

Two-dimensional Bose gases in systems with pointlike disorder



David Garrick
Keble College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity Term 2022

For my mother, Janet

Two-dimensional Bose gases in systems with pointlike disorder

David Garrick

Keble College
University of Oxford

*A thesis submitted for the degree of
Doctor of Philosophy*

Abstract

The effect of disorder on many-body physics is important for our understanding of real-world systems. In particular, the behaviour of phase transitions between superfluid, normal and insulating phases is an area with many unanswered questions and experimental evidence remains scarce in certain important cases. In this thesis, we focus on systems of two-dimensional (2D) Bose gases subjected to a pointlike disorder potential.

We trap atoms in radiofrequency dressed potentials and apply the technique of multiple-radiofrequency dressing to create a double-well potential into which we load the ultracold atoms. Matter-wave interferometry allows us to directly probe the local relative phase and its fluctuations between the clouds in the two potential wells. A pointlike disordered optical potential is produced by a high-intensity laser beam which is shaped by a programmable mirror array and passes through clouds in both of the wells. We adiabatically apply disordered potentials of increasing strength to observe their effect on the superfluids in equilibrium in a double-well. We observe that above a critical disorder strength the superfluid undergoes a transition to the normal phase at a lower temperature than that of the disorder-free case, and we characterise this disorder-induced transition.

We also study the flow of a 2D superfluid cloud around both a single defect and a random potential and measure the dissipation induced by the obstacles. The absence of dissipation is observed below a critical velocity in the case of a single defect but we find no such dissipationless regime for the disordered potential.

We describe the development of a software simulation package for modelling cold-atom systems and the implementation of particle collisions, magnetic trapping and evaporative cooling. Finally, we report on the initial work towards construction of a new cold-atom apparatus, capable of faster sequence times and greater atom numbers, to continue our investigation of many-body physics.

Acknowledgments

I would like to thank my supervisor Chris, whose experience and knowledge made him invaluable as he guided me through my DPhil. I also cannot give enough thanks to my friends in my lab group: Elliot, Kathrin, Adam and Shinichi who very patiently taught me so much about experimental physics when I started with no experience or competence whatsoever; and Abel, who selflessly helped me replace an ungodly number of vacuum flanges over the last few years, among so many other things. Thank you to all of the above as well as Abigail, Sean, Leo, and Tom for making the last four years fun when the physics itself was insufficient to sustain my soul.

Thank you to all my other friends in Oxford: Alex, Emi, Jess, Jake, Amy, Matt, Allison and so many others, who were effortlessly supportive, reminded me that a world outside the lab existed, and put up with my continuous claims that I was just so busy with my DPhil right now.

Thank you to my father Andrew, who has always taught me to be curious and interested in the world; to my sisters Eleanor and Katie, who supported and encouraged their little brother even when he was an annoying child; and thank you to Barbara and Max, whom I am thrilled to have in my family.

Thank you to my mother Janet, who always told me that she believed I could do anything I wanted. I have tried to live up to her faith in me; I hope she would be proud.

Finally, I must thank Ella, who more than any other person has made me happy and sane enough to keep me going for four years. I am so grateful for everything she has gone out of her way to do for me, when just her being there would have been enough.

Publications

This thesis presents novel work for which a publication is in preparation:

- **D. Garrick**, S. Sunami, V. Singh, A. Beregi, E. Bentine, L. Mathey and C. J. Foot. *The superfluid-normal transition in 2D Bose gases with pointlike disorder*. In preparation

The following related work to which I have contributed is detailed in the theses of previous group members:

- S. Sunami, V. Singh, **D. Garrick**, A. Beregi, E. Bentine, L. Mathey and C. J. Foot. *Density-noise correlation in 2D Bose gases across the BKT critical point*. In preparation
- S. Sunami, V. Singh, **D. Garrick**, A. Beregi, A. J. Barker, K. Luksch, E. Bentine, L. Mathey, C. J. Foot. *Universal scaling of the dynamic BKT transition in quenched 2D Bose gases*. Submitted for publication
- S. Sunami, V. Singh, **D. Garrick**, A. Beregi, A. J. Barker, K. Luksch, E. Bentine, L. Mathey, C. J. Foot. *Observation of the Berezinskii-Kosterlitz-Thouless Transition in a Two-Dimensional Bose Gas via Matter-Wave Interferometry*, Physical Review Letters **128**, 250402 (2022)
- A. J. Barker, S. Sunami, **D. Garrick**, A. Beregi, K. Luksch, E. Bentine, C. J. Foot. *Coherent splitting of two-dimensional Bose gases in magnetic potentials*, New Journal of Physics **22**, 103040 (2020)
- A. J. Barker, S. Sunami, **D. Garrick**, A. Beregi, K. Luksch, E. Bentine and C. J. Foot. *Realising a species-selective double-well using multiple-radiofrequency-dressed potentials*, J. Phys. B: At. Mol. Opt. Phys. **53** 155001 (2020)
- A. J. Barker, H. Style, K. Luksch, S. Sunami, **D. Garrick**, F. Hill, C. J. Foot and E. Bentine. *Applying machine learning optimization methods to the production of a quantum gas*, Machine Learning: Science and Technology **1**, 015007 (2020)

Contents

1	Introduction	1
2	Atoms in radiofrequency-dressed and optical dipole potentials	4
2.1	Atoms in static magnetic traps	4
2.1.1	Hyperfine structure of the atom	4
2.1.2	Energy levels in an external field	5
2.1.3	Magnetic trapping	6
2.1.3.1	Quadrupole trap	6
2.1.3.2	The Time-Orbiting potential	7
2.2	Atoms in radiofrequency-dressed potentials	8
2.2.1	The semi-classical approximation	9
2.2.2	Spatially varying static field and coupling strength	10
2.2.3	Multiple RF-dressing	13
2.2.4	Quasi-2D traps	14
2.3	Optical dipole potential	16
3	Experimental sequence and apparatus	18
3.1	Experimental sequence	18
3.1.1	Magneto-optical trap & magnetic transport	18
3.1.2	Magnetic traps	19
3.1.3	RF-dressed potentials	20
3.1.3.1	Circularly polarised shell trap	20
3.1.3.2	Single-well 2D clouds	21
3.1.3.3	MRF double-well splitting procedure	21
3.1.3.4	RF turnoff procedure	23
3.2	Vacuum system	23
3.2.1	Rubidium source	25
3.2.1.1	Rubidium dispensers	25

3.2.1.2	Installation of the rubidium chamber	25
3.2.1.3	Baking procedure	26
3.2.2	MOT and transport dependence on pressure	27
3.3	Laser system and imaging	29
3.3.1	Transitions in ^{87}Rb	29
3.3.2	Laser system	29
3.3.3	Imaging atomic clouds	31
3.3.3.1	Imaging systems	31
3.3.3.2	Selective repumping	32
3.4	Optical dipole potential	34
3.4.1	Focusing of dipole light	36
3.4.2	Calibrating potential depth	37
3.4.3	Point spread function	40
3.4.3.1	Measuring PSF of dipole light	41
4	Theory of disordered two-dimensional Bose gases	45
4.1	Phase transitions in the absence of disorder	45
4.1.1	Bose-Einstein condensation in a non-interacting Bose gas . . .	45
4.1.1.1	BECs in 2D harmonic traps	46
4.1.2	Weakly interacting Bose gases	47
4.1.2.1	Interactions in the quasi-2D limit	47
4.1.2.2	Density and phase fluctuations in the 2D gas	49
4.1.3	BKT transition	51
4.2	Disorder potentials	53
4.2.1	Forms of disorder	54
4.2.1.1	Speckle disorder	54
4.2.1.2	Pointlike disorder	57
4.2.2	Percolation theory	59
4.3	Many-body localisation in Bose gases	63
4.3.1	Thermalisation of quantum systems	63
4.3.2	Anderson localisation	65
4.3.3	Properties of MBL systems	67
4.3.4	Stability of MBL	68
4.3.5	Disordered phase transitions in Bose gases	70

5	The BKT transition in disordered systems	73
5.1	Experimental procedure	73
5.1.1	Pointlike disorder parameters	73
5.1.2	Loading the disordered double-well	76
5.2	Analysis of interference patterns	79
5.2.1	Fitting the interference pattern	79
5.2.2	Calculating correlation lengths	82
5.2.3	Accounting for inhomogeneity in the harmonic trap	84
5.2.4	Vortex detection	85
5.3	Observation of the disordered superfluid-normal transition	86
5.3.1	Percolation threshold	86
5.3.2	Depletion of the condensate	87
5.3.3	Correlations near the critical point	90
5.3.4	Vortex density	95
5.3.5	Superfluid regime at high disorder strengths	96
5.4	Conclusion	99
6	Damping of superfluid flow	100
6.1	Theory of superfluid flow	100
6.1.1	Landau criterion for superfluidity	100
6.1.2	Critical velocity for realistic obstacles	102
6.1.3	Excitations above the critical velocity	103
6.1.4	Superfluid flow in disordered potentials	104
6.2	Experimental procedure	106
6.2.1	Preparing the 2D cloud	106
6.2.2	Inducing flow around an obstacle	107
6.2.3	Measurement of damping rate	108
6.3	Superfluid flow around an obstacle	109
6.3.1	Observation of critical velocity	110
6.3.2	Frequency shift of oscillation	114
6.4	Damping in a disordered potential	116
6.4.1	Dependence on disorder strength	116
6.4.2	Dependence on initial flow velocity	118
6.5	Conclusion	120

7	AtomECS: Simulating atomic collisions and magnetic trapping	122
7.1	Overview of AtomECS	122
7.1.1	Entity-Component-System architecture	122
7.2	Direct Simulation Monte Carlo method	124
7.2.1	Kinetic theory of collisions	126
7.2.1.1	Partial-wave analysis of scattering	126
7.2.1.2	Scattering rate in a homogeneous gas	127
7.2.2	Implementation of collisions	129
7.2.2.1	Spatial partitioning	130
7.2.2.2	Collisions algorithm	131
7.2.3	Comparison with theory	132
7.2.3.1	Homogeneous trap	132
7.2.3.2	Harmonic trap	134
7.3	Magnetic trapping	136
7.3.1	Theory of magnetic trapping	136
7.3.2	Implementation of magnetic dipole force	137
7.3.3	The Time-Orbiting Potential	139
7.4	Simulating evaporative cooling	141
7.5	Conclusion	145
8	Towards a new experimental apparatus	146
8.1	Motivation	146
8.1.1	Sequence duration	146
8.1.2	Experimental complexity	148
8.1.3	Stability	148
8.1.4	New design concept	149
8.2	Laser system	149
8.2.1	Initial laser system	150
8.2.2	Current laser system	151
8.3	Recapture MOT	152
8.3.1	MOT distribution board	152
8.3.2	Magnetic fields	154
8.3.3	Optical layout	154
8.3.3.1	Periscope design	155
8.4	Vacuum system	157
8.4.1	Rubidium source	158

8.4.1.1	Initial rubidium chamber design	160
8.4.1.2	Current rubidium chamber design	161
8.4.2	Pyramid and cell vacuum systems	162
8.5	Characterising the MOTs	164
8.5.1	Operating the pyramid MOT source	164
8.5.2	Loading the recapture MOT	166
8.6	Towards RF-dressed trapping and BECs	168
8.7	Conclusion	171
9	Conclusion	172
9.1	Future experiments	173
9.1.1	Many-body localisation in 2D Bose gases	173
9.1.2	Potential shaping	174
9.1.3	Tunnelling in double-well potentials	175
	References	191

List of Figures

2.1	Uncoupled eigenstates of RF-dressed atoms	8
2.2	Eigenstates of of the RF-dressed atoms	11
2.3	MRF-dressed potentials	14
3.1	MRF double-well splitting procedure	22
3.2	3D model of current experimental apparatus	24
3.3	MOT and transport atom number after closing Rb source	27
3.4	Rb pressure and transport number after opening Rb source	28
3.5	Cooling and repumping transitions in ^{87}Rb	30
3.6	Optical dipole potential and thin repumping slice	33
3.7	Diagram of of the dipole potential, repumping and imaging optics .	35
3.8	Bisecting a cloud with the dipole potential	38
3.9	Calibrating dipole potential depth	40
3.10	Measuring point spread function of dipole optics	42
4.1	Simulated speckle disorder potential	55
4.2	Simulated pointlike disorder potential	58
4.3	Percolation in speckle disorder	60
4.4	Percolation in pointlike disorder	61
4.5	Percolation threshold as a function of pointlike filling factor	62
4.6	Phase diagram of disordered 2D Bose gas	71
5.1	Pointlike disorder shown on DMD array	74
5.2	BEC in pointlike disorder	75
5.3	Temperature and atom number for different disorder strengths	78
5.4	Fitting interference patterns and extracting phase profile	80
5.5	Two- and one-point correlation functions	83
5.6	Comparing algebraic and LCA fits to the correlation function	85
5.7	Bimodal fit of clouds in weak and strong disorder	88
5.8	Condensate depletion as a function of temperature for dilute disorder	89

5.9	Condensate depletion as a function of temperature for dense disorder	90
5.10	Condensate depletion as a function of disorder strength	91
5.11	Exponential and algebraic fits to correlation of clouds with disorder .	92
5.12	Superfluid-normal transition for dense disorder	93
5.13	Superfluid-normal transition for dilute disorder	94
5.14	Vortex density as a function of disorder strength	95
5.15	Superfluid correlations at high disorder	97
5.16	Possible mechanisms for high-disorder superfluidity	98
6.1	Damped oscillations of a superfluid around a barrier	110
6.2	Damping rates for barriers of different heights	111
6.3	Comparing oscillations below and above the critical velocity	112
6.4	Critical velocity dependence on barrier height	113
6.5	Frequency shift of dipole oscillations around a barrier	115
6.6	Damped oscillations in pointlike disorder	117
6.7	Damping rate as a function of disorder strength	118
6.8	Damping dependence on velocity in disorder	119
7.1	Simulated and predicted collision rate in a homogeneous trap	134
7.2	Simulated and predicted collision rate in a harmonic trap	135
7.3	Single-particle oscillation in a simulated TOP trap	141
7.4	Simulating TOP traps at low and high rotation frequencies	142
7.5	Simulating evaporative cooling to a BEC	144
8.1	MOT distribution board	153
8.2	MOT periscopes	157
8.3	Model of full vacuum system	158
8.4	Initial rubidium source design	160
8.5	Current rubidium source design	161
8.6	Model of pyramid MOT vacuum chamber	162
8.7	3D MOT in pyramid chamber	165
8.8	Pyramid MOT loading rate and bias field scan	166
8.9	Cell recapture MOT loading rate and bias field scan	167
8.10	Image of cell recapture MOT cloud	168
8.11	Loading a MOT in the cell from background vapour	169
8.12	Loading a shell trap directly from the MOT	170

Chapter 1

Introduction

Real physical systems typically have imperfections compared to idealised physical models. For example, condensed matter systems can feature random impurities or vacancies in a crystal formation; even in the absence of impurities, disorder may be intrinsic to the physics of certain systems such as high-temperature superconductors [1, 2]. Studying the effect of disorder on the physics of many-body systems is therefore crucial to understanding their behaviour.

Disorder can be expected to both quantitatively and qualitatively change important properties of the system. It has been predicted and observed to change the critical temperature of the transition from collective many-body states such as superfluids or superconductors to normal phases [3–7]. Usually disorder is expected to suppress such effects by reducing the coherence of the cloud, but in some cases the behaviour is complex and disorder may even enhance superfluidity [4, 8] and superconductivity [9].

More strikingly, disorder fundamentally alters our understanding of thermalisation in quantum systems and produces entirely new states of matter. The earliest known example of such a phase transition is the famous Anderson localisation [10], a disorder-induced transition to an insulating state in which transport is completely suppressed and the system fails to thermalise. Much experimental and theoretical effort has been devoted to expanding our understanding of this phenomenon, including the case of interacting particles [11–13] and the effect of the dimensionality of the system [14]. Anderson localisation is fundamentally a property of the ground state of a non-interacting system; the extension to many particles and finite temperatures led to the discovery of transitions to many-body localised (MBL) states [15, 16]. Within our current knowledge, disorder appears to be the primary mechanism through which these localised transitions occur, although it is not the only one [17, 18].

MBL states are not yet well understood, but it is clear that the dimensionality of the system has a very strong impact on the observed behaviour [19, 20]. Specifically in the case of Anderson localisation, in one dimension the transition is proven to occur for arbitrarily small disorder strengths [21]. In two dimensions (2D) a similar result is expected, but not proven; and in 3D no localisation occurs below a critical disorder strength.

In a many-body system the picture is more complicated. One issue is that the stability of the MBL insulating state is a matter of some debate, particularly the dependence of this stability on the dimensionality of the system [15, 19, 20, 22]. Furthermore, in many-body systems the insulating and normal phases may not be the only possible phases of the system. In the case of Bose gases in one dimension (1D) at non-zero temperature, only a transition between the insulating and normal fluid occurs. However, in 2D a transition to an algebraic superfluid is observed, known as the Berezinskii-Kosterlitz-Thouless (BKT) transition [23, 24]. The BKT transition has been observed in Bose gases [25, 26], He films [27], Fermi gases [28], and 2D superconductors [29, 30].

Importantly, the phenomenon of universality implies that these transitions are insensitive to the microscopic details of the system. The behaviour of the physics in the critical region near the phase transition is universal to all systems within the same universality class. The study of disorder and its effect on phase transitions in one system can therefore be applied to a range of systems and fields of study.

Ultracold atoms in particular are a powerful tool for studying these phenomena. Advances in trapping and manipulation techniques allow for the production of a system that is stable and highly isolated from the environment. A variety of Hamiltonians can be engineered and tuned precisely over a wide parameter range, and the timescale of dynamics is usually slow enough to be easily observed with standard experimental methods and for properties of the system measured with precision. In particular, matter-wave interferometry techniques are very effective for probing the phase dynamics of atomic clouds [25, 26, 31–33]. This allows the measurement of otherwise elusive properties such as phase correlations within a quantum gas.

While optical trapping methods are widely used and provide versatile trapping geometries, magnetic fields are another common technique and these provide smoother potentials and lower heating rates. Such a smooth, defect-free potential is important for the study of disorder, where we wish to control all aspects of the random potential experienced by the atoms. Flexible magnetic trapping potentials are formed by the technique of radiofrequency (RF) dressing [34–37]. RF-dressed potentials can be

sculpted into ring traps [38, 39], lattices [37], and the double-well potentials [32, 36, 40, 41] which are of primary interest to us in this work.

In this thesis, we present the experimental study of the effect of disorder on 2D Bose gases, and progress on the experimental techniques used to create them. We begin in Chapter 2 by presenting background theory relating to the techniques of magnetic trapping, radiofrequency-dressed potentials, and the optical dipole potential used to create the random disorder. In Chapter 3 we describe the current experimental apparatus used to produce and study disordered 2D Bose gases, which was used to obtain the experimental results for this thesis. Chapter 4 provides some explanatory theory of the BKT transition in 2D gases as well as reviewing some predicted and observed effects of disorder on quantum states, and discusses the nature of the many-body localised states that can arise from disorder. In Chapter 5 we present our experimental results on the effect of pointlike disorder on the BKT transition utilising the technique of matter-wave interferometry to probe the correlation function of the cloud. Chapter 6 details a preliminary investigation of the dynamics of 2D superfluid flow around both a single barrier and disordered potentials, by measuring the damping of dipole oscillations around such obstacles.

After these studies into disordered Bose gases, we turn to the design of future cold-atom experiments. Chapter 7 details contributions to the development of a cold-atom software simulation package, AtomECS, that aims to assist researchers in the design and modelling of experimental systems. Chapter 8 describes the design and initial steps towards the construction of a new, higher performance experimental apparatus. This system will be capable of faster sequence times and higher atom numbers, which will serve to accelerate our future research into the field of many-body physics, and is still undergoing development at the time of writing of this thesis. Finally, Chapter 9 summarises the work presented in this thesis and discusses the potential future experimental avenues of study.

Chapter 2

Atoms in radiofrequency-dressed and optical dipole potentials

This chapter introduces the theoretical foundation of trapping cold atoms. We discuss simple magnetic traps and the dressed-atom formalism and its application to trapping, including creating the highly anisotropic systems required to produce a two-dimensional gas. We expand this theory to multiple RF fields and thereby the production of controllable double-well potentials. Further detail on the theory of multiple radiofrequency-dressing and its application to our experiments can be found in previous group theses [40–42]. Finally we briefly discuss the optical dipole force and its use in trapping or shaping potentials.

2.1 Atoms in static magnetic traps

2.1.1 Hyperfine structure of the atom

The total magnetic moment of an atom is comprised of contributions from the angular momentum of both the nucleus and the electrons. The electronic magnetic moment itself contains contributions from the orbital angular momentum $\mathbf{L}$ and spin angular momentum $\mathbf{S}$ of the electrons. Associated with these angular momenta are the magnetic moments $\boldsymbol{\mu}_L, \boldsymbol{\mu}_S$, defined by

$$\boldsymbol{\mu}_L = -g_L \mu_B \mathbf{L}, \quad \boldsymbol{\mu}_S = -g_S \mu_B \mathbf{S}, \quad (2.1)$$

where g_L and g_S are constants of proportionality known as g-factors.

The first correction to the gross structure of the atom (where energy levels are labelled by the principal quantum number n and the electronic orbital angular momentum L) arises from the spin-orbit interaction between these two magnetic moments. This interaction is proportional to $\mathbf{S} \cdot \mathbf{L}$ and couples these momenta together such

that in general the projections along the z direction, m_L and m_S , are no longer good quantum numbers. Instead we label the new states by their total electronic angular momentum $\mathbf{J} = \mathbf{L} + \mathbf{S}$, which remains well-defined, and has an associated magnetic moment, $\boldsymbol{\mu}_J$. The additional splitting of energy spectra caused by this coupling is the fine structure of the atom.

In a similar manner, the angular momentum of the nucleus $\mathbf{I}$ gives rise to a nuclear magnetic moment $\boldsymbol{\mu}_I$. The interaction of the nuclear and electronic magnetic moments, of the form $\mathbf{I} \cdot \mathbf{J}$, again lead us to label the eigenstates of this magnetic-dipole interaction by the well-defined total angular momentum $\mathbf{F} = \mathbf{I} + \mathbf{J}$. This splitting (smaller than that of the fine structure) is called the hyperfine structure. The magnetic moment of our atoms is therefore given by

$$\boldsymbol{\mu} = -g_F \mu_B \mathbf{F}. \quad (2.2)$$

This is also the case under an externally applied magnetic field as long as the strength of the interaction between the magnetic moment and the external field does not exceed that of the interaction causing either the fine or hyperfine structure. Otherwise, the total angular momenta are also not well defined, and instead only the projection onto the external field is a good quantum number. This condition of a weak external field will usually be fulfilled in cold atoms experiments.

2.1.2 Energy levels in an external field

The interaction at the core of magnetic trapping is that of an atom's magnetic moment $\boldsymbol{\mu}$ with an external magnetic field $\mathbf{B}$, corresponding to the Hamiltonian

$$H = -\boldsymbol{\mu} \cdot \mathbf{B}. \quad (2.3)$$

The energy of an atom in an external magnetic field is determined by the orientation of the atom's magnetic moment with respect to that field. Consider a magnetic field $\mathbf{B} = B_0 \mathbf{e}_z$, where we have chosen the axis of the magnetic field to be our z -axis. Our energy levels then are the eigenstates of Eqn. 2.3, Zeeman states labelled by $|m_F\rangle$ where m_F is the projection of the angular momentum $\mathbf{F}$ onto the magnetic field axis $\mathbf{e}_z$. The energies of these states, for weak field strengths, are

$$E = m_F g_F \mu_B B_0. \quad (2.4)$$

At higher field strengths, higher order terms in the perturbation can play a part, but are not relevant for the magnetic fields strengths used in this work. In the absence of

an external field, these Zeeman states all have the same energy; thus the application of the field lifts the degeneracy of these states. The projection m_F can take one of $2F + 1$ values from $-F, \dots, 0, \dots, F$ with corresponding different eigenenergies. Note that the g-factor g_F can be positive or negative. Therefore the sign of the energy will be determined by the factor $m_F g_F$. Intuitively, this factor describes whether the magnetic moment is aligned or anti-aligned to the external field.

The Zeeman states can vary locally in a spatially varying magnetic field. Thus a magnetic field strength with a spatial dependence $B_0(\mathbf{r})$ will translate into a spatially varying energy $E(\mathbf{r}) = m_F g_F \mu_B B_0(\mathbf{r})$. States with $m_F g_F < 0$ have decreasing energy for increasing B_0 , and are described as ‘high-field seeking’. We are more interested in states with $m_F g_F > 0$, since these states have increasing energy for higher B_0 and are therefore ‘low-field seeking’. Atoms in such states will move towards a local minimum of magnetic field strength, and can be trapped at that minimum. If we wish to trap atoms in a specific region in space, we need to create a field with a minimum of field strength in that region.

2.1.3 Magnetic trapping

2.1.3.1 Quadrupole trap

Any field with a local minimum in field strength could be used to trap atoms. One commonly used example is the spherical quadrupole field, given by

$$\mathbf{B} = B'(x\mathbf{e}_x + y\mathbf{e}_y - 2z\mathbf{e}_z), \quad (2.5)$$

where B' is the quadrupole gradient. Such a field is produced by a pair of current-carrying coils in an anti-Helmholtz configuration. It can be easily seen that this field has a minimum (in fact, a zero) of field strength at the origin, and it is around this point that atoms will be trapped. The potential energy of the atoms in this field is

$$U(x, y, z) = m_F g_F \mu_B B' \sqrt{x^2 + y^2 + 4z^2}. \quad (2.6)$$

Here we should note that although the quadrupole field does have a minimum at the origin, the fact that the field strength is zero complicates matters. At or very near the origin itself the Zeeman states become degenerate, and atoms can transition into other m_F states and thus leave the trap. This can be understood as an example of non-adiabatic Majorana spin flip transitions [43]. The result of this is that the quadrupole trap has a ‘hole’ at its center, through which atoms are gradually lost.

2.1.3.2 The Time-Orbiting potential

The hole at the centre of the quadrupole field can be removed in a number of ways, including ‘plugging’ the hole with a repulsive optical dipole potential created by a blue detuned laser [44]. In our experiments, the hole is removed by the addition of a rotating uniform magnetic bias field to form what is known as the time-orbiting potential or TOP trap [45]. Indeed, this was the first trap in which an atomic Bose-Einstein condensate was observed [46].

The bias field rotates at a frequency greater than that of the atomic motion, but slower than that of the atom’s Larmor frequency. This results in atoms that follow the local Zeeman state adiabatically, but feel only the time averaged effect of the rotating uniform field. The resulting trap is an axially symmetric, approximately harmonic potential, with no zero of magnetic field to limit the trap lifetime. The residual effect of the rotating bias field does impart a small micro-motion to the atoms, which is usually negligible.

The time-averaged potential U experienced by the atoms is [47]

$$U(x, y, z) = m_F g_F \mu_B \frac{B'^2}{4B_{TOP}} (x^2 + y^2 + z^2), \quad (2.7)$$

where B' is the quadrupole field gradient and B_{TOP} is the amplitude of the TOP bias field. The radial and axial trap frequencies are therefore

$$\omega_r^2 = m_F g_F \mu_B \frac{B'^2}{2mB_{TOP}}, \quad \omega_z^2 = 8\omega_r^2. \quad (2.8)$$

From the trap frequencies it is clear that the TOP trap is anisotropic, with a tighter confinement along the axial direction. This is qualitatively how one produces a quasi-2D cloud, but TOP traps do not produce a sufficiently anisotropic trapping geometry to enter this regime, so it produces 3D BECs when the atoms within are sufficiently cooled.

While the time-average of the TOP eliminates the zero of the magnetic field at the origin, instantaneously the node is displaced to a radius $r_0 = B_{TOP}/B'$. The rotating bias field sweeping the location of this node in a circle faster than the atoms can move, creating a circular locus of points at which atoms will experience zero magnetic field at some point during a rotation period. The same spin-flip mechanism as the quadrupole trap can cause atom loss along this so-called ‘circle of death’. In our experiments, for typical values of $B_{TOP} \sim 10$ G and $B' \sim 150$ G/cm, the radius of the circle of death is $r_0 \sim 600 \mu\text{m}$. This is small enough that it may contribute to atom loss during initial loading and evaporative cooling in the TOP trap, but is much larger than the size of the BEC produced at the end of the cooling stage ($\sim 25 \mu\text{m}$).




Figure 2.1: The tensor product of Zeeman states (a) and Fock states (b) produces a ladder of Zeeman states (c). We use here a negative sign of g_F , as is the case for ^{87}Rb in the $F = 1$ hyperfine state. For B such that $g_F\mu_B B = \hbar\omega$ (red dashed line), Zeeman states from different N manifolds become degenerate (highlighted in blue).

2.2 Atoms in radiofrequency-dressed potentials

Magnetic traps have a few advantages over, for example, optical traps; they are very smooth and free from defects, leading to very low heating rates. A general disadvantage however is that their trapping geometries can be limited or inflexible so that they are comparatively less versatile. Dressing the magnetic traps with a radiofrequency (RF) oscillating field is a method of expanding the available geometries in a flexible and controllable way.

The addition of the RF field to the undressed atoms in a static magnetic field produces a Hamiltonian of the form

$$H = H_0 + H_{field} + H_{int}, \quad (2.9)$$

where H_0 is the interaction of the atom with the static magnetic field, H_{field} is the Hamiltonian of the RF field alone, and H_{int} is the interaction of the atom with the oscillating RF field.

In the absence of the interaction H_{int} , the eigenstates are straightforwardly the tensor product of the atom Zeeman states and RF field Fock states, labelled $|m_F, N\rangle = |F, m_F\rangle \otimes |N\rangle$. The energies of this system then are an overlapping ‘ladder’ of Zeeman states, with one set of Zeeman states per photon number N .

The Zeeman manifolds are separated by the photon energy $\hbar\omega$. This implies that at a resonant magnetic field strength B such that $g_F\mu_B B = \hbar\omega$, states from different Zeeman manifolds become degenerate.

Now we consider the coupling of the atom to the RF field. This takes the form of the exchange of angular momentum and energy between the atom and the field, when the RF field induces transitions between Zeeman manifolds at these resonant field strengths.

2.2.1 The semi-classical approximation

To derive the resulting potential experienced by the atoms, we will treat the RF-dressed atoms in the semi-classical picture. This is accurate in the case of strong fields where the atoms interact with many photons and we describe the field with a classical, coherent RF field [41, 48]. A full quantum treatment produces the same results but is also more generally applicable to both weak RF fields and strong coupling, beyond the rotating wave approximation (which we will employ in our treatment here) [49].

We consider a spatially homogeneous RF field of general polarisation,

$$\mathbf{B} = B_x \cos(\omega t + \phi_x) \mathbf{e}_x + B_y \cos(\omega t + \phi_y) \mathbf{e}_y + B_z \cos(\omega t + \phi_z) \mathbf{e}_z. \quad (2.10)$$

In our case, the atoms are also trapped in a static external magnetic field. Initially we will consider the static field $\mathbf{B}_0 = B_0 \mathbf{e}_z$ (where the direction of the field $\mathbf{e}_z$ defines the quantization axis) to also be homogeneous, and will later generalise our result to spatially varying static fields. We also only use RF fields polarised within the x - y plane, and so we consider RF fields with $B_z = 0$. In fact, any component of the field along this direction (being the same as the axis of the magnetic field) has no effect on the coupling between undressed states within the rotating wave approximation [48].

The interaction between the magnetic moment of the atom and the field $\mathbf{B}$ is given by

$$H_{int} = -\frac{g_F \mu_B}{\hbar} \mathbf{F} \cdot \mathbf{B}. \quad (2.11)$$

From here on we will take the sign of g_F to be negative, as in the case of ^{87}Rb in the $F = 1$ hyperfine state, so that $g_F = -|g_F|$. Using the raising and lowering operators of the angular momentum ($\hat{F}_+$ and $\hat{F}_-$) we can rewrite $\mathbf{F}$ as

$$\mathbf{F} = \frac{\hat{F}_+ + \hat{F}_-}{2} \mathbf{e}_x + \frac{\hat{F}_+ - \hat{F}_-}{2i} \mathbf{e}_y + \hat{F}_z \mathbf{e}_z, \quad (2.12)$$

so that Eqn. 2.11 becomes

$$H_{int} = -\frac{g_F \mu_B}{\hbar} \left[\frac{\hat{F}_+}{2} (B_x \cos(\omega t + \phi_x) + i B_y \cos(\omega t + \phi_y)) + \frac{\hat{F}_-}{2} (B_x \cos(\omega t + \phi_x) - i B_y \cos(\omega t + \phi_y)) \right]. \quad (2.13)$$

We now move into the frame rotating around z at a frequency ω . In this frame all co-rotating terms become static while counter-rotating terms oscillate at 2ω and have little effect on average, so can be neglected. In this frame the Hamiltonian becomes [48]

$$H_{int} = -\hbar\delta F_z + \hbar\Omega F_x, \quad (2.14)$$

where $\hbar\delta = \hbar\omega - |g_F|\mu_B B_0$ is the detuning from resonance and Ω is the Rabi frequency given by

$$\Omega = \frac{|g_F|\mu_B}{2\hbar} \sqrt{B_x^2 + B_y^2 + 2B_x B_y \sin(\phi_y - \phi_x)}. \quad (2.15)$$

The eigenvalues of this Hamiltonian are

$$E = \hbar\tilde{m}_F \sqrt{\delta^2 + \Omega^2}, \quad (2.16)$$

where $\tilde{m}_F$ labels the dressed Zeeman states, rather than the uncoupled ones. The quantity $\sqrt{\delta^2 + \Omega^2}$ is sometimes referred to as the generalised Rabi frequency.

We are most concerned with three cases: linear polarisation, and two circular polarisations of opposite handedness. For linear polarisation, one component (e.g. B_y) is zero, so that the field strength is equal to the amplitude of just one component $B_x = B$ and the Rabi frequency is:

$$\Omega_{lin} = \frac{|g_F|\mu_B B}{2\hbar}. \quad (2.17)$$

One circular polarisation meanwhile has $B_x = B_y = B$ and $\phi_y - \phi_x = \pi/2$, so that the Rabi frequency is

$$\Omega_{circ} = \frac{|g_F|\mu_B B}{\hbar}, \quad (2.18)$$

which is twice the Rabi frequency for the linear polarisation with the same field strength along one direction. For the other circular polarisation with phase $\phi_y - \phi_x = -\pi/2$, the coupling and Rabi frequency goes to zero.

2.2.2 Spatially varying static field and coupling strength

We emphasize again that the above expressions are derived in the case of an homogeneous external magnetic field $\mathbf{B}_0 = B_0 \mathbf{e}_z$. We now consider the effect of a spatially varying magnetic field.

The significance of the direction $\mathbf{e}_z$ is that it defined the quantization axis of the system; components of the field along this direction are unaffected by a transformation into the rotating frame and so continue to oscillate at ω and are thus neglected in the rotating wave approximation. We expect therefore that the coupling of the RF




Figure 2.2: a) Energies of the uncoupled (black) and coupled (blue) Zeeman states as a function of position. The addition of the RF field creates a ladder of bare Zeeman states separated by $\hbar\omega$ in energy. Different m_F states in each manifold intersect at points where the Zeeman splitting is resonant with $\hbar\omega$. At these points $\delta(\mathbf{r}) = 0$, indicated by the vertical dashed line. The coupling of the atoms to the RF field introduces an avoided crossing at this point and the coupled m_F states are relabelled $\tilde{m}_F$ in the new basis. The separation of the new states at the resonance is set by the Rabi frequency Ω . b) Coupling strength $\Omega(\mathbf{r})$ over the resonant ellipsoid on which the atoms are trapped for a circularly polarised RF field. The coupling is maximum at the bottom of the shell, where atoms collect under gravity.

field to the atoms will vary spatially as $\mathbf{B}$ does, even though the RF field amplitudes themselves are constant.

As long as the magnetic field varies sufficiently slowly, and the atoms travel sufficiently slowly, the atoms can adiabatically follow the local quantization axis at their position [50]. The effect of the varying magnetic field is then to add a spatial dependence to the coupling, as well as the detuning. This results in a spatially dependent RF-dressed potential $V(\mathbf{r})$, given by

$$V(\mathbf{r}) = \tilde{m}_F \hbar \sqrt{\delta(\mathbf{r})^2 + \Omega(\mathbf{r})^2}. \quad (2.19)$$

We here refer to the spatially dependent $\Omega(\mathbf{r})$ as the coupling of the field to the atoms. The static field used in this work is the quadrupole field:

$$\mathbf{B}_0 = B'(x\mathbf{e}_x + y\mathbf{e}_y - 2z\mathbf{e}_z). \quad (2.20)$$

The detuning $\hbar\delta(\mathbf{r}) = g_F\mu_B|\mathbf{B}_0(\mathbf{r})| - \hbar\omega$ now varies spatially with the magnetic field. The resonant surface defined by $\delta(\mathbf{r}) = 0$ is an ellipsoid. On this surface the energy

contribution from the detuning is minimised. The atoms therefore will be trapped on the surface (typically falling to the bottom under the effect of gravity). The potential due to the variation in $\delta(\mathbf{r})$ is generally the largest contributor to the shape of the potential. This method of trapping atoms on the resonant ellipsoid of an RF dressed quadrupole field is often referred to as a shell trap.

The coupling $\Omega(\mathbf{r})$ varies spatially due to the rotation of the local quantization axis around the resonant ellipsoid. A full derivation of its dependence on position can be found in Ref. [48].

For the case of circular polarisation that couples to the atoms at the bottom of the ellipsoid, where they fall under gravity, the coupling is [41]:

$$\Omega(\mathbf{r}) = \frac{g_F \mu_B B}{2\hbar} \left(1 - \frac{2z}{\sqrt{x^2 + y^2 + 4z^2}} \right). \quad (2.21)$$

At the bottom of the ellipsoid this takes its maximum value $g_F \mu_B B / \hbar = \Omega_{circ}$ as in Eqn. 2.18. This is expected, since the bottom of the shell trap resembles the simple case where the quantization axis is directly upwards. Similarly, at the top of the ellipsoid the coupling goes to zero, since the local quantization axis here points downwards and the polarisation is locally the opposite handedness. The coupling over the surface of the shell smoothly decreases to zero from the bottom to the top of the trap and is cylindrically symmetric. The maximum coupling at the bottom of the trap is what we generally refer to as the Rabi frequency Ω_0 , since it matches the case of a homogeneous static field.

For linear polarisation along x the coupling is [41]:

$$\Omega(\mathbf{r}) = \frac{g_F \mu_B B}{2\hbar} \left(\frac{y^2 + z^2}{\sqrt{x^2 + y^2 + 4z^2}} \right). \quad (2.22)$$

Again, at the bottom of the trap where the atoms lie this takes the same value as Eqn. 2.17. The coupling in this case is constant and maximal along the points $y^2 + z^2$ in the y - z plane, where the RF field is perpendicular to the quadrupole field. The coupling has two nodes at the equator along x , where the RF field is parallel to the quadrupole field.

The gradient of $\Omega(\mathbf{r})$ exerts an additional force on the atoms towards areas of low coupling. This force modifies the strength of the confinement at the bottom of the trap, but is usually much weaker than gravity and the confinement to the resonant surface. For particularly large Rabi frequencies and small shell radii, this force can become comparable to the transverse confinement allowing atoms to creep up the sides of the ellipsoid towards the nodes of the coupling strength. In general this is

undesirable since the nodes have a high rate of non-adiabatic loss, but we operate with large shell radii and usually low Rabi frequencies, so the effect of atoms moving up the shell to the nodes is negligible.

Finally, we can include the effect of gravity, which is usually significantly stronger than the coupling force so that the atoms fall to the bottom of the shell. This can be added straightforwardly to the potential created by the fields, so that our final potential is

$$U(\mathbf{r}) = V(\mathbf{r}) + mgz = \tilde{m}_F \hbar \sqrt{\delta(\mathbf{r})^2 + \Omega(\mathbf{r})^2} + mgz, \quad (2.23)$$

where m is the mass of the atom and g is the gravitational field strength.

2.2.3 Multiple RF-dressing

The application of a single RF field creates an avoided crossing between the Zeeman states at the point where it becomes resonant with the Zeeman splitting. The addition of multiple RF fields, at different frequencies and amplitudes, creates additional avoided crossings. This permits the creation of a wide variety of possible trapping geometries, in particular the double-well potential [32, 36, 51].

To produce a double well, three RF fields are needed. Each additional field creates an avoided crossing of alternating curvature: the addition of a second field creates, in addition to the potential minimum of the single RF field, a potential maximum. The addition of a third field creates another potential minimum. The transition spectrum of a multiple radiofrequency (MRF) dressed state is in general very dense, which can lead to many potential loss mechanisms for the trapped atoms [40, 41]. Specifically, weakly avoided crossings also appear at the resonances with sums and differences of the dressing frequencies. For incommensurate frequencies, there are many such crossings. Furthermore, mixing of the dressing RF signals in the RF amplifiers used to generate the fields will drive transitions at the sum and difference frequencies.

This problem is mitigated by stipulating that the three dressing frequencies used, $\omega_1, \omega_2, \omega_3$ share a common fundamental frequency $\Delta\omega$. This considerably simplifies the treatment of the MRF dressed states and reduces the number of weakly avoided crossings, as well as loss-inducing noise from the amplifiers.

In the quadrupole field, each potential well manifests as concentric resonant ellipsoids, with the barrier field forming a repulsive ellipsoid between them. The radius of each ellipsoid in space is determined by the dressing frequency of the respective field; while the depth of each well is set by the amplitude (Rabi frequency) of the field.




Figure 2.3: Energy of eigenstates in the MRF double well, not including the effects of gravity. The double well is formed by the combination of three RF fields at dressing frequencies $\omega_1, \omega_2, \omega_3$. The potential minima form close to ω_1 and ω_3 , but the energy is weakly affected by the other fields, displacing the minima slightly. The height of the wells and barrier are controlled by the Rabi frequencies of their respective fields (again except for small perturbations from the other fields). Each field is a multiple of a common factor $\Delta\omega$. Weak avoided crossings form at further multiples of $\Delta\omega$ away from the dressing frequencies.

Because atoms in one well still weakly feel the effects of the other fields, these variables are mostly, but not entirely, independent. For example, increasing the barrier amplitude pushes the barrier down in energy space, which also brings the potential wells closer together. Fig. 2.3 shows the energies in a double-well potential. However, the effects of gravity are neglected in Fig. 2.3.

2.2.4 Quasi-2D traps

Since atoms are trapped at the bottom of the resonant ellipsoid, their transverse confinement is relatively weak, being determined primarily by gravity pulling them down the curvature of the shell. Their axial confinement in contrast is high as it is caused mostly by the change in the detuning from resonance $\delta(\mathbf{r})$. A power series expansion in the shell trap potential shows that at the bottom of the trap the confinement is

harmonic to a good approximation. The harmonic trap frequencies for a circularly polarised shell are given by [35]:

$$\omega_r = \sqrt{\frac{g}{4R}} \left(1 - \frac{\tilde{m}_F \hbar \Omega_0}{2MgR} \sqrt{1 - \epsilon^2} \right)^{1/2}, \quad (2.24)$$

$$\omega_z = 2|g_F| \mu_B B' \sqrt{\frac{\tilde{m}_F}{M \hbar \Omega_0}} (1 - \epsilon)^{3/4}, \quad (2.25)$$

where $\epsilon = Mg/2g_F \mu_B B' \tilde{m}_F$ characterises the relative strength of the quadrupole gradient and gravity. For $\epsilon > 1$, the quadrupole field is insufficiently strong to hold the atoms up under gravity and the atoms fall out of the trap. R is the vertical location of the atoms beneath the centre of the quadrupole given by

$$R = \frac{r_0}{2} \left(1 + \frac{\epsilon}{\sqrt{1 - \epsilon^2}} \frac{\Omega_0}{\omega} \right), \quad (2.26)$$

where $r_0 = g_F \mu_B B' / \hbar \omega$ is the equatorial radius of the resonant ellipsoid. For $\epsilon \ll 1$, as it usually is, the atoms' position $R \approx r_0/2$. For a shell trap linearly polarised along x , ω_x remains the same but ω_y becomes

$$\omega_y = \sqrt{\frac{g}{4R}}. \quad (2.27)$$

The linear shell trap frequencies are then marginally different in the x and y directions. For most of the work presented here, the difference is very small, approximately 2%.

Typical values of the trap frequencies are tens of Hz for the transverse confinement, while the axial confinement can exceed 1 kHz. This strongly anisotropic potential compresses the cloud of atoms into a pancake shape, being much narrower in the vertical direction than the radial. This anisotropy is the key element to producing quasi-two-dimensional (2D) condensed clouds. By increasing ω_z , we increase the energy spacing between the lowest and first excited state in that direction, which is $\hbar \omega_z$. If this energy spacing exceeds the other relevant energy scales in the cloud, being temperature $k_B T$ and chemical potential μ , then states along the axial direction become energetically inaccessible. This freezes out possible dynamics along the z direction, resulting in a quasi-2D physical system. Details of the physics of 2D Bose gases are discussed later in Chapter 4.

To achieve this quasi-2D regime, we must increase the axial trapping frequency. From Eqn. 2.25 we see that is achieved by either increasing the quadrupole gradient B' or decreasing the Rabi frequency Ω . However, the radial trap frequency also increases under these conditions, and a tighter radial confinement increases the density of the

cloud and hence increases μ , pushing the cloud further out of the 2D regime. In practice, increasing the quadrupole gradient or decreasing Rabi frequency increases ω_z faster than it does μ , so that the ratio $\mu/\hbar\omega_z$ still decreases.

2.3 Optical dipole potential

Optical dipole traps are a commonly used alternative method to magnetic fields for trapping cold atoms. In this work we use the optical dipole force not to trap by itself, but to modify the potential seen by the atoms in some way, such as producing a random disordered potential.

Consider an electric dipole in the presence of an external electric field $\mathbf{E}$. The dipole moment is given by $-\mathbf{er} = \epsilon_0\chi\mathbf{E}$ where χ is the polarizability of the atom and e is the electronic charge. The force then originates from the interaction energy of the dipole with the electric field:

$$U = \frac{1}{2}\mathbf{er} \cdot \mathbf{E} = -\frac{1}{2}\epsilon_0\chi E(x, y, z)^2, \quad (2.28)$$

where $E(x, y, z)$ is the amplitude of the electric field. Since this energy is proportional to the intensity I of an external field, a gradient in the field will produce a force in the direction of the gradient. If we consider the external field to be a monochromatic plane wave, it has a Rabi frequency $\Omega = \langle 1 | \mathbf{er} \cdot \mathbf{E} | 2 \rangle / \hbar \propto \sqrt{I}$ where the states $|1\rangle, |2\rangle$ are the lower and upper states of the resonant transition [52].

The force arising from this interaction can be written as the sum of two terms. The first is proportional to the gradient of the intensity and is called the dipole force [52]:

$$\mathbf{F}_{dipole} = -\frac{\hbar\delta}{2} \frac{\Omega}{\delta^2 + \Omega^2/2 + \Gamma^2/4} \nabla\Omega, \quad (2.29)$$

where $\delta = \omega - \omega_0$ is the frequency detuning from resonance and Γ is the linewidth of the transition. Notice that the sign of this force depends on the sign of δ ; for blue-detuned light, the force is opposite to the gradient and atoms are pushed away from the intensity maximum. For red-detuned light, the atoms are pulled towards the maximum of the intensity.

The second term is the scattering force and is in the direction of propagation of the light. It represents the radiation pressure from absorption from the beam and is given by [52]:

$$F_{scatt} = \hbar k \frac{\Gamma}{2} \frac{\Omega^2/2}{\delta^2 + \Omega^2/2 + \Gamma^2/4}. \quad (2.30)$$

This scattering term is undesirable in optical dipole potentials, where it contributes to heating due to random fluctuations in the absorption and emission (the same mechanism that leads to the Doppler limit for laser cooling).

The dipole force is conservative and hence causes no change in the temperature of the cloud. It can be written as the gradient of a potential energy, the optical dipole potential. In the highly detuned case where $|\delta| \gg \Gamma, \Omega$, the potential is given by

$$U_{dipole} \approx \frac{\hbar\Omega^2}{4\delta} = \frac{\hbar^2\Gamma^2 I}{8\delta I_{sat}}, \quad (2.31)$$

where I_{sat} is the saturation intensity of the transition. By contrast, the scattering rate in this regime is

$$R_{scatt} \approx \frac{\Gamma^3 I}{8\delta^2 I_{sat}}. \quad (2.32)$$

The dipole force $F_{dipole} \propto I/\delta$, while the scattering rate $R_{scatt} \propto I/\delta^2$. This indicates that optical dipole potentials operate best at a very high detuning $|\delta| \gg \Gamma$ where the scattering rate is low, but at high enough intensity to still trap the cold atoms.

Chapter 3

Experimental sequence and apparatus

In this chapter we describe the apparatus used to conduct experiments on 2D gases in double-wells subject to dipole potentials (in particular, pointlike disorder). We give an overview of the general sequence for producing Bose-Einstein condensates and the process of loading them into single- and double-well 2D traps. We describe the vacuum components of the apparatus, including a new design of rubidium source, and the laser system used for laser cooling and imaging. Finally we discuss the optics used for the production and manipulation of the optical dipole beam, and characterise the resulting dipole potential.

3.1 Experimental sequence

Here we describe the procedure of a single experimental sequence, up to loading a double-well MRF dressed potential. Many more details on the experimental sequence and apparatus can be found in previous theses from the research group [40–42, 53, 54].

3.1.1 Magneto-optical trap & magnetic transport

The first step of any cold-atom experiment is laser cooling to obtain an initial cold cloud of atoms. We use a pyramid magneto-optical trap (MOT) to confine and cool atoms from a background vapour of rubidium. Later in this chapter we will discuss changes made to the source of rubidium for this experiment and details of the vacuum system. The atoms are confined by the combination of a 6 cm, ~ 300 mW beam of light resonant with the cooling transition in ^{87}Rb , and a weak quadrupole magnetic field. This field is generated by current-carrying coils around the pyramid driven by

a constant current supply¹. Additional magnetic bias fields align the node of the quadrupole field with the centre of the cooling beam for maximum capture rate.

The MOT quadrupole coils are mounted on a mechanical translation stage and also serve to create a purely magnetic trap that holds and moves the atoms during the transport stage. Once a large MOT cloud has been accumulated, the cloud is rapidly compressed by increasing the detuning and decreasing the quadrupole current. After the cloud is cooled and compressed, the current is increased to confine the atoms in a quadrupole potential and the mechanical translation stage moves the coils along with their trapped atoms along a tube to the glass cell. To facilitate this, a slot through which the atoms travel is cut in the side of the pyramid mirrors.

3.1.2 Magnetic traps

Once the atoms are deposited in the glass cell, a series of magnetic traps are employed to progress to the RF-dressed potential.

- **Quadrupole trap:** The atoms are first held in a pure quadrupole potential. This magnetic field is generated by two current-carrying coils in an anti-Helmholtz configuration. These copper coils are capable of carrying up to 300 A of current. They are hollow and cooled water (12 °C when the apparatus is inactive) is pumped through them to maintain an acceptably low temperature. Typical currents are in the range of 100 – 200 A.
- **Evaporative cooling:** When deposited in the quadrupole trap, the atoms have been heated by the transport sequence, so the first stage of evaporative cooling occurs in this trap, down to a temperature of $\sim 100 \mu\text{K}$. Evaporative cooling is implemented via a weak radio-frequency probe field resonant with the Zeeman splitting between the trapped and untrapped states in the trap. By sweeping the frequency of the probe, the field becomes resonant with the transition at different positions in the trap, ejecting atoms that approach those positions. This can be used to selectively remove the most energetic atoms, which travel the farthest from the minimum of the potential.
- **Time-orbiting potential (TOP) trap:** To remedy the short lifetime in the quadrupole trap which would prohibit the production of a condensate, a rotating bias field is turned on, producing a TOP trap. The rotating bias field is produced by an

¹Xantrex XFR 20V-60A

array of coils arranged laterally around the glass cell. A further stage of evaporative cooling occurs here, down to temperatures at which a Bose-Einstein condensate (BEC) can form. If progressing to an RF-dressed potential, evaporation is cut short before this stage and we load the TAAP with a thermal cloud.

- **Time-averaged adiabatic potential (TAAP) trap:** A circularly polarised RF-dressing field is switched on, with the rotating TOP bias field still active, produced by hand-wound current-carrying coils in close proximity to the glass cell. This produces a rotating resonant ellipsoid. The atoms experience a time averaged force from this rotating potential. At the initial high TOP fields, the atoms sit at the equator of the ellipsoid. As the TOP field is ramped down and the amplitude of the ellipsoids rotation shrinks, the atoms sink to the bottom of the ellipsoid under gravity. Evaporative cooling is carried out while this ramp occurs to finish with a BEC in the shell trap.

3.1.3 RF-dressed potentials

3.1.3.1 Circularly polarised shell trap

Different polarisations of dressing field have a spatially varying coupling strength to the atoms. RF fields that are linearly polarised horizontally have a minimum of this coupling at the equator, so that atoms can transition to untrapped states in this region. When the atoms are first loaded into the TAAP, they primarily sit at the equator of the resonant ellipsoid. We therefore must load the TAAP with a circularly polarised dressing field to avoid atom loss. Each dressing coil alone produces a linearly polarised field so that a circularly polarised shell requires driving two coils at a given amplitude and with a specific relative phase. These three degrees of freedom require a high degree of stability for reliable loading and evaporative cooling.

Particular details of the generation of stable RF fields are described in previous theses within our research group [53]. The RF fields are generated using the ARTIQ [55] software and hardware package for highly stable frequencies and phases. The RF signals are produced using Urukul units² which house four direct digital synthesis (DDS) chips³. The output from these generators is extremely stable, but low in power. An impedance match is inserted into the RF signal chain to maximise power

²<https://github.com/sinara-hw/Urukul/wiki>

³*Analog Devices AD9910*

transmission to the dressing coils and stability at a specific frequency. This frequency is then the RF-dressing frequency used in experiments.

Multiple RF dressing for the double well sequence is particularly sensitive. Amplitude and phase fluctuations of less than 1% are required for consistent and even splitting of a cloud into a double well. An equivalent stability is therefore required of the impedance match. A relatively high RF-dressing frequency of ~ 7.2 MHz was chosen for its favourable trapping frequencies for investigating 2D physics and because the RF signal meets the required stability at this frequency.

3.1.3.2 Single-well 2D clouds

We usually wish to carry out experiments in the linearly polarised shell. This is because the linear shell has only one relevant degree of freedom, the amplitude (or Rabi frequency Ω), whereas the circular shell has three (two RF amplitudes and the relative phase). This makes the linear shell more stable.

The circular shell is loaded with Rabi frequency $\Omega = 2\pi \times 380$ kHz. One of the coil amplitudes is slowly ramped down and off to produce a linear shell while the quadrupole current is increased to compress the cloud vertically. The ramp is performed adiabatically over the course of 1s to avoid exciting any collective modes. The quadrupole current is then further ramped to 147 A to vertically shift the cloud to the focus of the optical dipole potential. The other coil amplitude is reduced but kept on to produce a 2D cloud. 2D clouds typically have Rabi frequencies of $\Omega = 2\pi \times 20$ - 50 kHz depending on the axial trap frequency required. This produces a single 2D cloud with typical radial and axial trap frequencies $\omega_r, \omega_z = 2\pi \times 14.5$ Hz, $2\pi \times 1100$ Hz.

For loading double-wells potentials, we perform the same initial steps but instead ramp the Rabi frequency to $\Omega = 2\pi \times 140$ kHz, producing a 3D BEC in the linear shell.

3.1.3.3 MRF double-well splitting procedure

Once the linear shell is loaded, we begin the double-well splitting sequence. For our final double-well we need a relatively high spatial separation between the wells to ensure there is no unintentional coupling between the clouds and to produce many periods of interference fringes which improves the reliability of fitting a model to the fringe pattern during analysis of experimental data (see Section 5.2). The frequencies chosen for this work were 7.08, 7.2 and 7.32 MHz.




Figure 3.1: RF dressed potential over the 5 stages of the MRF splitting sequence i) The atoms are loaded into a single well with one RF field. ii) The primary and secondary fields are switched on with amplitudes chosen to preserve the potential shape. iii) The field amplitudes are modified to deform the potential into a flat well for even loading. iv) The barrier is raised gradually (and other fields slightly modified) to split the population evenly between the wells. v) The primary and secondary Rabi frequencies are ramped down to increase the axial trapping frequency and transition to 2D clouds.

The additional RF fields can not be smoothly ramped up to their final value from zero, since this would cause the Rabi frequency of these fields to cross over the low-frequency regime where significant noise sources exist, causing atom loss. They also can not be suddenly switched to their final value, since the significant difference in the single and double well potential would cause non-adiabatic atom loss or excitations.

The splitting sequence is performed over five distinct stages. The atoms are initially prepared in a 3D single well of Rabi frequency $\Omega = 2\pi \times 140$ kHz at the central dressing frequency 7.2 MHz. The primary (7.08 MHz) and secondary (7.32 MHz) fields are then switched on suddenly, over the course of $800 \mu\text{s}$. The amplitude of these fields

(including changing the central field amplitude) is chosen so that the MRF potential closely resembles that of the single RF case, a single 3D well. This causes little atom loss or heating to occur during the sudden switch on.

The amplitude of the fields is then smoothly modified over the next 2 stages to produce a balanced double well, as shown in Fig. 3.1. The final stage of the sequence occurs after the clouds are fully split, and ramps the Rabi frequencies to their final value to fix the trap frequencies of each cloud.

3.1.3.4 RF turnoff procedure

After each sequence we turn off the trapping potential and allow the atoms to fall and expand in time-of-flight before imaging. In this work we switch off the RF fields and wait 0.7 ms before turning off the quadrupole field. This sudden switch projects the atoms into the undressed Zeeman states of the bare quadrupole field. The atoms are initially trapped in the $F = 1$ hyperfine state so that the possible Zeeman states are $m_F = -1, 0, 1$. The $m_F = 0$ state is magnetically insensitive and is not perturbed by sudden changes in the quadrupole field as it is switched off. The magnetically sensitive clouds receive a ‘kick’ from the sudden changing of this potential, which can also be unpredictable due to eddy currents in the surrounding coils. The $m_F = 0$ state falls ballistically only under the effects of gravity. The separation between different m_F states after turnoff depends on the magnetic quadrupole gradient and the duration between the RF field and quadrupole turnoff. The wait duration of 0.7 ms after the RF fields are switched off is chosen to fully separate the different m_F states during the time-of-flight expansion.

The relative population of each Zeeman state is determined in part by the potential and in part by the phase of the fields at turnoff. We therefore switch off the RF fields at a well defined phase, chosen to maximise the population in the $m_F = 0$ state. In the double-well case, we not only wish for the maximum population in the $m_F = 0$ state, but equal populations from each well, since this maximises the contrast of the interference pattern. We have found in practice that these requirements are usually satisfied at the same phase, and maximising the population also maximises the contrast.

3.2 Vacuum system

The experimental apparatus consists primarily of two main vacuum chambers; the ultra-high vacuum (UHV) glass science cell, and the comparatively high pressure




Figure 3.2: 3D model of select components of the experimental apparatus. The MOT chamber is shown on the left, containing the pyramid-shaped mirrors that reflect incident light and trap the atoms (shown here as a red cloud inside the chamber). Attached to the MOT chamber is the rubidium source chamber containing an ampoule of 5 g of pure rubidium metal, with an all-metal angle valve for controlling the conductivity to the MOT chamber. The relatively high pressure MOT side is connected to the UHV side of the vacuum by a long thin differential pumping tube and separated by a gate valve (centre). The magnetic transport mechanically translates atoms along this tube into the glass cell (right) where the atoms are trapped by magnetic fields produced by an array of current-carrying coils. Model provided by Elliot Bentine.

MOT chamber. The UHV cell is typically held at pressures of $\sim 10^{-11}$ mbar, while the MOT chamber is at pressures of $\sim 10^{-9}$ mbar. This vacuum is sustained by two ion pumps, and the pressure difference maintained by means of a differential pumping tube connecting the two chambers. This long narrow tube produces a low conductance such that atoms leaking from the high pressure chamber are transmitted at a very low rate into the UHV side, and are removed by the ion pump.

The UHV side can be sealed off from the MOT chamber by a gate valve between the differential pumping tube and the glass cell, allowing the MOT chamber vacuum to be broken without disturbing the UHV portion of the system.

3.2.1 Rubidium source

The MOT chamber is the beginning of every experimental sequence, where atoms in a thermal background vapour are cooled and trapped by the MOT beams and magnetic field. The pyramid mirrors sit in a 200 mm diameter circular chamber with a 112 mm diameter viewport above them through which the laser light enters. The chamber also has a number of connecting flanges, to which ion pumps or rubidium sources are connected.

3.2.1.1 Rubidium dispensers

For a number of years, the sources of rubidium for the experiment were alkali metal dispensers⁴, each comprising a thin wire of metallic salt which produces rubidium when heated by a current, in conjunction with a getter material to adsorb other by-products of the reaction. As described in Chapter 8, these dispensers have a few disadvantages. In particular, the existing dispensers in the apparatus were old and had begun to produce a smaller amount of rubidium relative to other gases. This reduced the final atom number and produced instability in the atom number. Since a different design of rubidium source had been successfully trialled on a new experimental apparatus, on replacing the rubidium source we switched from the dispensers to an ingot of pure rubidium, contained in an adjoining chamber.

3.2.1.2 Installation of the rubidium chamber

The design of the new rubidium chamber is the same as the final version used on the new experimental apparatus, for which it was originally designed, with the exception that it includes a zero length adaptor to fit the DN40CF parts to the DN16CF flanges on the MOT chamber. For further details on its design see Chapter 8.

Before breaking vacuum, the gate valve was sealed to prevent the UHV side coming up to atmospheric pressure, since this part of the apparatus would be harder to bake *in situ* and the very low final pressures required could take a long time to achieve.

The dispensers were removed from the chamber and the new rubidium chamber attached to a previously unused flange on the MOT chamber. Unlike on the new apparatus design, where the chamber sits horizontally with its weight supported from underneath, the position of this flange meant that the chamber was oriented vertically. This weight placed too much strain on the small DN16CF flange by itself, so a metal

⁴SAES Rb/NF/3.4/12 FT10+10

former was implemented to support the rubidium chamber from below and was fixed to the optical table.

After attaching this piece, a turbomolecular pump⁵ was attached to the MOT chamber to pump it down to 10^{-6} mbar before the ion pumps were switched on and the system then baked in situ.

3.2.1.3 Baking procedure

Care must be taken when baking to not exceed safe temperatures and prevent significant temperature gradients from building up across sensitive components. In particular, the optical viewport on the MOT is sensitive to gradients, so it was brought up to temperature at a rate slower than $2\text{ }^{\circ}\text{C}/\text{min}$, and with a temperature difference of $< 3\text{ }^{\circ}\text{C}$ across its diameter. The maximum safe temperature of this viewport was not known, but previous bakes have not exceeded $80\text{ }^{\circ}\text{C}$ [42], so the baking procedure here does not exceed this temperature across the viewport.

In contrast, a temperature gradient must be maintained overall in the system, keeping the temperature of the differential pumping tube high, to drive atoms out of the low conductance tube. Since the cooled cloud is transported through this tube for every sequence, maintaining a low pressure in the tube is crucial to producing a high and reliable atom number in the glass cell.

The system was heated by being wrapped in heater tape, driven by variable AC supplies, and subsequently wrapped in aluminium foil to insulate the system and keep the temperature high. Several rolls of heater tape were used so that different parts of the system could be heated independently, since we required a temperature gradient and larger components had a much higher heat capacity. By distributing the heater tape across the system according to the temperatures required and driving them with different supplies, a constant temperature gradient could be maintained overall while keeping gradients across individual components small.

The differential pumping tube was brought to a temperature of $90\text{ }^{\circ}\text{C}$ and the MOT chamber a temperature of $70\text{ }^{\circ}\text{C}$. This temperature was held for 36 hours before closing the rubidium valve and cracking the rubidium ampoule. The rubidium ampoule was broken by closing the rubidium valve for the first time (See Chapter 8), and the rubidium chamber was heated to $70\text{ }^{\circ}\text{C}$ overnight to drive rubidium into the MOT chamber. As we will discuss, this temperature was unnecessarily high and produced an excessively high rubidium pressure.

⁵Pfeiffer HiCube 80 Eco




Figure 3.3: a) Change in equilibrium MOT number N_{EQ} and $1/e$ loading time τ as the rubidium vapour pressure drops after closing the rubidium source valve, where the pressure is initially very high. The linear relationship between the two variables can be fitted and used to extract the partial pressure of rubidium and other background gases [56]. b) The dependence of the atom number transported to the glass cell on the partial pressure of rubidium. Counterintuitively, a very high rubidium pressure reduces the final number, due to a greater collision rate between the vapour and the cooled atoms confined in the magnetic quadrupole trap as they are transported to the glass cell.

3.2.2 MOT and transport dependence on pressure

After heating the rubidium, a large MOT was present, with an equilibrium number of $\sim 2 \times 10^9$ atoms, but a very fast $1/e$ loading rate of ~ 1 s. This implies a high rubidium vapour pressure. The rubidium valve was closed, and the resulting pressure change observed for several hours via measuring the loading rate of the cloud [56], which is shown in Fig. 3.3a. The initial high pressure reduced the final atom number transported to the glass cell, as evidenced by the increase in atom number when the pressure dropped (Fig. 3.3b).

Due to the prohibitively high pressure, and slow reduction in the pressure after closing the valve, the chamber was gently baked and connected to the turbopump to remove the excess rubidium vapour. The differential pumping tube was maintained at a higher temperature than the rest of the chamber to prevent rubidium accumulating in the tube and drive the gas towards the pump.

After 2 days of pumping further, the $1/e$ loading rate increased to 13.6 s, and the MOT number dropped to 1.5×10^8 . This is consistent with a small rubidium pressure,




Figure 3.4: a) Pressure change in MOT chamber after opening the valve to the rubidium source chamber while the pressure is initially low. The pressure increases slowly, leading to increased MOT number and reducing the loading rate. b) Transported atom number in the glass cell against the measured rubidium pressure. The transported number increases but saturates at moderate rubidium pressures, where the increased collision rate begins to reduce the transport efficiency and offsets the advantage of larger MOT clouds.

and a lifetime limited mostly by background, non-rubidium gas, approximately 2×10^{-9} mbar. The rubidium valve was then opened a very small amount and the increase in pressure measured over time. The rubidium was not heated and remained at room temperature. The change in pressure is shown in Fig. 3.4a as it increased over a period of 7 hours. The measurements also give the pressure of background, non-rubidium gas as $\sim 1 \times 10^{-9}$ mbar. Fig. 3.4b shows the atom number of the cloud transported to the glass cell as a function of the measured rubidium pressure.

It can be seen that this transported number saturates at a partial rubidium pressure of $\sim 2 \times 10^{-9}$ mbar. The rubidium valve was adjusted to maintain this pressure in equilibrium. The final optimised MOT number is $\sim 1.6 \times 10^9$ atoms with a loading rate of 10.5 s. This is a significantly higher atom number than expected for this loading rate based on the linear fit in Fig. 3.3a. This is because the atom number and loading rate also depend on the background pressure. After the additional baking, the background pressure is lower so that a greater atom number and longer loading time is achieved. Furthermore, even at constant pressure the atom number varies with properties of the MOT such as laser power or alignment which may drift or be modified between experimental sweeps.

3.3 Laser system and imaging

Frequency-stabilised laser light is an essential component of any cold atoms experiment. It is required for both initial laser cooling of the thermal vapour, and for imaging the final atomic cloud.

The laser system of our apparatus remains much as described in previous theses [41, 42, 54]. A brief overview will be given here followed by the new additions that have since been made.

3.3.1 Transitions in ^{87}Rb

The relevant part of the level structure of ^{87}Rb is shown in Fig. 3.5. There are two transitions of interest: the $F = 2$ to $F' = 3$ ‘cooling’ transition, and the $F = 1$ to $F' = 2$ ‘repumping’ transition, both of which have a wavelength of approximately 780 nm and are separated by ~ 6.5 GHz.

The cooling transition is used to cool and trap the atoms, since it is almost a closed cycling transition. However, this transition still has a finite probability to transfer atoms to the $F' = 2$ state, from which they may decay to the $F = 1$ ground state. A small amount of repumping light is then required to transfer atoms back into the excited $F' = 2$ state, from which they can decay into the $F = 2$ state visible to the cooling light.

Both cooling and repumping are also required for imaging the trapped atomic cloud. Atoms are magnetically trapped in the $F = 1$ hyperfine ground state, but the cooling cycling transition is preferable for imaging. This requires us to repump the atoms into the $F = 2$ state before imaging. This can be utilised to image only a portion of the cloud, by selectively repumping a section with the desired shape, as we discuss later.

3.3.2 Laser system

To produce light of the desired wavelength and lock it to the correct transition, we use a combination of three *Toptica DL Pro* external-cavity diode lasers. The first is a reference laser, locked to the ^{87}Rb cooling transition via modulation transfer spectroscopy (MTS). This is a Doppler-free form of spectroscopy that provides a strong error signal for locking. This error signal is used to actively stabilise the laser frequency with a high degree of accuracy and reliability.

The light from the reference laser is combined with light from each of the other diode lasers and fed into an electronic beat-note offset lock [41]. This measures the




Figure 3.5: Hyperfine structure of the ^{87}Rb D2 line [57], showing the cooling and repumping transitions. The quantity δ indicates that light is typically red-detuned from the cooling transition during the laser cooling stage. Figure taken from Ref. [42].

difference in frequency between the diode lasers and the reference frequency, which is used as an error signal to lock the diode lasers to the desired detunings for both cooling light and repumping light.

This produces frequency-stabilised light but at insufficient power for laser cooling, which requires hundreds of mW of power. The cooling light is first amplified by an injection locked slave diode laser, producing ~ 80 mW of power. A small fraction of this is split off for imaging and the rest is combined with the repumping, where the repumping makes up 5-10% of the total power. This light seeds a tapered amplifier (TA) which produces ~ 500 mW of power, delivering approximately 300 mW to the

MOT after losses incurred by fibre coupling.

3.3.3 Imaging atomic clouds

The imaging of trapped atoms is performed by absorption imaging. This method of imaging provides accurate measurements of the atomic density distribution, either *in situ* or after some time-of-flight (TOF) expansion. In absorption imaging, the cloud is first optically pumped via a repumping beam from the trapped $F = 1$ hyperfine state to the $F = 2$ state, where it is resonant with the cooling light. A collimated cooling beam is then directed onto the atoms. The atoms absorb this light and leave a ‘shadow’ of reduced intensity in the transmitted beam (I_T), which can be captured by a charge-coupled device (CCD) camera. In addition to this image, another image is taken of the incident cooling beam with no atoms present (I_I). The optical density as a function of position can then be calculated using the Beer-Lambert law:

$$OD(x, y) = \ln\left(\frac{I_I}{I_T}\right) + \frac{I_I - I_T}{I_{sat}}, \quad (3.1)$$

where $I_{sat} = 1.67 \text{ mW/cm}^2$ is the saturation intensity for circularly polarised light on this transition [57]. This provides the optical density of a single pixel, related to the total atomic number integrated over the direction of propagation of the cloud within the area of that pixel. The number within the pixel can be calculated from the theoretical scattering cross section of the cloud by $N(x, y) = OD(x, y)/\sigma$. This calculation is then calibrated more precisely using the known relation between the atom number and the onset of Bose-Einstein condensation in a harmonic trap, as described in previous theses [53, 54].

Imaging can be performed with a beam propagating either perpendicular to the direction of gravity (‘horizontal imaging’) or parallel to the direction of gravity (‘vertical imaging’). The direction of propagation of the beam removes access to details or features that exist in that direction - for example, horizontal imaging can observe the radial profile only in the x -direction, not the y -direction, and vertical imaging can not observe the axial profile of the density at all.

3.3.3.1 Imaging systems

The horizontal imaging consists of two lenses: a $f = 75 \text{ mm}$ achromatic doublet objective lens, and an $f = 200 \text{ mm}$ eyepiece lens, producing a magnification of 2.71. In combination with the Ximea CCD camera pixel size of $4.54 \mu\text{m}$, this results in an image pixel size of $1.67 \mu\text{m}$. Almost all images taken in time-of-flight are taken with

the horizontal imaging, since the atoms stay in the focal plane of the camera as they fall. In contrast, the depth of field of the vertical imaging is narrow and atoms are only focused almost *in situ* (depending on the quadrupole gradient used, which shifts the location of the atoms vertically).

The vertical imaging optics consists of an objective lens with $f = 37.6$ mm and an eyepiece lens with $f = 500$ mm. The objective lens is a high numerical aperture (NA = 0.27), custom-made four lens system designed to correct for aberrations introduced by the walls of the glass cell [42]. The camera used is an *Andor* iKon-M electron-multiplying CCD (EMCCD), providing high sensitivity and an image pixel size of $1\text{ }\mu\text{m}$.

3.3.3.2 Selective repumping

Since the atoms are trapped in the long-lived hyperfine ground state $F = 1$, they are not initially resonant with the incident cooling light. An intermediate step of optically pumping the atoms into the $F = 2$ state is required, which is achieved with repumping light resonant with the $F = 1$ to $F' = 2$ transition. Under normal operation, the repumping light is a beam much wider than the profile of the atomic cloud, and contains enough power to repump all the atoms in a short duration ($\sim 300\text{ }\mu\text{s}$). This results in almost every atom being counted in the final image. The repump duration is kept short to minimise the effect of recoil as atoms scatter the repumping light in order to prevent distortion of the imaged atomic density. In $300\text{ }\mu\text{s}$ a recoiling atom moves by $\sim 1.8\text{ }\mu\text{m}$, comparable to the size of a single imaging pixel and therefore the optical resolution of the imaging system. As a result we choose this repump duration to fully repump all atoms without affecting the final image.

However, by imaging over the entire width of the cloud in one direction, short scale variations in density are averaged over and lost. In particular, for interference experiments, this can be problematic. Firstly, a finite thermal element, which is often much wider than the condensate in TOF, reduces the contrast of the interference pattern. Most significantly however, we are interested in local phase fluctuations between two identical clouds. This produces fluctuations in the phase profile of the fringes, which are equivalent to local fluctuations in density when considering a line along constant z . When the entire cloud is repumped and imaged, these fluctuations are averaged over and can be entirely lost, both reducing the contrast of the fringes and obscuring the local phase profile.

For previous work and the results presented here, we resolved this problem by selectively repumping a small slice of the cloud [26, 32]. The repumping light was




Figure 3.6: a) Two pancake shaped clouds in the double well, vertically stacked. Randomly placed beams of 532 nm light propagate vertically upwards and produce a repulsive pointlike disorder potential within the clouds. Each cloud feels the same random potential. b) After time-of-flight, the cloud expand vertically (due to their tighter confinement in the axial direction) and overlap, forming interference fringes in the observed density. Phase fluctuations in a 2D gas cause the relative phase at each point to vary, producing ‘wavy’ fringes. A thin sheet of repumping light sent in from the side optically pumps only a selected region of the atoms which are then imaged, while the rest remain invisible to the imaging light.

shaped by reflection off of a digital micromirror device (DMD). This device consists of an array of 1024×768 , $13.7 \mu\text{m}$ mirrors which can be switched between one of two angles. This creates a digital image with binary values from which a nearly uniform incident beam can be reflected, shaping the final repumping beam. This is not the only method of beam shaping - liquid crystal spatial light modulators (SLMs) modify the phase of an incident beam to create an image in the Fourier plane and have also been used in cold atoms experiments [58, 59]. These have the advantage of continuous amplitude control in the final image, where the DMD has only one of two amplitude values (0 or 100%). They also transfer all of the transmitted power into the final image, while DMDs discard any power incident on pixels that are inactive for the desired image. However, DMDs are less expensive pieces of equipment, and we require two: one for shaping the repumping beam, and one for shaping the 532 nm optical dipole light.

For interferometry experiments, we typically shape the repumping into a thin rectangular slice of thickness $\sim 4 \mu\text{m}$. The thin slice allows measurement of local phase profiles as long as the phase is correlated over a distance greater than the slice width. This condition is fulfilled for the experiments performed in this thesis. In previous work, the optical path directed the shaped repumping light vertically upwards through the vertical imaging objective lens. This allowed for beam profiles

that vary in the horizontal directions x and y . Here, we define y as the direction parallel to the propagation of light, which is integrated over during horizontal imaging, and x as the horizontal direction perpendicular to y (that is, the radial direction still visible on the final image). However, x is a less helpful direction to vary the shape of the repumping light, since the absorption image already captures all the information contained along x .

This configuration of the repumping light has no ability to selectively repump a horizontal slice. Selectively imaging a horizontal slice of a cloud in TOF gives access to the un-averaged horizontal density profile via the vertical imaging, a potentially valuable tool in measuring density fluctuations in 2D gases. Conversely, for the dipole potential, the ability to shape the potential in both horizontal axes is much more important to produce radial confinement or columnar disorder with small spot sizes.

To resolve this problem, an additional DMD was installed. The optics serving the previously installed DMD were reconfigured for the 532 nm dipole light, directing the light vertically upwards and allowing for potential shaping in the x - y plane (See Fig. 1.6a), while the new DMD shaped the repumping light. This second DMD directed repumping light from the side of the glass cell, not below it (Fig. 1.6b). This still gives it the ability to modify the light in the y -direction, providing a thin slice for interference fringes, but also to create a horizontal slice for vertical imaging. The optical layout for the repumping DMD, as well as dipole potential optics and vertical imaging, is shown in Fig. 3.7.

3.4 Optical dipole potential

To impose an additional potential onto our magnetic traps, such as a hard wall or disorder, we utilise high intensity, far detuned light to create an optical dipole potential. A very large detuning is desirable, since the rate at which atoms scatter light is inversely proportional to the square of the detuning. Such random scattering would result in prohibitively high heating rates at low detuning. At high detuning the dipole force reduces in strength so that we also require a large amount of power to achieve an appreciable potential depth.

We produce green 532 nm light at powers up to 10 W using a Coherent Verdi V10 laser. This wavelength is far blue-detuned from the cooling transition of 780 nm and 10 W is more than sufficient power after focusing. Blue detuned laser light produces a repulsive optical dipole potential. Since we wish to produce different potential shapes for each sequence (in particular, a random disorder pattern) we need a programmable




Figure 3.7: Optical layout of the repumping, dipole potential and vertical imaging optics for the experiment. The dipole beam propagates in free space along its full path after the output of the Verdi laser. It passes through an AOM which is used for shuttering and power control; the 1st diffracted order is sent to the DMD while the other orders are sent to high power beam dumps. A system of lenses demagnifies the image of the DMD by a factor of 49. All lenses are in a confocal configuration except for the $f = 500$ mm lens in the green path, which is ~ 100 mm closer to the objective to shift the focus of the green light closer to the imaging focus. The repumping is demagnified by a factor of 3.5, so that the smallest feature produced is $4 \mu\text{m}$.

DMD for the dipole light, in addition to the one used for repumping light. The DMD allows us to produce a range of possible potentials, including hard walls to eliminate the low density edge of a harmonically trapped cloud, individual barriers to test superfluid flow, or random pointlike disorder patterns.

The optical layout for the dipole light is shown in Fig. 3.7. The beam from the Verdi output is first passed through an optical isolator to protect the Verdi from backwards-propagating reflected light. Then a $\lambda/2$ waveplate and polarising beam-splitter combination is used to divert the majority of the power into a high power beam dump during alignment.

The light is passed through an acousto-optic modulator (AOM) in single pass

configuration. This high efficiency AOM⁶ is polarisation sensitive, requiring another $\lambda/2$ waveplate prior to the input. This allows diffraction efficiencies of $\sim 90\%$ to be achieved. This AOM is used for the purpose of shuttering and modifying the power delivered to the atoms during a sequence. The zeroth order beam from the AOM, i.e. the path taken by the light when not driving the AOM, is sent to another beam block. The first order light follows a path through a telescope to expand the beam before hitting the DMD. The relatively small amount of light diffracted into the second order is also sent to a beam block.

The change of frequency of the beam by ~ 100 MHz is irrelevant at such high detuning. This AOM is driven by a 5 W radiofrequency amplifier, which amplifies a 110 MHz signal from a voltage controlled oscillator (VCO). Modifying the amplitude of this VCO signal changes the RF driving power in the AOM and the diffraction into the first order beam, and hence allows control over the final dipole potential depth.

After magnification the dipole light joins the same path as the vertical imaging beam, but propagating in the opposite direction (upwards through the objective lens and cell instead of downwards). The two beam paths are combined by use of a dichroic mirror which reflects the green dipole light but transmits the 780 nm imaging light. After exiting the upper wall of the cell, the dipole light is reflected off another dichroic mirror in the imaging path and sent to a final beam block.

3.4.1 Focusing of dipole light

We use a direct imaging system to map the DMD image onto the plane of the atoms. Two telescopes of lenses in a confocal configuration demagnify the image. The final lens in this setup is the custom four-lens vertical imaging objective. This lens is AR-coated for light of wavelength 780 nm, the imaging wavelength. As a result, some 532 nm light is lost to reflections in this lens, resulting in an overall $\sim 20\%$ loss of power. The Verdi produces far more power than is required to manipulate condensed clouds, so this loss is not problematic.

A more significant problem is that the focal length of the objective varies as a function of wavelength. Ray optics calculations predict that the focus of the green 532 nm light is $\sim 600 \mu\text{m}$ above the focus of the 780 nm imaging light. This is a large offset that will produce very substantial aberrations for either the imaging or the dipole light while attempting to image the effect of the dipole potential on the

⁶G&H I-M110-2C10B6-3-GH26 400-540nm

cloud using vertical imaging. To verify the DMD pattern on the atoms, a much closer agreement of the focal point between the two wavelengths is required.

This can be achieved by relaxing the requirement that all the lenses be confocal. The eyepiece lens confocal with the vertical imaging objective is a 500 mm focal length, 2 in diameter achromatic doublet. This eyepiece lens, unlike the objective, is not in the path of the imaging light, so it can be moved freely to adjust the focus of the dipole light. Ray optics calculations further suggested that bringing the eyepiece lens closer to the objective would bring the focal position of the dipole light downwards, closer to the focus of the imaging, without causing prohibitive optical aberrations.

The eyepiece lens was mounted on a mechanical translation stage and the shape of the dipole potential observed by its effect on the density of the atomic cloud. The position of the eyepiece lens could then be adjusted to produce the most focused pattern on the atoms.

A perfect agreement between the imaging and dipole focus could not be achieved with the existing setup. The eyepiece was brought as close to the objective as it could be without bringing it into the path shared by the imaging light, approximately 8 cm closer than the confocal position. At this position, the imaging focuses $100\text{ }\mu\text{m}$ below the dipole light.

This difference in the focal position for the two wavelengths can be turned to an advantage if we wish to use both the dipole potential and image the cloud at short TOF using the vertical imaging system in the same sequence. For example, an experiment of interest is the correlation of density fluctuations in 2D clouds after short time-of-flight expansion (see Section 9.1.2). For this case, a small difference ($\sim 100\text{ }\mu\text{m}$) in the focal position between the dipole potential and the vertical imaging is actually required.

3.4.2 Calibrating potential depth

To characterise the disorder strength and compare its effect to theory, it is very important to have an absolute estimate of the depth of the dipole potential. Though this quantity could be approximately calculated from properties of the optical system, a more accurate estimate must be measured using the atoms themselves.

We can set the relative strength of the green light using either the power output of the Verdi laser, or by modifying the driving RF amplitude of the shutter AOM. This latter method allows us to easily and quickly ramp the green light intensity over the course of a single sequence. We can therefore sweep the dipole potential from




Figure 3.8: a) Absorption image of a cloud bisected by a $13\,\mu\text{m}$ width line-shaped dipole potential in the x - y plane taken on the vertical imaging system. The blue-detuned light produces a repulsive barrier, resulting in a reduction of the atomic density in a straight vertical line. b) Density profile of the image in a) integrated along the y direction. The density is fitted with the sum of a Thomas-Fermi profile for the harmonically trapped condensate and a negative Gaussian for the density dip. A potential depth equal to the chemical potential μ would result in a density of zero at the centre of the profile. The contrast here is less than one, indicating that the potential depth is less than μ .

0 – 100% of its maximum value, but this does not provide us with an estimate of its value in units of energy.

To estimate this, we impose a repulsive dipole potential in the shape of a wide line onto a cold, 3D condensate of $N \approx 65,000$ atoms. This repulsive potential causes the density in this region to drop. Using the Thomas-Fermi approximation the density distribution under an external dipole potential $V(r, z)$ is given by

$$n(r, z) = \max\left[0, \mu - \frac{1}{2}m\omega_r^2 r^2 - \frac{1}{2}m\omega_z^2 z^2 - V(r, z)\right]. \quad (3.2)$$

In our case, although the cloud is 3D, its axial width is much shorter than the Rayleigh range of the green light, so that V is approximately constant over z . It is clear that at $r, z = 0$ the density goes to zero if $V(0, 0) = \mu$. We can therefore sweep the potential depth of the line situated at the centre of the cloud until we observe the disappearance of atoms at the centre of the trap.

We image this cloud on the vertical imaging *in situ*. Recall that the focus of the vertical imaging is $\sim 100\,\mu\text{m}$ below the focus of the green light. We perform the measurement at the green light focus. We choose a wide line, $13\,\mu\text{m}$ wide, so that

the intensity is not reduced by the blurring of green light by the finite point spread function, and so that the resulting density pattern is still relatively well resolved by the vertical imaging.

We calculate the chemical potential of the cloud using the typical expression for a 3D BEC in a harmonic trap which is given by

$$\mu = \frac{1}{2} \hbar \bar{\omega} \left(\frac{15 N a_s}{a_0} \right)^{2/5}, \quad (3.3)$$

where $\bar{\omega} = (\omega_x \omega_y \omega_z)^{1/3}$ is the geometric mean of the trap frequencies, a_s is the s-wave scattering length of ^{87}Rb , and $a_0 = \sqrt{\hbar/m\bar{\omega}}$ [47]. However this expression is derived for a cloud in a harmonic trap without the dipole potential. It is plausible that the additional potential could significantly modify the chemical potential, requiring a new analytical expression for μ . This expression would be difficult to calculate, since it would require a precise knowledge of the shape of the dipole potential (and hence the point spread function). However, 3D GPE simulations show that the presence of this dipole potential modifies the chemical potential by $\sim 5\%$. This is a small error and so we proceed using the expression for a purely harmonic trap.

Fig. 3.8 shows an example image and density profile of the bisected cloud. We fit the density profile with the sum of a Thomas-Fermi distribution and a Gaussian profile for the depleted density in the centre of the cloud:

$$n(x) = n_{max} \left(1 - \frac{x^2}{x_0^2} - C e^{-x^2/w^2} \right). \quad (3.4)$$

Although the initial image is a rectangular intensity profile, the Gaussian profile is an accurate fit due to the blurring of the green light by the finite resolution of the optical system, some blurring due to the image being slightly defocused, and the smooth drop of density due to the finite healing length of the cloud. The drop in density due to the dipole potential is expressed as a contrast C , where the maximum value $C = 1$ occurs when the density at the centre of the cloud drops to zero.

Fig. 3.9 shows the fitted contrast as a function of relative green light intensity V/V_{max} , controlled by varying the driving RF amplitude of the shutter AOM. A relative intensity $V/V_{max} = 1$ represents the maximum transmission of the AOM. The experiment is performed at constant Verdi output power of 3 W.

We fit a piecewise linear function to the data to find the crossover point of the two regions, where the contrast reaches one. We neglect the lowest green intensities in the fit, which tend to underestimate the contrast due to image defocusing and pull




Figure 3.9: Plot of the fitted contrast C from Eqn. 3.4 measuring the reduction in density due to the dipole potential. For higher relative potential depths the contrast in the density increases. A contrast of 1 indicates that the density at the centre of the dipole potential has dropped to zero, from which we infer that the potential depth is equal to the known chemical potential μ . The value of V/V_{max} for which $C = 1$ estimated by a piecewise linear to the data, shown in red.

the fitted crossover point below $V/V_{max} = 0.25$, which can be seen to be the lower bound for the crossover.

This fitted critical $V/V_{max} = 0.27$ combined with the chemical potential of the cloud $\mu/h = 620(25)$ Hz provides an estimate of the maximum dipole depth at this power setting, $V_{max}/h = 2300(100)$ Hz, where the errors represent 1 standard deviation in the measured chemical potential of the cloud.

3.4.3 Point spread function

The dipole potential optical system, like all optical systems, also has a characteristic point spread function (PSF). This function describes the intensity pattern produced by propagating a single point source through the system. The PSF sets the minimum width of any image by blurring smaller features to larger sizes. Any image produced by the light at the location of the atoms will be the true image of the DMD convolved

with this PSF.

Neglecting all optical aberrations, a perfect lens will still produce a pattern with a finite width due to diffraction, referred to as an Airy pattern. The intensity pattern measured at an angle θ to the axis produced by a circular aperture of radius a is given by

$$I(\theta) = I_0 \left(\frac{2J_1(ka \sin \theta)}{ka \sin \theta} \right)^2, \quad (3.5)$$

where I_0 is the maximum intensity at the centre of the pattern, J_1 is a Bessel function of the first kind, and $k = 2\pi/\lambda$. The radius of the first zero of this oscillating function, here referred to as r_0 , is the radius of the Airy disc, where the majority of the power in the pattern is concentrated.

A system is commonly called diffraction limited if the width of the PSF is close to this theoretical minimum radius r_0 of the Airy disc. This implies that the smallest features produced by this system are not limited by optical aberrations, but instead fundamentally by diffraction of light. Assuming a perfect circular lens, r_0 is given by

$$r_0 = 1.22 \frac{\lambda f}{d}, \quad (3.6)$$

where f is the focal length of the lens and d is the diameter of the lens. In our case, the diffraction limited radius is $r_0 = 1.4 \mu\text{m}$.

Since the majority of the power is contained within this radius, it is a reasonably good approximation to describe this central peak with a Gaussian distribution and characterise it by the width of this Gaussian. In this case, the relationship between the σ of the Gaussian distribution and the radius of the Airy disc is $r_0 = 2.9\sigma$, or $\sigma \approx 0.5 \mu\text{m}$.

However, due to aberrations in the lenses and optical system, we can expect the true PSF to have a larger width than this ideal value.

3.4.3.1 Measuring PSF of dipole light

The point spread function of the dipole optics determines the minimum size of any features of the dipole potential, such as the pointlike disorder scatterers, as well as the image sharpness of any potential patterns larger than the PSF. This makes knowledge of the PSF crucial for comparison to theory, e.g. for calculating the percolation threshold of a given disordered potential. A method of measuring the PSF, as seen by the atoms, is therefore essential.

To do this we employ the fact that patterns smaller than the Airy disc (i.e. comparable in size to the width of the PSF) will have their peak intensity reduced, since




Figure 3.10: Plot of the fitted contrast C from Eqn. 3.4 for different values of the width of the line-shaped dipole potential. The potential depth is equal to the chemical potential μ so that for wide lines the density in the optical potential region drops to zero and $C = 1$. For narrower lines, the imaging system blurs the thin line and reduces the peak intensity of the light, so that $V < \mu$ and $C < 1$. The point spread function of the light can be inferred from the line width where C first drops below 1. Shown are a piecewise linear fit (black dashed line) and an exponential fit (red dashed line). Both give different estimates of the PSF; the 3σ point of the exponential is chosen as a more conservative estimate than the linear fit and is indicated by the vertical dashed line.

the power is conserved but spread out over the blurred area. With the potential depth now calibrated, we can sweep the width of the same line-shaped potential pattern used for that measurement at constant intensity and observe when the contrast begins to drop. This will indicate that the peak of the potential has dropped so that it is less than the chemical potential and atoms fill the area between the formerly bisected clouds. With the known width of the dipole potential we can infer the width of the Airy disc and hence the PSF.

This measurement requires imaging very small patterns, so the defocusing of the vertical imaging from the green light focus makes this impossible to do accurately

with this imaging system. Instead we will use the horizontal imaging system. The horizontal imaging is hindered by distortions to the density distribution at very short time-of-flight (TOF), due to the high density of the cloud and stray magnetic fields, so we image the cloud after ~ 14 ms TOF. The green intensity is set to equal the chemical potential, such that very wide patterns ($> 15 \mu\text{m}$) completely bisect the cloud and the density in the centre is zero. We then sweep the width of the line and observe a reduction in the fitted contrast at lower widths, as shown in Fig. 3.10. The dipole potential is switched off at the beginning of TOF, so the cloud may expand slightly to fill the gap as it expands during TOF. This would reduce contrast and hence lead us to overestimate the width of the point spread function.

Similarly, the finite resolution of the horizontal imaging system may reduce the contrast of the bisected clouds. This would also lead us to overestimate the resolution of the dipole optics. This method therefore provides an upper bound to the dipole optics resolution. This is preferable, since we will use this width in Chapter 5 to calculate relevant quantities such as the magnitude of a classical trapping effects in a disordered potential. An upper bound on the dipole potential width also provides an upper bound to the magnitude of classical trapping, so that we can be confident our results are not significantly changed by these effects.

The fitted contrast shows some variation, making a clear cut off hard to distinguish. This threshold where the contrast drops is when the diameter of the Airy disc equals the width of the line. Piecewise linear and exponential fits both produce the same root mean square error, indicating fits of a similar quality, but different estimates of the cut off. Taking the piecewise linear fit, the diameter of the Airy disc is fitted as $D_{\text{Airy}} = 8.25 \mu\text{m}$. However, a number of sequences above this width show a contrast of less than 1. It is expected however that the fitting algorithm will occasionally fit a contrast of almost, but slightly less than, 1. This is due to variations in the background density and image noise.

For the exponential fit, we choose to define the cut off as the 3σ point of the cloud, or when the fitted curve reaches a contrast of approximately 95%. This gives the diameter of the Airy disc to be $D_{\text{Airy}} = 9.66 \mu\text{m}$. For this work we choose to take the more conservative of these estimates, the exponential fit, which provides an upper bound on the estimated point spread function.

This diameter of Airy disc corresponds to a Gaussian PSF with width $\sigma_{\text{PSF}} = 2.50 \mu\text{m}$. This is the width of the PSF required to reduce the peak intensity of line patterns less than $9.66 \mu\text{m}$ wide, but lines larger than this retain the same peak intensity.

With this estimate of σ_{PSF} , we can calculate the true shape of any dipole potential actually seen by the atoms, as well as calculate the real peak potential height of any pointlike disorder with width smaller than this.

Chapter 4

Theory of disordered two-dimensional Bose gases

In this chapter we present the theoretical background necessary to describe our experiments into disordered 2D Bose gases. We discuss the existence and properties of phase transitions to the superfluid state in a gas without disorder. In particular we describe the Berezinskii-Kosterlitz-Thouless (BKT) transition, an interaction-driven transition to quasi-long-range order that exists in 2D systems, which is the primary subject of study in some of our experiments. We then discuss commonly used forms of disorder in experiment and theory, and clarify the differences between them which are particularly important in 2D systems. Finally, we discuss the effect of disorder on the phase transitions in the Bose gas and the phenomenon of many-body localisation.

4.1 Phase transitions in the absence of disorder

4.1.1 Bose-Einstein condensation in a non-interacting Bose gas

The phase transition to a Bose-Einstein condensate in three dimensions (3D) is driven by the saturation of excited states, which can hold a maximum number of atoms (dependent on their temperature). Any atoms added to the system that exceed this limit can only occupy the ground state, which can support an infinite number of atoms. Equivalently, for a system with a fixed number of atoms, as the temperature is reduced the maximum number of atoms that can occupy excited states drops, until eventually it is less than the total atom number and these surplus atoms condense into the ground state.

In two dimensions (2D), this picture is different. The density of states in a homogeneous 2D system is constant in energy, and as such the number of atoms N_{ex} that

can occupy the excited states at a temperature T and chemical potential μ is given by

$$N_{ex} = \int g(\epsilon) f(\epsilon) d\epsilon = \frac{mL^2}{2\pi\hbar^2} \int_0^\infty \frac{d\epsilon}{e^{\beta(\epsilon-\mu)} - 1}, \quad (4.1)$$

where $g(\epsilon)$ is the density of states with an energy ϵ , $f(\epsilon)$ is the Bose occupation distribution, m is the particle mass, L is the system size, and $\beta = 1/k_B T$.

It is required that $\mu < \epsilon$ for all ϵ , else the occupation number becomes negative and hence unphysical. For large N , the lowest ϵ can to a good approximation be considered to be $\epsilon_{min} = 0$, and so the maximum excited state occupation can be found by setting μ to 0 in Eqn. 4.1 [60]. This gives

$$N_{ex}(T_c, \mu = 0) \propto \int_0^\infty dx \frac{1}{e^x - 1}, \quad (4.2)$$

where $x = \epsilon/k_B T_c$. This integral diverges, so that in a 2D homogeneous system no saturation-driven transition to a BEC occurs. The key difference is that in 2D, the constant density of states allows an infinite number of excitations with arbitrarily long wavelengths to exist. This effect is quite generally described for infinite homogeneous systems in dimensions less than 3 by the Mermin-Wagner theorem [61]. It is important to note however that this theorem predicts only the absence of a symmetry-breaking transition; not the absence of any transition at all. Indeed, although 2D systems do not exhibit true long range order, they can display a transition to quasi-long range order such as the Berezinskii-Kosterlitz-Thouless (BKT) transition, as we shall see.

A further exception is that the Mermin-Wagner theorem describes infinite homogeneous systems; systems of finite size *do* display a BEC transition. In 2D systems, long wavelength excitations destroy phase coherence over distances $r \rightarrow \infty$, but if the phase is coherent over a length $\ell > L$ where L is the system size, the system will undergo a BEC transition.

4.1.1.1 BECs in 2D harmonic traps

Since the absence of a BEC transition in a 2D homogeneous system arose from the form of the density of states, it could be expected that the result differs in a harmonic system. In fact, 2D harmonic systems do show a saturation-driven BEC transition.

In a 2D harmonic trap, the analogous expression for N_{ex} to Eqn. 4.2 is

$$N_{ex}(T_c, \mu = 0) = \frac{k_B^2 T_c^2}{\hbar^2 \omega_r^2} \int_0^\infty dx \frac{x}{e^x - 1} = \frac{\pi^2}{6} \left(\frac{k_B T_c}{\hbar \omega_r} \right)^2, \quad (4.3)$$

where ω_r is the geometric mean of the in-plane trap frequencies. Equivalently, a system of constant particle number N has a critical temperature T_0 given by

$$T_0 = \frac{\sqrt{6N}\hbar\omega_r}{k_B\pi}. \quad (4.4)$$

This temperature describes the critical temperature for the onset of Bose-Einstein condensation in a 2D harmonic trap in a non-interacting gas.

This result can also be obtained by applying the local density approximation (LDA) to the case of a homogeneous system, by replacing a global chemical potential μ with a local $\mu - V(r)$, where $V(r)$ is the external harmonic potential. The total atom number is then

$$N(T) = -\frac{1}{\lambda_{th}^2} \int d^2\mathbf{r} \ln(1 - e^{\beta\mu - \beta V(\mathbf{r})}), \quad (4.5)$$

where $\lambda_{th} = h/\sqrt{2\pi m k_B T}$ is the thermal de Broglie wavelength. The critical atom number in the non-interacting case is obtained by setting $\mu = 0$ which reproduces Eqn. 4.4. Note also that this expression predicts a diverging local density (and phase-space density $n\lambda_{th}^2$) at the centre of the trap, since $n(r) \propto \ln(1 - e^{-\beta V(r)})$. We can conclude that the condensation transition occurs in a harmonic trap when the central phase-space density becomes infinite. In a gas with any finite repulsive interaction, the gas can never reach an infinite central phase-space density. As a result, no BEC transition occurs in an interacting gas in a 2D harmonic trap [62]. Instead, the interaction-driven BKT transition smoothly connects to the BEC transition at low interaction strengths [63].

4.1.2 Weakly interacting Bose gases

4.1.2.1 Interactions in the quasi-2D limit

The presence of interactions fundamentally and qualitatively changes the physics at play in 2D systems. Before we examine the resulting effects, we should note that a distinction is often made between 2D and quasi-2D systems, related to the length scales of scattering and the axial width of the system. At low temperatures the interactions between two atoms is characterised only by the 3D s-wave scattering length a_s . In most cold-atom experiments, and in all cases here discussed, interactions are repulsive and $a_s > 0$, since attractive interactions generally lead to instability in the condensate.

2D gases are experimentally realised by tightly confining atoms along one direction, in a highly anisotropic trap, such that the energy scale along the axial direction

exceeds other energy scales within the cloud such as the thermal or interaction energy. However, the spatial extent of the cloud in the axial direction l_z is generally much larger than a_s . The interactions are therefore described by 3D scattering. Such a system is often referred to as quasi-2D.

Most cold-atom experiments in 2D systems typically operate in the regime $l_z \gg a_s$. In our case, the relative lengths are $l_z/a_s \sim 200$. As such it is common to refer to these quasi-2D systems as simply 2D.

In 3D, the interaction strength is characterised by $g_{3D} = 4\pi\hbar^2 a_s/m$ which appears in the Gross-Pitaevskii equation. In the quasi-2D limit, we can integrate over the tightly confined axis to obtain

$$g_{2D} = g_{3D} \int \rho(z) dz = \frac{\hbar^2}{m} \frac{\sqrt{8\pi} a_s}{l_z}. \quad (4.6)$$

By writing this as $g_{2D} = (\hbar^2/m)\tilde{g}$ we obtain a dimensionless interaction strength

$$\tilde{g} = \frac{\sqrt{8\pi} a_s}{l_z}. \quad (4.7)$$

This dimensionless interaction strength is interesting because it contains no length scale related to the in-plane 2D physics, in contrast to 3D which depends on a_s . $\tilde{g}$ completely characterises the strength of the interactions in 2D, also unlike 3D systems. To see this, consider the interaction energy

$$E_{int} = \frac{g}{2} \int n^2(\mathbf{r}) d^2\mathbf{r} = \frac{\hbar^2}{2m} \tilde{g} N n, \quad (4.8)$$

where the right hand side has been obtained in the mean-field approximation by setting $n(\mathbf{r})$ to a constant density n . Comparing this to the kinetic energy of these N particles, spread evenly over the lowest N single-particle excited states, we find

$$E_K = \frac{\pi\hbar^2}{m} N n. \quad (4.9)$$

We can consider the strongly interacting limit to be the case when $E_{int} \geq E_K$, which is given by

$$\tilde{g} = 2\pi. \quad (4.10)$$

This result is independent of the density of the gas, in contrast to the case in 3D which is characterised by the parameter na_s^3 . In most experiments $\tilde{g} \sim 0.1$, and in the experiments performed here $\tilde{g} \sim 0.072$. These are well within the weakly interacting regime.

Another related phenomenon is that the 2D gas displays scale invariance - thermodynamic quantities of the system at different μ and T are related by a simple power-law scaling, allowing a universal description of the system that depends only on the ratio $\mu/k_B T$ [64–66], for a given dimensionless interaction strength $\tilde{g}$.

4.1.2.2 Density and phase fluctuations in the 2D gas

The full Hamiltonian of the 2D gas with contact interactions can be written in terms of a classical complex field $\psi = \sqrt{n}e^{i\theta}$, where n and θ are classical fields describing the density and phase of the cloud, respectively. The velocity field of the cloud can also be expressed as [47]

$$\mathbf{v} = \frac{\hbar}{m} \nabla \theta. \quad (4.11)$$

We neglect here the quantum model of the gas, as in Ref. [67]. In this classical field model, the effective long-wavelength Hamiltonian of the system is given by [67]

$$H = \int d\mathbf{r} \left(\frac{1}{2m} |\nabla \psi|^2 + \frac{g}{2} |\psi|^4 - \mu |\psi|^2 \right), \quad (4.12)$$

where μ is the chemical potential. At finite temperature, both of the fields $\sqrt{n}$ and θ will exhibit thermal fluctuations. To begin with, we consider fluctuations in the density of the cloud. Repulsive interactions cause any density fluctuations away from a constant average density $\langle n \rangle$ to increase the overall energy, and hence the fluctuations will be reduced. Using this mean field approach, we estimate the scales at which such fluctuations are strongly suppressed by comparing the thermal and interaction energy scales and require that $k_B T \ll g \langle n \rangle$, which can be rewritten as

$$\mathcal{D} \gg \frac{2\pi}{\tilde{g}}, \quad (4.13)$$

where $\mathcal{D} = \langle n \rangle \lambda_{th}^2$ is the phase-space density (PSD) of the system. In fact, it has been numerically found that $\mathcal{D} \gg 1$ is sufficient for suppressing density fluctuations [67].

We can expect then that as the PSD of an initially thermal cloud is increased it will cross into a regime where density fluctuations in the cloud are suppressed. This cloud has a wavefunction with a constant amplitude but phase fluctuations remain. This can be considered a ‘quasicondensate’, as a condensate with a varying phase. The quasicondensate can more specifically be defined using the correlator [67]

$$Q = 2\langle |\psi^2| \rangle^2 - \langle |\psi|^4 \rangle. \quad (4.14)$$

For a thermal cloud, $Q = 0$, and for a quasicondensate the density is $n_{qc} = \sqrt{Q}$. The transition to a quasicondensate is not the same as to the superfluid, which is defined by a change in correlations of the phase from short range to quasi-long range order. In general, the quasicondensate transition occurs at a lower phase-space density than the superfluid transition.

The absence of density fluctuations means that the density contributes only a constant term to the energy of the cloud and so at low temperatures the Hamiltonian from Eqn. 4.12 can be approximated as

$$H_{eff} = \frac{\hbar^2 n_s}{2m} \int (\nabla \theta)^2 d^2 \mathbf{r}, \quad (4.15)$$

where we have made the substitution of $n \rightarrow n_s$ where n_s is the superfluid density. This is accounting for the fact that we expect only the superfluid component of the cloud to be described by this Hamiltonian. At low T , $n_s \sim n$ and at higher T , $n_s < n$. This renormalisation of the density is essentially equivalent to ignoring the short-range physics, so that H_{eff} describes the cloud at distances $r \gg \xi, \lambda_{th}$ where ξ is the healing length and λ_{th} the thermal de Broglie wavelength [60].

Specifically, we are interested in the one-body correlation function $g_1(r) = \langle \psi^*(r) \psi(0) \rangle$ for $r \gg \xi$. For a BEC with long range order, g_1 tends to some non-zero constant as $r \rightarrow \infty$.

H_{eff} is equivalent to the spin-wave effective Hamiltonian of the 2D XY model. This model exhibits a BKT transition to a state with quasi-long-range order at a critical temperature T_c . This state is characterised by the correlation function $g_1(r)$ decaying to zero for $r \rightarrow \infty$, as expected from the Mermin-Wagner theorem, but algebraically slowly.

From H_{eff} the correlation function can be derived at low T . This is because the long distance behaviour of $g_1(r)$ depends primarily on the population of long-wavelength phonon modes ($k \ll \xi^{-1}$), which H_{eff} describes well. With knowledge of the energy of these long-wavelength modes and the population given by the Bose occupation number $(\exp(\beta \hbar \omega_k) - 1)^{-1}$, it is possible to calculate the behaviour of $g_1(r)$ at long distances [60].

One finds that the correlation function scales algebraically as

$$g_1(r) = n_s \left(\frac{\xi}{r} \right)^{1/(n_s \lambda_{th}^2)}, \quad (4.16)$$

and the exponent of this algebraic decay is often defined as

$$\eta = \frac{1}{n_s \lambda_{th}^2}. \quad (4.17)$$

This algebraic decay of correlations does not contradict the Mermin-Wagner theorem, since $g_1(r) \rightarrow 0$ for large r . However, since the exponent is $1/(n_s \lambda_{th}^2) = 1/\mathcal{D}$, and for the superfluid at low temperature $\mathcal{D} \gg 1$, the correlations still decay very slowly. In the thermal phase, by contrast, $g_1(r)$ decays exponentially,

$$g_1(r) \propto e^{-r/l}. \quad (4.18)$$

4.1.3 BKT transition

Although no spontaneous symmetry breaking transition occurs in 2D, given that the correlation function is seen to change from algebraically decaying at low temperatures to exponentially decaying at high temperatures we expect that there is a critical temperature at which a crossover occurs. Furthermore, since in the superfluid regime the effective Hamiltonian is the same as that of the 2D XY model, which displays a BKT transition, we expect a similar transition to exist. However, in the vicinity of the critical point, short range physics becomes relevant and perturbative approaches can not describe the microscopic physics or mechanism that causes the BKT transition, nor predict the exact critical point. Although it is conceivable that long wavelength modes (phonons) could cause a gradual transition from algebraic to exponentially decaying correlations, the BKT transition is in fact very sudden. The essential element that is absent from the above analysis is the presence of vortices.

In a vortex, the superfluid density vanishes and the phase in a path around the vortex changes by an integer multiple of 2π . From Eqn. 4.11 we see this can be interpreted as rotational flow of the gas around the vortex, with a defect of undefined phase at the centre. The direction of flow around the centre of rotation determines the sign of the vortex. The superfluid density vanishes only at the centre of the vortex, but a sharp discontinuity in the wavefunction is prohibited since it would correspond to an extremely high kinetic energy. Instead, the radial size of the vortex is set by the healing length $\xi \sim 1/\sqrt{\mu}$. Therefore the existence of vortices is consistent with density fluctuations being suppressed on a length scale greater than ξ , as this result was derived from the long-wavelength Hamiltonian in Eqn. 4.12.

If vortices exist in bound pairs of vortex and anti-vortex, then the pair creates no net circulation along paths that enclose the pair. They can affect the phase θ on a scale of order ξ , but do not affect phase fluctuations on longer length scales, which are caused solely by phonons. Free vortices meanwhile, not paired with a nearby anti-vortex, have a net circulation. This creates a sudden jump in the value of the phase, and the cloud is correlated only over a length scale less than the distance between

free vortices. The proliferation of free vortices at a critical phase-space density leads to the sudden change in the form of the correlation function g_1 .

A full description of the thermodynamics of the transition can be found in Refs. [68, 69], but an illustrative overview can be quickly made. The velocity field of the cloud at a distance r from a vortex, caused by the winding of the phase by 2π , is $\hbar/(mr)$. As such the kinetic energy associated with a vortex placed at the origin of a system of radius $R \gg \xi$ is given by

$$E = \int_{\xi}^R \frac{1}{2} n_s \left(\frac{\hbar}{mr} \right)^2 = \frac{\hbar^2 \pi}{m} n_s \ln \left(\frac{R}{\xi} \right). \quad (4.19)$$

The entropy associated with this vortex is given by the number of positions a vortex of radius ξ can be placed in the system of radius R :

$$S = k_B \ln \left(\frac{\pi R^2}{\pi \xi^2} \right) = 2k_B \ln \left(\frac{R}{\xi} \right). \quad (4.20)$$

Then we can combine these expressions to calculate the free energy $F = E - TS$ associated with an unbound vortex:

$$\beta F = \frac{1}{2} (n_s \lambda_{th}^2 - 4) \ln \left(\frac{R}{\xi} \right). \quad (4.21)$$

Recall that the sign of the free energy F determines whether a process can occur spontaneously. For $n_s \lambda_{th}^2 > 4$, $F > 0$ and so the creation of a new vortex can not occur spontaneously. Therefore in the superfluid regime, free vortices do not proliferate of their own accord and exist only in bound pairs, leaving the correlation length unaffected at long length scales. For $n_s \lambda_{th}^2 < 4$, the normal phase of the cloud, $F < 0$ and so vortices can proliferate freely and spontaneously. In addition, the creation of a new vortex reduces n_s , increasing the likelihood of additional free vortices and rapidly destroying the superfluid phase. For $R \gg \xi$, the magnitude of F is very large and so the point $n_s \lambda_{th}^2 = 4$ represents a sudden change between these two regimes.

The sharpness of this transition tells us that at the BKT critical point, the superfluid density suddenly jumps from 0 to $4/\lambda_{th}^2$. This also implies that the critical exponent of the algebraically decaying correlations, defined in Eqn. 4.17, is $\eta = 1/4$. Despite the simplicity of this calculation, a more detailed theory produces exactly the same result for the superfluid jump [68], and a rapid increase in the superfluid density has been experimentally observed [70]. This result also does not depend on the interaction strength $\tilde{g}$, even though this transition is driven by interactions. Importantly however, we note that in finite-size systems this sudden jump in the superfluid density

is broadened into a smooth increase [71, 72], and the width of this crossover region does depend on the interaction strength and atom number.

It may appear that this also predicts the critical phase-space density at which the jump occurs, but this is not the case. It tells us only that at the critical point, the superfluid phase-space density suddenly increases to 4. To calculate a critical temperature T_c , we would need a critical phase-space density for all components of the cloud.

The critical phase-space density for weak interactions was first predicted in 2001 by Prokofiev *et al.* [67] by use of classical field simulations. They obtained a critical phase-space density of

$$D_c = \ln\left(\frac{380}{\tilde{g}}\right). \quad (4.22)$$

For our experimental parameters, $D_c \approx 8.5$, which is easily accessible in cold-atom experiments. The authors also note that at the critical point

$$\frac{n_{qc}}{n_{tot}} = \frac{7.16}{\ln(380/\tilde{g})} \sim 1. \quad (4.23)$$

So the quasicondensate density is of the order of (but slightly less than) the total density at the transition point, and in fact the temperature at which a quasicondensate appears is always greater than the BKT critical temperature.

4.2 Disorder potentials

The specific form of disorder applied to a system can have a significant effect on the observed physics. Although in general the qualitative effects of disorder on systems depend on statistical properties of the potential and not on its microscopic configuration, certain practical differences can arise. It is possible that even in principle many-body localisation effects can only be seen for very short range, impurity-like disorder patterns [20]. Experimentally, classical trapping effects from strong disorder can overwhelm a potentially subtle indication of localisation arising from quantum interference, and quantitatively the importance of these effects can vary dramatically between disorder implementations [73]. Comparison to theoretical results requires a careful examination of properties of the disorder such as strength and correlation length, whose precise definitions also depend on the exact form of the potential. Here we discuss the forms of disorder most commonly considered experimentally and theoretically in quantum gases.

4.2.1 Forms of disorder

4.2.1.1 Speckle disorder

Perhaps the most commonly studied form of disorder in cold-atom experiments is that of an optical speckle pattern. Laser light reflected off the surface of a typical object and projected onto a plane is noticeable for having a distinctive, grainy appearance. This granular pattern is caused by the microscopic roughness of most objects on the scale of an optical wavelength λ . This roughness results in dramatically different path lengths for light rays propagating from different positions to the same point, and these rays interfere. Objects are usually very rough on the scale of λ , so that the full range of interference occurs from highly destructive to highly constructive, and a random intensity within that range can be observed, varying in 3D space. This intensity pattern $I(x, y, z)$ produces a potential landscape for atoms via the optical dipole force.

Experimentally, speckle patterns are not typically produced by reflection, but by transmission through an optical diffuser. These diffusive plates have a rough internal structure that applies a random phase to the propagating Fourier modes. These diffusive plates scatter or diffract light passing through them and the light rays interfere as they overlap. Additionally, the pattern is usually produced in an imaging geometry, where a beam may be focused by a converging lens and then passed through the plate to scatter the light.

Consider a diffusive plate subject to an incident blue-detuned laser beam. We define the size of the illuminated region of the plate D and consider this region to be comprised of many small light sources of size w . We assume these sources are independent and the light from each source has a random additional phase in the range $[-\pi, \pi]$. This is the assumption of a surface or internal structure that is rough compared to the wavelength λ . From each region the light diverges with an angular width $\sim \lambda/w$.

The amplitude of the field at each point is the sum of the contributions from these sources which each have a random phase. We usually consider ‘fully mixed’ or ‘fully developed’ speckle patterns, where the distance z from the diffuser to the observation plane is large enough that all of the source regions contribute to the interference, i.e. $z > Dw/\lambda$. This sum is equivalent to a random walk in the summed complex amplitude [74]. The value of the field, and hence the intensity I , is a random function of position, both transversally (in the $x - y$ plane) and longitudinally (z direction).




Figure 4.1: Example of a simulated disorder potential created by the speckle pattern of interfering laser light. The potential, proportional to the intensity of the light, is a random variable with an exponential distribution of mean $\langle V \rangle$. The final intensity pattern is composed of the sum of many Fourier modes with random phases, imposed by the roughness of the reflecting or transmitting object. The size of a grain of a speckle is determined by the wavelength of the light and the size of the rough object (see text).

One can realise different specific configurations of disorder by changing the distance z . Properties of this potential are therefore statistical in nature.

First we wish to know the distribution of the intensity, since this will determine the strength of the disorder. Being the sum of many random contributions, the real and imaginary parts of the field are given by Gaussian distributions, as described by the central limit theorem. As a result the intensity probability density distribution is exponentially distributed, and so is the potential:

$$P(V) = \frac{1}{\langle V \rangle} e^{-V/\langle V \rangle}, \quad (4.24)$$

where $\langle V \rangle$ is the mean value of the distribution. For an exponential distribution, the standard deviation of the potential is equal to the mean, $\sigma_V = \langle V \rangle$. This quantity gives a typical estimate of the amplitude of the modulation in the potential. It is clear

from Eqn. 4.24 that potentials of zero are possible, and arbitrarily high potentials are also possible with increasing rarity; 99% of the speckle peaks have an intensity less than or equal to $5\langle V \rangle$. Note that this distribution is not symmetric around its mean.

The average potential is proportional to the average intensity $\langle V \rangle \propto I_0$, which is determined by the intensity of the incident laser beam, the angle of divergence of the diffuser, and the distance from the diffuser to the observation plane z . As long as the condition $z > Dw/\lambda$ is fulfilled, z can be changed to modify the average intensity without changing the transverse correlations (the size of a single speckle grain).

The other crucial parameter of a disorder pattern is its correlation length. This is related to the width of a single grain of the speckle pattern and is given by

$$C_I(\delta\mathbf{r}) = \langle I(\mathbf{r})I(\mathbf{r} + \delta\mathbf{r}). \rangle \quad (4.25)$$

We may define the width of a speckle grain as the first zero of this function. We take the experimentally relevant case where we observe the speckle pattern at the focal plane of the converging lens with focal length f , with the diffuser a distance z_{diff} from this plane. Then in the simplest case of a uniform square intensity profile incident on the diffuser with width D , again in the far field, the transverse grain size is given by

$$\Delta x = \frac{\lambda z_{diff}}{D}. \quad (4.26)$$

Notice that if z_{diff} is changed, the size of the illuminated region D also changes and the width of the grain is constant. It is set by the half-angle $\alpha = D_{lens}/2f$ where D_{lens} is the width of the illuminated region on the lens. This shows that the transverse width of a speckle grain is in fact the diffraction limit of the imaging system. An anisotropic grain size can be produced by rectangular intensity profiles, rather than square (or radially symmetric) ones.

Alternatively the correlation length l_c is sometimes defined as the full width at half maximum (FWHM) of the peak of the correlation function [75]. In 1D the correlation function is particularly simple and is given by

$$C(x) = \left(\frac{\sin(\pi\Lambda x)}{\pi\Lambda x} \right)^2, \quad (4.27)$$

where Λ is a cutoff wavevector introduced by the finite numerical aperture of the optical system and the correlation length $l_c = 0.88/\Lambda$. This function can be approximated by a Gaussian profile, a particularly useful form of correlation for theoretical study.

The longitudinal correlation length, the distance z over which the speckle pattern remains similar, is given by

$$\Delta z \approx 7.6\lambda \frac{z_{diff}}{D^2} = \frac{7.6\Delta x^2}{\lambda}. \quad (4.28)$$

Speckle disorder can act as a disorder pattern in 1D, 2D or 3D systems by tuning these relative correlation lengths. If Δz is much greater than the axial width of the system l_z , the disorder pattern is essentially columnar and varies only in the transverse plane.

Speckle disorder is widely studied in theoretical investigations of disordered systems [4, 75–81] and is perhaps currently the most common experimental implementation of disorder in cold atoms systems [82–88]. This is due to its relative ease of production and simple parameterisation and therefore straightforward comparison between theory and experiment. Speckle disorder does have drawbacks, including that it is less tunable than disorder patterns created with a controllable device such as a DMD. It also poses more fundamental physical problems in two dimensions, as will be discussed in Section 4.2.2 when we discuss percolation theory.

4.2.1.2 Pointlike disorder

Pointlike disorder patterns are, like speckle patterns, generated by the optical dipole potential of incident blue-detuned laser light. Pointlike disorder is the sum of many independent, randomly placed scattering points in the space. A typical model is to consider each scatterer a Gaussian barrier of the same height, so that the disorder potential is given by:

$$V(\mathbf{r}) = \sum_i V_0 e^{-|\mathbf{r}-\mathbf{r}_i|^2/2\sigma^2}, \quad (4.29)$$

where V_0 is the peak disorder strength, $\mathbf{r}_i$ is the location of the i th scatterer, and σ is the root mean square width of a single barrier. It is also common to see the width of the barrier defined in terms of a width $w = \sqrt{2}\sigma$.

This potential is distinct from the speckle disorder in that it contains many barriers of similar or the same height, and it has an additional degree of freedom. In addition to changing the width σ of a single scatterer, comparable to the grain size of a speckle pattern, pointlike potentials can tune the density of scatterers n , introducing a second length scale $\sim n^{-1/2}$.

For $\sigma \ll \xi$ where ξ is the healing length of the cloud, pointlike disorder approaches a true impurity limit. In this context, an impurity disorder is equivalent to the kind imposed by randomly placed atoms of a different species to the surrounding gas.




Figure 4.2: Example of a simulated pointlike disorder potential. The pattern is the sum of $N = 500$ randomly placed Gaussian spots with width $\sigma/L = 0.01$ where L is the size of the system. The potential is shown in units of the average potential $\langle V \rangle$ for direct comparison to the case of speckle disorder. The peak potential $V_0 \sim 11 \langle V \rangle$.

These impurity atoms provide a typically very short range interaction via inter-atomic collisions. In the limit of small σ , this form of disorder can be expressed:

$$V(\mathbf{r}) = \sum_i g \delta(\mathbf{r} - \mathbf{r}_i), \quad (4.30)$$

where g is a coupling constant related to V_0 and σ . For small σ and a barrier potential of the same form as Eqn. 4.29, $g = V_0 \sigma \sqrt{2\pi}$. Of course, genuine impurity disorders can also be produced, particularly in lattice systems where rare, random lattice sites are occupied with a scattering atom.

Pointlike disorder is comparatively hard to produce in 3D, but in 1D and 2D is straightforwardly generated by specialist optical elements. This includes the use of a digital micromirror device (DMD) to shape the beam (which is the method employed in this thesis), or a liquid crystal spatial light modulator (SLM) to modify the phase of a beam such as those used to make optical boxes for homogeneous traps [59]. These devices allow for a highly tunable pointlike disorder pattern, easily randomised

between experimental sequences with a fine degree of control over the relevant parameters. Perhaps most importantly for two dimensional systems, classical trapping effects can be made much smaller than speckle disorder, permitting observation of localisation effects that are otherwise masked [89].

4.2.2 Percolation theory

When attempting to observe the localisation of atoms, particularly in transport experiments which measure the reduction in diffusion, it is essential to consider whether the disorder potential is merely classically trapping atoms as opposed to localising them via quantum interference.

To see how this occurs, consider that a random speckle pattern can create ‘valleys’ of low potential energy surrounded entirely by ‘peaks’ of high potential energy. If the path out of this valley must cross a potential barrier of at least energy E , then no atoms with an energy less than E can leave the valley (neglecting the effect of quantum tunnelling, which exponentially decreases as the threshold energy to escape increases). These atoms are classically trapped inside this region. If this energy exceeded the global chemical potential, $E > \mu$, then all atoms fall into this well and are unable to escape. This would classically trap atoms in a random low potential region, rather than localise atoms by the effects of quantum interference.

For this last situation to occur, it is necessary not only that the average potential depth $\langle V \rangle$ exceed the chemical potential μ , but that the potential completely cut off any paths around the potential with an energy less than μ . In other words, as long as the longest fully connected region the atoms can occupy is larger than the system size, the atoms can percolate through the disorder pattern freely. Equivalently, at low density of scatterers, we can consider a connected chain of scattering potentials. At a low density of scatterers, this chain is much shorter than the system size.

A fundamental result of percolation theory states that (for constant barrier strength) as the density of barriers in the system increases, the length of the connected region undergoes a sudden transition at a critical density [90, 91]. At this threshold, the length of connected potential chain increases dramatically and covers the entire system. Again, this is equivalent to the region in which the atoms travel freely becoming suddenly much smaller, corresponding to the cloud becoming classically trapped. This rapid transition (which in our case depends on the correlation length and the height of the potential) is referred to as the percolation threshold.

It is more experimentally relevant to restate this in terms of a fixed potential configuration and chemical potential with varying disorder strength V . There exists




Figure 4.3: Illustrations of the region over which atoms of an energy E can percolate, indicated by the uniformly teal coloured region, in the same 2D speckle potential as shown in Fig. 4.1. a) Region over which an atom with $E = 0.3 \langle V \rangle$ can travel. All points with a potential $V < E$ are shaded teal. In this case, $E < E_{perc} = 0.52 \langle V \rangle$ [76, 92], where E_{perc} is the critical energy at which an atom can traverse the full system size. The teal shaded regions are disconnected and do not cross the full system. b) For $E = 0.75 \langle V \rangle > E_{perc}$ the atom can traverse the entire system - the shaded area is fully connected across all space.

an energy $E_{perc}(V)$ where for all atoms with energy $E > E_{perc}(V)$, the atoms can freely move around the system. Atoms below this energy are locally trapped in connected regions of the potential.

This classical trapping effect is capable of masking or completely dominating localisation effects. As we will later discuss, localised systems also often have a mobility edge, an energy below which atoms are confined by quantum interference from random scatterers. If this edge is below the percolation threshold for a given disordered system, then the effect of localisation may be impossible to observe.

The percolation threshold depends very strongly on both the dimension of the system and the form of disorder. It is typically calculated numerically, and may be reported either as an absolute energy threshold or as a fraction α of the average disorder energy. In three dimensions, the percolation threshold of a blue-detuned speckle potential is $\sim 4 \times 10^{-4} \langle V \rangle$ [75]. This is very low, implying the suitability of speckle patterns for observing possible localisation effects in three dimensions. This can be thought of as resulting from the many possible paths an atom can take in 3D,




Figure 4.4: Illustrations of the region over which atoms of an energy E can percolate (teal shaded region) in the pointlike potential shown in Fig. 4.2. This potential has a filling factor $\rho = 0.1$ (see text) and a resulting percolation threshold of $E_{perc} \approx 0.2 \langle V \rangle$ [73]. a) Region over which an atom with $E = 0.05 \langle V \rangle < E_{perc}$, Shaded regions are disconnected and do not cross the full system. b) For $E = 0.3 \langle V \rangle > E_{perc}$ the shaded area accessible to the atom forms a fully connected area across the entire system. This is the same energy as Fig. 4.3a but in the case of pointlike disorder the atom can still cross the full system, unlike speckle disorder.

to wind its way around any barrier.

In two dimensions however, the possible paths are more geometrically constrained and the percolation threshold for speckle disorder is much higher, $E_{perc} \approx 0.52 \langle V \rangle$ [76, 92]. This implies that any atoms below half the average disorder strength are classically trapped. This limits experiments to low disorder strengths so that only a small fraction of atoms are trapped in this way. At low disorder strengths, the localisation length - the size of the state localised by interference from the disorder - is often larger than the system size, making the quantum effects of interest impossible to observe.

Pointlike disorder has been proposed and experimentally verified as a method of avoiding this problem in 2D [73, 89]. An important distinction is that because speckle disorder contains only one characteristic length scale, its percolation threshold is a fixed fraction of the average potential energy. Pointlike disorder has two length scales, the width of a scatterer σ and the average distance between points, $n_p^{-1/2}$ where n_p is the density of the scatterers. This allows pointlike disorder to tune the percolation




Figure 4.5: Percolation threshold as a fraction of average potential energy for pointlike disorder and speckle disorder. Pointlike disorder is plotted as a function of the filling factor ρ (blue solid line). Speckle disorder is constant as its correlation length is varied [73] so is shown as a constant threshold for comparison to the pointlike case (red dashed line). Plot based on data from Ref. [73].

fraction α .

The percolation threshold for pointlike disorder can be parameterised in terms of a filling factor $\rho = 2n_p\sigma^2$ (note the factor 2 here comes from the conversion between σ and $w = \sqrt{2}\sigma$ used in [73]). For low ρ , where localisation effects can still be relevant, the percolation threshold is much reduced compared to the 2D speckle disorder, and can even be less than 3D speckle patterns. A plot of the percolation fraction α as a function of ρ is shown in Fig. 4.5, using data from Ref. [73].

The low percolation threshold of pointlike disorder permits the use of higher disorder strengths without trapping atoms classically and hence dramatically reduces the expected localisation length in two dimensions. Pointlike disorder was used to observe Anderson localisation in two dimensional Bose gases for the first time [89]. It is the form of disorder we choose to experimentally implement in this work.

4.3 Many-body localisation in Bose gases

Many-body localisation (MBL) is a topic that has generated a great deal of theoretical and experimental interest. Here we will discuss what a many-body localised state is, what distinguishes these states from both thermal and other non-thermalising states, and some predicted effects of disorder and localisation on the superfluid phase transition in Bose gases. Further detail on the theoretical and experimental progress on MBL can be found in Refs. [93–95].

4.3.1 Thermalisation of quantum systems

Statistical mechanics that describes isolated classical systems tending towards a thermal state is founded upon the assumption of ergodicity. This is the condition that over an infinite period of time, an atom will access every available microstate with equal probability. For a large collection of atoms this assumption also implies that the infinite-time average of a single degree of freedom is equivalent to an ensemble average at a single time in the thermodynamic limit. Interactions in the system mediate exchange of information and energy between degrees of freedom and the system tends towards a thermal state where the system’s observables are distributed according to the microcanonical ensemble.

This picture is more complicated in a number of respects in quantum systems. There is in fact no similarly simple definition for ergodicity in a quantum context. A simple argument to illustrate this is made here following the same lines given in Ref. [93]. Consider a general isolated many-body quantum system with Hamiltonian $\hat{H}$. An initial state $|\psi(0)\rangle$ can be expanded in terms of the many-body eigenstates $|\alpha\rangle$:

$$|\psi(0)\rangle = \sum_{\alpha} A_{\alpha} |\alpha\rangle. \quad (4.31)$$

The time evolution of this state is given simply by the unitary evolution, where each eigenstate in the linear combination acquires a phase determined by its eigenenergy E_{α} :

$$|\psi(t)\rangle = e^{-i\hat{H}t} |\psi(0)\rangle = \sum_{\alpha} A_{\alpha} e^{-iE_{\alpha}t} |\alpha\rangle. \quad (4.32)$$

The probability of measuring the system in any of its eigenstates, $p_{\alpha} = |A_{\alpha}|^2$, is determined entirely by the initial state and is constant over time. The system then does not explore different states over time, in contrast to the classical case. It is also difficult to describe exploring states in phase space, since the position and momentum of the state can not be simultaneously known. This illustrates the essential difficulty

of defining ergodicity in a quantum system, and hence understanding the process of thermalisation in quantum systems. Reversible classical systems can irreversibly tend towards a thermal state, where most of the initial state's information is lost, due to the (correct) assumption of ergodicity. What then is the analogous mechanism for quantum systems?

It is most useful to discuss thermalisation in terms of a systems' observables. To say that a quantum system has thermalised means that, having started in some general initial state, the observables at long times have time-averaged values given by the microcanonical ensemble. The infinite-time average of an observable given by an operator $\hat{O}$ is given by

$$\langle \hat{O} \rangle_{\infty} = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \langle \psi(t) | \hat{O} | \psi(t) \rangle = \sum_{\alpha} p_{\alpha} \langle \alpha | \hat{O} | \alpha \rangle. \quad (4.33)$$

If the system has thermalised we require this expectation value, for any arbitrary local observable, to agree with microcanonical ensemble. That is to say that the system thermalises, assuming energy and particle number are conserved, if and only if

$$\langle \hat{O} \rangle_{\infty} = \text{Tr} \left(\hat{O} \frac{e^{-\beta(\hat{H} - \mu \hat{N})}}{Z} \right), \quad (4.34)$$

where the inverse temperature $\beta = 1/k_B T$ and the chemical potential μ are fixed by the constraint that energy and particle number are conserved, respectively (i.e, for $\hat{O} = \hat{H}$ and $\hat{O} = \hat{N}$).

Note that in Eqn. 4.33 the time average of an observable relies only on the initial p_{α} because in the average over long times the oscillating cross-terms integrate to zero. One can think of the progress towards a thermal state being the dephasing of these initial terms over time.

The question then is what is the condition for Eqn. 4.34 to hold? A general solution to this problem for interacting systems is not trivial, and indeed has not been found, but one possible solution that could be hypothesised is that the expectation value of $\hat{O}$ in each individual eigenstate already agree with the microcanonical ensemble. This postulated solution is the well-known and successful Eigenstate Thermalisation Hypothesis (ETH) [96, 97].

This solution implies that any subsystem of a full many-body state would experience the rest of the system as a thermal bath and explore many available states. This resembles ergodicity in classical systems, and hence the term is often used to describe this situation in a quantum context. In addition, ETH implies not only that the average value of an observable is thermal, but that the values of the observable

at long times are approximately equal to that average. This distinction is sometimes referred to as ‘weak’ versus ‘strong’ thermalisation, respectively.

The ETH has been numerically studied and verified in many systems including systems of hard-core bosons [98–100], spinless fermions in 1D [101] and the 2D transverse field Ising model [102]. All known thermalising systems obey the ETH, but since it is currently unproven whether it is a necessary condition for thermalisation it remains a hypothesis.

As powerful as the ETH is, systems can be found that do not obey it. Simple examples are integrable systems such as 1D lattices of non-interacting spinless fermions. The case that we are more concerned with here is systems that escape thermalisation by the addition of disorder. These systems may display a dynamical phase transition to a many-body localised state.

4.3.2 Anderson localisation

In 1958 Anderson proposed what he considered to be the simplest possible model that has an infinite number of degrees of freedom and yet does not approach thermal equilibrium [10]. While not an example of many-body localisation (MBL), it provides the earliest suggested example of a system that avoids thermalising due to the addition of disorder, and captures some related ideas that persist into MBL states. Below is a slightly simplified version of the scenario Anderson presented and the results he obtained, based on Ref. [21].

We consider a lattice model in the limit of low densities and very short range weak interactions, essentially neglecting their effect. We initialise a wavefunction localised to a specific lattice site $i = 0$ at a time $t = 0$. We are interested in the rate of diffusion away from this initial site in the case of a regular lattice, or a disordered lattice. Atoms can hop from site to site with amplitude g , and each site i is randomly occupied with an energy ϵ_i , uniformly distributed with width W . The Anderson Hamiltonian is then:

$$\hat{H} = -g \sum_{\langle ij \rangle} \left(c_i^\dagger c_j + c_j^\dagger c_i \right) + \sum_i \epsilon_i c_i^\dagger c_i. \quad (4.35)$$

In the case of no disorder, the width $W = 0$ and so all the site energies $\epsilon_i = 0$. In this case the Hamiltonian is translationally invariant and so the eigenstates are plane waves. Starting with a wavefunction entirely centred on the site $i = 0$, $\psi(t = 0, i) = \delta_{i,0}$. In the thermodynamic limit as the size of the system $L \rightarrow \infty$, the solution in 1D

converges to

$$\psi(t, j) = \int_{k \in [0, 2\pi]} \frac{dk}{2\pi} e^{ikj + 2itg \cos(k)} = iJ_j(2tg), \quad (4.36)$$

where $J_n(x)$ is a Bessel function of the first kind. The system is spherically symmetric and so in higher dimensions the solution is the product of the 1D solutions in each direction. The wavefunction remains centred on the initial site $j = 0$, while the mean square displacement is then given by:

$$\Delta x^2(t) = \sum_j j^2 |\psi(t, j)|^2 = 2g^2 t^2. \quad (4.37)$$

That is, the mean square displacement grows ballistically with time.

By contrast, in the limit where the disorder is much stronger than the hopping element, the wavepacket remains perfectly localised on its initial site:

$$\psi_i(j) = \delta_{i,j}, \quad (4.38)$$

and the mean square displacement $\Delta x^2(t) = 0$.

In the intermediate case, an analogy with classical motion moderated by elastic scattering would suggest a diffusive model where

$$\Delta x^2(t) = Dt, \quad (4.39)$$

where D is the diffusion coefficient, for $t \gg \tau$ where τ is a characteristic scattering time. However, Anderson predicted that for $W > W_c$ where W_c is some finite disorder strength, the diffusion coefficient D vanishes entirely. This is due to the destructive interference in the summation of all possible paths around the disordered lattice site. As a result the mean square displacement initially grows but converges to a finite value. The particle, without any additional external potential, has become localised within a finite region of space, characterised by a length scale referred to as the localisation length. As can be inferred from the limit of $W \gg g$, the localisation length decreases for increasing disorder strength.

Here we can introduce some intuitive terminology. Thermalising systems can exchange information and energy, so they are necessarily conducting to some extent. In contrast, the Anderson model exhibits a vanishing transport of atoms and energy, and such a state can therefore be referred to as insulating. Many MBL states are also referred to as insulating by the same analogy.

Anderson localisation has been proved to exist in any number of dimensions for a sufficiently strong, but finite, W_c . In fact, in 1D the critical disorder strength is

$W_c = 0$ - that is, the state is localised for even arbitrarily small disorder [21]. In 2D systems it is suspected, based on numerical investigations, that the critical disorder strength is again zero, but this has not been rigorously proved. In dimensions of 3 or greater, the critical disorder strength is finite but dependent on the energy of the atoms. This introduces a ‘mobility edge’ - an energy threshold below which atoms are localised by disorder of a given strength, but above which the atoms remain delocalised and propagate diffusively. In theory this mobility edge does not exist in 1D [103, 104], but experiments have observed an effective mobility edge due to the finite power spectrum of the disorder potential used [85].

Anderson localisation has been observed experimentally in 1D non-interacting [105] and weakly interacting [85] condensates, 2D systems [89, 106] and 3D BECs [107].

4.3.3 Properties of MBL systems

The addition of interactions between many particles in the presence of disorder takes us from Anderson localisation, in the limit of non-interacting particles, to many-body localisation. MBL features many similar properties to Anderson localisation, in that it has zero conductivity and hence is an insulating, non-thermalised phase. MBL states also display a low amount of entanglement. Thermal states have an entanglement entropy proportional to the volume of the system, since in thermal states the entanglement entropy is equal to the thermodynamic entropy which is an extensive property. In contrast, MBL states have an entropy that scales with the boundary of the system, referred to as area-law scaling [108, 109]. Dynamically, the entanglement in an MBL system grows logarithmically, where in thermal systems the entanglement propagates ballistically. This is often considered a crucial distinction between the single-particle Anderson localisation and MBL, since Anderson localisation can develop only specific kinds of entanglement such as the number of particles, and not, for example, entanglement in the configuration of those particles. This also provides positive evidence for a phase that is otherwise often defined only by the absence of thermalisation. This condition is commonly used in numerical studies as evidence of MBL [108]. Experimentally measuring this entanglement is difficult, but it has been achieved in a disordered 1D chain of atoms [110].

A powerful method of describing MBL systems is via the local integrals of motion (LIOMs) [111, 112]. Distant degrees of freedom in a fully MBL system are only weakly connected, so that the behaviour of a local observable is affected only by its nearby degrees of freedom. This gives rise to the emergent integrability of MBL

systems. This can also be thought of as the cause of the slow growth of entanglement in MBL systems [112]. The existence of the LIOMs also imply the existence of many associated conserved quantities which cannot decay unless the system is coupled to a thermal bath [111–113]. These quantities prevent the system from thermalising since its evolution is constrained by many conservation laws. The initial conditions for the LIOMs determine the evolution of the system over time and break ergodicity. They imply that the system retains some ‘memory’ of its initial state, in contrast to a thermalised system.

Perhaps most interestingly, the integrability of MBL systems is in some sense more robust than integrability in other systems. Typical simple examples of integrable systems include non-interacting spinless fermions in a one dimensional lattice [114], and 1D homogeneous Bose gases with point-like collisional interactions [115], which have been experimentally observed not to thermalise [116, 117]. However such systems are typically fragile in the sense that small perturbations break the perfect integrability and produce a system that is ‘weakly’ integrable. This implies that although they may have short-term dynamics that are integrable, over long times their conserved quantities do decay and the systems thermalise. This rapid initial relaxation after a quench to a non-thermal state, followed by slow relaxation to a fully thermalised state has come to be called prethermalisation [118]. By contrast, under weak perturbations the MBL LIOMs are deformed but not destroyed and the systems remain integrable even at long times. This may reflect a conceptual difference between MBL as a dynamical phase of matter, and other integrable Hamiltonians which represent distinct and unusual points in the space of possible Hamiltonians [93].

4.3.4 Stability of MBL

The existence and stability of localised states in the presence of interactions remains somewhat uncertain. For example, the exact effect of interactions on Anderson localisation specifically remains an open question; in some case interactions suppress localisation by providing additional methods of transport [12]. Others predict that critical disorder strength still exists, but is dependent on the strength of interactions [119, 120]. We also must now consider the effects of finite temperature. Once again, some have predicted this introduces an additional threshold, a critical temperature below which localisation is stable [15, 20].

In general, the stability of many-body localisation is still theoretically debated and experimentally difficult to prove. In particular, it has been suggested that rare regions where the disorder is low can facilitate locally thermalising parts of the cloud.

These thermal spots would thermalise their immediate neighbours. If these spots grow without bound, they would thermalise the entire system and destroy the MBL phase. Utilising random matrix theory (RMT), it has been shown that in dimensions greater than one, we do expect such thermal spots to grow and cause an avalanche of thermalisation for even arbitrarily weak interaction strengths [22], but the application of RMT theory here is not universally accepted [93].

In addition to such a mechanism, the existence of an energy dependent localisation length provides another mechanism for tendency towards delocalisation. Namely, that high energy delocalised states can interact with the localised phase and act to thermalise it [19]. The dependence of localisation length on energy, $\xi(E)$, is strongly affected by the dimensionality of the system. In general, $\xi(E)$ grows more quickly in higher dimensions; in 2D already the dependence is exponential [19, 20]. The resilience of an MBL phase to this effect relies on the competition between a large localisation length for high energy states, and the reduction of occupancy of those states at low T .

Yet, Basko *et al.* [15] show that MBL phases exist below a critical temperature T_c for any number of dimensions. Additionally, Nandkishore [20] when specifically studying continuum systems shows that in two dimensions the MBL phase can be stable below T_c , where T_c is dependent on the properties of the disorder. This insensitivity to the delocalised high energy states is thanks to their reduced occupancy at $T < T_c$, and for sufficiently small interaction strengths a perturbative approach can always be applied at these temperatures. In 1D, the MBL transition exists at any finite T , dependent only on the interaction strength.

There are conditions for this stability, however. Firstly, stability requires that the disorder is impurity-like, in that it consists of many very short range, independent scatterers. Secondly, we have not addressed the behaviour of the aforementioned thermal pockets, regions where the energy density is unusually high that could trigger an avalanche of delocalisation. Nandkishore states that true stability of the MBL phase relies on assuming these pockets do not grow, and we can otherwise only say that the state is apparently localised until thermalisation occurs at long timescales [20].

However, Gornyi *et al.* [19] disagree even that these states are stable in this sense. They claim that the effect of delocalised high energy states in a continuum system has been underestimated and destroys the stability of MBL phases in any dimension, in direct contrast to [15, 20]. They predict the absence of a non-zero critical temperature

and instead predict only a significant slowing in the dynamics of the system at low T .

Many-body localisation remains an area of intense theoretical and experimental work, with many unanswered questions remaining, particularly in the transition region. Experimental work is often difficult due to the long length or time scales involved, and have only more recently come to produce important results thanks to technical advances in creating highly controllable, isolated disordered quantum systems.

4.3.5 Disordered phase transitions in Bose gases

We now turn to focus on phase transitions that occur in disordered Bose gases specifically. The existence of a stable MBL phase remains an open and interesting question in this case, and at finite temperatures additional phases may exist.

Considering first the Bose gas at zero temperature, a disorder driven transition from a superfluid to insulating phase exists in all dimensions [3, 11, 75, 121]. This insulator is referred to as a Bose glass and is characterised by a finite compressibility and infinite superfluid susceptibility. Bose glasses are zero-temperature ground states, and yet appear to have been observed experimentally at finite temperatures in one [122] and three [82] dimensions. Being a phase of the ground state, Bose glasses are not strictly speaking an example of many-body localisation which is the absence of thermalisation in excited states. However, some predict that for weakly interacting bosons [123] and for strongly interacting bosons in one dimension [124], the Bose-glass phase at zero temperature is smoothly connected to an MBL phase at finite temperature.

Bertoli *et al.* [3] predict the stability of the MBL phase on the basis that in real traps, a finite trap depth truncates the energy distribution at high energies, preventing delocalised atoms above the mobility edge thermalising the cloud. In contrast, while Carleo *et al.* [5] also find that disorder affects the superfluid-normal transition, they find no evidence for an MBL phase.

We are for now most interested in the role of disorder near the superfluid-normal transition in 2D, since this transition is the most easily probed by our matter wave interference scheme. Though the normal phase is absent at zero temperature, it is possible that a tricritical point exists in two dimensions where the three phases meet [3]. In general it is found theoretically [3–5, 75] and experimentally [83, 125] that the presence of disorder reduces the transition temperature T_c for the superfluid-normal transition in 2D and 3D. It also depletes the condensate fraction [83, 126, 127] and




Figure 4.6: A qualitative illustration of the possible phases (at constant density) in the 2D disordered Bose gas. At zero disorder, the gas exhibits the usual superfluid-normal BKT transition at temperature T_c , shaded blue. On introducing disorder, it is generally expected that the transition temperature is reduced [3, 5]. In contrast, some predict that disorder can enhance superfluidity [4, 8], a possibility shaded here in teal and indicated by the dashed lines that extended to temperature $T > T_c$. Following a line of increasing disorder strength reaches the transition from superfluid to the normal phase (shaded red) at some critical disorder strength V_c , and this transition is predicted to be of the BKT type [5]. At much stronger disorder strengths lies a possible insulating many-body localised (MBL) phase (shaded green). At high disorder strengths and zero temperature there is a Bose-glass phase [3, 11, 75, 121], which may connect smoothly to the MBL phase.

marginally increases the speed of sound in the cloud [75]. Bourdel [4] predicts that a small amount of speckle disorder actually enhances superfluidity and can increase the transition temperature, owing to an average increase in the density of the cloud. This prediction is made in the local density approximation and does not consider any potential localisation effects. In two dimensions the enhancement of superfluidity is a small effect, but it is more pronounced in three dimensional BECs. Werner *et al.* [8] also predicted that in certain parameter regimes, disorder was beneficial to superfluidity on a 2D lattice.

The reduction in T_c is caused at least in part due to localisation effects destroying global coherence of the superfluid (but not fully preventing thermalisation). This is

evidenced often by the fact that the critical chemical potential occurs at a higher value than that predicted by the combination of the local density approximation and the percolation threshold of the disorder [5, 75]. Despite the additional role of localisation relative to the disorder-free case, Carleo *et al.* [5] found that the critical exponent (as defined in Eqn. 4.17) of the transition in two dimensions is $\eta = 0.25$ - in other words, that it is still of the BKT universality class. This result was valid for disorder strengths less than the chemical potential.

Evidently, the experimental mapping of the phase diagram in the presence of disorder is of great theoretical interest. Investigations of the critical region around the superfluid-normal transition and a possible MBL transition may provide a solid base of evidence for further theoretical work.

Chapter 5

The BKT transition in disordered systems

In this chapter we discuss the experimental measurements of a disordered 2D superfluid in the vicinity of the BKT critical point, and the effect of disorder on the phase coherence of the cloud. We describe how we extract phase correlations from the interference patterns between two 2D clouds and present results indicating that disorder induces the superfluid-normal transition at a critical disorder strength. We quantitatively estimate the percolation threshold for our disorder potentials and rule out the effect of classical trapping.

5.1 Experimental procedure

We here describe the experimental sequence used to study the BKT transition in the presence of disorder. This includes the specific parameters of the disorder potential used, the method of imposing the disorder on the double-well RF dressed potential, and the analysis of the interference patterns that appear after time-of-flight expansion.

5.1.1 Pointlike disorder parameters

As mentioned previously, in this work we produce disorder in the form of many randomly placed, small scattering points. These scatterers are generated using the digital micromirror device (DMD), which controllably and randomly creates a 2D pattern of many small discs. Fig. 5.1a shows a typical example of the DMD pattern used for pointlike disorder, comprising 300 total disorder points. The red circle indicates the region spanned by the 2D cloud. The area of the DMD displaying disorder is significantly larger than this so that firstly, the potential is insensitive to drifts in alignment and the disorder pattern can not drift off of the atoms during experimental runs, and




Figure 5.1: (a) Example pattern displayed on digital micromirror device (DMD) mirror array. Each disorder point is a circle of diameter $210\text{ }\mu\text{m}$ prior to demagnification. The shaded red circle indicates the region occupied by the atoms in the image plane, and the disorder points are scattered over a region larger than this to make the system robust to slow drifts in alignment. (b) Simulated potential experienced by the atoms after the point spread function of the optical system blurs the initial DMD pattern. Each circle is spread into an approximately Gaussian potential profile. When scatterers are nearby they form short connected regions. For the filling factor shown, the percolation threshold is low and atoms can travel freely throughout the entire space.

secondly so that the disorder is present even for large amplitude dipole oscillations of the cloud. These dipole oscillations can be either incidental or intentionally induced, as we will discuss in Chapter 6.

For the results shown here, the width of each scatterer is constant, with a diameter on the DMD of 16 pixels. After magnification, this would become a disc of diameter $4.5\text{ }\mu\text{m}$. However, as we have shown, the optical system blurs shapes to noticeably wider than this. Here this is in fact an advantage; since the shape is significantly smaller than the PSF, it will be completely smoothed into a good approximation of a Gaussian profile, a commonly studied theoretical model of pointlike disorder. We need only account for the reduction in the peak intensity for this size of scatterer and a $\sigma_{PSF} = 2.50\text{ }\mu\text{m}$, which is a reduction to 66% of the unblurred value. Each scatterer is a close approximation of a Gaussian profile with $\sigma = 2.7\text{ }\mu\text{m}$.

Fig. 5.1b shows a prediction of the same disorder pattern experienced by the atoms after the potential is convolved with the point spread function of the optical system. Each circle pattern becomes approximately a Gaussian potential. The binary nature




Figure 5.2: Absorption image of a cloud in a disordered potential, taken on the vertical imaging system. We show here a disorder pattern with fewer, larger scatterers than those used in experiments in order to produce a clearer image, since the imaging system has limited resolution. The image is the average of 5 images taken under the same disorder potential to eliminate imaging noise. Each disorder point is a circle on the DMD array of diameter $11\text{ }\mu\text{m}$ after magnification, which is blurred by the optical system to an approximate Gaussian of $\sigma = 4.3\text{ }\mu\text{m}$. The scatterers are distributed at a density of $n_p = 5 \times 10^{-3}\text{ }\mu\text{m}^{-2}$, with a potential depth of $V/\mu \sim 2$. The density is reduced at the location of the disorder sites due to the repulsive optical dipole potential.

of the initial DMD pattern (mirrors being either in the on or off position) means that the maximum value of the potential is approximately capped at the peak V of a single scatterer even where multiple disorder points overlap. If many points overlap such that they form a single shape significantly larger than the size of the PSF, the peak intensity is reduced by the blurring much less than a single scatterer; but at the disorder densities used in this work this is unlikely and every point has approximately the same peak potential height.

An example of a cloud in a disordered potential, imaged vertically in the $x - y$ plane, is shown in Fig. 5.2. The disorder potential shown here has a lower density

of scatterers that that used in experiments, comprised of 150 scattering points (i.e. a density $n_p = 5 \times 10^{-3} \mu\text{m}^{-2}$). For clearer observation on the vertical imaging system, the cloud has a relatively high temperature with large thermal wings and the width of each point on the DMD is increased to a disc of $11 \mu\text{m}$ diameter (after magnification). The blurring effect of the PSF spreads this pattern out to a Gaussian profile with a width $\sigma = 4.3 \mu\text{m}$. The strength of the potential $V \sim 2\mu$ so that the disorder reduces the density locally to zero. In our experiments we use smaller disorder widths and usually lower disorder strengths, which are difficult to image directly on the vertical imaging due to the difference in focal position for the imaging and dipole light wavelengths.

We also control the number (or density) of scatterers. For the pattern shown in Fig. 5.1 (one of the patterns used in experiments), there are 300 points, or a density of $n_p = 10 \times 10^{-3} \mu\text{m}^{-2}$. As in Ref. [73] we define a filling factor $\rho = 2n_p\sigma^2 = 0.146$ for this pattern.

5.1.2 Loading the disordered double-well

To load the double-well we follow a similar procedure to that described in Chapter 3, with a few significant differences. We initially load the linear shell at a quadrupole current of 120 A which provides optimum loading stability and atom number. The initial atom number is typically between 90,000 and 110,000 atoms at temperatures between 40 – 70 nK. After this we adiabatically ramp the current to 147 A over a period of 1s, to increase the anisotropy of the trap and lift the atoms up to the focus of the green light. Residual collective motion from the shell loading persists after this ramp, and would hinder the stability of the splitting sequence and analysis of the phase profile after matter-wave interference.

To eliminate this motion, we slowly ramp up the disorder potential over the course of 150ms to a moderate intensity, with a peak potential height of $V = 200 \text{ Hz} \approx 0.2\mu$ where μ is the chemical potential of the cloud. After the potential reaches its maximum value, it is then again ramped down over another 150 ms. As we will show in Chapter 6, the disordered potential has a strong damping effect on the motion of the cloud. This eliminates or strongly reduces any residual collective motion of the cloud before the splitting sequence. The duration of the ramp is selected as a compromise between any heating induced by a rapid change in the potential for short durations, and a constant heating rate in the presence of higher disorder strengths due to pointing instability. This damping procedure produces an insignificant amount of heating and no measurable atom loss.

The splitting procedure is then initiated, with the full duration of the split occurring adiabatically over 200 ms. The Rabi frequency of each well and the barrier is then further ramped over 50 ms to reach their final trap frequencies which are measured to be $\omega_r = 2\pi \times 14.5(2)$ Hz, $\omega_z = 2\pi \times 1030(20)$ Hz. Each well contains 40,000 – 50,000 atoms, giving them a chemical potential of $\mu \approx 700$ Hz. The wells are separated spatially by a distance of $d = 7\text{ }\mu\text{m}$ and a barrier of $V_b \sim 10$ kHz, so the clouds are highly decoupled and tunnelling effects are completely negligible.

The broad criterion for a 2D cloud is that the interaction and thermal energy be less than the energy spacing between the ground and first excited axial state, $\mu, k_B T < \hbar\omega_z$. More detailed studies find that the cloud is in the 2D regime for $\mu/\hbar\omega_z < 1.28(14)$ and $k_B T/\hbar\omega_z < 2.11(30)$ [128]. In the work presented here, $\mu/\hbar\omega_z \sim 0.7$ and $k_B T/\hbar\omega_z < 1.3$, comfortably satisfying the 2D criterion. That the clouds sit in the axial harmonic oscillator ground state is further confirmed by the Gaussian density profile along this direction in time-of-flight, as the radial profile has a Thomas-Fermi (inverted parabolic) distribution typical of condensates in harmonic traps. We therefore have a high degree of confidence that we have two independent, two-dimensional clouds.

From this stage we adiabatically ramp up the disorder potential of interest. The potential is ramped up over the course of 250 ms to avoid exciting any motion in the cloud and reducing heating that would arise from the rapid change in the potential experienced by the atoms. After the disorder reaches its maximum strength the cloud is held in the potential for 100 ms to allow the cloud to finish equilibrating. The RF fields are switched off in the manner described in Chapter 3 and the clouds fall under gravity. For the interference experiments here we use a TOF duration of 15.5 ms.

For high disorder strengths, the presence of the disorder potential heats the atoms, on the order of 10 nK over the duration of the sequence for the disorder strengths tested here. This is problematic, since a temperature change this large is easily enough to cause the cloud to cross from the superfluid to the normal phase by itself. If we intend to discern the effect of the disorder alone, the temperature must be the same at the point we compare different clouds - in interference during time-of-flight.

To address this we use the radiofrequency knife to cut deeper in energy while evaporatively cooling the cloud during the initial single shell loading. This produces a colder cloud at the point when the disorder is turned on. By varying this initial temperature for each value of disorder strength, we can offset the heating caused by the potential and produce clouds of constant temperature at the end of the sequence before turning off the trap and disorder. Since this evaporation cools at the expense of




Figure 5.3: Measured temperature T (red squares) and atom number N (blue dots) for different disorder strengths V/μ after compensating for heating induced by the disorder potential ($n_p = 7 \times 10^{-3} \mu\text{m}^{-2}$). Temperature is rescaled by the critical temperature T_0 of a non-interacting gas in a harmonic trap. Higher disorder strengths cause more heating after ramping up and holding at their maximum value, so the clouds are evaporatively cooled to initially lower temperatures before loading so that they are at a consistent temperature after the application of disorder. The red shaded region indicates the previously measured critical temperature for the BKT transition [26], $T_c/T_0 = 0.53(1)$. Evaporative cooling slightly reduces the final atom number, so that N decreases slightly for increasing disorder strength.

removing atoms, this also slightly reduces the atom number relative to the sequences at low disorder strengths. This effect is small, and furthermore can be accounted for during analysis by rescaling the temperature T by the transition temperature T_0 for a non-interacting Bose-Einstein condensate in a 2D harmonic trap,

$$T_0 = \frac{\sqrt{6N}\hbar\omega_r}{k_B\pi}, \quad (5.1)$$

which depends on atom number. This provides a dimensionless measure of the temperature which also facilitates comparison between this work and previous results from our group on the BKT transition in the absence of disorder [26, 53].

This method of stabilising the final temperature is less effective for higher dis-

order strengths, which heat the cloud more. This necessitates the evaporation of greater numbers of atoms, reducing the final atom number. For sufficiently high disorder strengths, the final atom number is reduced significantly which affects the stability and quality of the results. As a result, to conserve atom number some disorder strengths have a higher than average temperature, but they remain below the previously measured BKT transition, $T_c/T_0 = 0.53$ [26].

Similarly we can compare the measured phase-space density, given by

$$\mathcal{D}(x) = n(x)\lambda_{th}^2, \quad (5.2)$$

where $n(x)$ is the 2D density in the trap at the location x . Because the confining potential is harmonic and the density profile is inhomogeneous (even in the absence of disorder), the phase-space density varies spatially. We infer the *in situ* density of the cloud from the known temperature and condensate fraction in time-of-flight and calculate the average phase-space density $\bar{\mathcal{D}}$ within the Thomas-Fermi region of the cloud and find that $\bar{\mathcal{D}} \approx 14$. This is higher than the critical phase-space density $\mathcal{D}_c$ for the BKT transition, $\mathcal{D}_c = 8.5$ [67]. Measurements of the phase-space density and temperature therefore confirm that the cloud is expected to be in the superfluid regime based on the BKT critical point in the absence of disorder.

5.2 Analysis of interference patterns

Here we describe the methods used to analyse the fringe patterns produced by matter-wave interference of the clouds, and how we extract the phase profiles and relate these to the state of the cloud. Further detail on these methods is described in the thesis of Shinichi Sunami [53].

5.2.1 Fitting the interference pattern

In time-of-flight (TOF) the 2D condensed clouds are described by a Thomas-Fermi profile in the radial directions x, y but a Gaussian in the axial direction z . As each cloud expands they overlap and produce a fringe pattern.

We consider each cloud in TOF to be the expansion of an initial Gaussian wavepacket $\psi_j(z, t)$ centred at position $\pm d/2$. After TOF at a time t the wavepacket is then described by [47]

$$\psi_j(z, t) \propto \exp\left(i\phi_j(x, y) - \frac{1}{2} \frac{(z \pm d/2)^2 (1 + iht/ml_0^2)}{l_0^2 + (ht/ml_0)^2}\right), \quad (5.3)$$




Figure 5.4: a) Absorption image of interference fringes between the two clouds in the double-well. The clouds were in the superfluid regime so that phase coherence is long and the fringes are relatively flat. We fit the density profile $n_x(z)$ along the z axis at each location x in the radial direction, indicated by the blue dashed line. b) The fitted fringe distribution to the image in a). The fit provides estimates of the wavelength, contrast and phase of the fringes. The phase $\theta(x)$ is the local relative phase between the two clouds at the point x . c) From the fitting we construct the profile of the relative phase $\theta(x)$. From many of these profiles we extract correlation functions, and detect any vortices by sudden large jumps in the phase over short distances.

where $\phi_j(x, y)$ is the phase of the cloud j *in situ* and l_0 is the axial width of the initial Gaussian profile.

The density in TOF is given by

$$n(z, t) = |\psi_1(z, t) + \psi_2(z, t)|^2 \propto e^{-z^2/2l_t^2} \left[1 + \cos\left(\frac{md}{\hbar t}z + \phi_1(x, y) - \phi_2(x, y)\right) \right], \quad (5.4)$$

where $l_t^2 = l_0^2 + (\hbar t/ml_0)^2$, for sufficiently long TOF such that $l_t \gg l_0, d$. In other words, the interference produces oscillating patterns in the density with a wavelength $\lambda = 2\pi\hbar t/md$. This wavelength at fixed t is determined by the initial separation and so is constant. However the phase of the oscillation $\phi = \phi_1(x, y) - \phi_2(x, y)$ depends on the relative phase between the clouds at each radial position. For 3D BECs, the clouds have constant phase so that this relative phase is also constant and the fringes appear ‘flat’. For a 2D superfluid, with only algebraic phase coherence, this relative phase

varies (mostly) continuously along the radial direction, producing ‘wavy’ fringes. The exception to this is the presence of vortices, which give rise to sudden jumps in the relative phase because of their net circulation.

Note that we neglect the effects of radial expansion of the cloud here. This is a good approximation for the condensate, which expands only a small amount in the TOF duration used here. Small features in the wavefunction, such as the ‘holes’ in density produced by a pointlike potential, will expand more rapidly since they have a greater width in momentum space. These features produce density fluctuations after TOF expansion, and may also produce small phase fluctuations. However, this phase is common mode to both clouds, since each cloud experiences the same potential, and so does not affect our estimate of the relative phase. The relative phase is the essential parameter we use to analyse the properties of the cloud.

Eqn. 5.4 predicts that the contrast of the fringes is 100% - the density drops completely to zero at its minimum. In practice the observed interference fringes have a finite contrast less than 1, due to a residual thermal component with no phase coherence, imbalanced atom number between the wells, or aberrations in the imaging system.

We fit the observed patterns with the distribution

$$n_x(z) = n_0 \exp\left(-\frac{(z - z_0)^2}{2\sigma^2}\right) \left[1 + A \cos(kz + \theta(x))\right], \quad (5.5)$$

where n_0, z_0, σ, A, k and $\theta(x)$ are the fitted parameters. The fit is performed along the z direction for each pixel column x along the imaged radial direction. This produces a phase profile along this direction $\theta(x)$, representing the relative phase difference between the clouds along the line $y = 0$. An example of fitting interference fringes and extracting a phase profile is shown in Fig. 5.4.

If integrating along the entire cloud during imaging, the wavy fringes produced by a 2D cloud near the superfluid-normal transition reduce the overall fringe contrast, making determination of the local relative phase impossible. We therefore repump a thin slice of the cloud of width $L_y = 4\mu\text{m}$. As long as the thickness of the slice is less than the correlation length of the cloud, this allows us to measure directly the local relative phase between the clouds.

Imaging a slice of the cloud provides essential local information, but we also lose the ability to measure global properties of the cloud. This includes, most importantly, the atom number, the temperature, and the condensate fraction. As a result we run each experimental sweep twice, once imaging the thin slice of the cloud to measure

the relative phase of the clouds, and once imaging the full cloud to obtain statistics about the global cloud properties at that data point.

To estimate the condensed fraction and temperature of the clouds, we fit the radial density profile with a bimodal fitting formula to separately model the condensed and thermal parts. The quasi-condensate is fitted with an inverted parabolic Thomas-Fermi profile, and the thermal wings are fitted with a Gaussian profile as is commonly done [32, 35]. Near the quasi-condensate transition temperature, the thermal wings of the cloud are not exactly a pure Gaussian profile [125, 129, 130]. However, the error in estimating the temperature of our data with a Gaussian fit compared to the often used polylogarithmic profile [59, 131] is $\sim 10\%$, which is relatively small. Furthermore, we have obtained good agreement between observed temperatures inferred using a Gaussian fit and 2D Monte Carlo simulations of the same experimental setup in the absence of disorder [26, 53]. Therefore we also employ a Gaussian fitting profile in this analysis, which additionally maintains consistency in results between the disorder-free and disordered cases where it is useful to have like-for-like comparisons.

This measurement technique can also be inaccurate at very low temperatures where the thermal wings are small and hard to observe, but for the data in this work we remain close to the disorder-free BKT transition temperature, between 40 – 60 nK, where the fitting provides a reliable estimate of the temperature.

5.2.2 Calculating correlation lengths

The phase of the interference oscillations is determined by the relative phase between the clouds $\theta(x, y) = \phi_1(x, y) - \phi_2(x, y)$. The primary property of interest for identifying the BKT phase is the one-body correlation function $g_1(\mathbf{r}, \mathbf{r}')$ of each individual cloud. We measure the two-point correlation function $C(\mathbf{r}, \mathbf{r}')$ of the relative phase between two clouds. Making the assumptions of uniform and equal density n for each cloud and for small density fluctuations then the correlation function C is given by

$$C(\mathbf{r}, \mathbf{r}') = \langle e^{i(\theta(\mathbf{r}) - \theta(\mathbf{r}'))} \rangle. \quad (5.6)$$

If each cloud is identical then they have the same one-body correlation function and $g_1(\mathbf{r}, \mathbf{r}')$ is then

$$C(\mathbf{r}, \mathbf{r}') = \frac{g_1(\mathbf{r}, \mathbf{r}')^2}{n^2}. \quad (5.7)$$

Eqn. 5.7 shows that the correlation function g_1 of a single cloud is related to the correlation of the relative phases. The correlation function of relative phases is given




Figure 5.5: a) The measured two-point correlation function $C^r(x, x')$ for a superfluid cloud below the BKT critical point. Each point represents the average correlation over many ($N \approx 50$) phase profiles $\theta(x)$ and $\theta(x')$, a single example of which is shown below and left of the correlation map, respectively. The diagonal of $C^r(x, x' = x)$ is perfectly correlated and the off-diagonal values represent points at greater distances apart in the cloud, with lower correlation. b) Averaging the two-point correlation function over points with constant separation $\bar{x}$ produces the one point correlation function $C(\bar{x})$, related to the one-body correlation function in a single cloud $g_1(r)$. To this function we fit both algebraic or exponential models, the fit quality of which indicates the phase of the cloud. Error bars are 1σ deviation obtained by bootstrapping while calculating the correlation map.

by Eqn. 5.6. Here for discrete experimental data we calculate the equivalent quantity:

$$C(x, x') = \frac{1}{N} \sum_j^N e^{i(\theta(x) - \theta(x'))}, \quad (5.8)$$

where N is the number of repeats per data point and typically $N \sim 50$. We take the real part of C , $\text{Re}(C) = C^r(x, x')$ which equals 1 for perfectly correlated points and 0 for uncorrelated points. An example of $C^r(x, x')$ is shown in Fig. 5.5. Each x here represents a single pixel line for which the relative phase was fitted, and x is in units of image pixels unless otherwise indicated.

The correlation function in the BKT phase is typically expressed in terms of a single length, $g_1(r)$. To obtain a similar expression we average $C^r(x, x')$ over pairs of points with the same spatial separation, $x - x' = \bar{x}$. Not all pairs of points are averaged over, but rather a selection of $N_{\bar{x}} = 5$ points corresponding to $\bar{x}/2, \bar{x}/2 \pm 1, \bar{x}/2 \pm 2$

away from the centre of the cloud. This is because there are fewer combinations of points that are separated by a large distance comparable to the full cloud width, and this allows us to obtain an estimate of the error for the correlation at large $\bar{x}$. The averaged correlation function $C(\bar{x})$ is then given by

$$C(\bar{x}) = \frac{1}{N_{\bar{x}}} \sum_{\bar{x}/2-2}^{\bar{x}/2+2} \text{Re}[C(x, x + \bar{x})]. \quad (5.9)$$

5.2.3 Accounting for inhomogeneity in the harmonic trap

The theory of the BKT transition presented in Chapter 4 describes homogeneous systems. In inhomogeneous systems, the varying density leads to a spatial dependence of quantities such as the phase-space density or algebraic exponent η as defined in Eqn. 4.17. We account for this in a manner similar to Ref. [132], in which an accurate spin-wave description of the harmonically trapped gas is developed. In addition to the spin wave theory, the authors find that a good approximation of their result is obtained by applying the Local Density Approximation (LDA) to the homogeneous system and introducing a spatial dependence into the exponent of the correlation decay:

$$g_1(\mathbf{r}, \mathbf{r}') \propto |\mathbf{r} - \mathbf{r}'|^{-\eta \frac{n_0}{\sqrt{n(\mathbf{r})n(\mathbf{r}')}}}}, \quad (5.10)$$

where n_0 is the peak density. This is referred to as the local correlation approximation (LCA). We employ this technique here and fit the algebraic decay of correlations $C(\bar{x})$ with the function

$$f_{SF} = a\bar{x}^{-2\eta_0 \frac{\max(n(\bar{x}))}{n(\bar{x})}}, \quad (5.11)$$

where a and η_0 are the fit parameters and the factor of 2 comes from summing the fluctuations of two clouds in the measured correlation function [53]. This produces a good fit to the behaviour of the cloud (shown in Fig. 5.6) and the extracted η is the value at the centre of the trap.

In the thermal regime the exponential decay of the cloud is typically much more rapid than the length scale of density variations in the trap, and so the inhomogeneity has a comparatively much smaller effect. As a result we do not account for this in our analysis, but more investigation into the crossover region where the correlation length is comparable to the cloud may be needed.




Figure 5.6: Comparison between fitting the correlation function $C(\bar{x})$ in the superfluid regime ($T/T_0 = 0.43$) with the pure algebraic scaling applicable to a homogeneous system (red dashed line) and the algebraic scaling within the LCA (blue solid line) using the measured *in situ* density. Grey shaded region shows the 1σ error on the measured correlation function.

5.2.4 Vortex detection

As mentioned previously, the relative phase between fringes varies smoothly except in the vicinity of a free vortex. Exactly on opposite sides of a vortex the phase changes by an multiple of π and this manifests as a discontinuous jump in the relative phase in the interference fringes. We detect this only if the repumping slice happens to align with the centre of a vortex. If instead the slice is slight off centre then we may detect a large jump in the phase but not by exactly π .

We consider a vortex to be present if we measure a phase difference between $2\pi/3$ and $4\pi/3$ between next-nearest neighbour pixels. We take this criterion rather than adjacent pixels because one pixel width ($1.67 \mu\text{m}$) is comparable to the width of a vortex core, $\sim 1 \mu\text{m}$.

To estimate the average vortex density in the cloud from the number of detected vortices, we calculate the average number of vortices N_{col} observed per single pixel

column over a set of images. The detection area this single pixel corresponds to is $L_y \ell_p$, where $L_y = 4 \mu\text{m}$ is the width of the repumping slice and $\ell_p = 1.67 \mu\text{m}$ is the width of a single pixel in the image plane. The vortex density n_v is then

$$n_v = \frac{N_{col}}{2L_y \ell_p}, \quad (5.12)$$

where the factor of 2 comes from the fact that we can detect vortices in either of the two clouds and we define n_v to be the vortex density of a single cloud.

5.3 Observation of the disordered superfluid-normal transition

In this section we present experimental results on the effect of the pointlike disorder on the superfluid-normal transition in the harmonically trapped 2D Bose gas. First we demonstrate that the disorder reduces coherence by depleting the quasicondensate fraction even at low disorder strengths. We then show that the disorder can induce a transition from the superfluid to normal phase, and that this transition is consistent with being of the BKT type. We discuss some currently unexplained behaviour at high disorder strengths and scatterer density, and measure the effect of disorder on vortex density.

5.3.1 Percolation threshold

An important element of any analysis of the effect of the disorder is the relative importance of classical and quantum effects. Specifically, we calculate the percolation threshold E_{perc} for the potential used to determine the effect of classical trapping. Using the dependence of the percolation threshold on pointlike filling factor from Ref. [73], we can estimate the percolation threshold for pointlike scatterers of width $\sigma = 2.7 \mu\text{m}$ and a given density n_p . We experimentally test two values of density: $n_p = 10 \times 10^{-3} \mu\text{m}^{-2}$, and $n_p = 7 \times 10^{-3} \mu\text{m}^{-2}$. The percolation threshold in the first case is $E_{perc} = 0.40\bar{V}$, where $\bar{V}$ is the average disorder energy, and $E_{perc} = 0.20\bar{V}$ in the second case. Although we probe slightly higher V for the lower n_p , the percolation threshold remains below that of the denser disorder potential in all cases.

The optical system blurs the circular scatterers in the image of the DMD into Gaussian profiles but conserves energy, so we can straightforwardly calculate the average potential energy by

$$\bar{V} = V_0 n_p \pi r_0^2, \quad (5.13)$$

where V_0 is the peak disorder strength *before* taking the PSF into account, and r_0 is the radius of each disc in the unblurred image of the DMD, calculated using the known magnification of the optical system. This is checked for consistency against a simulated pattern of the disorder convolved with the PSF and provides a very similar result. Eqn. 5.13 slightly overestimates the real average potential because it counts overlapping scatterers twice, while in fact the potential from two overlapping circles on the DMD is capped at V_0 .

For the disorder strength $V = 1100(70)$ Hz and $n_p = 10 \times 10^{-3} \mu\text{m}^{-2}$ the average potential is $\bar{V} = 260(15)$ Hz, resulting in a percolation threshold energy of $E_{\text{perc}}/\mu = 0.16(2)$. This is significantly less than the chemical potential of the cloud and is also the highest percolation threshold that exists for any of the data presented here, indicating that classical trapping effects are not responsible for observed behaviour.

5.3.2 Depletion of the condensate

The condensate fraction is measured by the fitting of a bimodal distribution along the radial axis of the cloud (imaged along its full width, not within a thin slice). The condensate fraction supplies the simplest measure of how the coherence of the cloud is affected by disorder, as both the superfluid density n_s and quasicondensate fraction n_0 are expected to decrease in the presence of disorder [126, 127].

Fig. 5.7 shows an example of the density profile after time-of-flight for two different clouds, one in the absence of disorder and one at a high disorder strength $V/\mu = 2$. The clouds are fully repumped and imaged along the horizontal axis, integrating over the direction y . The density profile in x along a slice through the centre of the image is extracted and fitted with a bimodal distribution. The central parabolic profile of the condensate contains fewer atoms in the case of strong disorder when compared to the disorder-free case, with atoms being transferred instead into the thermal fraction even at the same temperature.

We sweep the amplitude of the disorder while varying the initial temperature of the atoms to produce clouds at approximately similar temperatures in TOF as described in Section 5.1.2. The density of scattering points for this data is $n_p = 7 \times 10^{-3} \mu\text{m}^{-2}$. The temperatures are quite consistent, with the exception of the highest disorder strength for which the heating was not fully mitigated.

To disentangle possible confounding effects from variation in temperature, we plot the condensate fraction as a function of temperature and disorder strength. Fig. 5.8 shows the measured condensate fraction as a function of T/T_0 for both the absence of disorder and increasing disorder strength. Although stronger disorder causes a small




Figure 5.7: Example of a bimodal fit for two clouds after 15.5 ms TOF at zero and strong disorder, expressed in units of chemical potential V/μ . The density profile is integrated along the imaging axis and a horizontal slice taken through the centre of the resulting 2D image. Each cloud contains ~ 40000 atoms at a temperature of 40 nK . The parabolic part of the fit describes the density profile of the condensed fraction and the thermal fraction is described by the Gaussian outer wings. The density of the scattering points in this example is $n_p = 7 \times 10^{-3} \mu\text{m}^{-2}$.

upward trend in the temperature, the condensate is depleted much faster by the increasing disorder strength than the equivalent amount of heating would otherwise cause. We also repeat the measurement for a higher density of scattering points, $n_p = 10 \times 10^{-3} \mu\text{m}^{-2}$, shown in Fig. 5.9. The same effect is observed for this higher density of points.

To examine whether a higher density of scatterers causes a stronger condensate depletion, we can compare the condensate fraction as a function of disorder strength directly for both, as shown in Fig. 5.10. The higher density disorder does appear to reduce the condensate fraction (coherence) of the cloud by a greater degree than the lower density case.




Figure 5.8: Condensate fraction as a function of temperature T/T_0 , for clouds free from disorder (grey circles) and for clouds subjected to increasingly strong disorder (squares). Increasing disorder strength causes the condensate fraction to drop much faster than the equivalent increase of temperature without disorder, indicating that the disorder is directly responsible. Linear fits are included as a guide to the eye. The density of scattering points used was $n_p = 7 \times 10^{-3} \mu\text{m}^{-2}$. Error bars are 1σ standard error on the measured condensate fraction and T/T_0 .

The temperatures of every point remain considerably below the measured quasi-condensate transition temperature, $T_{qc}/T_0 \approx 0.75$, defined as the temperature at which a condensed fraction appears in the absence of disorder. Furthermore each point is less than (though in some cases only slightly) the BKT transition temperature in the absence of disorder previously measured at $T_{BKT}/T_0 = 0.53$. The measurement of the condensed fraction alone does not provide any information as to the superfluid density or correlation length of the cloud, but we can consider ourselves to be in the vicinity of the critical point for all the clouds considered here, rather than in a low temperature regime.




Figure 5.9: Condensate fraction as a function of temperature T/T_0 , for clouds free from disorder (grey circles) and for clouds subjected to increasingly strong disorder (squares), as shown in Fig. 5.8. Linear fits are included as a guide to the eye. The density of scattering points used was $n_p = 10 \times 10^{-3} \mu\text{m}^{-2}$, higher than that in Fig. 5.8. Error bars are 1σ standard error on the measured condensate fraction and T/T_0 .

5.3.3 Correlations near the critical point

To determine the correlation statistics of the clouds in presence of disorder we now turn to the matter-wave interferometry technique. For many images ($N \approx 50$) we fit the interference pattern with Eqn. 5.5 at each pixel along the radial direction and obtain a relative phase $\theta(x)$. We calculate the two-point correlation function $C(x, x')$ as in Eqn. 5.8.

In order to fit the decaying correlations with an algebraic or exponential model and hence identify the superfluid or normal phase, we convert this two-point correlation to a one-point correlation function of separation $\bar{x}$ as in Eqn. 5.9 and plot the resulting correlation functions $C(\bar{x})$ in Fig. 5.11. Also shown are the fitted algebraic and exponential curves. The disorder density used for this data was $n_p = 10 \times 10^{-3} \mu\text{m}^{-2}$.

Some of the correlation functions display an apparent oscillatory behaviour, most clearly visible at disorder strengths just above the chemical potential. They display




Figure 5.10: Comparison of condensate depletion for different pointlike scatterer densities. Temperatures and atom numbers are approximately constant over the range of data. Vertical error bars are the standard error in the measurement of the quasicondensate fraction. Horizontal error bars are the combined uncertainty in the estimate of the potential depth and standard error in the measured chemical potential. The ‘bump’ in the condensate fraction at $V/\mu \approx 1$ is currently unexplained, since it does not correlate with a fluctuation in the measured temperature, atom number or correlation length. Further data is needed to eliminate the possibility of a statistical fluctuation in the measured condensate fraction.

an initial drop and then recovery of the phase correlation at short distances $\lesssim 6 \mu\text{m}$, and a similar drop and recovery at $\sim 15\text{--}20 \mu\text{m}$. Similar (albeit smaller) behaviour at short distances is also observed in the disorder-free case [26], but the behaviour at longer distances may be unique to the presence of disorder. A possible explanation for the short distance behaviour is the presence of a vortex pair (which are more likely to occur at higher disorder strengths since, as shown in Fig. 5.14), where the vortices cause an initial phase dislocation between themselves but screen each other at distances greater than the pair separation.

Possibly the reduction in correlations at $\sim 15 \mu\text{m}$ is similar to the short distance behaviour in that the presence of many vortices at that location causes a rapid re-




Figure 5.11: Correlation functions $C(\bar{x})$ of the cloud at different peak disorder strengths. Average temperature of each cloud T/T_0 rescaled by the non-interacting BEC transition temperature T_0 also included for comparison. Also plotted are the exponential fits (red), and algebraic fits (green) which include the effect of inhomogeneous density. Grey region indicates 1σ deviation obtained from bootstrapping.

duction in the correlation. These vortices may cluster around disorder points, so that correlations are reduced not only at the separation between individual pairs, but at typical separations between clusters. However, the mean separation between disorder points is slightly less than the observed feature, at $n_p^{-1/2} \approx 10 \mu\text{m}$. This picture therefore seems to be incomplete, and it could also be related to a gradual onset of the general superfluid phase at high disorder strength, which we will discuss later in Section 5.3.5.

We compare the relative quality of the algebraic and exponential models by the value of the reduced χ^2_ν statistic for goodness of fit, plotted in Fig. 5.12a. Lower χ^2_ν values indicate a better quality fit. The algebraic model remains a superior description of the correlation function for $V/\mu < 1$, indicating that the cloud remains in the BKT superfluid state. At $V/\mu \sim 1$ the fits become of a similar quality and the exponential model becomes a better description at disorder strengths a little higher than this, indicating the cloud has transitioned to the normal, thermal phase. The estimated critical disorder strength at which this crossover occurs is $V_c/\mu = 1.1(1)$. At the highest disorder strength however, the cloud appears to become better described by the algebraic model again.

In Fig. 5.12b we plot the fitted values of the algebraic exponent η for each disorder strength. Recall that η is related to the superfluid density of the cloud by $\eta =$




Figure 5.12: (a) Fit quality of algebraic and exponential models for correlation decay at pointlike disorder density $n_p = 10 \times 10^{-3} \mu\text{m}^{-2}$, as measured by the reduced χ^2_ν statistic. At low disorder strengths, the cloud is better described by algebraic scaling (green), indicating the superfluid phase. Above a critical strength $V_c/\mu = 1.1(1)$ (black dashed line) the exponential scaling becomes a better model for correlations, indicating transition to the normal phase. The cloud then appears to re-enter the superfluid phase for even higher disorder strengths, indicated by the blue shaded region. (b) Fitted algebraic exponent η as a function of disorder strength. Red dashed line indicates the critical exponent $\eta_c = 0.17(3)$, defined as the value of $\eta(V_c)$ as estimated from a quadratic fit (green dashed line). Black dashed line indicates estimated critical disorder strength V_c , defined as the point where the cloud transitions from algebraic to exponential decay of correlations. η decreases again as the correlations of the cloud appear to improve at higher disorder strengths.

$1/n_s \lambda_{th}^2$ and at the critical point the predicted critical exponent is $\eta_c = 0.25$. Previous experiments in our system on the BKT transition in the absence of disorder find a slightly lower value of $\eta_c = 0.17(3)$ mainly due to finite size effects [26].

We estimate the critical exponent heuristically by fitting points either side of V_c with a quadratic fit and finding $\eta_c = \eta(V_c)$. This gives $\eta_c = 0.17(3)$, similar to the critical exponent in the disorder-free case. The measured temperature of the cloud near this point is $T/T_0 = 0.44(1)$, significantly lower than the transition temperature in the absence of disorder, which is measured at $T/T_0 = 0.53(1)$ [26].

In Fig. 5.13a we plot the fit quality of algebraic and exponential scaling for a lower density of disorder, $n_p = 7 \times 10^{-3} \mu\text{m}^{-2}$. In this case the crossover from algebraic to exponential scaling occurs at significantly higher critical disorder strength, $V_c/\mu =$




Figure 5.13: (a) The fit quality of exponential (red) and algebraic (green) scaling to the decay of correlations for a disorder density $n_p = 7 \times 10^{-3} \mu\text{m}^{-2}$ measured by the reduced χ^2_ν statistic. The cloud is superfluid except for the highest strength tested where it crosses into the normal regime. The estimated critical disorder strength (black dashed line) at the crossover point is $V_c/\mu = 1.9(2)$. A large peak in the χ^2_ν for the exponential fit is seen for $V/\mu \sim 0.75$; the cloud is deeply in the superfluid regime, as evidenced by the low value of η at these points. (b) Fitted algebraic exponent η as a function of disorder strength. The estimated critical $\eta_c = \eta(V_c)$ (red dashed line) using the linear fit (green dashed line) is $\eta_c = 0.20(2)$.

1.9(2). Due to the large peak in the fitting error for the exponential fit around $V/\mu \sim 0.75$, the vertical scale of Fig. 5.13a appears compressed near the transition to the exponential scaling. However, the χ^2_ν for the algebraic fit is almost 50% higher than the exponential case, which is a significant difference. However this is only one data point, and more data is needed to confirm the superfluid behaviour of the cloud above this point.

Fig. 5.13b shows the fitted value of the algebraic exponent η for this case. Here we see that correlation appears to improve (lower η) for a small amount of disorder. Although such an effect has been predicted [4], it is also possible to attribute this to residual collective motion (‘sloshing’) in the cloud. The lower disorder density is less effective at eliminating this motion in the damping stage prior to the double-well loading, and so it may leave more residual sloshing in the cloud compared to the higher density ($n_p = 10 \times 10^{-3} \mu\text{m}^{-2}$). Higher disorder strengths may damp any remaining motion after the double-well splitting, and any residual sloshing at lower strengths will contribute to an apparently shorter correlation length.




Figure 5.14: Density of vortices for different pointlike disorder density n_p . Critical disorder strength $V_c(n_p)$ indicated by vertical dotted lines. A higher density of pointlike scatterers (bright green circles) produces more vortices at the same disorder strength than a lower density (dark green squares). Linear fits are included as a guide to the eye. Horizontal error bars are the combined uncertainty in the estimate of disorder strength and standard error in the chemical potential. Vertical error bars are the standard deviation in estimated vortex density.

Above this point, higher V/μ does appear to increase η and the measured critical exponent is $\eta_c = 0.20(2)$. Once again, this is consistent with the previously measured critical exponent in the absence of disorder. The temperature at this point is $T/T_0 = 0.47(2)$, also below the critical temperature for the BKT transition at zero disorder strength.

5.3.4 Vortex density

In Fig. 5.14 we plot the density of vortices as measured by sudden phase jumps in the relative phase profile of the clouds. It is clear that the number of vortices increases

significantly for increasing disorder strength and disorder density. No stark transition in the vortex density occurs at the critical disorder strength $V_c(n_p)$ for each density n_p . Additionally, at the highest disorder strengths for $n_p = 10 \times 10^{-3} \mu\text{m}^{-2}$ the correlation length was long and η was small, from which we would expect to observe fewer vortices; while we see a reduction in the vortex density for $V/\mu \sim 1.4$, the highest disorder strength $V/\mu \sim 1.8$ shows no such reduction.

The general increase of vortex density with disorder strength may be due to the disorder nucleating vortices at the disorder sites. At high amplitudes ($V \sim \mu$) each disorder point reduces the local density significantly, causing the cloud to locally cross the transition into the normal phase. This may cause vortices to proliferate spontaneously at these locations. That vortices are the driving mechanism behind the transition in the case of disorder is supported by the matching critical exponent with the disorder-free case (which comes from the free energy needed to produce vortices).

5.3.5 Superfluid regime at high disorder strengths

We now consider the feature shown in Fig. 5.12 where the superfluid appears to be enhanced at the highest disorder strengths (shaded blue region) for the case of $n_p = 10 \times 10^{-3} \mu\text{m}^{-2}$. This is evidence by the reduction in η below η_c and the crossing back to algebraic scaling. This counter-intuitive result requires more critical consideration.

These results are dependent on the fitting of both algebraic and exponential models to the correlation function $C(\bar{x})$. In Fig. 5.15 we plot the correlation function for low and high disorder strengths, for both values of scatterer density n_p tested. Fig. 5.15a shows correlation functions for a density $n_p = 7 \times 10^{-3} \mu\text{m}^{-2}$ and displays the expected behaviour - that higher disorder strength reduces the coherence of the cloud and shortens the correlation length. For $V < V_c$ the cloud shows approximately linear scaling on a log-log scale, indicating algebraically scaling correlations (with a deviation at large distances due to the inhomogeneity of the density). For $V > V_c$ the correlations drop rapidly on a log scale indicating exponentially decaying correlations.

Clouds in the higher disorder density $n_p = 10 \times 10^{-3} \mu\text{m}^{-2}$, shown in Fig. 5.15b display the same expected behaviour for $V = 0$ and $V = V_c$. However, while for the highest disorder strength tested the correlations do indeed decay faster than the disorder-free case for short distances, they plateau at a moderately high value, and scale algebraically with distance. This relatively high value of the correlation at long




Figure 5.15: Comparing correlation functions for different pointlike disorder densities and disorder strengths, shown on a log-log scale. Error bars are 1σ deviations obtained from bootstrapping. (a) For $V < V_c$ (blue) we see a slow, approximately algebraic decay of correlations. Deviation from a straight line occurs at long distances ($\bar{x} \gtrsim 10 \mu\text{m}$) due to the density (and hence spatial) dependence of the exponent, see Section 5.2.3. For $V \approx V_c$ (green), the correlation decays exponentially. (b) $V < V_c$ (blue) decays slowly with algebraic scaling, while for $V \approx V_c$ (green) the correlations decay exponentially and drop to zero at the edge of the cloud. The highest disorder strength $V/\mu = 1.76 > V_c$ (purple diamonds) initially shows a decay of correlation up to separation $\bar{x} \sim 10 \mu\text{m}$, but for larger separations levels off at a moderately high correlation comparable to the disorder-free case. This gives an approximately linear profile in the log-log scaling, indicating algebraic scaling with low value of η .

distances also produces a low apparent value of η , suggesting that the cloud has re-entered the superfluid regime.

One possible explanation for the unusually long-range coherence at high disorder strengths may be related to vortex pinning. Areas of reduced density represent a reduction in interaction energy and so vortices may cluster around (are ‘pinned’ to) the disorder sites. While these areas may be locally in the thermal phase, the region between the disorder sites (which spreads throughout the full system due to the low percolation threshold) may be a fully connected superfluid that is comparatively free from vortices. Vortices of opposite sign that are clustered around a disorder site may act like bound pairs, resulting in a net circulation of zero around the disorder sites. This mechanism is illustrated in Fig. 5.16a .

Another possibility is that while strong disorder locally reduces the density, the




Figure 5.16: Illustrations of possible mechanisms by which strong disorder could enhance superfluidity. a) Density reduction caused by the repulsive disorder potential causes vortices (red) to become pinned to the disorder sites (green). The free vortices which are usually responsible for the reduction in coherence over long length scales are absent or rare in most of the region outside the disorder sites, leading to a superfluid region that spreads over the whole cloud since the percolation threshold is low for pointlike disorder. Clustered vortices of opposite sign shield each other's effect on the rest of the cloud. b) The disorder-free value of density n_0 (black dashed line) may be below the critical density n_c for the superfluid transition (blue dashed line). The presence of a strong repulsive potential V (green line) reduces the local density, so that in the vicinity of the disorder the cloud is in the thermal phase (red shaded region). Atoms pushed out of this region into the rest of the cloud increase the density elsewhere to n_V , greater than n_c , producing a superfluid region (shaded blue) outside of the disorder sites with a density n_V . This illustration shows a homogeneous density for simplicity; a qualitatively similar effect can still occur in an inhomogeneous system like a harmonic trap [4], though the cloud may also slightly expand in size.

displaced atoms are pushed into the surrounding cloud and therefore increase the density in those regions around a disorder site, shown in Fig. 5.16b. Modelling of superfluids in speckle disorder has predicted that compression arising from the presence of disorder enhanced superfluidity, but that this effect competed with the increasing percolation threshold for higher disorder strengths resulting in a small effect in 2D harmonic systems [4]. The speckle potential does however have a much higher percolation threshold than the pointlike disorder used here. It may therefore be the case that the compressing effect of the disorder persists to higher disorder strengths than expected from the treatment of a speckle disorder. However, other studies of disorder predict no such effect [3, 75].

Further experiments are needed to clarify the mechanism at play. To identify vortex pinning it may be sufficient to observe many fringe patterns that feature multiple vortices that are separated by a distance comparable to the disorder size $\sigma = 2.7 \mu\text{m}$. In the data presented here, this method presents a problem. This distance is small and equivalent to approximately 2 pixels in the imaging plane. Our analysis identifies a vortex at the location of a phase jump that occurs over a distance of 2 pixels, so that we may incorrectly identify multiple close vortices as a single vortex (or no vortex at all if the vortices are of opposite sign so that the net circulation is zero). Varying additional parameters of the disorder, such as the scatterer size σ , might provide further insight.

5.4 Conclusion

In this chapter we described the experimental procedure for studying the effect of pointlike disorder on a condensed 2D cloud, and the transition from superfluid to normal state in the presence of this disorder. We investigated these effects at two different densities of the disorder and found that the higher density is in general more deleterious to the coherence of the cloud than the dilute case, as measured by the depletion of the condensate and the superfluid density, and the loss of coherence occurs at lower disorder strengths for denser disorder.

We observed that at a critical strength the disorder induces a transition from algebraic to exponential scaling of the phase correlations in the cloud in the same manner as the disorder-free BKT transition, and that the critical exponent of this transition matches the BKT critical value. The transition occurs at a lower temperature than in the absence of disorder, as predicted. However, unexpected behaviour is seen for the highest disorder strengths in the case of moderately dense disorder, where the phase correlation of the cloud improves and the superfluid density appears to increase. We suggest a possible explanation that this may be due to one or both of vortices clustering around the disorder sites leaving a fully connected superfluid region between the disorder sites, and compression of the cloud by the disorder increasing the local phase-space density over the critical value.

Chapter 6

Damping of superfluid flow

In this chapter we will discuss a preliminary investigation into superfluid flow in 2D systems. We begin with an overview of the classic theory of superfluidity and discuss the study of superfluid flow in a disordered potential. We then describe observations of damped motion around individual objects and pointlike disorder in our harmonic 2D system. We observe a critical velocity below which dipole oscillations of a superfluid around an object are undamped, and above which the damping rate grows with increasing velocity. This critical velocity is compared for different heights of obstacles. We examine the same damping effects in disordered potentials and find that the damping force is much increased.

6.1 Theory of superfluid flow

6.1.1 Landau criterion for superfluidity

Following the discovery of superfluidity in liquid helium-4 by Kapitza, Allen and Misener [133, 134], Landau introduced a simple yet powerful phenomenological argument for the existence of frictionless flow in condensed fluids. Here we describe this argument, following the formulation of Ref. [47].

Consider a non-interacting, uniform condensed fluid travelling at a constant velocity $\mathbf{v}$ past an obstacle. This situation is equivalent to an object travelling at velocity $-\mathbf{v}$ through a stationary fluid when viewed in the rest frame of the fluid. For a system with energy E and momentum $\mathbf{p}$ in a given frame of reference, the energy in a frame moving with relative velocity $\mathbf{v}$ is given by the Galilean transform:

$$E(\mathbf{v}) = E - \mathbf{p} \cdot \mathbf{v} + \frac{1}{2}Mv^2, \quad (6.1)$$

where M is the total mass of the system. Therefore in the rest frame of the obstacle the ground state energy is

$$E(\mathbf{v}) = E_0 + \frac{1}{2}Mv^2, \quad (6.2)$$

where E_0 is the ground state of the fluid in its rest frame. Now consider a state where the fluid has an excitation of momentum $\mathbf{p}$. In the rest frame of the fluid, the energy of this state is $E_{ex} = E_0 + \epsilon_p$, where ϵ_p is the energy of an excitation with momentum p , dependent on the spectrum of the system. Applying the Galilean transform, the energy of the excited state in the rest frame of the obstacle is

$$E_{ex}(\mathbf{v}) = E_0 + \epsilon_p - \mathbf{p} \cdot \mathbf{v} + \frac{1}{2}Mv^2. \quad (6.3)$$

The difference in energy between the ground state and excited state is the energy required to create an excitation in the rest frame of the obstacle. In this frame the obstacle is static and can not transfer any energy to the fluid. Therefore it only becomes possible to create this excitation when

$$v = \frac{\epsilon_p}{p}, \quad (6.4)$$

which produces an excitation travelling parallel to the fluid. Excitations which propagate at a different angle require greater fluid velocities. This indicates that an excitation can be made only if its phase velocity is greater than or equal to the velocity of the fluid relative to the obstacle. This implies the existence of a critical velocity v_c given by

$$v_c = \min\left(\frac{\epsilon_p}{p}\right). \quad (6.5)$$

This result is known as the Landau criterion for superfluidity, and this v_c the Landau critical velocity. In the case of a weakly interacting Bose-Einstein condensate, the energy of excitations of the condensate is given by the Bogoliubov spectrum [47, 135]:

$$\epsilon_p = \sqrt{\left(\frac{p^2}{2m}\right)^2 + 2\mu\frac{p^2}{2m}}, \quad (6.6)$$

where μ is the chemical potential of the uniform unperturbed fluid.

For this spectrum the minimum value of ϵ_p/p occurs at $p = 0$ where the phase velocity is $\sqrt{\mu/m} = c$ where c is the speed of sound. For fluid velocities below this velocity, no excitations can be produced, leading to no drag force and entirely frictionless flow. It is common to report the critical velocity of a system in units of the speed of sound, so that in this case Landau's theory predicts $v_c/c = 1$.

Although its prediction of the value of the critical velocity has generally proved to be an overestimate, Landau's argument for its existence is powerful and has proved correct in both numerical and experimental investigations of many systems [136–139].

6.1.2 Critical velocity for realistic obstacles

We can divide obstacles in a superfluid flow into categories based on their shape and size: narrow, where the size of the object is comparable to or less than the healing length ξ ; and wide, where the size is much greater than ξ . We may also consider obstacles to be ‘hard-walled’, where the potential increases very steeply as in a step function, or ‘soft-walled’, where the potential increases smoothly to a peak. This latter case is the more experimentally relevant for cold-atom experiments, where obstacles are frequently produced with optical dipole potentials [136–138] for which the potential typically has a Gaussian profile, with a width $\sigma > \xi$.

Narrow obstacles produce a sharp dip (or peak, in the case of an attractive potential) with a size $\sim \xi$ in the density of the fluid. For weak obstacles the critical velocity is equal to the speed of sound, as in Landau’s prediction. The critical velocity in the case of a repulsive potential decreases smoothly and asymptotically towards zero as the strength of the perturbation is increased (by increasing the height of the barrier or the width, so long as the width remains much less than ξ) [140]. Note that this implies $v_c > 0$ even for barrier heights $V_0 > \mu$.

In the case of wide obstacles, the density is smoothly reduced and we can introduce local state variables, $n(x)$, $\mu(x)$ and $c(x) = \sqrt{\mu(x)/n}$. In addition to changing the density and hence speed of sound, the presence of the obstacle modifies the velocity field of the fluid, increasing it up to a maximum value of twice the unperturbed value at the walls of the obstacle [139].

This suggests that we can apply the Landau criterion locally, inferring that excitations occur when the ratio $\alpha = v(x)/c(x)$ exceeds some critical fraction α_c at any location x (and Landau’s prediction is that $\alpha_c = 1$). In a system with such an obstacle, as the height of the barrier grows the density is further reduced. We therefore expect v_c , the velocity of the fluid far from the obstacle required to exceed this critical fraction at some position, to decrease as the height of the barrier V_0 increases. Indeed, in Ref. [140] the authors predict the critical velocity of a soft-walled object tends to zero as V_0 approaches μ , where the density vanishes.

This model is an extension of Landau’s criterion applied locally, in that it still predicts that the local sound speed is the critical velocity required to produce Bogoliubov excitations. If other excitations are possible, then the critical fraction can differ from $\alpha_c = 1$. In practice, Landau’s theory often overestimates the critical fraction. This is primarily because it is a perturbative treatment, which ignores the possibility of other non-linear excitations such as vortices. In practice, it is possible to create vortex excitations at lower velocities than this perturbative treatment would predict.

Feynman argued that a rough estimate for a critical velocity induced by vortex excitations within a channel of width D [141] is given by

$$v_c = \frac{\hbar}{MD} \ln\left(\frac{D}{\xi}\right). \quad (6.7)$$

In our experiments with a single obstacle, this would produce a critical velocity of $v_c \sim 0.15c$ which is typical of such estimates, but it can in principle approach zero for sufficiently large D . In experimental and numerical studies, the critical velocity is generally found to be slightly higher than this, since Feynman's expression is an argument based on energy scales without considering any dynamical mechanism for generating the vortices. The generation of vortex pairs around a circular obstacle as a mechanism of dissipation has been confirmed numerically [139] and experimentally [138].

In general, the critical velocity of a system depends on the size, height and shape of an obstacle and is typically hard to calculate analytically. Estimates of v_c vary depending on the system and obstacle being considered. Numerical studies have found, for example, $v_c \approx 0.4c$ for a hard-walled cylinder [139], and $v_c \approx 0.56c$ for a narrow Gaussian obstacle of width comparable to the healing length [142]. Experimental investigations have found values of $0.1c$ [138] and $0.25c$ [137] in 3D, and $0.3c - 0.6c$ in 2D [136]. These experimental investigations all used optical dipole potentials with a Gaussian profile wider than the healing length, in a similar manner to the work presented here.

For dipole oscillations of a ^7Li BEC in a 3D harmonic trap, the critical velocity was experimentally found to exhibit a power law dependence on the amplitude of the obstacle potential [143]:

$$\frac{v_c}{c} = \left(1 - \frac{V_0}{\mu}\right)^{5/2} \quad (6.8)$$

as was also predicted in [77]. These studies and others, both theoretical [140, 144] and experimental [145], found that at velocities significantly higher than the critical velocity, dissipation (and damping of the dipole oscillations) disappeared. This non-dissipative supersonic regime is due to the kinetic energy of the flow dominating other energy scales, including the interaction energy, and the rate of excitations produced drops almost to zero.

6.1.3 Excitations above the critical velocity

A number of studies have attempted to predict not only the value of the critical velocity, but the properties of the excitations produced at supercritical speeds.

Frisch *et al.* numerically predicted that the rate of vortex generations, and hence drag force on an obstacle, was proportional to $v - v_c$ [139]. This linear relation was used in Refs. [136, 137] to model the heating rates observed above the critical velocity.

However, more detailed studies have found that the rate of excitations produced shows a more complicated dependence on the velocity. Dries *et al.* [143] found that the damping rate of dipole oscillations in the presence of a Gaussian obstacle peaked in magnitude at approximately c before dropping again at higher velocities. As mentioned in the previous section, at sufficiently high velocities (several times higher than c) the flow transitioned to a supersonic non-dissipative regime as predicted [77, 146], with no observable damping. They also found that the critical velocities observed are consistent with Landau's theory when applied locally, implying that the damping occurs primarily due to Bogoliubov excitations and not, for example, vortex excitations.

Paris-Mandoki *et al.* [144] also found that the generated excitations peak in number between $1.5 - 2.0 v_c$ before dropping again. However, they predict an additional resonance effect. The scattering from a square repulsive potential in 1D shows resonances in the incident velocity where the transmission coefficient increases to exactly 1, referred to as Ramsauer-Townsend resonances. It was found that the same effect can occur in weakly interacting gases flowing past an obstacle in both 1D and 2D. These resonances produce specific velocities, greater than the critical velocity, where no excitations are produced and their resultant drag force vanishes entirely.

6.1.4 Superfluid flow in disordered potentials

The previous sections were concerned with superfluid flow around a single defect and the onset of dissipation above a critical velocity. We now consider the related but qualitatively distinct case of superfluid flow through a disordered potential.

Comparatively less investigation of such systems has been performed. Theoretical and numerical work by Albert *et al.* [147] on a system with slow-varying disorder potential with Gaussian correlations, such as a speckle pattern, found that the critical velocity was not a constant value but a random variable. The distribution of v_c is dependent on the strength of the disorder and the ratio of the system size to the correlation length. The dependence on the system size is introduced since a disorder potential of this kind has a randomly varying peak potential depth, and so larger systems have a higher chance of featuring stronger disorder amplitudes, limiting the critical velocity. They predict typical critical velocities of $0.3 - 0.6c$. A disorder

pattern of many independent but narrow peaks produce qualitatively similar results, but with a generally higher critical velocity.

In studying a speckle potential, Dries *et al.* [143] observed no critical velocity below which non-dissipative flow was observed. They attribute this to the presence of random defects even in the low density edge of the cloud, which was confined in a harmonic trap. They also observed a velocity-dependent damping rate. It was found that this damping rate was small for both low and high velocities, in similarity to the case of a single obstacle. In addition they found that the damping rate followed a power-law dependence on the amplitude of the disorder, where the exponent p exhibits a weak dependence on the velocity and has a mean value of $p = 2.1(5)$.

Work on speckle potentials such as those described above may differ significantly from our work which utilises pointlike disorder (although the obstacles are still wide in that the width of the scatterers is a few times larger than ξ). A more similar potential, albeit with smaller obstacle widths comparable to ξ , was numerically studied by Doran *et al.* [142]. They found that the critical velocity of a pair of obstacles depended on the relative angle between the flow velocity and the separation of the obstacles, as well as the magnitude of their separation. Typical critical velocities found were in the range of $0.4 - 0.65c$. Importantly, for large separations the critical velocity tended towards that of a single defect. In generalising this result to the case of many randomly placed obstacles, the critical velocity of the system was limited by the most closely-separated pair of obstacles. The scattering points were placed at least a certain minimum distance from each other, which is not the case in the experimental work we present here. As a result we again expect that the critical velocity for a disorder potential, if it exists, will depend on the random microscopic configuration of the disorder.

In all the studies mentioned previously, localisation is not expected to play a significant role. It has been predicted that Anderson localization could be observed in 1D pointlike disordered systems at fluid velocities significantly higher than the speed of sound ($\gtrsim 10c$) and in large systems [148]. Observation is difficult due to the long localisation length; the authors predict localisation may be observable as an exponential reduction in transmission for system sizes $L > 500\xi$. They also predict that interactions do not meaningfully change these results.

6.2 Experimental procedure

Here we describe the experimental sequence performed in order to produce a 2D cloud and observe the superfluid flow around obstacles of different forms.

6.2.1 Preparing the 2D cloud

As described in Chapter 3, we initially produce a Bose-Einstein condensate in a circularly polarised shell trap below the focus of the dipole potential. At the Rabi frequencies used for loading the atoms into the shell trap, $\Omega = 2\pi \times 180$ Hz, the cloud is initially a 3D BEC. To produce a 2D quantum gas we ramp the amplitude of one dressing field, and hence Rabi frequency, down to a final value of $\Omega = 2\pi \times 60$ Hz while the other field amplitude is reduced to zero. This, combined with increasing the quadrupole current to 147 A (a radial quadrupole gradient of 171 G/cm), produces a 2D condensate in a linearly polarised shell trap at the focus of the dipole potential.

For the experiments here, we produce a cloud of approximately 10^5 atoms at a temperature of 40 - 50 nK. The radial and axial trap frequencies are $\omega_r = 2\pi \times 14.5$ Hz and $\omega_z = 2\pi \times 1100$ Hz, respectively. This produces a cloud with a chemical potential of $\mu/h \approx 900$ Hz, so that the values of the standard criteria for 2D physics are $\mu/\hbar\omega_z \approx 0.8$ and $k_B T/\hbar\omega_z \approx 0.94$, comfortably within the 2D regime [128].

The flow of the cloud around the obstacles at high dipole potential strength and high flow velocities causes heating of the cloud as excitations are produced by the supercritical flow, up to temperatures of ~ 80 nK. This is still within the 2D condition and corresponds to a $T/T_0 = 0.47$, where T_0 is the critical temperature for the BEC transition in a non-interacting gas in a 2D harmonic trap. This temperature is also below the BKT transition temperature in the absence of disorder, which is $T_{BKT}/T_0 = 0.53$ [26], so that all the clouds are expected to be in the superfluid regime. In the case of strong disorder ($V/\mu > 1$), the transition temperature appears to be reduced to approximately $T_{BKT}/T_0 = 0.47$, as seen in Chapter 5. For the work presented here we study only comparatively weak disorder strengths ($V/\mu < 0.2$) where the cloud has been observed to remain in the superfluid phase at these temperatures.

We wish to study the flow of the superfluid around an obstacle by measuring dipole oscillations of the cloud, so any residual collective motion from the loading sequence will interfere with the results. To eliminate or reduce this motion we use the same protocol as in Chapter 5, switching on the dipole potential after loading the linear shell to damp any remaining motion. This damping is the exact effect we intend to study in this chapter; but we here apply some initial damping so that

we start our experiments from a stationary cloud and later induce intentional and controllable oscillations.

To damp the motion we smoothly ramp up the dipole potential from zero up to a value $V/\mu \sim 0.2 - 0.4$ over the course of 150 ms and then down to zero again over another 150 ms. At present the experimental control system does not allow the dipole potential shape to be modified mid-sequence, so we use the same shape for both the initial damping and the later experiment. Pointlike disorder is highly effective at eliminating any residual motion and so lower disorder strengths are used. For single obstacles the damping force is weaker and slightly higher potential strengths are used. In practice, the residual motion from loading is small, and while this effect is large enough to affect measurements of the phase profile in a matter-wave interference experiment, the velocities we impart to the cloud here are usually significantly greater. Any residual motion therefore is a small effect, though it contributes to increased variation in position for small oscillation amplitudes.

6.2.2 Inducing flow around an obstacle

Beginning with a stationary 2D cloud, we smoothly ramp up the potential being tested to its final value over 200 ms and then hold the cloud in the potential for 100 ms to allow the system to relax.

In some superfluid flow experiments the obstacle is moved through the stationary cloud [136, 137], but since our dipole potential is static over the course of a sequence, we instead force the condensate to flow past the obstacle. We impart a sudden velocity kick to the atoms by briefly translating the centre of the harmonic trap in the transverse x - y plane before returning it to its original position. This is achieved with the use of a current-carrying coil that produces a constant magnetic bias field in the y direction (horizontally and perpendicular to the imaging axis). We switch the bias field on and suddenly translate the centre of the cloud between 4-30 μm from its original position, and then smoothly ramp it back to its original value over 10 ms. This ramp profile was found to produce the cleanest velocity kick with minimal heating or atom loss from the sudden change in potential.

The initial velocity for a range of bias currents was measured by the amplitude of undamped dipole oscillations in time-of-flight. We test speeds up to $\sim 0.6c$, where

$$c = \sqrt{\mu/m} \tag{6.9}$$

is the speed of sound at the centre of the cloud in the unperturbed fluid. In our experiments here, $c \approx 2 \mu\text{m/ms}$. This was the strongest velocity kick that could be imparted to the cloud without significant heating.

We test flow around two different kinds of obstacles: a single Gaussian barrier, and a pointlike disordered potential. The single Gaussian barrier is produced by a single circular pattern on the digital micromirror device (DMD) which is blurred by the finite point spread function of the optical system into a Gaussian barrier. This single obstacle is positioned at the centre of the initially static cloud. The pointlike disorder potential is composed of many such Gaussian barriers randomly and uniformly distributed throughout space (see Section 5.1.1). For both the single obstacle and disorder we use barriers with a Gaussian width of $\sigma = 2.7 \mu\text{m}$. The healing length at the centre of the cloud is approximately $\xi = 0.6 \mu\text{m}$ and varies spatially due to the inhomogeneous density profile, leading to increased healing length towards the edge of the cloud. This implies that the ratio of the barrier size to the healing length ξ is approximately $\sigma/\xi \approx 4$. As a result we occupy an intermediate regime where the obstacle is larger, but not much greater, than the healing length.

6.2.3 Measurement of damping rate

We observe the motion induced by the velocity kick as oscillations in the radial position of the cloud after time-of-flight (TOF). To prevent any additional momentum kick from switching off the magnetic trapping fields, we turn off the dressing field and dipole potential before the magnetic trap, projecting the atoms into the undressed Zeeman states. The $m_F = 0$ state is magnetically insensitive and drops ballistically under gravity.

We allow each condensate to oscillate for varying periods of time, measured relative to the end of the velocity kick. We fit the radial position after 15.5 ms TOF with the formula

$$x = A e^{-\Gamma t} \cos(2\pi f t + \phi), \quad (6.10)$$

where A, Γ, f and ϕ are the fit parameters. A is the amplitude of the oscillations after TOF. Because the cloud drops freely, its position $x_{TOF}(t)$ measured in TOF after t seconds hold duration in the trap is given by $x_{TOF}(t) = x_0(t) + v(t)t_{TOF}$ where $x_0(t)$ is the position of the cloud *in situ*. The velocity varies sinusoidally in time and $\pi/2$ out of phase with the position so that their sum in time-of-flight gives the position

$$x_{TOF}(t) = x_0(t) \sqrt{1 + t_{TOF}^2 \omega_r^2}. \quad (6.11)$$

The amplitude in situ A_0 is then

$$A_0 = \frac{A}{\sqrt{1 + t_{TOF}^2 \omega_r^2}}, \quad (6.12)$$

and the maximum (initial) velocity of the cloud is $v = A_0 \omega_r$.

Our primary parameter of investigation is the damping rate Γ . The damping of the oscillations is caused by excitations produced within the superfluid by supercritical flow. The damping rate therefore indicates whether the flow velocity exceeds the critical velocity for the fluid, and acts as a relative measure of the rate of excitations produced.

Since we expect the damping rate to depend on the velocity of the flow, and for the maximum flow velocity to drop over time as the excitations generated dissipate energy, in principle the damping rate should be velocity (and hence time) dependent. Ref. [143] observed this effect experimentally, by dividing the full hold duration into many intervals of a few periods each and treating the damping rate as approximately constant within each interval. They observed that the damping rate had a non-linear dependence on the velocity with a peak at approximately the speed of sound. In this work we model the damping rate Γ as constant over the oscillation, because we find that this simpler model still provides a good fit to the observed oscillations.

The damping rate of oscillations in the trap in the absence of any obstacle is very low. No damping is observable over the course of 1.6s, significantly longer than the durations tested here.

6.3 Superfluid flow around an obstacle

Here we present preliminary measurements of the damping caused by flow around a single obstacle. We measure the position of the cloud at times up to 125 ms and again at times between 400 - 550 ms to allow an improved estimate of the fitting parameters, since the damping rates in this case are relatively low.

Fig. 6.1 shows an example of damped oscillations in the position of the condensate after time-of-flight. The obstacle here is a single Gaussian barrier with peak potential height of $V/\mu = 0.7$. We expect this obstacle to have a low critical velocity due to its relatively high barrier height. For example, using Eqn. 6.8 we can estimate that the critical velocity should be roughly $0.05c$. The fitted initial flow velocity is $v_0/c = 0.28$, measured in units of the speed of sound at the centre of the trap $c = 2\mu\text{m/ms}$. Damping is clearly visible and we infer that it is caused by excitations produced by the flow dissipating energy, indicating that this v_0 is above the critical velocity for




Figure 6.1: Example of damped oscillations around a single Gaussian barrier of height $V/\mu = 0.7$ after an initial velocity kick. The initial velocity is calculated based on the fitted initial amplitude of the oscillation, which models the motion as exponentially decaying with a constant damping rate. The velocity is reported in units of the speed of sound c defined in Eqn. 6.9.

this cloud if such a threshold exists. The fitted damping rate is $\Gamma = 1.5(2) \text{ s}^{-1}$. From the fitted amplitude $A = 10.7 \mu\text{m}$ in TOF we infer an *in situ* oscillation amplitude of $A_0 = 6.3 \mu\text{m}$. This small compared to the Thomas-Fermi radius of $28 \mu\text{m}$, so the obstacle remains within the condensate region for the full duration of the oscillation.

6.3.1 Observation of critical velocity

To identify a possible critical velocity, defined as a threshold above which the damping rate significantly increases, we sweep the initial flow velocity around the obstacle for different values of the obstacle height $V/\mu = 0.4, 0.6, 0.7$ and extract a damping rate at each flow velocity. We expect that since the higher barriers reduce the local density more, they will have a lower critical velocity.




Figure 6.2: a) Fitted damping rate Γ as a function of initial flow velocity for a single barrier of height $V/\mu = 0.4$. Error bars show 1σ confidence intervals on the fitted value. Lower velocities have larger error bars because they produce smaller oscillations where random variation in the cloud position contributes to the uncertainty. We see an apparent increase in the damping rate for $v_0/c > 0.3$. The midpoint between the two points nearest this increase is identified as a critical velocity v_c/c , indicated by the blue dashed line. Values of $\Gamma < 0.2 \text{ s}^{-1}$ are consistent with no damping, indicated by the red shaded region. b) The damping rate increases for higher flow velocities but also shows a unexpected sharp peak at $v_0/c \approx 0.1$, and drops to zero again for $v_0/c = 0.17$. Again the critical velocity is identified as the midpoint between the absence and onset of damping. The point of zero damping at $v_0/c = 0.17$ is unexpected and may be related to a resonance effect, and we do not associate it with being below a critical velocity.

Fig. 6.2a shows the damping rate for a barrier height of $V/\mu = 0.4$. Error bars are 1σ confidence intervals on the fit parameter. Lower velocities typically have larger errors because the amplitude of the oscillations are much smaller, the equivalent of a few pixels on the absorption image of the cloud from which position is extracted. There is a small random variation of the measured cloud position of $0.5 - 1 \mu\text{m}$ presumably due to some remaining initial residual motion of the cloud after loading. A damping rate of less than $\sim 0.2 \text{ s}^{-1}$ is consistent with zero damping for the hold durations observed. This is because, for low to moderate oscillation amplitudes, the resulting change in amplitude for this damping rate is less than the random variation in the cloud position.

We see in this case that the damping rates do appear to be either absent or so




Figure 6.3: a) Oscillations for $v_0 < v_c$, which show no evidence of damping. b) Damped oscillations for $v_0 > v_c$, with $\Gamma = 0.53(12) \text{ s}^{-1}$.

small that they are undetectable for $v_0/c \lesssim 0.3$. We estimate v_0/c to be the midpoint between the two velocities either side of the sudden damping rate increase, and the error half the difference between them. This provides a rough estimate of the critical velocity for this barrier as $v_c/c < 0.31(5)$. Example plots for the undamped and damped oscillations below and above this apparent threshold are shown in Fig. 6.3. However, the damping rates remain small in absolute terms.

Fig. 6.2b shows the damping rate as a function of velocity for $V/\mu = 0.6$. We again see an increasing damping rate for higher velocities but we also observe an initial narrow peak in the damping rate at a velocity $v_0/c = 0.13$, above which the damping rate decreases to zero again at $v_0/c = 0.17$.

One possible explanation for the vanishing damping rate is that of resonances in the transmission through the barrier [144]. The authors of Ref. [144] predict that the rate of excitations produced is maximal at values slightly higher than the critical velocity, which broadly agrees with the peak in damping in Fig. 6.2b. Above this point, the wavevector of the flow periodically becomes resonant with the barrier width, with the result that the transmission is perfect and the rate of excitations produced vanishes. We do not expect our results here to quantitatively match that of Ref. [144], which considers a one-dimensional hard walled barrier, but it is possible a similar kind of resonant behaviour may be at play. The wavevector of the flow at velocity v_0 is $k = mv_0/\hbar$ and we may roughly expect a resonance to occur when $k\sigma \sim 1$




Figure 6.4: a) Damping rate as a function of initial flow velocity around a single Gaussian barrier of height $V/\mu = 0.7$. The red shaded region indicates values of Γ consistent with the absence of damping. The lowest velocity tested shows no visible damping and the estimated critical velocity is indicated by a blue dashed line. b) Estimated critical velocity v_c as a function of the barrier height. Vertical error bars are the interval within which the critical velocity is estimated to lie; horizontal error bars are 1σ uncertainty in the measured value of V and variation in μ . Red line shows the fit for an algebraic scaling, with exponent $\beta = 2.29(13)$, as described in the text.

where σ is the width of the barrier. For our Gaussian barrier of $\sigma = 2.7 \mu\text{m}$ and a velocity of $v_0/c = 0.17$, $k\sigma = 1.2$. It seems possible then that a similar resonance effect is responsible for the absence of damping at $v_0/c = 0.17$.

If we identify a critical velocity at the transition from no identifiable damping to observable damping, we can plot an estimate of the critical velocity v_c/c against the barrier height V/μ . This estimate is shown in Fig. 6.4b. Also shown is an algebraic fit to the data of the form

$$\frac{v_c}{c} = \left(1 - \frac{V}{\mu}\right)^\beta, \quad (6.13)$$

where β is the fitted parameter. Algebraic dependence of this kind has been predicted [77, 140] and observed [143] for 1D and 3D gases with exponents $3/2$ and $5/2$, respectively. Here we fit a value of $\beta = 2.29(13)$, close to but slightly lower than the 3D case. We estimate the error in the algebraic fit by iterating over many fit attempts ($N = 200$) with randomly distributed values of the critical velocities and barrier heights within their respective error estimates. We then plot the mean value of β and quote its error as one standard deviation from the mean.

The value of $5/2$ is derived for the case of applying Landau's theory locally with a critical fraction $\alpha_c = 1$, so the near agreement with this model may indicate that this is the primary form of dissipation. However, further data points are required to improve the estimate of β and the confidence in the fit model. For example, Feynman's estimate of a critical velocity for vortex excitation (Eqn. 6.7) gives an estimate of $v_c/c \sim 0.15$. This is broadly consistent with the observed $v_c \approx 0.1$ for higher barrier heights. Further investigation may be able to clarify the dissipation mechanisms at play.

6.3.2 Frequency shift of oscillation

We also attempt to observe a frequency shift in the oscillation. In general, the frequency of damped harmonic oscillation is not equal to the oscillation frequency of the undamped system, which here is the radial trap frequency measured to be $\omega_r = 2\pi \times 14.46(3)$ Hz.

A small frequency shift occurs in the case of a damping force that is linear in the velocity, $F_d = -kv$. Such a system with an undamped natural frequency of ω_0 exhibits motion of the form

$$x(t) = Ae^{-\gamma t} \cos(\omega' t + \phi), \quad (6.14)$$

where the frequency $\omega' = \sqrt{\omega_0^2 - \gamma^2}$ and γ is the damping rate. The frequency shift is then dependent on the damping force. This model was used in fitting the oscillation frequencies in previous experimental work [143]. However, a small negative frequency shift for low barrier heights and oscillation amplitudes much smaller than the Thomas-Fermi radius has been analytically calculated in 1D Bose gases [77], and this value is different to the prediction of Eqn. 6.14. The prediction for the frequency shift is

$$\delta\omega = -\frac{1}{2m\omega_r} \int_{-\infty}^{\infty} dx \frac{dn(x)}{dx} \frac{dV(x)}{dx}, \quad (6.15)$$

where $n(x)$ is the unperturbed density at the point x , and $V(x)$ is the profile of the barrier (in this case, Gaussian). For the case of a Gaussian barrier, this Eqn. 6.15 becomes

$$\delta\omega = -\frac{3V\sigma}{8\mu^{3/2}} \omega_r^2 \sqrt{\pi m}. \quad (6.16)$$

However, this value is derived from a perturbative treatment in 1D and so is unlikely to be quantitatively accurate in our case where $V/\mu \sim 1$ and the cloud is two-dimensional.




Figure 6.5: The percentage deviation in the frequency of the oscillations from the radial trap frequency (black dashed line), which is the observed frequency in the absence of damping. The blue shaded region indicates the error in the measured radial trap frequency. We observe a small but significant negative frequency shift in almost all cases. Higher barrier heights cause a more significant shift at lower velocities, but at high velocities the shift appears to converge a similar value independent of the barrier height. Frequency shifts are still observed at velocities below the critical velocity for each respective barrier height.

We leave the frequency of the oscillation as a free fit parameter and plot it as a function of the flow velocity for each barrier height, shown in Fig. 6.5. The presence of the barrier appears to shift the frequency of oscillation down by a small but measurable amount, in the range of 1-3%, with greater barrier heights causing a larger shift. Due to the low damping rates the expected frequency shift from Eqn. 6.14 is extremely small, less than 0.02% even for the highest damping. This implies either that the change in frequency of the oscillation is not purely from a classical damping effect, or that the damping does not exhibit a linear dependence on the velocity. The observed shift is similar in magnitude to those predicted by Eqn. 6.16, and has the same sign (negative). This produces an expected frequency shift of 3 – 5% for the barrier heights used here. This is an overestimate over the shifts observed, likely due

to the fact that even the lowest barrier height is too large to be perturbative and Eqn. 6.16 applies to 1D systems. In the perturbative treatment there is also no dependence on the flow velocity since it is assumed to be very small.

Importantly, this frequency shift is expected to occur even in the superfluid regime where no damping occurs. We appear to observe this same effect, with a negative shift occurring even for points that display no visible damping. Beyond this, the frequency shift shows some dependence on the velocity; for low v_0/c , the barrier height causes a significant difference in the frequency of the oscillation, while the shifts for different barrier heights appear to converge to a similar value for higher flow velocities.

6.4 Damping in a disordered potential

We now consider the damping caused by flow through a disordered potential. In general we expect that the disorder provides much greater damping, for several reasons. Firstly, there is on average always some potential barrier present in low density parts of the cloud, where the local v/c is higher and therefore likely to exceed the local critical velocity. Secondly, pointlike disorder can produce multiple barriers close together which act like a larger obstacle and hence more easily produce excitations [142].

In a random potential both the damping rate and the critical velocity (if it exists) are random variables, so for these experiments we take 3 repeats at each data point to improve our confidence in fitting the oscillations.

6.4.1 Dependence on disorder strength

Fig. 6.6 shows the oscillation of the cloud in TOF for different peak disorder heights, $V/\mu = 0.13$ and 0.17 for approximately the same initial velocity $v_0/c \approx 0.35$. The disorder is pointlike with density $n_p = 7 \times 10^{-3} \mu\text{m}^{-2}$. It is evident that the damping rates are much greater than that of a single obstacle even at lower potential heights. This strong damping makes the disorder effective at eliminating the residual motion from loading the shell prior to the experimental sequences used here and in Chapter 5.

We measure the damping rate as a function of disorder strength for two values of velocity v_0/c . For disorder strengths $V/\mu \gtrsim 0.25$, the damping is very strong and produces overdamped motion (a similar value to that found in [84]). In addition, the disorder potential causes the centre of mass of the cloud to vary randomly with every




Figure 6.6: Examples of damped oscillation for motion through a pointlike disordered potential with $n_p = 7 \times 10^{-3} \mu\text{m}^{-2}$. The damping is much stronger than the case of a single Gaussian barrier even for lower peak potentials. Each single data point uses a different random configuration of the disorder.

disorder configuration (i.e at every data point while varying hold duration). This produces noisy fits where only the initial decay to equilibrium is clearly observable.

Additionally at high disorders the fitted velocity is significantly less than that expected for the same initial kick determined by the bias field current. This is likely due to the damping rate being so strong that it significantly affects the dynamics over the period of the kick (10 ms). This makes it difficult to compare data points like-for-like since they have different v_0/c , where we want to examine the dependence on V/μ at constant v_0 .

For the values of V/μ in which oscillation is clearly observable and damping rates reasonably fitted we plot the damping rate Γ as a function of disorder strength, as shown in Fig. 6.7. We fit the data to a power law as in [143], with

$$\Gamma \propto \left(\frac{V}{\mu}\right)^p, \quad (6.17)$$

where p is the fitted exponent which can in principle depend on the velocity v_0/c . We find that within the 1σ uncertainty in the fit the exponent $p(v_0/c)$ is very similar for each value of v_0/c , being $p(0.21) = 2.4(5)$ and $p(0.47) = 2.7(8)$. Ref. [143] found a weak dependence of p on the velocity with a mean value of $p = 2.1(5)$ (for $v_0/c \approx 0.4$, similar to the values used here, $p \approx 1.6(6)$) in 3D, while Ref. [84] found good agreement




Figure 6.7: Fitted damping rate Γ as a function of peak disorder strength V/μ for two different initial flow velocities v_0/c plotted on a log scale. Vertical error bars are 1σ uncertainties in the fitted damping rate, and horizontal error bars are the combined 1σ error in the estimate of the disorder strength and the variation in the chemical potential μ . The data are fitted with a power-law dependence with exponent $p(v_0/c)$ that depends on the velocity. The fitted exponents are $p(0.21) = 2.4(5)$ and $p(0.47) = 2.7(8)$.

with a quadratic fit for fermionic condensates, i.e. $p = 2$. The observed p in this 2D case appear slightly higher but are consistent with these values within the uncertainty of the fit.

6.4.2 Dependence on initial flow velocity

Fig. 6.8 shows the fitted damping rate as a function of the initial flow velocity for disordered potential with a low disorder strength, $V/\mu < 0.1$. For the range of velocities tested we see no evidence of a critical velocity below which the damping is negligible. This is consistent with the findings of [143], although a subsonic non-




Figure 6.8: a) Damping rate of oscillations in a disordered potential as a function of initial flow velocity for disorder strength $V/\mu = 0.04$. A peak similar to that in the case of a single barrier appears at low velocities. b) Damping rates for disorder strength $V/\mu = 0.09$. The damping rate increases significantly with flow velocity, and another small peak seems to appear at a similar velocity to the lower disorder strength case.

dissipative regime was predicted by [77]. It may be the case that a finite critical velocity exists at lower velocities than those tested, but it may be too low to be detected with the resolution or random variation present in our experiment. Broadly we see the damping rate increase with greater flow velocities. The exception is a narrow peak in the damping rate that appears at relatively low v for both values of V/μ .

This peak is unexpected and no similar feature is observed in other experimental or theoretical works on disordered superfluid flow. Further data is needed to improve the estimate of the damping rate and probe the shape of the peaks. Note that this peak resembles the similar feature that appears in the damping rate of a single obstacle of barrier height $V/\mu = 0.6$, as shown in Fig. 6.2b. The damping rate in Fig. 6.8a also drops to almost zero at $v_0/c \sim 0.2$, suggesting again a resonance in the transmission through the disorder, although no similar feature is clearly seen in Fig. 6.8b.

Finally, we briefly note that no significant shift in the oscillation frequency is observed for oscillations in the disordered potential. This is likely because disorder is capable of shifting the oscillation to both increased and decreased frequencies and so may cause no discernible shift on average [77, 149]. In measuring the position of the

cloud over time, from which we extract frequency and damping rate parameters, each data point represents a different random configuration of the disorder. Any shift in the frequency then manifests only as additional noise in the position of the cloud.

6.5 Conclusion

In this chapter we presented a preliminary investigation into the behaviour of a 2D superfluid performing sinusoidal oscillations around both a single obstacle and a point-like disordered potential. We found that in the case of a single obstacle the oscillation was damped and that this damping rate increased for higher initial flow velocities; with the exception that below a certain velocity, dependent on the height of the barrier, no damping was observed. We identify this as a critical velocity below which the superfluid exhibits frictionless flow, and find that this critical velocity is considerably lower than the speed of sound, as is usually found experimentally. We find that the scaling of the critical velocity with barrier height is consistent with behaviour in 3D in which Bogoliubov excitations are the primary form of dissipation and the critical velocity is determined by the local speed of sound. We can not however definitively exclude the production of other excitations without probing a wider range of velocities and barrier heights.

We further observe an unexpected narrow peak in the damping rate at certain barrier heights and velocities, for which we have no obvious explanation. A possible cause is a form of resonance between the wavevector of the flow velocity and the barrier width, though more investigation is needed to characterise the specific behaviour seen and explain it. We measured the frequency of the damped oscillations and find that they are shifted to below the radial trap frequency, and that this shift is significantly larger than that predicted by a classical damped harmonic oscillator.

In the case of flow through a disordered potential we found a much stronger damping effect even at peak potential heights much lower than that of a single obstacle. The damping rate was found to scale with the disorder strength as a power law with the same exponent for two different flow velocities, and this scaling appears to be consistent with values measured in other experimental setups in 3D BECs. We again see a peak in the damping rate at similar flow velocities to the case of a single obstacle and consider the same possibility of a resonance phenomenon. No average frequency shift is observed in the presence of disorder, since each experimental sweep comprises an average of many realisations of disorder which can shift the frequency in both positive and negative directions.

Further data may improve estimates of the damping rate, and in particular probe more velocities and dipole potential strengths to fully map the phase diagram of 2D superfluid flow around individual and random obstacles. In particular, it would be interesting to investigate further the unexpected peaks in the damping rate, which may provide evidence of a novel dissipation mechanism or interference effect. The combination of this flow around a barrier with the double-well potential and matter-wave interference scheme may also allow us to more directly observe some of the dissipative excitations produced, such as vortices.

Chapter 7

AtomECS: Simulating atomic collisions and magnetic trapping

AtomECS [150] is software package developed in our research group for simulating laser cooling and magneto-optical trapping of cold atoms. It adopts an Entity-Component-System (ECS) architecture implemented in the Rust programming language for a high degree of flexibility, scalability and performance. It has been tested and successfully applied to model 3D magneto-optical traps (MOTs), optical molasses, and 2D-MOT atomic sources. It is open source and licensed under GPLv3.

The first release of the package did not include several important physical effects, including atom-atom interactions and magnetic trapping of neutral atoms. In this chapter we will discuss the implementation of these effects and compare to analytic results to verify the accuracy of the simulations.

7.1 Overview of AtomECS

7.1.1 Entity-Component-System architecture

An advantage of the AtomECS package is its use of Entity-Component-System (ECS) architecture. This is a data-oriented design pattern originally popularised for game development; both the popular game engine Unity [151] and the new Unreal Engine 5 [152] utilise ECS architecture, along with a host of newer engines [153, 154].

The ECS pattern contains three main ingredients: Entities, Components, and Systems. Entities represent a ‘thing’, such as an atom or a harmonic trap in which the atom is held. The entity contains no data itself other than a unique identifier, with which other data is associated. A component holds a particular type of data associated with an entity; for example, an atom’s position and velocity, or the spring constant of a harmonic trap. Systems act on entities or their components, such as

changing the force on an atom or adding new components to entities. Systems process all the logic within the ECS pattern; components and entities are functionless.

In the ECS pattern, systems process data by querying or operating only on entities with specified components. This means the behaviour of an entity is determined by the components associated with it. This makes ECS an example of behaviour by composition, rather than inheritance; the behaviour of an entity is determined by data that it is composed of, rather than by inheriting traits from some superclass as is typically done in object oriented programming.

A simple example would be the aforementioned atom in a harmonic trap. To simulate this, one would create an entity that represents the atom. To this we add an **Atom** component, to label the nature of the entity. We also need our atom to have **Position**, **Velocity** and **Force** components, which the harmonic trap and integrator systems will read and change. We then create a harmonic trap entity, which is given the component **HarmonicTrap**, and this component contains the data about the harmonic trap such as the spring constant. This entity also has a **Position** component to describe where the centre of the trap is located.

A system, **ApplyHarmonicTrapForceSystem**, is then run which operates on all entities with the **HarmonicTrap** component, and for each harmonic trap the system then operates on entities with the **Atom** and **Position** components. Based on the data in the components, this system calculates the force exerted on the atom, and adds this to the entity's **Force** component. The integrator systems then calculate the new position and velocity of the atom based on the **Position**, **Velocity** and **Force** components, and updates the data stored in these components.

One great advantage of the ECS architecture is that it is necessarily highly modular and loosely coupled - all functionality is decoupled which makes unit testing and development efficient and easy. Another is that since systems in ECS act only on components, they can be agnostic about what entities those components belong to. This produces a flexibility in code beyond the initial scope. For example, a **Velocity** component can be added to any entity (such as a trap or laser beam) to allow it to move during the simulation, even if the entity had been originally conceived to be stationary.

Furthermore, entities can be stored entirely separately from the data associated with those entities. The **Position** component, for example, acts as a container for all **Position** data so that this data may be stored in a contiguous block, rather than scattered throughout memory. This is well suited to modern CPU design and provides a boost to performance, since all the relevant data can be shuffled in and out of the

memory cache of the processor as needed, as opposed to the randomly allocated heaps typically used in other design patterns.

ECS also lends itself to easier parallelisation for high performance, since systems explicitly state the data they require so that dependencies are easily resolvable. This makes it straightforward to allow different systems to run on different component storages independently and simultaneously without conflicts (‘outer parallelism’). In addition, each system can itself often run in parallel over the entities it operates on (‘inner parallelism’). Rust provides a high performance language with the required degree of memory control in which to implement ECS.

7.2 Direct Simulation Monte Carlo method

The current release of AtomECS [150] does not include atom-atom interactions, but these are important in many cold-atom experiments. Although Bose-Einstein condensates (BECs) arise from the statistical mechanics of the many-body system, collisions between atoms are essential for thermalising the cloud during evaporative cooling. This is the most common method of cooling below the limit imposed by laser cooling techniques. To simulate the loading of our traps and subsequent evaporative cooling stages, it is therefore required that we have an accurate implementation of atomic collisions.

The simulation of collisions is a necessary problem to solve for anyone wishing to understand the flow of gases on a microscopic level. In the continuum limit, where the density of molecules is so high that it can be treated as a continuous medium, this flow is modelled macroscopically by the Navier-Stokes equations. However at lower densities or collision rates, this continuum model breaks down. The validity of the continuum description depends on the Knudsen number

$$\text{Kn} = \frac{\lambda}{L}, \quad (7.1)$$

where λ is the average distance travelled with no collisions, called the mean free path, of an atom, and L is a characteristic length scale of the flow. The Navier-Stokes equations are generally considered to break down at $\text{Kn} > 0.1$ [155]. In cold atoms experiments, Kn is generally larger than this even at comparatively high densities. Therefore a microscopic simulation of the atomic collisions is required to model the dynamics of the gas.

A valid first approach to simulating particle collisions would be to deterministically simulate every particle, and test at each step of the simulation whether two particles

have approached close enough to perform a collision. Indeed, this was among the first methods used to study elastic collisions between hard spheres [156]. Two problems with this method are that first, its computation time scales with the square of the particle number; and secondly, it is only accurate and practical for dense gases [157].

Another solution is the Direct Simulation Monte Carlo (DSMC) approach for gas flow, developed by G. A. Bird [157]. DSMC has become a widely used tool for modelling fluid flow outside of the continuum regime, and has been applied to subjects such as the flow of rarefied gases [158], aerodynamics of spacecraft [159], and the collisions of ultra-cold gaseous clouds [160–162], among many others.

DSMC generally refers to any method that models collisions of particles on a probabilistic basis, in which the dynamical motion and interatomic collisions are decoupled within a simulation step. This approximation is valid if the timescale of a step is much less than the mean time between collisions for an individual atom. In addition, it is common to divide the simulation space into a grid of collision cells within which the collisions occur. This ensures that only particles near each other (i.e. within the width of one cell) can collide. Real atoms are replaced with simulated particles which usually represent a large number of identical particles. In this way, the simulation is a statistical sampling of the true atomic distribution. Sometimes the reverse is true; that many simulated particles represent one real particle, to supersample the true probability distribution, useful in highly non-equilibrium use cases [160]. It has been proven that in the limit of small timesteps, small cells, and large simulated particle numbers, the DSMC method converges to the Boltzmann equation [163]. The probabilistic modelling of collisions is based on classical kinetic theory, like the Boltzmann equation, and is therefore subject to the same limitations; namely, the assumption of molecular chaos and dilute gases. These assumptions are, respectively, that the velocities of the particles are independent of each other and of their positions, and that only a very small fraction of space is occupied by a particle.

Since DSMC is a probabilistic method, it can display statistical fluctuations. In addition, since it represents an extremely large number of atoms with a small number of representative particles, the statistical noise is much greater than in the real gas. This noise is Poissonian in nature, and therefore is of the order $1/\sqrt{N_P}$ where N_P is the number of simulated particles in a collision cell. This implies that for few particles per cell, the fluctuations may be very significant, but when averaged over a sufficiently long timescale (or over sufficiently many identical cells) the method produces accurate collision rates.

7.2.1 Kinetic theory of collisions

For the purposes of modelling collisions in a cold gas, we will henceforth assume that the collisions are entirely elastic and we consider a single atomic species and do not include multiple internal states of the atom. In DSMC, we also make the assumption of a dilute gas, in which the collisions will be dominated by two-body processes. In many cold atoms experiments this is often a good approximation. In ^{87}Rb , used in our experiments, this approximation is very good.

Loss rates due to two-body inelastic collisions depend on the internal states of the atom, but in the commonly used case of the $F = 1, m_F = -1$ state of ^{87}Rb are negligible in comparison to the background $1/e$ lifetime of the atoms. Another common trapped state of ^{87}Rb is the $F = 2, m_F = 1$ state, which has been measured to have a two-body loss rate of $\gamma \sim 10^{-13} \text{ cm}^3\text{s}^{-1}$ [164], leading to a small or negligible loss rate for densities up to the condensed regime ($n \lesssim 10^{12} \text{ cm}^{-3}$). For some species, such as ^{133}Cs , two-body inelastic processes are much more significant and may not be negligible. Three body processes are also generally neglected except in dense Bose-Einstein condensates, where they are typically treated only as a loss mechanism.

7.2.1.1 Partial-wave analysis of scattering

To apply DSMC we must model the inter-atomic collisions probabilistically. It is evident that the probability of a collision between two atoms depends in some sense on the ‘size’ of the atom. It is not immediately obvious how to define this quantity; atoms are not objects of a fixed shape and volume. Rather, scattering occurs due to a distance-dependent interaction between two closely approaching particles. A general treatment of scattering is complicated, but a simplifying approximation can be made at low energies. This is achieved by expanding the scattered wavefunction in terms of a linear combination of waves with different angular momenta, known as partial wave analysis [165].

The problem of two-body scattering is most effectively solved by restating the degrees of freedom of each particle ($\mathbf{p}_1, \mathbf{r}_1, \mathbf{p}_2, \mathbf{r}_2$) as the motion of the system’s centre of mass ($\mathbf{p}_{CM} = \mathbf{p}_1 + \mathbf{p}_2, \mathbf{r}_{CM} = \mathbf{r}_1 + \mathbf{r}_2$), and the relative motion between the two particles ($\mathbf{r} = \mathbf{r}_1 - \mathbf{r}_2, \mathbf{p} = (\mathbf{p}_1 - \mathbf{p}_2)/2$) [165]. The centre of mass of the system moves like a free particle, so that all the scattering effects of interest are contained in the relative motion. The relative momentum and position can be modelled as a ‘relative particle’ with a mass equal to the reduced mass of the system, which travels through the potential $V(\mathbf{r}_1 - \mathbf{r}_2) = V(r)$.

The cross section σ for scattering off a spherically symmetric potential is given by

$$\sigma(k) = \sum_{l=0}^{\infty} \frac{4\pi}{k^2} (2l+1) \sin^2 \delta_l(k), \quad (7.2)$$

where $k = 2\pi/\lambda$ is the wavevector of the incident wave and δ_l is the phase shift imposed on each partial wave by the potential. It can be shown that as $k \rightarrow 0$, $\delta_l \rightarrow k^{2l+1}$, so that at low k only the $l = 0$ term contributes, equivalent to the scattered wave with no angular momentum [166]. This isotropic reflection is referred to as s-wave scattering.

The s-wave scattering cross section is then

$$\sigma = \lim_{k \rightarrow 0} \frac{4\pi}{k^2} \sin^2 \delta_0(k) = 4\pi a^2, \quad (7.3)$$

where a is the s-wave scattering length

$$a = -\lim_{k \rightarrow 0} \frac{\tan \delta_0(k)}{k}. \quad (7.4)$$

Note that this cross section at low energies is independent of momentum. It is important to note that as long as the gas remains dilute ($na^3 \ll 1$, where n is the atomic density), the collisional dynamics of the cloud depend only on the scattering length, not the microscopic form of the scattering potential [165]. Considering a as a characteristic range of the potential, the condition for this model to be valid is that $ka \ll 1$. ^{87}Rb has a measured s-wave scattering length of $\sim 100a_0$ [164], where a_0 is the Bohr radius, so that this condition corresponds to a temperature much less than 3 mK. Almost all cold-atom experiments operate at temperatures less than this; for example, the laser cooling Doppler limit of ^{87}Rb is $\sim 140 \mu\text{K}$.

7.2.1.2 Scattering rate in a homogeneous gas

This scattering resembles a classical model where an atom is a hard sphere with a radius a . A collision between two spherical particles occurs when the path of one particle intersects another. This occurs if a particle comes within a distance of $2a$ of another; so that we can consider a single particle sweeping out a volume of $\pi(2a)^2 \Delta x = \sigma \Delta x$ as it travels a distance Δx . The probability of a collision is the probability that any other particles have a position that lies within this volume.

Consider an atom moving through a gas of uniform density n at a speed v , and assume for now that all the other particles are static. The volume covered by the atom in a short time δt is $\sigma v \delta t$, and therefore the average number of atoms within

that volume is $n\sigma v\delta t$. For sufficiently short time δt , such that the average number of particles in the volume is much less than one, this average tends to the probability of a collision

$$P_{coll} = n\sigma v\delta t, \quad (7.5)$$

and the collision rate per particle is therefore

$$\nu_{coll} = n\sigma v. \quad (7.6)$$

However, our initial approximation that all other atoms in the gas are static is clearly incorrect. We instead should perform this analysis in the centre of momentum frame of the colliding particles, and replace the speed of one particle v by the average relative velocity of atoms, v_r .

Useful quantities to describe the collision dynamics include the mean free time and mean free path of a particle. The mean free time τ is the average amount of time taken for a particular particle to undergo one collision; that is, it is the time such that $n\sigma v_r\tau$ is equal to one. The mean free time is therefore

$$\tau = \frac{1}{n\sigma v_r} = \frac{1}{\nu_{coll}}. \quad (7.7)$$

The mean free path λ is similarly the average distance travelled by a particle before colliding, and is therefore easily given by

$$\lambda = \bar{v}\tau, \quad (7.8)$$

where $\bar{v}$ is the average speed, not the relative speed.

To apply this theory, we will need to know the average relative velocity of particles in a homogeneous gas. This quantity could be calculated directly, but the size of this calculation would grow with the number of pairs of atoms, which incurs an intensive $\mathcal{O}(N^2)$ complexity cost. We can instead straightforwardly calculate the average speed of particles within the gas, and relate this to the average relative velocity. To do this we must make an assumption about the velocity distribution within the gas. It is reasonable to assume that in many relevant cases the distribution will be approximately thermal. For example, in simulating evaporative cooling (which we discuss later in the chapter) the velocity distribution is approximately, but not exactly, thermal. Evaporative cooling truncates the velocity distribution by removing atoms with high energies. However, only a few atoms at the low probability tail of the distribution are removed. This does not significantly affect the relationship between the average velocity $\bar{v}$ and the average relative velocity v_r .

The distribution of speeds is obtained from integrating the Maxwell-Boltzmann velocity distribution over all solid angles to give

$$f(v) = \left(\frac{m}{2\pi k_B T}\right)^{3/2} 4\pi v^2 e^{-\frac{mv^2}{2k_B T}}, \quad (7.9)$$

where m is the mass of a particle, k_B is Boltzmann constant, and T is the temperature of the gas. The average speed $\bar{v}$ is then

$$\bar{v} = \int_0^\infty v f(v) dv = \sqrt{\frac{8k_B T}{\pi m}}. \quad (7.10)$$

The average relative velocity is similarly calculated by

$$v_r = \langle |\mathbf{v}_1 - \mathbf{v}_2| \rangle = \int d^3v_1 d^3v_2 |\mathbf{v}_1 - \mathbf{v}_2| f(\mathbf{v}_1) f(\mathbf{v}_2) = \frac{4}{\sqrt{\pi}} \sqrt{\frac{k_B T}{m}}, \quad (7.11)$$

and so for a thermal gas

$$v_r = \sqrt{2}\bar{v}. \quad (7.12)$$

The collision rate per atom therefore is

$$\nu_{coll} = \sqrt{2}n\sigma\bar{v}. \quad (7.13)$$

The total number of collisions per unit time in the gas then would appear to be $N\nu_{coll}$, where N is the number of particles. However, this would actually overestimate the collisions by a factor of two, since it would count the same collision twice, once from the perspective of each particle. The total number of collisions per unit time R_{coll} in the homogeneous gas is therefore [167]

$$R_{coll} = \frac{1}{\sqrt{2}} N n \sigma \bar{v}. \quad (7.14)$$

7.2.2 Implementation of collisions

We do not simulate every real atom, since it would be unnecessary and prohibitively slow. Instead we only require a statistically representative sample of atoms, and each simulated particle represents α real atoms, where α is the ‘macroparticle’ factor. This does not significantly modify our previous theory. For accurate physical behaviour we only require that the collision rate per particle is conserved with respect to changing α . The collision rate per particle after partitioning space into cells of volume V_{cell} is therefore

$$\nu_{coll} = \frac{\sqrt{2}\alpha N_p \sigma \bar{v}}{V_{cell}}, \quad (7.15)$$

where N_p is the number of simulated particles in the cell. The total number of collisions that occur in a cell in a given timestep of length Δt is

$$N_{coll} = \frac{1}{\sqrt{2}} \frac{\alpha N_p^2 \sigma \bar{v} \Delta t}{V_{cell}}. \quad (7.16)$$

7.2.2.1 Spatial partitioning

As described in the above discussion of DSMC, we first wish to divide the simulation space into a grid of cells within which the collisions will occur, and we model the particles inside the cell as a thermal, uniform gas. Since our kinetic theory is applicable to homogeneous gases, we need our cell size to be small relative to gradients in the macroscopic properties of the gas, such as density, so that each cell is approximately homogeneous. However, we also need the cell size to be large enough to contain a statistically representative number of particles. Too few particles in a cell causes large statistical fluctuations in the simulated number of collisions, or in the extreme case of only one particle per cell, no collisions can occur at all.

The simplest method of spatial partitioning is to use a fixed cell width and number, determined by the user, and divide space evenly according to these cubic cells. This is the fastest method since the partitioning is static and requires no additional overhead other than assigning the atoms to their cells at each time step. However, static partitioning deals poorly with cases where the density or width of the cloud changes over time, particularly if it changes by many orders of magnitude (such as evaporative cooling).

In the case of evaporative cooling, a partition cell size must be chosen that is smaller than the final width of the cooled cloud, so that our kinetic theory still applies. However at the beginning of the simulation, with the hot dilute cloud, this small cell size may contain extremely few particles on average. In some cases this therefore requires increasing the initial simulated particle number, and hence computation time; in others it is possible that no feasible set of parameters that meet both requirements can be found.

If the density of the cloud is expected to change significantly over time, the spatial partition must be rescaled at every time step. There are even more sophisticated methods of this rescaling that rescale the partition cells locally [160]; but we choose the simpler but still effective method of rescaling every cell width to maintain a similar particle density over the course of the simulation. An initial guess of the cell width constructs the first spatial partition. This partition is used to calculate the first frame of collisions, and also provides an estimate of the density distribution. The average

number of particles per cell is calculated and the cell width w is rescaled to achieve a target density ρ_0 , which is set by the user. To ensure that the cell width remains smaller than the cloud width, a maximum cell width w_{max} is imposed. w_{max} is defined as a fixed fraction of the full cloud width along its narrowest axis, so that the cell is always at most, for example 10% of the cloud size. From timestep n to $n + 1$ the cell width is rescaled as

$$w_{n+1} = \min\left\{w_n \left(\frac{\rho_0}{\rho_n}\right)^{1/3}, w_{max}\right\}. \quad (7.17)$$

This partition rescaling incurs a little additional computational cost; however, it allows certain simulations to be performed at a much smaller and more efficient total particle number, saving time.

7.2.2.2 Collisions algorithm

The basic algorithm for implementing particle collisions is:

1. Update the positions and velocities of all particles according to their evolution under all other dynamical forces. The essential assumption of DSMC is that dynamical and collisional effects can be decoupled. We therefore ignore collisions while calculating the effect of other forces such as photon absorption, or a magnetic force.
2. Create a grid of collision cells based on either a user-specified cell width and number for a static grid, or a dynamically scaled grid to achieve a target cell density.
3. Using the newly updated positions of the particles, assign each particle's velocity to a given key in a list of key-value pairs called a hashmap. The keys of this hashmap are the index of each collision cell, and the values associated with that key are the velocities of particles in that cell. Since DSMC does not care about the position of the particles, we do not need to store the particle's position or any other data about the particle; only its velocity is needed.
4. Once the hashmap is created, the collision cells are entirely independent. No particle interacts with a particle from any other cell, so we can iterate in parallel over each key of the hashmap (each collision cell), providing an increase in computational speed. Additionally, only collision cells containing particles are added to the hashmap, so we do not iterate over empty cells.
5. If there is only one particle in the cell, no collisions occur.

6. Calculate the average velocity of the particles in the cell.
7. Calculate the number of collisions that occur in the cell according to Eq. 7.16.
8. Randomly select N_{coll} pairs of atoms to collide, and mutate their velocities. The direction of travel after a collision is chosen randomly. This is done in the centre of mass frame to conserve energy and momentum and then transformed back to the lab frame.
9. If $N_{coll} < 1$, then select one pair of atoms and collide them with probability of N_{coll} , otherwise no collisions occur.
10. Store the new values of velocities and advance to the next simulation step.

In AtomECS, this algorithm is implemented by the `ApplyCollisionsSystem` and resulting data and parameters are stored in components. A list of the components, systems and resources used are listed in Table 7.1. Resource here is a term that refers to data that exists only once in the simulation, such as the simulation timestep. They are global parameters that are not associated with any particular entity.

7.2.3 Comparison with theory

7.2.3.1 Homogeneous trap

We now compare the collision rates measured in AtomECS to the theoretically predicted values. For a homogeneous box trap we expect a collision rate given by Eqn. 7.13. We simulate 10^8 ^{87}Rb atoms at a temperature of $1\text{ }\mu\text{K}$ in a homogeneous trap of width $100\text{ }\mu\text{m}$ for a period of 1 ms . We choose collision cells of size $10\text{ }\mu\text{m}$ with a timestep of $0.1\text{ }\mu\text{s}$ and keep the partition static, since the cloud does not change over time. We vary the macroparticle factor α to vary the simulated particle number while keeping the total atom number constant. This demonstrates the effects of statistical fluctuations on the measured collision rate. DSMC predicts that the time (and in a homogeneous gas, spatially) averaged collision rate is accurate, but statistical fluctuations are present and decrease with increasing simulated particle number. The varying particle number and constant cell size lead to a varying mean number of particles per cell. We plot the average and standard error of the simulated collision rates over multiple repeats. Also plotted are dashed lines denoting a $\pm 10\%$ deviation from the theoretical value.

The agreement is quite good ($< 10\%$ discrepancy) even for ~ 10 particles per cell and converges to the exact value for higher numbers. Therefore we can see that it

Name	Type	Description
<code>CollisionParameters</code>	Resource	Contains parameters for the DSMC algorithm, such as the macroparticle factor and collisional cross-section
<code>CollisionsTracker</code>	Resource	Stores data about collisions e.g. number of collisions, number of particles per cell
<code>ApplyCollisionsSystem</code>	System	Performs the DSMC collisions algorithm on all atoms using the constructed spatial partition
<code>BoxID</code>	Component	Key number of the collision cell containing the atom
<code>ApplyCollisionsOption</code>	Resource	Flag to set the collision system on or off
<code>BuildSpatialPartitionSystem</code>	System	Constructs a spatial partition at each timestep and assigns atoms to partition cells
<code>RescaleSpatialPartitionSystem</code>	System	Rescales the width and number of partition cells based on calculated density
<code>PartitionParameters</code>	Resource	Contains parameters for the spatial partition such as cell width and number

Table 7.1: Table of components, systems and resources used in the collisions implementation in AtomECS.

is desirable to choose either the cell size or α to have at least this many particles per cell. The cell size sets a coarse graining of position and velocity correlations, and also needs to be smaller than significant density variations in the trap. These requirements may impose an upper limit on the cell size, beyond which reducing α (thereby increasing the number of simulated particles for a fixed atom number) can achieve reduced statistical fluctuation.




Figure 7.1: Agreement between simulated and predicted time-averaged collision rate in a homogeneous trap of 10^8 ^{87}Rb atoms at $1\,\mu\text{K}$. Errors show the standard error over many time steps. Lower mean particles per cell lead to larger statistical fluctuations.

7.2.3.2 Harmonic trap

Another important test case is that of a harmonic trap. A harmonically trapped cloud has a varying density profile, and so it provides a more general test of the collision accuracy while still being analytically solvable. The theoretical collision rate per particle in a harmonic trap is given by [160]:

$$R = \frac{N_P m \omega^3 \sigma}{2\pi^2 k_B T}, \quad (7.18)$$

where N_P is the number of particles, ω is the geometric mean of the trapping frequencies, and T is the temperature of the cloud.

We simulate an isotropic harmonic trap of trapping frequency $\omega = 2\pi \times 200\,\text{Hz}$ containing 10^8 atoms at a temperature of $2\,\mu\text{K}$. We again vary the simulated particle number and macroparticle factor to vary the particles per collision cell at constant total atom number. The density profile of the cloud is a Gaussian with a width at this




Figure 7.2: Agreement between simulated and predicted time-averaged collision rate in a harmonic trap containing 10^8 ^{87}Rb atoms at $2\,\mu\text{K}$. Error bars show standard deviation of the collision rate over many time steps. Results for two fixed partition cell sizes are shown. Small cell size (light blue) leads to smaller systematic error in the collision rate for sufficiently high particle density than large cell size (dark blue).

temperature of $15.5\,\mu\text{m}$ ($\sim 60\,\mu\text{m}$ total width). It is illustrative here to first compare the collision rates obtained from using a static spatial partitioning with different cell widths, shown in Fig. 7.2.

It is evident that although for the cell widths shown the collision rate converges to a reasonable agreement, there is a systematic error in the larger cell widths even after convergence. This can be attributed to a consistent underestimate in the collision number arising from applying kinetic theory for a homogeneous gas to an inhomogeneous cloud. For larger cell widths, the density will show an appreciable gradient across the cell. This non-uniformity will always lead to a greater number of collisions than a uniform density distribution, so that the total collision rate is underestimated.

Smaller cell widths are a closer approximation to a truly uniform gas inside the cell, so the total error is reduced.

7.3 Magnetic trapping

The forms of trapping previously implemented in the current version of AtomECS are optical dipole trapping and magneto-optical trapping (MOTs). In this latter example, magnetic fields are present but are too weak to trap the atoms; instead the confining force is provided by the laser light. Confining forces originating from magnetic fields were therefore not implemented in AtomECS. However, cold atoms can be confined in purely magnetic traps, and it is important to be able to simulate these.

7.3.1 Theory of magnetic trapping

The interaction at the core of magnetic trapping is that of an atom's magnetic dipole moment $\boldsymbol{\mu}$ with an external magnetic field $\mathbf{B}$. The Hamiltonian for this interaction is

$$H = -\boldsymbol{\mu} \cdot \mathbf{B}. \quad (7.19)$$

In a weak external field the magnetic moment $\boldsymbol{\mu}$ is given by

$$\boldsymbol{\mu} = -g_F \mu_B \mathbf{F}, \quad (7.20)$$

where g_F is the g-factor of the hyperfine state F , μ_B is the Bohr magneton, and $\mathbf{F}$ is the total angular momentum of the atom.

The energy of an atom in an external magnetic field is determined by the orientation of the atom's magnetic moment with respect to that field. Consider a magnetic field $\mathbf{B} = B_0 \mathbf{e}_z$, where we have chosen the axis of the magnetic field to be our z -axis. Our energy levels then are the eigenstates of Eqn. 7.19, Zeeman states labelled by $|m_F\rangle$ where m_F is the projection of the angular momentum $\mathbf{F}$ onto the magnetic field axis $\mathbf{e}_z$. The energies of these states, for weak field strengths, are

$$E = m_F g_F \mu_B B_0. \quad (7.21)$$

States with $m_F g_F < 0$ have decreasing energy for increasing B_0 , and are described as 'high-field seeking'. We are more interested in states with $m_F g_F > 0$, since these states have increasing energy for higher B_0 and are therefore 'low-field seeking'. Atoms in such states will move towards a local minimum of magnetic field strength, and can therefore be trapped in such regions.

AtomECS simulates the dynamics of physical systems by applying forces to atoms, rather than by minimising potential energy. We therefore need an expression of the force exerted by this magnetic dipole interaction.

The force exerted on a particle moving in a potential field V is given by

$$\mathbf{F} = -\nabla V, \quad (7.22)$$

and hence for the magnetic dipole force,

$$\mathbf{F}_{dip} = -m_F g_F \mu_B \nabla |B|, \quad (7.23)$$

where $|B|$ is the magnitude of the local magnetic field.

One underlying assumption here is that the motion of the atom is adiabatic so that the atom remains in the same local Zeeman state. This assumption could be violated if the magnitude of the magnetic field were small or if the atom were moving very quickly.

7.3.2 Implementation of magnetic dipole force

Magnetic fields in AtomECS are represented by entities with specific components. For example, a quadrupole field is an entity with a `QuadrupoleField3D` component, containing the direction and quadrupole gradient of the field, and a `Position` component, containing the origin of the magnetic field. Multiple magnetic fields can be defined within one simulation.

An atom (or any entity that needs to know the values of magnetic fields) is given a component `MagneticFieldSampler`. This component stores the value of the field at the position of its parent entity. Each field has an associated system (e.g. `Sample3DQuadrupoleFieldSystem`) which calculates (samples) the value of that magnetic field at the position of every entity with the `MagneticFieldSampler` component. Since fields can be added in linear superposition, the total field is just the sum of every separate field sampler. This is a fast and modular way to calculate total magnetic fields, since these fields need only be calculated at the specific location that they must be known, rather than being calculated everywhere and sampled later.

However, this presents a small complication for the magnetic dipole force which depends on the gradient of the magnetic field magnitude. A gradient requires knowledge of the magnetic field at nearby positions, but each magnetic sampler can only calculate the value of the field at its exact position. We must add to each sampler system a calculation of the gradient near the position of the atom. The gradient

of the total magnetic field magnitude is not a linear superposition of the individual magnetic magnitude gradients. We require an expression for ∇B_{tot} in terms of the gradients and values of the individual fields.

Firstly we define the total magnetic field $\mathbf{B}_{tot}$ in terms of its constituent magnetic fields $\mathbf{B}_i$,

$$\mathbf{B}_{tot} = \sum_i \mathbf{B}_i, \quad (7.24)$$

and the magnitude of the field $B_{tot} = |\mathbf{B}_{tot}|$ is

$$B_{tot} = \left(\sum_j B_{tot,j}^2 \right)^{\frac{1}{2}}, \quad (7.25)$$

where $j = x, y, z$ is the index of each component of the total magnetic field.

The i th component of the magnetic magnitude gradient $\partial_i B_{tot}$ is then given by

$$\begin{aligned} \partial_i \mathbf{B}_{tot} &= \partial_i \left(\sum_j B_{tot,j}^2 \right)^{\frac{1}{2}} \\ &= \frac{1}{2} \left(\sum_j B_{tot,j}^2 \right)^{-\frac{1}{2}} \left(\sum_j 2B_{tot,j} \partial_i B_{tot,j} \right) \\ &= \frac{1}{B_{tot}} \mathbf{B}_{tot} \cdot \partial_i \mathbf{B}_{tot}. \end{aligned} \quad (7.26)$$

From this we can see that the gradient of the field magnitude contains information about the gradient in every direction of each component of the constituent field. We can write this conveniently as

$$\nabla B_{tot} = \frac{1}{B_{tot}} \mathcal{J} \mathbf{B}_{tot}, \quad (7.27)$$

where $\mathcal{J}$ is the Jacobian matrix of the total magnetic field, with elements given by

$$\mathcal{J}_{ij} = \partial_i B_{tot,j}. \quad (7.28)$$

Some authors conventionally define the Jacobian as the transpose of this form. The purpose of writing the magnitude gradient in this form is that the Jacobian of the total magnetic field is a linear superposition of the individual field Jacobians. To implement this we can therefore add some extra data to the **MagneticFieldSampler** component - namely, the Jacobian and the total magnetic magnitude gradient. Each sampler system for the constituent fields calculates, in addition to the value of the field, the Jacobian of the field at that point in space, and adds it to the value stored in the sampler. A new system, the **CalculateMagneticMagnitudeGradientSystem**,

runs after all other magnetic sampling systems and calculates the gradient according to Eqn. 7.27, and adds it to the `MagneticFieldSampler`, ready to be used by the `ApplyMagneticDipoleForceSystem`. A list of these and further systems and components used in implementing the magnetic trapping are given in Table 7.2.

If it were known that the only field was, by way of example, a 3D quadrupole field, the field magnitude gradient could be calculated analytically and applied in an explicit form. This may be more straightforward for simple magnetic fields; but the above method is completely general for any superposition of arbitrarily defined fields. In particular, this renders AtomECS highly future proof. For example, the magnetic force can be calculated for the magnetic field from an arbitrary array of current carrying coils, for which the field's magnitude and gradient are in general not known analytically. This gradient calculating system is independent of the form of the constituent magnetic fields, in keeping with the ECS philosophy of decoupled systems.

7.3.3 The Time-Orbiting Potential

A very experimentally relevant form of magnetic trap is the time-orbiting potential, or TOP trap, which was discussed in Chapter 1. We quickly summarise here that the TOP is formed by the addition of a rotating bias field, which rotates faster than the atomic motion so that the atoms experience a time-averaged potential [45].

The trap frequencies of the TOP trap are given by

$$\omega_r^2 = m_F g_F \mu_B \frac{B'^2}{2mB_{TOP}}, \quad \omega_z^2 = 8\omega_r^2, \quad (7.29)$$

where B' is the quadrupole gradient and B_{TOP} is the amplitude of the rotating bias field.

We can simulate magnetic trapping in the TOP directly. Since it is the superposition of a static 3D quadrupole field and a rotating uniform bias field, we create two entities with these magnetic fields. A new component, the `TimeOrbitingPotential` will store the amplitude, rotation frequency and axis of rotation of the TOP. The `TimeOrbitingPotentialSystem` will, based on the data in this component, add a uniform bias field to the magnetic sampler and rotate it every timestep. Since the field is uniform in space, it has no Jacobian to calculate, though it will still affect the magnetic dipole force through its modification of the field magnitude.

We can compare the motion of a particle in the simulated TOP trap to the oscillation frequencies in Eqn. 7.29. We test a TOP trap with $B_{TOP} = 10$ G, $B' = 100$ G/cm,

Name	Type	Description
MagneticFieldSampler	Component	Stores values of field, field magnitude, Jacobian and gradient at entity's position
Quadrupole3D	Component	Stores quadrupole gradient and direction
SampleQuadrupole3D System	System	Calculates field and Jacobian at position of entity
CalculateMagnetic MagnitudeGradient System	System	Runs after other magnetic systems and calculates field magnitude gradient at position of entities
TimeOrbitingPotential	Component	Stores TOP amplitude, rotation frequency and axis of rotation
TimeOrbitingPotential System	System	Adds a uniform bias field and rotates at frequency given by TOP component
MagneticDipole	Component	Stores product of Zeeman state and g factor $m_F g_F$
ApplyMagneticForce System	System	Applies magnetic dipole force to atoms with a magnetic dipole

Table 7.2: Table of components and systems used in the implementation of magnetic dipole trapping.

and $\omega_{TOP} = 2\pi \times 7$ kHz. We give a single particle situated at the origin an initial velocity of 2 cm/s in the x direction. The position of the particle is plotted in Fig. 7.3.

The predicted radial trap frequency from Eqn. 7.29 is $\omega_r = 2\pi \times 20.16$ Hz. Fitting a sinusoidal curve to the position of the particle, the fitted trap frequency is $\omega_{fit} = 20.186(1)$ Hz, where the uncertainty represents a 95% confidence interval of the fit, a relative difference of 0.13%. In this case, we can expect that the simulated trap frequency is the more accurate estimate, since Eqn. 7.29 is an approximation that neglects the fast dynamical effects that occur on the timescale of ω_{TOP} . The simulation also shows a small motion in the y -direction that the time-averaged theory neglects. The basic assumption of the TOP theory is that the TOP rotates much faster than the oscillation frequency of the atoms. We can simulate the effect of a lower rotation frequency. In Fig 7.4 we compare a TOP rotation frequency of 7 kHz




Figure 7.3: Simulated single particle oscillation in the TOP. Note that the rotating field in the $x - y$ plane causes a small motion along the y -direction, despite zero initial y -velocity.

and 200 Hz in an otherwise identical trap.

It is clearly seen that the lower rotation frequency (even though it is still 10 times larger than ω_r) results in a significant micromotion superimposed on an envelope that closely resembles the 7 kHz curve. The trap frequency is also slightly reduced, to a fitted frequency of $\omega_r = 2\pi \times 20.050(15)$ Hz.

7.4 Simulating evaporative cooling

An essential part of cold atoms experiments is evaporative cooling. Laser cooling methods often only cool to the Doppler limit of the transition on which they operate. This is typically in range of a few hundred μK for alkali atoms [168]. In some cases, cooling below this limit (sub-Doppler cooling) is possible. However even in this optimal scenario the lower limit is on the order of μK , which is generally insufficient for achieving the critical phase-space density of the BEC transition. In our experiment, the critical temperature for the BEC transition is typically of the order of 100 nK. To




Figure 7.4: Simulation of the effect of different TOP rotation frequencies. A lower rotation frequency, closer to the oscillation frequency of the atoms, results in significant micromotion on the envelope of the sinusoidal motion predicted in Eqn. 7.29, and a slightly reduced trap frequency.

achieve these temperatures, evaporative cooling is generally employed. This method of cooling is the process of selectively removing the most energetic atoms, reducing the average energy per atom, and upon allowing the cloud to rethermalise the new temperature will be lower than before the evaporation. By successively (or continuously) removing the upper tail of the Boltzmann distribution, the cloud equilibrates to lower temperatures. This process reduces atom number, but can result in an increased phase-space density overall, sometimes referred to as ‘run-away’ evaporation.

Since evaporative cooling relies on thermalisation to operate, it is dependent on inter-atomic interactions. The implementation of s-wave collisions allows us simulate this process in AtomECS.

Evaporative cooling can be experimentally implemented in a number of ways. One of the most common is by means of an ‘RF knife’: an oscillating radiofrequency field resonant with a sub-class of the atoms, which transitions them to untrapped states. In the case of a magnetic trap, an rf field of frequency ω will be resonant with the

Zeeman splitting at all positions $\mathbf{r}$ such that

$$\hbar\omega = m_F g_F \mu_B B(\mathbf{r}). \quad (7.30)$$

The cut in energy space is therefore equivalent to a cut in position space; since atoms are trapped at a local minimum of magnetic field strength, higher field strengths and hence Zeeman splitting are reached farther in position from the trap minimum, so the RF knife cuts from the outside in as the frequency is gradually decreased.

This RF knife is easy to implement in AtomECS. We define a component `RFKnife` which stores the frequency of the knife. We also define a system, `ApplyRFKnifeSystem`, which reads the `MagneticFieldSampler` of an atom and marks it for removal from the simulation if its Zeeman splitting is greater than the knife frequency. This removes all atoms farther out than the resonance with the knife.

This is not exactly the same as a real RF knife, which only removes atoms close to the resonance, but in practice an RF knife usually removes atoms at a greater Zeeman splitting once the atoms cross the resonance in either direction. An RF knife is also usually started at a high frequency and monotonically decreased to sweep into the atoms from the outside in, producing the same effect as the AtomECS implementation.

The sequence we simulate is that of evaporation in the TOP trap. This is a particularly relevant example, since simulating the achievement of BEC densities by evaporative cooling is a desirable use case. A typical experimental ramp profile is shown in Fig. 7.5a, which we use as the basis of the simulation. This ramp includes modifying the TOP field amplitude which changes the trap frequency ω . This is included in the experimental sequence to compress the cloud in order to improve the collision rate. It has the additional effect of changing the temperature of the cloud due to adiabatic compression, such that

$$T_f = T_i \frac{\omega_f}{\omega_i}, \quad (7.31)$$

where T_f and T_i are the temperatures after and before the compression, respectively.

That AtomECS correctly simulates this effect has been checked and the results are in excellent quantitative agreement with the theory [169]. Furthermore it can be proved that while this changes the temperature, the change in density exactly cancels this effect when calculating the phase-space density so that it remains unchanged by the compression.




Figure 7.5: (a) RF Knife frequency and TOP amplitude ramps that are used in a typical experimental sequence. TOP amplitude converted to an energy in terms of Zeeman splitting of the atom. (b) Phase-space density vs atom number curve for a simulated evaporative cooling ramp using the parameters in a).

We plot the phase-space density (PSD) against the total atom number. This is a useful metric because the phase-space density is a dimensionless measure that combines the requirement to be both cold and have high density. It is given by

$$PSD = n\lambda_{th}^3 = n\left(\frac{h^2}{2\pi mk_B T}\right)^{\frac{3}{2}}, \quad (7.32)$$

where n is the atomic number density and λ_{th} is the thermal de Broglie wavelength. It is also an easy way to define the BEC transition, since it is equivalent to exceeding a phase-space density of $PSD_{crit} > 2.612$. This critical phase-space density is marked on the plot by a red dashed line.

It should be noted there are a few effects in the experiment that are not modelled by AtomECS in this version. One is the finite lifetime of the trap, limited by vacuum pressure. This causes some atom loss and reduces the efficiency of the evaporative cooling ramp, but if the vacuum-limited lifetime is long the effect is small. Similarly, heating rates in magnetic traps are generally very low, on the order of a few nK/s, and so neglecting them is justifiable here. Perhaps the biggest unaccounted for effect is the so-called ‘circle-of-death’. This is the locus of the displaced zero of magnetic field, pushed away from the centre of the trap by the uniform bias field. This magnetic

field node orbits at a radius of $r_0 = B_{TOP}/B'$. Atoms that approach it undergo the same spin-flip transitions that caused atom loss in the plain quadrupole trap, thus acting as an additional form of evaporation. These spin-flips are not modelled by AtomECS. For the duration of this sequence, r_0 is a few times larger than the size of the cloud, but it may still cause a small amount of atom loss.

The difficulty in obtaining data from a ‘pure’ system that matches the simulation very closely means that a quantitative comparison to experiment has not been performed. However, the simulation does produce results that are qualitatively similar to observations from experiment, reaching the critical PSD at a similar atom number to experiments ($\sim 10^6$.)

7.5 Conclusion

We have presented the implementation and demonstration of simulating two new effects in the AtomECS package; Direct Simulation Monte Carlo simulation of atomic s-wave collisions, and trapping via the magnetic dipole interaction with an external field. The simulated collision rates and magnetic trap frequencies match those predicted from theory. These effects were combined to simulate the process of evaporatively cooling a cloud of atoms to BEC density. The code has been unit tested and the ECS architecture permits easy future development and testing.

Useful further upgrades to the simulation of collisions would be to generalise to mixtures of atomic species, velocity dependent models of collision cross-section, and relax the assumption of thermal distributions of velocities within a cell. These could be implemented relatively straightforwardly, but the gain in versatility would come with a cost in performance. In particular, to generalise to mixtures and velocity-dependent cross-sections, the collision algorithm would have to loop over each individual pair of atoms and calculate a probability of collision, rather than the pre-calculated total number of collisions. This would take computational time in both calculating these probabilities, and iterating over pairs of atoms that do not end up colliding.

Chapter 8

Towards a new experimental apparatus

This chapter is concerned with the design and construction of a new experimental apparatus for producing Bose-Einstein condensates, which will feature a number of improvements over the existing apparatus.

8.1 Motivation

The current experimental apparatus for producing cold atomic clouds has evolved over more than a decade, based on the original design and construction carried out mainly by Marcus Gildemeister, Benjamin Sherlock and Benjamin Sheard [170–172]. This apparatus has performed its intended purpose well but over this long period has aged and shows a number of flaws, particularly in comparison to the current state-of-the-art cold atom experiments. These shortcomings can be broadly described in three categories.

8.1.1 Sequence duration

Our apparatus has sequence times of over one minute, longer than is usually achievable by other apparatuses (~ 30 s [44]) and dramatically slower than the fastest time to produce BECs of just one second [173]. Significant limitations include the relatively slow magneto-optical trap (MOT) loading rate, and the evaporative cooling required to compensate for heating during the transport to the glass science cell.

It is very difficult to produce a Bose-Einstein condensate (BEC) and load a MOT from a thermal background vapour at the same site because the pressures required are incompatible. A BEC, particularly if long lifetimes are desired for experiments,

requires ultra-high vacuum pressures at approximately 10^{-10} mbar, while MOTs require a vapour pressure of the order of 10^{-8} mbar to reach an appreciable loading rate. As such, at least two different vacuum chambers are usually maintained at different pressures by means of differential pumping. This often takes the form of a low conductance tube, with separate pumps on the high pressure and low pressure side. Some method is then employed to transport the laser cooled atoms from the high pressure chamber to the ultra-high vacuum (UHV) section. The current apparatus implements this by initially trapping atoms in a 3D pyramid MOT, after which the current-carrying coils which produce a confining quadrupole field are mechanically translated along a transport rail, carrying the atoms through the differential pumping tube to the science cell.

This mechanical transportation of the atoms out of the pyramid MOT through the high pressure chamber means that the cooled cloud must travel through the very same background gas that it loads from. A high background pressure may lead to a high MOT loading rate, but it causes more atoms to be lost or heated via collisions during the transport process. This imposes a ceiling on the rubidium vapour pressure in the MOT chamber, limiting the fastest MOT loading rate achievable. This is improved considerably by the use of light-induced atomic desorption (LIAD) [174] to briefly raise the pressure at the beginning of the MOT loading, but it is still at least 10 s.

The magnetic transport also causes significant heating of the atoms during their rapid acceleration. This is exacerbated by the additional magnetic push given to the atoms to steer them through the differential pumping tube, which due to a design flaw is not perfectly straight. This results in comparatively hot atoms being released into the science cell before being captured in the quadrupole trap. After ~ 20 s of evaporative cooling, the atoms reach of the order of $\sim 100 \mu\text{K}$ - similar to the Doppler temperature for Rb and therefore barely colder than when they started in the MOT. Already almost 30 s of sequence time have elapsed. The current system of initial cooling and transport is therefore a significant and inefficient use of time.

Furthermore, there are a number of intermediate trapping systems between loading atoms into the science cell and our final desired RF-dressed shell trap. We load first a quadrupole trap, then a TOP trap, then a TAAP trap, and finally the shell trap. This necessitates mode matching between these traps and further evaporative cooling to offset additional heating and atom loss caused. This also adds significantly to the sequence duration. If the process of loading the shell trap could be initiated earlier in the sequence, more time could be saved.

8.1.2 Experimental complexity

Not only does this sequence of trapping systems prolong the duration of the sequence, it considerably complicates it. All the traps require additional hardware, such as coils, supplies, amplifiers, signal generators, RF switches, and digital and analog channels. This provides many points of failure, and since this hardware has aged along with the apparatus it serves, failures do occur. On the sequence side, transitions between different trapping geometries must be optimised and can result in not only atom loss and heating, but collective excitations in the cloud such as dipole oscillations. These excitations can interfere with the physics being investigated. If they occur early in the sequence, they can dephase over time, resulting in an apparent random instability in final atom number or temperature.

The intricacy of the apparatus and sequence leads to difficulty diagnosing and debugging problems, costing lab time and effort and further slowing down the process of research.

8.1.3 Stability

The complexity of the apparatus can also manifest as instability of the experiment; often as fluctuations in the atom number or temperature, or as collective excitations in the final cloud. In the case of loading double-well potentials, splitting the condensate is also highly sensitive to the amplitude and phase of the RF fields.

There are many stages to the experiment for which calibrations or optimal parameters can drift over time. The low repetition rate makes us particularly susceptible to drifts that occur on the timescale of a few hours or days, since this can occur over the duration of a significant data taking sweep. These are often thermal in nature and the passive temperature stability of the laboratory is insufficient to eliminate drifts.

The RF coils are three dimensional, hand-wound coils. This makes them somewhat susceptible to drifts in their impedance, inductance and mutual coupling. The MRF sequences are sensitive to relative drifts in amplitude of just a few parts in a thousand, and it is possible for our calibrations to drift by this much over the course of days.

If we wish to study many body phenomena (which by their statistical nature often require a large quantity of data), or more sensitive effects such as quantum tunnelling, then a more stable, faster, simpler apparatus would be greatly advantageous.

8.1.4 New design concept

A number of these problems can be solved and further advantages gained by switching to a high flux pyramid MOT source. We again load a MOT from a thermal background vapour in a pyramid MOT, but rather than a 3D MOT that is magnetically transported, this pyramid has an aperture at its apex which produces a ‘hole’ in the reflected axial beam, forming a low-velocity intense source [175]. Away from the apex the pyramid acts a typical 3D MOT, collecting atoms from background vapour and cooling them. As the cooled atoms diffuse into the centre of the MOT, they enter the shadow of the aperture and are accelerated out of the pyramid through its apex by the input cooling beam. This produces a high flux of cooled atoms with a narrow velocity distribution. For this new apparatus, we used a pyramid source designed and constructed by Sean Ravenhall, another group member, intended for use as a high-flux rubidium source [176, 177].

This atomic beam is recaptured by a six-beam, 3D MOT situated in a glass cell. From here it would be typical to further transport the atoms recaptured by the 3D MOT to a final science cell, but this risks replicating some of the problems with the current transport design. Instead, we will attempt to load the shell trap directly from the recapture MOT and perform evaporative cooling down to BEC.

The rubidium for the pyramid MOT is provided not by the previous and common method of resistively heated alkali metal dispensers, but by a separate chamber containing a pure ingot of rubidium. This provides a stable, pure and practically limitless source of atoms.

8.2 Laser system

Frequency-stabilised laser light is needed for two primary purposes. Laser cooling requires cooling and repumping light resonant with the transitions in the D2 line of ^{87}Rb (within some frequency detuning comparable to the linewidth of the transition), just as with the existing apparatus (see Section 3.3.1). Cooling and repumping light is also used for absorption imaging.

This apparatus initially produced recapture MOTs in the cell using one diode laser as the source of light and a single round of amplification for both MOTs. This was later upgraded to afford additional capability and improved power. Below we describe the initial system and then how this was modified to the current laser system.

8.2.1 Initial laser system

All the light we use must be frequency stabilised close to the resonance of the ^{87}Rb cooling transition. This is generally achieved by the use of a stable, but comparatively low power, reference laser, which provides light of the correct frequency. This is then later amplified to reach the desired power for the experiment.

Our initial reference laser was a *Toptica DL100* external cavity diode laser (ECDL) locked to the ^{87}Rb cooling line via modulation transfer spectroscopy (MTS). At the output of the laser the light was split into two beams. One was sent to be amplified and the other to the spectroscopy cell. Prior to the spectroscopy cell, the light was sent through two acousto-optic modulators (AOMs) both in a double-pass configuration.

AOMs typically modify the frequency of light by a minimum of 10s or 100s of MHz, too much to be useful in our case. By utilising two AOMs detuning the frequency in opposite directions, small detunings could be achieved. This allowed the frequency of the laser light to be modified by up to several multiples of the linewidth $\Gamma = 2\pi \times 6\text{ MHz}$ in either direction, to find the optimum detuning for operating the MOTs. The double-pass configuration ensures minimal displacement of the beam under modification of the frequency and maintains a strong MTS signal, which is dependent on the beam alignment, over the bandwidth range of the AOMs.

We detuned the light used for the spectroscopy to a higher frequency than that sent to the MOT, and lock the blue-detuned light to the cooling transition. This results in the light sent to the MOT being red-detuned with respect to the cooling transition, as is usually done for MOTs.

After diverting power to the MTS, only about 10 mW was left, far too little to run the experiment. The remaining power was used to seed a tapered amplifier (TA) to provide the requisite power for the experiment.

The two MOTs require significant amounts of power due to the relatively large final beam sizes and numerous fibre couplings involved. To achieve appreciable loading rates and maximum atom number, the TA produced $\sim 650\text{ mW}$ of light. This required running the TA near its maximum current of 2.5 A and is also above the optical fibre manufacturer's recommended power for coupling.

After the TA the light was sent to a distribution board where it was passed through an electro-optical modulator (EOM) to produce the repumper sidebands. The requirements of a high optical damage threshold and resonance at the sideband frequency of $\sim 6.5\text{ GHz}$ excluded most commercially available EOMs at the time.

As a result a prototype EOM developed by Qubig GmbH¹ was used, which had an adjustable frequency range of 6,550 – 6,691 MHz and an optical power threshold of 5 W/mm².

This EOM produced two sidebands at $\sim \pm 6.5$ GHz, one of which was resonant with the repumper transition. The maximum power in each of the sidebands was 6% of the cooling power.

This light was further split three ways - one fibre to the pyramid chamber, one to a splitting board to divide the light into six MOT beams for the recapture MOT, and one to an etalon to monitor the frequency of the cooling and repumper light.

This laser setup was functional and produced large MOT clouds, but had a number of disadvantages. First, the *Toptica DL100* is a generally obsolete model of ECDL, and provides relatively poor stability. It was prone to unlocking from the resonance frequency, or operating in a multi-mode regime.

Secondly, the TA was being driven very hard to produce the required power, and this power exceeded the safe threshold for fibre coupling. This would shorten the lifetime of the TA and risked damaging the optical fibre in addition to reducing fibre coupling.

Finally, this setup only allowed detuning the global reference frequency, changing the detuning of both pyramid and 3D recapture MOT simultaneously. It is likely both MOTs could have different optimal detunings, which did indeed prove to be the case. To do this with the first setup would have required another fibre coupling and a second board to detune the 3D recapture MOT with AOMs, at a further cost to total available power.

8.2.2 Current laser system

To remedy these disadvantages, the laser system was redesigned to concurrently operate with the *Toptica DL Pro* ECDLs used by the old apparatus, since both apparatuses were simultaneously operational.

To produce cooling light a small amount of power was taken from the reference *DL Pro*, which is locked on resonance with the cooling transition. The reference light is split two ways (one for the pyramid MOT and one for the glass cell recapture MOT) and each beam is passed through AOMs in the double-pass configuration to control the detuning for each MOT independently. These low power cooling (1 mW)

¹Qubig Rb87-6.6G

beams are used to seed two separate injection-locked slave diode lasers². These output ~ 100 mW of power.

Instead of producing repumping light via an EOM, the repumping light was also borrowed from the repumping *DL Pro*, itself detuned 6.5 GHz from the reference laser using a beat-note offset lock. The repumping light frequency is not typically ramped via the beat-note offset lock in either of the apparatuses' experiments and so this frequency is static and can be used for both apparatuses in parallel. The borrowed repumping light seeds another injection-locked slave diode laser and is also split in two, one beam for each MOT.

These cooling and repumping beams are then recombined via coupling into the same fibre. The combined cooling and repumping beams then each seed a separate tapered amplifier (TA). This produces two distinct, high power beams, which can have their detuning independently modified. Having two TAs also reduced the power output required of either single TA, improving the lifetime of the diodes and reducing the power coupled into a fibre below a safe threshold.

The output of one TA is then fibre coupled to the pyramid input board. This light is cleaned by a polarising beam splitter, then circularly polarised by a $\lambda/4$ wave plate before being expanded to a beam of $1/e^2$ diameter 12 cm and launched through the pyramid viewport. The other TA output is sent to the recapture MOT optics via the six-way splitting board. Further detail on the current laser system can be found in the thesis of Abel Beregi [178].

8.3 Recapture MOT

8.3.1 MOT distribution board

The 3D recapture MOT is a standard six beam MOT configuration, with 6 optical paths each fed by a fibre. This means a single, moderately high power input beam must be split into six roughly equal (but adjustable to correct for varying fibre coupling) ways. The polarisation and amplitude stability of this splitting is highly important, since power imbalances in the final MOT will reduce atom number, increase temperature, and cause the shape and position of the laser cooled cloud to drift. This will be particularly problematic for loading into the final shell potential, where small changes in mean position or atomic distribution can lead to significant

²*Panasonic* IR Laser Diode 783 nm 200 mW




Figure 8.1: 3D render of the MOT beam distribution board. A single combined cooling and repumping beam enters through the fibre on the left. Its polarisation is cleaned and then adjusted by a succession of half-wave plates to split off light in reflection from a PBS producing six beams of similar power. The polarisation of each beam is further cleaned by a second PBS before being steered into one of six optical fibres to be sent to the MOT board. Model provided by Abel Beregi.

increases in the potential energy. A standard optical layout may not achieve the stability required, and would also take up a significant footprint on the limited optical table space.

To conserve space a custom, compact splitting board was built. This board was made of a single piece of aluminium, with holes tapped and bored for optical components. The input beam is split by a succession of half-inch waveplates and polarising beam splitters (PBS). The PBSs are fixed by an epoxy resin for mechanical stability. The reflected component from each PBS is used for the MOT beams, but since PBSs generally have cleaner polarisation in transmission each beam is passed through a second PBS to minimise polarisation drifts after fibre coupling.

Each beam then is steered by a pair of half-inch mirrors, screwed into the board so that their orientation can be manually adjusted if necessary, into a fibre collimator. The collimators are set in machined holes in the walls of the board and secured with a screw from above. For improved stability, the holes are not circular to match the profile of the collimators, but rather double-bored shaped, so that the collimators sit

and are wedged between well defined axes; this prevents them from rocking on an uneven 2D surface. Threaded holes are set in the corners to allow an aluminium lid to be fixed to the board.

This board provides six MOT beams that are highly alignment and polarisation stable. The optical fibres lead to the MOT board around the glass cell. The glass cell is the final destination of the atoms, where the atomic beam will be recaptured by a standard, 3D six beam MOT, and then transferred into its final magnetic trap.

8.3.2 Magnetic fields

The MOT and the shell trap both require a quadrupole magnetic field. Shell traps often require high quadrupole gradients, corresponding to high currents. The field is produced by two hand wound cylindrical coils, composed of hollow, square copper tubing through which cold water is run to cool the coils. The coils have outer radius 35 mm and inner radius 18 mm, and are 14 mm thick and are separated by 50 mm. These coils are an almost identical copy of those used in the old apparatus.

These coils are glued into 3D printed square formers, mounted on four threaded stainless steel rods and fixed in place by nuts from below and above. This allows the position of the formers to be adjusted to modify the separation between the quadrupole coils. The two formers are separate pieces such that horizontal optical access is maximised, only impeded by the rods, and vertical optical access is permitted by circular clearance holes through the centre of the coils. The quadrupole coil current is provided by a *Delta Elektronika* SM15-400 DC power supply.

8.3.3 Optical layout

The 3D recapture MOT is generated by the combination of the magnetic fields and the trapping laser beams. Six beams of balanced power are needed along each cardinal direction to trap the atoms with optimal efficiency. Imbalanced powers along each direction leads to a reduced capture efficiency from the pyramid and an uneven trapping region due to interference fringes and net force in one direction, resulting in a much smaller MOT. Four beams travel along the horizontal plane; two more travel vertically. To produce the upwards travelling vertical beam, light is sent underneath the lower quadrupole former to a 45° mirror which reflects the beam upwards.

The downward travelling vertical beam is placed on an optical breadboard which sits above the rest of the optics, mounted in place by four columns at the edges of the board and further fixed by the stainless steel rods that hold the quadrupole formers.

The upper board has a machined clearance hole positioned relative to the 4 clearance holes in which the steel rods sit. This aligns the clearance hole of the board to the clearance hole of the formers and coils, so that the downwards beam travels through the axis of the quadrupole. The upper level provides additional optical footprint for a future vertical imaging system, selective repumper optics, or optical dipole trapping.

In addition to the MOT beams, the layout requires a camera for fluorescence imaging while optimising the MOT. Fluorescence imaging is standard for observing MOTs because it provides a non-destructive, continuous real time image with high sensitivity. Future experiments in the magnetic trap require absorption imaging. Absorption imaging is destructive but captures a very accurate snapshot of the atomic density distribution at a moment in time. The absorption imaging necessitates an independent optical path with input cooling light, repumper light, and a system of lenses to magnify the image for improved resolution. Good optical access is crucial for this last element to allow the largest NA objective lens possible for maximum magnification.

To minimise the footprint taken up by six laser beams and associated optics being spread out on an optical breadboard, a cage mounted periscope for each beam was designed and constructed to compactly produce the required beam at the height of the glass cell.

8.3.3.1 Periscope design

The MOT beams have the following requirements.

- **Power stability:** The beams must all have approximately equal power. In a pyramid MOT this requirement is obtained by having only one incident beam which is reflected with very high efficiency to confine the atoms from every direction. For six beams produced independently, and sent through different fibres, it is easy for power instability to occur from numerous optical elements, polarisation drifts and fibre coupling drifts.
- **Alignment stability:** The MOT beams can recapture and trap atoms only within the intersection region of every beam. Beams that are misaligned to each other, or to the quadrupole field node, trap inefficiently and unevenly. In particular a variance in MOT position or distribution could have significant knock-on effects when attempting to load into a magnetic trap.

- Compactness: Six copies of all the optical elements required to achieve the above goals could occupy a large footprint on the relatively small optical board that geometric constraints allow us to place around the glass cell.

A compact, 30 mm cage mount periscope was designed to meet these requirements. The cage mount optics provide stable alignment and a small optical footprint. The glass cell is 80 mm off the board, where the fibre inputs sit. This distance can not be reduced, since space is needed to position the upwards travelling MOT beam underneath the coil formers beneath the cell. Positioning as many optical elements close to the board, on short posts, improves stability. A periscope then adjusts the vertical position of the horizontal MOT beams while also providing the two mirrors needed for beam alignment.

A 40 mm achromatic lens collimates the light from the fibre output and produces a MOT beam of diameter approximately 1 cm. Wider MOT beams have a larger recapture volume, increasing capture rate, and are less sensitive to small drifts in alignment. A 1 cm beam is the largest that still provides adequate clearance through the optics for adjusting alignment, and that has a high enough intensity (compared to the saturation intensity of the transition) given power limitations.

A linear polariser positioned after the collimating lens cleans the polarisation out of the fibre. This converts polarisation drifts in the optical fibre into power drifts. These power drifts then need to be removed.

A 10:90 beam sampler diverts 10% of the light after the polariser to a photodiode to monitor the beam power. This measurement is used to actively stabilise the beam power in each periscope with a PID controller. Details of this PID design and its implementation will be described in the thesis of Abel Beregi [178]. The PID modifies the power in each beam via a voltage controlled, LCD variable retarder positioned before the linear polariser, after the collimating lens. This variable-wavelength ‘waveplate’ adjusts the polarisation into the linear polariser and stabilises the transmitted power.

The periscope beam is reflected through two elliptical mirrors in right-angle mirror mounts to adjust the vertical position of the beam and provide fine alignment of the beam angle and position going into the glass cell. A $\lambda/4$ waveplate circularly polarises the light.

These four periscopes are positioned at orthogonal directions and counter-propagating around the glass cell. Similar designs are used for the vertical travelling beams, but with 90° reflections in the plane to direct the beams vertically without adjusting the beam height.



Figure 8.2: 3D render of the MOT periscope for one of the four horizontally propagating beams. Light from an optical fibre is collimated by an achromatic lens. The light passes through a voltage-controlled LCD variable retarder, which rotates the polarisation of the light before it is passed through a linear polarising plate. This allows a controllable amount of power after the linear polariser. A 10:90 beam sampler splits off light to be measured by a photodiode, which provides an error signal to actively stabilise the beam power using the LCD waveplate. Light is then steered with two right-angled elliptical mirrors to elevate the beam to the height of the glass cell and passes through a circular polariser to obtain the correct σ polarisation for the cooling transition. Model provided by Abel Beregi.

8.4 Vacuum system

The pyramid mirrors, chamber and glass cell vacuum parts were constructed by Sean Ravenhall, for his study on an adjustable pyramid design to serve as a high-flux atom source [176, 177]. Further additions have been made to this existing setup. The three components to the vacuum system are the rubidium source, the pyramid chamber, and the UHV chamber and glass cell. All the metal vacuum parts are 304L stainless




Figure 8.3: 3D model of the full vacuum system. The left ion pump serves the pyramid MOT chamber behind the glass cell in the middle of the apparatus. Attached to the top of the ion pump is an angle valve for connecting a roughing or turbomolecular pump during the bakeout and pumping down procedure. The T-shaped rubidium chamber on the top of the pyramid chamber provides thermal rubidium vapour for the MOTs. The flux of atoms out of the pyramid are recaptured in the glass cell by a six-beam MOT (not shown). On the right a second ion pump and titanium sublimation pump maintain vacuum in the UHV chamber, and an angle valve allows for attaching a turbomolecular pump. Model provided by Abel Beregi.

steel. A 3D rendering of the full vacuum system is shown in Fig. 8.3.

8.4.1 Rubidium source

Many rubidium experiments source their ^{87}Rb from resistively heated alkali metal dispensers, which was also the original method of producing ^{87}Rb for both this new apparatus and the previous apparatus. Dispensers have a few advantages. Their emission is controllable, dependent on the current run through them. They are also unaffected by contact with atmospheric pressure, at least for short periods of time. This makes breaking vacuum an easier and less risky process.

They also have some disadvantages. We rarely, if ever, wish to break vacuum once our apparatus is operational, so this flexibility is of limited use. Dispensers have a relatively short lifetime, depending on how aggressively they are used, and their advantage of controllability is also in some respects a problem. Dispensers emit not

only ^{87}Rb but also background gases, due to impurities in the metallic salt they are composed of. The ratio of ^{87}Rb to background gas emitted depends on the current run through the dispenser. Higher ratios of ^{87}Rb are generally achieved towards the higher end of the dispensers operating current. This can produce excessively high rubidium pressures, and shortens the lifetime of the dispensers.

On the previous experiment, this was solved by running the dispensers periodically, turning them on at 2-2.5 A for a few minutes every day or every few days. This led to a time-dependence of the rubidium pressure, and thus MOT number, over this period while rubidium was produced and then gradually pumped out of the system. As the remaining ^{87}Rb content of the dispensers decreased, this operating cycle had to be adjusted, and is difficult to optimise manually. This operating cycle of the dispensers represents an additional form of complexity in running the experiment.

Here we employ a much simpler rubidium source. A 5 g ampoule of rubidium (99.7% purity, with a natural ratio $\sim 3 : 1$ of ^{85}Rb to ^{87}Rb) under argon gas at atmospheric pressure is contained in an adjoining chamber. After sealing the vacuum system and pumping down to low pressure, the ampoule is broken and the rubidium vaporises, eventually increasing the vapour pressure in the pyramid chamber after a period of a day or two. This is possible for rubidium because it has a high enough vapour pressure even at room temperature.

The vaporisation rate of rubidium, and hence equilibrium pressure, is controllable by two means. Firstly, the rubidium chamber can be mildly heated to increase the vapour pressure of the rubidium. This is done after the ampoule is first cracked, to accelerate the rise to equilibrium pressure in the pyramid and coat the walls of the chamber.

The second means of pressure control is an adjustable angle valve that connects the rubidium chamber to the pyramid chamber. This can be closed to reduce the conductivity of the connection or completely seal the rubidium chamber.

The advantages to this method are that the rubidium is very pure, and that once equilibrium has been reached it acts as a constant supply of rubidium and maintains a constant pressure over time. It also has a much longer lifetime than dispensers because it has a much higher mass of rubidium - 5 g of rubidium, compared to typically 5 – 10 mg in one rubidium dispenser [179]. Typically multiple dispensers are installed but the total amount of rubidium is still orders of magnitude less than the rubidium ampoule.

8.4.1.1 Initial rubidium chamber design

The first iteration of the rubidium chamber was made with DN16CF flanges and cylinders, shown in Fig. 8.4. This was intended to be compact and light, and included an adaptor flange to fit the existing vacuum parts, which were generally DN40CF. The primary component of the chamber was a cross shaped piece, in which the rubidium sat. The rubidium ampoule conveniently fit the DN16CF cylinder almost exactly, so no additional parts were needed to hold the ampoule in position. The ampoule was cracked by inserting a long, thin, hollow breaking rod with an angled beak that ‘bites’ onto the ampoule. The other end of the rod was in contact with the sealing face of the angle valve, such that when the valve was fully closed the rod pushed through the ampoule and break it open. The rod was hollow to allow the newly freed rubidium to travel down it, through the valve and into the pyramid chamber.




Figure 8.4: 3D model of the first rubidium chamber design. This design had DN16CF flanges and all parts were stainless steel, including all-metal angle valve.

A viewport was attached to one of the flanges of the cross piece. This was intended to allow visible access to the rubidium ampoule in situ, to identify that it had successfully cracked open and view the state of the metal inside.

This design had two problems. Firstly, the small design was indeed compact and light, but the narrow radius of the cylinders resulted in a limited conductivity (which scales with the cube of the radius) of the rubidium through to the pyramid chamber.

This meant the equilibrium pressure was low and the pyramid pressure was slow to respond to changes in the temperature or valve position. After a week of heating the rubidium to 40°C we still observed almost no rubidium in the chamber so to accelerate the pressure rise we heated the rubidium to $\sim 100^{\circ}\text{C}$. This did produce enough rubidium to observe a MOT after a day of heating.

However, the vacuum seal on the initial design broke after a few more days passed. The specific cause of this is unknown, but through the viewport a large buildup of condensed rubidium was visible. Such a large concentration of heated rubidium after being heated may have corroded the viewport seal, as rubidium is known to be corrosive.

8.4.1.2 Current rubidium chamber design

Based off of these disadvantages we upgraded to a DN40CF design as shown in Fig. 8.5. Though bulkier and heavier, it provides much improved conductance. This version has no viewport as a potential weak point, so uses a T-shaped piece instead of a cross.




Figure 8.5: 3D model of the internal design of the second rubidium chamber that is currently used. The walls of the T-shaped piece and one of the cylindrical formers are hidden. The rubidium is held in a glass ampoule, suspended by two cylindrical steel formers (one shown). The ampoule is broken by closing the angle valve, which pushes the breaking rod through the ampoule. The inner diameter of the angle valve narrows prior to the sealing face (not visible), which supports the breaking rod.

The rubidium ampoule is much narrower than DN40CF vacuum pieces and no longer sits in the path of the breaking rod of its own accord, so two cylindrical formers were machined to fit the inner diameter of the chamber snugly and hold the ampoule in place, suspended in the centre of the DN40CF cross piece. None of the pieces are fixed mechanically to others; only tightly geometrically constrained in their motion.

This design has a much higher conductance due to its larger radius, and gives an improved equilibrium pressure in the pyramid chamber without aggressive heating.

8.4.2 Pyramid and cell vacuum systems




Figure 8.6: 3D model of the pyramid chamber, rubidium source and attached ion pump. The pyramid has reflective metal mirrors, and a small hole at its apex through which the cooled atom flux travels. Model provided by Abel Beregi.

As stated previously, the pyramid chamber contains a pyramid with adjustable aperture designed by Sean Ravenhall [176]. The aperture of this pyramid is set to 2 mm diameter, which provides a high flux. The pyramid chamber was initially

pumped by a 40 l/s ion pump³; this has since been changed to a 55 l/s pump⁴. Current carrying coils are wrapped around the pyramid chamber to provide the quadrupole field for the MOT. The two coils are wrapped in an anti-Helmholtz configuration with radius 56 mm, separation 50 mm, and each contains 15 turns of 3 mm diameter copper tubing with an inner diameter of 1.5 mm. These coils carry approximately 24 A of current for the optimal flux.

Additionally, three circular coils are arranged along orthogonal directions to provide translational bias fields. Each coil has 100 turns of 1 mm diameter copper wire and an inner diameter of 150 mm. These coils are important for translating the node of the quadrupole field, which is the centre of the MOT. If this node is located above the aperture of the pyramid, the atoms are pushed through and the pyramid acts as a source of cold atoms. If the bias fields are adjusted to position the node above the mirrors, away from the aperture, the pyramid acts as a confining 3D MOT.

The cooling laser light enters through a DN100CF uncoated viewport directly onto the pyramid mirrors. The incident beam has a $1/e^2$ diameter of approximately 12 cm, completely covering the pyramid mirrors, with a power of ~ 175 mW. From this chamber the generated atomic flux passes through the pyramid aperture through a differential pumping tube, 5 mm in diameter and 17 mm long, and directly into the glass science cell.

The UHV chamber and cell are pumped by a 55 l/s ion pump⁵. The UHV chamber also contains an angle valve, to which a turbomolecular pump⁶ can be connected. Before breaking vacuum, the system was backfilled with argon through this valve to minimise exposure of the interior to air. The new rubidium source was connected, and then the turbo pump brought the system from atmospheric pressure down to pressures of around $\sim 10^{-6}$ mbar, at which point the ion pumps were turned on. From this point the vacuum system was wrapped in foil for insulation and baked in situ with heater tape. The temperature was brought up slowly to minimise gradients across sensitive parts such as the pyramid viewport or glass cell, less than 2 °C/min as per manufacturers' recommendations. The maximum temperature of the pyramid and UHV chamber was 70 – 80 °C, and the ion pumps were heated to ~ 150 °C. This temperature was maintained for two days. The valve to the turbo pump was closed and the heater tape turned off, and the system was left to cool and the pressure to

³Varian VacIon Plus Starcell 40 l/s

⁴Varian VacIon Plus Diode 55 l/s

⁵Varian VacIon Plus Starcell 55 l/s

⁶Pfeiffer HiCube 80 Eco

stabilise for another day. Both the pyramid and science chamber pressures were at the floor of what the ion pumps can measure, approximately 10^{-9} mbar.

Once the pressure had equalised, the rubidium valve was fully closed, cracking the ampoule. The breaking of the ampoule is loud and so confirmation that it has broken is easily audible. The ampoule released the contained argon under atmospheric pressure. After a few minutes the initial pressure spike subsided and the rubidium chamber was slowly heated to drive rubidium atoms into the pyramid chamber. The presence of rubidium in the chamber can be detected by sweeping the input laser light around the resonance and looking for fluorescence. For the first smaller rubidium chamber, DN16CF, no rubidium was observed until a few days passed, and the signal was weak. The new, DN40CF chamber produced a strong fluorescence signal after a day of milder heating, approximately 40°C . Once a sufficient rubidium pressure has been reached, as indicated by a strong fluorescence signal, the quadrupole is switched on and the bias fields adjusted to produce either a confined 3D MOT or a MOT source.

8.5 Characterising the MOTs

8.5.1 Operating the pyramid MOT source

An example picture of a small MOT cloud is shown in Fig. 8.7. The MOT coil current is 24 A which generates an axial magnetic field gradient of 12 G/cm which is empirically found to produce the largest recapture MOT when the pyramid is operated as a source.

It should be noted that the system operates as a confined pyramid MOT for a wide range of bias currents - because the aperture is very small, only a small range of the three bias coil currents will place the cloud over the aperture so that the system operates as a source.

The strong background scatter and fluorescence makes an accurate estimate of atom number difficult in the pyramid. We can however still plot the total image pixel counts vs time, which once normalised is proportional to the atom number. This plot is shown in Fig. 8.8b. At the pressures used to produce a high flux for loading the recapture MOT, the fitted $1/e$ loading time is 54 ms, producing a practically fully loaded MOT in just 250 ms. A MOT loading rate can be used to estimate the total pressure P in the chamber to within approximately a factor of 2 [180] using the formula

$$P = \frac{C}{\tau}, \quad (8.1)$$




Figure 8.7: Image of a 3D MOT in the pyramid chamber. Note the aperture of the pyramid as seen in the central dark spot. The central cloud can be seen just beneath the aperture and reflective pyramid mirrors produce four mirror images of the MOT cloud in each quadrant of the image. Fluorescing ^{87}Rb vapour produces a strong background signal.

where τ is the measured loading rate and the constant $C = 2.6 \times 10^{-8} \text{ mbar s}$. This gives an estimate of the pyramid chamber pressure as $\sim 5 \times 10^{-7} \text{ mbar}$. This is consistent with the vapour pressure of rubidium at room temperature [57]. This pressure would be much too high for the old apparatus, in which the cold cloud is translated relatively slowly through this background gas. In the pyramid source, the atoms are laser cooled and ejected into the science cell at a velocity of 20-30 m/s [177] and hence stay in the pyramid only $\sim 1 \text{ ms}$ after cooling, permitting a much higher pressure and hence loading rate.

To operate as a MOT source, the correct bias fields must be found. This can be done by sweeping over the values for the coil currents, translating the quadrupole node, and therefore the cloud. The aperture is small and when the cloud is positioned over it, the fluorescence of the MOT will disappear, and then return as the cloud moves past the aperture. An example of this sweep can be seen in Fig. 8.8a. The position of the cloud is most sensitive to the two bias fields that move the cloud perpendicular




Figure 8.8: a) Pixel counts from the fluorescing MOT cloud as a function of the two bias fields, which translate the cooled MOT cloud around within the pyramid. The direction of the bias fields are labelled Up/Down (U/D) and North/South (N/S), which are the directions perpendicular to the propagation of the MOT beam. The counts disappear when the cloud is positioned over the aperture and the atoms are ejected out of the pyramid. Since the cloud is larger than the aperture, the reduction in counts is continuous. b) Loading rate of the pyramid MOT. Accurate atom number counts are difficult to extract for the pyramid images so plotted instead are the normalised image pixel counts, proportional to atom number. The fitted $1/e$ loading time $\tau = 54.4$ ms.

to the direction of the cooling beam. The other bias field moves the cloud in and out of the pyramid. This axis is roughly aligned with the direction of the input beam, so has a weaker effect on the mode of operation of the MOT, but was still varied to optimise the atom number in the cell recapture MOT.

8.5.2 Loading the recapture MOT

Once the pyramid is operating as a MOT source, atoms can be recaptured in the glass cell 3D MOT and their fluorescence measured. The flux out of the pyramid has a moderately wide angle of divergence [176], allowing significant leniency with the alignment of the quadrupole node to the axis of the pyramid. In fact, no recapture MOT in the cell could be found until after displacing the quadrupole field with a small (~ 2 mm) offset to the pyramid axis. This is due to the small amount of light that propagates through the pyramid aperture, which pushes atoms out of the recapture MOT. Initially this displacement was implemented by shifting the entire breadboard,




Figure 8.9: a) Sweeping the pyramid bias fields to maximise the atom number (measured normalised total pixel counts) in the recapture MOT. Maximum atom number indicates the centre of the pyramid MOT is placed over the pyramid aperture. b) Plot of the loading curve of the recapture MOT in the glass cell. The fitted exponential fit gives a maximum atom number of 3.8 billion atoms with a $1/e$ loading rate of 1.95 s.

including optics and quadrupole coils, but this has subsequently been replaced by the use of magnetic bias fields to translate the node of the quadrupole field.

The operation of the source can be optimised by sweeping the pyramid bias fields (or other parameters such as quadrupole gradient and detuning) and maximising the atom number seen in the recapture MOT. An example of such a sweep is shown in Fig. 8.9a.

At maximum size, the 3D MOT cloud is sufficiently bright to be seen by the naked eye. The cloud reaches a maximum atom number of 3.8 billion atoms, with a $1/e$ loading time of 1.95 s, or a total duration of 6 s. This is both faster and larger than the 3D MOT on the current apparatus, which loads between 1 and 1.5 billion atoms in a total of 12 s. Fig. 8.10 shows a picture of a large MOT as taken by fluorescence imaging on the CCD camera, along with a picture taken by a mobile phone digital camera. Note that the brightness of an image taken with a digital camera is dependent on the IR filter in the camera; a high quality SLR camera could not detect the MOT, even though it is clearly visible to the human eye.

However, based on the properties of the cell MOT cloud, it was discovered that the pressure in the cell is relatively high. The lifetime of the MOT after loading stopped was much shorter than a background pressure at UHV should produce. In




Figure 8.10: a) Fluorescence image of the 3D MOT in the glass cell. b) Image of the 3D MOT, visible to the human eye. The amorphous shape of the MOT is caused by imbalances in the power or misalignment of the beams in these early observations.

particular, when the pyramid source is not operated (by, for example, shuttering the light to the pyramid), a small MOT can be observed loading from background vapour in the cell. The loading curve of this cloud is plotted in Fig. 8.11.

From the fitted loading rate of ~ 2 s, we can infer (using Eqn. 8.1) that the pressure is approximately 1×10^{-8} mbar. This is low enough to form a large MOT but much too high for useful experiments with a Bose-Einstein condensate, due to the reduced lifetime of the trapped cloud. The formation of a MOT from background gas suggests that a significant proportion of this gas is ^{87}Rb leaking through from the pyramid. The pressure ratio produced by the relatively short, wide tube connecting the pyramid and UHV chambers is approximately 100. This final glass cell pressure is consistent with this differential pumping ratio combined with a pyramid at the measured pressure of $\sim 5 \times 10^{-7}$ mbar, close to the vapour pressure of rubidium. The pyramid pressure could be decreased, but this would decrease the flux and hence the overall sequence time, hindering one of the goals of the new apparatus. To improve the differential pumping ratio, a new differential pumping tube with a lower conductance has been attached by Abel Beregi, as described in his thesis [178].

8.6 Towards RF-dressed trapping and BECs

The ultimate purpose of the new design was to provide a fast, reliable way to load atoms into a magnetic trap and produce a Bose-Einstein condensate. An impor-




Figure 8.11: Loading curve of a MOT in the glass cell with the pyramid source switched off. The resulting loading is solely from undesirable background ^{87}Rb in the glass cell.

tant and uncertain element of the improved experimental sequence is loading an RF-dressed shell trap directly from the MOT. Usually such experiments employ some intermediate stage trap, such as the Ioffe-Pritchard trap sometimes used in atom-chip experiments [181].

The mode-matching between the large, spherical MOT and the typically small, anisotropic shell trap is, at first impression, very poor. Additionally, the MOT is situated at the node of the quadrupole field, while the shell trap minimum sits at the resonance of the field with the dressing frequency, significantly below the node. This mismatch between the potentials might suggest a poor loading efficiency between the traps. However, the initial MOT number is so high that even a mediocre loading efficiency could provide enough atoms for runaway evaporation to a condensate.

The offset in position can be fixed by employing bias fields that are quickly switched, faster than atoms can respond. The bias fields shift the position of the MOT downwards, to the minimum of the intended shell potential. If the bias fields and MOT beams are switched off and the shell trap switched on faster than the atoms can move a significant distance, the atoms will already be situated at or around the




Figure 8.12: Absorption image of cloud of atoms loaded into the shell trap directly from the MOT. The distinctive ellipsoidal shape of the shell is clearly visible and matches the predicted resonant ellipsoid shown (black dashed line), with a secondary cloud in a different dressed $\tilde{m}_F$ state trapped at the centre of the quadrupole trap (blue dot). The vertical shell radius is $400\ \mu\text{m}$.

shell trap minimum. The atoms will move only under their own velocity and the effect of gravity. Atoms at the Doppler temperature of $146\ \mu\text{K}$ move at a velocity of approximately $0.15\ \text{m/s}$. If the fields are switched off in less than $1\ \text{ms}$, the atoms will travel less than $\sim 150\ \mu\text{m}$. For a vertical trap frequency of $\omega_z = 300\ \text{kHz}$, the direction of tightest confinement, this would correspond to adding in energy comparable to the existing thermal energy of the cloud, so that we do not cause dramatic heating. Since switching times of less than $1\ \text{ms}$ are achievable, this order of magnitude estimate suggests that this condition will be easily met, particularly since the shell will preferentially trap the cooler atoms which travel slower.

To improve the matching between the potentials, a relatively large resonant ellipsoid is needed to create a trapping potential of comparable size to the MOT cloud. To produce such a trap, low quadrupole gradients are required, and a large Rabi frequency (i.e. RF field amplitude) further reduces the trap frequency.

Confirmation that loading a shell trap directly from a 3D MOT in the cell is possible was recently obtained by Abel Beregi [178]. An absorption image of atoms in the shell trap is shown in Fig. 8.12. Note the curvature and anisotropy of the cloud that is typical of atomic clouds in shell traps with gravity.

Further work is needed to maximise the loading efficiency, and evaporative cooling must be performed to reach the condensate stage. As mentioned previously, the unexpectedly high pressure in the cell limits the lifetime of the shell to $< 10\ \text{s}$, which significantly hinders the effectiveness of evaporative cooling, and interferes with the

intended experiments. Changes to the vacuum system, such as the improved differential pumping, have been implemented by others to improve the lifetime of atoms in the shell trap.

8.7 Conclusion

This chapter described the design, construction and operation of a new cold-atoms apparatus, based on a pre-existing pyramid rubidium source and glass cell. The apparatus is intended to have much shorter MOT loading time and overall sequence time, and have a generally simpler experimental sequence and design. Some of these goals have already been demonstrated, as the MOT loads much faster than the previous apparatus and to a much larger atom number. The experimental sequence has already been simplified, loading directly from the MOT to the RF-dressed potential with no intermediate traps. As a result the new apparatus is faster and requires far less hardware and software complexity. At the time of the completion of this thesis the apparatus did not reach the stage of gathering scientific data; further improvements will continue to be made to the apparatus by others, with the expectation that it will shortly be capable of producing Bose-Einstein condensates and be ready for cold-atom research.

Chapter 9

Conclusion

This thesis describes the experimental study of 2D Bose gases in disordered potentials. Firstly, using the technique of MRF-dressed potentials, we load atoms into a 2D double-well and subject the cloud to a controllable pointlike disorder. We describe the method of implementing this disorder, using an optical dipole potential created by a high-intensity laser beam. This light is shaped by reflection off of a digital micromirror device consisting of an array of many small mirrors that can create many different potentials, including a pointlike disorder potential that can be created with a different random configuration for every sequence.

After splitting the quantum gas into two separate clouds and allowing them to equilibrate, we release the clouds and observe the resulting interference pattern. This experimental procedure measures the relative phase between the clouds and the phase fluctuations within each condensate. The correlation function of the condensate derived from analysis of the experimental data identifies the state of the cloud as a superfluid or normal fluid.

We vary the strength of the disorder potential and observe that at a critical disorder strength the cloud crosses from the superfluid to the normal phase at a temperature lower than that of the disorder-free case. We estimate the critical exponent of the correlation function at the transition point and find that it is consistent with the exponent in the absence of the disorder, indicating that the transition is still of the BKT type and that, in this parameter range at least, disorder does not modify the universality class of the transition. We find the same result for two different scatterer densities of pointlike disorder. We further observe a smooth reduction in the coherence of the cloud as measured by the condensate fraction with increasing disorder strength and disorder density.

At disorder strengths higher than the critical disorder, we find that the critical exponent is reduced again and drops below the critical value, unexpectedly. However,

we observe no concurrent reduction in the vortex density, nor an increase in the condensate fraction at this disorder strength. We suggest a possible explanation may be related to the effect of vortex pinning, which is stronger at higher disorder strengths. The vortices may cluster around the scattering sites and leave a connected superfluid region which contains fewer vortices. In addition, the presence of strong disorder may increase the density of the cloud outside the disorder sites, increasing the phase-space density and therefore the coherence. This conjecture requires further experimental and theoretical consideration.

We then examine superfluid flow around both single barriers and disordered potentials in 2D condensates. By inducing collective dipole oscillations of the cloud around the obstacle, we measure the rate of dissipation by observing damped sinusoidal oscillations of the cloud in time-of-flight expansion. We find that no observable dissipation occurs for flow around a single obstacle below a critical velocity, which is dependent on the height of the barrier. Faster flow velocities above this point lead to higher damping rates, except for an apparent resonance at certain velocities where the damping is much higher than other similar velocities. In the case of flow around disordered potentials, we observe no critical velocity below which the dissipation vanishes. The damping rate of the disordered potential is much greater than a single barrier even for lower peak potential heights. Once again, we observe unexpected sharp peaks in the damping rate.

We then discussed the development of a software simulation package, AtomECS, which can be utilised in the design of cold atoms experiments. This package was used to assist with the design of a new cold atoms apparatus, which is intended to be faster, more stable, and capable of higher atoms numbers than our existing apparatus on which the above experiments were performed. Finally, we described the first stages of its design and construction, including the vacuum work and initial laser cooled clouds in magneto-optical traps. In the near future this apparatus will be capable of running many experimental sequences in a short period of time, a particularly useful feature for studying many-body physics which typically require large amounts of data to reduce statistical uncertainties.

9.1 Future experiments

9.1.1 Many-body localisation in 2D Bose gases

As previously mentioned, MBL is an area of intense theoretical and experimental research [15, 17, 19, 93, 110, 112]. In 2D Bose gases, it is predicted that, at sufficiently

high disorder strengths, a transition to an insulating state exists at finite temperature [3, 5, 20, 182]. However, the existence of this state and its stability remain open questions [15, 19, 20, 22]. The experimental techniques developed here provide the means of experimentally studying this transition.

An example of a possible experimental approach would be quenching the 2D superfluid so that it is above the superfluid-normal transition. This can be done by rapidly splitting the cloud and loading into the double-well, reducing the density while maintaining a constant temperature, as we have previously performed [53]. By conducting this quench in the number density in the presence of disorder, we may observe how the relaxation and thermalisation dynamics are affected. We may expect that for high enough disorder strengths, above the localisation transition, the cloud does not thermalise and the correlation function does not evolve over time. Such absence of thermalisation is a hallmark of MBL states and could constitute powerful evidence for its existence. Furthermore, the lifetime of this insulating state might be capable of determining the stability of the MBL phase if it can be distinguished from finite lifetime in the trap due to other considerations, such as RF noise or non-adiabatic loss in the tight 2D trap.

9.1.2 Potential shaping

The high degree of controllability of the optical dipole potential permits potential shapes other than pointlike disorder or individual Gaussian scatterers. For example, by applying a ring potential around the atoms, we can create a hard wall and cut off the low-density edge of the usual density distribution in a harmonic trap. In particular, for small radii of ring potentials, the remaining inhomogeneity inside the ring can be further reduced. In this way we can produce a ‘quasi-homogeneous’ potential, with an approximately uniform density. Homogeneous systems are more easily compared to theoretical models, since inhomogeneity in the density often obscures both theoretical predictions and clear observation of thermodynamic properties in experiment.

One experiment of interest is the measurement of the correlations in density-density fluctuations for 2D clouds in short time-of-flight. The phase fluctuations in the 2D cloud translate into density fluctuations in TOF and the behaviour of these correlations can be related to the scaling exponent of the superfluid phase [183]. The inhomogeneity of a harmonic trap complicates this analysis. A hard wall imposed by the dipole potential can significantly reduce this inhomogeneity and thus produce an ‘quasi-homogeneous’ cloud. The cloud can then be allowed to expand as it falls in

TOF. After ~ 5 ms, the cloud has dropped to the focus of the imaging system (see Section 3.4.1) and can be imaged with a thin horizontal repumping slice. This method can produce clean and high-quality estimates of the density noise correlations at short time-of-flight, providing another method to probe the physics of the 2D condensed gas. Experimental work on this measurement scheme was started after the data for this thesis was gathered and is providing another valuable probe of the internal state of 2D clouds.

The hard-walled ring potential can also be combined with the MRF dressing potentials to produce a double-well, quasi-homogeneous potential. This may produce clearer results for both the experimental schemes already discussed and others.

Other possible applications of potential shaping are studying superfluid flow around a variety of differently shaped obstacles, or comparing the effect of randomly placed scattering sites with scatterers ordered into a regular lattice. This latter case may clarify what effect localisation precisely has on a system, since localisation occurs only in the case of disordered potentials.

9.1.3 Tunnelling in double-well potentials

Quantum tunnelling between two wells will lead to another wide range of interesting physics. Coupled quantum gases can give rise to Josephson oscillations, which have been observed in both 1D [184] and 3D [185] Bose gases, and 2D Fermi gases [186]. It has also been predicted that a coupling between two 2D superfluids modifies the phase diagram in the region of the superfluid-normal transition [187, 188].

The creation of 2D double-well potentials with a high coupling between the two clouds is not entirely straightforward in our MRF-dressed potentials. The creation of 2D double-wells necessitates a high double-well barrier, since this barrier provides half of the tight axial confinement felt by each well. This same condition also reduces the coupling strength between the wells, which must therefore be compensated for by having a very small separation between the wells. These close separations lead to widely spaced fringes during matter-wave interference, so that the analysis methods used in this thesis may struggle to identify the local relative phase with confidence when the fringe spacing becomes comparable to the size of the cloud in time-of-flight expansion. This problem was considered in more detail in the thesis of Shinichi Sunami [53].

Furthermore, the close well separation leads to a significant atom density in the barrier region, and the gas can not be accurately considered to be two distinct clouds,

which may invalidate some of the more commonly used theoretical models of tunnelling in Bose gases such as the two-mode model [185, 189]. Nevertheless, it is an area of research into 2D Bose gases which remains relatively rarely studied experimentally, and there are a number of mechanisms at play in two dimensions - such as unbound vortices or phase fluctuations across the cloud - that are likely to produce rich and interesting dynamics.

References

1. Pan, S. *et al.* Microscopic electronic inhomogeneity in the high- T_c superconductor $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+x}$. *Nature* **413**, 282–285 (2001).
2. Cren, T., Roditchev, D., Sacks, W. & Klein, J. Nanometer scale mapping of the density of states in an inhomogeneous superconductor. *Europhysics Letters (EPL)* **54**, 84–90. <https://doi.org/10.1209/epl/i2001-00232-4> (Apr. 2001).
3. Bertoli, G., Michal, V. P., Altshuler, B. L. & Shlyapnikov, G. V. Finite-Temperature Disordered Bosons in Two Dimensions. *Phys. Rev. Lett.* **121**, 030403. <https://link.aps.org/doi/10.1103/PhysRevLett.121.030403> (3 July 2018).
4. Bourdel, T. Phase diagrams of two-dimensional and three-dimensional disordered Bose gases in the local density approximation. *Phys. Rev. A* **86**, 063626. <https://link.aps.org/doi/10.1103/PhysRevA.86.063626> (6 Dec. 2012).
5. Carleo, G., Bo  ris, G., Holzmann, M. & Sanchez-Palencia, L. Universal Superfluid Transition and Transport Properties of Two-Dimensional Dirty Bosons. *Phys. Rev. Lett.* **111**, 050406. <https://link.aps.org/doi/10.1103/PhysRevLett.111.050406> (5 Aug. 2013).
6. Sac  p  , B. *et al.* Disorder-Induced Inhomogeneities of the Superconducting State Close to the Superconductor-Insulator Transition. *Phys. Rev. Lett.* **101**, 157006. <https://link.aps.org/doi/10.1103/PhysRevLett.101.157006> (15 Oct. 2008).
7. Sambandamurthy, G., Engel, L. W., Johansson, A., Peled, E. & Shahar, D. Experimental Evidence for a Collective Insulating State in Two-Dimensional Superconductors. *Phys. Rev. Lett.* **94**, 017003. <https://link.aps.org/doi/10.1103/PhysRevLett.94.017003> (1 Jan. 2005).
8. Krauth, W., Trivedi, N. & Ceperley, D. Superfluid-insulator transition in disordered boson systems. *Phys. Rev. Lett.* **67**, 2307–2310. <https://link.aps.org/doi/10.1103/PhysRevLett.67.2307> (17 Oct. 1991).
9. Neverov, V., A.E., L., Krasavin, A., Vagov, A. & Croitoru, M. Correlated disorder as a way towards robust superconductivity. *Communication Physics* **5**, 177. <https://doi.org/10.1038/s42005-022-00933-z> (2022).
10. Anderson, P. W. Absence of Diffusion in Certain Random Lattices. *Phys. Rev.* **109**, 1492–1505. <https://link.aps.org/doi/10.1103/PhysRev.109.1492> (5 Mar. 1958).

11. Giamarchi, T. & Schulz, H. J. Anderson localization and interactions in one-dimensional metals. *Phys. Rev. B* **37**, 325–340. <https://link.aps.org/doi/10.1103/PhysRevB.37.325> (1 Jan. 1988).
12. Lee, C., Rai, A., Noh, C. & Angelakis, D. G. Probing the effect of interaction in Anderson localization using linear photonic lattices. *Phys. Rev. A* **89**, 023823. <https://link.aps.org/doi/10.1103/PhysRevA.89.023823> (2 Feb. 2014).
13. Shepelyansky, D. L. Delocalization of quantum chaos by weak nonlinearity. *Phys. Rev. Lett.* **70**, 1787–1790. <https://link.aps.org/doi/10.1103/PhysRevLett.70.1787> (12 Mar. 1993).
14. Abrahams, E., Anderson, P. W., Licciardello, D. C. & Ramakrishnan, T. V. Scaling Theory of Localization: Absence of Quantum Diffusion in Two Dimensions. *Phys. Rev. Lett.* **42**, 673–676. <https://link.aps.org/doi/10.1103/PhysRevLett.42.673> (10 Mar. 1979).
15. Basko, D., Aleiner, I. & Altshuler, B. Metal–insulator transition in a weakly interacting many-electron system with localized single-particle states. *Annals of Physics* **321**, 1126–1205. ISSN: 0003-4916. <https://www.sciencedirect.com/science/article/pii/S0003491605002630> (2006).
16. Gornyi, I. V., Mirlin, A. D. & Polyakov, D. G. Interacting Electrons in Disordered Wires: Anderson Localization and Low- T Transport. *Phys. Rev. Lett.* **95**, 206603. <https://link.aps.org/doi/10.1103/PhysRevLett.95.206603> (20 Nov. 2005).
17. Carleo, G., Becca, F., Scihro, M. & Fabrizio, M. Localization and Glassy Dynamics Of Many-Body Quantum Systems. *Sci Rep* **2**, 243. <https://doi.org/10.1038/srep00243> (2012).
18. Schiulaz, M. & Müller, M. Ideal quantum glass transitions: Many-body localization without quenched disorder. *AIP Conference Proceedings* **1610**, 11–23. eprint: <https://aip.scitation.org/doi/pdf/10.1063/1.4893505>. <https://aip.scitation.org/doi/abs/10.1063/1.4893505> (2014).
19. Gornyi, I. V., Mirlin, A. D., Müller, M. & Polyakov, D. G. Absence of many-body localization in a continuum. *Annalen der Physik* **529**, 1600365. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.201600365>. <https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.201600365> (2017).
20. Nandkishore, R. Many-body localization and delocalization in the two-dimensional continuum. *Phys. Rev. B* **90**, 184204. <https://link.aps.org/doi/10.1103/PhysRevB.90.184204> (18 Nov. 2014).
21. Scardicchio, A. & Thiery, T. *Perturbation theory approaches to Anderson and Many-Body Localization: some lecture notes* 2017. <https://arxiv.org/abs/1710.01234>.

22. De Roeck, W. & Huveneers, F. ç. Stability and instability towards delocalization in many-body localization systems. *Phys. Rev. B* **95**, 155129. <https://link.aps.org/doi/10.1103/PhysRevB.95.155129> (15 Apr. 2017).
23. Berezinskii, V. L. Destruction of Long-range Order in One-dimensional and Two-dimensional Systems Possessing a Continuous Symmetry Group. II. Quantum Systems. **34**, 610 (1971).
24. Kosterlitz, J. M. & Thouless, D. J. Ordering, metastability and phase transitions in two-dimensional systems. *Journal of Physics C Solid State Physics* **6**, 1181–1203 (Apr. 1973).
25. Hadzibabic, Z., Krüger, P., Cheneau, M., Battelier, B. & Dalibard, J. Berezinskii-Kosterlitz-Thouless crossover in a trapped atomic gas. *Nature* **441**, 1118–1121 (June 2006).
26. Sunami, S. *et al.* Observation of the Berezinskii-Kosterlitz-Thouless Transition in a Two-Dimensional Bose Gas via Matter-Wave Interferometry. *Phys. Rev. Lett.* **128**, 250402. <https://link.aps.org/doi/10.1103/PhysRevLett.128.250402> (25 June 2022).
27. McQueeney, D., Agnolet, G. & Reppy, J. D. Surface Superfluidity in Dilute ^4He - ^3He Mixtures. *Phys. Rev. Lett.* **52**, 1325–1328. <https://link.aps.org/doi/10.1103/PhysRevLett.52.1325> (15 Apr. 1984).
28. Murthy, P. A. *et al.* Observation of the Berezinskii-Kosterlitz-Thouless Phase Transition in an Ultracold Fermi Gas. *Phys. Rev. Lett.* **115**, 010401. <https://link.aps.org/doi/10.1103/PhysRevLett.115.010401> (1 June 2015).
29. Zhao, W. *et al.* Evidence for Berezinskii–Kosterlitz–Thouless transition in atomically flat two-dimensional Pb superconducting films. *Solid State Communications* **165**, 59–63. ISSN: 0038-1098. <https://www.sciencedirect.com/science/article/pii/S0038109813001993> (2013).
30. Baturina, T. I. *et al.* Superconducting phase transitions in ultrathin TiN films. *EPL (Europhysics Letters)* **97**, 17012. <https://doi.org/10.1209/0295-5075/97/17012> (Jan. 2012).
31. Cronin, A. D., Schmiedmayer, J. & Pritchard, D. E. Optics and interferometry with atoms and molecules. *Rev. Mod. Phys.* **81**, 1051–1129. <https://link.aps.org/doi/10.1103/RevModPhys.81.1051> (3 July 2009).
32. Barker, A. J. *et al.* Coherent splitting of two-dimensional Bose gases in magnetic potentials. *New Journal of Physics* **22**, 103040. <https://doi.org/10.1088/1367-2630/abbcd> (Oct. 2020).
33. Schumm, T. *et al.* Matter-wave interferometry in a double well on an atom chip. *Nature Physics* **1**, 57–62. <https://doi.org/10.1038/nphys125> (2005).
34. Zobay, O. & Garraway, B. M. Two-Dimensional Atom Trapping in Field-Induced Adiabatic Potentials. *Phys. Rev. Lett.* **86**, 1195–1198. <https://link.aps.org/doi/10.1103/PhysRevLett.86.1195> (7 Feb. 2001).

35. Merloti, K. *et al.* A two-dimensional quantum gas in a magnetic trap. *New Journal of Physics* **15**, 033007 (Mar. 2013).
36. Harte, T. L. *et al.* Ultracold atoms in multiple radio-frequency dressed adiabatic potentials. *Phys. Rev. A* **97**, 013616. <https://link.aps.org/doi/10.1103/PhysRevA.97.013616> (1 Jan. 2018).
37. Courteille, P. W. *et al.* Highly versatile atomic micro traps generated by multi-frequency magnetic field modulation. *Journal of Physics B: Atomic, Molecular and Optical Physics* **39**, 1055–1064. <https://doi.org/10.1088/0953-4075/39/5/005> (Feb. 2006).
38. Heathcote, W. H., Nugent, E., Sheard, B. T. & Foot, C. J. A ring trap for ultracold atoms in an RF-dressed state. *New Journal of Physics* **10**, 043012. <https://doi.org/10.1088/1367-2630/10/4/043012> (Apr. 2008).
39. Morizot, O., Colombe, Y., Lorent, V., Perrin, H. & Garraway, B. M. Ring trap for ultracold atoms. *Phys. Rev. A* **74**, 023617. <https://link.aps.org/doi/10.1103/PhysRevA.74.023617> (2 Aug. 2006).
40. Luksch, K. *Investigating the spectrum of atoms in multiple-radiofrequency dressed potentials for coherent splitting of quantum gases* PhD thesis (University of Oxford, 2019).
41. Bentine, E. *Atomic mixtures in radiofrequency dressed potentials* PhD thesis (University of Oxford, 2018).
42. Harte, T. *Ultracold atoms in dressed potentials* PhD thesis (University of Oxford, 2017).
43. Majorana, E. Atomi orientati in campo magnetico variabile. *Il Nuovo Cimento* **9**, 43–50 (1932).
44. Dubessy, R. *et al.* Rubidium-87 Bose-Einstein condensate in an optically plugged quadrupole trap. *Phys. Rev. A* **85**, 013643. <https://link.aps.org/doi/10.1103/PhysRevA.85.013643> (1 Jan. 2012).
45. Petrich, W., Anderson, M. H., Ensher, J. R. & Cornell, E. A. Stable, Tightly Confining Magnetic Trap for Evaporative Cooling of Neutral Atoms. *Phys. Rev. Lett.* **74**, 3352–3355. <https://link.aps.org/doi/10.1103/PhysRevLett.74.3352> (17 Apr. 1995).
46. Anderson, M. H., Ensher, J. R., Matthews, M. R., Wieman, C. E. & Cornell, E. A. Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor. *Science* **269**, 198–201. ISSN: 0036-8075. eprint: <https://science.sciencemag.org/content/269/5221/198.full.pdf>. <https://science.sciencemag.org/content/269/5221/198> (1995).
47. Pethick, C. J. & Smith, H. *Bose–Einstein Condensation in Dilute Gases* 2nd ed. (Cambridge University Press, 2008).
48. Perrin, H. & Garraway, B. M. Chapter Four - Trapping Atoms With Radio Frequency Adiabatic Potentials. *Advances in Atomic Molecular and Optical Physics* **66**, 181–262 (2017).

49. Cohen-Tannoudji, C. & Reynaud, S. Dressed-atom description of resonance fluorescence and absorption spectra of a multi-level atom in an intense laser beam. *Journal of Physics B: Atomic and Molecular Physics* **10**, 345–363. <https://doi.org/10.1088/0022-3700/10/3/005> (Feb. 1977).
50. Messiah, A. in *Quantum mechanics* (Wiley New York, 1966).
51. Barker, A. J. *et al.* Realising a species-selective double well with multiple-radiofrequency-dressed potentials. *Journal of Physics B: Atomic, Molecular and Optical Physics* **53**, 155001. <https://doi.org/10.1088/1361-6455/ab9152> (June 2020).
52. Foot, C. J. *Atomic physics* (Oxford University Press, 2005).
53. Sunami, S. *Non-equilibrium dynamics in two-dimensional quantum systems* PhD thesis (University of Oxford, 2022).
54. Barker, A. J. *Bose gases in double-well potentials* PhD thesis (University of Oxford, 2020).
55. R. Jördens D. Nadlinger, e. a. *ARTIQ* Online. <https://m-labs.hk/experiment-control/artiq/>.
56. Moore, R. W. G. *et al.* Measurement of vacuum pressure with a magneto-optical trap: A pressure-rise method. *Review of Scientific Instruments* **86**, 093108. eprint: <https://doi.org/10.1063/1.4928154>. <https://doi.org/10.1063/1.4928154> (2015).
57. Steck, D. A. 'Rubidium 87 D Line Data' Online at: 2021. <http://steck.us/alkalidata>.
58. Rhodes, D. P. *et al.* Atom guiding along high order Laguerre–Gaussian light beams formed by spatial light modulation. *Journal of Modern Optics* **53**, 547–556. eprint: <https://doi.org/10.1080/09500340500429960>. <https://doi.org/10.1080/09500340500429960> (2006).
59. Gaunt, A. L., Schmidutz, T. F., Gotlibovych, I., Smith, R. P. & Hadzibabic, Z. Bose-Einstein Condensation of Atoms in a Uniform Potential. *Phys. Rev. Lett.* **110**, 200406. <https://link.aps.org/doi/10.1103/PhysRevLett.110.200406> (20 May 2013).
60. Hadzibabic, Z. & Dalibard, J. Two-dimensional Bose fluids: An atomic physics perspective. *arXiv*, arXiv:0912.1490 (Dec. 2009).
61. Mermin, N. D. & Wagner, H. Absence of Ferromagnetism or Antiferromagnetism in One- or Two-Dimensional Isotropic Heisenberg Models. *Phys. Rev. Lett.* **17**, 1133–1136 (22 Nov. 1966).
62. Bhaduri, R. K., Reimann, S. M., Viefers, S., Choudhury, A. G. & Srivastava, M. K. The effect of interactions on Bose-Einstein condensation in a quasi two-dimensional harmonic trap. *Journal of Physics B: Atomic, Molecular and Optical Physics* **33**, 3895–3903. <https://doi.org/10.1088/0953-4075/33/19/304> (Sept. 2000).

63. Fletcher, R. J. *et al.* Connecting Berezinskii-Kosterlitz-Thouless and BEC Phase Transitions by Tuning Interactions in a Trapped Gas. *Phys. Rev. Lett.* **114**, 255302. <https://link.aps.org/doi/10.1103/PhysRevLett.114.255302> (25 June 2015).
64. Holzmann, M., Chevallier, M. & Krauth, W. Universal correlations and coherence in quasi-two-dimensional trapped Bose gases. *Phys. Rev. A* **81**, 043622. <https://link.aps.org/doi/10.1103/PhysRevA.81.043622> (4 Apr. 2010).
65. Prokof'ev, N. & Svistunov, B. Two-dimensional weakly interacting Bose gas in the fluctuation region. *Phys. Rev. A* **66**, 043608. <https://link.aps.org/doi/10.1103/PhysRevA.66.043608> (4 Oct. 2002).
66. Hung, C., Zhang, X., Gemelke, N. & Chin, C. Observation of scale invariance and universality in two-dimensional Bose gases. *Nature* **470**, 236239. <https://doi.org/10.1038/nature09722> (2011).
67. Prokof'ev, N., Ruebenacker, O. & Svistunov, B. Critical Point of a Weakly Interacting Two-Dimensional Bose Gas. *Phys. Rev. Lett.* **87**, 270402. <https://link.aps.org/doi/10.1103/PhysRevLett.87.270402> (27 Dec. 2001).
68. Nelson, D. R. & Kosterlitz, J. M. Universal Jump in the Superfluid Density of Two-Dimensional Superfluids. *Phys. Rev. Lett.* **39**, 1201–1205. <https://link.aps.org/doi/10.1103/PhysRevLett.39.1201> (19 Nov. 1977).
69. Minnhagen, P. & Warren, G. G. Superfluid density of a two-dimensional fluid. *Phys. Rev. B* **24**, 2526–2532. <https://link.aps.org/doi/10.1103/PhysRevB.24.2526> (5 Sept. 1981).
70. Christodoulou, P. *et al.* Observation of first and second sound in a BKT superfluid. *Nature* **594**, 191–194. <https://doi.org/10.1038/s41586-021-03537-9> (2021).
71. Pilati, S., Giorgini, S. & Prokof'ev, N. Critical Temperature of Interacting Bose Gases in Two and Three Dimensions. *Phys. Rev. Lett.* **100**, 140405. <https://link.aps.org/doi/10.1103/PhysRevLett.100.140405> (14 Apr. 2008).
72. Gawryluk, K. & Brewczyk, M. Signatures of a universal jump in the superfluid density of a two-dimensional Bose gas with a finite number of particles. *Phys. Rev. A* **99**, 033615. <https://link.aps.org/doi/10.1103/PhysRevA.99.033615> (3 Mar. 2019).
73. Morong, W. & DeMarco, B. Simulation of Anderson localization in two-dimensional ultracold gases for pointlike disorder. *Phys. Rev. A* **92**, 023625. <https://link.aps.org/doi/10.1103/PhysRevA.92.023625> (2 Aug. 2015).
74. Goodman, J. W. Some fundamental properties of speckle*. *J. Opt. Soc. Am.* **66**, 1145–1150. <http://opg.optica.org/abstract.cfm?URI=josa-66-11-1145> (Nov. 1976).

75. Pilati, S., Giorgini, S., Modugno, M. & Prokof'ev, N. Dilute Bose gas with correlated disorder: a path integral Monte Carlo study. *New Journal of Physics* **12**, 073003. <https://doi.org/10.1088/1367-2630/12/7/073003> (July 2010).
76. Pezzé, L. *et al.* Regimes of classical transport of cold gases in a two-dimensional anisotropic disorder. *New Journal of Physics* **13**, 095015. <https://doi.org/10.1088/1367-2630/13/9/095015> (Sept. 2011).
77. Albert, M., Paul, T., Pavloff, N. & Leboeuf, P. Dipole Oscillations of a Bose-Einstein Condensate in the Presence of Defects and Disorder. *Phys. Rev. Lett.* **100**, 250405. <https://link.aps.org/doi/10.1103/PhysRevLett.100.250405> (25 June 2008).
78. Delande, D. & Orso, G. Mobility Edge for Cold Atoms in Laser Speckle Potentials. *Phys. Rev. Lett.* **113**, 060601. <https://link.aps.org/doi/10.1103/PhysRevLett.113.060601> (6 Aug. 2014).
79. Sucu, S., Aktas, S., Okan, S. E., Akdeniz, Z. & Vignolo, P. Anderson localization in optical lattices with speckle disorder. *Phys. Rev. A* **84**, 065602. <https://link.aps.org/doi/10.1103/PhysRevA.84.065602> (6 Dec. 2011).
80. Semmler, D., Wernsdorfer, J., Bissbort, U., Byczuk, K. & Hofstetter, W. Localization of correlated fermions in optical lattices with speckle disorder. *Phys. Rev. B* **82**, 235115. <https://link.aps.org/doi/10.1103/PhysRevB.82.235115> (23 Dec. 2010).
81. Pilati, S. & Pieri, P. Supervised machine learning of ultracold atoms with speckle disorder. *Scientific Reports* **9**, 5613. <https://doi.org/10.1038/s41598-019-42125-w> (2019).
82. Meldgin, C. *et al.* Probing the Bose glass–superfluid transition using quantum quenches of disorder. *Nature Physics* **12**, 646649. <https://doi.org/10.1038/nphys3695> (7 2016).
83. Allard, B. *et al.* Effect of disorder close to the superfluid transition in a two-dimensional Bose gas. *Phys. Rev. A* **85**, 033602. <https://link.aps.org/doi/10.1103/PhysRevA.85.033602> (3 Mar. 2012).
84. Nagler, B. *et al.* Dipole oscillations of fermionic quantum gases along the BEC-BCS crossover in disordered potentials. *Phys. Rev. A* **101**, 053633. <https://link.aps.org/doi/10.1103/PhysRevA.101.053633> (5 May 2020).
85. Billy, J., Josse, V. & Zuo, Z. *e. a.* Direct observation of Anderson localization of matter waves in a controlled disorder. *Nature* **453**, 891894. <https://doi.org/10.1038/nature07000> (2008).
86. Volchkov, V. V. *et al.* Measurement of Spectral Functions of Ultracold Atoms in Disordered Potentials. *Phys. Rev. Lett.* **120**, 060404. <https://link.aps.org/doi/10.1103/PhysRevLett.120.060404> (6 Feb. 2018).

87. Robert-de-Saint-Vincent, M. *et al.* Anisotropic 2D Diffusive Expansion of Ultracold Atoms in a Disordered Potential. *Phys. Rev. Lett.* **104**, 220602. <https://link.aps.org/doi/10.1103/PhysRevLett.104.220602> (22 June 2010).
88. Kondov, S. S., McGehee, W. R., Zirbel, J. J. & DeMarco, B. Three-Dimensional Anderson Localization of Ultracold Matter. *Science* **334**, 66–68. eprint: <https://www.science.org/doi/pdf/10.1126/science.1209019>. <https://www.science.org/doi/abs/10.1126/science.1209019> (2011).
89. White, D. *et al.* Observation of two-dimensional Anderson localisation of ultracold atoms. *Nat Commun* **11**, 4942. <https://doi.org/10.1038/s41467-020-18652-w> (1 2020).
90. Mendelson, K. S. Percolation threshold of a class of correlated lattices. *Phys. Rev. E* **56**, 6586–6588. <https://link.aps.org/doi/10.1103/PhysRevE.56.6586> (6 Dec. 1997).
91. Essam, J. W. Percolation theory. *Rep. Prog. Phys* **43**, 833. <https://doi.org/10.1088/0034-4885/43/7/001> (1980).
92. Weinrib, A. Percolation threshold of a two-dimensional continuum system. *Phys. Rev. B* **26**, 1352–1361. <https://link.aps.org/doi/10.1103/PhysRevB.26.1352> (3 Aug. 1982).
93. Abanin, D. A., Altman, E., Bloch, I. & Serbyn, M. Colloquium: Many-body localization, thermalization, and entanglement. *Rev. Mod. Phys.* **91**, 021001. <https://link.aps.org/doi/10.1103/RevModPhys.91.021001> (2 May 2019).
94. Nandkishore, R. & Huse, D. A. Many-Body Localization and Thermalization in Quantum Statistical Mechanics. *Annual Review of Condensed Matter Physics* **6**, 15–38. eprint: <https://doi.org/10.1146/annurev-conmatphys-031214-014726>. <https://doi.org/10.1146/annurev-conmatphys-031214-014726> (2015).
95. Alet, F. & Laflorencie, N. Many-body localization: An introduction and selected topics. *Comptes Rendus Physique* **19**. Quantum simulation / Simulation quantique, 498–525. ISSN: 1631-0705. <https://www.sciencedirect.com/science/article/pii/S163107051830032X> (2018).
96. Deutsch, J. M. Quantum statistical mechanics in a closed system. *Phys. Rev. A* **43**, 2046–2049. <https://link.aps.org/doi/10.1103/PhysRevA.43.2046> (4 Feb. 1991).
97. Srednicki, M. Chaos and quantum thermalization. *Phys. Rev. E* **50**, 888–901. <https://link.aps.org/doi/10.1103/PhysRevE.50.888> (2 Aug. 1994).
98. Rigol, M., Dunjko, V. & Olshanii, M. Thermalization and its mechanism for generic isolated quantum systems. *Nature* **452**, 854–858 (2008).
99. Santos, L. F. & Rigol, M. Onset of quantum chaos in one-dimensional bosonic and fermionic systems and its relation to thermalization. *Phys. Rev. E* **81**, 036206. <https://link.aps.org/doi/10.1103/PhysRevE.81.036206> (3 Mar. 2010).

100. Rigol, M. & Santos, L. F. Quantum chaos and thermalization in gapped systems. *Phys. Rev. A* **82**, 011604. <https://link.aps.org/doi/10.1103/PhysRevA.82.011604> (1 July 2010).
101. Neuenhahn, C. & Marquardt, F. Thermalization of interacting fermions and delocalization in Fock space. *Phys. Rev. E* **85**, 060101. <https://link.aps.org/doi/10.1103/PhysRevE.85.060101> (6 June 2012).
102. Mondaini, R., Fratus, K. R., Srednicki, M. & Rigol, M. Eigenstate thermalization in the two-dimensional transverse field Ising model. *Phys. Rev. E* **93**, 032104. <https://link.aps.org/doi/10.1103/PhysRevE.93.032104> (3 Mar. 2016).
103. Thouless, D. J. A relation between the density of states and range of localization for one dimensional random systems. *Journal of Physics C: Solid State Physics* **5**, 77–81. <https://doi.org/10.1088/0022-3719/5/1/010> (Jan. 1972).
104. Herbert, D. C. & Jones, R. Localized states in disordered systems. *Journal of Physics C: Solid State Physics* **4**, 1145–1161 (1971).
105. Roati, G. *et al.* Anderson localization of a non-interacting Bose–Einstein condensate. *Nature* **453**, 895898. <https://doi.org/10.1038/nature07071> (2008).
106. Manai, I. *et al.* Experimental Observation of Two-Dimensional Anderson Localization with the Atomic Kicked Rotor. *Phys. Rev. Lett.* **115**, 240603. <https://link.aps.org/doi/10.1103/PhysRevLett.115.240603> (24 Dec. 2015).
107. Jendrzejewski, F. *et al.* Three-dimensional localization of ultracold atoms in an optical disordered potential. *Nature Physics* **8**, 398403. <https://doi.org/10.1038/nphys2256> (2012).
108. Serbyn, M., Papić, Z. & Abanin, D. A. Universal Slow Growth of Entanglement in Interacting Strongly Disordered Systems. *Phys. Rev. Lett.* **110**, 260601. <https://link.aps.org/doi/10.1103/PhysRevLett.110.260601> (26 June 2013).
109. Bauer, B. & Nayak, C. Area laws in a many-body localized state and its implications for topological order. *J. Stat. Mech.* <https://doi.org/10.1088/1742-5468/2013/09/P09005> (2013).
110. Lukin, A. *et al.* Probing entanglement in a many-body localized system. *Science* **364**, 256–260. eprint: <https://www.science.org/doi/pdf/10.1126/science.aau0818>. <https://www.science.org/doi/abs/10.1126/science.aau0818> (2019).
111. Serbyn, M., Papić, Z. & Abanin, D. A. Local Conservation Laws and the Structure of the Many-Body Localized States. *Phys. Rev. Lett.* **111**, 127201. <https://link.aps.org/doi/10.1103/PhysRevLett.111.127201> (12 Sept. 2013).

112. Huse, D. A., Nandkishore, R. & Oganesyan, V. Phenomenology of fully many-body-localized systems. *Phys. Rev. B* **90**, 174202. <https://link.aps.org/doi/10.1103/PhysRevB.90.174202> (17 Nov. 2014).
113. Pekker, D., Refael, G., Altman, E., Demler, E. & Oganesyan, V. Hilbert-Glass Transition: New Universality of Temperature-Tuned Many-Body Dynamical Quantum Criticality. *Phys. Rev. X* **4**, 011052. <https://link.aps.org/doi/10.1103/PhysRevX.4.011052> (1 Mar. 2014).
114. He, K., Santos, L. F., Wright, T. M. & Rigol, M. Single-particle and many-body analyses of a quasiperiodic integrable system after a quench. *Phys. Rev. A* **87**, 063637. <https://link.aps.org/doi/10.1103/PhysRevA.87.063637> (6 June 2013).
115. Lieb, E. H. & Liniger, W. Exact Analysis of an Interacting Bose Gas. I. The General Solution and the Ground State. *Phys. Rev.* **130**, 1605–1616. <https://link.aps.org/doi/10.1103/PhysRev.130.1605> (4 May 1963).
116. Kinoshita, T., Wenger, T. & Weiss, D. A quantum Newton’s cradle. *Nature* **440**, 900903. <https://doi.org/10.1038/nature04693> (2006).
117. Langen, T. *et al.* Experimental observation of a generalized Gibbs ensemble. *Science* **348**, 207–211. eprint: <https://www.science.org/doi/pdf/10.1126/science.1257026>. <https://www.science.org/doi/abs/10.1126/science.1257026> (2015).
118. Gring, M. *et al.* Relaxation and Prethermalization in an Isolated Quantum System. *Science* **337**, 1318–1322. ISSN: 0036-8075. eprint: <https://science.sciencemag.org/content/337/6100/1318.full.pdf>. <https://science.sciencemag.org/content/337/6100/1318> (2012).
119. Žnidarič, M., Prosen, T. ž. & Prelovšek, P. Many-body localization in the Heisenberg XXZ magnet in a random field. *Phys. Rev. B* **77**, 064426. <https://link.aps.org/doi/10.1103/PhysRevB.77.064426> (6 Feb. 2008).
120. Pal, A. & Huse, D. A. Many-body localization phase transition. *Phys. Rev. B* **82**, 174411. <https://link.aps.org/doi/10.1103/PhysRevB.82.174411> (17 Nov. 2010).
121. Fisher, M. P. A., Weichman, P. B., Grinstein, G. & Fisher, D. S. Boson localization and the superfluid-insulator transition. *Phys. Rev. B* **40**, 546–570. <https://link.aps.org/doi/10.1103/PhysRevB.40.546> (1 July 1989).
122. Fallani, L., Lye, J. E., Guarrera, V., Fort, C. & Inguscio, M. Ultracold Atoms in a Disordered Crystal of Light: Towards a Bose Glass. *Phys. Rev. Lett.* **98**, 130404. <https://link.aps.org/doi/10.1103/PhysRevLett.98.130404> (13 Mar. 2007).
123. Aleiner, I., Altshuler, B. & Shlyapnikov, G. A finite-temperature phase transition for disordered weakly interacting bosons in one dimension. *Nature Physics* **6**, 900904. <https://doi.org/10.1038/nphys1758> (11 2010).

124. Michal, V. P., Aleiner, I. L., Altshuler, B. L. & Shlyapnikov, G. V. Finite-temperature fluid-insulator transition of strongly interacting 1D disordered bosons. *Proceedings of the National Academy of Sciences* **113**, E4455–E4459. eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.1606908113>. <https://www.pnas.org/doi/abs/10.1073/pnas.1606908113> (2016).
125. Plisson, T. *et al.* Coherence properties of a two-dimensional trapped Bose gas around the superfluid transition. *Phys. Rev. A* **84**, 061606. <https://link.aps.org/doi/10.1103/PhysRevA.84.061606> (6 Dec. 2011).
126. Giorgini, S., Pitaevskii, L. & Stringari, S. Effects of disorder in a dilute Bose gas. *Phys. Rev. B* **49**, 12938–12944. <https://link.aps.org/doi/10.1103/PhysRevB.49.12938> (18 May 1994).
127. Huang, K. & Meng, H.-F. Hard-sphere Bose gas in random external potentials. *Phys. Rev. Lett.* **69**, 644–647. <https://link.aps.org/doi/10.1103/PhysRevLett.69.644> (4 July 1992).
128. Noh, J., Lee, J. & Mun, J. Mean-field interaction-induced dimensional crossover from two to three dimensions in a weakly interacting Bose gas. *Phys. Rev. A* **101**, 023619. <https://link.aps.org/doi/10.1103/PhysRevA.101.023619> (2 Feb. 2020).
129. Lewandowski, H., Harber, D., Whitaker, D. & Cornell, E. Simplified System for Creating a Bose–Einstein Condensate. *Journal of Low Temperature Physics* **132**, 309367 (2003).
130. Holzmann, M. & Krauth, W. Kosterlitz-Thouless Transition of the Quasi-Two-Dimensional Trapped Bose Gas. *Phys. Rev. Lett.* **100**, 190402. <https://link.aps.org/doi/10.1103/PhysRevLett.100.190402> (19 May 2008).
131. Campbell, R. *hermodynamic properties of a Bose gas with tuneable interactions* PhD thesis (University of Cambridge, UK, 2011).
132. Boettcher, I. & Holzmann, M. Quasi-long-range order in trapped two-dimensional Bose gases. *Phys. Rev. A* **94**, 011602. <https://link.aps.org/doi/10.1103/PhysRevA.94.011602> (1 July 2016).
133. Kapitza, P. Viscosity of Liquid Helium below the λ -Point. *Nature* **141**, 74 (1938).
134. Allen, J. & Misener, A. Flow of Liquid Helium II. *Nature* **141**, 75 (1938).
135. Bogoliubov, N. N. On the theory of superfluidity. *J. Phys. USSR* **11**, 23 (1947).
136. Desbuquois, R. *et al.* Superfluid behaviour of a two-dimensional Bose gas. *Nature Physics* **8**, 645 (2012).
137. Raman, C. *et al.* Evidence for a Critical Velocity in a Bose-Einstein Condensed Gas. *Phys. Rev. Lett.* **83**, 2502–2505. <https://link.aps.org/doi/10.1103/PhysRevLett.83.2502> (13 Sept. 1999).

138. Neely, T. W., Samson, E. C., Bradley, A. S., Davis, M. J. & Anderson, B. P. Observation of Vortex Dipoles in an Oblate Bose-Einstein Condensate. *Phys. Rev. Lett.* **104**, 160401. <https://link.aps.org/doi/10.1103/PhysRevLett.104.160401> (16 Apr. 2010).
139. Frisch, T., Pomeau, Y. & Rica, S. Transition to dissipation in a model of superflow. *Phys. Rev. Lett.* **69**, 1644–1647. <https://link.aps.org/doi/10.1103/PhysRevLett.69.1644> (11 Sept. 1992).
140. Huynh, J., Albert, M. & Larré, P.-É. Critical velocity for superfluidity in the one-dimensional mean-field regime: From matter to light quantum fluids. *Phys. Rev. A* **105**, 023305. <https://link.aps.org/doi/10.1103/PhysRevA.105.023305> (2 Feb. 2022).
141. Feynman, R. in (ed Gorter, C.) 17–53 (Elsevier, 1955). <https://www.sciencedirect.com/science/article/pii/S0079641708600773>.
142. Doran, R., Groszek, A. J. & Billam, T. P. *Critical Velocity and Arrest of a Superfluid in a Point-Like Disordered Potential* 2022. <https://arxiv.org/abs/2201.11887>.
143. Dries, D., Pollack, S. E., Hitchcock, J. M. & Hulet, R. G. Dissipative transport of a Bose-Einstein condensate. *Phys. Rev. A* **82**, 033603. <https://link.aps.org/doi/10.1103/PhysRevA.82.033603> (3 Sept. 2010).
144. Paris-Mandoki, A. *et al.* Superfluid flow above the critical velocity. *Sci Rep* **7**, 9070 (2017).
145. Engels, P. & Atherton, C. Stationary and Nonstationary Fluid Flow of a Bose-Einstein Condensate Through a Penetrable Barrier. *Phys. Rev. Lett.* **99**, 160405. <https://link.aps.org/doi/10.1103/PhysRevLett.99.160405> (16 Oct. 2007).
146. Leboeuf, P. & Pavloff, N. Bose-Einstein beams: Coherent propagation through a guide. *Phys. Rev. A* **64**, 033602. <https://link.aps.org/doi/10.1103/PhysRevA.64.033602> (3 Aug. 2001).
147. Albert, M., Paul, T., Pavloff, N. & Leboeuf, P. Breakdown of the superfluidity of a matter wave in a random environment. *Phys. Rev. A* **82**, 011602. <https://link.aps.org/doi/10.1103/PhysRevA.82.011602> (1 July 2010).
148. Paul, T., Schlagheck, P., Leboeuf, P. & Pavloff, N. Superfluidity versus Anderson Localization in a Dilute Bose Gas. *Phys. Rev. Lett.* **98**, 210602. <https://link.aps.org/doi/10.1103/PhysRevLett.98.210602> (21 May 2007).
149. Lye, J. E. *et al.* Bose-Einstein Condensate in a Random Potential. *Phys. Rev. Lett.* **95**, 070401. <https://link.aps.org/doi/10.1103/PhysRevLett.95.070401> (7 Aug. 2005).
150. Chen, X. *et al.* *AtomECS: Simulate laser cooling and magneto-optical traps* 2021. arXiv: 2105.06447 [cond-mat.quant-gas].
151. Acton, M. *Democratizing data-oriented design a data-oriented approach to using component systems* Unity at GDC 2018. 2018.

152. *Unreal Engine 5 MassEntity* <https://docs.unrealengine.com/5.0/en-US/overview-of-mass-entity-in-unreal-engine/>.
153. *ECS* <https://bevyengine.org/learn/book/getting-started/ecs/>.
154. *The open source, data-driven game engine* <https://amethyst.rs/>.
155. Karniadakis, G. E., Ali, B., Ali, B. & Aluru, N. R. *Microflows and Nanoflows: Fundamentals and Simulation* (Springer New York, 2005).
156. Alder, B. & Wainwright, T. Phase Transition for a Hard Sphere System. *The Journal of Chemical Physics* **27**, 1208–1209 (1957).
157. Bird, G. A. *Molecular gas dynamics and the direct simulation of gas flows* (Clarendon Press, 1998).
158. Lord, R. G. Direct simulation Monte Carlo calculations of rarefied flows with incomplete surface accommodation. *Journal of Fluid Mechanics* **239**, 449 (1992).
159. Oran, E., Oh, C. & Cybyk, B. DIRECT SIMULATION MONTE CARLO: Recent Advances and Applications. *Annual Review of Fluid Mechanics* **30**, 403–441 (1998).
160. Wade, A. C. J., Baillie, D. & Blakie, P. B. Direct simulation Monte Carlo method for cold-atom dynamics: Classical Boltzmann equation in the quantum collision regime. *Phys. Rev. A* **84**, 023612. <https://link.aps.org/doi/10.1103/PhysRevA.84.023612> (2 Aug. 2011).
161. Ma, Z.-Y., Foot, C. J. & Cornish, S. L. Optimized evaporative cooling using a dimple potential: an efficient route to Bose–Einstein condensation. *Journal of Physics B: Atomic, Molecular and Optical Physics* **37**, 3187–3195. <https://doi.org/10.1088/0953-4075/37/15/013> (July 2004).
162. Ma, Z.-Y., Thomas, A. M., Foot, C. J. & Cornish, S. L. The evaporative cooling of a gas of caesium atoms in the hydrodynamic regime. *Journal of Physics B: Atomic, Molecular and Optical Physics* **36**, 3533–3540. <https://doi.org/10.1088/0953-4075/36/16/313> (Aug. 2003).
163. Babovsky, H. & Illner, R. A Convergence Proof for Nanbu’s Simulation Method for the Full Boltzmann Equation. *SIAM Journal on Numerical Analysis* **26**, 45–65. ISSN: 00361429. <http://www.jstor.org/stable/2157706> (1989).
164. Egorov, M. *et al.* Measurement of *s*-wave scattering lengths in a two-component Bose-Einstein condensate. *Phys. Rev. A* **87**, 053614. <https://link.aps.org/doi/10.1103/PhysRevA.87.053614> (5 May 2013).
165. Dalibard, J. *Collisional dynamics of ultracold atomic gases* in (2002).
166. Landau, L. & Lifshitz, E. *Quantum Mechanics (Third Edition)* Third Edition. ISBN: 978-0-08-020940-1 (Pergamon, 1977).
167. Blundell, S. & Blundell, K. *Concepts in thermal physics* 2nd ed. eng. ISBN: 9780191718236 (Oxford University Press, Oxford, 2010).
168. Metcalf, H. J. & Straten, P. v. d. *Laser cooling and trapping* (Springer, 2002).

169. Ligare, M. Classical thermodynamics of particles in harmonic traps. *American Journal of Physics* **78**, 815–819 (Aug. 2010).
170. Gildemeister, M. *Trapping ultracold atoms in time-averaged adiabatic potentials* PhD thesis (University of Oxford, 2016).
171. Sherlock, B. *Ultracold quantum gases in time-averaged adiabatic potentials* PhD thesis (Oxford University, UK, 2016).
172. Sheard, B. *Magnetic transport and Bose-Einstein condensation of rubidium atoms* PhD thesis (Oxford University, UK, 2010).
173. Farkas, D. M., Salim, E. A. & Ramirez-Serrano, J. *Production of Rubidium Bose-Einstein Condensates at a 1 Hz Rate* 2014. <https://arxiv.org/abs/1403.4641>.
174. Torralbo-Campo, L., Bruce, G., Smirne, G. & Cassettari, D. Light-induced atomic desorption in a compact system for ultracold atoms. *Sci Rep* **5**. <https://www.nature.com/articles/srep14729> (2015).
175. Lu, Z. T. *et al.* Low-Velocity Intense Source of Atoms from a Magneto-optical Trap. *Phys. Rev. Lett.* **77**, 3331–3334. <https://link.aps.org/doi/10.1103/PhysRevLett.77.3331> (16 Oct. 1996).
176. Ravenhall, S. *A compact, high-flux source of cold atoms* PhD thesis (University of Oxford, 2019).
177. Ravenhall, S., Yuen, B. & Foot, C. High-flux, adjustable, compact cold-atom source. *Opt. Express* **29**, 21143–21159. <http://opg.optica.org/oe/abstract.cfm?URI=oe-29-14-21143> (July 2021).
178. Beregi, A. *In Preparation* PhD thesis (Oxford University, UK).
179. Getters, S. *Alkali metal dispensers brochure* Online.
180. Arpornthip, T., Sackett, C. A. & Hughes, K. J. Vacuum-pressure measurement using a magneto-optical trap. *Phys. Rev. A* **85**, 033420. <https://link.aps.org/doi/10.1103/PhysRevA.85.033420> (3 Mar. 2012).
181. Ott, H., Fortagh, J., Schlotterbeck, G., Grossmann, A. & Zimmermann, C. Bose-Einstein Condensation in a Surface Microtrap. *Phys. Rev. Lett.* **87**, 230401. <https://link.aps.org/doi/10.1103/PhysRevLett.87.230401> (23 Nov. 2001).
182. Bertoli, G., Altshuler, B. L. & Shlyapnikov, G. V. Many-body localization in continuum systems: Two-dimensional bosons. *Phys. Rev. A* **100**, 013628. <https://link.aps.org/doi/10.1103/PhysRevA.100.013628> (1 July 2019).
183. Singh, V. P. & Mathey, L. Noise correlations of two-dimensional Bose gases. *Phys. Rev. A* **89**, 053612. <https://link.aps.org/doi/10.1103/PhysRevA.89.053612> (5 May 2014).
184. Pigneur, M. *et al.* Relaxation to a Phase-Locked Equilibrium State in a One-Dimensional Bosonic Josephson Junction. *Phys. Rev. Lett.* **120**, 173601. <https://link.aps.org/doi/10.1103/PhysRevLett.120.173601> (17 Apr. 2018).

185. Gati, R. & Oberthaler, M. K. A bosonic Josephson junction. *Journal of Physics B: Atomic, Molecular and Optical Physics* **40**, R61–R89. <https://doi.org/10.1088/0953-4075/40/10/r01> (May 2007).
186. Luick, N. *et al.* An ideal Josephson junction in an ultracold two-dimensional Fermi gas. *Science* **369**, 89–91. eprint: <https://www.science.org/doi/pdf/10.1126/science.aaz2342>. <https://www.science.org/doi/abs/10.1126/science.aaz2342> (2020).
187. Song, F.-F. & Zhang, G.-M. Phase Coherence of Pairs of Cooper Pairs as Quasi-Long-Range Order of Half-Vortex Pairs in a Two-Dimensional Bilayer System. *Phys. Rev. Lett.* **128**, 195301. <https://link.aps.org/doi/10.1103/PhysRevLett.128.195301> (19 May 2022).
188. Bighin, G., Defenu, N., Nándori, I., Salasnich, L. & Trombettoni, A. Berezinskii-Kosterlitz-Thouless Paired Phase in Coupled XY Models. *Phys. Rev. Lett.* **123**, 100601. <https://link.aps.org/doi/10.1103/PhysRevLett.123.100601> (10 Sept. 2019).
189. Milburn, G. J., Corney, J., Wright, E. M. & Walls, D. F. Quantum dynamics of an atomic Bose-Einstein condensate in a double-well potential. *Phys. Rev. A* **55**, 4318–4324. <https://link.aps.org/doi/10.1103/PhysRevA.55.4318> (6 June 1997).