedure is illustrated in Fig. 5.3, where all the steps of the fitting procedure are demonstrated. The 1D procedure is performed both for the horizontal and vertical axes of the image and the equations of the bimodal fits are used as a starting point for the two dimensional bimodal fitting procedure. The two profiles are used to generate a 2D image of a BEC, then a reiterative nonlinear least squares fitting procedure takes place




Figure 5.4: Example 2D fit of bimodal density profiles a) The original absorption image of the Bose gas, b) the final two-dimensional fit to the gas, c) the residuals between the fit and the original image. d) Density profile of the original image integrated along the vertical axis (solid black line), and the profile of the 2D fit, also integrated (red dotted line).

to minimise the error of the 2D fit with respect to the original image. The final two-dimensional bimodal fit is demonstrated in Fig. 5.4, alongside the original image and residuals of the fit. Because imaging noise is not taken into account in our model, the residuals are thought to reflect the speckle pattern on the original image. Since there are no prominent regions that appear in the error of the cloud, it is considered that our model closely matches what we image, so the properties of both the thermal and condensed components of the gas can be determined.

5.1.4 Microwave Spectroscopy

Microwave spectroscopy is the technique we use to calibrate the magnetic field experienced by the atoms throughout various stages of our experiment.

The two ground states of caesium are separated by 9.19 GHz, so a microwave of this frequency can flip the spin of the atoms from $F=3 \rightarrow F=4$ and vice versa. Each ground state is split into multiple Zeeman levels, which are degenerate when the atoms experience no magnetic field. In the presence of a magnetic field, this degeneracy is lifted, and each level is shifted by

$$\Delta E = h\Delta f = \mu_B g_F m_F B_z, \quad (5.6)$$

where μ_B is the Bohr magneton, g_F is the Landé g-factor for the hyperfine state being probed, m_F is the Zeeman level being shifted, B_z is the magnetic field in the position of the atom on the z axis and Δf is the detuning from the 9.19 GHz resonance. The Landé g-factor is $g_F = -1/4$ for $F=3$, and $g_F = 1/4$ for $F=4$, resulting in the corresponding m_F levels being shifted by the same amount but in opposite directions given the same magnetic field strength.

Given that both $m_F=0$ levels remain 9.19 GHz apart, and each subsequent level is split by $\pm 0.35 \text{ MHz } m_F / \text{G}$, it is possible to probe the magnetic field experienced by the atoms by scanning over a range of frequencies centred around 9.19 GHz, measuring the spacing between the shifted levels. This is illustrated in Fig. 5.5 a) and b).

We begin the microwave spectroscopy by pumping the atoms into the $F=3$ ground state before using the microwave so that if the microwave is resonant with a transition, some of our atoms then populate the $F=4$ ground state. During our normal imaging procedure, we image $F=4$ atoms, so must use repump light to transfer atoms from the $F=3$ ground state to the $F=4$ ground state, however for microwave spectroscopy we use no repump light. This ensures that we only

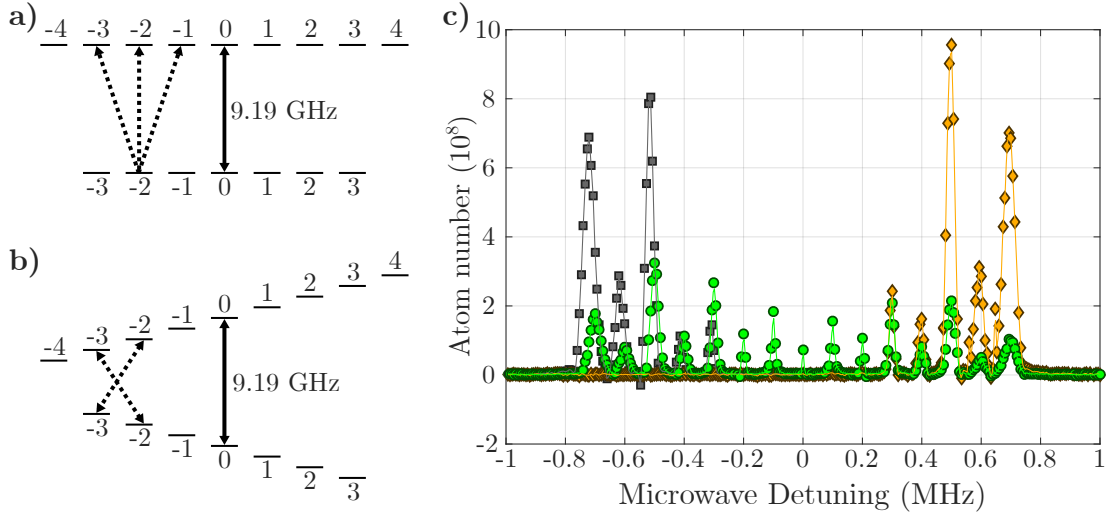


Figure 5.5: An illustration of the principle behind microwave spectroscopy. a) The m_F levels of the ground states of caesium in zero magnetic field. The two ground states are separated by 9.19 GHz, and all the m_F levels in each state are degenerate. The levels occupied after a microwave transition are determined by the selection rule $\Delta m_F = 0, \pm 1$ (an example being indicated by the dotted arrows). b) The m_F levels of the caesium ground states in the presence of a magnetic field. The spacing between the ground states stays the same for the $m_F = 0 \rightarrow m_F = 0$ transition, but is reduced for $m_F < 0$ transitions, and increased for $m_F > 0$ transitions. However, some degeneracies still exist (an example being indicated by the dotted arrows). c) Atom number vs detuning from the 9.19 GHz resonance, with different polarisations of light used to pump the atoms - linear polarisation (green circles), and circular polarisations stimulating σ^- (grey squares) and σ^+ (yellow diamonds) transitions. This measurement indicates that the atoms experience a magnetic field of 0.29 G.

image the atoms that have undergone the microwave transition. This method can also be used to determine the polarisation of the beam used to pump the atoms into the $F = 3$ state - if the polarisation induces σ^+ transitions, we will see a much higher population of the $|3, 2\rangle$ and $|3, 3\rangle$ states than if we use linearly polarised light.

Such a measurement is included in Fig. 5.5 c), where it can be observed that the peaks in atom number due to each microwave transition occur approximately 100 kHz apart, corresponding to a magnetic field of 0.29 G. In the linearly polarised result (green circles), peaks corresponding to all transitions can be ob-

served, indicating approximately equal occupations in all states. When circularly polarised light is used, in the case of σ^+ transitions (yellow diamonds), the high m_F states are more populated so the peaks due to the $m_F = 2 \rightarrow m'_F = 3$, $m_F = 3 \rightarrow m'_F = 2$ and $m_F = 3 \rightarrow m'_F = 4$ transitions are the most pronounced, with a smaller peak for the $m_F = 3 \rightarrow m'_F = 3$ transition. This effect is mirrored for the case of the light causing σ^- transitions (grey squares), where the lowest m_F states are populated.

5.1.5 Optical Trap Frequency

The dipole trap frequency, calculated using Eq. 2.22, can be observed experimentally by perturbing an atom cloud while it is trapped, and tracking its centre of mass position against hold duration in the trap. The perturbation is usually applied by adjusting the magnetic field gradient on the axis we wish to measure using shim coils, and quickly restoring it to its original value, causing centre-of-mass oscillations.

As well as allowing us to more accurately calculate the phase space density of our sample, the trap frequency can also allow us to better learn the trapping beam parameters. If the frequency of light used to trap the atoms is well known, by knowing either the optical power used or the beam waist, measuring the trap frequency allows us to determine the other.

Fig. 5.6 shows the variation in the centre of mass position when horizontal oscillations are induced, where we have used a 5 mW beam, and the damped sine fit applied to the data gives a frequency of 6.97(4) Hz, corresponding to a beam waist of 101 μm . This analysis was performed under the assumption that there were no other forces acting on the atoms, however magnetic levitation was active throughout, which inflicts an anti-trapping potential in the horizontal plane and lowers the total trap frequency experienced by the atoms. The stated trap frequency of 6.97(4) Hz is actually the total frequency of the trap and the




Figure 5.6: A measurement of trap frequency in a vertical dipole beam containing 5 mW of optical power, giving a value of 6.97(4) Hz.

optical trap itself has a higher frequency, so the calculated beam waist value is an overestimation. The following Section describes the analysis taking the anti-trapping effect into account in pursuit of more accurate measurements of optical trap frequency.

5.1.6 Levitation Gradient Curvature

As illustrated in Fig. 2.6, the magnetic levitation field has a curvature which applies an anti-trapping potential on the atoms. This levitation curvature has an anti-trap frequency which competes against the dipole trap frequency. The anti-trap frequency, calculated in Eq. 2.7, depends on the magnetic field offset as $\alpha \propto 1/\sqrt{B_0}$. With a magnetic field offset of 17.2 G at the position of the atoms, the calculated anti-trap frequency is $\alpha = 2\pi \times 3.35$ Hz.

An experiment was performed to determine the anti-trapping frequency of the levitation gradient, using the same method as described in Section 5.1.5 and varying the power of the trapping beam. The atoms were levitated inside a vertical dipole beam at different optical powers, and a centre-of-mass oscillation




Figure 5.7: A comparison of trap frequency vs trap power with and without the anti-trapping potential. The blue squares show results of trap frequency measurements in a vertically aligned beam with magnetic levitation active. Equation 2.22 is used to indicate the trap frequencies that would be expected without levitation (dashed line), and is fitted to the experimental data with an offset (solid line). This fit indicates that the anti-trapping frequency $\alpha = 2\pi \times 3.63(17)$ Hz, and that the dipole beam balances this trapping frequency at a power of 1.05 mW.

was induced. The results of the experiment are presented in Fig. 5.7. Because we measure the total trap frequency which includes the anti-trapping potential, the measured frequencies diverge from expected values for lower optical powers. We still observe the \sqrt{P} dependence on the trap frequency from Eq. 2.22, with a power offset. The anti-trapping frequency is found by calculating the expected dipole trap frequency for the power at which the experimentally determined total trap frequency decreases to 0 Hz.

These results indicate an anti-trapping frequency of $2\pi \times 3.63(17)$ Hz, in good agreement with the previous calculation, and that this anti-trapping frequency can be cancelled by the trapping frequency of the beam at a power of 1.05 mW, corresponding to a beam waist of 95 μm .

5.2 Route to a Bose-Einstein condensate

This Section describes our cooling scheme to generate a BEC of caesium atoms. We first laser cool our atoms and confine them in a cigar-shaped 2D magneto-optical trap, which is used as a cold atom source for the main experimental chamber. We catch the atoms in a 3D MOT, then we perform degenerate Raman sideband cooling before magnetically levitating them. The levitated atoms are captured in a dipole trap with ‘dimple’ beams through its centre, the larger beams serving as a reservoir for the smaller dimple beams as their intensities are ramped down. Finally, the dimple beams are ramped down, evaporating the hottest atoms and resulting in a BEC.

5.2.1 2D+ MOT

When we load the main experimental chamber, it is important to reduce the rate of atom loss from as early as possible. One of the methods used to achieve this is the pre-cooling of atoms from room temperature to temperatures on the order of mK so they are more likely to be caught by the 3D MOT after being accelerated through the vacuum apparatus. For this purpose, we use a 2D+ MOT because it can load atoms from the background gas and increase its density, and it can be treated as a high-flux cold atom source for the experiment.

As with traditional 2D MOTs, the trap is formed using a combination of laser beams and magnetic field gradients in two dimensions (see Fig. 5.8), trapping atoms along the centre of the transverse axes of the glass cell. The added benefit to a 2D+ MOT, is that laser beams in the third dimension provide additional cooling in the form of optical molasses. These features are shown in Fig. 5.8, where we can see the arrangement of optics that allow the MOT to form.

The optics used for the 2D+ MOT are arranged linearly and mounted to aluminium brackets designed for compactness and simplicity of alignment (see

Fig. 5.8). Optical fibers are secured to two of the brackets, each delivering 140 mW of optical power - 137 mW is 4Γ detuned from the $F = 4 \rightarrow F' = 5$ cooling transition of caesium, and 3 mW is tuned to the $F = 3 \rightarrow F' = 4$ repump transition.




Figure 5.8: a) A computer generated render of the 2D+ MOT setup, top-down, including the push beam. Vertical beams are not shown. b) A diagram of the beam polarisation and magnetic field configuration in the transverse plane of the 2D MOT.

The beam is collimated to a 25 mm diameter in a cage-mount system. Three half-wave plates and three beamsplitter cubes are also mounted to each bracket so that the beam can be split into three equal components on the horizontal transverse axis, and three on the vertical. Mirrors are secured in two more brackets, each mirror facing an exit port of the bracket containing beamsplitter cubes, so the beams are retro-reflected over themselves with complete overlap. Within each exit port of all four brackets, there is also a quarter-wave plate to produce the circular polarisations required for trapping in all beams. The result is that there are three counter-propagating pairs of beams propagating horizontally and three propagating vertically, each set overlapping in the centre of the trap. Because these are aligned with the magnetic field gradients, these are called ‘trapping’ beams. The magnetic field gradient is empirically set to ~ 5 G/cm.

One last 25 mm diameter beam, with 2.7 mW optical power, propagates along

the longitudinal axis of the trap, coinciding with the three overlap volumes from the trapping beams. This is reflected from inside the cell using a gold-plated mirror, and reflected back over itself, forming optical molasses in one dimension. This is referred to as the ‘plus’ beam. Through its centre is an additional 2.5 mm diameter beam with 430 μW of power which is aimed directly through the hole in the centre of the gold mirror, through the differential pumping tube, into the main experimental cell. This is the ‘push’ beam, which accelerates atoms to the main experimental chamber once they have been cooled.

5.2.2 3D MOT

The 3D MOT catches atoms that have been transferred from the 2D+ MOT, and cools them further so they can be contained in weaker traps for sub-Doppler cooling. The 3D MOT also confines the atoms in the centre of the glass cell, at a position where all subsequent cooling and trapping beams will coincide.

The configuration of laser beams and magnetic fields that forms the 3D MOT is similar to that of the 2D+ MOT, but with a magnetic field gradient in three dimensions. This results in trapping in 3D, forming an almost spherical trap rather than a cigar-shaped one. We accomplish this using the main experimental coils, with six beams overlapping directly between the centre of the coils. Using six separate beams rather than retro-reflecting three beams has the advantage that the forces from all beams on the atoms can be evenly balanced. The cooling beams are each 25 mm in diameter with 80 mW of optical power resulting in a peak intensity of $29.6 I/I_{\text{SAT}}$. The light is 20 MHz ($\sim 4\Gamma$) red-detuned from the $F = 4 \rightarrow F' = 5$ transition, and the magnetic field gradient is 12 G/cm. After catching atoms for one second under these conditions, the 3D MOT contains on the order of 6×10^9 atoms, but we allow for a loading duration of 3 s for a higher atom number (although the atom number does not seem to be increasing in Fig. 5.9 because very little imaging light escapes the atom cloud). After 1 s

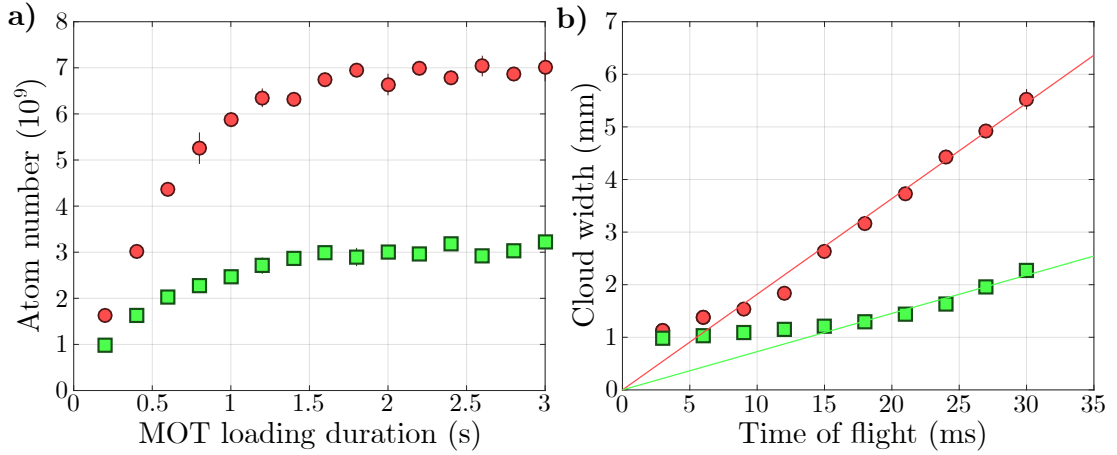


Figure 5.9: Loading curves and temperatures in the 3D MOT. a) Atom number vs loading duration with compression (green squares) and without compression (red circles). b) Cloud width vs expansion time after 100 ms of MOT loading with compression (green squares) and without compression (red circles). The solid lines illustrate fit results using Eq. 5.3 to extract the temperature ($264(3) \mu\text{K}$ without compression, $42(1) \mu\text{K}$ with compression). The points for times of flight below 15 ms were excluded from the fit because the samples were optically dense, rendering the results unreliable.

loading duration and 15 ms time of flight, the cloud has a size of $\sigma_x = 4.06(5) \text{ mm}$, $\sigma_z = 6.15(22) \text{ mm}$ before compression, and $\sigma_x = 2.25(15) \text{ mm}$, $\sigma_z = 2.18(20) \text{ mm}$ after compression.

A compression phase is implemented in our 3D MOT process to increase the phase space density of our gas before Raman sideband cooling. Raman cooling does not increase the density of the atoms as it cools, and the Raman beams are narrower than the cloud after 3 s of MOT loading. This problem is exacerbated by the fact we switch off MOT beams and magnetic fields for 2 ms before ramping up Raman cooling beams, during which time the cloud will expand. To reduce the number of atoms that will be lost between processes, the cloud needs to be compressed. Typically, compression is caused by an increase in magnetic field gradient, however we reduce ours to 10.5 G/cm to avoid potential complications caused by eddy currents when we begin Raman cooling. The detuning of our MOT beams is increased to 115 MHz and power is reduced to 50 mW per beam,

which reduces the rate of reabsorption in the gas. Usually compression is a trade-off between high density and low temperature; our compression was optimised for lower temperatures while keeping a high atom number. The reduction in temperature is possibly due to the loss of the hottest atoms from the trap as the density of the gas increases.

The loading curves for the MOT without compression, and with 50 ms of compression can be seen in Fig. 5.9. The measured initial load rate without compression is $7 \times 10^9 \text{ s}^{-1}$. This initial load rate does not include atoms that were not caught in the 3D MOT, so only gives a lower bound on the atom flux from the 2D+ MOT. After one second loading duration the atom number appears to plateau, but it is in fact still increasing - the cloud is so dense that it prevents any light from reaching some of the pixels on the camera sensor, so increases in density cannot be detected. We also performed temperature measurements comparing the MOT with and without compression. Because we optimised our compression for low temperatures rather than maximum atom density, we achieve temperatures at a fraction of our 3D MOT without compression - without compression, we reach a temperature of $264(3) \mu\text{K}$, or $42(1) \mu\text{K}$ with compression. By the end of the MOT sequence the atoms are mostly in the $F=4$ state due to the repump light used, and in a mixture of Zeeman states, ready to undergo degenerate Raman sideband cooling.

5.2.3 Degenerate Raman Sideband Cooling

Degenerate Raman sideband cooling is used to cool our atoms so that their energies are comparable to dipole trap depths we can produce, with the added advantage of polarising the atoms into a high-field seeking state, enabling magnetic levitation.

The Raman cooling beams are configured as in Refs [18, 72, 113] with four lattice beams and a single polarising beam, as shown in Fig. 5.10. Our lattice




Figure 5.10: A diagram of the beam geometry used during degenerate Raman sideband cooling. The laser beams that form the Raman lattice are represented in red, and the polariser beam in blue. The white arrows indicate the propagation direction of each beam.

beams are tuned to the $F = 4 \rightarrow F' = 4$ resonance, which pumps the atoms into the $F = 3$ state and provides tight confinement for, and couples the degenerate m_F levels of the atoms in the $F = 3$ state, as described in Section 2.3. These beams are linearly ramped up to 25 mW in 2 ms, because switching them on suddenly could cause heating in the sample.

The polarising beam is 7.3 MHz blue-detuned from the $F = 3 \rightarrow F' = 2$ transition, which means the atoms tend to decay into $F = 3$ with a higher m_F state, with their vibrational states conserved. This beam contains 2 mW of optical power, polarised mostly circularly with a small π component to perform the cooling cycle as described in Section 2.3.

Because only a minor energy shift of the m_F states is required for Raman cooling, a relatively small magnetic field is generated. Shim coils are used to produce this magnetic field, and fine-tune the quantisation axis to ensure the correct ratio of σ^+ transitions to π transitions. We examined the effects of magnetic field direction on the atom number during Raman cooling (see Fig. 5.11) by adjusting the current supplied to a pair of shim coils affecting horizontal magnetic fields. We can see that the number of atoms that survive 3 ms of Raman cooling is highly




Figure 5.11: Atom number in the $|3, 3\rangle$ state (red squares) and the $|3, 2\rangle$ state (green squares) after 3 ms Raman cooling vs the current supplied to the horizontal shim coils at either side of the experimental chamber. Adjusting the shim coil current affects the direction of the magnetic field, which determines the ratio of σ^+ to π transitions.

dependent on the quantisation axis of the magnetic field used because this alters the ratio of σ^+ transitions to σ^- and π transitions. At the maximum in atom number in the $|3, 3\rangle$ state, there is also a maximum in the number of atoms in the $|3, 2\rangle$ state. This is possibly due to a higher atom density in the Raman lattice, which leads to higher levels of reabsorption of scattered light. We optimised this across all shim coils for a maximum of $|3, 3\rangle$ atoms in the cloud.

Results of our Raman sideband cooling after shim coil optimisation can be seen in Fig. 5.12, from which we optimised the cooling process for high atom number and low temperature. For our temperature measurements we levitated the $|3, 3\rangle$ atoms and tracked the expansion of the cloud over various durations. The results for lower expansion durations are somewhat problematic because there is a smaller cloud of $|3, 2\rangle$ atoms falling from the main cloud at $1/3g$, overlapping with it. This deforms the 1D Gaussian profiles of the cloud rendering a fit to the profile impossible until after the clouds have separated (see Fig. 5.13). We apply temperature fits only to expansion durations exceeding 60 ms to ensure more reliable results.

We selected a cooling duration of 3 ms because it achieves a temperature

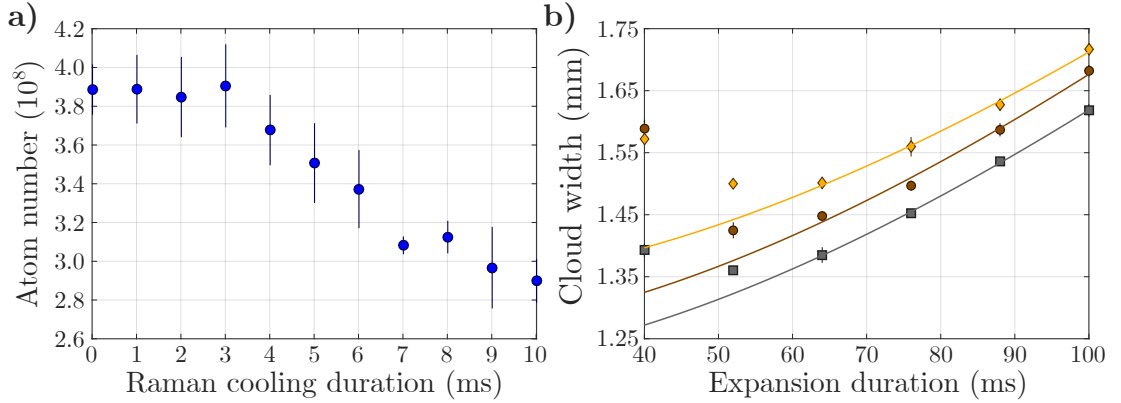


Figure 5.12: Atom number and temperature after degenerate Raman sideband cooling. a) Atom number vs Raman cooling duration. b) Cloud width vs expansion duration after 1 ms (brown circles), 4 ms (grey squares) and 10 ms (yellow diamonds) of Raman cooling. The solid lines illustrate our fit results using Eq. 5.3 to determine temperatures of 974(12) nK after 1 ms, 959(15) nK after 4 ms and 897(19) nK after 10 ms. We only fit to data points that occur after 60 ms expansion, because before then the cloud is optically dense and contains $|3, 2\rangle$ and $|3, 1\rangle$ atoms, yielding unreliable results.

below $1 \mu\text{K}$ without sacrificing many atoms. Longer durations mainly serve to diminish the atom number with only incremental improvements in temperature. The atom loss is thought to be due to collisions, reabsorption and occasional σ^- transitions, all of which result in heating the atoms, causing them to escape the trap. Reabsorption and collisions are both consequences of high atom density, however at this stage higher densities take priority over lower temperatures. A previous implementation of Raman sideband cooling reported 3D temperatures of 290 nK with an atom density of $1 \times 10^{10} \text{ cm}^{-3}$ [72], which is estimated to be a third of the densities we observe.

After Raman cooling is complete, the beams are ramped down over $600 \mu\text{s}$ to avoid the cloud expanding after suddenly being released. There are now $\sim 3.9 \times 10^8$ spin polarised atoms with a temperature of $\lesssim 1 \mu\text{K}$, corresponding to a phase space density of 9×10^{-5} , that are ready to be magnetically levitated.

5.2.4 Magnetic Levitation

Although the temperature of our atoms is now well below dipole trap depths we can produce, the lifetime of this trap would be extremely limited because the downward force due to gravity is much stronger than our dipole force. We use magnetic levitation to counteract the effects of gravity, which enhances the trap depth in the z direction and eliminates gravitational sag in the trap.

Levitating one Zeeman state brings with it some added benefits. Different spin states are separated using this method, so we levitate only atoms in the $|3, 3\rangle$ state and under-levitate all other $|3, m_F\rangle$ states, achieving spin purity in our cloud. In addition, we can observe our cloud over expansion durations far longer than simple time of flight measurements would permit because unlevitated atoms would fall from the field of view of our imaging system. In turn, this gives us the ability to measure atoms in momentum space and in very low energy regimes.

Magnetic levitation also presents some challenges. Because Raman cooling is very magnetically sensitive, beginning levitation before releasing the atoms from the Raman lattice is not an option. We instead need to quickly produce the required magnetic field offset and gradient after the atoms have been released, which cannot be done before the atoms have been accelerated by gravity. Our solution is to use a 400 V discharge capacitor which drives a section of our coils (the blue section in Fig. 4.8 c)). This produces 20% of the required gradient in 5 ms while the gradient coils ramp up. The discharge coils and gradient coils over-levitate the atoms for a brief period to bring the atoms to rest, before settling at the magnetic field gradient that levitates them.

The levitation potential also has a curvature (illustrated in Fig. 2.6) which can accelerate the expansion of the cloud in the (x, y) plane and cause atoms to fall if they spread too far from the centre of the curvature. Figure 5.13 shows an image of a cloud of atoms separating in the presence of a magnetic field gradient,




Figure 5.13: An absorption image of atom cloud separation after 40 ms of levitation following Raman sideband cooling. The top cloud is made up of atoms in the levitated $|3, 3\rangle$ state, the middle cloud of the $|3, 2\rangle$ state, and the faint bottom cloud of the $|3, 1\rangle$ state.

much like a Stern-Gerlach experiment. The separation of states increases with the expansion duration, so eventually only the $|3, 3\rangle$ cloud remains in a position where the reservoir beams will overlap.

5.2.5 Dipole Trap

As our atoms approach their recoil temperature, it becomes much more difficult to cool them further using near resonant light. While degenerate Raman sideband cooling has gotten to within a factor of two of the recoil temperature for caesium [59], the lower temperature limit is set by reabsorption. To cross the recoil temperature limit, we need to use far off-resonant light to provide conservative trapping forces. For this purpose, we have a 1070 nm, 200 W IPG ytterbium fiber laser that we use to form a crossed-beam dipole trap of two ~ 100 W ‘reservoir’ beams with overlapping foci. Each of these beams is focused down to a $1/e^2$ waist




Figure 5.14: A diagram of dipole trap and optical lattice (not described in main text) beam geometries. The reservoir beams are represented in red, the dimple beams in blue, and the lattice beams are signified by the yellow ‘pancake’ structure. White arrows indicate the propagation direction of each laser beam.

of $800\ \mu\text{m}$, resulting in a trap depth of $25.4\ \mu\text{K}$. A diagram of the geometry used for these beams is shown in Fig. 5.14.

In addition, we use light from a 1064 nm Coherent Mephisto Nd:YAG laser to inject light into a laser diode, which in turn provides seed light for a 50 W fiber amplifier for narrow ‘dimple’ trap beams. The Mephisto laser has a frequency linewidth of $< 10\ \text{kHz}$, which persists through the amplification process and is important for the lifetime of the trap. The so-called horizontal and vertical dimple beams overlap in the centre of the aforementioned crossed-beam dipole trap. The dimple beam waists at the position of the atoms are $50\ \mu\text{m}$ and $90\ \mu\text{m}$ respectively, so these beams are often set to have different optical powers to ensure similar trap frequencies in all axes. Typically, in the dipole trap the horizontal dimple power is set to 120 mW, and the vertical power 260 mW. These beams facilitate an order of magnitude increase in the phase space density of the gas, without affecting its temperature, which is a good starting point for evaporating to a Bose-Einstein condensate.

We begin by switching on the reservoir beams at 20 W (the minimum the laser controller will allow), before ramping them up to 200 W over 40 ms, so not




Figure 5.15: Lifetime and temperature measurements of atoms in dipole trap. a) Lifetime measurement of atoms in the dipole trap. There are two distinct rates of loss - a steep loss rate for the first 150 ms, which settles thereafter. The dotted line is a fit to the first 100 ms of hold time, and the solid line is a fit to the data from 200 ms onwards. The latter gives a trap $1/e$ lifetime of 788(73) ms. b) Temperature measurements of the atoms in the trap at 200 W after 50 ms (red squares), at 200 W after 200 ms (green squares), and having ramped down to 130 W over 1300 ms. The solid red, green and blue lines represent fits to the data used to extract the temperature of the cloud, giving temperatures of 8.95(9) μ K, 8.95(2) μ K, and 3.77(1) μ K, respectively.

to perturb the atoms, before introducing the dimple beams over 2 ms. We allow 100 ms for the required magnetic offset field to be generated for loading the dimple trap. Under normal circumstances we then ramp the reservoir beams down over the course of 1300 ms, but for the purpose of this experiment we kept them on full power for varying hold durations to determine the lifetime and the temperature of the atoms in this trap. The results of this experiment are found in Fig. 5.15. It appears that there are two different loss regimes in our lifetime measurement because for the first 150 ms, atoms that have not been caught in the dipole trap are still being imaged. These atoms dissipate quickly, resulting in a sharp drop in the atom number, and the losses observed afterwards are nominal for a dipole trap so a valid measure of lifetime can be obtained. The $1/e$ lifetime of our trap at 200 W is 788(73) ms, still relatively short for a dipole trap. This is possibly due to the fact that the IPG laser generates light in several modes, resulting in a large spectral linewidth (~ 3 nm), which grows for high powers [114]. At 200 W

this linewidth is presumed to be very broad, and the trap depth is much deeper than the temperature of the atoms, both of which can cause heating [115, 116]. However, this is not a cause for concern because this regime is short-lived in the experimental sequence.

We performed temperature measurements on the atoms after hold times of 50 ms and 200 ms (red and green squares in Fig. 5.15 b) at 200 W, a trap depth of $25.4 \mu\text{K}$, for which we measure respective atom numbers of 8.4×10^7 and 3.8×10^7 . Both of these measurements yield a temperature of $8.95 \mu\text{K}$. The temperature of atoms in the trap stays the same over long hold durations, because the temperature is limited by the finite trap depth. We can only reduce their temperature by ramping down the optical power used in the trap, which we elect to do as soon as possible after ramping up the dimple beams. We performed one more temperature measurement after reducing the reservoir beams to 130 W over 1300 ms. With 1.3×10^7 atoms left, we measure a temperature of $3.77(1) \mu\text{K}$. We usually end our dipole trapping procedure with a phase space density on the order of 10^{-4} .

5.2.6 Evaporative Cooling

Evaporative cooling is the most common approach to Bose-Einstein condensation. Where previously the goal was to lose as few atoms as possible, we now deliberately remove the hottest ones to stimulate the cooling of the cloud as a whole. Evaporative cooling is often performed using rf-evaporation which for some alkali atoms results in runaway evaporation due to increases in the atom density and collision rate. However, because of its scattering properties, the phase space densities achieved from rf-evaporation of caesium were only fractions of that required for Bose-Einstein condensation to occur. We use an all-optical approach which involves reducing the depth of our dipole trap in stages, starting with the evaporation parameters in [18], then optimising them empirically.

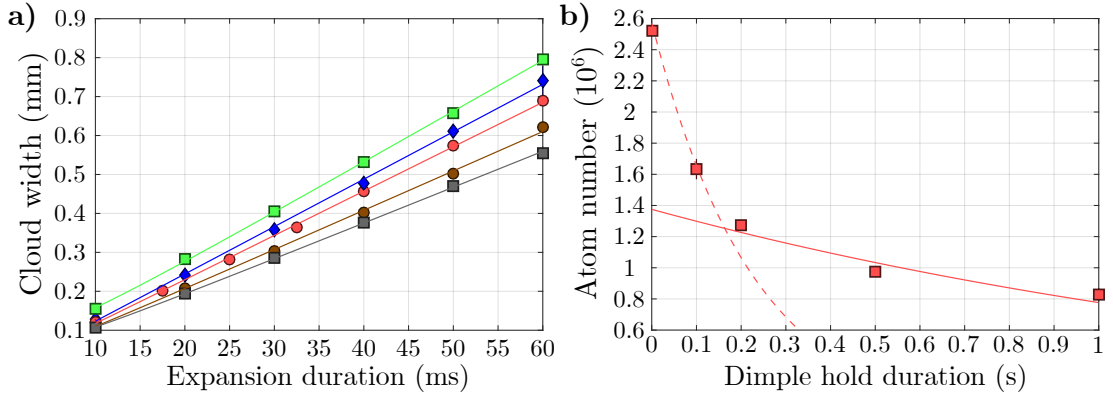


Figure 5.16: Characterisation of thermalisation in our dimple beams after the larger reservoir beams have been ramped down. a) Temperature measurements after various hold durations in the dimple beams. Green squares show a temperature measurement taken after 1 ms, blue diamonds after 100 ms, red circles after 200 ms, brown circles after 500 ms, and grey squares after 1 s. The solid line fits to these respective measurements indicate temperatures of $1.379(7) \mu\text{K}$, $1.19(1) \mu\text{K}$, $1.041(4) \mu\text{K}$, $823(5) \text{nK}$ and $688(6) \text{nK}$. b) Lifetime measurement of the dimple trap after the pre-ramp. The atom loss rate is very high for the first 100 ms as the atoms that were not loaded into the dimple beams dissipate. Both lines are exponential decay fits to the data - the dashed line fit tracks the initial loss of the atoms, and the solid line fit gives a measure of the $1/e$ trap lifetime, $1.76(24) \text{s}$. This lifetime is still somewhat short, thought to be because the high atom density incurs three-body losses.

For each of the four stages in our evaporative cooling process, we are able to tune the scattering length of the atoms and the trap depth at varying rates. To avoid heating or sloshing the atoms and to allow longer times for thermalisation, each step in the process is more gradual than the last, the full evaporation taking about 10 s. There are four main evaporation stages (pre-ramp and ramps 1-3). A timing diagram of ramps 1-3 and the changes of scattering length and beam powers that occur during these ramps is shown in Fig. 5.17.

Pre-ramp The main goal of this stage is to transfer atoms from the reservoir beams to the dimple beams. The reservoir beams are linearly reduced from 130 W to 20 W (the minimum allowed by the laser controller) in 1.4 s. During the first 50 ms of this sweep, the magnetic field offset current is ramped to 12.4 A, cor-

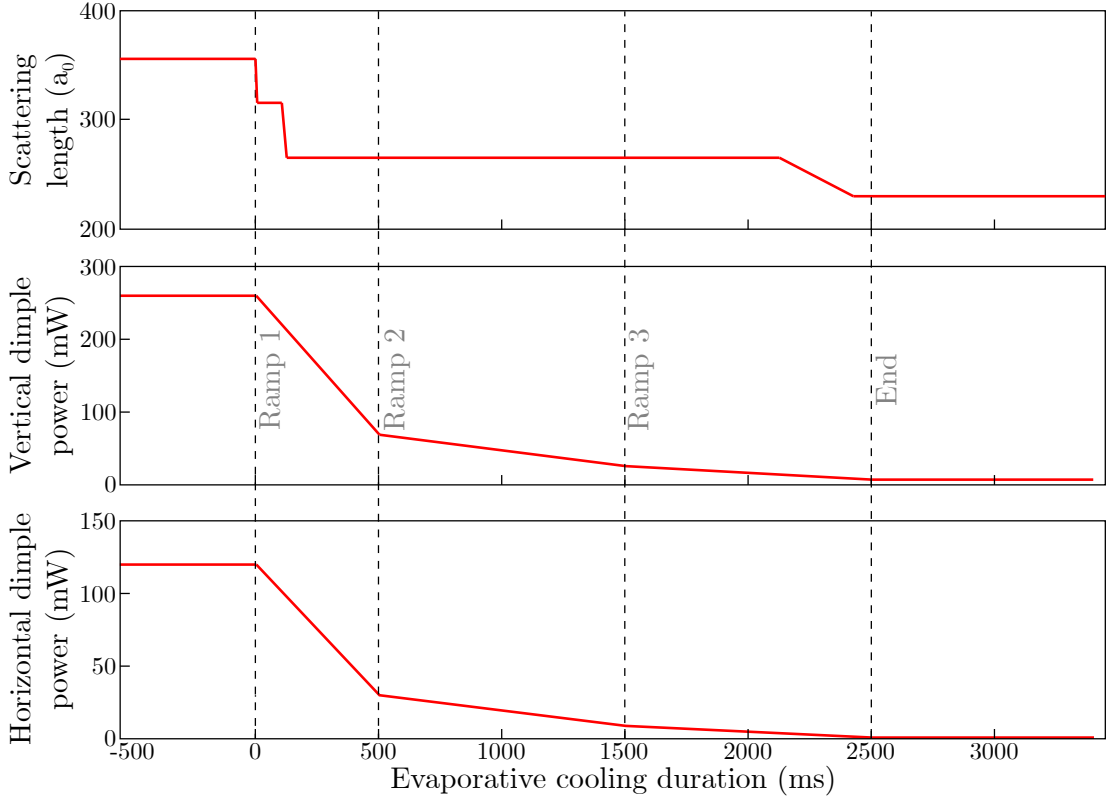


Figure 5.17: Timing diagram for the evaporative cooling sequence. Plots of the scattering length and horizontal and vertical dimple beam powers are shown, with dashed lines indicating the beginning of each stage.

responding to a magnetic field offset of ~ 23 G, to reduce the scattering length to $\sim 340 a_0$, which improves the loading of the dimple trap and facilitates fast thermalisation [17, 39]. Dimple beam powers are kept constant at 120 mW horizontal power, and 260 mW vertical. After the linear ramp, the scattering length is reduced to $300 a_0$.

A measurement of the dimple trap temperature and lifetime after the pre-ramp is shown in Fig. 5.16. We begin with 2.5×10^6 atoms and a trap depth of $5.2 \mu\text{K}$, but the atom number quickly drops and the temperature does not go far below $1 \mu\text{K}$ - after 1 s in this trap, only 8×10^5 atoms remain with a temperature of $688(6)$ nK. We must allow some time for plain evaporation, because continuing to lower our trap depth before the atoms thermalise will result in high levels

of loss. However, with peak atom densities on the order of 10^{12} cm^{-3} , waiting too long could also result in losses due to three-body collisions - we strike a balance with 100 ms hold duration before beginning evaporation ramp 1. Now that the reservoir beams have been switched off, it is more accurate to use the trap parameters to calculate the phase space density of the gas, from Eq. 3.8. After these 100 ms, before we begin evaporative cooling, the phase space density of the gas is 6.04×10^{-2} .

Ramp 1 - At this stage, it is crucial to reduce the atom density quickly to reduce three-body losses. The dimple beams are initially linearly ramped down to 30 mW horizontal and 70 mW vertical in 500 ms, the trap depth reduced to $1.4 \mu\text{K}$, after which 8×10^5 atoms remain at a temperature of $234.6(9) \text{ nK}$, as shown in Fig. 5.18. The scattering length is ramped to $250 a_0$ in the first 20 ms of the sweep, to limit possible 3 body losses. The phase space density of our sample is 0.53, we are now within an order of magnitude of the phase space density required for Bose-Einstein condensation.

Ramp 2 - Only the dimple beam powers are adjusted in this ramp. As we approach Bose-Einstein condensation, we need to reduce the rate at which we lower the trapping potential to account for a lower thermalisation rate due to a reduced collision rate, so we ramp at $\sim 10\%$ the rate of ramp 1. The beams are linearly ramped down to 10 mW and 26 mW in 1 s. We now have a trap depth of 500 nK , and 5×10^5 atoms left at a temperature of $62.7(4) \text{ nK}$ (see Fig. 5.18). With a phase space density of 3.72, we have finally surpassed the phase space density required for Bose-Einstein condensation to occur. A small condensate fraction in the cloud can now be observed (see Fig. 5.19). Although the definition of phase space density becomes less relevant once there are non-thermal atoms in the gas, it is still calculated for completeness, and so that an evaporation efficiency can be determined from the results (see Fig. 5.20).




Figure 5.18: Measurements for each stage in our evaporative cooling process. a) Temperature measurements after the first ramp (red circles), second ramp (green squares), and third ramp (blue diamonds). The latter result were obtained by bimodal fits, as the third evaporation ramp results in a condensate. The solid line temperature fits yield results of 234.6(9) nK, 62.7(4) nK and 14.0(3) nK. b) The atom number after each of the ramps detailed in a). Through evaporative cooling we lose an order of magnitude in our atom number from the dimple trap, but we cool the gas by a factor of 100, and increase its phase space density by a factor of 300.

Ramp 3 - Now that we observe a BEC, the goal is to condense as many of the remaining thermal atoms as possible. The dimple beams are ramped down once more at a further reduced rate ($\sim 15\%$ the ramp rate of ramp 2) to their final powers. Typically we sweep down linearly to 1 mW horizontal and 8 mW vertical in 3 s, giving a final trap depth of 100 nK. For the first 300 ms of this ramp, we change the magnetic field offset for a scattering length of approximately $210 a_0$. We now have a cloud of 2.6×10^5 atoms at 14.0(3) nK (see Fig. 5.18), a BEC with a condensate fraction $N/N_0 = 0.48$. While it is possible to ramp the trap down further for lower temperatures, we would do so at the expense of atom number for minimal gain.

Absorption images taken after each of the three evaporation ramps are shown

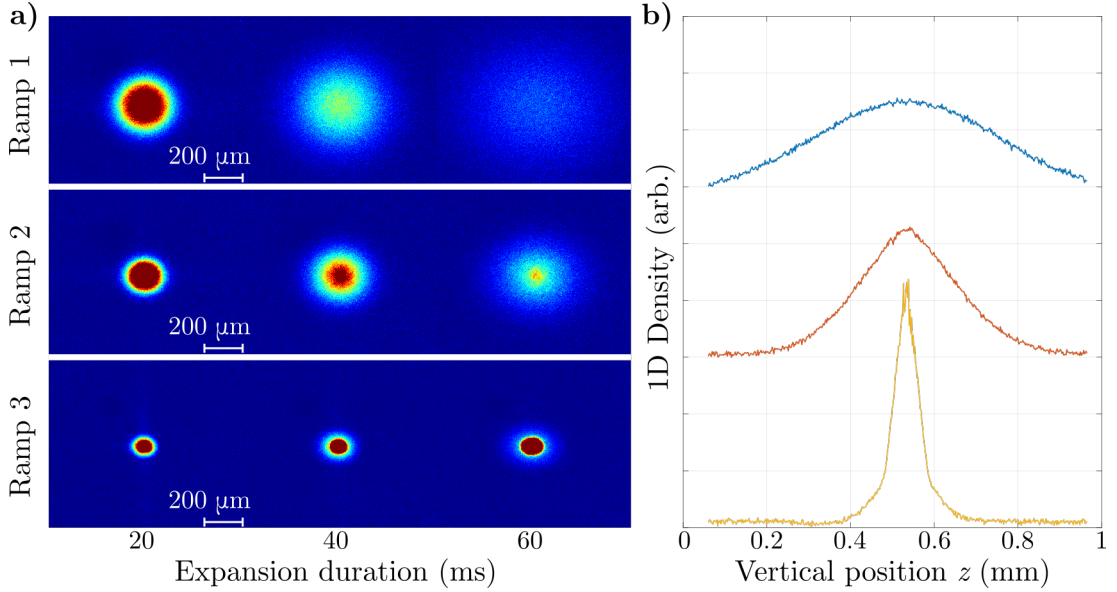


Figure 5.19: a) Absorption images used for the temperature measurement in Fig. 5.18 a) - each image is an average of two. b) The 1D density profiles in the z axis for each of the 60 ms expansion duration images in a) - the peak characteristic of the existence of a condensate is clearly visible in the orange (ramp 2) and yellow (ramp 3) profiles. The profiles are each offset by the same amount.

in Fig. 5.19. The top row is an expansion measurement performed after evaporation ramp 1; the centre and bottom rows show expansion results after ramps 2 and 3, respectively, all for the same expansion durations. Although the difference is not really clear after 20 ms while even the thermal sample is optically dense, after 40 ms it can be clearly seen that there are only thermal atoms after ramp 1, there is some condensate after ramp 2 and much more condensate after ramp 3. In fact, after ramp 3 once the sample has been given 60 ms to expand, the sample is still optically dense except for some thermal atoms which can be observed.

Having cooled our atoms from room temperature to BEC, we examine the phase space densities and atom numbers that result from each stage of the cooling process in Fig. 5.20. It appears that the efficiency falters at certain points in the process, such as from Raman cooling to the dipole trap. It is worth noting that in the process of transferring atoms from one trap to another, some initial heating

is completely normal and these transfers are necessary to be able to cool further. The loss of atoms from the lowering of the dipole trap from 200 W to 130 W does not seem worth it for the minimal increase in phase space density, but it is a necessary evil in a process that produces a BEC.

The leap in phase space density from the dipole trap at 130 W to the evaporation pre-ramp results from the change in regimes for phase space density calculations. The evaporation pre-ramp was the first stage for which Eq. 3.8 was used, instead of Eq. 3.4 to calculate the phase space density. The dimple pre-ramp stage was deemed the most suitable to begin using the trap parameters to calculate the phase space density, which is more accurate than using the properties of the gas once a condensate starts forming. The efficiency of the evaporation stages as in Eq. 2.34, is obtained from the gradient of the log-log plot in Fig. 5.20. The mean evaporation efficiency is 3.2, which indicates that the fractional gain in phase space density was far greater than the fractional loss of atoms.




Figure 5.20: Phase space density vs atom number after each stage in the cooling process on a log-log scale. All points represent stages during our experimental sequence. The hold times for the dipole and dimple traps are 200 ms for the ‘Dipole 200 W stage’, 1.3 s for the ‘Dipole 130 W’ stage and 100 ms for the ‘Dimple pre-ramp’ stage. All evaporation ramps are included for completeness, although phase space density readings could be inaccurate after the BEC transition occurs.

5.3 Bose-Einstein Condensate Characterisation

Once we have achieved a BEC, we can examine the effects described in Chapter 3, and the properties of a quantum gas. We obtain results using the bimodal fitting procedure described in Section 5.1.3, and characterise our condensate under various conditions. I will present low energy temperature measurements, our method of thermal atom reduction, the collapse of a Bose-Einstein Condensate under attractive interactions and our ability to observe a BEC in a guiding beam for durations up to 1 s.

5.3.1 BEC Benchmarks

At various stages throughout the cooling process we measured the temperature of the gas. After our second evaporative cooling ramp we observed the onset of Bose-Einstein condensation, so began to use our bimodal fitting procedure to measure the temperature of the thermal atoms. We adjusted the evaporation

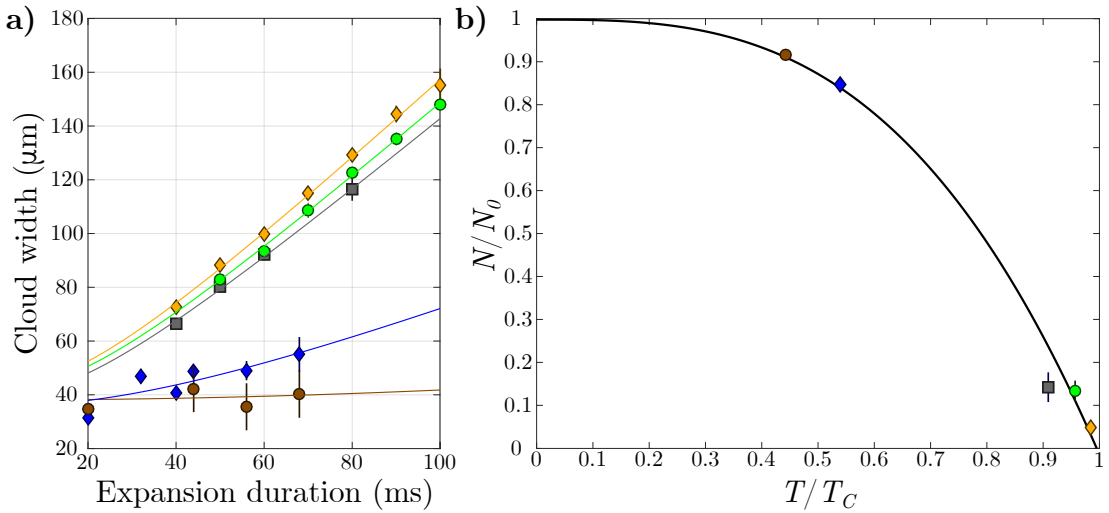


Figure 5.21: Temperature measurements of the thermal component of the Bose gas after various evaporation ramps as a way of calculating their critical temperatures. a) Temperature measurements after each ramp, yielding temperatures of 18.3(3) nK (yellow diamonds), 16.3(2) nK (green circles), 15.0(3) nK (grey squares), 3.1(4) nK (blue diamonds) and 676(625) pK (brown circles). b) A plot of Eq. 3.10, overlapped with our experimental data.




Figure 5.22: Atom number vs hold duration in the dimple beams after evaporation (red squares). The solid red lines are exponential decay fits associated with trap lifetime, for 0 s - 2.5 s, and for 2.5 s onwards.

procedure to have samples with high and low condensate fractions, for analysis of both regimes. These results are presented in Fig. 5.21, where the temperature fits begin to break down for high condensate fractions, represented by blue diamonds and brown circles. From these two fits, it is indicated that the hotter 3.1(4) nK measurement (blue diamonds) appears to begin with a more compressed cloud than the colder 676(625) pK measurement (brown circles). This could be explained partly by an error in the temperature fit due to optically dense results at 20 ms, and partly by the possibility that the initial expansion of the thermal component is more driven at 676 pK as the higher condensate fraction pushes the thermal atoms outwards. Despite this discrepancy, we compare the results to Eq. 3.10 to identify whether the temperatures we observe are accurate. The temperatures we obtained are used in conjunction with the total atom number N_0 and condensate number N from our fit to calculate T_C using Eq. 3.10, and N/N_0 is plotted against T/T_C as in Fig. 3.2. We see in Fig. 5.21 that the T/T_C obtained from the bimodal fit to the atom cloud in both the low and high con-

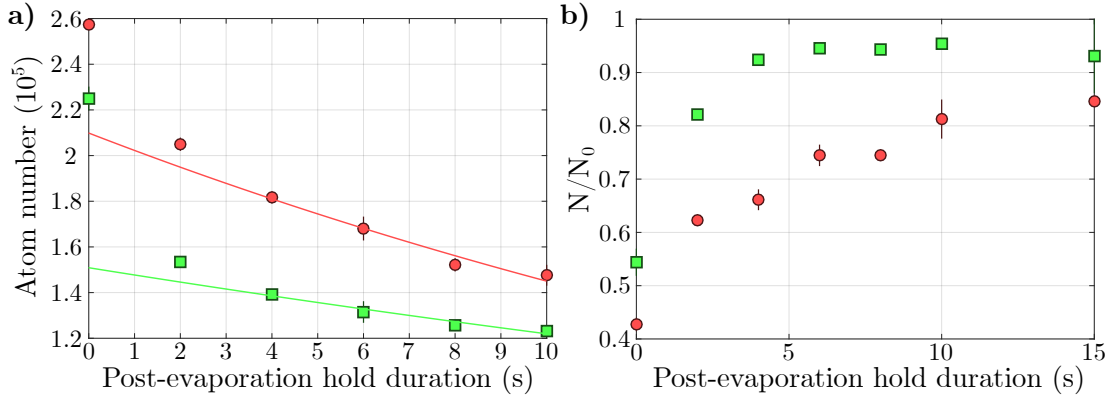


Figure 5.23: A study of the trap lifetime and condensate fraction against hold duration in the trap under different conditions. The red circles in both plots indicate data taken for dimple beam powers of 1 mW horizontal and 8 mW vertical, and the green squares for powers of 0.3 mW horizontal and 6 mW vertical. a) A measurement of the lifetime in these traps. The solid red line is a typical exponential fit to the 1 mW, 8 mW data from 2.5 s hold duration onwards, and the green solid line is a fit to the 0.3 mW, 6 mW data for the same range of hold durations. b) Condensate fraction with respect to hold duration in these traps.

condensate fraction regimes agree well with the theory.

Examining the trap lifetime following the formation of the BEC in Fig. 5.22, we see two distinct rates of loss - for the first 2.5 s the rate is much higher as the hottest atoms that cannot be contained in the trap are lost, and the atoms that remain thermalise; from 2.5 s onwards, the rate has settled and mainly three-body loss is observed. The trap lifetime obtained from the second slope is 24.6 s, an order of magnitude gain on the lifetime measurement before evaporation seen in Fig. 5.16. The increase in lifetime follows from the decrease in density resulting in a lower rate of associated inelastic scattering, in addition to the massive reduction of kinetic energy in the gas.

The density and kinetic energy are reduced again for a further reduction of dimple beam powers, resulting in an even longer trap lifetime as seen in Fig. 5.23. For beam powers of 1 mW horizontal and 8 mW vertical, after the first 2.5 s an exponential decay fit to the data yields a trap lifetime of 27(2) s. However, by decreasing the powers to 0.3 mW horizontal and 6 mW vertical, we achieve a

lifetime of 47(8) s. The increase of condensate fractions observed for the same images demonstrate that most of the atoms that escape the trap are thermal. For 1 mW and 8 mW, the condensate fraction never exceeds 85%, but for the 0.3 mW horizontal and 6 mW vertical configuration we achieve an almost pure condensate of 1.4×10^5 atoms. It is thought that the drop in condensate fraction after 15 s is because our bimodal fitting procedure breaks down for condensates above 95% purity. In general, we use 1 mW horizontal and 8 mW vertical to trap our BEC and a different method to increase the purity of the condensate.

5.3.2 Thermal Atom Removal

It might happen that after evaporation, there is still a thermal fraction in the cloud that is not desired. We typically reduce this thermal component by lowering the dimple traps linearly to 3 mW horizontal and 6 mW vertical in 400 ms and reducing the magnetic field levitation gradient over the course of 1.4 s to somewhere between 98% and 99% of its usual value. This lowers the trap depth further, just enough that the atoms lost are mainly thermal. The levitation is

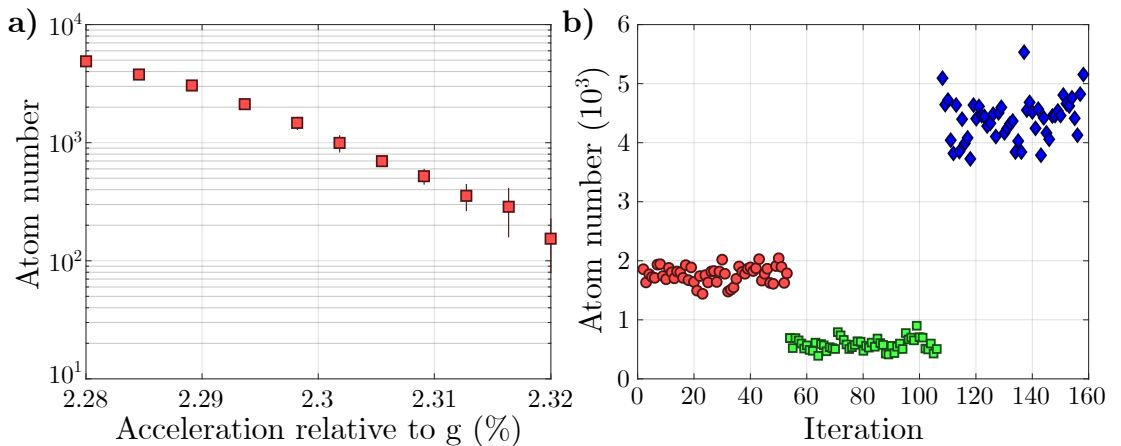


Figure 5.24: A study of our atom number reduction method. a) Atom number vs the percentage of the acceleration due to gravity the atoms experience in the trap. b) A demonstration of the reproducibility of this atom number reduction. The red circles indicate measurements taken at 2.30% of g , the green squares at 2.31% and the blue diamonds at 2.28%.

held steady at its reduced value for 1 s, and is then ramped back to its original value in 1.4 s. Where most other sweeps of magnetic levitation current would be linear, this sweep takes the form of a hyperbolic tangent for a softer transition, as any sudden changes could result in most atoms being launched or dropped from the trap. This process results in a drop from 2.5×10^5 atoms in our sample, to 1.5×10^5 atoms with no detectable thermal component. Using this method, by varying the adjusted magnetic field gradient value, we can reproducibly reduce the atom number down to about 200 atoms. Our atom number using this method can fluctuate by ± 100 atoms for attempts to reduce our atom number to 600, or ± 350 atoms for reduction to 4500 atoms.

5.3.3 BEC Collapse

As noted in Section 3.3, a BEC will collapse if it contains a large number of atoms and is then tuned to have a negative scattering length [93]. In this case, the condensate implodes and suffers heavy atom losses, as seen in Fig. 5.25. From the absorption images in this figure, we can see that the condensate is stable at $5.8 a_0$, with a condensate fraction of $N/N_0 = 0.45$, which shrinks as the positive scattering length is reduced. In Fig. 5.25 d) and e), we see the condensate with attractive interactions. In both cases, about half of the atoms have been lost, heating has occurred and the clouds take strange shapes, as the collapse has likely induced unusual oscillations and breathing modes.

5.3.4 Long expansion durations

The ability to observe BECs over long expansion durations (> 100 ms) is limited by the curvature of the magnetic field gradient, magnetic field fluctuations in the horizontal axes and the expansion rate of the condensate. Using a vertical guiding beam, we are able to circumvent the anti-trapping force introduced by magnetic levitation, and by tuning the interactions of the condensate we limit its initial




Figure 5.25: The effects of the change of interaction strength after the formation of a BEC. a) Atom number vs. scattering length clearly shows the atom loss caused by attractive interactions. b)-d) single-shot absorption images of a condensate in various interaction regimes. b) $5.8 a_0$, c) $0.56 a_0$, d) $-2.08 a_0$, e) $-6.04 a_0$

expansion energy. These measures enable the observation of condensates for durations of up to 1 s. We demonstrate one such observation in Fig. 5.26, where several averaged absorption images are shown for expansion durations ranging from 50 ms to 1 s. The atoms, tuned to $15 a_0$, are levitated and held in a vertical guiding beam with a radial trapping frequency of $2\pi \times 2.1$ Hz.

We plot the vertical RMS radius of the cloud against expansion duration to measure the expansion velocity. The RMS radius is calculated with $\Delta Z = \sqrt{\frac{1}{N} \int n(z)(z - \bar{z})^2}$, where \bar{z} is the centre of mass position of the cloud, $n(z)$ is the atom density at the position z , and N is the atom number. We observe an initial expansion velocity of $v_{\text{RMS}} = 0.128(5)$ mm/s, corresponding to an expansion energy of 260(20) pK. The initial expansion is determined by the set interaction strength, however after 400 ms expansion duration in the guiding beam, the expansion begins to accelerate. There are some contributing factors in this

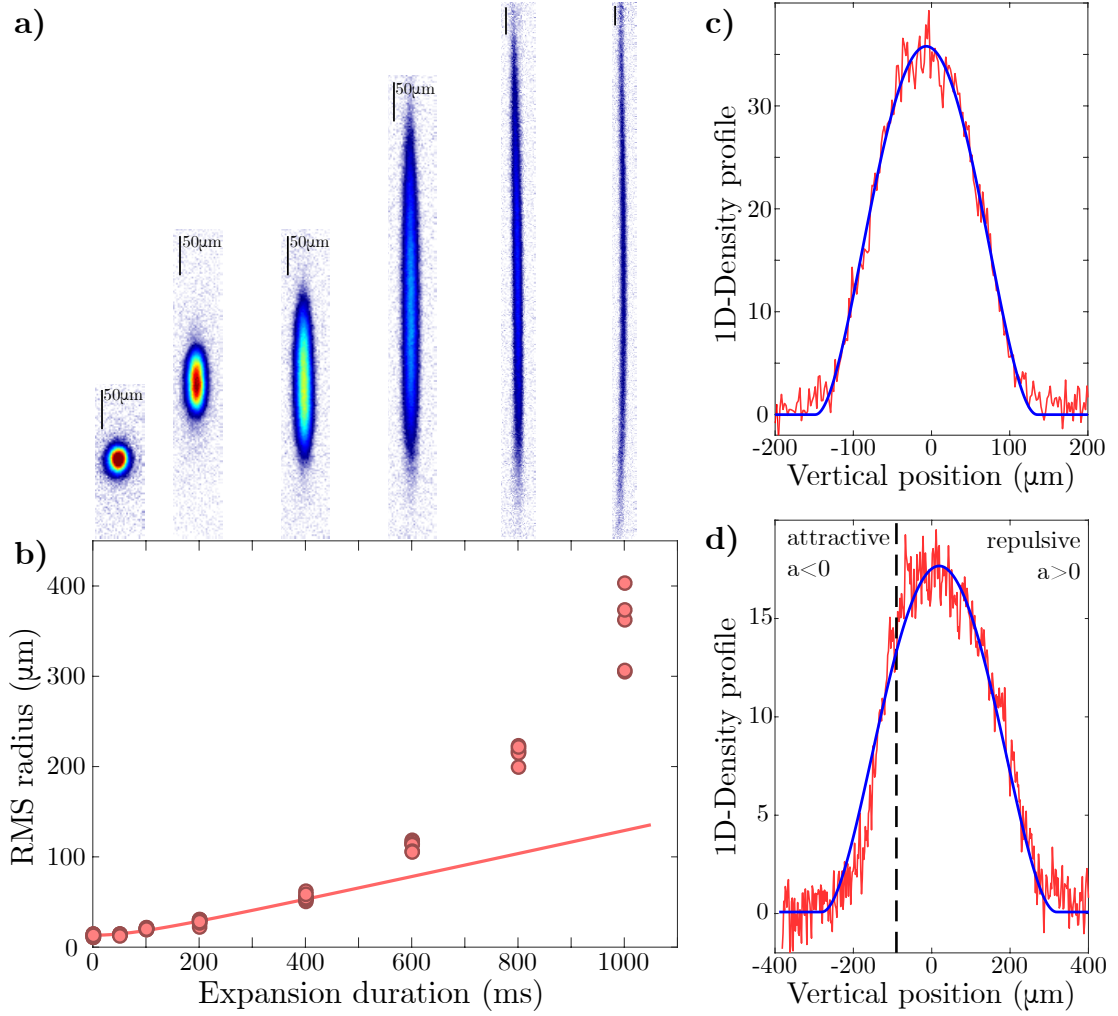


Figure 5.26: Expansion results for a BEC of 1.1×10^4 atoms in a guiding beam, published in Ref. [51], and found in Appendix B of this thesis. a) Averages of 8 images taken after expansion durations of 50 ms, 200 ms, 400 ms, 600 ms, 800 ms and 1 s. The scale bars represent $50 \mu\text{m}$ on the vertical axis of each image. b) Vertical RMS radius of the cloud vs expansion duration for the images, from which we observe two expansion regimes. The solid red line is a fit to the initial expansion from which we obtain a velocity. c)-d) Horizontally integrated density profiles for expansion durations of c) 400 ms and d) 600 ms, with blue lines representing fits to the data. In d), the dashed line represents the position at which zero scattering length occurs.

acceleration - an effect of the transition from a BEC to a Tonks gas [117, 118]; slight fluctuations of the magnetic field strength and gradient throughout the expansion; and possibly coupling between horizontal oscillations, which were de-

tected radially in the beam, and vertical motion.

Also included in Fig. 5.26 are horizontally integrated 1D profiles of the absorption images after 400 ms and 600 ms expansion durations, where it is apparent that the expansion of the cloud is asymmetric over longer durations. This is considered to be because the BEC is expanding along the vertical axis where a considerable magnetic field gradient exists, so the condensate extends to regions where the interactions are strongly attractive, so the expansion is hindered on the lower portion of the cloud.

Chapter 6

Outlook and Conclusions

During the course of this work, I contributed to developing and building a new, versatile apparatus designed for the production of caesium Bose-Einstein condensates. We built a vacuum chamber which maintains a pressure of 10^{-8} mbar in the section that houses a 2D+ MOT, and 10^{-11} mbar in the chamber in which experiments are performed.

We routinely trap 10^8 atoms in our MOT, cool them close to their lowest vibrational state using Raman sideband cooling and evaporatively cool them to degeneracy, tuning interactions to favour elastic collisions. From this process we produce BECs of 2.5×10^5 atoms, with the ability to further increase the purity and reproducibly reduce the atom number down to 200 atoms, if desired. We are capable of producing BECs in the sub-nK regime, where time-of-flight thermometry becomes inaccurate. By using a vertical guiding beam, we can observe the expansion of our condensates for durations up to 1 second, and reduce their dimensionality to 1D [51].

6.1 Future Work

Studies on measuring the temperatures of Bose gases with minute thermal components will continue. In Ref. [57], the margin of error on the temperature measured by time-of-flight was $\sim 20\%$. It is thought that the accuracy of time-of-flight thermometry methods can be improved by adjusting the interaction regime before the gas is released. With slightly attractive interactions, the condensate component of the gas could shrink as the thermal component expands. In this case, more of the small ‘wings’ of the thermal component could be revealed to obtain temperature measurements in very low energy regimes.

In previous work, spin excitations of a rubidium BEC, or magnons, were used to measure the temperature of the main gas [53] down to temperatures of $0.02 T_C$. Magnons were created by altering the Zeeman states of up to 15% of their atoms using rf-pulses, keeping them dilute enough to avoid the magnons condensing. The magnons were then allowed to thermalise in the trap with the condensate before being released. Finally, a spin-selective imaging technique was performed to image only the magnons, so they would not be obscured by the main cloud containing the condensate. In principle, this thermometry technique can be implemented in our current setup with no modifications to the apparatus. Using rf-pulses, we can transfer some of the atoms of our condensate from the $|3, 3\rangle$ to the $|3, 2\rangle$ state while they are still in the trap, allow them to thermalise with the condensate and release them. Upon being released, both states would be separated with a magnetic field gradient, as was demonstrated in Fig. 5.13. This would make it simple to image both clouds simultaneously and compare the temperature measurements in a regime where it could be measured from both to confirm our magnons have thermalised properly, before using the method for lower temperature regimes.

Our experimental apparatus could also be used to implement less destructive methods of thermometry than time-of-flight imaging for the sub-nK regime, that

have been proposed in recent years. In Ref. [55], it was suggested that the phase readout from different methods of atom interferometry could give an indication of the temperature of the gas. Schemes for both Mach-Zehnder interferometry and Ramsey interferometry were suggested, although the latter is considered better for the purpose of thermometry. Spatial separation of the clouds does not occur with the Ramsey method, and the phase measured from it is capable of being twice as accurate as the phase measured using the Mach-Zehnder method. Although both methods are viable using our apparatus, it would take fewer adjustments to implement the Mach-Zehnder method. While this thesis was being written an optical lattice was introduced into the experiment, permitting this type of interferometry. From this, we were able to determine the phase shift introduced by having the atoms on one arm of the interferometer pass through a dipole beam before recombination [51]. If we performed a similar experiment with atoms trapped in the dipole beam before the interferometry sequence, the phase readout could indicate their temperature.

Finally, a proposal [56] indicates that thermometry can be performed non-destructively by repeated measurements of the momentum and position spread of impurities immersed in a BEC. The proposal suggests trapping a small number of probe atoms in the centre of a larger atom cloud, allowing them some time to thermalise with the cloud, and imaging the probes to determine their position and momentum spread after thermalisation. Although the example given in this publication involves using ytterbium atoms to measure the temperature of a potassium BEC, the calculations are general. The best regime for accuracy at low temperatures is predicted to be where the coupling strength between the impurity and the BEC (g_{IB}) is almost equal to the coupling strength of the BEC atoms with each other (g_B), i.e. $g_{IB}/g_B \approx 1$. Under these circumstances, 100 repeated measurements of the probe position spread could have an uncertainty of <14% for temperatures around 200 pK, which decreases for higher temperatures




Figure 6.1: An illustration of how non-destructive measurements could be performed using our experimental apparatus. a) An optical lattice creates flat ‘pancake’ lattice sites, which allows a small number of atoms in each site to be addressed by a Raman beam (blue), inducing a two-photon transition. b) The proposed energy levels that could be used. The atoms start off in the $|3, 3\rangle$ state (brown circle), and undergo a Raman transition to the $|3, 2\rangle$ state (blue circle) where they remain as an impurity until they are imaged. A microwave transition is used to transfer the impurity atoms to the $|4, 1\rangle$ state, so they can be imaged with minimal disturbance to the $|3, 3\rangle$ atoms.

or more measurements. In our experiment, an initial study of this method is feasible if we use an optical lattice to create a 2D system, which would allow us to implement the analysis of different spin components which would otherwise be under-levitated or anti-levitated [110, 119]. We could address the atoms in the centre of the 2D ‘pancakes’ with a two photon Raman transition, to transfer them from the $|3, 3\rangle$ state to the $|3, 2\rangle$ state. These atoms should be dilute enough to avoid condensation, and should have a similar coupling strength with the $|3, 3\rangle$ atoms as other $|3, 3\rangle$ atoms. After being allowed enough time to thermalise, a microwave transition could be used to transfer only the $|3, 2\rangle$ atoms to the $F=4$ manifold, preferably to the $|4, 1\rangle$ state to be far away from any microwave transitions available to the $|3, 3\rangle$ atoms, before being imaged selectively. In this way, the position spread of the impurity could be measured with minimal disturbance to the BEC. The scheme for this proposed experiment is illustrated in Fig. 6.1.

This non-time-of-flight method of thermometry could be used to observe inter-

6.1. FUTURE WORK

esting temperature regimes which occur in Bose gases of lower dimensionalities. In our setup, we have implemented beams which form an optical lattice, enabling us to create a quasi-2D Bose gas, which presents the possibility of measuring temperature in a 2D quasi-condensate regime [117, 120], which could be studied using the non time-of-flight methods of thermometry presented above. In addition, modulation of this lattice would allow for studies of excitations, band structure, and interactions [121].

Appendix A

Publication: A simple laser shutter with protective shielding for beam powers up to 1 W

Published in Rev. Sci. Instrum.

DOI: [10.1063/1.5053212](https://doi.org/10.1063/1.5053212)

Note: A simple laser shutter with protective shielding for beam powers up to 1 W

Craig D. Colquhoun, Andrea Di Carli, Stefan Kuhr, and Elmar Haller^{a)}

Department of Physics, SUPA, University of Strathclyde, Glasgow G4 0NG, United Kingdom

(Received 22 August 2018; accepted 10 November 2018; published online 5 December 2018)

We present the design of an inexpensive and reliable mechanical laser shutter and its electronic driver. A camera diaphragm shutter unit with several sets of blades is utilized to provide fast blocking of laser light and protective shielding of the shutter mechanism up to a laser beam power of 1 W. The driver unit is based on an Arduino microcontroller with a motor-shield. Our objective was to strongly reduce construction effort and expenditure by limiting ourselves to a small number of modular parts, which are readily available. We measured opening and closing durations of less than 800 μs , and a timing jitter of less than 25 μs for the fastest set of blades. No degradation of the shutter performance was observed over $5 \cdot 10^4$ cycles. © 2018 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/1.5053212>

Mechanical optical shutter units have become indispensable in modern optics laboratories to provide a time-dependent extinction of laser light. Depending on the application, a multitude of desirable properties can be identified, such as low extinction ratios, fast switching times, low time jitter, high reliability, high repetition rates, small sizes, or long operation lifetimes.

Commercial products are currently available that fulfil most of those design requirements at high costs,¹ but laboratories often need dozens of shutter units and commercial solutions can quickly become unaffordable. As a result, many experimental groups have developed their own shutter designs with varying design goals and technical approaches,² e.g., based on loudspeakers,³ computer hard drives,⁴ or piezoelectric devices.^{5,6}

In this article, we present the design of a mechanical shutter and its driver unit with two design objectives. The first objective is to strongly reduce construction effort and costs while preserving fast switching times and a high reliability. We do so by limiting ourselves to a small number of modular parts which are readily available.^{7,8} The shutter unit utilizes a small diaphragm shutter with multiple blades as is normally used in compact digital cameras, and the driver unit is based on an Arduino microcontroller with a motor-shield.⁸ The second design objective is a protection mechanism that facilitates the blocking of laser beams up to a continuous power of 1 W. We implement the protection with a shielding blade that reflects the laser light.

Details and additional materials for the construction are available in the [supplementary material](#). Here, we give an outline to the design of the shutter blades, the driver unit, the microcontroller software, and the enclosure of the shutter. An experimental characterisation of the switching time, the jitter, and the reliability is provided.

Figure 1 illustrates the design of the **shutter blades**. The shutter contains three sets of blades—a light pair of blades B1 that close from opposite sides of the aperture in a “scissor” motion, overlapping in the centre and blocking light; a sturdy filter blade B2 originally intended to attenuate the light in a camera; and an unused blade with a hole which only limits the aperture size B3. We utilize blades B1 for fast switching operations and blade B2 for protection and dispersive reflection of laser light. Typically, the blades of small diaphragm shutters are optimized for low weight and friction, and they start to bend or melt when absorbing laser powers of more than 50 mW. We managed to increase the beam power up to 1 W⁹ by adhering a small strip of aluminum foil to filter blade B2 that dispersive reflects the laser light and dissipates heat. By our




FIG. 1. Images of the (a) back and (b) front of the shutter before modifications and (c) within its 3D-printed enclosure. Red dashed lines indicate removed parts, and red arrows point toward the (a) connection terminals of the solenoids and (b) sliders of the shutter blades. The labels indicate that the connection terminals and sliders attach to scissor blades B1, filter blade B2, and unused aperture blade B3.

^{a)}Electronic mail: elmar.haller@strath.ac.uk

design, most of the reflected light is trapped in the enclosure of the shutter.

The positions of the blades are controlled by small solenoids with independent connection terminals as indicated by red arrows in Fig. 1(a). The shutter blades are bistable without any springs or other self-restoring elements, and a short current pulse of ± 200 mA for a duration of 3 ms is sufficient to flip the position. The final state of the shutter is determined by the direction of the current. For simplicity, we typically connect both blades, B1 and B2, in series by soldering thin wires to the connection terminals, but an independent control of the blades is used for the purpose of testing the shutter for this note.

The **shutter driver** consists of an Arduino microcontroller with a motor-shield (Fig. 2). The microcontroller monitors a digital (TTL) input signal that indicates the state of the shutter—a low (high) signal corresponds to a closed (open) state. The detection of a signal change triggers the short current pulse of the motor-shield with the required current direction to flip the blades. An operation of both solenoids in series requires a supply voltage of 5–6 V for the motor-shield to generate the correct current pulse. It is possible to supply the shield by the regulated 5 V output of the Arduino microcontroller, but a direct connection to the main power supply is advisable for the simultaneous control of 4 shutters units. For convenience, we added to the circuit a toggle switch to open the shutter manually, and a light emitting diode (LED) to indicate the shutter status. We intentionally limited the circuit to include only essential elements, and all components except for the microcontroller can be integrated into the front panel without the need of an additional circuit board.

Our **microcontroller software** is provided in the [supplementary material](#). The tasks of the program are the tracking of the shutter status, the detection of a change of the TTL input signal, and the control of the motor-shield. Timer interrupts are included for the parallel control of several shutter units. The use of interrupts allows us to generate current pulses of well-defined duration without blocking the program flow. We measured a response delay between the input signal and the current pulse of $230(30)$ μ s for a simultaneous use of 4 shutters units.

A plastic **enclosure** is used for the shutter unit to reduce the coupling of vibrations. The casing is 3D-printed using




FIG. 2. Circuit diagram of the shutter controller. The design is based on an Arduino microcontroller with a motor-shield. We added a toggle switch to open the shutter manually by bypassing the TTL input signal, and an LED to indicate the status of the shutter blades. One motor-shield facilitates the simultaneous control of 4 shutter units.

fused deposition modeling of polylactide (PLA) plastic. We sandwich the shutter between rubber “O” rings and a black anodised aluminum disc with a small hole of 4 mm diameter to further dampen vibrations. The disc reduces possible backscattering from the aluminum foil adhered to the shutter blades, and it prevents a melting of the casing material due to a misaligned laser beam. The corresponding computer-aided design (CAD)-model files of the enclosure can be found in the [supplementary material](#).

The final part of this note describes an experiment to benchmark the speed, time jitter, and robustness of the shutter and driver unit. We used a photodiode¹⁰ and an oscilloscope to measure the power of a laser beam after it propagated through the shutter. Timings for the shutter and for the acquisition oscilloscope were provided by an NI-multifunction IO device.¹¹ The shutter aperture is 4 mm in diameter, and the laser beam was collimated with a $1/e^2$ waist of 1.1 mm. As opposed to our normal operation, we connected each shutter blade to the shutter driver separately to study the timed opening and closing of the blades independently of one another. The intensity profiles of 500 consecutive opening and closing cycles were recorded and analyzed (Fig. 3). No degradation was detected over the course of $5 \cdot 10^4$ additional cycles.

Figure 3 shows the photodiode signal for a time t after the trigger signal to (a) open or (b) close the shutter with scissor blades B1 (blue) and filter blade B2 (red). The photodiode voltage is normalized for each data set to the signal of an open shutter. We determine an opening delay between the trigger and an increase to 5% of the full photodiode signal of $2.29(2)$ ms and $3.71(3)$ ms for blades B1 and B2. The opening durations, measured by an increase from 5% to 95% of the total signal, is $790(10)$ μ s and $1.51(3)$ ms for the two sets of blades. The closing procedure is slightly faster with a closing delay of $2.73(2)$ ms and $2.71(3)$ ms and a closing duration of $573(7)$ μ s and $1.46(2)$ ms for blades B1 and B2, respectively. Opening and closing delays are longer than the electronic response time, and we expect most of the delay time to be used to overcome




FIG. 3. Transmission signals of the two blades being (a) opened and (b) closed. The red (blue) lines show the transmission throughout the operation of filter blade B2 (scissor blades B1). The gray lines show the operation of B2 for a reduced beam waist (see text). The time scale indicates the delay time t after the change of the TTL input signal.




FIG. 4. Histograms of opening times for the transmission signals. (a) half-opening times and (b) half-closing times of filter blade B2. (c) half-opening times and (d) half-closing times of scissor blades B1.

friction and to separate the overlapping blades. We presume that the scissor blades are faster than the filter blade because they close in from both sides and meet in the centre of the aperture, thus traveling half the distance. Both blades have a velocity of approximately 1.2 m/s.

Another important property to characterise a shutter is the reproducibility of operation times. The histograms in Fig. 4 show the variation of the half-opening and half-closing times, i.e., the time ΔT to reach 50% of the total beam power after a change of the trigger signal. Our histograms display a low jitter time with no significant outliers. Filter blade B2 shows a positive (negative) skew of the distribution for the opening (closing) process with standard deviations of 60 μs (40 μs). The distributions of the timing of scissor blades B1 show the opposite skews with the standard deviations of 21 μs (24 μs). We speculate that this skewing is due a position dependent variation of the friction between blades, and the details of the skewing might vary from device to device. The difference in opening and closing times for the same blade might be due to a small misalignment between the center of the shutter aperture and the laser beam.

For a better comparison with other publications, we reduce the $1/e^2$ -waist of the beam to 140 μm and repeat the measurements. 500 data sets for the opening and closing of blade B2 are represented by gray lines in Fig. 3. The reduced beam waist results in a reduction in the time taken for the

photodiode signal to change between 5% and 95% of the total signal. For blades B1 (B2), we measure an opening duration of 137(7) μs [220(5) μs] and a closing duration of 100(4) μs [155(5) μs], which are in agreement with previous measurements and with the scaling of the beam waist.

In conclusion, we implemented and benchmarked a simple and robust shutter design based on a diaphragm shutter with multiple pairs of blades. A lightweight pair of blades is utilized for fast shutter operation while being protected by a slower and sturdier blade. The shutter can operate up to a continuous laser beam power of 1 W. For the opening and closing of fast blades B1, we measured delays of less than 3 ms, opening and closing durations of less than 800 μs (140 μs for the smaller waist) and a timing jitter of less than 25 μs . Our design goal for the shutter and driver units was to strongly reduce construction effort and costs while preserving robustness and high power operation.

Please see [supplementary material](#) for the software of the microcontroller and for the CAD-model files for the casing of the shutter.

We acknowledge the financial support by the EU through the Collaborative Project QuProCS (Grant Agreement No. 641277). A.D.C. acknowledges the financial support by EPSRC and SFC via the International Max-Planck Partnership. This work was supported in part by the EPSRC Programme Grant DesOEQ (Grant No EP/P009565/1).

¹Examples for commercial products: Vincent Associates, ES6B Laser Shutter, VED24 Shutter Driver, or Thorlabs, SHB025 Shutter.

²S. Martínez, L. Hernández, D. Reyes, E. Gomez, M. Ivory, C. Davidson, and S. Aubin, *Rev. Sci. Instrum.* **82**, 046102 (2011).

³K. Singer, S. Jochim, M. Mudrich, A. Mosk, and M. Weidemüller, *Rev. Sci. Instrum.* **73**, 4402 (2002).

⁴L. Maguire, S. Szilagy, and R. Sholten, *Rev. Sci. Instrum.* **75**, 3077 (2004).

⁵C. S. Adams, *Rev. Sci. Instrum.* **71**, 59 (2000).

⁶W. Bowden, I. R. Hill, P. E. G. Baird, and P. Gill, *Rev. Sci. Instrum.* **88**, 016102 (2017).

⁷Images show the Q1 diaphragm shutter used for digital cameras ES65, ES70, ES71 by Samsung. This shutter and similar shutters by other manufacturers are available on ebay and aliexpress.

⁸Arduino Mega Microcontroller, and Arduino Motor-Shield by Adafruit, version 1. Inexpensive units can be found on ebay and aliexpress.

⁹Test conditions: power 1 W, wavelength 1064 nm, duration 1 hour.

¹⁰Photodiode, Thorlabs PDA100A-EC, bandwidth 2.4 MHz.

¹¹I/O-Device, National Instruments USB-6366.

Appendix B

**Publication: Interferometric
measurement of micro-g
acceleration with levitated atoms**

Published in New J. Phys.

DOI: [10.1088/1367-2630/ab1bbd](https://doi.org/10.1088/1367-2630/ab1bbd)

**PAPER**

Interferometric measurement of micro-g acceleration with levitated atoms

OPEN ACCESS**RECEIVED**

9 November 2018

REVISED

27 March 2019

ACCEPTED FOR PUBLICATION

23 April 2019

PUBLISHED

22 May 2019

A Di Carli, C D Colquhoun, S Kuhr and E Haller

University of Strathclyde, Department of Physics, Scottish Universities Physics Alliance (SUPA), Glasgow G4 0NG, United Kingdom

E-mail: elmar.haller@strath.ac.uk**Keywords:** interferometry, BEC, metrology, acceleration measurement

Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/4.0/).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

**Abstract**

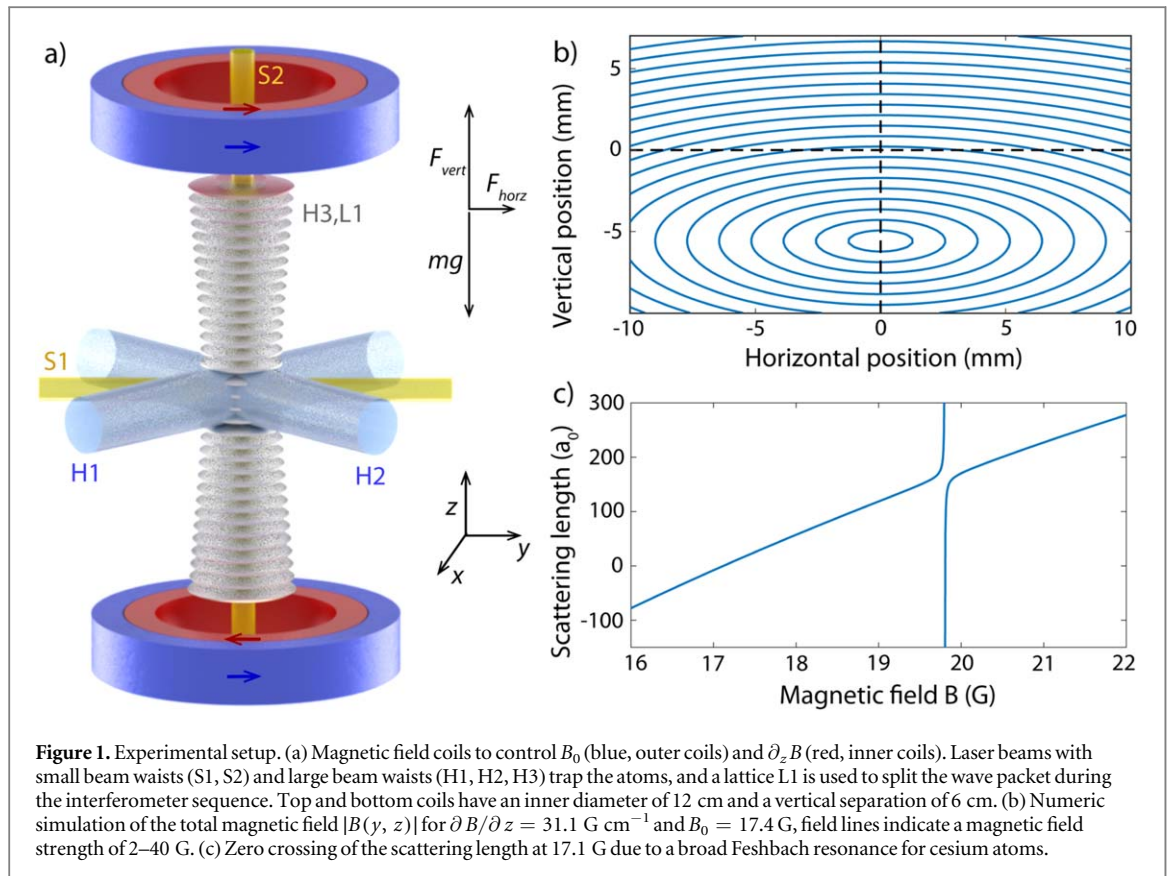
The sensitivity of atom interferometers is usually limited by the observation time of a free falling cloud of atoms in Earth's gravitational field. Considerable efforts are currently made to increase this observation time, e.g. in fountain experiments, drop towers and in space. In this article, we experimentally study and discuss the use of magnetic levitation for interferometric precision measurements. We employ a Bose–Einstein condensate of cesium atoms with tuneable interaction and a Michelson interferometer scheme for the detection of micro-g acceleration. In addition, we demonstrate observation times of 1s, which are comparable to current drop-tower experiments, we study the curvature of our force field, and we observe the effects of a phase-shifting element in the interferometer paths.

1. Introduction

Precision measurements with matter waves have shown tremendous advances over the last decades. In particular, atomic matter wave interferometers demonstrated a ground-breaking increase of the measurement precision of inertial effects, such as rotation [1, 2] and acceleration [3, 4]. In addition, atomic matter wave interferometers have been used to determine the fine-structure constant [5], Newton's gravitational constant [6, 7], and constraints on dark energy [8]. Similar to optical interferometers, atom interferometers split a matter wave into two parts, evolve the parts independently along different paths, and finally recombine the waves to form an interference pattern [9]. The interference pattern depends on the accumulated phase shift of the wave packets during the independent evolution, and the measured quantity is typically inferred from the shape and time evolution of the pattern. The sensitivity of interferometers increases with the accumulated phase shift, which again depends on the evolution time [10]. However, the evolution time of a free falling atom cloud is limited by Earth's gravitational acceleration in most experimental setups, and considerable efforts are made to increase the duration, e.g. in fountain experiments [11], drop towers [12, 13], parabolic flights [14, 15] and in space [16].

In this article, we employ magnetic levitation as a different method to extend the evolution time in earthbound laboratories. Magnetic levitation relies on the use of magnetic forces to cancel the gravitational acceleration and to levitate the particles in space. The method is well established for experiments with ultracold atoms [17–19], and its experimental implementation, i.e. using a pair of current-carrying coils, is significantly simpler and smaller than an atomic fountain apparatus or a drop-tower experiment. Here, we study the advantages and limitations of magnetic levitation for matter wave interferometry with the motional states of Bose–Einstein condensates (BECs), and we demonstrate that magnetic levitation can be employed to reach an expansion time of 1s, which is comparable to current drop-tower experiments [12]. Furthermore, we utilize magnetic levitation to create and to interferometrically measure micro-g acceleration in free expansion, and we show that the negligible center-of-mass motion of levitated atoms facilitates a direct study of phase-shifting elements in the interferometer paths.

Other interferometer schemes use external trapping potentials to prevent the gravitational acceleration by channelling the wave packets along magnetic [20, 21] and optical [22, 23] waveguides. External guiding and



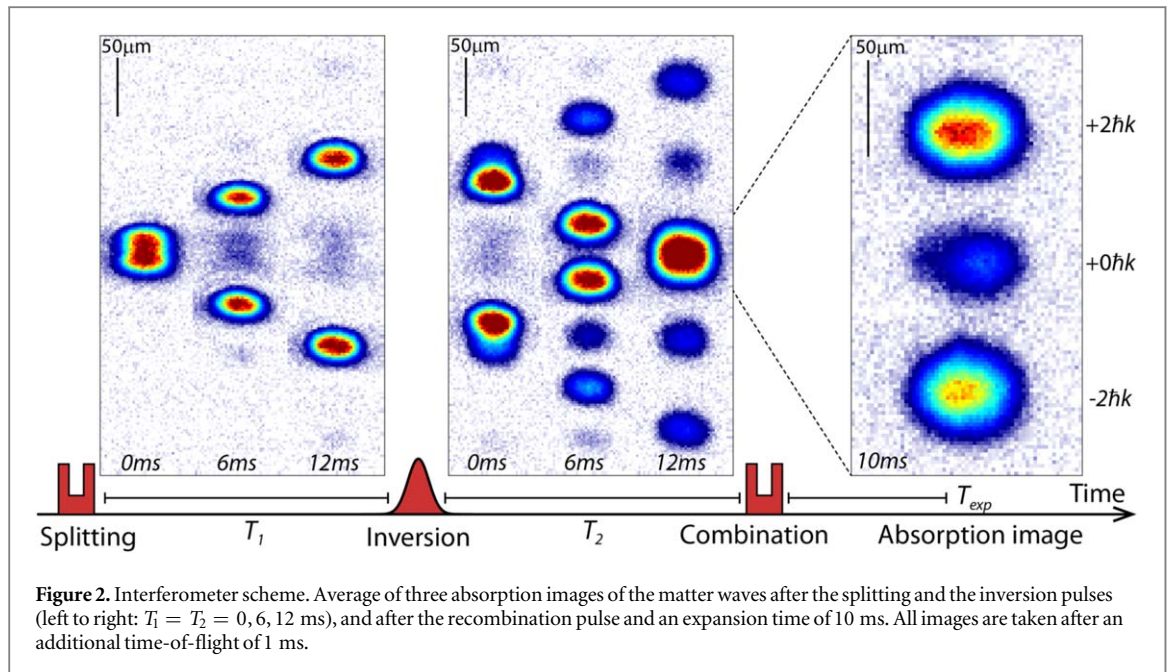
trapping potentials allow for equally long observation times [24], however, they introduce additional challenges. External potentials can cause spatially varying phase shifts and undesired excitations of the wave packets [23, 24], which limit the measurement precision. Our levitation scheme avoids trapping potentials along the gravitational axis, and it facilitates a tuneable scattering length for future studies of interaction effects in atom interferometers.

This article is structured as follows: section 2 provides an overview of our experimental setup, magnetic levitation scheme, and the use of a magnetic Feshbach resonance to control the interaction strength of cesium atoms. Section 3 is used to illustrate the interferometer scheme, and in section 4 we evaluate our measurement precision. Small changes to the magnetic levitation gradient allow us to create marginal accelerations of milli-g (section 4.1) and micro-g (section 4.2). An additional laser beam in one of the interferometer paths constitutes a phase-shifting element in section 4.3. In section 5, we measure features of the magnetic field distribution, such as the transversal curvature of the force field. Finally, using a combination of low interaction strength, low trapping frequencies, and magnetic levitation we demonstrate long expansion and observation times in section 6.

2. Magnetic levitation scheme and experimental apparatus

Our experimental apparatus is designed to independently control two parameters of the magnetic field. The magnetic field strength $B_0 = |\mathbf{B}(x, y, z)|$, at the position of the atoms ($x = y = z = 0 \text{ mm}$) is used to tune atomic interactions by means of a broad magnetic Feshbach resonance for cesium atoms in the strong-field-seeking Zeeman state $|F = 3, m_F = 3\rangle$. We reduce the effects of interaction by setting B_0 to 17.4 G with an s-wave scattering length, a , of approximately $65 a_0$ during the interferometer sequences (figure 1(c)), where a_0 is Bohr's radius. The second controlled parameter is the vertical gradient of the magnetic field, $\partial_z B$, which can be adjusted to exert a vertical pull on the atoms and cancel the gravitational acceleration. Due to the Zeeman effect, cesium atoms in the given state experience a vertical force that is proportional to the magnetic field gradient, $F_{\text{vert}} = \frac{3}{4} \mu_B \partial_z B$. For a mass m of a cesium atom, the levitation gradient can be calculated as $\partial_z B = 4mg/(3\mu_B) = 31.1 \text{ G cm}^{-1}$ [19, 25]. Here, μ_B represents the Bohr magneton and g the gravitational acceleration.

Our coil configuration is based on established designs [18, 19, 25]. It consists of two vertical coils above and below the atoms (inner diameter 12 cm, separation 6 cm), with 5 independently controllable sections. We generate B_0 and $\partial_z B$ by means of two vertical pairs of coil sections with co- and counter-propagating currents (outer and inner sections in figure 1(a)). Pairs of shim coils on each axis at distances of approximately 20 cm



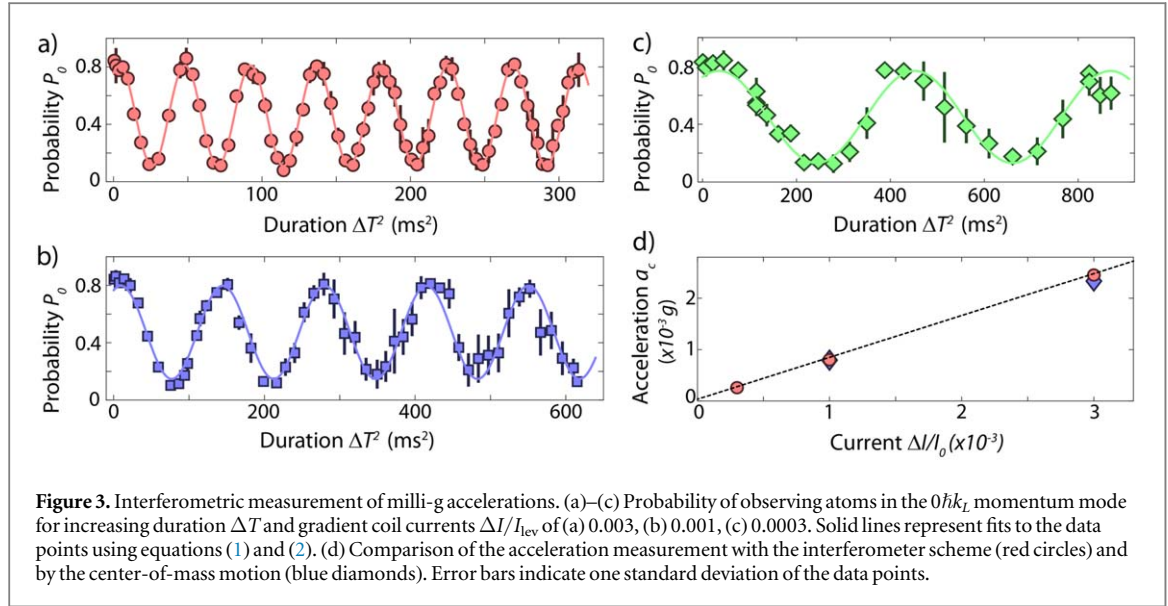
from the atoms allow for additional fine control of the magnetic field. Figure 1(b) shows the total magnetic field strength $B(y, z)$ in the vertical plane as calculated by a numerical simulation of our coils with finite wire elements. The field can be approximated by a magnetic quadrupole field with a shifted minimum at a few millimetres below the atom cloud. Experimentally, we determine B_0 by microwave spectroscopy and we optimize the levitation gradient $\partial_z B$ by varying the levitation current I_{lev} and minimizing position drifts of a BEC during free levitated expansion. Additional effects due to horizontal field curvature and limitations of the levitations scheme for precision measurements are discussed in section 5.

The matter waves of our interferometer are provided by Bose–Einstein condensates. In our setup, 2×10^9 cesium atoms are loaded from a 2D+ magneto optical trap (MOT) into a 3D MOT within 3 s. The atoms are cooled by degenerate Raman sideband cooling [26], and then sequentially transferred into two pairs of crossed optical dipole traps, the first with wavelength 1070 nm, total power 200 W, waists $700 \mu\text{m}$, and the second with wavelength $\lambda = 1064.495(1)$ nm, power 400 mW, waists $90 \mu\text{m}$ (labels S1, S2 in figure 1(a)). Bose–Einstein condensation is reached after 6 s of evaporative cooling, and the density distribution of the atoms is detected by means of resonant absorption imaging after a variable time of levitated expansion and after 1 ms of unlevitated time-of-flight. One cooling cycle has a duration of 15 s and it is similar to [25].

We generate BECs of 2.5×10^5 atoms in the Zeeman sub-state $|F = 3, m_F = 3\rangle$ at a scattering length of $a = 210 a_0$, trapped in the crossed laser beams S1, S2 with trap frequencies of $\omega_{x, y, z} = 2\pi \times (23.5, 17.7, 15.4)$ Hz. To reduce interactions during the interferometric measurement, we tune the scattering length to $65 a_0$ and remove atoms by forced evaporation with a non-levitating magnetic field gradient. The BECs for the interferometer measurements in this work consist of approximately 8×10^4 atoms with a thermal fraction below 5%. Vibrational isolation and damping of the optical table is achieved by a pneumatic isolation system (Newport S-2000A).

3. Interferometer scheme

We employ a Michelson interferometer scheme that is based on three Kapitza–Dirac pulses with a standing light wave (figure 1(a), beam L1) [27]. The pulses change the motional states of the matter waves but leave the internal states of the atoms unchanged [28]. Our pulse sequence and the resulting motion of the matter wave packets are illustrated in figure 2. A first pulse splits the BEC into two wave packets with opposite momenta $\pm 2\hbar k_L$. Here, $k_L = 2\pi/\lambda$ is the wavenumber of the lattice beam and \hbar is Planck’s constant. The wave packets propagate freely for an evolution time T_1 until we apply a second pulse that inverts the direction of the wave packets and changes their momentum by $4\hbar k_L$. A third pulse is used after an evolution time T_2 to recombine the two wave packets. It is identical to the first pulse and generates three wave packets with momenta $p_0 = 0, p_{\pm} = \pm 2\hbar k_L$. The relative population of the recombined wave packets depends on the acquired phase difference $\Delta\Phi$, resulting in a probability P_0 of finding an atom in the p_0 momentum mode



$$P_0 = P_m + \frac{C}{2} \cos(\Delta\Phi). \quad (1)$$

Here, C is the interference contrast and P_m is the offset of the interference signal. We determine P_0 from the ratio of atoms in the p_0 mode to the total atom number in all momentum modes.

Several factors can contribute to the phase difference $\Delta\Phi$. For falling wave packets with spatially homogeneous acceleration a_c , the phase difference is directly proportional to the center-of-mass displacement Δz that was acquired during the total interferometer time $\Delta T = T_1 + T_2 + T_{\text{pulse}}$. Here, T_{pulse} represents the total duration of the pulses. The total phase difference is given by [29]

$$\Delta\Phi = 2k_L \Delta z + \Phi_0 = 2k_L \frac{1}{2} a_c (\Delta T)^2 + \Phi_0, \quad (2)$$

with a term Φ_0 that accounts for additional phase shifts introduced during the initialization process, by noise such as lattice vibrations [15], or by interactions (see section 4.2).

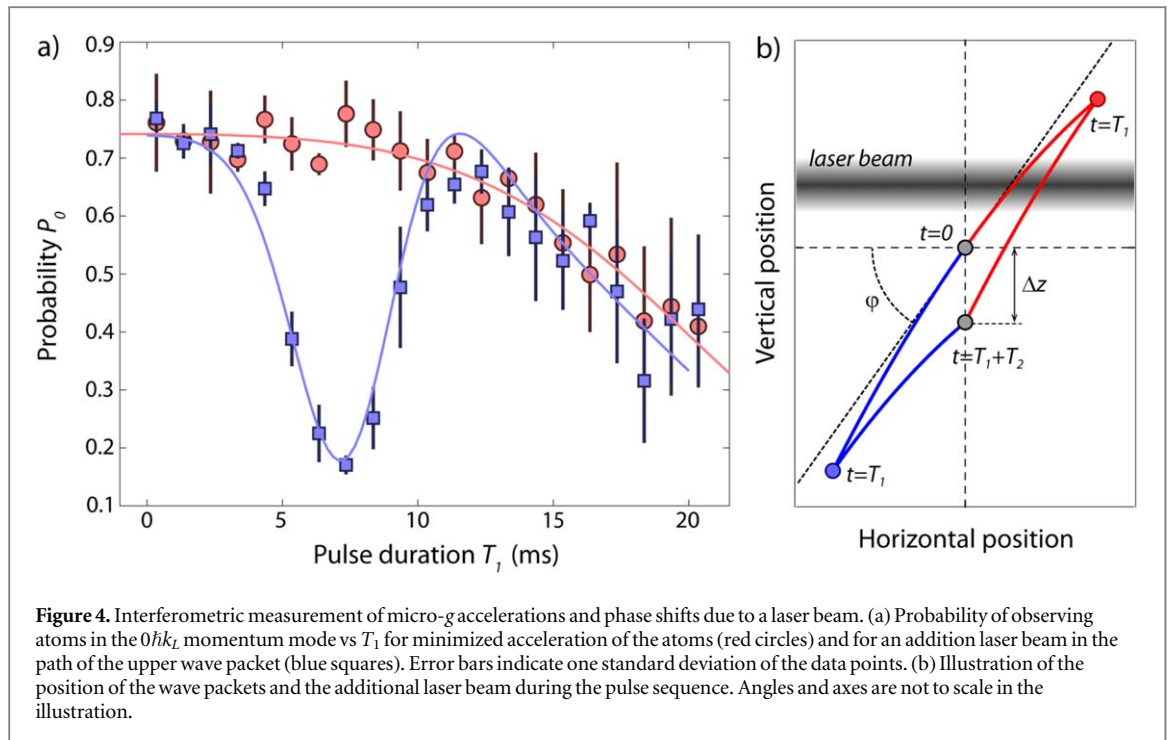
The pulse sequence used in this experiment is based on previous work [20, 30, 31]. Our splitting and recombination pulses consist of three sub-pulses of lattice beam $L1$ with durations 60, 110 and 60 μs , and lattice intensities of $6.6 E_r$, $0.2 E_r$, and $6.6 E_r$. Here, $E_r = \hbar^2 k_L^2 / (2m)$ is the recoil energy for cesium at a lattice wavelength of 1064 nm. Our inversion pulse has a Gaussian intensity distribution with a maximum of $17 E_r$, and a $1/e$ -duration of 35 μs . The sub-pulse scheme allows us to reach a splitting efficiency of 96% of the atoms in the $\pm 2k_L$ modes, and we speculate that the limit of the efficiency is given by the thermal component of our BEC. The efficiency of the inversion pulse is lower, 83%, and residual atoms are clearly visible in figure 2 in the 0 and $\pm 2\hbar k_L$ modes. We suspect that this is due to the velocity selectivity of the inversion pulse and the velocity difference of the accelerated wave packets.

4. Interferometric measurements

4.1. Measuring milli-g acceleration

Our magnetic levitation scheme allows us to apply small forces to the atoms by changing the levitation current I_{lev} in the vertical coils with counter-propagating currents. We use this approach to characterize our interferometer setup for non-zero accelerations. After the preparation of the BEC we increase the current I in the coils, which create the magnetic field gradient, in 75 ms to the ratios I/I_{lev} of 1.003, 1.001, and 1.0003. The acceleration of the BEC is measured with our interferometer scheme. Figures 3(a)–(c) show the corresponding measurements of P_0 for varying evolution times ΔT^2 with $T_1 = T_2$. As expected, we observe sinusoidal oscillations of P_0 , which are fitted using equations (1) and (2) (solid lines) to determine the accelerations a_c (red circles, figure 3(d)).

An independent measurement of a_c , based on the free motion of the BEC, is provided for comparison. We measure the shift of the center-of-mass position for an expansion time T_{exp} of an untrapped BEC in our magnetic field gradients, $z(T_{\text{exp}}) = 1/2 a_c T_{\text{exp}}^2$, with a fit parameter a_c (blue diamonds, figure 3(d)). We find excellent agreement within two standard deviations between the two methods. However, the sensitivity of the free expansion measurement is limited by the observation time. Although our levitation scheme allows for very



long observation times (section 6), it also induces a horizontal dispersion of the BEC in free space, which will be discussed in section 5. Here, we limit the observation time to 200 ms, which allows us to measure the acceleration for $I/I_{lev} = 1.001, 1.003$, but not for 1.0003. The measurement results in figure 3(d) have relative uncertainties of approximately 4% for the free expansion measurement and 0.5% for the interferometric approach.

4.2. Measuring micro-g acceleration

In a second measurement, we utilize the interferometer scheme to minimize the forces on the atoms. We vary the currents in our shim coils and I_{lev} with the goal to maximize the oscillation period of P_0 (red circles figure 4). For optimal current values, we observe a slow drop of the value of P_0 from approximately 0.75 to 0.45 over $\Delta T^2 \approx 1600 \text{ ms}^2$. This reduction is not necessarily caused by a residual acceleration of the wave packets, as it can also originate from dephasing mechanisms that are discussed in the next paragraph. However, fitting $P_0(t)$ with equation (1) provides an upper limit to the acceleration experienced by the atoms. We determine an upper limit for the acceleration of the atoms of $a_c = 70(10) \times 10^{-6} g$. Atomic fountain interferometers facilitate the measurement of significantly smaller differential accelerations and reach staggering precisions of the order $\Delta g/g \sim 10^{-10}$ [3, 4, 32]. Our measurement, however, provides, to the best of our knowledge, the smallest absolute value for an acceleration that is measured directly with ultracold atom interferometry.

We estimate possible sources of measurement errors, fluctuations and dephasing mechanisms. Fluctuations of a homogeneous magnetic field will only slightly change the interaction strength of our BEC, but deviations of the magnetic field gradient can induce additional accelerations and alter the measurement result. In our setup, small deviations of the magnetic field gradient can occur as the wave packets move during an interferometer sequence away from the original position with optimized levitation. We estimate from our numerical magnetic field simulation that our coil design causes a relative increase of the field gradient of 2×10^{-6} for a vertical position shift of $50 \mu\text{m}$. In addition, the quadratic Zeeman effect induces another deviation of the levitation force of 6×10^{-6} for the same position shift. As a result, the upper and lower wave packets experience a position-dependent acceleration, which increases the separation of the wave packets before the inversion pulse, and which reduces the convergence after the inversion pulse. Similar to our measurements in section 5, we would expect the final displacement of the wave packets to cause horizontal fringes in the absorption images, which we do not observe. As a result, we conclude that the vertical force gradients are negligible for the time scales of our interferometer.

In addition, the position-dependent magnetic field strength causes an almost linear change of the scattering length of approximately $\pm 10 a_0$ over $50 \mu\text{m}$ (see also section 6). As a result, the atoms in the upper wave packet experience a stronger interaction and faster phase evolution than atoms in the lower wave packet. Assuming constant densities and a linear change of the scattering length, we would expect the phase shift between the wave packets to increase with ΔT^2 , and it would be difficult to distinguish this effect from a phase evolution due to

acceleration. However, in our setup the wave packets expand after release and the densities decrease strongly over a timescale of $1/\omega_{x,y,z} \approx 10$ ms. The position-dependent scattering length would result in a change of the oscillation frequencies within 10–15 ms in figures 3(a)–(c), which we do not observe, and we conclude that the phase shift due to a position-dependent scattering length is below our sensitivity for this measurement.

Fluctuations of the acceleration of the BEC can be caused by time-dependent changes of B_0 and $\partial_z B$, either due to external magnetic fields or due to the finite stability of the currents in our coils. We determine a current reproducibility of 1.4×10^{-6} by measuring the standard deviation of the current during the interferometer sequence over 60 consecutive cycles. For each cycle, the current measurement averages over 80 ms. We believe that the current reproducibility will eventually set the limiting precision for our interferometric measurements with levitated atoms. While it is in principle possible to increase the current reproducibility by 1–2 orders of magnitude by improving our current regulation electronics, it would be very hard to reach the precision of atomic fountain experiments. Nonetheless, we believe that magnetic levitation schemes will provide a valuable technological addition for precision measurements with ultracold atoms. Reducing gravitational acceleration to micro- g effectively removes the center-of-mass motion of the atoms, and it allows for a direct measurement of phase-shifts due to additional elements in the interferometer path. We demonstrate this approach in the next section by adding a focused laser beam in the upper path of the interferometer and by measuring its position-dependent phase shift on the atoms.

4.3. Detection of phase-shifting elements

Compared to fountain experiments, the center-of-mass motion of our wave packets is contained within a small spatial region of a few hundreds of μm , and it is straightforward to add additional phase shifting elements in the path of the wave packets. As a result, it is possible to use the levitated interferometer scheme to analyze additional potentials for the atoms with high precision. We demonstrate this approach by adding a horizontal laser beam (wavelength 1064 nm, waist 40 μm , power 29 μW) approximately 50 μm above the initial position of the atoms (figure 4(b)). This beam creates a Gaussian dipole potential with a depth of approximately 3 nK, and it introduces between the upper and lower wave packets a differential phase shift, which can be detected by the interferometer. In addition to a measurement of the AC Stark shift of the light field as in reference [33], our setup facilitates the study of the spatial dependence of the potential.

The effect of the laser beam on $P_0(t)$ is clearly visible in figure 4(a) when comparing the data sets with the beam (blue squares) and without the beam (red circles). For increasing duration T_1 , the upper wave packet passes twice through the laser beam and it samples increasing spatial sections of the potential. We adjusted the power of the beam to create a single oscillation of the phase for a wave packet that fully transverses the beam, resulting in a minimum of $P_0(t)$ at an evolution time $T_1 = 7$ ms in figure 4(a).

Constant propagation velocities of the wave packets during the evolution times T_1 and T_2 make it easy to relate the time to the position of the atoms. We use a numerical model to integrate the phase shift of the upper wave packet in the dipole potential of the laser beam over the interferometer path $z(t)$ and include the unperturbed phase shift as measured in section 4.2. Fitting the model parameters to our data set (blue line figure 4(a)), we determine a beam position of 45(1) μm , a waist of 37(4) μm and a beam power of 25(3) μW , which are in excellent agreement with the independently measured values.

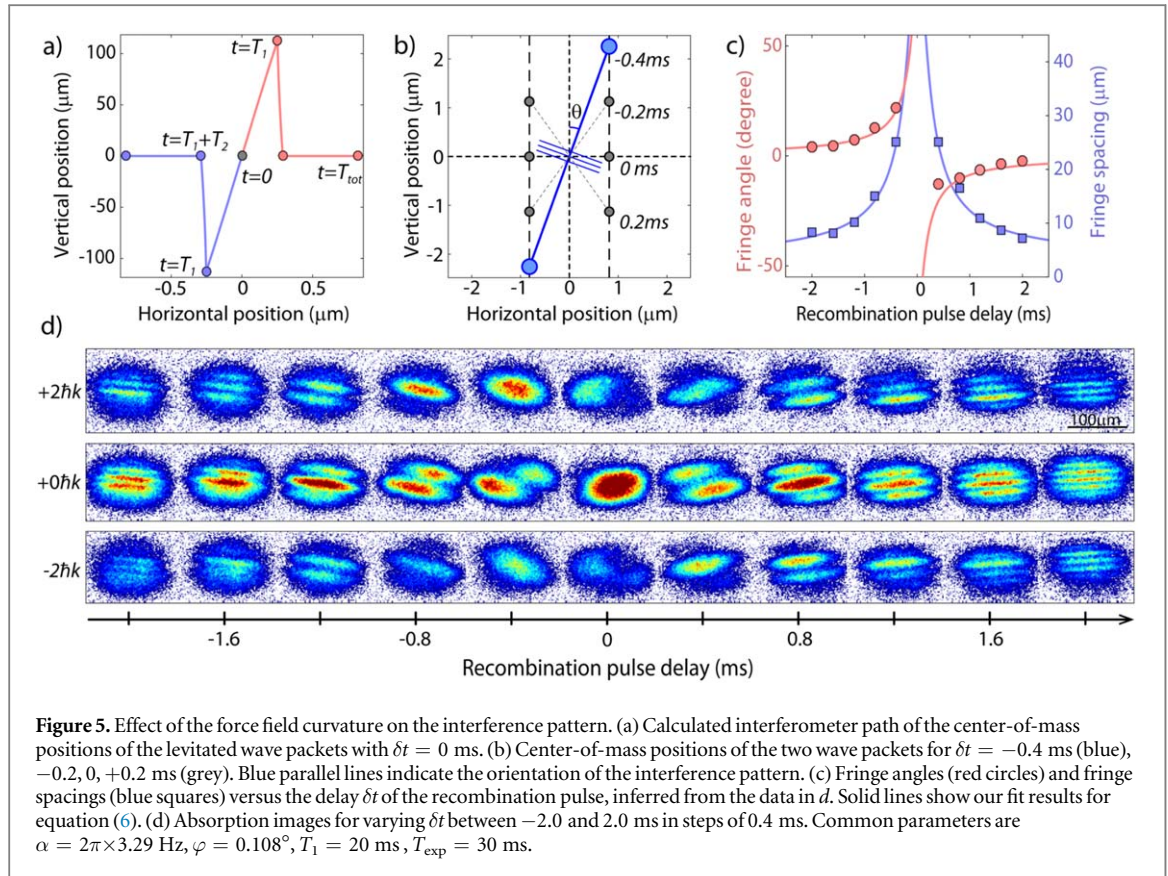
Our model neglects the spatial extent of the wave packets and we determine the phase shift at the center-of-mass position, whereas our experimental sequence averages over local phase shifts within the upper matter wave packet. Local phase shifts result in density variations in the profiles of the momentum modes in our absorption images, but measuring the total atom number in the momentum modes provides only the average phase shift of the wave packet.

5. Spatial curvature of the force field

Our magnetic field configuration does not only provide a vertical magnetic field gradient to levitate the atoms, but it also generates a weak, horizontal anti-trapping potential. This potential is a result of the spatial curvature of our quadrupole-like distribution of the magnetic field (see figure 1(b)). In this section, we demonstrate that the anti-trapping potential causes an additional interference pattern, which can be employed to measure the anti-trapping frequency or the angle between the lattice beam and the vertical field axis.

Within the quadrupole approximation it is possible to derive simple equations for the magnetic field and for the forces along the dashed horizontal line in figure 1(b) [19, 25, 34]

$$B_{\text{horz}}(r) = B_0 + \frac{2}{9} \frac{m^2}{\mu_B} \frac{g^2}{B_0} r^2, \quad F_{\text{horz}}(r) = m\alpha^2 r \quad \text{with} \quad \alpha = g \sqrt{\frac{m}{3\mu_B B_0}}. \quad (3)$$



Here, $r = \sqrt{x^2 + y^2}$ is the horizontal displacement of the atoms from the origin. The quadratic scaling of $B_{\text{horz}}(r)$ with r results in a weak, outwards-directed force in the horizontal plane. This anti-trapping effect can be associated with frequency α , and it causes a weak, position-dependent acceleration with a time-dependent horizontal position $r(t)$ and horizontal velocity $v_r(t)$ [35]:

$$\begin{aligned} r(t) &= r(0) \cosh(\alpha t) + \alpha^{-1} v_r(0) \sinh(\alpha t) \\ v_r(t) &= v_r(0) \cosh(\alpha t) + \alpha r_0 \sinh(\alpha t) \\ z(t) &= v_z(0)t + z(0). \end{aligned} \quad (4)$$

For this calculation we assume perfect levitation and linear vertical motion $z(t)$ during the interferometer sequence.

In an experimental setup there will always be a small angle φ between the lattice beam L1 and the vertical axis of the magnetic field, and a splitting pulse will always imprint a small velocity component $v_r(0) = (\hbar k_L / m) \sin(\varphi)$ along the horizontal direction. Consequently, a small horizontal displacement due to $v_r(0)$ results in an outwards-directed force on the wave packets in the anti-trapping potential, and in a finite horizontal displacement at the end of the interferometer sequence as illustrated in figure 5(a). The horizontal distance between the wave packets is typically two orders of magnitude smaller than the vertical displacement during the interferometer sequence, and both distances become comparable only in the proximity of the recombination pulse and during the expansion time. We illustrate the positions of the wave packets in figure 5(b) for small delay times of the recombination pulse $\delta t = T_2 - T_1$ with $T_1 = 20$ ms. Depending on δt , the orientation of the blue line connecting the wave packets changes from almost vertical for $\delta t = \pm 0.4$ ms to horizontal for $\delta t = 0$ ms. We define an angle θ , which is chosen to be positive clockwise and in the interval $[-90^\circ, 90^\circ]$, to indicate the orientation of the line, and we define $d(\delta t)$ to be the distance between the two wave packets.

In analogy to Young's double slit experiment [13, 36], the interference pattern of two wave packets at distance $d(\delta t)$ shows a fringe spacing d_F of

$$d_F = \pi \hbar t / (m d) + d_0. \quad (5)$$

Here, t is the total duration of the interferometer sequence with $t = T_1 + T_2 + T_{\text{pulse}} - \delta t + T_{\text{exp}}$, and $d_0 \geq 0$ is a constant phase shift that depends on the initial conditions such as the density distribution [37–39]. In our absorption images of the interfering wave packets for constant times T_1 , T_{exp} and varying delay δt (figure 5(d)), interference fringes with varying separation d_F and angle θ are clearly visible for all momentum modes p_0, p_{\pm} .

From the evolution of the fringes as a function of time delay δt , we infer properties of the curvature α and the angle φ . We simultaneously fit the fringe spacing in equation (5) and the fringe angle θ with $\theta(\delta t) = \arctan(z(\delta t)/r(\delta t))$. Here $z(\delta t)$ and $r(\delta t)$ are the vertical and horizontal positions of the wave packets for varying δt . We integrate the center-of-mass motion of the wave packets in equation (4) with starting conditions $r(0) = z(0) = 0$ over all steps of the interferometer sequence to determine $z(\delta t)$ and $r(\delta t)$

$$\begin{aligned} z(\delta t) &= -v_z(0) \delta t \\ r(\delta t) &= \frac{v_r(0)}{\alpha} \cosh(\alpha T_{\text{exp}}) [\sinh(\alpha T_1) \cosh(\alpha(T_1 + \delta t)) \\ &\quad + (\cosh(\alpha T_1) - 1) \sinh(\alpha(T_1 + \delta t))] \\ &\quad + \frac{v_r(0)}{\alpha} \sinh(\alpha T_{\text{exp}}) ([\sinh(\alpha T_1) \sinh(\alpha(T_1 + \delta t)) \\ &\quad + (\cosh(\alpha T_1) - 1) \cosh(\alpha(T_1 + \delta t))] + 1). \end{aligned} \quad (6)$$

Equations (6) contain two free parameters, the anti-trapping frequency α and the lattice angle φ , which can both be used to fit our data points in figure 5(c). We choose to constrain α and vary φ during the fitting procedure, as it is experimentally difficult to determine the laser beam angle with milliradian precision, and we independently measured α by observing center-of-mass oscillations of BECs in optical dipole traps. The fit results, represented by solid lines in figure 5(c), show good agreement with our data points, and we measure a lattice angle of $\varphi = 0.108(7)^\circ$ for $\alpha = 2\pi \times 3.29(5)$ Hz.

Note that α scales with $1/\sqrt{B_0}$ in equation (3), and we can use larger values for B_0 to reduce the anti-trapping effect, e.g. by tuning the interaction strength with a broad magnetic Feshbach resonance at 800 G [40]. However, it will be difficult to reduce α significantly due to its square-root dependence on B_0 . Instead, it is easier to compensate the anti-trapping effect with an additional dipole trap, as demonstrated in the next section.

6. Long expansion times

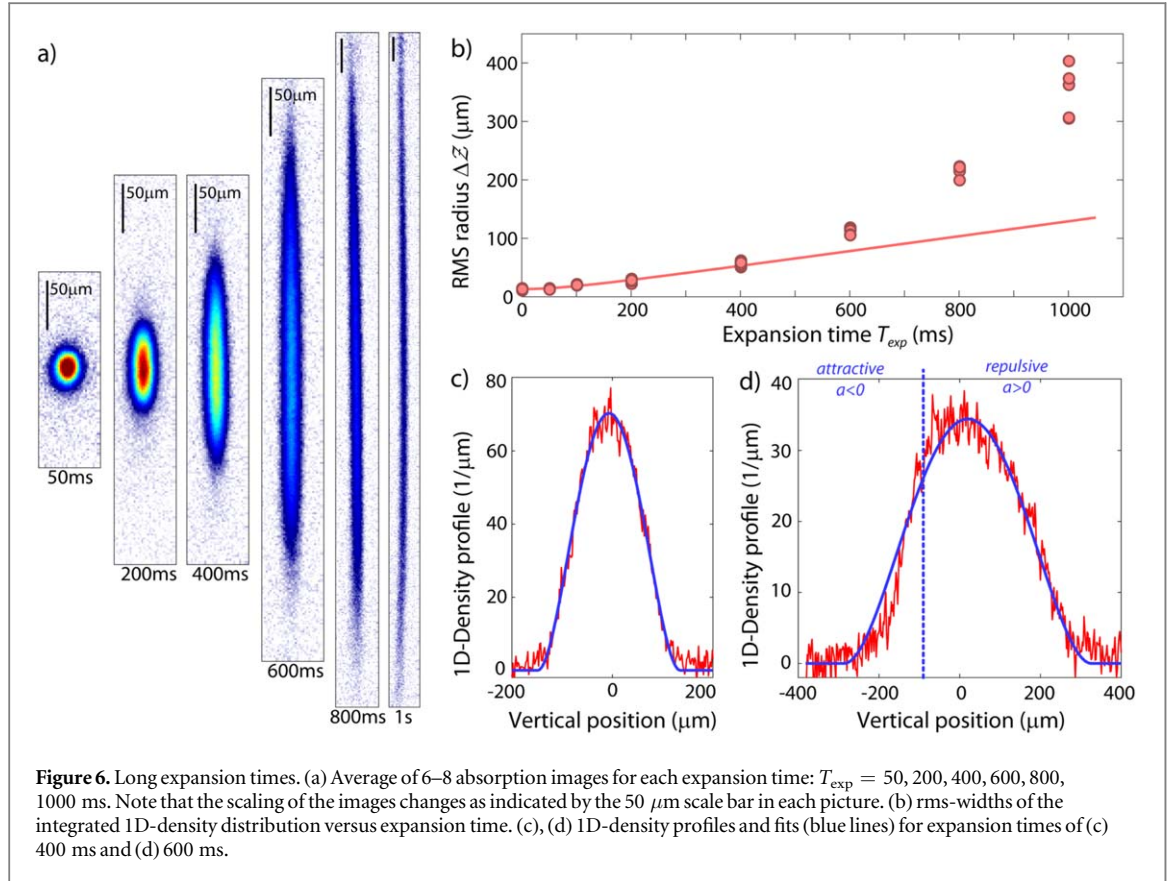
The sensitivity of an interferometric measurement increases with the evolution time of the wave packets [12], but even without the implementation of an interferometer scheme, long observation times of an expanding BEC facilitate a sensitive acceleration measurement. In this section, we demonstrate that magnetic levitation allows us to extend the expansion time of a BEC to 1 s, and we evaluate advantages and limitations of this scheme for precision measurements.

Typical expansion times for falling BECs are on the order of tens of milliseconds, often limited by the detection area of the imaging system, by the gravitational acceleration and by the expansion velocity of the gas. Usually, the expansion velocity of a quantum gas is not caused by the temperature of the gas but by repulsive interaction during the initial spreading. The current record for long observation times under milli-g acceleration is 1 s [12] with an expansion energy of 9 nK. The experiment was performed in a drop tower, and ballistic expansion was observed over approximately 500 ms, limited by stray magnetic fields.

In our experiment, we can reduce the interaction energy of the BEC by tuning the scattering length close to 0 a_0 by means of a magnetic Feshbach resonance (figure 1(c)). Further reduction of the expansion energy has been demonstrated by rapidly changing the scattering length from a positive value to 0 a_0 during trap release [25], but we refrain from using this trick to avoid excitations of the BEC during release. Our horizontal magnetic field curvature (section 5) introduces another limitation. During long observation times, the BEC expands horizontally into regions with a lower magnetic field gradient, causing a position-dependent sag of the density profile. In addition, small fluctuations of the horizontal magnetic field can break the symmetry and introduce slow horizontal drifts. We suppress both effects by keeping a vertical laser beam (H3 in figure 1(a)) on during the expansion time, thus observing free expansion only in the vertical direction.

In detail, we reduce the trap frequency by slowly transferring the atoms from a crossed dipole trap of beams S1, S2 to a crossed dipole trap of beams H1, H2, and H3 with final trap frequencies of $\omega_{x,y,z} = 2\pi \times (3.2, 3.4, 2.1)$ Hz, a scattering length of 15 a_0 and atom numbers of approximately 1.1×10^4 . Excitations of the BEC during the transfer are suppressed by smooth changes of the potential with a total transfer duration of 4 s. After an additional settling time of 1 s we switch off the horizontal beams H1 and H2 and study the expansion of the BEC in the vertical beam H3. The vertical trapping frequency of the laser beam H3 is approximately 25 mHz, and the resulting fractional reduction of the expansion width after 1 s is 6×10^{-4} , which is far below our measurement sensitivity for the width of the BEC.

The expansion of the BEC in the vertical direction is clearly visible on absorption images (figure 6(a)) for expansion times 0–1000 ms, and horizontally-integrated 1D density profiles for expansion times of 400 ms and 600 ms are given in figures 6(c) and (d). Although the trapped BEC is initially only weakly confined with almost



symmetric trap frequencies, it changes dimensionality during the expansion process in the vertical beam. The density of the BEC decreases strongly during the vertical expansion, and the chemical potential becomes smaller than the transversal harmonic oscillator energy $\hbar\omega_{x,y}$, as required for a quasi-1D description [41]. As a result, we do not expect a shape-preserving spreading of the density distribution for a 1D expansion because the BEC passes through various interaction regimes as its density decreases [42, 43]. For illustration, we show a fit to the upper 80% of the 1D-density profiles $n(z)$ for the ‘3D cigar’-regime [44] (figure 6(c)), but we refrain from a complete analysis of the density profiles, which is beyond the scope of this article. Instead, we quantify the width of the expanding BEC with the root-mean-square (rms) radius $\Delta Z = \left(\frac{1}{N} \int n(z)(z - \bar{z})^2\right)^{1/2}$ to provide an estimate of the expansion velocity (red circles figure 6(b)). Here, \bar{z} is the center-of-mass position of the atoms. We observe an initial interaction driven expansion and a ballistic flight for $T_{\text{exp}} \leq 400$ ms with an rms expansion velocity of $v_{\text{rms}} = 0.128(5) \text{ mm s}^{-1}$ and a corresponding kinetic energy of $mv_{\text{rms}}^2/2 = 1/2 k_B \times 260(20) \text{ pK}$. We note that this is the expansion energy of the BEC component, but not the initial temperature of the trapped quantum gas.

Similar to reference [12], we find an accelerated expansion for longer expansion times, $T_{\text{exp}} > 500$ ms. We expect that the dominant source of the accelerated expansion is the curvature of our levitation gradient due to the quadratic Zeeman effect and due to our coil design, as discussed in section 4.2. However, the density profiles of the atoms on the absorption images indicate two other contributions. We observe small radial oscillations for long expansion times after release from the trap in the guiding beam H3 (see image $T_{\text{exp}} = 1$ s in figure 6(a)). Those oscillations can couple to the vertical motion or they can distort the radially integrated density distribution. In addition, we observe asymmetric 1D density profiles $n(z)$ for $T_{\text{exp}} > 500$ ms (figure 6(d)). The profiles show a slower expansion velocity for the lower part of the cloud than for the upper part. We assume that this effect is caused by the position-dependent scattering length due to our magnetic field gradient. The zero-crossing of a is indicated in figure 6(d) by a dashed blue line. This asymmetric expansion of a BEC with position-dependent scattering length requires further investigation that is beyond the scope of this article. We find small position fluctuations for long expansion times $T_{\text{exp}} > 400$ ms of the BEC due to the finite current stability for the magnetic field gradient (section 4.2). For illustration, we re-centered the center-of-mass position in the absorption images for the averaging process in figure 6(a), but all other data in figures 6(b)–(c) results from the analysis of individual absorption images.

7. Conclusion

In conclusion, we experimentally studied the benefits and challenges of the use of magnetic levitation schemes for interferometric precision measurements with ultracold atoms. We employed a Michelson-type interferometer setup with BECs with tuneable interaction and magnetic levitation to demonstrate absolute acceleration measurements in the micro-*g* regime and we used the negligible center-of-mass motion of levitated atoms to study the position-dependent phase shift of the dipole potential of a focused laser beam. Moreover, we demonstrated expansion times of 1 s for a BEC, which is comparable to current drop tower experiments, and we used an extrapolation method for the fringe patterns to study the curvature of a force field that acts perpendicularly to our interferometer setup.

In our setup, limitations of the sensitivity arise from magnetic field fluctuations due to the current regulation, and from position-dependent interactions and magnetic field gradients. Although the sensitivity in our setup is significantly lower than the sensitivity of atomic fountain experiments, we believe that levitation schemes provide interesting features with the prospect of technical applications. Cancelling gravitational acceleration offers the possibility to combine long observation times with compact interferometer setups. Interesting applications are the measurement of local variations of electric and magnetic fields, and of mean field effects due to atomic interactions.

Acknowledgments

The authors would like to thank E Riis and PF Griffin for helpful discussions. We acknowledge financial support by the EU through the Collaborative Project QuProCS (Grant Agreement 641277). AdC acknowledges financial support by EPSRC and SFC via the International Max-Planck Partnership. The work was also supported in part by the EPSRC Programme Grant DesOEQ (Grant No EP/P009565/1).

References

- [1] Gustavson T L, Bouyer P and Kasevich M A 1997 *Phys. Rev. Lett.* **78** 2046
- [2] Dutta I, Savoie D, Fang B, Venon B, Garrido Alzar C L, Geiger R and Landragin A 2016 *Phys. Rev. Lett.* **116** 183003
- [3] Peters A, Chung K Y and Chu S 1999 *Nature* **400** 849
- [4] Hardman K S, Everitt P J, McDonald G D, Manju P, Wigley P B, Sooriyabandara M A, Kuhn C C N, Debs J E, Close J D and Robins N P 2016 *Phys. Rev. Lett.* **117** 138501
- [5] Bouchendira R, Cladé P, Guellati-Khélifa S, Nez F and Biraben F 2011 *Phys. Rev. Lett.* **106** 080801
- [6] Rosi G, Sorrentino F, Cacciapuoti L, Prevedelli M and Tino G M 2014 *Nature* **510** 518
- [7] Prevedelli M, Cacciapuoti L, Rosi G, Sorrentino F and Tino G M 2014 *Phil. Trans. R. Soc. A* **372** 20140030
- [8] Hamilton P, Jaffe M, Haslinger P, Simmons Q, Müller H and Khoury J 2015 *Science* **349** 849
- [9] Cronin A D, Schmiedmayer J and Pritchard D E 2009 *Rev. Mod. Phys.* **81** 1051
- [10] Debs J E, Altin P A, Barter T H, Döring D, Dennis G R, McDonald G, Anderson R P, Close J D and Robins N P 2011 *Phys. Rev. A* **84** 033610
- [11] Chiow S-W, Kovachy T, Chien H-C and Kasevich M A 2011 *Phys. Rev. Lett.* **107** 130403
- [12] van Zoest T et al 2010 *Science* **329** 1540
- [13] Müntinga H 2013 *Phys. Rev. Lett.* **110** 093602
- [14] Geiger R et al 2011 *Nat. Commun.* **2** 474
- [15] Barrett B, Antoni-Micollier L, Chichet L, Battelier B, Lévêque T, Landragin A and Bouyer P 2016 *Nat. Commun.* **7** 13786
- [16] Becker D et al 2018 *Nature* **562** 391
- [17] Anderson M H, Ensher J R, Matthews M R, Wieman C E and Cornell E A 1995 *Science* **269** 198
- [18] Han D J, DePue M T and Weiss D S 2001 *Phys. Rev. A* **63** 023405
- [19] Weber T, Herbig J, Mark M, Nägerl H-C and Grimm R 2003 *Science* **299** 232
- [20] Wang Y-J, Anderson D Z, Bright V M, Cornell E A, Diot Q, Kishimoto T, Prentiss M, Saravanan R A, Segal S R and Wu S 2005 *Phys. Rev. Lett.* **94** 090405
- [21] Garcia O, Deissler B, Hughes K J, Reeves J M and Sackett C A 2006 *Phys. Rev. A* **74** 031601(R)
- [22] McDonald G D, Kuhn C C N, Bennetts S, Debs J E, Hardman K S, Johnsson M, Close J D and Robins N P 2013 *Phys. Rev. A* **88** 053620
- [23] Marti G E, Olf R and Stamper-Kurn D M 2015 *Phys. Rev. A* **91** 013602
- [24] Burke J H T, Deissler B, Hughes K J and Sackett C A 2008 *Phys. Rev. A* **78** 023619
- [25] Kraemer T, Herbig J, Mark M, Weber T, Chin C, Nägerl H-C and Grimm R 2004 *Appl. Phys. B* **79** 1013
- [26] Kerman A J, Vuletic V, Chin C and Chu S 2000 *Phys. Rev. Lett.* **84** 439
- [27] Gould P L, Ruff G A and Pritchard D E 1986 *Phys. Rev. Lett.* **56** 827
- [28] Rasel E M, Oberthaler M K, Batelaan H, Schmiedmayer J and Zeilinger A 1995 *Phys. Rev. Lett.* **75** 2633
- [29] Storey P and Cohen-Tannoudji C 1994 *J. Phys. II* **4** 1999
- [30] Wu S, Wang Y-J, Diot Q and Prentiss M 2005 *Phys. Rev. A* **71** 043602
- [31] Robertson B I, MacKellar A R, Halket J, Gribbon A, Pritchard J D, Arnold A S, Riis E and Griffin P F 2017 *Phys. Rev. A* **96** 053622
- [32] Asenbaum P, Overstreet C, Kovachy T, Brown D D, Hogan J M and Kasevich M A 2017 *Phys. Rev. Lett.* **118** 183602
- [33] Deissler B, Hughes K J, Burke J H T and Sackett C A 2008 *Phys. Rev. A* **77** 031604
- [34] Sackett C A 2006 *Phys. Rev. A* **73** 013626
- [35] Herbig J, Kraemer T, Mark M, Weber T, Chin C, Nägerl H-C and Grimm R 2003 *Science* **301** 1510
- [36] Andrews M R, Townsend C G, Miesner H-J, Durfee D S, Kurn D M and Ketterle W 1997 *Science* **275** 637

- [37] Wallis H, Röhr A, Naraschewski M and Schenzle A 1997 *Phys. Rev. A* **55** 2109
- [38] Röhr A, Naraschewski M, Schenzle A and Wallis H 1997 *Phys. Rev. Lett.* **78** 4143
- [39] Simsarian J E, Denschlag J, Edwards M, Clark C W, Deng L, Hagley E W, Helmerson K, Rolston S L and Phillips W D 2000 *Phys. Rev. Lett.* **85** 2040
- [40] Berninger M, Zenesini A, Huang B, Harm W, Nägerl H-C, Ferlaino F, Grimm R, Julienne P S and Hutson J M 2013 *Phys. Rev. A* **87** 032517
- [41] Petrov D S, Gangardt D M and Shlyapnikov G V 2004 *J. Phys. IV* **116** 5
- [42] Öhberg P and Santos L 2002 *Phys. Rev. Lett.* **89** 240402
- [43] Pedri P, Santos L, Öhberg P and Stringari S 2003 *Phys. Rev. A* **68** 043601
- [44] Menotti C and Stringari S 2002 *Phys. Rev. A* **66** 043610

Appendix C

Publication: Excitation modes of bright matter wave solitons

Published in Phys. Rev. Lett.

DOI: [10.1103/PhysRevLett.123.123602](https://doi.org/10.1103/PhysRevLett.123.123602)

Excitation Modes of Bright Matter-Wave Solitons

Andrea Di Carli, Craig D. Colquhoun, Grant Henderson, Stuart Flannigan, Gian-Luca Oppo,
Andrew J. Daley, Stefan Kuhr, and Elmar Haller[✉]
Department of Physics and SUPA, University of Strathclyde, Glasgow G4 0NG, United Kingdom

 (Received 2 May 2019; revised manuscript received 18 July 2019; published 17 September 2019)

We experimentally study the excitation modes of bright matter-wave solitons in a quasi-one-dimensional geometry. The solitons are created by quenching the interactions of a Bose-Einstein condensate of cesium atoms from repulsive to attractive in combination with a rapid reduction of the longitudinal confinement. A deliberate mismatch of quench parameters allows for the excitation of breathing modes of the emerging soliton and for the determination of its breathing frequency as a function of atom number and confinement. In addition, we observe signatures of higher-order solitons and the splitting of the wave packet after the quench. Our experimental results are compared to analytical predictions and to numerical simulations of the one-dimensional Gross-Pitaevskii equation.

DOI: 10.1103/PhysRevLett.123.123602

The dispersionless propagation of solitary waves is one of the most striking features of nonlinear dynamics, with multiple applications in hydrodynamics, nonlinear optics, and broadband long-distance communications [1]. In fiber optics, one-dimensional (1D) “bright” solitons, i.e., solitons presenting a local electric field maximum with one-dimensional propagation, have been observed [2]. They exhibit a dispersionless flow and excitation modes such as breathing or higher-order modes [2–4]. Matter waves can also display solitary dispersion properties. Typically, bright matter-wave solitons are created in quasi-1D systems by quenching the particle interaction in a Bose-Einstein condensate (BEC) from repulsive to attractive [5]. Recent experiments demonstrated the collapse [6], collisions [7], reflection from a barrier [8], and the formation of trains [9–11] of bright solitons.

In this Letter, we experimentally study the excitation modes of a single bright matter-wave soliton. In previous studies, other dynamical properties have been observed, such as the center-of-mass oscillation of solitons in an external trap [7] and excitations following the collapse of attractive BECs [6,12]. Here, we probe the fundamental breathing mode of a single soliton by measuring its oscillation frequency and the time evolution of its density profile. In addition, we observe signatures of higher-order matter-wave solitons, which can be interpreted as stable excitations with periodic oscillations of the density profile and phase, or as a bound state of overlapping modes [3,13].

The shape-preserving evolution of a matter-wave soliton is due to a balancing of dispersive and attractive terms in the underlying 3D Gross-Pitaevskii equation (GPE) [14]. For quasi-1D systems with tight radial confinement, we can approximate the matter wave in the 3D GPE by the product of a Gaussian wave function for the radial direction and a function $f(z)$ for the longitudinal direction (see Ref. [15]). Depending on the ansatz for the Gaussian with either constant or varying radial sizes, $f(z)$ satisfies either the 1D GPE or the nonpolynomial Schrödinger equation [18]. We make reference to the analytical solutions of the 1D GPE in the Letter, but use both equations in our numerical simulations [15].

For the 1D GPE, an ansatz for the normalized longitudinal wave function $f(z)$ is of the form

$$f(z) = \frac{1}{\sqrt{2}l_z} \operatorname{sech}\left(\frac{z}{l_z}\right), \quad (1)$$

with a single parameter l_z that determines both the longitudinal size and the amplitude of the soliton. Solitons form with a value of l_z that minimizes the total energy and that provides a compromise between the kinetic and the interaction energies. This is illustrated in Fig. 1(b), which shows the energy of the wave packet for varying sizes l_z [19]. The kinetic energy provides a potential barrier for small l_z that prevents the collapse of the soliton, while its spreading is inhibited by the interaction energy, which increases for large l_z .

Even without an external longitudinal potential, the soliton is stable against small perturbations of l_z . In a way, a bright matter-wave soliton creates its own trapping potential, which defines its size and excitation modes. Variational methods provide accurate predictions of its size

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.




FIG. 1. Experimental setup and oscillation measurements. (a) Sketch of the experimental setup. Inset: Density profiles for a BEC (solid red line) and for a soliton (dashed blue line). (b) Total energy of a soliton, $a = -5.2a_0$, $\omega_r = 2\pi \times 95$ Hz, $N = 2000$, with an external trap, $\omega_z = 2\pi \times 5$ Hz (dashed blue line), and without external trap, $\omega_z = 0$ Hz (solid red line). (c) Absorption images after a free-expansion time of 16 ms [from dataset with circles in (d)], integrated density profile for $t = 60$ ms (blue line) and fit (dashed red line). (d) Oscillations of a quantum gas after the quench procedure. Blue diamonds, quench of only ω_z for a BEC (Q1); red circles, additional interaction quench to create soliton (Q3); green squares, optimized quench parameters to minimize breathing of the soliton (Q2). Uncertainty intervals indicate ± 1 standard error.

at the energy minimum which can be calculated analytically [13,20] or numerically [19]. For the fundamental solution (order $n = 1$) of the 1D GPE with an atom number N , s -wave scattering length a , and radial trapping frequency ω_r , the size l_z corresponds to the healing length at the peak density of the soliton, i.e., $l_z^{(n=1)} = a_r^2 / (N|a|)$ [13,19]. Here, $a_r = \sqrt{\hbar / (m\omega_r)}$ is the radial harmonic oscillator length. Small deviations of l_z close to the energy minimum lead to oscillations of the soliton size. We use those oscillations resulting from an initial mismatch of l_z to experimentally measure the self-trapping frequency of the soliton potential.

Our experimental starting point is a Bose-Einstein condensate of 500–2000 cesium (Cs) atoms in the state $|F = 3, m_F = 3\rangle$ at scattering length of $a = +7a_0$, where a_0 is Bohr's radius. The BEC is levitated by a magnetic field gradient, and it is confined by an optical dipole trap formed by the horizontal and vertical laser beams L_H and L_V [Fig. 1(a)]. An additional magnetic offset field allows us to tune the scattering length by means of a broad magnetic Feshbach resonance [21]. Details about our experimental setup, the levitation scheme, and the removal of atoms can be found in Refs. [15,22].

Our matter-wave solitons are confined to a quasi-1D geometry with almost free propagation along the horizontal direction and strong radial confinement of $\omega_r = 2\pi \times 95$ Hz provided by laser beam L_H . They are generated with a quench of the scattering length towards attractive interaction ($a_i \rightarrow a_f$), and by a reduction of the longitudinal trap frequency ($\omega_{z,i} \rightarrow \omega_{z,f}$). When changing a and ω_z independently, the quenches excite inward and outward motions, respectively. Usually, it is desirable to minimize the excitations of the soliton by matching the initial Thomas-Fermi density profile of the BEC closely to the density profile of the soliton [inset of Fig. 1(a)]. However, we deliberately mismatch the quench parameters to create breathing oscillations of the soliton in order to study its self-trapping potential. Quenches with different parameters are labeled by the symbols $Q1$ – $Q7$ (see Ref. [15]). Following an evolution time t in quasi-1D and after a short period of 16 ms of expansion in free space, we take absorption images to determine the density profile of the atoms [Fig. 1(c)]. The cloud size $l_z(t)$ is determined by fitting the function $A(\text{sech}(z/B))^2$ to the integrated 1D-density profiles with fit parameters A and B [15].

The response of the atomic cloud to the different quenches is presented in Fig. 1(d). We first quench only the longitudinal confinement by 25% to $\omega_{z,f} = 2\pi \times 4.3(2)$ Hz (quench $Q1$ in Ref. [15]) while keeping the repulsive interaction strength constant [Fig. 1(d), diamonds]. The BEC starts an outwards motion with an oscillation frequency of $2\pi \times 7.5(1)$ Hz $\approx \sqrt{3}\omega_{z,f}$ as expected for a BEC in the Thomas-Fermi regime [23,24]. In a second measurement, we additionally quench the interaction strength a_f to $-5.4a_0$ and increase $\omega_{z,i}$ to match the initial size of the BEC to the expected size of the soliton [Q2, Fig. 1(d), squares]. As a result, we observe almost dispersionless solitons with a linear increase of the cloud size of $0.7(3)$ $\mu\text{m/s}$ [Fig. 1(d), green line]. Finally, we deliberately mismatch the initial size of the BEC by reducing $\omega_{z,i}$ (Q3), and generate small-amplitude oscillations of the soliton with a frequency ω_{sol} of $2\pi \times 12.8(4)$ Hz [Fig. 1(d), circles]. This breathing frequency of the soliton is significantly larger than any breathing frequency of a BEC or of noninteracting atoms, $2\omega_{z,f} = 2\pi \times 8.6(3)$ Hz. We observe no discernible oscillation in the radial direction after the quenches.

In a second experiment, we demonstrate that the breathing frequency ω_{sol} depends on the interaction term Na in the 1D GPE, a property typical of the nonlinear character of the soliton. We choose to change N , since the initial removal process is independent of the interaction quench, and we can study ω_{sol} without changing the quench protocol [Q4, Fig. 2(a), circles]. The measured values of ω_{sol} decrease for lower N , and they approach the breathing frequency $2\omega_{z,f}$ for noninteracting atoms in a harmonic trap [Fig. 2(a), dashed line].

We compare our experimental data points to two theoretical models. In a numerical simulation of the 1D




FIG. 2. Breathing frequency ω_{sol} of the soliton. (a) Atom number dependence (Q4). Red circles, experimental data; the uncertainty bars for the atom number indicate the standard deviation of N over the first 100 ms of each frequency measurement. Blue triangles, simulation of the 1D GPE [15]. Red line, analytical approximation [13,15]. Dashed gray line, oscillation frequency of a noninteracting gas, $2\omega_{z,f}$. (b) Dependence of ω_{sol} on the trap frequency (Q5). Red circles, experimental data points for $N \approx 1450$. Blue area, simulation of the 1D GPE for $N = 1300-1500$. Red line, analytical approximation. Dashed gray line, $2\omega_{z,f}$.

GPE, we use the ansatz in Eq. (1) to set the starting conditions, and we determine the breathing frequency from a spectral analysis of the time evolution of the wave function [15] [Fig. 2(a), triangles]. In addition, we use an analytical approximation for the breathing frequency (red line) calculated with a Lagrangian variational analysis at the energy minimum of the 3D GPE [13,15]. We find that both models agree well with the trend of the measurements of ω_{sol} , although our experimental data points are systematically lower for large N than our theoretical predictions. We speculate that this is due to nonharmonic contributions to the energy of the soliton on the breathing oscillations for finite oscillation amplitudes [Fig. 1(b)].

To determine the influence of the trapping potential, we measure the variation of ω_{sol} as we reduce the longitudinal trapping frequency $\omega_{z,f}$ (Q5). Two regimes of ω_{sol} can be identified in Fig. 2(b) for varying the values of $\omega_{z,f}$. For large values of $\omega_{z,f}$, the trap dominates the breathing of the soliton and ω_{sol} increases like $2\omega_{z,f}$. For small values of $\omega_{z,f}$, interactions dominate the breathing of the soliton and ω_{sol} reaches a constant value. This offset of the breathing frequency is a result of the “self-trapping” potential of a free soliton.

Again, we compare the experimental results with our theoretical model [Fig. 2(b), red line] and the numerical

simulations of the 1D GPE. The blue band in Fig. 2(b) indicates the simulated frequencies for $N = 1300$ to $N = 1500$. The simulation predicts a lower breathing frequency for the free soliton than the analytical approximation, but all curves are within the uncertainty range of the experimental data.

External trapping potentials can in principle alter the soliton dynamics [7,25,26], causing, e.g., modulations of the soliton’s tails due to residual nonautonomous terms of the 1D GPE in a harmonic potential [27]. For the following experiments, however, we employ trap frequencies that are significantly smaller than the observed oscillation frequencies of the soliton ($2\omega_z < \omega_{\text{sol}}$) to decouple the influence of the trapping potential. In summary, for small-amplitude oscillations we find good agreement of ω_{sol} between our experimental results and analytical and numerical predictions based on the 1D GPE (and nonpolynomial Schrödinger equation [15]).

Breathing oscillations of l_z close to the equilibrium size are not the only possible excitation modes of solitons. The existence of higher-order solitons has been predicted in the nonlinear Schrödinger equation [3], and has been observed for optical solitons in silica-glass fibers [2,4]. A soliton of order n can be interpreted as a bound state of n strongly overlapping solitons [13]. By exploiting the equivalence of the nonlinear Schrödinger equation and 1D GPE, similar effects were later proposed for bright matter-wave solitons [13,28], where it was suggested that n th-order solitons can be generated by a rapid increase of the attractive interaction strength by a factor n^2 . Similarly, our simulations of the 1D GPE show that higher-order solitons can be created for an increased initial size of the wave packet. An n th-order soliton forms for a sech-shaped wave function with an initial size $l_z^{(n)}$ that is the n^2 multiple of the healing length $l_z^{(1)}$, i.e., $l_z^{(n)} = n^2 l_z^{(1)}$ [15].

Within the 1D GPE theory, both creation methods result in the periodic development of multi-peaked structures for higher-order solitons [3,29]; e.g., they create a sharp central peak with side wings for a second-order soliton [Fig. 3(a)] and a double peak for a third-order soliton [15]. Sizes and interaction quenches that do not fulfil the previous conditions lead to a “shedding” of the atomic density in the z direction. The wave packet oscillates and loses particles until its size and shape match the next (lower n) higher-order soliton [3]. For a second-order soliton, the predicted oscillation period $T^{(2)}$ is [13]

$$T^{(2)} = \frac{8\pi}{\hbar} m \left(\frac{a_r^2}{N|a_f|} \right)^2. \quad (2)$$

Recently, excitation modes of higher order have also been used as a test bed for various theoretical models beyond GP theory. The fragmentation of solitons with an increased initial width was predicted within the multi-configurational time-dependent Hartree method for bosons

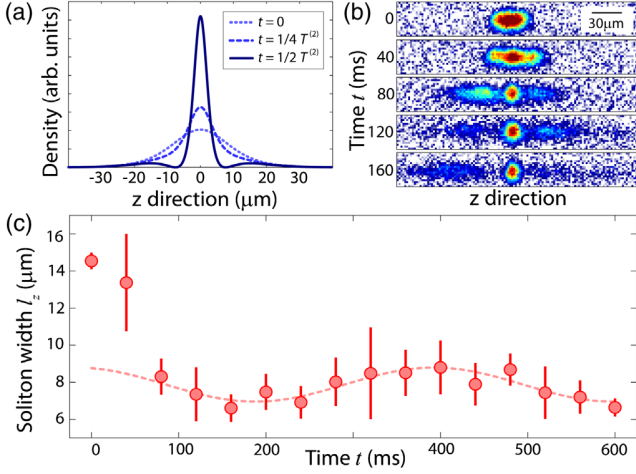


FIG. 3. Time evolution after a strong quench of interactions and trap frequency ($Q6$). (a) 1D GPE simulation of the density profiles for a second-order soliton with 1100 atoms, $a_f = -5.3a_0$, and with an oscillation period $T^{(2)}$ of 432 ms. (b) Absorption images at time t after the quench and after 11 ms of free expansion. (c) Time evolution of the measured width l_z of the central wave packet (red circles), sinusoidal fit with period 420(30) ms (dashed red line). The uncertainty intervals indicate ± 1 standard deviation.

[30] and critically discussed [31], and the influence of quantum effects on the dissociation process was investigated [32–34].

Here, we apply two different quench protocols to study the evolution of strongly excited solitons. Depending on the initial size and the quench parameters, we observe shedding and fragmentation of the wave packet, and we measure oscillation frequencies that indicate the creation of higher-order solitons. To demonstrate the effect of a strong quench of an elongated BEC, we increase a_i and reduce $\omega_{z,i}$ before ramping a and ω_z to $-5.3a_0$ and $2\pi \times 0.0(6)$ Hz in 13 ms ($Q6$). Our quench induces an initial spreading of the wave packet, followed first by a strong shedding of atoms, and then by the formation of a soliton that contains approximately 1/3 of the initial atom number [Fig. 3(b)]. We determine the soliton width and find a slow oscillation of $l_z(t)$ with a frequency of $2\pi \times 2.4(2)$ Hz [Fig. 3(c)]. This frequency is significantly smaller than the expected breathing frequency of first-order solitons, $2\pi \times 6.0$ Hz, and it matches well to the expected frequency of $2\pi \times 2.3$ Hz for second-order solitons in Eq. (2).

Observing shedding and oscillations agrees with the predictions for higher-order solitons within the 1D GPE [3]; however, we find a strong dependence on details of the quench protocol and on the dynamical evolution during the quench. For a closer match to theoretical works [13], we implement a double-quench protocol, with a first quench to generate a soliton with weak attractive interaction, $a_f = -0.8a_0$, $\omega_f = 2\pi \times 1.4(2)$ Hz, and, after a settling time of 25 ms, a second quench of only the interaction

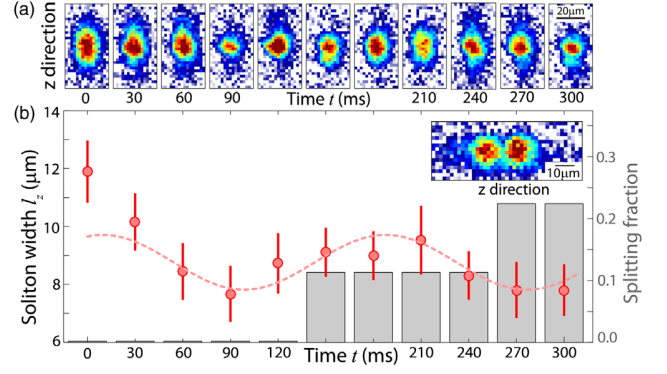


FIG. 4. Second-order soliton and splitting after the double quench $Q7$. (a) Absorption images at time t after the quench and after 7 ms of free expansion. (b) Time evolution of the measured width l_z of the central wave packet (red circles), sinusoidal fit with period 180(20) ms (dashed red line). The expected period from the 1D GPE simulations is 192 ms. The histogram counts the fraction of images showing a splitting of the wave function (9 repetitions per time step). Inset: Absorption image of a split matter wave for $t = 210$ ms.

strength, $a_f = -4.6a_0$ ($Q7$). Starting with approximately 2200 atoms, we observe no shedding but a small loss of 300 atoms during the first 60 ms. The vertical density profiles in our absorption images [Fig. 4(a)] resemble the expected profiles of a second-order soliton [Fig. 3(a)], and the vertical width of the wave packet oscillates with a frequency of $2\pi \times 5.6(6)$ Hz [Fig. 4(b)], which matches the expected frequency of $2\pi \times 5.2$ Hz for second-order solitons ($2\pi \times 13.2$ Hz for the first-order solitons).

For both measurements [Figs. 3(c) and 4(b)], a small percentage of absorption images show a splitting of the soliton into two fragments [inset of Fig. 4(b)], and they are omitted from the fitting procedure. Because of the destructive nature of our absorption images it is difficult to conclude on the evolution and on the cause of the splitting process. A double-peak structure in the density profile can indicate the generation of a third-order soliton, fragmentation due to quantum effects, or simply an insufficient technical control of our quench parameters. For our setup, the control of horizontal magnetic field gradients to avoid longitudinal accelerations is especially challenging [22]. The percentage of images that show a splitting of the wave packet increases for longer evolution times, and we indicate their fraction in Fig. 4(b) with a histogram.

In conclusion, we experimentally studied the creation and the excitation of breathing modes of bright matter-waves solitons in a quasi-one-dimensional geometry after a quench of interaction and longitudinal confinement. We measured the “self-trapping” frequency ω_{sol} for first-order solitons and its dependence on N and ω_z . For stronger excitations and for a double-quench protocol, we observed signatures of second-order solitons and the shedding and splitting of the wave function. Further measurements of the

splitting process and the damping of the oscillations due to shedding are necessary to distinguish technical fluctuations from higher-order solitons and fragmentation due to quantum effects [32–34].

The data used in this publication are openly available at the University of Strathclyde KnowledgeBase [35].

We thank L. D. Carr for helpful initial discussions during his visit. We acknowledge support by the EU through “QuProCS” (GA 641277) and the ETN “ColOpt” (GA 721465), and by the EPSRC Programme Grant “DesOEQ” (EP/P009565/1). A. D. C. acknowledges financial support by EPSRC and SFC via the IMPP. G. H. acknowledges support from a Carnegie Undergraduate Vacation Scholarship.

-
- [1] M. P. T. Dauxois, *Physics of Solitons* (Cambridge University Press, Cambridge, England, 2006).
- [2] L. F. Mollenauer, R. H. Stolen, and J. P. Gordon, *Phys. Rev. Lett.* **45**, 1095 (1980).
- [3] J. Satsuma and N. Yajima, *Prog. Theor. Phys. Suppl.* **55**, 284 (1974).
- [4] R. H. Stolen, W. J. Tomlinson, and L. F. Mollenauer, *Opt. Lett.* **8**, 186 (1983).
- [5] L. Khaykovich, F. Schreck, G. Ferrari, T. Bourdel, J. Cubizolles, L. D. Carr, Y. Castin, and C. Salomon, *Science* **296**, 1290 (2002).
- [6] E. A. Donley, N. R. Claussen, S. L. Cornish, J. L. Roberts, E. A. Cornell, and C. E. Wieman, *Nature (London)* **412**, 295 (2001).
- [7] J. H. V. Nguyen, P. Dyke, D. Luo, B. A. Malomed, and R. G. Hulet, *Nat. Phys.* **10**, 918 (2014).
- [8] A. L. Marchant, T. P. Billam, T. P. Wiles, M. M. H. Yu, S. A. Gardiner, and S. L. Cornish, *Nat. Commun.* **4**, 1865 (2013).
- [9] K. E. Strecker, G. B. Partridge, A. G. Truscott, and R. G. Hulet, *Nature (London)* **417**, 150 (2002).
- [10] J. H. V. Nguyen, D. Luo, and R. G. Hulet, *Science* **356**, 422 (2017).
- [11] T. Mežnaršič, T. Arh, J. Brenc, J. Pišljari, K. Gosar, Ž. Gosar, R. Žitko, E. Zupanič, and P. Jeglič, *Phys. Rev. A* **99**, 033625 (2019).
- [12] S. L. Cornish, S. T. Thompson, and C. E. Wieman, *Phys. Rev. Lett.* **96**, 170401 (2006).
- [13] L. D. Carr and Y. Castin, *Phys. Rev. A* **66**, 063602 (2002).
- [14] F. Dalfovo, S. Giorgini, L. P. Pitaevskii, and S. Stringari, *Rev. Mod. Phys.* **71**, 463 (1999).
- [15] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.123.123602> for additional information, which includes Refs. [16,17].
- [16] S. Cowell, H. Heiselberg, I. E. Mazets, J. Morales, V. R. Pandharipande, and C. J. Pethick, *Phys. Rev. Lett.* **88**, 210403 (2002).
- [17] P. J. Everitt, M. A. Sooriyabandara, M. Guasoni, P. B. Wigley, C. H. Wei, G. D. McDonald, K. S. Hardman, P. Manju, J. D. Close, C. C. N. Kuhn, S. S. Szigeti, Y. S. Kivshar, and N. P. Robins, *Phys. Rev. A* **96**, 041601(R) (2017).
- [18] L. Salasnich, A. Parola, and L. Reatto, *Phys. Rev. A* **65**, 043614 (2002).
- [19] N. G. Parker, S. L. Cornish, C. S. Adams, and A. M. Martin, *J. Phys. B* **40**, 3127 (2007).
- [20] D. Longenecker and E. J. Mueller, *Phys. Rev. A* **99**, 053618 (2019).
- [21] M. Berninger, A. Zenesini, B. Huang, W. Harm, H.-C. Nägerl, F. Ferlaino, R. Grimm, P. S. Julienne, and J. M. Hutson, *Phys. Rev. A* **87**, 032517 (2013).
- [22] A. D. Carli, C. D. Colquhoun, S. Kuhr, and E. Haller, *New J. Phys.* **21**, 053028 (2019).
- [23] C. Menotti and S. Stringari, *Phys. Rev. A* **66**, 043610 (2002).
- [24] E. Haller, M. Gustavsson, M. J. Mark, J. G. Danzl, R. Hart, G. Pupillo, and H.-C. Nägerl, *Science* **325**, 1224 (2009).
- [25] V. Serkin, A. Hasegawa, and T. Belyaeva, in *Odyssey of Light in Nonlinear Optical Fibers: Theory and Applications*, edited by K. Porsezian and R. Ganapathy (CRC Press, Boca Raton, FL, 2015), pp. 145–187.
- [26] V. Serkin and T. Belyaeva, *Optik (Stuttgart)* **176**, 38 (2019).
- [27] C. H. Tenorio, E. V. Vargas, V. N. Serkin, M. A. Granados, T. L. Belyaeva, R. P. Moreno, and L. M. Lara, *Quantum Electron.* **35**, 929 (2005).
- [28] J. Golde, J. Ruhl, M. Olshanii, V. Dunjko, S. Datta, and B. A. Malomed, *Phys. Rev. A* **97**, 053604 (2018).
- [29] O. Gamayun and M. Semenyakin, *J. Phys. A* **49**, 335201 (2016).
- [30] A. I. Streltsov, O. E. Alon, and L. S. Cederbaum, *Phys. Rev. Lett.* **100**, 130401 (2008).
- [31] J. G. Cosme, C. Weiss, and J. Brand, *Phys. Rev. A* **94**, 043603 (2016).
- [32] C. Weiss and L. D. Carr, [arXiv:1612.05545](https://arxiv.org/abs/1612.05545).
- [33] V. A. Yurovsky, B. A. Malomed, R. G. Hulet, and M. Olshanii, *Phys. Rev. Lett.* **119**, 220401 (2017).
- [34] K. L. Ng, B. Opanchuk, M. D. Reid, and P. D. Drummond, *Phys. Rev. Lett.* **122**, 203604 (2019).
- [35] A. Di Carli, C. D. Colquhoun, G. Henderson, S. Flannigan, G.-L. Oppo, A. J. Daley, S. Kuhr, and E. Haller, Data for ‘Excitation Modes of Bright Matter-Wave Solitons’, University of Strathclyde KnowledgeBase, <https://doi.org/10.15129/699a004b-fbf3-4f37-b2f3-9c1c97f44f0f>.

Supplemental Material

“Excitation modes of bright matter-wave solitons”

Andrea Di Carli, Craig D. Colquhoun, Grant Henderson, Stuart Flannigan,
 Gian-Luca Oppo, Andrew J. Daley, Stefan Kuhr, Elmar Haller
Department of Physics and SUPA, University of Strathclyde, Glasgow G4 0NG, Scotland, United Kingdom

I. EXPERIMENTAL METHODS

A. Controlling the atom number in the BEC

The solitons are confined to a quasi-1D geometry with almost free propagation along the horizontal direction and strong radial confinement of $\omega_r = 2\pi \times 95$ Hz provided by laser beam L_H . In quasi-1D geometry, bright matter-wave solitons collapse for large densities and interactions [1], which for our typical experimental scattering length of approximately $-5a_0$ corresponds to a critical atom number of 2500 [2]. As a result, we need to strongly reduce the atom number to avoid collapse, modulation instabilities [3] and three-body loss [4] for a deterministic and reproducible creation of the soliton. We remove atoms with a small additional magnetic field gradient, which pushes the atoms over the edge of the optical dipole trap. Our precise control of magnetic field strengths allows us to reduce the atom number down to 200 atoms, with a reproducibility of ± 100 for 600 atoms and ± 350 for 4500 atoms, measured as the standard deviation of the atom number in 50 consecutive runs. A removal period of 4 s and smooth ramps of the magnetic field strength are necessary to minimize excitations of the BEC. Following the removal procedure we measure residual fluctuations of the width of the BEC below 3.5%.

B. Quench parameters

Several different quench protocols are employed for the measurements. The quenches are labeled by the symbols Q1-Q7 in the main article:

- Q1 We quench only the trap frequency from $\omega_{z,i} = 2\pi \times 5.8(2)$ Hz to $\omega_{z,f} = 2\pi \times 4.3(2)$ Hz with a linear ramp of the laser power of beam L_V over 4 ms. Atom number $N \approx 1800$, constant interaction strength $a_i = +7a_0$, $\omega_r = 2\pi \times 95$ Hz.
- Q2 In addition to the quench Q1 of the trap frequency, we also quench the interaction strength from $a_i = +7a_0$ to $a_f = -5.4a_0$ in 4 ms. We minimize oscillations of the width of the soliton by reducing the initial size of the BEC with $\omega_{z,i} = 2\pi \times 11.2(2)$ Hz. Atom number $N \approx 1800$, $\omega_{z,f} = 2\pi \times 4.3(2)$ Hz, $\omega_r = 2\pi \times 95$ Hz.
- Q3 We mismatch the initial size of the BEC before

the quench with $\omega_{z,i} = 2\pi \times 12.8(4)$ Hz to generate small amplitude oscillations of the width of the soliton. Atom number $N \approx 1700$, $\omega_{z,f} = 2\pi \times 4.3(2)$ Hz, $\omega_r = 2\pi \times 95$ Hz.

- Q4 Quench to determine the atom-number dependence of ω_{sol} . We vary the atom number N from 500 to 1700 for the measurement. $\omega_{z,f} = 2\pi \times 4.3(2)$ Hz, $a_i = +7a_0$, $a_f = -5.4a_0$, ramp duration 4 ms, $\omega_r = 2\pi \times 95$ Hz.
- Q5 Quench to determine the dependence of ω_{sol} on the trap frequency $\omega_{z,f}$. We vary $\omega_{z,f}$ from approximately 1 Hz to 9 Hz. Smaller values of $\omega_{z,f}$ result in larger equilibrium sizes of the soliton, and we need to reduce the initial trap frequencies $\omega_{z,i}$ to keep the oscillation amplitudes comparable during the measurements. The typical difference between $\omega_{z,i}$ and $\omega_{z,f}$ is approximately 3 Hz. $N \approx 1500$, $a_f = -5.4a_0$, $\omega_r = 2\pi \times 95$ Hz.
- Q6 Strong quench starting from an elongated BEC to excite higher-order oscillations and shedding. The ratio between the calculated initial Thomas-Fermi radius of the BEC and the expected width l_z of the soliton is 24. $a_i = 56a_0$, $a_f = -5.3a_0$, $\omega_{z,i} = 2\pi \times 4.9(2)$ Hz, $\omega_{z,f} = 2\pi \times 0.0(6)$ Hz, initial atom number $N \approx 3000$ drops to 1100 after shedding of atoms, quench duration 13 ms, $\omega_r = 2\pi \times 86$ Hz.
- Q7 Double quench to create a stable soliton in step 1 and quench the scattering length by approximately a factor of 4 in step 2. Step 1: $\omega_{z,i} = 2\pi \times 4.9(2)$ Hz, $\omega_{z,f} = 2\pi \times 1.4(2)$ Hz, $a_i = 29a_0$, $a_f = -0.8a_0$, $\omega_r = 2\pi \times 86$ Hz, quench duration 15 ms, $N \approx 2200$. Settling delay of 25 ms between quenches. Step 2: reduce interaction strength in 2 ms to $a_f = -4.6a_0$, no change of other parameters.

C. Fit of density profiles

We employ absorption imaging to measure the 2D-density profile of the soliton, and we integrate over one radial axis to determine the 1D-density profile (red line in Fig. 5). The width l_z of the soliton in Eq. 1 of the main article, is determined by fitting the function $A(\text{sech}(z/B))^2$, with fit-parameters A and B , to the integrated 1D-density profiles (dotted blue line in Fig. 5).




FIG. 5. 1D-density profile of a soliton. Red line: integrated density profile of the absorption image for $t = 60$ ms in Fig. 1c (main article). Dotted blue line: fitted profile according to Eq. 1 in the main article.

II. THEORETICAL METHODS

A. The Model

The time evolution of the collective wave function of N atoms in an external potential with the 3D Gross-Pitaevski equation (GPE) for a time and space-dependent collective atomic wave-function, $\psi(\mathbf{r}, t)$, is given by,

$$i\hbar \frac{\partial}{\partial t} \psi(\mathbf{r}, t) = \left[-\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) + gN |\psi(\mathbf{r}, t)|^2 \right] \psi(\mathbf{r}, t), \quad (1)$$

where $g = 4\pi\hbar^2 a/m$, m is the atomic mass, and a is the two-body s-wave scattering length. This semi-classical field equation can be seen as a mean-field computation, and describes the dynamics of many weakly interacting particles at low temperatures when the condition $n|a|^3 \ll 1$ is satisfied [5], where n is the particle density. Our external potential $V(\mathbf{r})$ is given by a 3D (anisotropic) harmonic trap.

For tight radial trapping potentials, $\omega_r \gg \omega_z$, we can approximate the 3D wave function with a Gaussian solution in the radial directions and an arbitrary component, $f(z, t)$, in the longitudinal direction,

$$\psi(\mathbf{r}, t) = f(z, t) \frac{1}{\sqrt{\pi} a_r \sigma(z, t)} \exp \left[-\frac{(x^2 + y^2)}{2a_r^2 \sigma(z, t)^2} \right], \quad (2)$$

where a_r is the harmonic oscillator length in the radial direction and $\sigma(z, t)$ is a free parameter that dictates the width of the radial wavefunction. Substituting this ansatz into the 3D-GPE, and integrating over the radial directions, we arrive at the so called non-polynomial

Schödinger equation (NPSE) [6]

$$i\hbar \frac{\partial}{\partial t} f(z, t) = \left[-\frac{\hbar^2}{2m} \frac{\partial^2}{\partial z^2} + V(z) + \frac{gN}{2\pi a_r^2 \sigma(z, t)^2} |f(z, t)|^2 + \frac{\hbar\omega_r}{2} \left(\sigma(z, t)^2 + \frac{1}{\sigma(z, t)^2} \right) \right] f(z, t), \quad (3)$$

where $\omega_r = \hbar/ma_r^2$. The condition for $\sigma(z, t)$ that minimizes the action functional integrated along the trajectories in phase space is [6],

$$\sigma(z, t)^2 = \sqrt{1 + 2aN|f(z, t)|^2}. \quad (4)$$

For $\sigma(z, t) = 1$, we obtain the ground state of a harmonic oscillator in the radial directions, and we recover the usual 1D-GPE

$$i\hbar \frac{\partial}{\partial t} f(z, t) = \left[-\frac{\hbar^2}{2m} \frac{\partial^2}{\partial z^2} + V(z) + \frac{gN}{2\pi a_r^2} |f(z, t)|^2 + \hbar\omega_r \right] f(z, t). \quad (5)$$

We have numerically integrated Eqs 3 and 5 using the split-step Fourier transform method [7], where we exploit the fact that the kinetic and potential terms in the Hamiltonian are diagonal in momentum and real space, respectively.

III. SOLITON BREATHING FREQUENCY

In this section we explain how the numerical calculations of the soliton breathing frequencies shown in Fig. 2 of the main text were carried out. We begin with the order 1 soliton solution,

$$f(z, 0) = \frac{1}{\sqrt{2}l_z} \operatorname{sech} \left(\frac{z}{l_z} \right) \quad (6)$$

where $l_z = a_r^2/(N|a_i|)$, and we have used $a_i = 7a_0$. We then evolve this initial state either with the 1D-GPE or NPSE to a simulation time of 4000 ms and evaluate the frequency spectrum of the oscillation of the soliton's centre ($z = 0$). In Fig. 6 we present the frequency spectrum for the GPE and a longitudinal frequency of $\omega_z = 2\pi \times 5$ Hz and atom number $N = 1300$, which is characteristic of the behaviour for all other ω_z data points. We observe several prominent frequency modes in the signal, but we select the lowest frequency peak to compare to the experimental measurements, because the resolution in the experiment is restricted to low frequency components.

Fig. 6b also shows the results of the simulation using both the 1D-GPE and the NPSE (compare with Fig. 2 of the main text). We can see that for these atom numbers there are differences between the predictions of the 1D-GPE and NPSE. However these differences are small compared to the uncertainty in the experimental results.

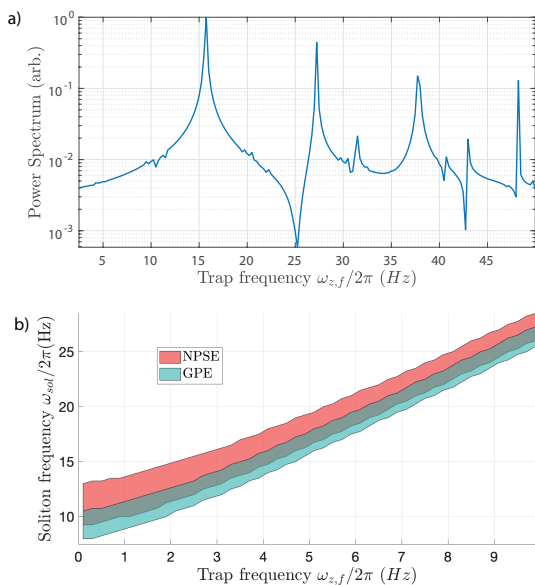


FIG. 6. Simulation results for the soliton breathing frequency, for comparison with Fig. 2 in the main text. (a) Frequency spectrum calculated using the 1D-GPE with a longitudinal frequency of $\omega_z = 2\pi \times 5$ Hz and atom number $N = 1300$. (b) Breathing frequency (first peak in the spectrum as in a) vs. trap frequency. $N \approx 1300 - 1500$ atoms, $a_f = -5.4 a_0$ for the NPSE (red) and the GPE (green). The simulations were evolved in time to 4000 ms.

IV. HIGHER ORDER SOLITONS

Figure 7 shows numerical simulations of the 1D-GPE for the time evolution of second- and third-order solitons with initial sizes $l_z^{(2)}$ and $l_z^{(3)}$. Large initial soliton sizes lead to the periodic formation of local maxima and minima of the density profile. Striking characteristics of the time evolution are the periodic development of a sharp central peak with side wings for the second-order soliton (Fig. 7a,b), and the periodic formation of a broad double-peak structure for the third-order soliton (Fig. 7c,d).

We also simulate the time evolution of solitons with the same start conditions using the NPSE and analyse the results using the root mean square width of the wave packet for a quantitative comparison (Fig. 8)

$$z_{rms}(t) = \left(\frac{1}{N} \int n(z,t)(z - \bar{z})^2 dz \right)^{1/2}. \quad (7)$$

Here, \bar{z} is the mean position of the wave packet and $n(z,t)$ is the 1D-density. We observe small quantitative differences between the two equations but the overall behaviour is very similar.




FIG. 7. Simulation of higher-order solitons in the 1D-GPE. Temporal snapshots (a) and temporal evolution (b) of the atomic density profile of an $n = 2$ soliton for $N = 1800$, $a = -3.7 a_0$, $l_z^{(2)} = 10.2 \mu\text{m} = 4l_z^{(1)}$, and an oscillation period of $T_2 = 271$ ms. Temporal snapshots (c) and temporal evolution (d) of the atomic density profile of an $n = 3$ soliton for the same values of N, a , but with $l_z^{(3)} = 22.8 \mu\text{m} = 9l_z^{(1)}$, and with a period $T_3 = 1373$ ms. The density profiles in (a) and (c) are plotted at $t = 0$ (dotted lines), $t = 1/4T$ (dashed lines), $t = 1/2T$ (solid lines). The dashed lines in (b) and (d) display the temporal evolution of the size of the soliton wavepacket z_{rms} (right scale).




FIG. 8. Simulation results for the root mean square width of the soliton as it undergoes second order solitary behaviour, for the NPSE (blue) and the 1D-GPE (red). Here, $\omega_z = 0$ Hz, with an atom number $N = 1800$ and a scattering length $a = -3.7 a_0$.

V. VARIATIONAL APPROACH FOR THE BREATHING FREQUENCY

In this section, we show how the longitudinal breathing frequency plotted in Fig. 2 of the main article can be determined from the variational ansatz for the soliton. For a cylindrical cigar-shaped potential the energy functional of Eq. 1 is given by [2, 8]

$$E[\psi] = \int d^3 \mathbf{r} \left[\frac{\hbar^2}{2m} |\nabla \psi(\mathbf{r})|^2 + \frac{1}{2} m (\omega_r^2 r^2 + \omega_z^2 z^2) |\psi(\mathbf{r})|^2 + \frac{gN}{2} |\psi(\mathbf{r})|^4 \right] \quad (8)$$

The energy of a soliton can be determined with a variational method using the following ansatz for the wave function

$$\psi(r, z) = \frac{1}{\sqrt{2}l_z} \operatorname{sech}\left(\frac{z}{l_z}\right) \cdot \frac{1}{\sqrt{\pi}l_r} \exp\left(-\frac{r^2}{2l_r^2}\right), \quad (9)$$

where the transverse width l_r and longitudinal width l_z are the variational parameters [2, 8]. Combining Eqs. 8 and 9, and rescaling the variables by the transverse frequency ω_r , provides an equation for the normalized energy of the soliton [2]

$$\varepsilon_{GP} = \frac{1}{2\gamma_r^2} + \frac{\gamma_r^2}{2} + \frac{1}{6\gamma_z^2} + \frac{\pi^2}{24}\lambda^2\gamma_z^2 + \frac{\alpha}{3\gamma_r^2\gamma_z}, \quad (10)$$

with $\varepsilon_{GP} = E/\hbar\omega_r$, $\gamma_r = l_r/\sigma_r$, $\gamma_z = l_z/\sigma_r$, $\lambda = \omega_z/\omega_r$, $\alpha = Na/\sigma_r$, and $\sigma_r = \sqrt{\hbar/m\omega_r}$. We can simplify Eq. 10 for our system with weak interactions and strong transverse confinement by neglecting variations of the radial soliton size, i.e. $\gamma_r = 1$. The energy minimum is found by calculating the zero-crossing of the first derivative of Eq. 10 with respect to γ_z

$$\frac{\pi^2}{4}\lambda^2\gamma_z^4 + \sqrt{\zeta}\gamma_z - 1 = 0, \quad (11)$$

where $\alpha = -|\alpha| = -\sqrt{\zeta}$. Eq. 11 has been solved for an expulsive potential with $\omega_z^2 < 0$ [2]. Here, we provide the solution for a trapping potential with $\omega_z^2 > 0$. The longitudinal size of the soliton γ_z^* at the energy minimum is

$$\gamma_z^* = \frac{F}{\sqrt{\zeta}}, \quad (12)$$

with

$$F = -\sqrt{\frac{G}{2}} + \frac{1}{2}\sqrt{-2G + \frac{4\sqrt{2}}{\pi^2\sqrt{G}}\left(\frac{\zeta}{\lambda}\right)^2}, \quad (13)$$

where

$$G = \frac{\Delta}{\pi^{\frac{4}{3}}}\left(\frac{\zeta}{\lambda}\right)^{\frac{4}{3}} - \frac{4}{3\pi^{\frac{2}{3}}}\frac{1}{\Delta}\left(\frac{\zeta}{\lambda}\right)^{\frac{2}{3}}, \quad (14)$$

with

$$\Delta = \sqrt[3]{1 + \sqrt{1 + \frac{64\pi^2}{27}\left(\frac{\lambda}{\zeta}\right)^2}}. \quad (15)$$

In order to find the oscillation frequency ω_z of the soliton, the equations of motion for the variational parameters are determined with a Lagrangian variational analysis [2]

$$\left(\frac{\pi^2}{12}\right)\ddot{\gamma}_z = \frac{\gamma_z^{-3}}{3} - \frac{\pi^2}{12}\lambda^2\gamma_z + \frac{\alpha}{3}\gamma_z^{-2}, \quad (16)$$

where the time derivative is calculated with respect to the normalised time $\tau = \omega_r t$. Again, we have assumed that the radial size of the soliton is constant, i.e. $\gamma_r = 1$.

For small deviations of the soliton size from its equilibrium value, we can write the solution as $\gamma_z = \gamma_z^* + \delta\gamma_z$, where γ_z^* is the minimum given by Eq. 12 and $\delta\gamma_z$ is a small deviation. A linear expansion of Eq. 16 leads to the expression for the longitudinal breathing frequency

$$\omega_z = \omega_r \sqrt{\frac{12}{\pi^2}\left(\gamma_z^{*-4} + \frac{\pi^2}{12}\lambda^2 + \frac{2\alpha}{3}\gamma_z^{*-3}\right)}. \quad (17)$$

We compare our experimental measurements of the breathing frequency of the soliton to the predictions of Eq. 17 in Fig. 2 of the main article (red line).

-
- [1] E. A. Donley, N. R. Claussen, S. L. Cornish, J. L. Roberts, E. A. Cornell, and C. E. Wieman, *Nature* **412**, 295 (2001).
 [2] L. D. Carr and Y. Castin, *Phys. Rev. A* **66**, 063602 (2002).
 [3] K. E. Strecker, G. B. Partridge, A. G. Truscott, and R. G. Hulet, *Nature* **417**, 150 (2002).
 [4] P. J. Everitt, M. A. Sooriyabandara, M. Guasoni, P. B. Wigley, C. H. Wei, G. D. McDonald, K. S. Hardman, P. Manju, J. D. Close, C. C. N. Kuhn, S. S. Szigeti, Y. S. Kivshar, and N. P. Robins, *Phys. Rev. A* **96**, 041601 (2017).
 [5] S. Cowell, H. Heiselberg, I. E. Mazets, J. Morales, V. R. Pandharipande, and C. J. Pethick, *Phys. Rev. Lett.* **88**, 210403 (2002).
 [6] L. Salasnich, A. Parola, and L. Reatto, *Phys. Rev. A* **65**, 043614 (2002).
 [7] P. Suarez, "An introduction to the split step fourier method using matlab." (2015).
 [8] N. G. Parker, S. L. Cornish, C. S. Adams, and A. M.

Martin, *J. Phys. B At. Mol. Opt. Phys.* **40**, 3127 (2007).

Appendix D

DDS Driver Code (Arduino Uno)

```
1      /*
2          ****
3          * ddsAD9959
4          * Control DDS AD9958 and AD9959
5          *
6          * Craig D. Colquhoun, Glasgow 2017
7
8          TO CHANGE DEFAULT FREQUENCY, CHANGE BOTH THE VALUE OF defaultFreq (
9              LINE 65 AT TIME OF WRITING) AND THE FOUR commandVect[i].
10             startFreqWrite VALUES IN THE clearCommandList FUNCTION.
11             TAKE FREQUENCY IN HZ, DIVIDE BY CLOCK FREQUENCY (1 GHZ AT TIME OF
12             WRITING) AND MULTIPLY BY 2^32, BEFORE CONVERTING TO HEXCODE!
13             ****
14             */
15
16 #include <SPI.h>
17 #include <Wire.h>
18 #include <LiquidCrystal_I2C.h>
19 #include <Ethernet.h>
20 #include <EthernetUdp.h>
21 #define UDP_TX_PACKET_MAX_SIZE 256 //increase UDP size
22
23 LiquidCrystal_I2C lcd(0x3F,20,21); // set the LCD address to 0x27
24     for a 16 chars and 2 line
25
26 /*****
27 // network stuff
28 *****/
29
30 // local MAC address, fake
31 byte mac[] = { 0x04, 0x7D, 0x4B, 0x28, 0x51, 0x33 };
32 // local ip
33 IPAddress localIPNum(172, 16, 1, 105);
34 byte gatewayIPNum[] = { 172, 16, 0, 1 };
35 byte dnsIPNum[] = { 8, 8, 8, 8 };
36 byte maskIPNum[] = { 255, 255, 0, 0 };
37
```

```

32 // An EthernetUDP instance to let us send and receive packets over
    UDP
33 EthernetUDP Udp;
34
35 //*****
36 // Set Pins for arduino
37 //*****
38 const int eth_CS      = 10; //cable select pin for ethernet shield
39 const int pin_CS      = 22; // cable select pin
40 const int pin_IOWrite = 23; // update pin
41 const int pin_Reset   = 24; // reset DDS pin
42 const int pin_RcvCmd  = 30; // TTL to receive commands over
    Ethernet connection
43 const int pin_NextCmd = 32; // TTL to select next command from
    list of received commands
44 const int pin_ExeCmd  = 31; // TTL to execute the selected command
45 const int pin_P0     = 40; // input pin P0
46 const int pin_P1     = 41; // input pin P1
47 const int pin_P2     = 42; // input pin P2
48 const int pin_P3     = 43; // input pin P3
49 const int pin_IO1    = 44; // input output pin IO1
50 const int pin_IO2    = 45; // input output pin IO2
51 const int pin_IO3    = 46; // input output pin IO3
52 int profilePin;
53 bool pinStatus = 0;
54
55 int idleDelay = 100;
56
57 //unsigned long startSweepTimer;
58 //unsigned long sweepTimer;
59 //unsigned long durt;
60 String commandString; //String used to store our command list
    later
61 bool bHasReset = 0; //Used to check whether the DDS was
    reset during the DDSinit function. This is for AOM protection, if
    the DDS has begun outputting and is reset it will stop
    outputting!
62 static float sineTable[] = {0.00, 0.02, 0.03, 0.05, 0.07, 0.09,
    0.10, 0.12, 0.14, 0.16, 0.17, 0.19, 0.21, 0.22, 0.24, 0.26, 0.28,
    0.29, 0.31, 0.33, 0.34, 0.36, 0.37, 0.39, 0.41, 0.42, 0.44,
    0.45, 0.47, 0.48, 0.50, 0.51, 0.53, 0.54, 0.56, 0.57, 0.59, 0.60,
    0.62, 0.63, 0.64, 0.66, 0.67, 0.68, 0.69, 0.71, 0.72, 0.73,
    0.74, 0.75, 0.77, 0.78, 0.79, 0.80, 0.81, 0.82, 0.83, 0.84, 0.85,
    0.86, 0.87, 0.87, 0.88, 0.89, 0.90, 0.91, 0.91, 0.92, 0.93,
    0.93, 0.94, 0.95, 0.95, 0.96, 0.96, 0.97, 0.97, 0.97, 0.98, 0.98,
    0.98, 0.99, 0.99, 0.99, 0.99, 1.00, 1.00, 1.00, 1.00, 1.00,
    1.00};
63 byte timeStep[90];
64 unsigned int delayTime;
65 double defaultFreq = 72.0;
66
67 typedef enum {
68 CH0      = 0x10, //Unused
69 CH1      = 0x20, //Unused
70 CH2      = 0x40, //Channel 0
71 CH3      = 0x80, //Channel 1

```

```

72     BOTH                = 0xC0    //Channels 0 and 1 simultaneously
73     } ad9959_channels;
74
75     typedef enum {
76     CSR                  = 0x00, //Channel select register
77     FR1                  = 0x01, //Function Register 1
78     FR2                  = 0x02, //Function Register 2
79     CFR                  = 0x03, //Channel Function Register
80     CTW0                 = 0x04, //Channel Tuning Word 0
81     CPW0                 = 0x05, //Channel Phase Tuning Word 0
82     ACR                  = 0x06, //Amplitude Control Register
83     LSR                  = 0x07, //Linear Sweep Ramp Rate
84     RDW                  = 0x08, //Rising Delta Word
85     FDW                  = 0x09, //Falling Delta Word
86     CTW1                 = 0x0A, //Channel Tuning Register 1
87     CTW2                 = 0x0B, //...and so on
88     CTW3                 = 0x0C,
89     CTW4                 = 0x0D,
90     CTW5                 = 0x0E,
91     CTW6                 = 0x0F,
92     CTW7                 = 0x10,
93     CTW8                 = 0x11,
94     CTW9                 = 0x12,
95     CTW10                = 0x13,
96     CTW11                = 0x14,
97     CTW12                = 0x15,
98     CTW13                = 0x16,
99     CTW14                = 0x17,
100    CTW15                = 0x18,
101    READ                  = 0x80 // not really a register
102    } ad9959_registers;
103
104    byte CSRbyte          = 0x00;
105    byte FR1byte[3]      = {0x00, 0x00, 0x00};
106    byte FR2byte[2]      = {0x00, 0x00};
107    byte CFRbyte[3]      = {0x00, 0x00, 0x00};
108    byte CTW0byte[4]     = {0x00, 0x00, 0x00, 0x00};
109    byte CPW0byte[2]     = {0x00, 0x00};
110    byte ACRbyte[3]      = {0x00, 0x00, 0x00};
111    byte LSRbyte[2]      = {0x00, 0x00};
112    byte RDWbyte[4]      = {0x00, 0x00, 0x00, 0x00};
113    byte FDWbyte[4]      = {0x00, 0x00, 0x00, 0x00};
114    byte CTW1byte[4]     = {0x00, 0x00, 0x00, 0x00};
115    //byte default72MHz[4] = {0x24, 0xdd, 0x2f, 0x1b}; SETTING DEFAULT
116    //byte default80MHz[4] = {0x28, 0xf5, 0xc2, 0x8f}; SETTING DEFAULT
117    //byte default80MHz[4] = {0x28, 0xf5, 0xc2, 0x8f}; SETTING DEFAULT
118    byte chan = 0;
119
120    struct tCommand {
121    int rcvCmd;
122    ad9959_channels chan;
123    double startFreq;
124    double stopFreq;
125    double duration;
126    byte startFreqWrite[4];

```

```

126     byte stopFreqWrite[4];
127     byte riseWord[4];
128     byte fallWord[4];
129     byte rampRate[2];
130 };
131
132     #define maxNCommands 5
133     tCommand commandVect[maxNCommands];
134     int usedNCommands = 0;           // number of commands used
135     int actNCommand = 0;           // index of the active command
136
137     //*****
138     // setup
139     //*****
140     void setup() {
141         // configure pins
142         pinMode(pin_IOUpdate, OUTPUT);
143         pinMode(pin_Reset, OUTPUT);
144         pinMode(pin_CS, OUTPUT);
145         pinMode(pin_P0, OUTPUT);
146         pinMode(pin_P1, OUTPUT);
147         pinMode(pin_P2, OUTPUT);
148         pinMode(pin_P3, OUTPUT);
149         pinMode(pin_IO1, OUTPUT);
150         pinMode(pin_IO2, OUTPUT);
151         pinMode(pin_IO3, OUTPUT);
152         pinMode(pin_RcvCmd, INPUT);
153         pinMode(pin_ExeCmd, INPUT);
154         pinMode(pin_NextCmd, INPUT);
155         pinMode(eth_CS, OUTPUT);
156         digitalWrite(eth_CS, LOW);
157         Ethernet.begin(mac, localIPNum, dnsIPNum, gatewayIPNum, maskIPNum );
158         Udp.begin(8020);
159         digitalWrite(eth_CS, HIGH);
160
161         lcd.init();
162         lcd.backlight();
163         lcd.setCursor(0,0);
164         lcd.print("0");
165         lcd.setCursor(0,1);
166         lcd.print("1");
167         setDefaultSetting();
168         clearCommandList();
169     }
170
171     void loop() {
172         int retValue;
173         bool bExeFlag = 0;
174         bool doIDisplay = 1;
175
176         if (!digitalRead(pin_RcvCmd))
177         {
178             delay(idleDelay);
179             return;
180         }

```

```

181     digitalWrite(eth_CS, LOW);                                     //Select ethernet
        shield so that it can be used without interfering with DDS
182     retValue = receiveCommand();                                  //Receive
        commands through ethernet port
183     digitalWrite(eth_CS, HIGH);                                   //Deselect ethernet
        shield in order to send commands to DDS again
184     if (!retValue) return;
185     if (!parseCommandString()) return;
186     if (commandVect[actNCommand].chan == CH2) lcd.setCursor(1,0);
187     else if (commandVect[actNCommand].chan == CH3) lcd.setCursor(1,1);
188     lcd.print("COMMANDS PARSED");
189
190     while(digitalRead(pin_RcvCmd)) delay(1);
191     while(actNCommand<=usedNCommands)
192     {
193     if(doIDisplay == 1)
194     {
195     if (commandVect[actNCommand].chan == CH2) lcd.setCursor(1,0);
196     else if (commandVect[actNCommand].chan == CH3) lcd.setCursor(1,1);
197     lcd.print("COMMAND ");
198     lcd.print(actNCommand);
199     lcd.print(" READY");
200     doIDisplay = 0;
201     }
202
203     if (digitalRead(pin_RcvCmd)) return;
204     else if (!bExeFlag && digitalRead(pin_ExeCmd))
205     {
206     executeCommand();
207     actNCommand++;
208     bExeFlag = 1;
209     }
210     else if (bExeFlag && !digitalRead(pin_ExeCmd))
211     {
212     bExeFlag = 0;
213     doIDisplay = 1;
214     if (actNCommand == usedNCommands) break;
215     }
216     }
217     setDefaultSetting();
218     return;
219     }
220
221     //*****
222     // clear command structure
223     //*****
224     void clearCommandList() {
225     for (int i = 0; i < 5; i++){
226     commandVect[i].rcvCmd = 1;
227     commandVect[i].chan = BOTH;
228     commandVect[i].startFreq = defaultFreq;
229     commandVect[i].stopFreq = defaultFreq;
230     commandVect[i].duration = 0;
231     // 80 MHz hex value used as default below
232     // commandVect[i].startFreqWrite[0] = 0x28;
233     // commandVect[i].startFreqWrite[1] = 0xf5;

```

```

234 //     commandVect[i].startFreqWrite[2] = 0xc2;
235 //     commandVect[i].startFreqWrite[3] = 0x8f;
236 // 72 MHz hex value used as default below
237 commandVect[i].startFreqWrite[0] = 0x24;
238 commandVect[i].startFreqWrite[1] = 0xdd;
239 commandVect[i].startFreqWrite[2] = 0x2f;
240 commandVect[i].startFreqWrite[3] = 0x1b;
241 //
242 for(int j = 0; j<4; j++){
243     commandVect[i].stopFreqWrite[j] = 0x00;
244     commandVect[i].riseWord[j] = 0x00;
245     commandVect[i].fallWord[j] = 0x00;
246 }
247 for (int j=0;j<2;j++){
248     commandVect[i].rampRate[j] = 0x00;
249 }
250 }
251 usedNCommands = 0;
252 actNCommand   = 0;
253 }
254 //*****
255 // Frequency
256 //*****
257 void calcParam(int cIndex, ad9959_registers REG, double param){
258     if ((REG == 4) || (REG > 7))
259     {
260         unsigned long paramWord = (unsigned long)(pow(2,32)*(param/500));
261         //Convert from the value given to the word value understood by
262         //the DDS
263         if (REG == 4)
264         {
265             commandVect[cIndex].startFreqWrite[3] = (byte) paramWord;
266             commandVect[cIndex].startFreqWrite[2] = (byte) (paramWord >> 8);
267             commandVect[cIndex].startFreqWrite[1] = (byte) (paramWord >> 16);
268             commandVect[cIndex].startFreqWrite[0] = (byte) (paramWord >> 24);
269         }
270         else if (REG == 8)
271         {
272             commandVect[cIndex].riseWord[3] = (byte) paramWord;
273             commandVect[cIndex].riseWord[2] = (byte) (paramWord >> 8);
274             commandVect[cIndex].riseWord[1] = (byte) (paramWord >> 16);
275             commandVect[cIndex].riseWord[0] = (byte) (paramWord >> 24);
276         }
277         else if (REG == 9)
278         {
279             commandVect[cIndex].fallWord[3] = (byte) paramWord;
280             commandVect[cIndex].fallWord[2] = (byte) (paramWord >> 8);
281             commandVect[cIndex].fallWord[1] = (byte) (paramWord >> 16);
282             commandVect[cIndex].fallWord[0] = (byte) (paramWord >> 24);
283         }
284         else if (REG == 10)
285         {
286             commandVect[cIndex].stopFreqWrite[3] = (byte) paramWord;
287             commandVect[cIndex].stopFreqWrite[2] = (byte) (paramWord >> 8);
288             commandVect[cIndex].stopFreqWrite[1] = (byte) (paramWord >> 16);
289             commandVect[cIndex].stopFreqWrite[0] = (byte) (paramWord >> 24);

```

```

288     }
289     else if (REG > 10)
290     {
291     //NOT PROGRAMMED YET
292     return;
293     }
294     }
295     else if (REG == 7)
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340

```

//If the
register is Linear Sweep Ramp Rate

```

    {
    unsigned long paramWord = (unsigned long) param;
    commandVect[cIndex].rampRate[1] = (byte) paramWord;
    commandVect[cIndex].rampRate[0] = (byte) paramWord;
    }
}

void programDDS()
{
if (commandVect[actNCommand].rcvCmd == 1)
{
writeReg(CSR, commandVect[actNCommand].chan);
writeReg(CTW0, commandVect[actNCommand].startFreqWrite, 4);
}
else if (commandVect[actNCommand].rcvCmd == 2) //if linear sweep
{
if (commandVect[actNCommand].chan == CH3) profilePin = pin_P3;
else if (commandVect[actNCommand].chan == CH2) profilePin = pin_P2;
writeReg(CSR, commandVect[actNCommand].chan);
if (commandVect[actNCommand].startFreq > commandVect[actNCommand].
stopFreq) //if sweep down
{
pinStatus = 1;
writeReg(RDW, commandVect[actNCommand].riseWord, 4);
writeReg(FDW, commandVect[actNCommand].fallWord, 4);
writeReg(LSR, commandVect[actNCommand].rampRate, 2);
writeReg(CTW1, commandVect[actNCommand].startFreqWrite, 4);
writeReg(CTW0, commandVect[actNCommand].startFreqWrite-1, 4);
digitalWrite(profilePin, pinStatus);
writeReg(CTW0, commandVect[actNCommand].stopFreqWrite, 4);
}
else //assumed sweeping up
{
pinStatus = 0;
writeReg(CTW0, commandVect[actNCommand].startFreqWrite, 4);
writeReg(CTW1, commandVect[actNCommand].stopFreqWrite, 4);
writeReg(RDW, commandVect[actNCommand].riseWord, 4);
writeReg(LSR, commandVect[actNCommand].rampRate, 2);
}
}
else //procedures for new commands will be written here...
{
return;
}
}
}

```

```

341 void calcFreqSweep(int cIndex)
342 {
343     double hzStartFreq;
344     double hzEndFreq;
345     if (commandVect[cIndex].startFreq > commandVect[cIndex].stopFreq)
346     {
347         hzStartFreq = commandVect[cIndex].stopFreq * 1000000;
348         hzEndFreq = commandVect[cIndex].startFreq * 1000000;
349     }
350     else
351     {
352         hzStartFreq = commandVect[cIndex].startFreq * 1000000;
353         hzEndFreq = commandVect[cIndex].stopFreq * 1000000;
354     }
355     unsigned long n = 65536;
356     bool foundFlag = 0;
357     float dTime, dFreq;
358     int k = 10000;
359     int l = 0;
360     int m = 0;
361     int minNValue = 1;
362     unsigned long maxNValue = min(commandVect[cIndex].duration
363         /0.000000008, (hzEndFreq-hzStartFreq)/0.12);
364     do
365     {
366         dTime = (125000000.0*commandVect[cIndex].duration)/n;
367         dFreq = 1.0*(hzEndFreq-hzStartFreq)/n;
368         if ((dFreq < 0.12) || (dTime < 1))
369         {
370             l = 1;
371             n = n - k;
372         }
373         else if ((dFreq > (hzEndFreq-hzStartFreq)) || (dTime > 255))
374         {
375             l = -1;
376             n = n + k;
377         }
378         else if((dFreq <= (hzEndFreq-hzStartFreq)) && (dTime <= 255) && (
379             dFreq >= 0.12) && (dTime >= 1))
380         {
381             if (k > 1)
382             {
383                 m = l*k;
384                 n = n + m;
385                 k = k/10;
386             }
387             else foundFlag = 1;
388         }
389         else
390         {
391             //error
392             return;
393         }
394         if (((n<minNValue) || (n>maxNValue)) && (k != 1))
395         {
396             m = l*k;

```

```

395     n = n + m;
396     k = k/10;
397     }
398     else if (((n<minNValue) || (n>maxNValue)) && (k == 1)) break;
399 }while (foundFlag == 0);
400
401
402     if(foundFlag==0)
403     {
404     if (commandVect[cIndex].chan == CH2) lcd.setCursor(1,0);
405     else if (commandVect[cIndex].chan == CH3) lcd.setCursor(1,1);
406     lcd.print("PROBLEMi");
407     lcd.print(cIndex);
408     lcd.print("NOnVAL");
409
410     return;
411     }
412
413     dFreq = dFreq / 1000000.0;
414
415     calcParam(cIndex, CTW0, commandVect[cIndex].startFreq);
416     calcParam(cIndex, CTW1, commandVect[cIndex].stopFreq);
417     calcParam(cIndex, RDW, dFreq);
418     calcParam(cIndex, FDW, dFreq);
419     calcParam(cIndex, LSR, dTime);
420     }
421
422     void setDefaultSetting(){
423     clearCommandList();
424     CFRbyte[0] = 0x00; CFRbyte[1] = 0x03; CFRbyte[2] = 0x00;
425     initDDS(BOTH);
426     programDDS();
427     lcd.setCursor(1,0);
428     lcd.print("DEFAULT      ");
429     lcd.setCursor(1,1);
430     lcd.print("DEFAULT      ");
431     }
432
433     void initDDS(byte initChan) {
434     digitalWrite(pin_P0, LOW);
435     digitalWrite(pin_P1, LOW);
436     digitalWrite(pin_P2, LOW);
437     digitalWrite(pin_P3, LOW);
438     digitalWrite(pin_IO1, LOW);
439     digitalWrite(pin_IO2, LOW);
440     digitalWrite(pin_IO3, LOW);
441
442     digitalWrite(pin_CS, HIGH); // ensure SS stays high
443     SPI.begin();
444     SPI.setClockDivider(2); //Can be integer from 1 -
445     // 255 using Due, divides the 84MHz clock rate by this number
446     SPI.setBitOrder(MSBFIRST);
447     SPI.setDataMode(SPI_MODE0);
448
449     if (!bHasReset){
450     resetDDS();

```

```

450     bHasReset = 1;
451     }
452     writeReg(CSR, initChan);
453     FR1byte[0] = 0xD0; FR1byte[1] = 0x54;
454     writeReg(FR1, FR1byte, sizeof(FR1byte));
455     writeReg(FR2, FR2byte, sizeof(FR2byte));
456     writeReg(CFR, CFRbyte, sizeof(CFRbyte));
457     }
458
459     // flip reset pin
460     void resetDDS() {
461         digitalWrite(pin_Reset, HIGH);
462         digitalWrite(pin_Reset, LOW);
463     }
464
465     // trigger update pin
466     void updateDDS() {
467         digitalWrite(pin_IOUpdate, HIGH);
468         digitalWrite(pin_IOUpdate, LOW);
469     }
470
471     void writeReg(ad9959_registers REG, byte value) {
472         digitalWrite(pin_CS, LOW);
473         SPI.transfer(REG);
474         SPI.transfer(value);
475         updateDDS();
476         digitalWrite(pin_CS, HIGH);
477     }
478
479     void writeReg(ad9959_registers REG, byte *buffer, int len){
480         digitalWrite(pin_CS, LOW);
481         SPI.transfer(REG);
482         SPI.transfer(buffer, len);
483         updateDDS();
484         digitalWrite(pin_CS, HIGH);
485     }
486
487     int receiveCommand()
488     {
489         // buffer to receive string with UDP
490         char UDPBuffer[UDP_TX_PACKET_MAX_SIZE];
491         int ret = 0;
492         int packLength = Udp.parsePacket();
493         int expectedResponse = 12;
494         commandString = ""; // clear the commandString
495
496         while (packLength>0) {
497             if (packLength >= expectedResponse) {
498                 Udp.read(UDPBuffer, UDP_TX_PACKET_MAX_SIZE); // read packet into
499                     the buffer
500                 for (int i = 0; i < packLength; i++){
501                     commandString = commandString + UDPBuffer[i];
502                 }
503                 ret = 1;
504             }
505             packLength = Udp.parsePacket();

```

```

505     }
506     return ret;
507 }
508
509
510 int parseCommandString() {
511     int ret = 0;
512     String stringParse = ""; // dummy string
513     String stringParseAT = ""; // dummy string
514
515     int chSelect; // channel
516     int rcvCmd; // command mode
517     double startFreq, stopFreq, duration;
518
519     int cIndexAT;
520     int cIndexHASH, cIndexStart;
521     int cIndexCounter = 0;
522
523     clearCommandList();
524
525     commandString.trim();
526     if (commandString.length()==0) return ret;
527     cIndexAT = commandString.indexOf('@');
528
529     // parse the commandStrings
530     while (cIndexAT!=-1 && cIndexCounter<maxNCommands) {
531
532         // get the substring up to next @
533         stringParseAT = commandString.substring(0, cIndexAT);
534
535         // get command number
536         cIndexHASH = stringParseAT.indexOf('#', 0);
537         if (cIndexHASH==-1) return ret;
538         stringParse = stringParseAT.substring(0, cIndexHASH);
539         rcvCmd = stringParse.toInt(); //setting
540         if (rcvCmd<1 || rcvCmd>3) return ret;
541         commandVect[cIndexCounter].rcvCmd = rcvCmd;
542         cIndexStart = cIndexHASH+1;
543
544         // get device number
545         cIndexHASH = stringParseAT.indexOf('#', cIndexStart);
546         if (cIndexHASH==-1) return ret;
547         stringParse = stringParseAT.substring(cIndexStart, cIndexHASH);
548         chSelect = stringParse.toInt(); //setting
549         switch (chSelect) {
550             case 1:
551                 commandVect[cIndexCounter].chan = CH2; break;
552             case 2:
553                 commandVect[cIndexCounter].chan = CH3; break;
554             default:
555                 return ret;
556         }
557         cIndexStart = cIndexHASH+1;
558
559         // get start freq
560         cIndexHASH = stringParseAT.indexOf('#', cIndexStart);

```

```

561     if (cIndexHASH==-1) return ret;
562     stringParse = stringParseAT.substring(cIndexStart ,cIndexHASH);
563     startFreq = stringParse.toDouble();
564     if (startFreq<60 || startFreq>110) return ret;
565     commandVect[cIndexCounter].startFreq = startFreq;
566     cIndexStart = cIndexHASH+1;
567
568     // get stop freq
569     cIndexHASH = stringParseAT.indexOf('#',cIndexStart);
570     if (cIndexHASH==-1) commandVect[cIndexCounter].stopFreq =
        defaultFreq;
571     else {
572     stringParse = stringParseAT.substring(cIndexStart ,cIndexHASH);
573     stopFreq = stringParse.toDouble();
574     if (stopFreq<60 || stopFreq>100) return ret;
575     }
576     commandVect[cIndexCounter].stopFreq = stopFreq;
577     cIndexStart = cIndexHASH+1;
578
579     stringParse = stringParseAT.substring(cIndexStart);
580     if (stringParse.length()<1) commandVect[cIndexCounter].duration =
        0.0;
581     else {
582     duration = stringParse.toDouble();
583     if (duration<0 || duration>10) return ret;
584     if (duration==0) return ret;
585     }
586     commandVect[cIndexCounter].duration = duration;
587     calcFreqSweep(cIndexCounter);
588
589     commandString.remove(0 ,cIndexAT+1);
590     cIndexAT = commandString.indexOf('@');
591     cIndexCounter++;
592     usedNCommands++;
593     }
594
595     ret = 1;
596     return ret;
597 }
598
599 void executeCommand() {
600
601     if (actNCommand>=usedNCommands) return;
602
603     String str;
604     str = String(commandVect[actNCommand].rcvCmd) + " " + String(
        commandVect[actNCommand].chan) + " " + String(commandVect[
        actNCommand].startFreq,6) + " " + String(commandVect[actNCommand
        ].stopFreq,6) + " " + String(commandVect[actNCommand].duration,4)
        ;
605
606     // nothing to be done?
607     if (commandVect[actNCommand].rcvCmd == 0){
608     setDefaultSetting();
609     return;
610     }

```

```

611
612 // single frequency
613 if (commandVect[actNCommand].rcvCmd == 1){
614 // display
615 if(commandVect[actNCommand].chan == CH2) lcd.setCursor(1,0);
616 else if(commandVect[actNCommand].chan == CH3) lcd.setCursor(1,1);
617 lcd.print("C");
618 lcd.print(commandVect[actNCommand].startFreq);
619 lcd.print(" ");
620 CFRbyte[0] = 0x00; CFRbyte[1] = 0x03; CFRbyte[2] = 0x00;
621 initDDS(commandVect[actNCommand].chan);
622 programDDS();
623 return;
624 }
625
626 // linear sweep up or down
627 if (commandVect[actNCommand].rcvCmd == 2){
628 if(commandVect[actNCommand].chan == CH2) lcd.setCursor(1,0);
629 else if(commandVect[actNCommand].chan == CH3) lcd.setCursor(1,1);
630 lcd.print("L");
631 lcd.print(commandVect[actNCommand].startFreq,1);
632 lcd.print("-");
633 lcd.print(commandVect[actNCommand].stopFreq,1);
634 lcd.print(",");
635 lcd.print(commandVect[actNCommand].duration);
636 lcd.print(" ");
637 CFRbyte[0] = 0x80; CFRbyte[1] = 0x43; CFRbyte[2] = 0x00;
        //Set DDS to output frequency sweep, and enable linear sweep mode
        (for linear or sinusoidal sweeps)
638 initDDS(commandVect[actNCommand].chan);
639 programDDS();
640
641 digitalWrite(profilePin, !pinStatus); //Set the profile pin to
        whatever it isn't right now
642 return;
643 }
644
645 //linear sweep up and down
646 if (commandVect[actNCommand].rcvCmd == 3){
647 if(commandVect[actNCommand].chan == CH2) lcd.setCursor(1,0);
648 else if(commandVect[actNCommand].chan == CH3) lcd.setCursor(1,1);
649 lcd.print("Z");
650 lcd.print(commandVect[actNCommand].stopFreq,1);
651 lcd.print("-");
652 lcd.print(commandVect[actNCommand].startFreq,1);
653 lcd.print(",");
654 lcd.print(commandVect[actNCommand].duration);
655 lcd.print(" ");
656 CFRbyte[0] = 0x80; CFRbyte[1] = 0x43; CFRbyte[2] = 0x00;
        //Set DDS to output frequency sweep, and enable linear sweep mode
        (for linear or sinusoidal sweeps)
657 initDDS(commandVect[actNCommand].chan);
658 programDDS();
659
660 digitalWrite(profilePin, !pinStatus); //Switch the profile
        pin to whatever it isn't set to right now

```

```
661     return;
662 }
663
664 // sine frequency
665 if (commandVect[actNCommand].rcvCmd == 4){
666     if(commandVect[actNCommand].chan == CH2) lcd.setCursor(1,0);
667     else if(commandVect[actNCommand].chan == CH3) lcd.setCursor(1,1);
668     lcd.print("S");
669     lcd.print(commandVect[actNCommand].startFreq,1);
670     lcd.print("-");
671     lcd.print(commandVect[actNCommand].stopFreq,1);
672     lcd.print(",");
673     lcd.print(commandVect[actNCommand].duration);
674     lcd.print("      ");
675     CFRbyte[0] = 0x80; CFRbyte[1] = 0x43; CFRbyte[2] = 0x00;
        //Set DDS to output frequency sweep, and enable linear sweep mode
        (for linear or sinusoidal sweeps)
676     initDDS(commandVect[actNCommand].chan);
677 }
678 }
```

Bibliography

- [1] S. Chu, L. W. Hollberg, J. E. Bjorkholm, A. Cable, and A. Ashkin, “Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure,” *Phys. Rev. Lett.*, vol. 55, pp. 48–51, 1985.
- [2] 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, no. 11, pp. 2023–2045, 1989.
- [3] 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.
- [4] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle, “Bose-Einstein condensation in a gas of sodium atoms,” *Phys. Rev. Lett.*, vol. 75, pp. 3969–3973, 1995.
- [5] C. C. Bradley, C. A. Sackett, J. J. Tollett, and R. G. Hulet, “Evidence of Bose-Einstein condensation in an atomic gas with attractive interactions,” *Phys. Rev. Lett.*, vol. 75, pp. 1687–1690, 1995.
- [6] W. D. Phillips, “Nobel lecture: Laser cooling and trapping of neutral atoms,” *Rev. Mod. Phys.*, vol. 70, pp. 721–741, 1998.
- [7] S. Chu, “Nobel lecture: The manipulation of neutral particles,” *Rev. Mod. Phys.*, vol. 70, pp. 685–706, 1998.
- [8] C. N. Cohen-Tannoudji, “Nobel lecture: Manipulating atoms with photons,” *Rev. Mod. Phys.*, vol. 70, pp. 707–719, 1998.
- [9] W. Ketterle, “Nobel lecture: When atoms behave as waves: Bose-Einstein condensation and the atom laser,” *Rev. Mod. Phys.*, vol. 74, pp. 1131–1151, 2002.
- [10] E. A. Cornell and C. E. Wieman, “Bose-Einstein condensation in a dilute gas: The first 70 years and some recent experiments (Nobel lecture),” *ChemPhysChem*, vol. 3, no. 6, pp. 476–493, 2002.

- [11] M.-O. Mewes, M. R. Andrews, D. M. Kurn, D. S. Durfee, C. G. Townsend, and W. Ketterle, “Output coupler for Bose-Einstein condensed atoms,” *Phys. Rev. Lett.*, vol. 78, pp. 582–585, 1997.
- [12] L. Khaykovich, F. Schreck, G. Ferrari, T. Bourdel, J. Cubizolles, L. D. Carr, Y. Castin, and C. Salomon, “Formation of a matter-wave bright soliton,” *Science*, vol. 296, no. 5571, pp. 1290–1293, 2002.
- [13] K. E. Strecker, G. B. Partridge, A. G. Truscott, and R. G. Hulet, “Formation and propagation of matter-wave soliton trains,” *Nature*, vol. 417, no. 6885, pp. 150–153, 2002.
- [14] M. R. Matthews, B. P. Anderson, P. C. Haljan, D. S. Hall, C. E. Wieman, and E. A. Cornell, “Vortices in a Bose-Einstein condensate,” *Phys. Rev. Lett.*, vol. 83, pp. 2498–2501, 1999.
- [15] K. W. Madison, F. Chevy, W. Wohlleben, and J. Dalibard, “Vortex formation in a stirred Bose-Einstein condensate,” *Phys. Rev. Lett.*, vol. 84, pp. 806–809, 2000.
- [16] M. R. Andrews, C. G. Townsend, H.-J. Miesner, D. S. Durfee, D. M. Kurn, and W. Ketterle, “Observation of interference between two Bose condensates,” *Science*, vol. 275, no. 5300, pp. 637–641, 1997.
- [17] T. Weber, *Bose-Einstein Condensation of Optically Trapped Cesium*. PhD thesis, University of Innsbruck, 2003.
- [18] J. Herbig, *Quantum-Degenerate Cesium: Atoms and Molecules*. PhD thesis, University of Innsbruck, 2005.
- [19] L. Essen and J. V. L. Parry, “An atomic standard of frequency and time interval: A caesium resonator,” *Nature*, vol. 176, no. 4476, pp. 280–282, 1955.
- [20] M. A. Kasevich, E. Riis, S. Chu, and R. G. DeVoe, “rf spectroscopy in an atomic fountain,” *Phys. Rev. Lett.*, vol. 63, pp. 612–615, 1989.
- [21] A. Clairon, C. Salomon, S. Guellati, and W. D. Phillips, “Ramsey resonance in a Zacharias fountain,” *Europhysics Letters (EPL)*, vol. 16, no. 2, pp. 165–170, 1991.
- [22] E. Tiesinga, B. J. Verhaar, H. T. C. Stoof, and D. van Bragt, “Spin-exchange frequency shift in a cesium atomic fountain,” *Phys. Rev. A*, vol. 45, pp. R2671–R2673, 1992.

- [23] E. Tiesinga, A. J. Moerdijk, B. J. Verhaar, and H. T. C. Stoof, “Conditions for Bose-Einstein condensation in magnetically trapped atomic cesium,” *Phys. Rev. A*, vol. 46, pp. R1167–R1170, 1992.
- [24] A. Kastberg, W. D. Phillips, S. L. Rolston, R. J. C. Spreeuw, and P. S. Jessen, “Adiabatic cooling of cesium to 700 nk in an optical lattice,” *Phys. Rev. Lett.*, vol. 74, pp. 1542–1545, 1995.
- [25] J. Arlt, P. Bance, S. Hopkins, J. Martin, S. Webster, A. Wilson, K. Zetie, and C. J. Foot, “Suppression of collisional loss from a magnetic trap,” *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 31, no. 7, pp. L321–L327, 1998.
- [26] J. Söding, D. Guéry-Odelin, P. Desbiolles, G. Ferrari, and J. Dalibard, “Giant spin relaxation of an ultracold cesium gas,” *Phys. Rev. Lett.*, vol. 80, pp. 1869–1872, 1998.
- [27] P. J. Leo, E. Tiesinga, P. S. Julienne, D. K. Walter, S. Kadlecsek, and T. G. Walker, “Elastic and inelastic collisions of cold spin-polarized ^{133}Cs atoms,” *Phys. Rev. Lett.*, vol. 81, pp. 1389–1392, 1998.
- [28] D. Guéry-Odelin, J. Söding, P. Desbiolles, and J. Dalibard, “Is Bose-Einstein condensation of atomic cesium possible?,” *Europhysics Letters (EPL)*, vol. 44, no. 1, pp. 25–30, 1998.
- [29] S. J. J. M. F. Kokkelmans, B. J. Verhaar, and K. Gibble, “Prospects for Bose-Einstein condensation in cesium,” *Phys. Rev. Lett.*, vol. 81, pp. 951–954, 1998.
- [30] D. Guéry-Odelin, J. Söding, P. Desbiolles, and J. Dalibard, “Strong evaporative cooling of a trapped cesium gas,” *Opt. Express*, vol. 2, no. 8, pp. 323–329, 1998.
- [31] S. A. Hopkins, S. Webster, J. Arlt, P. Bance, S. Cornish, O. Maragò, and C. J. Foot, “Measurement of elastic cross section for cold cesium collisions,” *Phys. Rev. A*, vol. 61, p. 032707, 2000.
- [32] D. Guéry-Odelin, J. Söding, P. Desbiolles, and J. Dalibard, “Is Bose-Einstein condensation of atomic cesium possible?,” *Europhysics Letters (EPL)*, vol. 44, no. 1, pp. 25–30, 1998.
- [33] E. Tiesinga, B. J. Verhaar, and H. T. C. Stoof, “Threshold and resonance phenomena in ultracold ground-state collisions,” *Phys. Rev. A*, vol. 47, pp. 4114–4122, 1993.

- [34] 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, 1993.
- [35] M. Arndt, M. Ben Dahan, D. Guéry-Odelin, M. W. Reynolds, and J. Dalibard, “Observation of a zero-energy resonance in Cs-Cs collisions,” *Phys. Rev. Lett.*, vol. 79, pp. 625–628, 1997.
- [36] V. Vuletić, A. J. Kerman, C. Chin, and S. Chu, “Observation of low-field Feshbach resonances in collisions of cesium atoms,” *Phys. Rev. Lett.*, vol. 82, pp. 1406–1409, 1999.
- [37] C. Chin, V. Vuletić, A. J. Kerman, and S. Chu, “High resolution Feshbach spectroscopy of cesium,” *Phys. Rev. Lett.*, vol. 85, pp. 2717–2720, 2000.
- [38] A. J. Kerman, C. Chin, V. Vuletić, S. Chu, P. J. Leo, C. J. Williams, and P. S. Julienne, “Determination of Cs–Cs interaction parameters using Feshbach spectroscopy,” *Comptes Rendus de l’Académie des Sciences - Series IV - Physics*, vol. 2, no. 4, pp. 633 – 639, 2001.
- [39] T. Weber, J. Herbig, M. Mark, H.-C. Nägerl, and R. Grimm, “Bose-Einstein condensation of cesium,” *Science*, vol. 299, no. 5604, pp. 232–235, 2003.
- [40] T. Kraemer, J. Herbig, M. Mark, T. Weber, C. Chin, H.-C. Nägerl, and R. Grimm, “Optimized production of a cesium Bose–Einstein condensate,” *Applied Physics B*, vol. 79, no. 8, pp. 1013–1019, 2004.
- [41] D. Boiron, A. Michaud, P. Lemonde, Y. Castin, C. Salomon, S. Weyers, K. Szymaniec, L. Cognet, and A. Clairon, “Laser cooling of cesium atoms in gray optical molasses down to $1.1\ \mu\text{K}$,” *Phys. Rev. A*, vol. 53, pp. R3734–R3737, 1996.
- [42] O. Mandel, M. Greiner, A. Widera, T. Rom, T. W. Hänsch, and I. Bloch, “Coherent transport of neutral atoms in spin-dependent optical lattice potentials,” *Phys. Rev. Lett.*, vol. 91, p. 010407, 2003.
- [43] D. C. McKay, C. Meldgin, D. Chen, and B. DeMarco, “Slow thermalization between a lattice and free Bose gas,” *Phys. Rev. Lett.*, vol. 111, p. 063002, 2013.
- [44] C. Robens, J. Zopes, W. Alt, S. Brakhane, D. Meschede, and A. Alberti, “Low-entropy states of neutral atoms in polarization-synthesized optical lattices,” *Phys. Rev. Lett.*, vol. 118, p. 065302, 2017.
- [45] M. Berninger, A. Zenesini, B. Huang, W. Harm, H.-C. Nägerl, F. Ferlaino, R. Grimm, P. S. Julienne, and J. M. Hutson, “Feshbach resonances, weakly

- bound molecular states, and coupled-channel potentials for cesium at high magnetic fields,” *Phys. Rev. A*, vol. 87, p. 032517, 2013.
- [46] A. Marte, T. Volz, J. Schuster, S. Dürr, G. Rempe, E. G. M. van Kempen, and B. J. Verhaar, “Feshbach resonances in rubidium 87: Precision measurement and analysis,” *Phys. Rev. Lett.*, vol. 89, p. 283202, 2002.
- [47] S. Inouye, M. R. Andrews, J. Stenger, H. J. Miesner, D. M. Stamper-Kurn, and W. Ketterle, “Observation of Feshbach resonances in a Bose–Einstein condensate,” *Nature*, vol. 392, p. 151, 1998.
- [48] C.-L. Hung, X. Zhang, N. Gemelke, and C. Chin, “Accelerating evaporative cooling of atoms into Bose-Einstein condensation in optical traps,” *Phys. Rev. A*, vol. 78, p. 011604, 2008.
- [49] M. P. Köppinger, P. D. Gregory, D. L. Jenkin, D. J. McCarron, A. L. Marchant, and S. L. Cornish, “Repeated output coupling of ultracold Feshbach molecules from a Cs BEC,” *New Journal of Physics*, vol. 16, no. 11, p. 115016, 2014.
- [50] T. Mežnaršič, T. Arh, J. Brence, J. Pišljarič, K. Gosar, i. c. v. Gosar, R. Žitko, E. Zupanič, and P. Jeglič, “Cesium bright matter-wave solitons and soliton trains,” *Phys. Rev. A*, vol. 99, p. 033625, 2019.
- [51] A. Di Carli, C. D. Colquhoun, S. Kuhr, and E. Haller, “Interferometric measurement of micro-g acceleration with levitated atoms,” *New Journal of Physics*, vol. 666, pp. 666–666, 2019.
- [52] W. Ketterle, D. S. Durfee, and D. M. Stamper-Kurn, “Making, probing and understanding Bose-Einstein condensates,” in *Proceedings of the International School of Physics Enrico Fermi Course CXL*, IOS Press, Amsterdam, 1999. arXiv:cond-mat/9904034.
- [53] R. Olf, F. Fang, G. E. Marti, A. MacRae, and D. M. Stamper-Kurn, “Thermometry and cooling of a Bose gas to 0.02 times the condensation temperature,” *Nature Physics*, vol. 11, pp. 720–723, 2015.
- [54] R. S. Lous, I. Fritsche, M. Jag, B. Huang, and R. Grimm, “Thermometry of a deeply degenerate Fermi gas with a Bose-Einstein condensate,” *Phys. Rev. A*, vol. 95, p. 053627, 2017.
- [55] C. Sabín, A. White, L. Hackermüller, and I. Fuentes, “Impurities as a quantum thermometer for a Bose-Einstein condensate,” *Scientific Reports*, vol. 4, no. 6436, 2014.

- [56] M. Mehboudi, A. Lampo, C. Charalambous, L. A. Correa, M. A. García-March, and M. Lewenstein, “Using polarons for sub-nk quantum nondestruction thermometry in a Bose-Einstein condensate,” *Phys. Rev. Lett.*, vol. 122, p. 030403, 2019.
- [57] A. E. Leanhardt, T. A. Pasquini, M. Saba, A. Schirotzek, Y. Shin, D. Kielpinski, D. E. Pritchard, and W. Ketterle, “Cooling Bose-Einstein condensates below 500 picokelvin,” *Science*, vol. 301, no. 5639, pp. 1513–1515, 2003.
- [58] A. Di Carli, *A tuneable quantum gas for matter wave interferometry and soliton experiments*. PhD thesis, University of Strathclyde, 2019.
- [59] D.-J. Han, S. Wolf, S. Oliver, C. McCormick, M. T. DePue, and D. S. Weiss, “3D Raman sideband cooling of cesium atoms at high density,” *Phys. Rev. Lett.*, vol. 85, pp. 724–727, 2000.
- [60] J. Kepler, *De cometis libelli tres ... De cometis libelli tres*, Typis Andreae Apergeri, sumptibus Sebastiani Mylii bibliopolæ Augustani, 1619.
- [61] J. C. Maxwell, *A Treatise on Electricity and Magnetism*, vol. 2 of *Cambridge Library Collection - Physical Sciences*. Cambridge University Press, 2010.
- [62] O. R. Frisch, “Experimenteller nachweis des einsteinschen strahlungsrückstoßes,” *Zeitschrift für Physik*, vol. 86, no. 1, pp. 42–48, 1933.
- [63] A. Ashkin, “Acceleration and trapping of particles by radiation pressure,” *Phys. Rev. Lett.*, vol. 24, pp. 156–159, 1970.
- [64] T. Hänsch and A. Schawlow, “Cooling of gases by laser radiation,” *Optics Communications*, vol. 13, no. 1, pp. 68 – 69, 1975.
- [65] 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, 1978.
- [66] W. D. Phillips and H. Metcalf, “Laser deceleration of an atomic beam,” *Phys. Rev. Lett.*, vol. 48, pp. 596–599, 1982.
- [67] 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, 1988.
- [68] D. A. Steck, “Cesium D line data,” Available online at <https://steck.us/alkalidata/>, 1998.

- [69] 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, 1987.
- [70] E. Riis, D. S. Weiss, K. A. Moler, and S. Chu, “Atom funnel for the production of a slow, high-density atomic beam,” *Phys. Rev. Lett.*, vol. 64, pp. 1658–1661, 1990.
- [71] C. Adams and E. Riis, “Laser cooling and manipulation of neutral particles,” *Progress in Quantum Electronics*, vol. 21, 1997.
- [72] A. J. Kerman, V. Vuletić, C. Chin, and S. Chu, “Beyond optical molasses: 3D Raman sideband cooling of atomic cesium to high phase-space density,” *Phys. Rev. Lett.*, vol. 84, pp. 439–442, 2000.
- [73] V. Vuletić, C. Chin, A. J. Kerman, and S. Chu, “Degenerate Raman sideband cooling of trapped cesium atoms at very high atomic densities,” *Phys. Rev. Lett.*, vol. 81, pp. 5768–5771, 1998.
- [74] D. J. Han, M. T. DePue, and D. S. Weiss, “Loading and compressing Cs atoms in a very far-off-resonant light trap,” *Phys. Rev. A*, vol. 63, p. 023405, 2001.
- [75] J. Herbig, T. Kraemer, M. Mark, T. Weber, C. Chin, H.-C. Nägerl, and R. Grimm, “Preparation of a pure molecular quantum gas,” *Science*, vol. 301, no. 5639, pp. 1510–1513, 2003.
- [76] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov, “Optical dipole traps for neutral atoms,” *Advances In Atomic, Molecular, and Optical Physics*, vol. 42, pp. 95 – 170, 2000.
- [77] M. Greiner, *Ultracold quantum gases in three-dimensional optical lattice potentials*. PhD thesis, Ludwig-Maximilians-Universität München, 2003.
- [78] P. W. H. Pinkse, A. Mosk, M. Weidemüller, M. W. Reynolds, T. W. Hijmans, and J. T. M. Walraven, “Adiabatically changing the phase-space density of a trapped Bose gas,” *Phys. Rev. Lett.*, vol. 78, pp. 990–993, 1997.
- [79] D. M. Stamper-Kurn, H.-J. Miesner, A. P. Chikkatur, S. Inouye, J. Stenger, and W. Ketterle, “Reversible formation of a Bose-Einstein condensate,” *Phys. Rev. Lett.*, vol. 81, pp. 2194–2197, 1998.
- [80] T. Weber, J. Herbig, M. Mark, H.-C. Nägerl, and R. Grimm, “Three-body recombination at large scattering lengths in an ultracold atomic gas,” *Phys. Rev. Lett.*, vol. 91, p. 123201, 2003.

- [81] A. Mair Thomas, *Ultracold collisions and evaporative cooling of caesium in a magnetic trap*. PhD thesis, University of Oxford, 2004.
- [82] W. Ketterle and N. V. Druten, “Evaporative cooling of trapped atoms,” *Advances In Atomic, Molecular, and Optical Physics*, vol. 37, pp. 181 – 236, 1996.
- [83] S. N. Bose, “Plancks gesetz und lichtquantenhypothese,” *Zeitschrift für Physik*, vol. 26, no. 1, pp. 178–181, 1924.
- [84] A. Einstein, “Quantentheorie des einatomigen idealen gases,” *Sitzungsber. Preuss. Akad. Wiss., Phys. Math. Kl. Bericht 3*, vol. 18, 1925.
- [85] S. L. Cornish, N. R. Claussen, J. L. Roberts, E. A. Cornell, and C. E. Wieman, “Stable ^{85}Rb Bose-Einstein condensates with widely tunable interactions,” *Phys. Rev. Lett.*, vol. 85, pp. 1795–1798, 2000.
- [86] E. A. Donley, N. R. Claussen, S. L. Cornish, J. L. Roberts, E. A. Cornell, and C. E. Wieman, “Dynamics of collapsing and exploding Bose–Einstein condensates,” *Nature*, vol. 412, pp. 295–299, 2001.
- [87] R. Wynar, R. S. Freeland, D. J. Han, C. Ryu, and D. J. Heinzen, “Molecules in a Bose-Einstein condensate,” *Science*, vol. 287, no. 5455, pp. 1016–1019, 2000.
- [88] T. Kraemer, M. Mark, P. Waldburger, J. G. Danzl, C. Chin, B. Engeser, A. D. Lange, K. Pich, A. Jaakkola, H.-C. Nägerl, and R. Grimm, “Evidence for Efimov quantum states in an ultracold gas of caesium atoms,” *Nature*, vol. 440, pp. 315–318, 2006.
- [89] C. J. Pethick and H. Smith, *Bose–Einstein Condensation in Dilute Gases*. Cambridge University Press, 2 ed., 2008.
- [90] E. P. Gross, “Structure of a quantized vortex in boson systems,” *Il Nuovo Cimento (1955-1965)*, vol. 20, no. 3, pp. 454–477, 1961.
- [91] L. P. Pitaevskii, “Vortex lines in an imperfect bose gas,” *Journal of Experimental and Theoretical Physics*, vol. 13, pp. 451–454, 1961.
- [92] P. A. Ruprecht, M. J. Holland, K. Burnett, and M. Edwards, “Time-dependent solution of the nonlinear Schrödinger equation for Bose-condensed trapped neutral atoms,” *Phys. Rev. A*, vol. 51, pp. 4704–4711, 1995.
- [93] J. L. Roberts, N. R. Claussen, S. L. Cornish, E. A. Donley, E. A. Cornell, and C. E. Wieman, “Controlled collapse of a Bose-Einstein condensate,” *Phys. Rev. Lett.*, vol. 86, pp. 4211–4214, 2001.

- [94] Y. Castin and R. Dum, “Bose-Einstein condensates in time dependent traps,” *Phys. Rev. Lett.*, vol. 77, pp. 5315–5319, 1996.
- [95] Y. Kagan, E. L. Surkov, and G. V. Shlyapnikov, “Evolution of a Bose-condensed gas under variations of the confining potential,” *Phys. Rev. A*, vol. 54, pp. R1753–R1756, 1996.
- [96] F. Dalfovo, C. Minniti, S. Stringari, and L. Pitaevskii, “Nonlinear dynamics of a Bose condensed gas,” *Physics Letters A*, vol. 227, no. 3, pp. 259 – 264, 1997.
- [97] M. Naraschewski and D. M. Stamper-Kurn, “Analytical description of a trapped semi-ideal bose gas at finite temperature,” *Phys. Rev. A*, vol. 58, pp. 2423–2426, 1998.
- [98] M.-S. Chang, *Coherent Spin Dynamics of a Spin-1 Bose-Einstein Condensate*. PhD thesis, Georgia Institute of Technology, 2006.
- [99] A. S. Arnold, J. S. Wilson, and M. G. Boshier, “A simple extended-cavity diode laser,” *Review of Scientific Instruments*, vol. 69, no. 3, pp. 1236–1239, 1998.
- [100] A. L. Schawlow, “Spectroscopy in a new light,” *Rev. Mod. Phys.*, vol. 54, pp. 697–707, 1982.
- [101] D. W. Preston, “Doppler-free saturated absorption: Laser spectroscopy,” *American Journal of Physics*, vol. 64, no. 11, pp. 1432–1436, 1996.
- [102] B. E. Sherlock and I. G. Hughes, “How weak is a weak probe in laser spectroscopy?,” *American Journal of Physics*, vol. 77, no. 2, pp. 111–115, 2009.
- [103] C. Wieman and T. W. Hänsch, “Doppler-free laser polarization spectroscopy,” *Phys. Rev. Lett.*, vol. 36, pp. 1170–1173, 1976.
- [104] G. D. Rovera, G. Santarelli, and A. Clairon, “A laser diode system stabilized on the caesium D2 line,” *Review of Scientific Instruments*, vol. 65, no. 5, pp. 1502–1505, 1994.
- [105] C. P. Pearman, C. S. Adams, S. G. Cox, P. F. Griffin, D. A. Smith, and I. G. Hughes, “Polarization spectroscopy of a closed atomic transition: applications to laser frequency locking,” *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 35, no. 24, p. 5141, 2002.
- [106] C. S. Adams, “A mechanical shutter for light using piezoelectric actuators,” *Review of Scientific Instruments*, vol. 71, no. 1, pp. 59–60, 2000.

- [107] K. Singer, S. Jochim, M. Mudrich, A. Mosk, and M. Weidemüller, “Low-cost mechanical shutter for light beams,” *Review of Scientific Instruments*, vol. 73, no. 12, pp. 4402–4404, 2002.
- [108] L. P. Maguire, S. Szilagyi, and R. E. Scholten, “High performance laser shutter using a hard disk drive voice-coil actuator,” *Review of Scientific Instruments*, vol. 75, no. 9, pp. 3077–3079, 2004.
- [109] C. D. Colquhoun, A. Di Carli, S. Kuhr, and E. Haller, “Note: A simple laser shutter with protective shielding for beam powers up to 1 W,” *Review of Scientific Instruments*, vol. 89, no. 12, p. 126102, 2018.
- [110] O. Morsch and M. Oberthaler, “Dynamics of Bose-Einstein condensates in optical lattices,” *Rev. Mod. Phys.*, vol. 78, pp. 179–215, 2006.
- [111] Analog Devices, AD9958 Datasheet: *2-Channel, 500 MSPS DDS with 10-Bit DACs*, 2016. Rev. C.
- [112] 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, no. 11, pp. 2084–2107, 1989.
- [113] P. Treutlein, K. Y. Chung, and S. Chu, “High-brightness atom source for atomic fountains,” *Phys. Rev. A*, vol. 63, p. 051401, 2001.
- [114] M. Gustavsson, *A quantum gas with tunable interactions in an optical lattice*. PhD thesis, University of Innsbruck, 2008.
- [115] T. A. Savard, K. M. O’Hara, and J. E. Thomas, “Laser-noise-induced heating in far-off resonance optical traps,” *Phys. Rev. A*, vol. 56, pp. R1095–R1098, 1997.
- [116] M. E. Gehm, K. M. O’Hara, T. A. Savard, and J. E. Thomas, “Dynamics of noise-induced heating in atom traps,” *Phys. Rev. A*, vol. 58, pp. 3914–3921, 1998.
- [117] D. S. Petrov, D. M. Gangardt, and G. V. Shlyapnikov, “Low-dimensional trapped gases,” *J. Phys. IV France*, vol. 116, pp. 5–44, 2004.
- [118] D. S. Petrov, G. V. Shlyapnikov, and J. T. M. Walraven, “Regimes of quantum degeneracy in trapped 1D gases,” *Phys. Rev. Lett.*, vol. 85, pp. 3745–3749, 2000.
- [119] I. Bloch, “Ultracold quantum gases in optical lattices,” *Nature Physics*, vol. 1, no. 1, pp. 23–30, 2005.

BIBLIOGRAPHY

- [120] D. S. Petrov, M. Holzmann, and G. V. Shlyapnikov, “Bose-Einstein condensation in quasi-2D trapped gases,” *Phys. Rev. Lett.*, vol. 84, pp. 2551–2555, 2000.
- [121] J. H. Denschlag, J. E. Simsarian, H. Häffner, C. McKenzie, A. Browaeys, D. Cho, K. Helmerson, S. L. Rolston, and W. D. Phillips, “A Bose-Einstein condensate in an optical lattice,” *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 35, no. 14, pp. 3095–3110, 2002.