




Probing universality of 2D quantum systems
with bilayer Bose gases

Abel Beregi
Wolfson College



University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

Hilary 2024

Probing universality of 2D quantum systems with bilayer Bose gases

Abel Beregi

Wolfson College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Hilary 2024

The universal nature of phase transitions allows one to develop a unified description of microscopically different systems. A particularly interesting universality class, which includes the 2D Bose gas, is that of the 2D XY model. In these systems there exists a phase transition between a superfluid and a normal phase. While 2D Bose gases have been studied extensively theoretically and experimentally, there remain many open questions. This thesis explores the various aspects of universal behaviour of a 2D Bose gas in and out of equilibrium.

In this thesis, we prepare ultracold quantum gases of ^{87}Rb in radio frequency (RF) dressed potentials. We describe a new experimental apparatus which utilises this technique and allows for a simple trap loading procedure from a cloud of laser-cooled atoms. The creation of a degenerate quantum gas is demonstrated, with evaporative cooling performed entirely in RF-dressed potential. We create a double well potential using multiple radio frequencies and perform coherent splitting to form bilayer 2D Bose gases. We use matter-wave interferometry to access two important quantities, the local relative phase of the clouds and the interference contrast. Using these methods, we characterise the full counting statistics of interference contrast across the superfluid transition and demonstrate its convergence to a universal functional form. We also study the effect of disordered potentials on 2D Bose gases and show that the system loses coherence with increasing disorder strength and crosses into the normal phase. Our results suggest that this disorder-induced transition is likely to be still within the 2D XY universality class. We show preliminary results for experiments where the BKT transition is crossed dynamically and compare the time-evolution of the system with universal scaling laws. Finally, we propose an RF dressed trap in which homogeneous 2D systems can be created.

Table of Contents

1	Introduction	2
2	Theory of two-dimensional superfluids	5
2.1	The 2D XY model	5
2.2	The absence of true long-range order in the 2D XY model	6
2.3	Mapping 2D Bose gases onto the 2D XY model	7
2.4	Vortices in 2D Bose gases	9
2.5	The BKT transition	11
2.6	Renormalisation group treatment of the BKT transition	12
2.7	2D Bose gases in disordered environments	15
2.8	Bilayer 2D Bose gases	17
2.9	The effect of interlayer coupling	19
2.9.1	RG theory of bilayer gases	21
3	RF-dressed adiabatic potentials	23
3.1	Magnetic trapping in DC fields	23
3.2	Dressed atoms	24
3.2.1	The dressed Hamiltonian	25
3.2.2	An approximate solution to the dressed Hamiltonian	26
3.2.3	The rotating-wave Hamiltonian	28
3.3	Confining cold atoms in RF-dressed potentials	30
3.3.1	The dressed quadrupole trap	31
3.4	Multiple RF-dressing	34
4	Experimental Apparatus	37
4.1	Overview of the apparatus	37
4.1.1	The vacuum system	39
4.1.2	Modelling the vacuum system	39
4.1.3	The history of the vacuum system	42
4.2	The laser system	43
4.2.1	Overview	44
4.2.2	Injection locked diode lasers	45
4.2.3	Tapered amplifiers	46
4.2.3.1	Eliminating spurious TA temperature drifts	48
4.2.4	The six-beam MOT distribution board	50
4.3	The MOT beam launching periscopes	52
4.3.1	Liquid crystal variable waveplates	54
4.3.2	Photodetectors	55

4.4	PID controllers and LCD drivers	56
4.4.1	PID controller hardware	57
4.4.1.1	The microcontroller	58
4.4.1.2	The input signal chain	59
4.4.1.3	The output signal chain	59
4.4.2	Performance evaluation	60
4.5	The glass cell MOT setup	61
4.6	The RF signal chain	62
4.6.1	The RF signal generation and amplification	63
4.6.2	PCB-based RF coils	63
4.6.3	RF amplitude and phase measurement	66
4.6.3.1	The resistive pickups	67
4.6.3.2	Processing of pickup signals	68
4.7	The control system	69
4.7.1	Example of ARTIQ sequences	70
4.7.2	The apparatus module	71
4.7.3	The sequence module	72
4.7.4	The control interface	75
4.7.5	Machine learning optimisation	79
4.7.6	Interfacing the PID controllers to ARTIQ	83
4.8	Future upgrades of the new experiment	85
4.8.1	Laser system	85
4.8.2	PCB RF coils	86
4.8.3	RF-dressed dimple trap	87
5	Direct Loading and Bose-Einstein condensation in the shell trap	88
5.1	RF dressed trap loading schemes	88
5.2	Loading a shell trap from optical molasses	90
5.2.1	Cell MOT loading	90
5.2.2	Compressed MOT and molasses	91
5.2.3	Shell trap loading	93
5.2.4	RF spectroscopy of the shell trap	95
5.2.5	Calibration of the RF amplitude and polarisation	96
5.3	Evaporative cooling and Bose-Einstein condensation	97
5.4	Conclusion and outlook	104
6	Full counting statistics of interference contrast of bilayer 2D Bose gases	105
6.1	Related work	105
6.2	Experimental protocol	106
6.3	Correlation function analysis	111

6.4	Contrast analysis	115
6.5	Simulating contrast full counting statistics	116
6.6	Comparison of simulation and experiment	118
6.7	Universal properties of contrast statistics in 2D	121
6.8	Moments of contrast distribution	123
6.9	Conclusion and outlook	124
7	Superfluid transition in disordered systems	126
7.1	Generation of disordered optical potential	126
7.2	Experimental protocol	127
7.3	Estimating the strength of disorder	128
7.4	Effects of optical potential on bulk parameters	129
7.5	Correlation function analysis	131
7.6	Vortex number analysis	133
7.7	The effective coupling model	135
7.8	Simulating the effect of disorder on Bose gases	137
7.9	Conclusion and future plans	139
8	Dynamical control of bilayers across the BKT transition	142
8.1	Tunnel coupling strength quench	142
8.1.1	Experimental methods	143
8.1.2	Relaxation dynamics after quench	144
8.1.3	Conclusion and outlook	145
8.2	Disorder strength quench	146
8.2.1	Experimental methods	147
8.2.2	Relaxation dynamics after a reverse quench	147
8.2.3	Relaxation dynamics after a forward quench	148
8.2.4	Conclusion and outlook	150
9	Creating uniform two-dimensional RF-dressed potentials	151
9.1	The advantages of uniform potentials	151
9.2	Realising uniform 2d potentials in a magnetic dipole trap	152
9.2.1	Magnetic field of a single wire loop	152
9.2.2	Quasi-uniform trap in absence of gravity	156
9.2.3	PCB-based uniform trap in absence of gravity	157
9.2.4	Quasi-uniform trap on Earth	159
10	Conclusion and outlook	162
10.1	Conclusion	162
10.2	Future experimental directions	163

10.2.1	Matter-wave interferometry with the next-generation experiment	163
10.2.2	Phase diagram of disordered 2D Bose gases	163
Appendix A	XY model correlation function integral	165
Appendix B	PID controller firmware and control interface	166
Appendix C	FFT-based analysis of interference phase and contrast	169
C.1	FFT analysis for clean data	169
C.2	FFT analysis of noisy data	170
Appendix D	Simulating quantum gases	172
D.1	Coupled Gross-Pitaevskii equations	172
D.2	Finding the ground state at zero temperature	173
D.2.1	The case of rotational symmetry	175
D.3	Finding the initial state at a finite temperature	177
D.4	Finding ground state with a fluctuating phase	178
D.5	Simulating dynamics	179
D.6	Case study 1: Simulating matter-wave interferometry	180
D.7	Case study 2: Matter-wave lensing	183
D.8	Classical field Monte Carlo simulations	188
D.9	Case study 3: Effect of box potential on correlation functions . . .	189

List of Figures

2.1	Illustration of the role of vortices in the BKT transition	12
2.2	BKT transition RG flow	14
2.3	BKT critical temperature in a disordered system	18
2.4	BKT critical temperature in a coupled bilayer system	22
3.1	RF-dressed eigenstates.	28
3.2	Coupling strength on the resonant spheroid	33
3.3	MRF dressed adiabatic potentials	36
4.1	Comparison of the previous generation and new experimental setups	38
4.2	The vacuum system	40
4.3	Lumped-element model of vacuum system	41
4.4	Laser system schematic	45
4.5	ILD setup schematic	47
4.6	TA setup schematic	47
4.7	Modifying TED8040 controllers	49
4.8	MOT light distribution	51
4.9	MOT light distribution schematic	53
4.10	Beam launching periscope	54
4.11	Amplified photodetector	56
4.12	PID controller	58
4.13	Glass cell MOT setup	62
4.14	PCB RF coils	65
4.15	PCB RF coil AC characteristics	67
4.16	ARTIQ dashboard	78
4.17	ML optimisation evolution plot	84
4.18	ML result analysis	84
4.19	PID interface PC and GUI	85
5.1	Atoms in the shell trap	94
5.2	RF spectroscopy in the shell trap	96
5.3	Evaporative cooling in the shell trap	103
6.1	Effect of ring-shaped blue-detuned beam on trapped atoms	108
6.2	Contrast statistics experiment thermometry	110
6.3	Illustration of correlation function analysis	113
6.4	Contrast statistics experiment correlation functions	114
6.5	Contrast statistics experiment correlation analysis results	115
6.6	Simulated and observed contrast statistics	119
6.7	Comparison of noisy simulated and observed contrast statistics	121

6.8	Convergence of contrast statistics to universal form	122
6.9	Analysis of moments of contrast statistics	124
7.1	Disorder experiment thermometry	131
7.2	BKT critical point identification	133
7.3	Disorder experiment vortex number analysis	135
7.4	Scaling of algebraic exponent with disorder strength	137
7.5	Comparison of MC simulations of a 2D Bose gas in a disordered environment and the experimental results.	140
8.1	Tunnel coupling strength quench	145
8.2	Universal relaxation dynamics following a quench	146
8.3	Disorder reverse quench	148
8.4	Disorder forward quench	149
9.1	Isomagnetic surfaces of a single wire loop.	153
9.2	Numerical results for RF-dressed potential by a single wire loop.	155
9.3	PCB implementation of DC coil for uniform potentials	159
9.4	Coil geometry for Earth-based uniform RF-dressed trap	161
C.1	Illustration of FFT-based phase and contrast analysis	170
C.2	Reduction of phase and contrast noise with FFT analysis	171
D.1	pyGPE ground state calculations	176
D.2	Ground state calculations with non-uniform phase	179
D.3	pyGPE simulations of condensate dynamics	180
D.4	Matter-wave interference simulation protocol	182
D.5	Investigating systematic effects on matter-wave interferometry	184
D.6	Matter-wave lensing simulation protocol	187
D.7	Investigation of matter-wave lensing with pyGPE	188
D.8	Classical field MC simulations	191

List of Tables

5.1	Parameters along the evaporative cooling ramp	100
-----	---	-----

Acknowledgements

Firstly, I would like to express my gratitude to Prof. Chris Foot for the unwavering support I have received over the past four years. He has been incredibly supportive and was always available for discussions on any topic.

This work would not have been possible without the support of the research group. I am sincerely thankful to Dr Elliot Bentine, who provided immense support with the experiments over the past four years and it has always been a pleasure to have technical discussions with him. I also appreciate the help from Dr Adam Barker during my first year at Oxford. He was always keen to answer my questions, even if I asked the same for the third time. I consider myself fortunate to have worked with Dr Shinichi Sunami, with whom I spent memorable time in the lab, finding the weirdest bugs with the old experiment. I am also thankful for the meticulous feedback he provided on this work. It was great to work with Dr David Garrick in the lab, I will never forget how exciting was the first realisation of a MOT with the new experiment. Special thanks to Oscar Chang who helped me to learn again how to operate the old experiment after a 2 year long break. Working with Erik Rydow over the past year on the new experiment has been great and I am sure he will take it much further. I would also like to thank for the productive discussions with our theory collaborators, Dr. Vijay Singh and Prof. Ludwig Mathey. Their enthusiasm for new experimental results always inspires me.

I would like to thank for the support of the other members of the research group, Dr Sean Ravenhall, Dr Dongyang Xu, Dr Tom Hird, Dr Kenneth Hughes, Abigail James, Jesse Schelfhout, and Ghadah Alshalan. Also, it was great supervise Josh Greensmith, Zhengze Li, Quentin Leterrier, Simai Jia, Amarin Bienvenue and most recently, Magnus Sesodia, who all made contributions to this work.

Finally, I am extremely thankful for all the support I have received from my partner, Magdalena. She always listened patiently when I complained about the experiment (which must have been a significant time over the past 4 years) and always believed that I will succeed with this degree.

1 | Introduction

Exploration of the universal nature of phase transitions is a cornerstone of condensed matter physics since universality facilitates a unified description of systems which are microscopically different or are of vastly different scale. One example for this is the two-dimensional (2D) XY universality class which describes ^4He films [1], 2D superconductors [2], Josephson junction arrays [3], or 2D Bose gases [4]. Systems within this universality class exhibit a phase transition between a quasi-ordered superfluid and a disordered normal phase. The microscopic origin of this transition, the unbinding of vortex-antivortex pairs, was identified by Berezinskii, Kosterlitz and Thouless (BKT) [5, 6]. For this and related discoveries, Kosterlitz and Thouless were awarded a share of the Nobel Prize in Physics in 2016.

Investigations of these phenomena are often limited by computation speeds or the uncertainties of sample preparation in solid-state systems. Since the first realization of a degenerate quantum gas in 1995 [7], ultracold atomic systems have been used extensively for investigating many-body physics, as these offer reproducible sample preparation, high-degree of isolation from the environment as well as exceptional level of control for manipulation [8, 9]. One remarkable feature of ultracold Bose gas experiments is the ability to perform matter-wave interferometry, giving access to the relative phase between two atomic clouds [10]. This experimental method was key to the first observation of the BKT phase transition in 2D Bose gases by Hadzibabic *et al.* [11] as well as to the comprehensive measurements of Sunami *et al.* [12]. The high degree of isolation from the environment in ultracold atomic systems also allows probing the relaxation dynamics after a dynamical crossing of the superfluid transition in 2D systems [13].

In this work, we investigate universality of 2D bilayer Bose gases from three different aspects. Firstly, we investigate the full counting statistics of interference contrast; the probability distribution of this observable is expected to take the form of a universal function at the critical point and deep in the superfluid regime [14]. Secondly, we explore the effect of a disordered external potential on the superfluid transition and compare our results with theoretical predictions which show that for weak disorder, the phase transition is topologically protected and is still within the 2D XY universality class [15]. Finally, we observe the time evolution of the bilayer system following a quench of the interlayer coupling strength and compare the results with the universal scaling law predicted in [16]. We also show preliminary results of experiments where we realise a quench by changing the strength of the disorder potential.

This thesis is structured as follows. In Chapter 2, we give a brief review of the theory of 2D Bose gases and the BKT phase transition. We also discuss the effects of disorder potentials and interlayer coupling on the superfluid transition. In Chapter 3 we present the theory of radio-frequency (RF) dressed adiabatic potentials which serves as the main experimental method for trapping atoms in this work. Chapter 4 gives a detailed description of the next-generation Oxford apparatus. Chapter 5 is devoted to a novel loading scheme for RF-dressed quadrupole traps. We also demonstrate the creation of a degenerate quantum gas using the next-generation apparatus. Chapter 6 presents experimental results of full counting statistics of interference contrast of 2D Bose gases. Here, also describe in detail the main analysis methods which are used to extract correlation functions from matter-wave interference patterns. In Chapter 7 we explore the effect of a disordered potential on the superfluid transition. In Chapter 8, we show results of experiments in which the superfluid transition is crossed dynamically by various methods. Chapter 9 explores the possibility of realising a uniform potential

along the horizontal plane using RF-dressed techniques. Finally, we summarise the work and provide an outlook for future experiments in Chapter 10.

2 | Theory of two-dimensional superfluids

In this chapter the theoretical description of 2D superfluids is reviewed. We start with introducing the 2D XY model which serves as a fundamental starting point for describing the behaviour of such systems. We show that the 2D XY model does not exhibit true long-range order at any non-zero temperature. Next, we map the problem of a cold 2D Bose gas onto the XY model. The role of vortices in 2D Bose gases is discussed and we highlight how these topological defects play a crucial role in the BKT phase transition by simple statistical arguments as well as renormalisation group techniques. Finally, we review the effects of disordered potentials and coherent coupling in bilayers on the BKT transition.

2.1 The 2D XY model

The classical 2D XY model considers magnetic moments on a square lattice, oriented along the system plane which interact through nearest neighbour interactions. In the absence of an external magnetic field, the Hamiltonian therefore reads as

$$\mathcal{H} = -J \sum_{\langle i,j \rangle} \mathbf{S}_i \cdot \mathbf{S}_j = -J \sum_{\langle i,j \rangle} \cos(\theta_i - \theta_j), \quad (2.1)$$

where we considered unit magnetic moments for simplicity. For slow spatial variations in the orientation of the magnetic moments between neighbouring sites ($\theta_i - \theta_j \ll 1$), a Taylor expansion of $\cos(\theta_i - \theta_j)$ gives an approximate continuum Hamiltonian

$$\mathcal{H} \approx E_0 + \frac{J}{2} \int (\nabla\theta)^2 d\mathbf{r}, \quad (2.2)$$

where E_0 is the ground-state energy. Within these approximations, the system is completely described by the classical field $\theta(\mathbf{r})$ which we refer to as the phase

field as it is analogous to the complex phase of the condensate wavefunction in the case of a 2D Bose gas.

2.2 The absence of true long-range order in the 2D XY model

By virtue of the Mermin-Wagner theorem [17], we expect true long-range order to be absent in 2D systems at any non-zero temperature. To see this, we can expand the phase field into Fourier modes as

$$\theta(\mathbf{r}) = \sum_{\mathbf{k}} c_{\mathbf{k}} e^{i\mathbf{k}\cdot\mathbf{r}}, \quad (2.3)$$

where the coefficients $c_{\mathbf{k}}$ obey Hermitian symmetry since $\theta(\mathbf{r})$ is real. This allows us to rewrite the Hamiltonian as a sum over energies of spin-wave modes

$$\mathcal{H} = \frac{L^2 J}{2} \sum_{\mathbf{k}} k^2 |c_{\mathbf{k}}|^2, \quad (2.4)$$

where L is the system size. By applying the equipartition theorem, we find the mode amplitudes in thermal equilibrium [18]

$$\langle |c_{\mathbf{k}}|^2 \rangle = \frac{k_B T}{L^2 J k^2}. \quad (2.5)$$

To quantify the degree of order in the system, we introduce the phase correlation function as

$$g_1(\mathbf{r}) = \langle \mathbf{S}(\mathbf{r}) \cdot \mathbf{S}(0) \rangle = \langle e^{i(\theta(\mathbf{r}) - \theta(0))} \rangle. \quad (2.6)$$

Using mode amplitudes, we arrive at

$$g_1(\mathbf{r}) = \exp\left(-\frac{k_B T}{4\pi^2 J} \int \frac{1 - \exp(i\mathbf{k} \cdot \mathbf{r})}{k^2} d\mathbf{k}\right), \quad (2.7)$$

where we transformed the discrete sum over \mathbf{k} into an integral. The details of evaluating this integral is given in Appendix A. Imposing a higher limit of $k_{max} = 1/\xi$ which is equivalent to a short-distance cutoff as well as a lower limit $k_{min} = 1/r$ ¹

$$g_1(r) = \left(\frac{\xi}{r}\right)^\eta. \quad (2.8)$$

Therefore the correlation function exhibits a slow, algebraic decay characterised by the exponent $\eta = \frac{k_B T}{2\pi J}$. The absence of true long-range order is clear as $g_1(r \rightarrow \infty) = 0$. However, at sufficiently low temperatures, the decay of correlations is extremely slow and can be negligible over the size of the system. This is referred to as quasi long range order [19].

2.3 Mapping 2D Bose gases onto the 2D XY model

Despite obvious differences on the microscopic level between a 2D XY model magnet and interacting 2D Bose gases, they both belong to the XY universality class [4]. Here, we show that under certain conditions, the Hamiltonian of a 2D Bose gas can be reduced to the form of 2.2, which allows us to realise a mapping between the two problems. We begin with the assumption the gas is accurately described by a classical Bose field, which we write in the density-phase representation as

$$\psi(\mathbf{r}, t) = \sqrt{n(\mathbf{r}, t)} e^{i\theta(\mathbf{r}, t)}. \quad (2.9)$$

¹Ultimately, we are interested at the behaviour of the correlation function at large distances, therefore $1/r$ is the lower limit of the integral.

It is convenient to separate the density into a mean value n and fluctuations δn as $|\psi(\mathbf{r}, t)|^2 = n(1 + 2\delta n(\mathbf{r}, t))$. This allows us to write the classical field Hamiltonian of the system as [20, 21]

$$\mathcal{H} = \frac{\hbar^2 n}{2m} \int (\nabla\theta(\mathbf{r}))^2 + (\nabla\delta n(\mathbf{r}))^2 + 2gn(\delta n(\mathbf{r}))^2 d\mathbf{r}, \quad (2.10)$$

where g is the 2D interaction strength. The first term is of the same form as 2.2. Therefore, by demonstrating that the energy cost of the density fluctuations is large compared to that of phase fluctuations, we also show that the 2D XY model provides an approximate description of 2D Bose gases. By using the Fourier decompositions of the phase and density fluctuations, from 2.10, we obtain

$$\frac{|c_{\mathbf{k}}|^2}{|d_{\mathbf{k}}|^2} = \frac{\hbar^2 k^2 / 2m + 2gn}{\hbar^2 k^2 / 2m} \quad (2.11)$$

where $|c_{\mathbf{k}}|$ and $|d_{\mathbf{k}}|$ are the amplitudes of phase and density fluctuations of the phonon mode with wavevector \mathbf{k} . Therefore, for long-wavelength fluctuations (which destroy true long-range order as shown earlier for the 2D XY model), the density fluctuations are suppressed. This has been confirmed both by simulations [22] and experiments [23]. In this case, we can approximate 2.10 with an effective low-energy Hamiltonian

$$\mathcal{H} \approx \frac{\hbar^2 n_s}{2m} \int (\nabla\theta(\mathbf{r}))^2 d\mathbf{r}. \quad (2.12)$$

which is of the same form as the 2D XY model continuum Hamiltonian (2.2). Note that we have replaced the density n with the superfluid density n_s to heuristically account for short-wavelength phonons. This is justified as only for the superfluid component it is meaningful to define the local velocity as $\mathbf{v}_s = (\hbar/m)\nabla\theta$ while the normal fluid density is non-zero at any finite temperature thus contribut-

ing to the total density. Comparing this to 2.2, the correlation function 2.8 becomes

$$g_1(r) = \langle \psi^*(\mathbf{r})\psi(0) \rangle = n_s \left(\frac{\xi}{r} \right)^{1/n_s \lambda^2} \propto r^{-\eta}, \quad (2.13)$$

where $\xi = 1/\sqrt{\tilde{g}n}$ is the 2D healing length [24] and we have introduced the dimensionless interactions strength as $\tilde{g} = mg/\hbar^2$. For quasi 2D Bose gases confined in a harmonic potential along z , characterised by the oscillator length $a_z = \sqrt{\hbar/m\omega_z}$, the dimensionless interaction strength is $\tilde{g} = \sqrt{8\pi}a/a_z$ where a is the s-wave scattering length. Comparing 2.13 with 2.8 allows us to realise a mapping between the spin-spin coupling J and properties of the 2D Bose gas.

2.4 Vortices in 2D Bose gases

The previous treatment of the 2D XY model and the 2D Bose gas assumed a low temperature. At higher temperatures, it becomes important to consider another class of excitations which are called vortices by analogy with a circulating flow in fluid dynamics. Since the wavefunction of the condensate must be single-valued, the change in phase over a closed loop must obey

$$\Delta\theta = \oint \nabla\theta \cdot d\mathbf{l} = 2\pi l, \quad (2.14)$$

where $l \in \mathbb{Z}$. Using the hydrodynamic equations [25], the phase gradient is directly proportional to the velocity field. Therefore, the velocity of the condensate is inversely proportional to the distance from the vortex core similar to the case of an irrotational vortex in fluid mechanics. From the condition for the wavefunction to be single-valued, one can also see that the density must vanish at the vortex core. In a condensate, the characteristic length over which the density retains its

normal value is the healing length (or equivalently the minimum distance for a substantial change in the wavefunction), giving the characteristic radius of a vortex core. For simplicity, we only consider vortices with $\pm 2\pi$ phase winding², as vortices with higher phase winding are energetically unstable in most scenarios [27]. It is important to note that vortices are topological excitations as there is no continuous deformation between the phase field of a vortex and a uniform phase field. This is in distinct contrast to the case of phonons; as their amplitude decreases to zero, the corresponding phase field becomes uniform. To estimate the energy of a vortex, we consider a uniform system of radius R , and calculate the kinetic energy of the superfluid mass flow of a single vortex in an annular region between the healing length and the system size, where the superfluid density can be approximated as constant. This gives

$$E_v \approx \frac{\hbar^2 n_s}{2m} \int_{\xi}^R \frac{1}{r^2} 2\pi r dr = \frac{\hbar^2 n_s \pi}{m} \ln \left(\frac{R}{\xi} \right). \quad (2.15)$$

To estimate the configurational entropy of a single vortex in this system, we assign an area of $\pi\xi^2$ to a vortex such that there are $\pi R^2/\pi\xi^2$ microstates, giving a Boltzmann entropy of

$$S_v \approx 2k_B \ln \left(\frac{R}{\xi} \right). \quad (2.16)$$

Hence the approximate free energy of a vortex is given by

$$F_v = E_v - TS_v \approx \left(\frac{\hbar^2 n_s \pi}{m} - 2k_B T \right) \ln \left(\frac{R}{\xi} \right). \quad (2.17)$$

As the temperature increases, the free energy changes sign from positive to negative at the critical point, therefore the proliferation of vortex excitations is energetically favourable. Since vortices disrupt the phase field over short distances,

²The phase winding is often referred to as vortex charge since vortices in a 2D Bose gas can be modelled as a Coulomb gas [26].

their widespread occurrence in a 2D system is inconsistent with the slow algebraic decay of $g_1(r)$ that is characteristic of quasi long range order, indicating a crossover to a disordered phase with a short-ranged correlation function. This energetic argument was first mentioned by Kosterlitz in 1974 [28].

2.5 The BKT transition

To resolve the contradiction presented earlier, Berezinskii, Kosterlitz and Thouless proposed a microscopic theory which explains the connection between the two regimes [5, 6]. The key idea is that the transition between the quasi-ordered and the disordered phases is mediated by the unbinding of vortex pairs of opposite circulation. Such a vortex pair has no net circulation for contours which enclose both. Therefore, significant phase fluctuations are only present over length scales comparable with the separation between the vortices. We illustrate this in Figure 2.1 a, where the phase field of seven randomly placed vortex pairs is shown, with each pair having no net circulation. Since the vortices which constitute a pair are close compared to the system size, the phase stays near zero in the majority of the system. Figure 2.1 b shows the phase field of the same number of vortices, but here the locations of the individual vortices are chosen randomly. In contrast with the previous case, the phase exhibits strong variations throughout the system. This is analogous to how the electric potential of a dipole is zero to first order at distances larger than the separation between the charges. In fact, the expression 2.17 is of the same form as the free energy of an isolated charge in a Coulomb gas [6, 26]. From 2.17 we can calculate the critical superfluid density at which the BKT transition occurs by the simple relation

$$n_{s,c}\lambda^2 = 4, \tag{2.18}$$

which gives a universal value for the critical exponent $\eta_c = 0.25$. We can also conclude from 2.17 that at any temperature, the superfluid density must be at least that of the critical value when the system is in the quasi-ordered phase. This universal jump in superfluid density is one of the main indicators of the BKT phase transition. However, it is important to note that in finite-size systems, the jump becomes a smooth crossover.

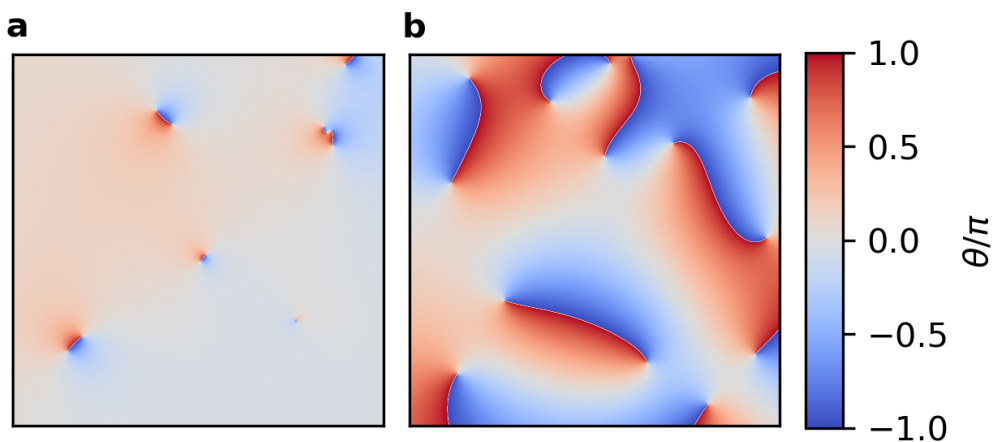


Figure 2.1: A cartoon picture of the role of vortices in the BKT transition. **a:** The phase field of 7 pairs of vortices, each pair consist of closely separated (bound) vortices of opposite circulation. The vortices have little effect on the phase field at distances much greater than the characteristic separation of vortices forming a pair. **b:** The phase field of 14 vortices at random locations with no net circulation.

2.6 Renormalisation group treatment of the BKT transition

In order to gain a full understanding of the BKT phase transition, a powerful theoretical tool is the renormalisation group (RG) which is widely used in quantum field theory and condensed matter physics. The key idea behind RG methods is observing how the coupling constants of a system's Hamiltonian, such as J in

2.2, change as we view the system on increasing length scales. This can be motivated by various arguments. Firstly, one usually wants to describe the behaviour of a system in the thermodynamic limit starting from microscopic models, such as nearest-neighbour spin-spin interactions. Secondly, within the scope of condensed matter physics, RG theory gave significant insights into second-order phase transitions which are characterised by a diverging correlation length as the critical point is approached from the disordered phase [29, 30]. This implies that microscopically different systems, such as the XY model magnet and a 2D Bose gas, may share universal properties on long length scales and it is desirable to develop a formal mathematical procedure to ignore short-distance physics. For the BKT transition, the typical RG treatment is as follows. The partition function of the system has two important parameters, the vortex fugacity $y = e^{-\beta\epsilon_c}$, where ϵ_c is the energy of a vortex core (given by 2.15), and the shifted ratio of temperature and spin-spin coupling strength $x = 1/(\beta J) - \pi/2$ ³. In a simplified picture, y serves as a measure of vortex density [13] while x is a measure of the algebraic exponent η . At every "iteration", we integrate out vortex pairs with separation $r_0 \rightarrow r_0 e^{dl}$ starting at the healing length ξ , where l is typically referred to as the RG flow parameter. This gives a partition function which resembles the one before but with modified x and y parameters. This leads to the RG flow equations

$$\begin{aligned}\frac{dx}{dl} &= 4\pi^3 y^2, \\ \frac{dy}{dl} &= \frac{4}{\pi} xy.\end{aligned}\tag{2.19}$$

These equations can be visually represented in a flow diagram which is shown in Figure 2.2. From the equations, it is clear that the $y = 0$ line is a fixed point, which is stable for negative values and unstable for positive values of x . Furthermore, a separatrix can be identified which is given by $y = x/\pi^2$. Anywhere left of this

³This choice yields simpler equations later.

line, the RG flow drives the system towards a state without vortices, as the stable fixed point corresponds to an infinite energy cost of creating a single vortex. While anywhere right of this line, the system is driven to a state with a high number of vortices as the energy cost of creating vortices decreases. Therefore, we can identify the critical temperature as

$$T_c = \frac{\pi J}{2 k_B}, \quad (2.20)$$

which is identical to the earlier result 2.17 based on free energy considerations after mapping the spin-spin coupling to properties of the 2D Bose gas.




Figure 2.2: RG flow diagram describing the equilibrium BKT transition. The variable x relates to η while y is a measure of vortex density (see text for detailed definitions).

Apart from identifying the BKT phase transition in the thermodynamic limit, the RG theory can be used to predict how parameters such as temperature affect observables, for instance, the algebraic exponent η in finite size systems. Firstly, the initial point for the RG flow must be identified, which depends on the system

parameters. For a 2D Bose gas, the superfluid coupling constant is temperature dependent due to the renormalisation of superfluid density. Therefore, the initial point is temperature dependent through both β and J . For accurate results, one can either interpolate numerical results for the superfluid density from high-precision Monte Carlo (MC) simulations [31] or use an approximate analytical ansatz [32]. After this, the RG equations can be integrated numerically until the flow parameter reaches $\ln(L/\xi)$ where L is the system size. From the final value of x , the renormalised superfluid coupling J' can be calculated which gives the algebraic exponent of the finite-size system.

2.7 2D Bose gases in disordered environments

The effect of imperfections in condensed matter is of significant interest as experimental samples are impure. Examples include the Kondo effect [33] or Anderson localisation [34]. The interplay between disorder and interactions in these systems leads to complex and not fully understood phenomena such as phase transitions to insulating states [35] or high- T_c superconductivity [36]. The high degree of control over the potential and the repeatability of sample preparation makes ultra-cold atom systems ideal to investigate such effects. The case of a 2D Bose gas in a random environment is particularly interesting because of the existence of the BKT phase transition between superfluid and normal phases in systems without disorder.

A common starting point of theoretical descriptions of disordered 2D Bose gases is the local density approximation (LDA), which is valid if the disorder correlation length is larger than the healing length and the thermal de Broglie wavelength of the gas. This condition is satisfied for optical speckle potentials in typical experimental conditions. Using the LDA and percolation theory, Carleo *et al.*

argue that the system is expected to remain in the superfluid phase when there exist percolating clusters where the density is higher than the critical density in a disorder-free system [15]. Using the percolation threshold of a laser speckle potential in 2D systems ($\alpha = 0.52$) [37], this results in a linear reduction of the critical temperature, shown in Figure 2.3 (blue line), given by

$$T_c(\bar{V}) = T_c^0 \left(1 - \alpha \frac{\bar{V}}{\mu} \right), \quad (2.21)$$

where T_c^0 is the critical temperature in a non-disordered system, \bar{V} is the mean value of the disordered potential and μ is the chemical potential of the gas. Since the equation of state of a 2D Bose gas is not known analytically, Bourdel analysed the same problem [38] using numerical results for the equation of state from high-precision MC simulations [31]. Furthermore, the analysis in this work considers the addition of the disorder potential after the preparation of the gas as often done in experiments [39]. The predicted shift in the critical temperature is shown in Figure 2.3 (orange line). While the two classical field theory results agree at zero disorder strength and at the critical disorder strength which destroys the superfluid phase transition, Bourdel's results show an increase in the critical temperature for weak disorder, which is explained by the compression of the gas. Therefore, at any intermediate disorder strength, the decrease in the critical temperature is smaller compared to the simple model of Carleo *et al.*

The problem was further analysed by Carleo *et al.* with Quantum Monte Carlo (QMC) methods [15]. It is important to note that this work also considers parameters for which the LDA is valid. The results demonstrate a universal linear scaling of the critical temperature given by

$$T_c(\bar{V}) = T_c^0 \left(1 - \frac{\bar{V}}{\mu} \right), \quad (2.22)$$

where \bar{V} is the mean disorder strength and T_c^0 is the critical temperature in absence of disorder. An intriguing feature of this result is that as opposed to [38], T_c^0 is known at different values of the dimensionless interaction strength \tilde{g} [40]. The QMC results also show that the nature of the phase transition is topologically protected; the correlation functions remain identical to those in a non-disordered system but with a reduced critical temperature up to $\bar{V} = \mu$. Above critical disorder strength $\bar{V} = \mu$, the superfluid phase vanishes even at zero temperature. Therefore, by comparing 2.22 and 2.17, one possible interpretation for the effects of the disordered potential is the renormalisation of the superfluid density to a lower value. At higher disorder strengths a transition from the normal to an insulating phase with a three-phase coexistence at zero temperature is predicted [41]. Exploring the transitions between these phases is of great fundamental interest. Experiments with liquid ^4He films adsorbed on disordered substrates [42], measurements of momentum distributions [39] and matter-wave interference of ultracold atomic gases [43] qualitatively showed a BKT-type transition with a loss of coherence with increasing disorder. However, a detailed characterisation of *local* phase fluctuations is required for a better insight into the physics of these systems.

2.8 Bilayer 2D Bose gases

For the experimental investigations in this work, we typically use bilayer 2D Bose gases. While this split quantum gas approach may initially complicate the theoretical description and the experimental techniques, it offers access to an important observable: the complex amplitude of interference fringes A_Q , which we measure by matter-wave interferometry. Here, $Q = md/\hbar t_{TOF}$ is the wavevector of the fringes that form for an initial cloud separation of d and time of flight t_{TOF} . For

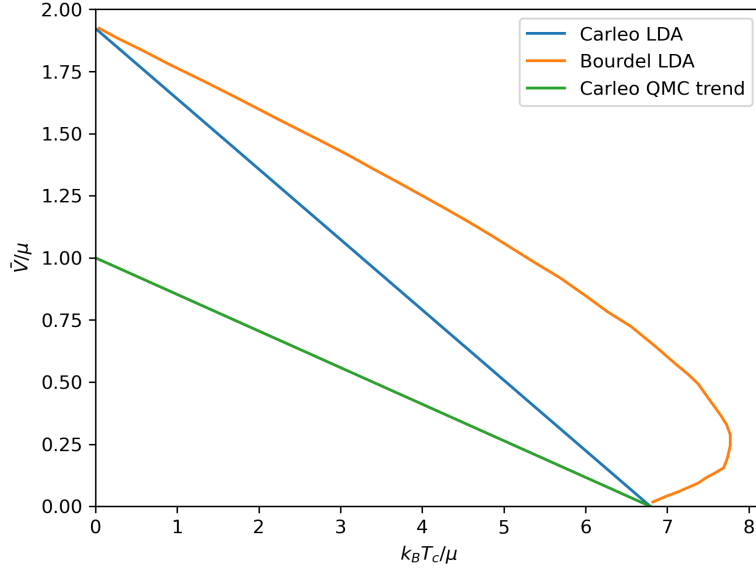


Figure 2.3: Comparison of different theoretical results of the BKT critical temperature in a disordered system.

our highly anisotropic system and realistic values of t_{TOF} , it is safe to assume that that the clouds do not expand significantly perpendicular to their separation. In this case, it can be shown, assuming Gaussian wavepacket evolution, that [44, 45]

$$\hat{A}_Q(x, y) = \hat{\Psi}_1^\dagger(x, y) \hat{\Psi}_2(x, y) \quad (2.23)$$

where $\hat{\Psi}_i$ are Bosonic field operators which describe the *in-situ* systems. Experimentally, we detect atoms in a region of finite width of Δy along the imaging axis (y) and Δx perpendicular to the imaging axis due to the finite imaging resolution, therefore the measured quantity is

$$\hat{A}'_Q(x_0) = \int_{x_0 - \Delta x/2}^{x_0 + \Delta x/2} \int_{-\Delta y/2}^{\Delta y/2} \hat{\Psi}_1^\dagger(x, y) \hat{\Psi}_2(x, y) dx dy. \quad (2.24)$$

From the measured eigenvalue of $\hat{A}'_Q(x_0)$ we can extract the amplitude and phase of the fringes which contain information about the atom number difference and

phase difference between the two in-situ fields around location x_0 . Intuitively, choosing a small Δy results in a more local measurement as the integration effects in 2.24 are less significant. From a large data set of many measurements, it is possible to extract the correlation function $g_1(x)$ from the spatial variations of the phase difference and sample the distribution of interference contrast. Both of these quantities are of great importance in characterising quantum many-body systems in and out of equilibrium [12, 13, 46, 47].

2.9 The effect of interlayer coupling

The interaction between two layers of bilayer 2D Bose gases is described by the Josephson-like Hamiltonian [48, 49]

$$\mathcal{H}_{int} = -J_{\perp} \int \cos(\theta_1(\mathbf{r}) - \theta_2(\mathbf{r})) d\mathbf{r}. \quad (2.25)$$

Here, $J_{\perp} \geq 0$ corresponding to an interaction term that attempts to reduce the local phase difference between the two gases. It is important to consider that the coupling constant J_{\perp} scales linearly with the mean density of the layers such that the interaction energy is extensive. If the two gases have the same superfluid fraction such that the prefactors in the low energy Hamiltonian (2.12) of the individual gases are equal, the Hamiltonian of the entire system can be expressed in terms of the relative $\phi = \theta_1 - \theta_2$ and common-mode phase $\Phi = \theta_1 + \theta_2$ as

$$\mathcal{H} = 2K \int (\nabla\phi(\mathbf{r}))^2 + (\nabla\Phi(\mathbf{r}))^2 d\mathbf{r} - J_{\perp} \int \cos(\phi(\mathbf{r})) d\mathbf{r}, \quad (2.26)$$

where $K = \hbar^2 n_s / 2m$.

In the limit of $J_{\perp} \rightarrow 0$, which is easily achievable due to the fast decay of J_{\perp} with cloud separation, the Hamiltonian reduces to that of two independent gases. In

this scenario, the experimentally detected g_1 correlation function is the square of the correlation function in the corresponding single-layer system, since the fluctuations of both gases contribute to the measured interferometric phase profile. This means that for exponential correlation functions we measure half the correlation length while for algebraic correlation functions, we measure double the algebraic exponent compared to a single layer system.

In the limit of strong coupling ($J_{\perp} \rightarrow \infty$), the local phases of the two gases are locked as any fluctuation would have an infinite energy cost. In this case, the Hamiltonian reduces to that of a single cloud but with twice the superfluid density. Intuitively, infinite coupling strength corresponds to no separation between the clouds. This results in a BKT critical temperature that is twice that of a single cloud.

For finite coupling strengths, we can argue that the coupling tries to minimise the relative phase but at the same time the common-mode phase exhibits fluctuations independent from the interlayer coupling rate and can undergo the BKT transition. Under the BKT critical point in the corresponding non-interacting system, both the relative and common-mode phases decay algebraically. For intermediate temperatures and coupling strengths, a phase appears in which the coupling between the layers causes the correlations in the common-mode phase to decay exponentially but correlations in the relative phase decay algebraically due to the phase-locking effect of coupling. At higher temperatures correlations of both decay exponentially.

In addition to exploring the effects of coupling on bilayer 2D Bose gases in equilibrium, dynamic control over J_{\perp} provides an important tool to study out of equilibrium dynamics. A sudden reduction of the interlayer coupling strength can quench the system across the BKT critical point by reducing the critical temper-

ature. The behaviour of such a system is described by the reverse Kibble-Zurek mechanism which predicts a slow proliferation of topological defects. This can be realised by rapid splitting of a single layer 2D Bose gas which was investigated theoretically in [50] and experimentally in [13]. The exponential scaling of J_{\perp} with well separation allows for similar experiments with greater precision involving only small changes in the trapping potential. Similarly, a sudden increase in T_c can probe the Kibble-Zurek process, where the system is expected to evolve towards a more correlated state. Finally, by applying periodic modulations of J_{\perp} such that the system would repeatedly cross the BKT critical point in equilibrium, diverse dynamical features such as time-averaged non-equilibrium states with enhanced coherence are predicted [51].

2.9.1 RG theory of bilayer gases

The rich physics of bilayer 2D Bose gases is captured in a set of RG equations which were first derived by Benfatto *et al.* [52]. Mathey *et al.* extended these to treat all relevant couplings and their interactions with one another. The full set of RG equations are given by [48]

$$\frac{dJ_{\perp}}{dl} = \left(2 - \frac{T}{2\pi J_a}\right) J_{\perp}, \quad (2.27)$$

$$\frac{dJ_a}{dl} = \alpha_2 \left(\frac{J_{\perp}^2}{4\pi^4 J_a} - 4 \frac{A_a^2}{T^4} J_a^3 - \frac{A_1^2}{2T^4} (J_s + J_a) J_a^2 \right), \quad (2.28)$$

$$\frac{dJ_s}{dl} = -\alpha_2 \left(2 \frac{A_s^2}{T^4} J_s^2 + \frac{A_1^2}{4T^4} (J_s + J_a) J_s \right) 2J_s, \quad (2.29)$$

$$\frac{dA_1}{dl} = \left(2 - \frac{\pi(J_s + J_a)}{2T} + \alpha_3 \frac{A_s J_s + A_a J_a}{T^2} \right) A_1, \quad (2.30)$$

$$\frac{dA_a}{dl} = \left(2 - 2\pi \frac{J_a}{T} \right) A_a + \alpha_3 \frac{A_1^2 (J_s - J_a)}{2T^2}, \quad (2.31)$$

$$\frac{dA_s}{dl} = \left(2 - 2\pi \frac{J_s}{T} \right) A_s + \alpha_3 \frac{A_1^2 (J_a - J_s)}{2T^2}. \quad (2.32)$$

where $J_{a/s}$ and $A_{a/s}$ are the (anti)symmetric superfluid coupling and vortex fugacities respectively, while the single-vortex fugacity is A_1 . $\alpha_{2/3}$ are model-dependent non-universal constants of order unity and they do not alter the qualitative results. From these equations, it is clear that the quantity of primary interest J_a , which we observe by matter-wave interferometry, depends on all other RG couplings. Therefore, these equations are important for interpreting experimental results at finite interlayer coupling strengths. As an example, we show the BKT critical temperature of the relative phase as a function of experimentally relevant values of J_\perp in Figure 2.4. For this calculation, we use a typical experimental density of $10 \mu\text{m}^{-2}$, however, we simulate a large system ($l_{max} = 10$) to suppress finite size effects and detect the critical temperature as a sharp divergence of $1/J_a$.




Figure 2.4: RG calculation of the BKT critical temperature of the relative phase mode as a function of J_\perp . The density used for this calculation is $10 \mu\text{m}^{-2}$ and we consider a large system to identify the phase transition as a divergence of $1/J_a$.

3 | RF-dressed adiabatic potentials

In this chapter, a brief overview of RF-dressed adiabatic potentials is given, which is the primary method for trapping (ultra)cold atoms in this work. We start by describing the interaction between atoms and external magnetic fields. We derive the RF-dressed Hamiltonian and we provide approximate solutions using the Rotating Wave Approximation (RWA). We use these to explain how a highly anisotropic trap can be created by the combination of RF and DC magnetic fields. Finally, dressing with multiple RFs is described which is used in this work to create a controllable double-well potential.

3.1 Magnetic trapping in DC fields

The effect of external magnetic fields on atomic energy levels is described by the Zeeman effect. The energy level shifts are derived from the interaction Hamiltonian

$$\hat{\mathcal{H}} = -\hat{\boldsymbol{\mu}} \cdot \mathbf{B}, \quad (3.1)$$

where $\hat{\boldsymbol{\mu}}$ is the magnetic moment operator of the atom and \mathbf{B} is the external magnetic field vector. For weak external fields, the interaction can be treated perturbatively and to first order, the eigenenergies can be expressed as

$$E_{m_F} = m_F g_F \mu_B |\mathbf{B}|, \quad (3.2)$$

where m_F is the quantum number labelling the state, g_F is the Landé g-factor, and μ_B is the Bohr magneton. This treatment is valid for $E_{m_F} \ll \Delta E_{HFS}$, which is true for typical RF-dressed setups, as $E_{m_F} \sim 1\text{-}10$ MHz while $\Delta E_{HFS} \sim 1\text{-}10$ GHz. For ^{87}Rb in the $|F = 1\rangle$ hyperfine level which is the atomic species used in

this work, $m_F = 0, \pm 1$ and $g_F = -\frac{1}{2}$ which results in a Zeeman splitting of -0.7 MHz/G¹. The shift in energy levels for the magnetically sensitive states can be exploited to confine atoms around spatial extrema of the magnetic field strength. However, only a local minimum of $|\mathbf{B}|$ can be created in a current-free region [55] hence only the low-field seeking $|m_F = -1\rangle$ can be trapped in practise by DC magnetic fields. A common magnetic field configuration for achieving this is the quadrupole field created in the middle of an anti-Helmholtz coil pair, widely employed in experiments. The quadrupole magnetic field is given by

$$\mathbf{B} = b'(x\mathbf{e}_x + y\mathbf{e}_y - 2z\mathbf{e}_z) \quad (3.3)$$

where b' is the quadrupole gradient, typically measured in G/cm and \mathbf{e}_i are the Cartesian unit vectors. Magnetic quadrupole traps are simple to realise and have a trap depth in the mK range, therefore they are ideal for storing laser-cooled atoms and have the ability to further cool them by RF-induced evaporation [56]. However, Majorana spin-flip losses affect the atoms near the zero of the magnetic field strength which prevents cooling to ultracold temperatures and prohibits the creation of a degenerate quantum gas in the magnetic quadrupole trap [57]. To circumvent this issue, various methods have been developed which include the Ioffe-Pritchard trap [58], the TOP trap [59], or the plugged quadrupole trap [60].

3.2 Dressed atoms

While the semi-classical two-level atom model can be used to describe the interaction between atoms and electromagnetic (EM) radiation, it treats the interaction perturbatively and thus is valid only for weak coupling (where the Rabi frequency

¹These values are approximate, the multi-electron structure of rubidium [53] and quantum electrodynamic effects [54] result in small corrections.

is significantly less than the natural linewidth of the transition). To overcome the limitations of this model, the concept of dressed atoms was introduced in 1977 by Cohen-Tannoudji to explain atomic fluorescence in intense laser fields [61].

3.2.1 The dressed Hamiltonian

The key idea behind the dressed-atom picture is to consider the Hamiltonian of the combined atom-field system, which for an external magnetic field pointing in the z -direction reads as

$$\hat{\mathcal{H}} = \underbrace{\mu_B g_F B_0 \hat{F}_z}_{\hat{\mathcal{H}}_{atom}} + \underbrace{\hbar \omega_{RF} (\hat{a}^\dagger \hat{a} + 1/2)}_{\hat{\mathcal{H}}_{RF}} + \hat{\mathcal{H}}_{int}, \quad (3.4)$$

where \hat{F}_z is the z -component of the atomic generalised angular momentum operator, \hat{a}^\dagger is the creation operator for the RF electromagnetic field, $\hat{\mathcal{H}}_{int}$ is the interaction Hamiltonian. Without the interaction term, it is straightforward to diagonalise the Hamiltonian as the eigenstates are tensor products of the atomic Zeeman states and the Fock states describing the electromagnetic field. We denote such a state with photon number N and angular momentum projection m_F as $|N, m_F\rangle$. For the lower hyperfine level in the ground configuration of rubidium, the interaction Hamiltonian is given by

$$\hat{\mathcal{H}}_{int} = \lambda_+ (\hat{F}_+ \hat{a} + \hat{F}_- \hat{a}^\dagger) + \lambda_- (\hat{F}_+ \hat{a}^\dagger + \hat{F}_- \hat{a}) + \lambda_z \hat{F}_z (\hat{a}^\dagger + \hat{a}) \quad (3.5)$$

where $\lambda_{\pm,z}$ describe the coupling strength between the atom and the EM field for a complete set of polarisations (left and right-handed circular polarisation in the x - y plane as well as linear polarisation along z) and \hat{F}_\pm are the atomic angular momentum raising and lowering operators [62]. The different terms in the interaction Hamiltonian can be understood by considering the conservation of angular

momentum of the entire system. The first term describes a photon's absorption (emission) from (into) the σ_+ mode of the EM field, which causes the atomic angular momentum to increase (decrease) about the quantisation axis. The second term describes a similar interaction with the σ_- mode of the EM field. The last term captures the interaction with a π polarised photon leaving the atomic angular momentum along the quantisation axis unchanged. It is not possible to diagonalise this Hamiltonian analytically. In the next section, we find an approximate solutions to the eigenvalues and eigenstates.

3.2.2 An approximate solution to the dressed Hamiltonian

Since the $|F = 1\rangle$ level has a negative g-factor, the set of states

$$\mathbb{D} = \{|N - 1, -1\rangle, |N, 0\rangle, |N + 1, 1\rangle\} \quad (3.6)$$

are degenerate in energy if $\hbar\omega_{RF} = \mu_B g_F B_0$. These states are coupled by the σ_- mode of the EM field, therefore we can neglect the first and last terms in 3.5 (those proportional to λ_+ and λ_z). Furthermore, we can neglect any other states since they are well-separated in energy from the states in \mathbb{D} . Considering a coherent state with a large amplitude, $\langle N - 1 | \hat{a} | N \rangle = \sqrt{\langle N \rangle}$ and $\langle N + 1 | \hat{a}^\dagger | N \rangle = \sqrt{\langle N \rangle + 1} \approx \sqrt{\langle N \rangle}$ for large $\langle N \rangle$ where $\langle N \rangle$ is the mean photon number. Within this approximation, the RF-dressed Hamiltonian reduces to

$$\mathcal{H} = \begin{pmatrix} \hbar N \omega_{RF} - \hbar \delta & \sqrt{2\langle N \rangle} \lambda_- & 0 \\ \sqrt{2\langle N \rangle} \lambda_- & \hbar N \omega_{RF} & \sqrt{2\langle N \rangle} \lambda_- \\ 0 & \sqrt{2\langle N \rangle} \lambda_- & \hbar N \omega_{RF} + \hbar \delta \end{pmatrix} \quad (3.7)$$

represented in the basis \mathbb{D} where we have introduced the detuning from magnetic resonance as

$$\delta = \omega_{RF} - |g_F| \mu_B B_0 / \hbar. \quad (3.8)$$

On resonance, $\delta = 0$ and the splitting between subsequent energy eigenvalues is $2\sqrt{\langle N \rangle} \lambda_-$ therefore we identify the Rabi frequency $\Omega = 2\sqrt{\langle N \rangle} \lambda_-$ which is the natural unit for the amplitude of the EM field. For a general detuning, the eigenvalues of 3.7 are

$$E_{\tilde{m}_F, \tilde{N}} = \hbar \omega_{RF} \tilde{N} + \hbar \tilde{m}_F \sqrt{\delta^2 + \Omega^2} \quad (3.9)$$

where $\tilde{m}_F = 0, \pm 1$ and \tilde{N} are new quantum numbers labelling the state of the atom and the EM field. These eigenvalues are shown as a function of normalised Zeeman splitting in Figure 3.1 a for three different manifolds given by consecutive values of \tilde{N} . Note that while the adiabatic energies in a given manifold do not cross, states belonging to different manifolds may cross. The state with $\tilde{m}_F = 1$ has a local minimum at a non-zero Zeeman splitting therefore atoms in this state can be trapped without suffering from Majorana losses and we consider this to be the trapping state. Conversely, the state with $\tilde{m}_F = -1$ has a local minimum at zero Zeeman splitting. Atoms in this state are affected by Majorana losses, therefore we call this state anti-trapping. The RF-dressed eigenstates can be expressed as a linear combination of the bare Zeeman states at various detunings since they form a complete set for the atom. This is shown in Figure 3.1 b. For large blue detuning, the state is well-approximated with the $m_F = 1$ state while for large red detuning, the state is well-approximated with the $m_F = -1$ state.

The above description is no longer valid for RF polarisations other than circular. For a general polarisation of the EM field, the Hamiltonian needs to be solved numerically with a truncated basis. In practice, one has to rely on numerical

methods even in the case of circular polarisation. While the polarisation may be well-defined and uniform in the lab frame, in the frame of the atoms (given by the direction of the DC magnetic field), the polarisation is typically non-uniform.

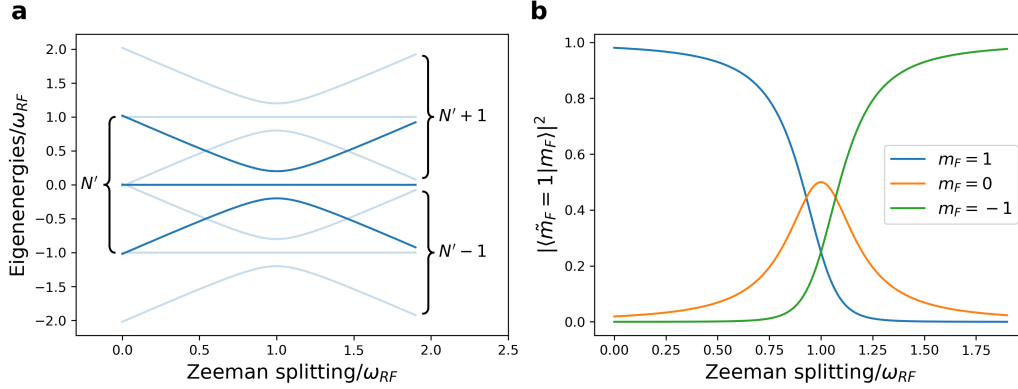


Figure 3.1: Dressed eigenstates and their relation to the bare Zeeman states. **a**: Dressed-atom energy eigenvalues as a function of normalised Zeeman splitting. **b**: Probability of finding the $\tilde{m}_F = 1$ trapping state in the Zeeman states as a function of Zeeman splitting.

3.2.3 The rotating-wave Hamiltonian

An alternative approach to finding the eigenvalues and eigenstates of the RF-dressed system is by using a semi-classical Hamiltonian. Within this approximation, we treat the atoms as three-level quantum systems interacting with a classical time-dependent magnetic field. Analytic solutions to the Hamiltonian for this system can be found by using the RWA. We define an oscillating magnetic field of arbitrary polarisation as

$$\mathbf{B} = \frac{B_x}{2} e^{-i(\omega t + \phi_x)} \mathbf{e}_x + \frac{B_y}{2} e^{-i(\omega t + \phi_y)} \mathbf{e}_y + \frac{B_z}{2} e^{-i\omega t} \mathbf{e}_z + c.c. \quad (3.10)$$

where we use $\omega \equiv \omega_{RF}$ for the sake of clarity. In this case, the interaction Hamiltonian is given by

$$\hat{\mathcal{H}}_{int} = \frac{g_F \mu_B}{\hbar} \mathbf{B} \cdot \hat{\mathbf{F}} = \frac{g_F \mu_B}{\hbar} \left(\frac{\hat{F}_+ e^{-i\omega t}}{4} (B_x e^{-i\phi_x} - iB_y e^{-i\phi_y}) + \frac{i\hat{F}_- e^{-i\omega t}}{4} (B_x e^{-i\phi_x} + iB_y e^{-i\phi_y}) + \frac{\hat{F}_z B_z e^{-i\omega t}}{2} + c.c. \right) \quad (3.11)$$

Using the well-known representation of the $\hat{F}_{\pm,z}$ operators in the basis of Zeeman states, we can transform the full semi-classical Hamiltonian $\hat{\mathcal{H}} = \hat{\mathcal{H}}_{atom} + \hat{\mathcal{H}}_{int}$ to a frame rotating about z with frequency ω

$$\hat{\mathcal{H}} \rightarrow \hat{U} \hat{\mathcal{H}} \hat{U}^\dagger - i\hbar \hat{U} \frac{\partial \hat{U}^\dagger}{\partial t}, \quad (3.12)$$

where the unitary transformation is given by $\hat{U} = e^{i\hat{F}_z \omega t / \hbar}$. The transformation results in time-independent terms as well as oscillatory terms which we ignore within the RWA. Furthermore, the approximation ignores the oscillating magnetic field along \mathbf{e}_z , therefore we set $\phi_x = 0$ and redefine ϕ_y to be the phase shift relative to ϕ_x . This yields

$$\mathcal{H} = \begin{pmatrix} -\hbar\delta & \frac{g_F \mu_B}{2\sqrt{2}} (B_x + iB_y e^{-i\phi_y}) & 0 \\ \frac{g_F \mu_B}{2\sqrt{2}} (B_x - iB_y e^{i\phi_y}) & 0 & \frac{g_F \mu_B}{2\sqrt{2}} (B_x + iB_y e^{-i\phi_y}) \\ 0 & \frac{g_F \mu_B}{2\sqrt{2}} (B_x - iB_y e^{i\phi_y}) & \hbar\delta \end{pmatrix}. \quad (3.13)$$

This matrix resembles 3.7 however, the off-diagonal elements differ. This can be fixed by applying a further, time-independent rotation about \mathbf{e}_z by the unitary operator $e^{i\hat{F}_z \alpha / \hbar}$. The rotation angle α can be found by equating the different off-diagonal matrix elements giving $\alpha = \arctan \left(\frac{B_y \cos(\phi_y)}{B_y \sin(\phi_y) - B_x} \right)$. This transformation

allows us to define the Rabi frequency as

$$\Omega = \frac{g_F \mu_B}{2\hbar} \sqrt{B_x^2 + B_y^2 - 2B_x B_y \sin(\phi_y)} \quad (3.14)$$

and the rotating wave Hamiltonian as

$$\mathcal{H}_{RWA} = \begin{pmatrix} -\hbar\delta & \Omega/\sqrt{2} & 0 \\ \Omega/\sqrt{2} & 0 & \Omega/\sqrt{2} \\ 0 & \Omega/\sqrt{2} & \hbar\delta \end{pmatrix}. \quad (3.15)$$

This matrix is straightforward to diagonalise and yields eigenvalues which are in agreement with the approximate solution to the dressed Hamiltonian. Note that [3.14](#) implies that the Rabi frequency is maximum for $\phi_y = -\pi/2$ and zero for $\phi_y = \pi/2$. This demonstrates the importance of the polarisation of the dressing magnetic field in RF-dressed traps.

3.3 Confining cold atoms in RF-dressed potentials

In the previous section, the eigenenergies of atoms dressed by an oscillating magnetic field were calculated. Here, we describe how this can be exploited to trap atoms. The eigenenergies are governed by two terms in [3.9](#) which, in general, are position-dependent: the detuning and the Rabi frequency. The spatial variation in the detuning arises from non-uniform static magnetic field strength while the spatial variation in the Rabi frequency comes from non-uniform polarisation of the dressing field in the frame of the atoms. Provided that the RF field is strong enough and the atoms traverse the avoided crossing sufficiently slowly, there is a high probability of remaining in the initial \tilde{m}_F dressed state by virtue of the

adiabatic theorem² [64]. Therefore, we can consider that the atom is moving in a conservative potential given by

$$V(\mathbf{r}) = \hbar\tilde{m}_F\sqrt{\delta(\mathbf{r})^2 + \Omega(\mathbf{r})^2}. \quad (3.16)$$

The spatial dependence of the detuning is straightforward to evaluate using 3.8 for a given magnetic field. However, the previous expression for the Rabi-frequency (3.14) was derived for a DC magnetic field along the z-direction, which provides the quantisation axis. A more general expression is [65]

$$\Omega(\mathbf{r}) = \frac{|g_F|\mu_B|B|}{\hbar}|\boldsymbol{\epsilon} \times \mathbf{u} + is\mathbf{u} \times \boldsymbol{\epsilon} \times \mathbf{u}|, \quad (3.17)$$

where, $\boldsymbol{\epsilon}$ is the RF polarisation vector, \mathbf{u} is the unit vector pointing along the DC magnetic field, and s is the sign of g_F . While both 3.14 and 3.17 hold within the RWA, which assumes circular polarisation about the quantisation axis, these are good approximations for arbitrary polarisations. The most important effect beyond the RWA is the Bloch-Siegert shift [66] which moves the minimum of the RF-dressed potential. However, the shift remains small as long as $\Omega(\mathbf{r}) \ll \omega_{RF}$, which is true in the scope of this work. Beyond RWA effects can be accounted for in numerical solutions based on Floquet's theorem [67].

3.3.1 The dressed quadrupole trap

One of the most commonly used DC magnetic field configuration for RF-dressed traps is the quadrupole field. Around the zero of the magnetic field the isomagnetic surfaces are oblate spheroids of aspect ratio 2 (see 3.3). Considering a circularly

²For the simplified, analytically solvable Landau-Zener model of an avoided crossing, the probability of not remaining in the initial adiabatic state is suppressed exponentially with coupling strength [63].

polarised dressing field in the x - y plane, the spatial variation of the Rabi frequency can be calculated from 3.17, giving

$$\Omega(x, y, z) = \frac{|g_F| \mu_B B}{\hbar} \left(1 \pm \frac{2z}{\sqrt{x^2 + y^2 + 4z^2}} \right), \quad (3.18)$$

where the sign of the second term depends on the sign of $g_F b'$. Therefore, by choosing the correct handedness of circular polarisation, the coupling strength is maximised at the bottom of the shell and decays to zero linearly with increasing height. Since $2R \equiv \sqrt{x^2 + y^2 + 4z^2}$ is a constant, this results in an axially symmetric trap. This is visualised in Figure 3.2 a. Vertically, confinement is strong because of the increase in detuning perpendicular to the surface of the spheroid. Horizontally, the atoms are weakly confined by gravity as the isomagnetic surface curves upwards, and this is made even weaker by the reduction of Rabi frequency away from the axis. Considering the effects of gravitational sag, around the minimum of the potential the radial and axial trapping frequencies are given by [68]

$$\omega_r = \sqrt{\frac{g}{4R}} \left(1 - \frac{\tilde{m}_F \hbar \Omega_0}{2mgR} \sqrt{1 - \beta^2} \right)^{1/2} \quad (3.19)$$

and

$$\omega_z = \frac{2|g_F b'| \mu_B}{\hbar} \sqrt{\frac{\tilde{m}_F \hbar}{m \Omega_0}} (1 - \beta^2)^{3/4}, \quad (3.20)$$

where g is the gravitational acceleration, Ω_0 is the Rabi frequency at the bottom of the shell and $\beta = \frac{mg}{2g_F \tilde{m}_F \mu_B b'}$ is the ratio of gravitational and magnetic quadrupole forces along z .

For a linearly polarised dressing field along the y -direction, the Rabi frequency is

$$\Omega(x, y, z) = \frac{|g_F| \mu_B B}{\hbar} \left(1 - \frac{y^2}{\sqrt{x^2 + y^2 + 4z^2}} \right) \quad (3.21)$$

having nodes where the RF field is parallel with the DC magnetic field as shown in Figure 3.2 b. The mechanism for confinement is similar to the case of circular polarisation. The trapping frequencies along z and y are identical to the axial and radial trapping frequencies in the circular dressing case. However, the trapping frequency along x is [69]

$$\omega_x = \sqrt{\frac{g}{4R}} \quad (3.22)$$

since the confining mechanism is purely gravitational.




Figure 3.2: Rabi frequency on the resonant spheroid normalised to the maximum value at the bottom for circular RF polarisation in the x - y plane (**a**) and linear polarisation along y (**b**).

For typical experimental parameters, using circular or linear polarisation, the trapping frequencies in the x - y plane are in the 8-15 Hz range while along z it is in the 300-1200 Hz range. Therefore, the dressed quadrupole trap is an ideal platform for realising quasi-2D quantum gases [68]. The choice of RF polarisation is determined by the loading scheme used and experiment one wishes to perform. Frequency-swept loading schemes [70] or the method developed in this work (Chapter 5) can be used with either circular or linear RF polarisation. Loading from a TOP trap on the other hand is only possible with circular RF polarisation [71]. Regarding the experimental considerations, due to the spatial dependence

of the Rabi frequency and its magnitude at a given RF amplitude, circular shells offer longer lifetimes. On the other hand, circular polarisation requires more control parameters and a more complicated setup. Therefore, linear polarisation is preferred when stability is the priority or multiple radio-frequencies are needed.

3.4 Multiple RF-dressing

The RF-dressed formalism can be extended to RF-signals which contain multiple frequencies [72]. Using a quadrupole DC magnetic field, this results in multiple concentric shells. Since adiabatic energy levels belonging to the same manifold cannot cross, these concentric shells alternate between attracting and repelling the atoms. Each applied frequency creates an avoided crossing and the shape of the potential along the direction of the DC magnetic field gradient can be controlled by the amplitudes of each frequency component. For appropriately chosen amplitudes, $2n - 1$ frequencies create an n -well potential, for example a 3-frequency signal is used to create a controllable double-well potential in this work. The non-linear behaviour of RF-amplifiers leads to intermodulation products in their the output spectrum. To prevent these spurious RF components from causing atom losses, we use commensurate frequencies [73]. Synthesising such frequencies is straightforward using Direct Digital Synthesis (DDS) technology. Moreover, with commensurate frequencies the RF-dressed Hamiltonian is periodic in time and Floquet theory can be used to calculate the energy eigenvalues. A 3-RF potential created with 7.0, 7.1 and 7.2 MHz fields and 35 kHz Rabi frequencies is displayed in Figure 3.3 a for three manifolds of dressed states. By convention, we refer to the well realised by the lowest RF frequency as *upper* and the well realised by the highest RF frequency as *lower* as this reflects their position in the lab frame. In contrast with the case of single RF dressing, the spacing of the manifolds is given

by the frequency separation between the RF fields. Typically, the separation is of the order of the Rabi frequencies. This results in the positions of the wells being slightly offset from resonance. Moreover, the energy separation between the $|\tilde{m}_F = 0\rangle$ and the other states at the centre of the wells and barrier is only approximately the Rabi frequency, demonstrating the importance of numerical methods in calculating the RF-dressed potential. Using multiple radio-frequencies also creates second-order crossings close to the potential wells as indicated by the inset in Figure 3.3 a. Since higher order processes are weak, atoms are very likely to be lost via non-adiabatic processes around these points. Furthermore, losses were observed where adiabatic energy levels belonging to different manifolds cross [74]. For these reasons, MRF potentials have a significantly lower trap depth compared to single RF potentials, but in practice this is not a problem except for loading the trap. A significant advantage of MRF potentials is the dynamic control by modulating the amplitudes of the RF fields as demonstrated in Figure 3.3 b. The amplitude of the central "barrier" field is changed between 37.5 kHz and 87.5 kHz, which makes it possible to split a quantum gas. This protocol was first demonstrated in [75] and serves as the basis for detecting local phase fluctuations of Bose gases by matter-wave interferometry [12].

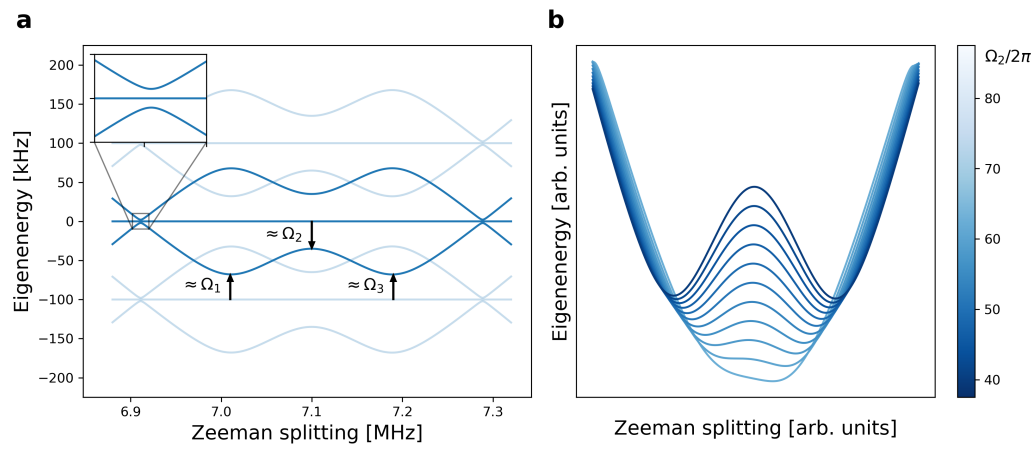


Figure 3.3: **a**: MRF double-well potential with frequencies 7.0, 7.1, and 7.2 MHz and 35 kHz Rabi frequencies. Inset shows a weak, second-order crossing. **b**: MRF potentials at various amplitudes of the middle frequency that creates the central barrier.

4 | Experimental Apparatus

This chapter describes the next generation of rubidium BEC apparatus in Oxford. In the initial phase, David Garrick made significant contributions and these are described in detail in his thesis [76]. This chapter starts with a brief overview of the experimental apparatus followed by a detailed account of the parts of the setup developed by the author.

4.1 Overview of the apparatus

The design of the new apparatus was motivated by the significant disadvantages of our existing apparatus. The previous generation Oxford experiment relies on successive transfers of (ultra)cold atoms between different traps until finally loading the atoms into an RF-dressed trap. This results in a low repetition rate, about one minute per run, because of the large number of steps that cannot be performed quickly without causing excessive heating or atom loss. Furthermore, a high number of control parameters are necessary many of which are coupled to each other, leading to lengthy optimisation sequences. In contrast, state-of-the-art BEC machines can produce a degenerate quantum gas at rates up to 1 Hz [77]. These high repetition rate experiments have various advantages. Firstly, data collection time is reduced which also means that the experiment is less susceptible to slow drifts in the control parameters. Secondly, the main focus of our research has been investigating 2D quantum gases, where quantum and thermal fluctuations play a crucial role and large amounts of data must be recorded to faithfully reconstruct the statistics of experimental observables. A typical data run can last as long as two weeks over which the experiment must be closely monitored. Over this timescale, external factors such as issues with cooling water or forced soft-

ware updates can interfere with the experiment. Thirdly, due to the complexity of typical experimental setups, optimising the experiment has always been key to conducting ultracold-atom research. Recently, Machine Learning (ML) algorithms have gained popularity as they can provide less biased results to optimisation problems and can discover correlations in high dimensional spaces [78]. The effectiveness of these methods often depends on sufficient exploration of the control parameter hyperspace. This means that a high repetition rate experiment has a natural advantage, as ML algorithms can make it function even better than a low repetition rate one. The two experimental setups are compared in Figure 4.1, highlighting the main differences.

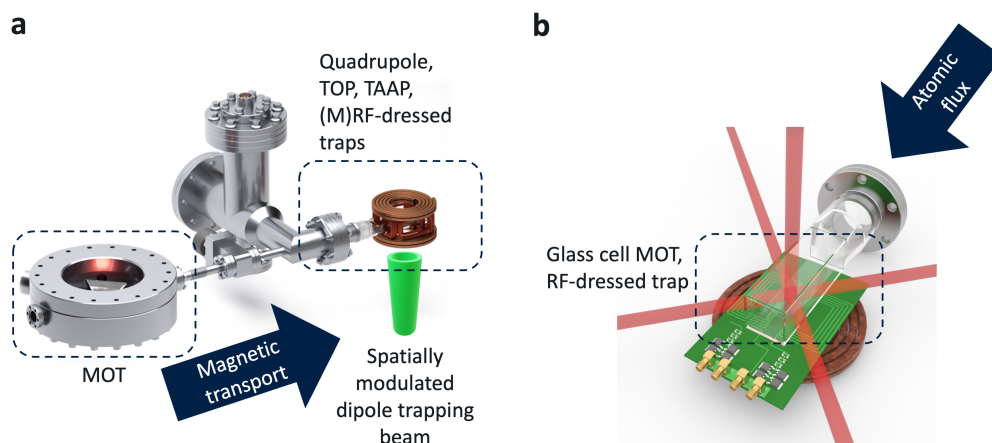


Figure 4.1: **a**: Model of the previous generation experimental setup. Atoms are captured in a MOT from a room-temperature vapour followed by a magnetic transport to the ultra-high vacuum glass cell. The atoms are transferred between various traps before the final RF-dressed potential. **b**: Model of the new experimental setup. A MOT is loaded from a high-flux source of cold atoms in the UHV glass cell. The laser-cooled atoms are transferred directly to the RF-dressed quadrupole trap without any intermediate magnetic traps. The quadrupole coil above the glass cell is not shown for illustration purposes.

The first stage of the experimental apparatus is a high-flux pyramidal source of cold ^{87}Rb atoms which was built and characterised by Sean Ravenhall [79]. The atoms are emitted from the source towards the second stage of the experiment,

an Ultra-High Vacuum (UHV) science chamber. This eliminates the need for transportation of the cold atoms, which often results in loss, heating, and long sequence times as in the old apparatus. In the science cell, the atoms are recaptured by a 3D Magneto-Optical Trap (MOT) and they are loaded into an RF-dressed quadrupole trap. Subsequently, they are further cooled to quantum degeneracy with RF-induced evaporative cooling. In the following sections, specific parts of the experimental apparatus are presented in detail.

4.1.1 The vacuum system

The vacuum system of the apparatus is a standard two-chamber design with differential pumping, shown in Figure 4.2 from three views: the pyramid source and glass cell sides as well as along a horizontal section. The vacuum system has undergone various upgrades over time and its history is described in Section 4.1.3.

4.1.2 Modelling the vacuum system

To estimate the equilibrium pressures in the vacuum system, we present a simple lumped element model as shown in Figure 4.3 which consists of chambers connected by pipes with certain conductances and pumps with specified pumping speeds. For pipes, the conductance in the molecular flow regime is given by

$$C = \frac{\bar{v}\pi d^3}{12l}, \quad (4.1)$$

where \bar{v} is the mean thermal velocity (267 m/s for ^{87}Rb at 20 degrees Celsius), d is the pipe diameter and l is the pipe length [80]. Using this formula we get $C_1 = 1.9$ l/s (considering a fully open valve and ignoring the right angle bend) and $C_3 = 0.5$ l/s. We neglect C_2 and C_4 as they are short CF64 and CF40 tubes. These values can be used to calculate the equilibrium pressures using the definition of




Figure 4.2: The vacuum system viewed from the pyramid source (top), from the glass cell (middle) and across a horizontal section (bottom). The rubidium source (a) is built from off-the-shelf DN16CF components. In the pyramid chamber (b) a VacIon 55 l/s Diode pump maintains vacuum (c). Between the pyramid and UHV chambers, a low conductance pipe of inner diameter 2.5 mm and length 25 mm ensures a pressure differential. A VacIon 55 l/s StarCell pump (d) and D-400 NEG pump (e) are responsible for maintaining ultra-high vacuum in the glass cell (f).

conductance

$$C = \frac{Q}{\Delta p}, \quad (4.2)$$

where $Q = k_B T \frac{dN}{dt}$ is the volumetric flow rate and Δp is the pressure difference between the two ends of the pipe. Assuming that the pressure in the source chamber is the vapour pressure of rubidium at 20 degrees Celsius (2.6×10^{-7} mbar [81]) and that the pyramid chamber is only pumped by the pyramid pump (since $C_3/C_2 \ll 1$), we obtain $p_{MOT} = 1 \times 10^{-8}$ mbar. This is an ideal pressure for a MOT and it can be controlled by heating the source chamber (increasing the vapour pressure of rubidium) or by partially closing the valve between the source chamber and the pyramid chamber (reducing C_1).

The pressure in the UHV chamber can be calculated similarly giving $p_{UHV} = 1 \times 10^{-10}$ mbar. Note that only rubidium was considered in this simple model. Moreover, we neglected the pumping rate of the NEG pump as its pumping speed depends strongly on the pumped gas species.

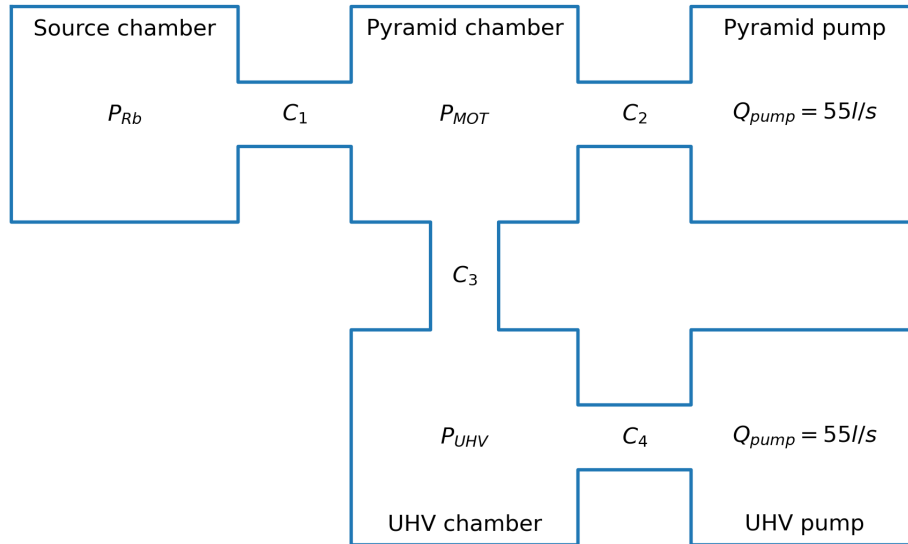


Figure 4.3: Simple lumped-element model of the vacuum system used to estimate equilibrium pressures.

4.1.3 The history of the vacuum system

The first two iterations of the vacuum system are described in David Garrick's thesis [76]. As the project progressed and atoms were successfully loaded into the RF-dressed trap, lifetime measurements indicated that the pressure in the UHV chamber was too high. To increase the pressure differential, a stainless steel low conductance pipe was placed in the original differential aperture. The inner diameter and length of the pipe were 2 mm and 67.5 mm respectively, giving a factor of 60 lower conductance between the pyramid chamber and the UHV glass cell compared to the previous differential aperture. This upgrade allowed very low pressures in the UHV chamber. The UHV ion pump controller displayed the lowest possible current which corresponds to 3×10^{-9} mbar during the bakeout when the chamber's temperature was 150 degrees Celsius. Therefore, the pressure was significantly lower with the chamber cooled back to room temperature. In the RF-dressed trap, 1/e lifetimes of 80 seconds were measured, approximately twice as long as in the old apparatus [82]. Despite the increased lifetime, the number of atoms trapped in the glass cell MOT was a factor of 5 lower than without the low conductance pipe. We attribute this to the pipe skimming the atomic flux. The low conductance pipe also significantly reduced the angular tolerance of the pyramid MOT source laser beam necessitating realignment every week. To increase the loading rate to the glass cell MOT, the temperature of the rubidium source was gradually increased which doubled the atom number in the glass cell MOT. However, the increased rubidium pressure in the pyramid chamber contaminated the pyramid ion pump giving a significant leakage current and a reduced pumping speed even after the rubidium source was cooled back to room temperature. To solve this issue, the vacuum chamber was opened and the pyramid ion pump was firstly rinsed with distilled water and then it was regenerated by baking at a high

temperature (400 degrees Celsius) without the magnets¹. The dimensions of the low conductance pipe were changed to an inner diameter of 2.5 mm and a length of 25 mm. During the machining it was noticed that the hole in the original pipe was slightly bent, which explains the low atomic flux and the difficulty of aligning the pyramid MOT source laser beam. The upgrade also involved changing the rubidium source to a modified version of the initial CF16 design described in David Garrick's thesis [76] which caused a vacuum leak. Based on anecdotal evidence, it was believed that the molten rubidium on viewport of the initial CF16 source caused the leak. The viewport which was replaced by a blank flange thereby providing a solution to the issues of the initial design. Moreover, the cross piece of the source was attached tilted (see Figure 4.2) to ensure that the lump of rubidium stays in the source chamber even when molten. Reducing the source pipe diameter was advantageous as the lower conductance makes it more difficult to contaminate the pyramid ion pump. This upgrade to the vacuum setup proved to provide a low enough pressure in the UHV chamber giving a 1/e lifetime of 35 seconds while maintaining a high loading rate into the glass cell MOT with a room-temperature rubidium source, and has been in use since December 2022.

4.2 The laser system

The laser system of this apparatus has undergone multiple iterations and this section describes the final state. The system utilises the existing lasers of the old apparatus as references, thus avoiding the need for expenditure on new items such as a new diode laser system or beatnote locking electronics. However, it was constructed to be independently tuneable through Acousto-Optic Modulators (AOMs) which offer sufficient freedom to set and modulate the laser frequencies.

¹This procedure was recommended by the manufacturer.

4.2.1 Overview

The overall scheme of the laser system is illustrated in Figure 4.4. A Toptica DL pro External Cavity Diode Laser (ECDL) is locked to the cycling transition of the D2 spectral line of ^{87}Rb using Modulation Transfer Spectroscopy [83] and the output from this laser is used as reference light. The reference light is used to lock another Toptica DL Pro ECDL with a beat note locking technique detuned 6476 MHz from the reference laser to provide repumping light². More details about this part of the laser system can be found in [73]. A small fraction of the reference reference laser output (1 mW) is frequency shifted with a double-pass AOM setup and split two ways to provide seed light for two Injection Locked Diode Lasers (ILDs). Injection locking allows a laser resonator to oscillate at the frequency of the weak drive provided that the seed frequency is close to one of the resonant frequencies of the laser cavity and there is sufficient gain [84]. The two ILDs are responsible for providing about 70 mW cooling light for the pyramid source and the glass cell MOT, which are further amplified later. The output of the ILDs is shifted back to the desired frequency by two independent double-pass AOM setups. While this shift could be performed before the ILDs, due to the finite diffraction efficiency of AOMs and the limited reference laser power, a single double-pass stage yields higher seed powers, which result in a more stable injection lock. To derive resonant imaging light which can be shuttered fast, a small fraction of the cell cooling ILD output is frequency shifted by a third double-pass AOM setup. Similarly, the repumper ECDL is frequency shifted with a single-pass AOM setup and seeds an ILD. The output is split three ways, providing the repumper light for the pyramid source, for the glass cell MOT, and for absorption imaging. After splitting, the imaging repumper and the MOT

²The light needs another 92 MHz shift to be resonant with the repumping transition which is achieved by a single pass through an AOM. This allows for fast shuttering of repumping light.

repumper beams are passed through a single-pass AOM setup allowing the light to be shuttered quickly. The cooling and repumping laser beams for the pyramid source and glass cell MOT are mixed using polarisation optics and amplified with Tapered Amplifiers (TAs) as shown in Figure 4.4.

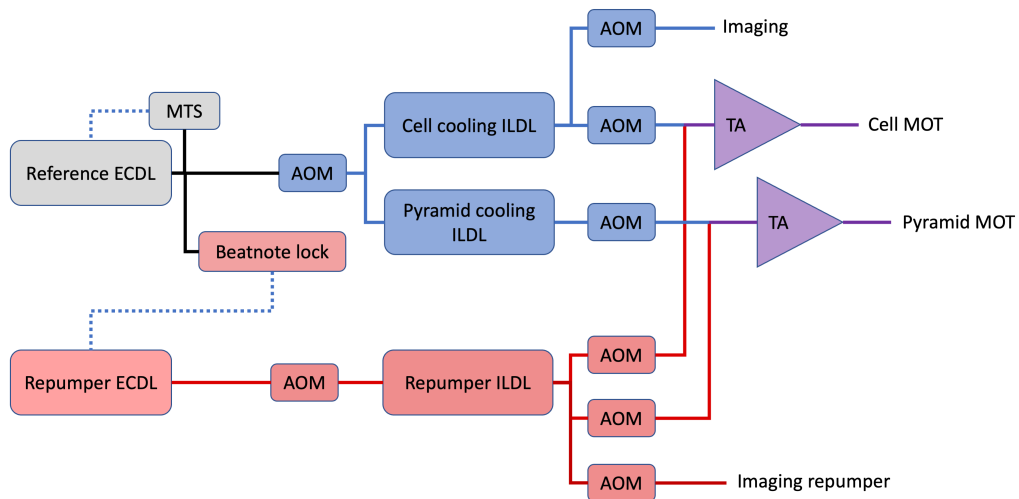


Figure 4.4: Schematic of the laser system. Solid lines indicate optical signals and dashed lines indicate electronic signals.

4.2.2 Injection locked diode lasers

The basic ILDL setup can be seen in Figure 4.5. A 783 nm, 200 mW laser diode (Panasonic, LNC728PS01WW) is used in a custom-made, temperature-controlled mount which includes a laser diode collimation lens tube (Thorlabs, LT230P-B). The laser diode current and temperature are controlled by Toptica DCC110 and DTC110 controllers. The DTC110 controller was originally designed for the DL100 ECDL setups. Due to the lower thermal mass of the laser diode mount compared to the DL100, the PID gains of the temperature controller had to be adjusted to prevent self-oscillations of the control loop. The laser output is passed through a two-stage Faraday isolator and coupled into a Polarisation Maintaining

Single Mode optical fibre³. Seed light is delivered through another optical fibre and injected through the side port of the isolator. Its polarisation is adjusted with a half waveplate to ensure that maximum intensity emerges from the input side of the isolator towards the laser diode. The seed beam must be accurately aligned to maximise the overlap between the laser output mode and the seed mode. Since a small fraction of the laser diode output emerges from the side port, optimal alignment is achieved by maximising the transmission of the laser output through the seed fibre. Optimal seed alignment not only improves the long-term stability of the injection lock but also extends the frequency tuning range. Injection locking is achieved by changing the laser diode current until an internal mode is sufficiently close to resonance with the seed light. To achieve this, a Fabry-Perot etalon is used to monitor the output spectrum of the ILDL. When far from resonance, the etalon signal shows a broad background and a small, sharp peak corresponding to the seed laser. As the laser diode current is tuned closer to the optimal value, the height of the sharp peak increases rapidly and saturates, indicating a locked laser. We found that the stability of the injection lock is excellent as the diode current needs small adjustments approximately once per two months. This is even true if the reference laser goes out of lock for an extended period before relocking. Furthermore, the frequency tuning range extends to 40 MHz towards the red detuning which is sufficient for the required applications.

4.2.3 Tapered amplifiers

While injection locked diode lasers can amplify the ~ 1 mW power of the seed light to ~ 100 mW, this output power is not sufficient for laser cooling applications on our setup. The pyramid source requires around 200 mW of optical power while

³Since optical fibres of this type are used predominantly in the laboratory, unless specified otherwise, the term optical fibre refers to this type throughout this work.




Figure 4.5: Schematic of the basic ILDL setup.

the glass cell MOT requires 100 mW in total. Finite AOM diffraction efficiencies, optical losses from fibre coupling, polarisation cleaning, and reflections from optical elements further reduce this power. Therefore, we further amplify the light by TAs, which are near-infrared semiconductor amplifiers capable of boosting the optical power from tens of mW to a few W. Such a high output power is possible as these devices have a gain medium of gradually increasing width to avoid damage due to excessive optical intensities. The schematic of the TA setups used can be seen in Figure 4.6.




Figure 4.6: Schematic of the tapered amplifier setup.

The seed light is delivered to the TA setup with an optical fibre which carries cooling and repumper light simultaneously in the same polarisation mode. This eliminates the need for precision alignment of cooling and repumper beams individually into the amplifier chip which is a long and iterative process. The light is passed through a Faraday isolator to ensure that no backwards propagating am-

plified spontaneous emission is coupled back into the optical fibre, which would cause unwanted optical feedback. The polarisation of the seed light is aligned with a half waveplate and the seed light is focused with an aspheric lens onto the single mode input facet of the TA diode⁴. The current through the TA diode is controlled by a Thorlabs LDC8040 laser diode driver and its temperature is regulated by a Peltier element driven by a Thorlabs TED8040 temperature controller. The output mode of the TA is highly astigmatic and divergent; the output beam is collimated first by an aspheric lens in the vertical direction and then by a cylindrical lens in the horizontal direction. The collimated output beam is then passed through a Faraday isolator to avoid damaging the TA gain medium by back reflections. Finally, the laser beam is coupled into an optical fibre for spatial mode cleaning and beam delivery. With careful adjustment of the beam diameter by a telescope, a coupling efficiency of 70 percent was achieved.

4.2.3.1 Eliminating spurious TA temperature drifts

Temperature control is crucial for tapered amplifiers as small changes in temperature can significantly affect the output power. The relationship between temperature and output power is highly non-linear as the TA can jump between different internal modes. It was observed that the strong radio-frequency signals used in the experiment (and other nearby experiments) were picked up by the long cables of the temperature controller. These AC signals were rectified by non-linear circuit elements, spuriously altering the temperature reading. Due to closed loop temperature regulation, this caused a periodic drift of temperature and unstable laser powers. To fix this issue, two methods were developed.

In the TA setups, an AD590 sensor was used to measure the temperature. This IC acts as a temperature-dependent current source with $1 \mu\text{A}/\text{K}$ output current. The

⁴Optica-Eagleyard EYP-TPA-0780-01000-3006-CMT03-0000

temperature controller converts the current to a voltage by a resistor which serves as the input signal for a PID controller. To filter the input signal, a capacitor must be placed in parallel with the current to voltage converting resistor. The circuit diagram of the TED8040 is proprietary, so the circuit had to be manually debugged. It was found that an empty pad labelled as R99 (as shown in Figure 4.7) which is parallel with the appropriate resistor was present. Adding a 100 nF capacitor significantly reduced drifts of the temperature reading in the presence of strong AC signals, improving the stability of the output power from the TA setups.




Figure 4.7: Location of the R99 pad on the TED8040 circuit board with a filtering capacitor, which improves the temperature stability in environments with strong RF signals.

The second solution involves replacing the AD590 with a thermistor. Since these components act as temperature-dependent resistors, the temperature is typically read out by driving a constant current through the thermistor and measuring the corresponding voltage drop. For this approach, the filtering can be achieved with a capacitor in parallel with the sensor, which can be implemented without modifying the controller circuit. However, the temperature controller uses two different

pins on its output connector for thermistors and therefore the control cable must be modified when employing this method.

4.2.4 The six-beam MOT distribution board

The standard configuration of a 3D MOT consists of three pairs of counterpropagating laser beams, each pair mutually perpendicular to the other two. In practise, the six beams can be reduced to three by retro-reflections, but this can lead to difficulties in alignment as well a small optical power imbalance in the counter-propagating beams due to optical losses and atomic absorption. More complicated schemes can further reduce the number of required beams to one [85], however, the previously mentioned effects are even stronger in this case. The power imbalance would be significant in our case as the UHV glass cell has no anti-reflection coating. For an estimated 4 percent Fresnel loss at every air to glass interface, the retroreflected beam would have 15 percent less optical power than the incident beam at the centre of the glass cell. Therefore, we decided to have six independent beams to allow full optical power control as we had sufficient optical power available. Using this approach, the beam positions and propagation directions are fully decoupled from each other, which helps the alignment process.

To split a single beam six ways, a custom-made, optical fibre-based distribution setup was designed as shown in Figure 4.8. The distribution board consists of a monolithic aluminium breadboard which was fabricated in the Oxford Physics Workshop. All components are directly mounted without the use of optical posts. Although this greatly limits modification possibilities, the requirements for such an optical setup are well-established. Being monolithic, the breadboard offers exceptional long-term stability; fibre coupling adjustments were only needed when output fibre connectors were moved accidentally. Despite using affordable Newport M05 mirror mounts, stability remains excellent. These mounts have several

benefits in this particular setup over other half-inch models such as the Thorlabs KM05. Firstly, the mounting hole is a through-hole which allows mounting from above with a screw. Secondly, a large contact area between the mount and the breadboard improves stability, while the contact area for other mounts is designed for half-inch posts. The footprint of this assembly, measuring 390mm x 220mm, is considerably less than a setup with post-mounted optics. Furthermore, the cost of the optomechanical components for an identical, post-mount setup would exceed £1000 with off-the-shelf components while the total production cost of the custom-made board was £750.




Figure 4.8: Rendering of the six-beam MOT distribution board.

The schematic of the optical layout can be seen in Figure 4.9. Light arrives to the distribution board via an optical fibre. It is collimated with a Thorlabs F240APC-780 collimator which allows optimal fibre coupling efficiency without using a relay lens system, even for the longest beam path on the distribution board. The collimators are clamped in double-bored holes by two nylon-tipped setscrews. This mounting method, inspired by the Thorlabs AD12F adapter, provides high

stability due to two well-defined parallel edges on which the collimator rests. The polarisation of the input beam is cleaned with a Polarising Beam Splitter (PBS) and a mirror pair is used to align the beam along the axis of six successive half-waveplate and PBS pairs. These are used to coarsely adjust the optical power in each of the six MOT beams. A mirror pair is used to couple the beams into the output optical fibres. Since the beams going towards the collimators are reflected from PBSs, their polarisation is slightly elliptical (considering a typical measured power extinction ratio of 100:1, this results in a polarisation ellipse of aspect ratio 10:1). This resulted in significant polarisation drifts at the outputs of the six MOT fibres. Therefore, PBSs were added on rotation mounts to clean and align the polarisation to the slow axis of each MOT fibre. Even though this method of polarisation alignment is at the expense of optical power, the effect is small as the collimators can be aligned with a few degrees precision manually. An earlier version of the setup used Thorlabs LPNIRE050-B economy film polarisers instead of PBSs. While they offer a good extinction ratio, we found that they introduce some fringing on large beams and their bad parallelism significantly reduces the fibre coupling efficiency as the polariser is rotated.

4.3 The MOT beam launching periscopes

The six MOT fibres connect to a cage-mount, periscope-based beam launching setup which prepares the free-space MOT beams and ensures optimal spatial alignment. The setup can be seen in Figure 4.10. This system, originally designed by David Garrick [76], was later modified to incorporate PID control of the optical beam powers. In each launching periscope, the optical fibre from the distribution board is connected to a Thorlabs SM1FCA fibre adapter plate. The emerging beam is collimated by a Thorlabs AC254-040-B 40 mm achromatic dou-

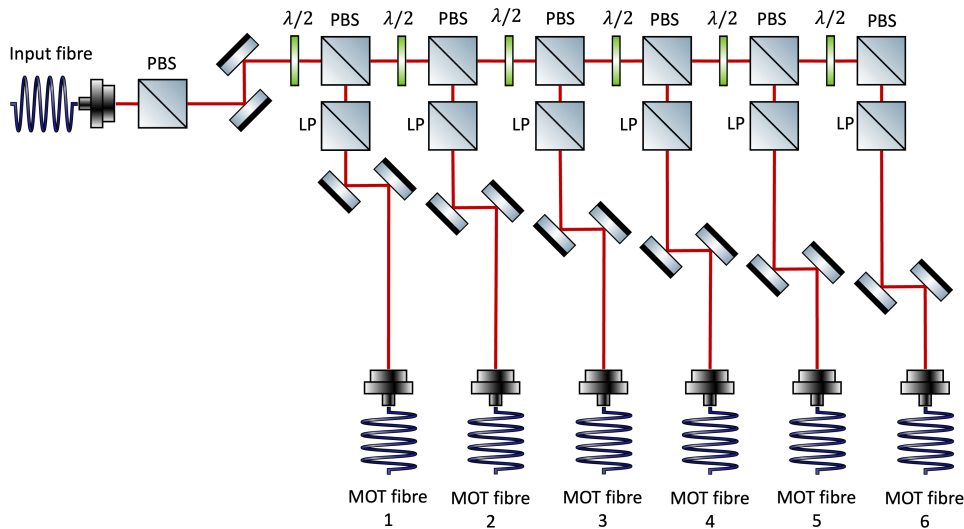


Figure 4.9: Optical layout schematic of the six-beam MOT distribution board.

plet lens which results in a beam diameter of ≈ 1 cm. Previously, singlet lenses were used but they introduced spherical aberration which caused circular fringes in the transverse intensity profiles when the beams were propagated through a few tens of cm which is the typical distance between the periscopes and the atomic cloud in the MOT. After collimation, the beam passes through a liquid crystal variable retarder and a PBS, which modifies the intensity of the transmitted light. After the PBS, a beam sampler is used to reflect 5 percent of the optical power towards a photodetector, allowing for feedback control of the beam powers. A pair of elliptical mirrors in kinematic mounts ensures optimal alignment of the laser beam. Finally, a QWP modifies the polarisation state of the light to circular with the appropriate chirality for the MOT. The two vertical MOT beams use conceptually similar setups.




Figure 4.10: Beam launching periscope setup used for the four horizontal MOT beams. A conceptually similar setup is used for the two vertical MOT beams.

4.3.1 Liquid crystal variable waveplates

Instead of scientific-grade LCD waveplates, we use modified inexpensive (£3 per piece) LCD light valves. These devices are widely available in robotics hobbyist stores. They consist of an LCD cell sandwiched between two film polarisers with aligned axes. Upon the application of a voltage, the LCD cell acts as a half-wave retarder with a diagonal fast axis such that the polarisation of the light incident on the output polariser is at right angles to its transmission axis. The film polarisers are easy to peel off and the glue residue is removed with a solvent. The remaining LCD cell is used as a variable retarder. We encountered three drawbacks with these cost-effective LCD devices. Firstly, the glass plates of the cell are uncoated leading to a total Fresnel loss of around 20 percent at 780 nm. This is not an issue

in our application because of the available optical power. However, the setup had to be modified as the LCD waveplates were initially mounted on the six-beam MOT distribution board. Since they were in series, the cumulative Fresnel loss was very large. Secondly, these devices come with minimal and partially wrong documentation. In particular, the manufacturer advertised that these devices require a DC driving voltage. According to measurements, the device does respond to DC voltages but the retardance decays to zero over a few minutes. This can be fixed by driving the LCD cell with an AC voltage such that the retardance depends on the RMS value of the driving voltage. Due to the digital nature of the PID controller developed, changing the driving signal from DC to AC was possible by software modifications only. Finally, with the currently used driving method, it is not possible to achieve half-wave retardance which results in a low extinction ratio. It is possible that by applying a higher RMS voltage and a sinusoidal drive would improve this (commercial LCD waveplates typically require 0 to 25 VRMS to cover the 0 to half wave retardation range), but our existing driver does not allow such high output voltages. In our application, high extinction ratios are not required since a single mechanical shutter is used to turn off the MOT beams simultaneously, before they are split six-ways.

4.3.2 Photodetectors

The optical power in each of the 6 MOT beams is detected by home-built amplified photodetectors. The assembly and the schematic drawing of the circuit are shown in Figure 4.11. The detector is based on a reverse biased Si PIN photodiode (Hamamatsu S1223). The photocurrent is converted to voltage with a resistor, the value of which can be changed with a set of 4 DIP switches. By changing the termination resistor, the gain of the photodetector can be increased at the expense of reduced bandwidth. The photodiode voltage is then filtered with a passive

low-pass RC filter and amplified by an operational amplifier in the single-supply, non-inverting configuration. The feedback resistor of the amplifier is set by 4 DIP switches similar to the photodiode termination resistor. The output of the amplifier is filtered with a low-pass RC filter and it is available through a BNC connector. The circuit was designed such that it is easy to modify if a different application is desired. The circuit can be powered externally or by a battery. In case of battery operation, a micro power, low bandwidth operational amplifier such as the OPA333 (350 kHz gain bandwidth product) can be used. When an external supply is used, a higher bandwidth op-amp such as the OPA211 (45 MHz gain bandwidth product) can be utilised. The PCB is compatible with both the TO-5 and TO-18 photodiode packages which are among the most common.

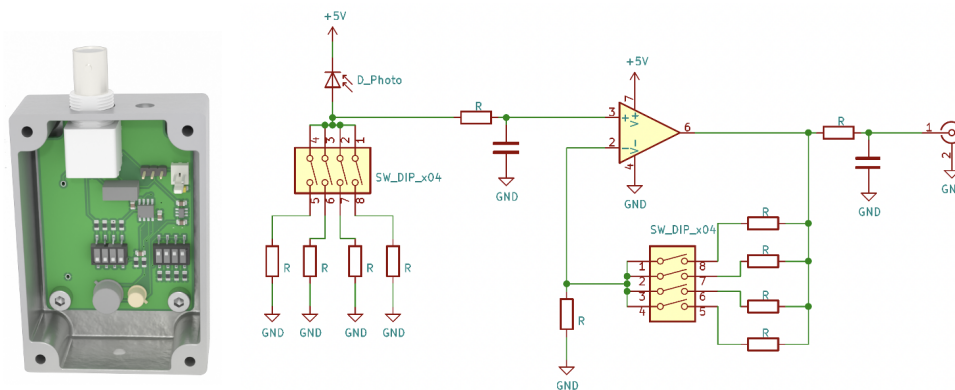


Figure 4.11: Amplified photodetector and its schematic.

4.4 PID controllers and LCD drivers

Active PID control of the optical powers in the six MOT beams might seem excessive for most ultracold atom experiments but yet it offers several benefits. Firstly, the PID control removes the optical power drifts caused by slow, small drifts in the output polarisation of the optical fibres. Even though these power drifts are small (1-2 percent), the stability of the atom count in the MOT improved with the

implementation of active power stabilisation. Furthermore, stable optical powers are crucial for later stages of the experiment such as optical molasses. Secondly, a stable setup enhances the efficiency of machine learning algorithms. Slow drifts that affect the cost function can decelerate the convergence of machine learning optimisations. Thirdly, active feedback control permits dynamic adjustments of the optical powers, enabling the manipulation of the atomic cloud's shape in the MOT. This can be exploited when loading the atoms into an RF-dressed trap which is highly anisotropic and loading without changing the shape of the cloud would cause excessive heating and atom loss. Lastly, the high degree of remote control offers advantages. During the COVID-19 pandemic, it was possible to remotely operate the previous generation Oxford experiment due to a similar design. Extending remote control to the MOT beam powers results in further reduction of on-site maintenance.

4.4.1 PID controller hardware

Rather than designing an analog PID controller for active stabilisation of the optical powers, we adopted a microcontroller (MCU) based digital approach. This has various advantages over analog controllers. Firstly, the parameters of the control loop can be software-tuned, enabling integration with the existing control system. On analog variants, this would be typically done with changing resistors or turning trim pots, complicating the implementation of remote control. Secondly, digital circuits permit substantial signal chain modifications solely by software updates. As previously noted, the LCD waveplates were initially driven by a DC voltage following the manufacturer's guidance, which turned out to be incorrect. It was possible to reconfigure the outputs to be pulse-width modulated signals, which improved the performance of the LCD waveplates. In case of analog circuitry, this would have required adding new components and the development of a

new PCB, resulting in a lengthy process. Finally, with digital electronics, it is easy to compensate for the non-linear response of the actuator by an interpolated lookup table, enhancing the performance of the control loop. The hardware was designed such that it fits in a standard 3U tall, 8 HP wide plug-in unit allowing for 22 PID channels to occupy a single 19-inch subrack. The total cost of a dual-channel unit is £150 which is far cheaper than commercial PID controllers. A single unit is shown in Figure 4.12. Details of the firmware and control interface are described in Appendix B.




Figure 4.12: PID controller front panel and hardware. The controllers are used to regulate the optical power in each of the six beams of the glass cell MOT.

4.4.1.1 The microcontroller

The choice of the microcontroller was motivated by an existing open-source design known as the Pyboard [86], which is the flagship development board for the MicroPython project [87]. The MCU used belongs to the STM32 F4⁵ family and

⁵STM32F405RGT6

features a 168 MHz microprocessor, 192 kB RAM and peripherals required for this project such as the SPI bus or data converters. While MicroPython is not commonly used in industry, it is widely used among enthusiasts because of the simplicity of the Python programming language.

4.4.1.2 The input signal chain

The analog input voltage of the PIDs is processed through a signal chain which consist of a series of operational amplifiers that map the range of input voltages (-10 V to 10 V) to the range of the Analog to Digital Converter (ADC) of the MCU (0 V to 3.3 V). Firstly, a unity-gain instrumentation amplifier is used to subtract an analog modulation voltage from the input voltage. This modulation input can be used for fast control of the setpoint of the controller. Two inverting amplifiers in series amplify and offset the voltage. Finally, another set of two operational amplifiers in series clamp the output voltage using diodes to ensure that it remains within the ADC's specified range.

4.4.1.3 The output signal chain

The output signal is derived from a Pulse-Width Modulation (PWM) capable output of the MCU. On this output, pulses of variable duty-cycle (16 bit resolution) and repetition rate can be programmed. The output signal is amplified and offset by a series of 2 inverting operational amplifiers and sent to the output BNC connector. PWM signals are not ideally suited for controlling LCD devices as they have a non-zero mean which can degrade LCD cells through electrolysis. However, we found no evidence of degradation and the output power varied monotonously with the duty cycle of the control signal.

A second monitor port is also available which can be configured to output auxiliary signals such as the error signal. Alternatively, this output could be used to

drive a second actuator with a different gain and bandwidth. Due to the limited number of built-in DACs in the MCU, this output is controlled by a separate DAC. This DAC communicates with the MCU through an I2C bus offering a 12-bit resolution and 25 kHz update rate. The output of this DAC is processed by a circuit identical to the main output signal chain.

The PID controller is also equipped with four indicator LEDs, which display the status of the controller, the status of the control loop as well as input and output saturation. Light from the indicator LEDs is channeled to the front panel by light pipes. These LEDs are useful tools for monitoring the status of the PID controllers on-site.

4.4.2 Performance evaluation

The performance of the PID controllers was assessed using two different metrics. Firstly, the stability of the optical power controlled by the PID controllers was measured and compared to readings of an optical power meter. It was found that the optical power stayed constant within measurement errors (the least significant digit of a three-digit optical power meter stayed constant over half an hour). When compared to unregulated operation, on average a one percent optical power drift was detected due to slow drifts in the polarisation. Secondly, the AC characteristics of the controller was measured. We introduced a small sinusoidal setpoint modulation and the behaviour of the control loop was compared at different modulation frequencies. The control was considered optimal when the optical power followed the modulation such that the magnitude of the error signal relative to the input stayed under one percent. Following fine-tuning of the PID gains, optimal performance was sustained up to a modulation frequency of 22 Hz. Consequently, the characteristic time the optical powers could be ramped from one value to another is about 15 ms, a quarter of the above modulation period. This response time

is typical of LCD waveplates, indicating that the bandwidth of the PID controller is higher.

4.5 The glass cell MOT setup

The glass cell MOT setup is shown in Figure 4.13. The bottom vertical beam launcher and the four horizontal beam launchers are mounted on the main breadboard. The top vertical launcher is mounted on a separate, elevated breadboard which is supported by four 1.5" posts for stability. The quadrupole magnetic field is produced by a pair of water-cooled anti-Helmholtz coils (upper coil is not shown for illustration purposes). The coils are driven by a Delta Elektronika SM15-400 high-current power supply which is capable of outputting currents up to 400 A. Due to the fast response time of the supply (≈ 1 ms) we control the output current directly with an analog voltage. The coils are glued into 3D-printed coil formers made of the mechanically robust PA12 nylon material, which are attached to four vertical M12 threaded rods. These coil formers also serve as mounts the PCB-based RF and bias coils, providing a flat surface. The entire coil assembly is rotated approximately 10 degrees about the vertical axis. This is necessary to ensure that the MOT beams are mutually perpendicular and intersect at the node of the quadrupole magnetic field. Since the only coils carrying high-currents are the anti-Helmholtz coils, there is no requirement to use multiple high-current power supplies or a MOSFET bank to manage the current distribution as in the previous generation Oxford experiment.

In addition to the MOT beam launching periscopes, the glass cell MOT apparatus also includes the absorption imaging setup and the imaging repumper beam. The absorption imaging setup consists of a beam launching periscope similar to the MOT beam launchers but it does not include a beam sampler. The imaging beam




Figure 4.13: The glass cell MOT setup. The elevated breadboard and the downwards propagating MOT beam is not shown for clarity.

is aligned to pass through the location of the trapped atoms. The atoms are imaged with a simple one-to-one telescope which consists of two 100 mm achromatic doublet lenses of 2-inch diameter. The image is captured by a Thorlabs CS505MU CMOS camera. In the future, the objective achromat will be replaced by a custom-made four lens objective which corrects for the aberrations introduced by the glass cell walls as well as increasing the magnification of the imaging system. This update is necessary when imaging a matter-wave interference pattern as important features are on the length scale of a few μm .

4.6 The RF signal chain

At the final stage of each experimental sequence, the atoms are confined in an RF-dressed trap. The parameters of the potential are highly influenced by the frequency and amplitude of the applied EM fields therefore precise control over these parameters is essential. Moreover, when a double-well trap is created by

applying three radio frequencies, the potential is especially sensitive to the relative phases and amplitudes of the three fields. To achieve low heating rates and good repeatability, the stability and high degree of control over these EM fields are very important.

4.6.1 The RF signal generation and amplification

The RF signals are derived from an AD9910-based four-channel DDS synthesiser called Urukul which is a part of the ARTIQ ecosystem. This device can produce signals with frequencies in the range of 1 MHz to 400 MHz with 0.25 Hz resolution and it features a 14-bit output DAC. We use all four channels to synthesise the three radio-frequencies for the double-well trap as well as the signal for evaporative cooling. The RF signals are combined with an RF power splitter (MiniCircuits ZMSC-4-1+), then passed through a high-pass filter (MiniCircuits ZFHP-1R2-S+, 1.2 MHz-800 MHz passband) and subsequently amplified by a MiniCircuits LZY22+ 30W RF amplifier. The output of the amplifier is connected to a PCB based, resonantly matched RF coil which generates the RF magnetic fields. The coil is a two-port device; its output is terminated with a small resistor which serves as a pick-up to measure the RF current in the coil.

4.6.2 PCB-based RF coils

Experiments employing magnetic trapping techniques typically use macroscopic, wire-wound coils [88] or atom chips [89] to generate the required magnetic fields. The previous generation RF-dressed experiment in Oxford uses the first approach with the significant benefit of the smoothness of the generated potential, since the atoms are located far (about 1-1.5 cm) from the coils. However, the turn-to-turn parasitic capacitance as well as other capacitive couplings between the coil and its environment can lead to RF amplitude drifts which result in the need for

frequent calibrations through trap spectroscopy. The atom-chip approach exploits the fact that the magnetic field strength generated by a current carrying wire scales inversely with distance. By placing a wire close to the atoms, the required current is lower. Atom chips can also be fabricated to a high precision to realise various magnetic field geometries. However, the proximity of atoms to the chip can lead to magnetic field noise and stray fields, which causes heating and imperfections in the trap [90].

We decided to adopt a third method for RF magnetic field generation which involves macroscopic PCB-based coils that slide under the science chamber. The atoms are located similarly far as with the wire-wound coil approach, leading to much less magnetic noise compared to atom chips. The coil geometry is more rigid due to the laminated copper traces on the PCB substrate, therefore we expect the parasitic couplings to be more stable. The PCB coils can be easily exchanged which allows for changing the number of turns or the resonant frequency of the coils. This is in contrast with the old experiment, the coils were wound around the UHV glass cell with tight tolerances. Replacing the RF coils would risk damaging the UHV glass cell and therefore this task was never performed. Finally, the matching components are mounted on the PCB itself, which reduces the number of RF connectors and coaxial cable length which is expected to further enhance passive stability.

The geometry of the PCB coils was inspired by the magnetic field created by two bar magnets with vertical axes but opposite polarity. Above and between the opposite poles, the magnetic field is horizontal, which is the required direction for the dressed quadrupole trap. The same magnetic field geometry can be realised by two wire loops next to each other one with clockwise and one with counterclockwise currents. The coil geometry was optimised using Magpylib, an open-source magnetic field solver package [91]. The PCB traces were modelled

as one-dimensional current lines and the entire geometry was parametrised with five variables, the height and width of the outermost loop, the separation between coils, trace separations, and offset relative to the geometrical centre of the coils. We sampled the magnetic field at 1000 points 12 mm above the plane of the coil within a cuboid region comparable in size to the atomic cloud in the RF-dressed trap. The standard deviation along the desired direction, relative to the mean, served as a cost function for an optimisation algorithm to determine the best coil geometry. This process was repeated for different numbers of turns. After evaluating the strength of the magnetic field produced at a current of 1 A (approximate peak output current of the RF amplifier used considering a 50Ω load), we found that four- and five-turn designs are the best. In Figure 4.14 a, we show the simulated magnetic field lines for a prototype coil, plotted along an x - z plane bisecting the coils along the y -axis. The PCB implementation of this coil geometry is shown in Figure 4.14 b, showing the matching components and RF connectors. Note that a similar coil geometry, rotated 90 degrees, is present at the bottom copper layer of the coil which creates a magnetic field along the y direction.




Figure 4.14: **a**: Numerical simulation of the magnetic field of the proposed geometry for PCB RF coils. **b**: Implementation of the PCB RF coils.

To characterise the RF properties of the PCB coils, a vector network analyser (Omicron Lab Bode 100) was used. Firstly, the impedance of each coil was measured across a range of frequencies. During these measurements, the second port of the coil under test was terminated at 50 Ohms while the other coil was left unterminated. By analysing the imaginary part of the impedance, we were able to deduce the inductance and stray capacitance of the coils using a simple lumped element model of a resistor in series with an inductor (whose stray capacitance is modelled with a parallel capacitor). For the x coil, the values are $3.51 \mu\text{H}$ and 12.6 pF , while for the y coil the values are $3.42 \mu\text{H}$ and 12.9 pF . The self-resonant frequency of the coils is therefore 24 MHz. The impedance measurement results are shown in Figure 4.15 a/b. These findings demonstrate the high manufacturing precision of the PCB industry and highlight the applicability of PCB-based coils in future commercial applications of RF-dressed atomic systems.

To measure the cross-coupling between the coils, the x coil was driven by the VNA's output port while the y coil was connected to the input port. The second port of both coils were terminated at 50 Ohms. The cross-coupling as a function of frequency is shown in Figure 4.15 c. At the typical dressing frequency of 4.2 MHz, the cross-coupling is -30 dB which corresponds to an amplitude ratio of 31.6 which demonstrates sufficient isolation between the two coils.

4.6.3 RF amplitude and phase measurement

The amplitudes and phases of the RF signals influence strongly the shape of the RF dressed trap, making their precise measurement essential. High accuracy calibration of RF amplitudes is possible through RF trap spectroscopy [92] but this process takes significantly longer than a simple amplitude measurement. The previous generation Oxford experiment uses pickup coils co-wound with the RF coils which were used to measure the RF amplitudes and relative phases of MRF sig-

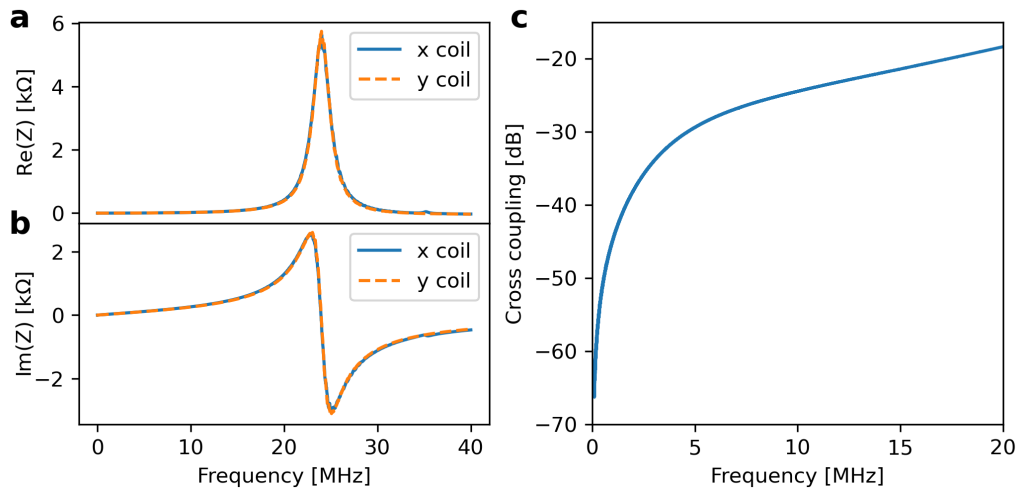


Figure 4.15: Measurement of the real (a) and imaginary (b) part of PCB coil impedance as well as cross-coupling between the coils (c).

nals. As these coils are inductively coupled, the voltage measured across the pickup coil is proportional to the strength of the oscillating magnetic field and it was recorded with an oscilloscope. The amplitudes and phases were then extracted with an FFT-based algorithm. This measurement method has two disadvantages. Firstly, it was found that just as the parasitic capacitive coupling between the RF coils and their environment caused slow amplitude and relative phase drifts, the inductive coupling between the pickup coil and the RF coils drifted as well. Only RF trap spectroscopy gave a result accurate enough to re-calibrate the RF amplitudes and phases accurate enough to perform splitting of Bose gases. Secondly, the saved waveforms significantly increased the total amount of generated data by the experiment, even though in most cases, the analysed data consists of 3 amplitudes and 2 relative phases.

4.6.3.1 The resistive pickups

To improve the RF amplitude and phase measurement a new scheme was designed. The first improvement is using a small current sense termination resistor

as a pickup. The voltage across the pickup resistor is directly related to the current flowing in the circuit, instead of being inductively coupled. The main factor which limits the performance is the temperature coefficient of the resistor, which is typically 50 ppm/K. To improve the accuracy, a Peltier element based analog temperature controller was designed. The temperature stability was of the order of 50 mK which means that the theoretical precision of the pickup resistor is 2.5 ppm. The voltage across the pickup resistor can be accessed directly or through a unity-gain precision instrumentation amplifier (AD8429) which acts as a buffer.

4.6.3.2 Processing of pickup signals

The output signal of the pickup circuit can be further processed three different ways. Firstly, the signal can be acquired by a digital oscilloscope and can be analysed similar to the case of the previous generation Oxford experiment. Secondly, a linear RF detector IC (ADL5511) can be used to measure the RF power of the signal. This IC outputs a voltage proportional to the RMS input voltage as well as the envelope of the measured signal on two different pins. While this technique is suitable for measuring the amplitude of a single-frequency RF signal, it is not useful when three radio frequencies are used. Thirdly, a lock-in amplifier was developed which can measure the amplitudes and relative phases of the three RF signals simultaneously. The lock-in amplifier design is based on the AD8333 I/Q demodulator IC. The input pickup signal is split three ways by a resistive 50 Ohm splitter made of low temperature coefficient resistors (10 ppm/K) for high stability. The three signals are converted from single-ended to differential signals with a 2.5 V common mode voltage as per the recommendation of the demodulator datasheet using the AD8139 low-noise differential amplifier. Three additional RF inputs are used to provide the local oscillator signals at four times the desired frequencies. The demodulator IC generates the in-phase and in-quadrature local

oscillator signals by a divide-by-four logic circuit and multiplies them with the input signal. The output signals are available as currents proportional to the I and Q demodulated signals. These currents are converted to voltage by transimpedance amplifiers. In case of 3 radio frequencies, the zero frequency components contain the desired amplitudes but the signal also contains components at all the possible beat frequencies. To filter these, a 3-stage, 6th-order Sallen-Key filter was designed with a 20 kHz bandwidth. The bandwidth was chosen considering the typically used minimum RF spacing of 40 kHz. The filtered analog signals are subsequently converted to digital signals by a 16-bit ADC (ADAS3023) which sends the I and Q amplitudes to a microprocessor to extract the amplitudes and phases trigonometrically. The measured values can then be used to monitor the RF amplitudes and phases.

4.7 The control system

The control system of the experiment is based on the Sinara hardware family⁶ running the ARTIQ control software which was developed specifically for quantum physics experiments [93]. This ecosystem has various advantages to alternatives, for instance, open-source hardware and software, nanosecond temporal resolution and the ability to use the Python programming language for compiling and executing experimental sequences.

The core device of the hardware is the Kasli FPGA Carrier module. This module is responsible for executing various input-output operations accurately to nanosecond timescales and it communicates via ethernet with a host computer which compiles and dispatches tasks. The core device has 12 extension buses for connecting

⁶In the following, the equipment will be referred to as ARTIQ hardware as often used by the community.

various external modules. Simple extension modules such as a Digital Input Output module serves as a breakout for FPGA pins. More complex modules such as the Zotino 32-channel ADC or the Urukul 4-channel DDS synthesiser are interfaced via a serial interface to the core device, utilising Low Voltage Differential Signalling. While the ARTIQ ecosystem is readily available, actual implementations vary significantly. This section describes the implementation of the ARTIQ-based control system of the new Oxford experiment.

4.7.1 Example of ARTIQ sequences

To demonstrate how the control system works, consider a simple sequence: generating ten pulses on a digital output, each lasting 2 microseconds with an 8 microsecond gap between them. On the host PC, the following function is used to compile the sequence.

```
@kernel
def test():
    for i in range(10):
        ttl.on()
        delay(2*us)
        ttl.off()
        delay(8*us)
```

The `@kernel` decorator tells the host PC that this code is to be executed by the core device in real-time. The host PC's interpreter goes through the code line-by-line and submits Real Time Input Output (RTIO) events to the high-precision RTIO timeline such as turning on or off a TTL channel. The delay statements advance the timeline by the specified amount. It is very important for the time of all RTIO events to be always ahead of the actual RTIO time, otherwise causality would be violated. In this case, an RTIO underflow error is raised. If there is no such issue, the programmed events are executed to a nanosecond temporal

precision. This all or nothing accuracy is an intriguing feature of the control system.

4.7.2 The apparatus module

At the lowest level, the control system is interfaced to various parts of the experimental hardware such as power supplies through analog or digital input and output channels. To directly control these elements of the experimental hardware, an apparatus module was developed. In the following, an example for a DDS synthesizer channel which controls the detuning of the glass cell MOT cooling light is introduced.

```
class CellCoolingDetuning(EnvExperiment, Apparatus):

    def build(self):
        self.setattr_device('cell_cooling_detuning_aom')
        self.setattr_device('core')

    @kernel
    def set_detuning(self, detuning, amplitude):
        freq = (102.5-0.5*detuning)*MHz
        self.cell_cooling_detuning_aom.set(freq, amplitude)

    @kernel
    def ramp_detuning(self, detuning1, detuning2, dt_us, N):
        ddet = (detuning2-detuning1)/N
        for i in range(int(N)):
            self.set_detuning(detuning=detuning1+(i+1)*ddet)
            delay(dt_us*us)

    @kernel
    def run(self):
        self.core.reset()
        delay(0.1*ms)
        self.ramp_detuning(-16.0, -30.0, 5000, 1000)
        delay(0.1*s)
        self.core.wait_until_mu(now_mu())
```

For each distinct hardware component, a separate class is created, in this case,

`CellCoolingDetuning` which controls the detuning of the glass cell MOT light using an AOM. This class inherits from the `EnvExperiment` class which is the built-in base class for the entire experimental environment and the `Apparatus` class which signals that the current code describes the apparatus and not a sequence. The first, `build` method initialises the devices that will be controlled as class attributes. In this case, `cell_cooling_detuning_aom` is a memorable alias for the given DDS output channel and `core` is the core device which runs the experiment. After this, 2 kernel methods are defined to be called by sequences. For example, `set_detuning` changes the detuning of the glass cell MOT cooling beams. Note that this function converts between the desired detuning and the output frequency of the DDS channel. This way, the control parameters are more meaningful and physically relevant. The `ramp_detuning` function performs a linear ramp of the detuning between two values, `detuning1` and `detuning2` in `N` steps, separated by `dt_us` microseconds of time. Finally, the `run` method performs the detuning ramp. This method is only for debugging and test purposes and it is executed with the ARTIQ dashboard GUI.

4.7.3 The sequence module

The sequence module contains code that generates experimental sequences for individual runs, one- or multi-dimensional parameter sweeps, and machine learning optimisations. Given the complexity of typical sequences, only a brief overview of a typical sequence module is provided here. This example involves loading the glass cell MOT and capturing a fluorescence image of the atomic cloud. Every sequence module contains a worker class with a name starting with an underscore as shown below.

```
class _LoadCellMOT(EnvExperiment):
```

This worker class contains all necessary functions to perform a single experiment, therefore it is easily reconfigurable to run sweeps or machine learning optimisations. Every worker class contains a `build` method which initialises all the apparatus and subsequence objects for the sequence, defines all parameters, and performs all the conversions to machine units if necessary.

```
def build(self) :
    #Build apparatus
    self.CellBias = CellBias(self)
    self.CellCamera = CellCamera(self)
    self.PyramidShutter = PyramidShutter(self)
    #Define parameters
    self._p_pyramid_quad_current = 24.3
    self._p_pyramid_ud_bias = -0.7
    self._p_cell_mot_y_bias = 0.5
    self._p_cell_mot_cooling_detuning = -26.0
```

When another sequence is inherited as a subsequence (for instance, more complicated sequences such as loading the shell trap would inherit the `LoadCellMOT` sequence), there is no need to redefine the parameters or apparatus classes. All parameters must be prefixed with `_p_`. This helps the control interface to identify parameters and distinguish them from other variables. The worker class also includes kernel methods which execute RTIO events in small logical blocks. For instance, the function below compiles instructions which empty the glass cell MOT.

```
@kernel
def empty_cell_mot(self) :
    self.PyramidShutter.close()
    self.CellShutter.close()
    delay(2*ms) #Shutter activation time
    self.CellQuad.set_current(0.0)
    delay(100*ms)
```

A combination of these methods is used to create a main kernel method called

`kernel_jobs` that compiles all the events for execution on the high-precision timeline.

```
@kernel
def kernel_jobs(self):
    self.core.reset()
    delay(1*ms)
    self.empty_cell_mot()
    self.load_cell_mot()
    delay(self._p_t_first_img*ms)
    self.CellCamera.trigger_images()
    self.core.wait_until_mu(now_mu())
```

This function, by default, starts with resetting the core (hence the high precision timeline), followed by a small delay to ensure that events are always submitted for execution in the future. After this, methods such as the previously described `empty_cell_mot` are executed. Finally, the core device waits until all the events have been completed. Certain components of the experimental apparatus are not interfaced to the ARTIQ hardware directly, but communicate with the host PC or the core device via Ethernet. An example for this is the camera: while triggering images is an RTIO event with precise timing, other camera-related functions must be performed through its USB interface. These include arming the camera or downloading the images from the camera memory to the PC. Therefore, a function `do_single` is defined which executes these non-kernel methods before or after the `kernel_jobs` method, that requires a continuous, high-precision timeline.

```
def do_single(self, savepath, run_num):
    self.CellCamera.glass_cell_camera.arm()
    time.sleep(0.25) #Let the PIDs lock
    self.kernel_jobs()
    time.sleep(0.5)
    self.CellCamera.save_imgs(savepath, run_num)
```

This function has 2 important input arguments which describe the file path for saving images and the run number which is used to identify individual sequences in a parameter sweep. To transform the worker class `_LoadCellMOT` into normal class which executes an experiment, the following code is needed in the same file.

```
class LoadCellMOT(EnvExperiment):

    def build(self):
        build_from_worker(self, _LoadCellMOT(self))

    def run(self):
        run_from_worker(self)
```

The presence of the `run` method indicates that this experiment can be executed by ARTIQ unlike the worker class. The functions `build_from_worker` and `run_from_worker` are important utility functions of the control interface which is described later in detail. They expose all parameters to be modified by the control interface as well as compile a set of parameter lists in case a parameter sweep is executed.

4.7.4 The control interface

The default ARTIQ control interface is the ARTIQ Dashboard, shown in Figure 4.16. This tool has the ability to run experiments (provided that their class inherits

from `EnvExperiment` and they have a `run` method, gives manual control over the outputs, and provides widgets for data visualisation. While the dashboard is a useful tool for testing outputs, it can become overcrowded with more complicated experiments involving many parameters. A simple experiment, loading the glass cell MOT, requires 24 control parameters which must be set individually on the dashboard. Furthermore, if sweeps are needed for some of these parameters, the code of the experiment would have to be modified such that a scannable input variable appears on the dashboard. Alternatively, by making all parameters scannable (with the option of a single value), the dashboard interface would become even more complicated. In practice, most of the control parameters are left at the default value and only one or two are changed or scanned. Therefore, it is an unnecessary complication to display all the parameters that can be changed or scanned.

To overcome the limitations of the ARTIQ dashboard, a custom, code-editor based control interface was developed. This interface compiles a Python dictionary with all the parameters necessary to run the experiment and communicates with the host PC over an RPC server. The control interface is introduced with a simple example which involves a 2D parameter sweep.

```
E = Experiment('LoadCellMOT')
E.p_cell_mot_bias_x = 1.23
E.p_cell_mot_bias_y = np.linspace(-2, 2, 10)
E.p_cell_mot_bias_z = np.linspace(-2, 2, 10)
E.run()
```

Initially, a `LoadCellMOT` experiment is created. Then, the value of the control parameter `p_cell_mot_bias_x` is set to a non-default value. Then, two other control parameters `p_cell_mot_bias_y` and `p_cell_mot_bias_z` are set to a linearly spaced array of 10 elements between -2 and 2, indicating a

two-dimensional parameter sweep. Finally, the experiment is submitted to the host PC.

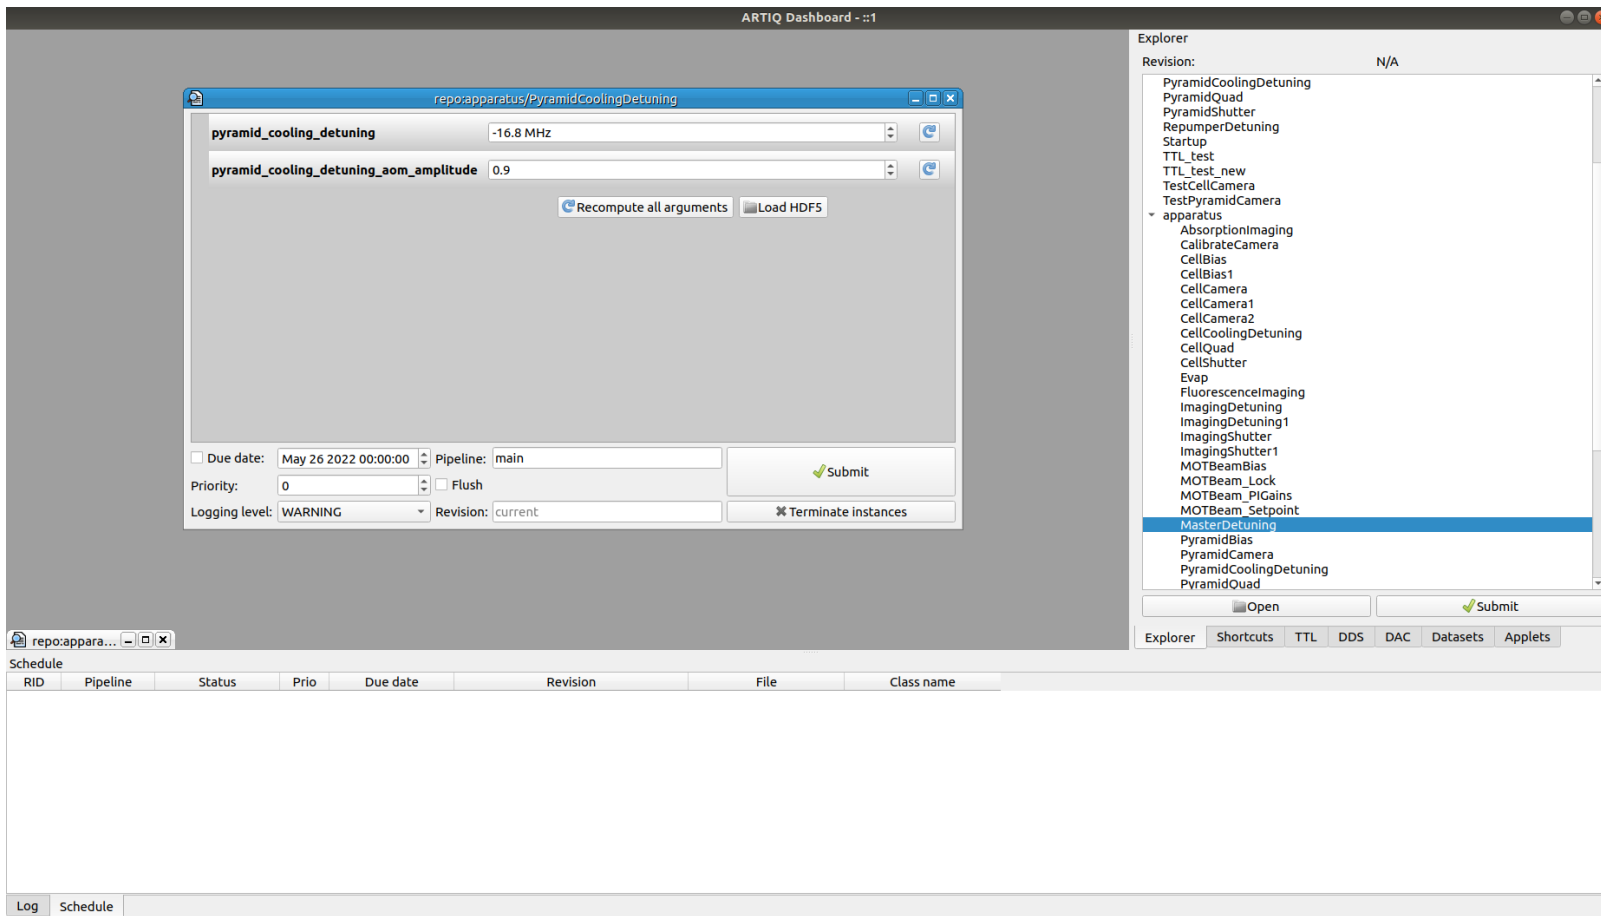


Figure 4.16: ARTIQ dashboard, displaying a simple experiment where the detuning and power of the pyramid MOT cooling light can be changed.

If a machine learning optimisation is required over the same two parameters instead of a two-dimensional parameter sweep, the following code is used.

```
E = Experiment('LoadCellMOT')
E.ML = True
E.p_cell_mot_bias_x = 1.23
E.p_cell_mot_bias_y_ml = [-2, 2]
E.p_cell_mot_bias_z_ml = [-2, 2]
E.run()
```

By appending `_ml` to the parameter names and providing a list of 2 values, the control interface interprets these as the lower and upper bound of the optimisation. The absence of `_ml` in `p_cell_mot_bias_x` indicates the optimisation is performed at the non-default value for this parameter. Omitting this line would result in the use of the default value.

4.7.5 Machine learning optimisation

One of the primary applications of machine learning is finding extrema of black box functions with numerous input variables, therefore, it is well-suited for optimising quantum gas experiments as demonstrated in [78]. Since the optimiser has no a priori knowledge of the experiment, it can provide a less biased result compared to a human optimiser. Moreover, it can efficiently discover correlations in high dimensional spaces, while manual optimisation is usually limited to 2D parameter scans.

A typical implementation of an ML optimiser takes as input a function of n parameters as well as the corresponding bounds. Therefore, the previously introduced worker classes can be easily reconfigured to perform machine learning optimisations as follows.

```
class MLLoadShell(EnvExperiment, ML):

    def build(self):
        build_from_worker_ml(self, _LoadShell(self))
        self.initialise_ml(self.ml_param_names, 1)

    def errorfunc(self, params):
        respath = run_single_ml(self, params)
        atomN = self.analyse()
        self.update_ml(atomN, params, None)
        return -atomN

    def run(self):
        optimize(self.errorfunc, bounds=self.bnds)
        self.finaliseML()
```

The `MLLoadShell` class inherits from the experimental environment as well as from `ML` which is a class containing machine learning helper routines, for instance parameter saving at every iteration of the optimisation or visualising the learner's progress. The `build` method, similar to the previous example, calls a utility function which builds the worker class specific to the needs of the machine learner and it initialises the machine learning helper routines. The `errorfunc` function serves as the cost function of the optimiser. It takes as inputs a list of parameters which are the variables of the optimisation. The function then executes an experiment with these parameters. The result of the experiment is analysed, examples of analysis include the MOT fluorescence captured by a photodiode or the number of atoms in a region of interest from an image. The `run` is the highest-level function which performs the optimisation and saves the result.

We tested various open-source optimisers. The first attempt involved a Bayesian optimiser [94] which performed poorly in general. The algorithm was fitting to the noise of the experiment and therefore could not converge to the global maximum. Moreover, the optimiser often failed during execution due to the lack of proper

exception handling in the code. There exists an option to tune the optimiser to increase susceptibility to noise in the cost function but this did not improve the results.

To ensure that the experimental uncertainties do not affect the optimisation the SciPy [95] implementation of the differential evolution algorithm [96] was tested. The choice for this algorithm was motivated by its gradient-independent approach in finding the optimal parameters. The algorithm randomly samples the parameter space during the first iteration and then uses the results to sample in the neighbourhood of the best parameters. While this algorithm did succeed in finding better experimental parameters, its convergence was slow.

The most effective optimisation package found was scikit-optimize [97] which is based on the widely used machine learning package scikit-learn [98]. We use the `gp_minimize` function of this package which performs Bayesian optimisation using a Gaussian process. This algorithm tolerates the noise present in the experiment and converges to the optimal parameters significantly faster than differential evolution. Furthermore, this package contains useful utility functions which visualise the optimisation result. This helps in identifying the parameters to which the cost function is particularly sensitive.

An example of a machine learning optimisation is shown in Figures 4.17 and 4.18. In this instance, the cost function being maximised is the atom number in the shell trap shortly after loading, before the evaporative cooling procedure starts. The parameters of the optimisation are the currents in the pyramid MOT quadrupole and bias coils and the pyramid MOT cooling detuning. These parameters influence the characteristics of the atomic flux from the pyramid source and therefore the number of atoms trapped in the glass cell MOT. Figure 4.17 displays the evolution of the parameters and the value of the cost function. For the first 64 runs,

the hyperparameter space is sampled using the Sobol method [99] which provides 2^N low discrepancy quasi-random points. The number of initial points is typically chosen by a heuristic rule, about 10 samples per parameter rounded to the closest power of 2. Due to wide bounds on the initial parameters, only a few of these random samples resulted in a reasonable atom number. However, after the initial exploration phase, the optimiser was able to quickly converge to the best values. Variations in the cost function are likely to be caused by long term drifts in the apparatus. In Figure 4.18 one of the utility functions of scikit-optimize is used to visualise the optimisation result. For clarity, only three parameters, the three bias fields of the pyramid MOT source are presented. The diagonal plots show the effect of varying a single parameter around the optimum using the learnt objective function. The vertical red line shows the best observed point, which demonstrates the success of this optimisation run as it is close to the best estimated point. The off-diagonal plots display the same quantity but as a function of two parameters. The black dots show each function evaluation while the red star corresponds to the best observed set of parameters. The shape of the contours on the off-diagonal plots reveal the underlying symmetries of the experimental system. Since pyramid source emits atoms in the "EW" direction, the location of the node of the magnetic field along this direction is not as important as along the plane transverse to it (given by the "UD" and "NS" directions). The slight asymmetry in the UD-NS plot can be attributed to the bias coils being at different distances from the atoms in the 2D MOT. The advantages of the machine learning optimisations are further exemplified by a scenario where the experimental hardware was significantly modified by changing the quadrupole coils around the glass cell which required realigning the beams of the glass cell MOT. Despite trapping a normal number of atoms in the glass cell MOT, the transfer efficiency into the shell trap was very low. Four hours of manual optimisation with one dimensional sweeps

did not improve the loading efficiency significantly. However, a machine learning optimisation of the bias fields and detuning during the molasses stage with only 64 trials during the exploration phase improved the number of trapped atoms by a factor of 10. This example demonstrates the ability of machine learning algorithms to discover correlations in high-dimensional spaces and exploit them for optimisation.

4.7.6 Interfacing the PID controllers to ARTIQ

The only parts of hardware which were developed in-house and require a mixed analog/digital interface to the control system are the PID controllers which regulate the glass cell MOT beam powers. The setpoint of the PID controllers can be modulated through an analog channel, which is straightforward to implement. This is used during sequences when fast and high temporal precision is required. However, several other parameters, while important, do not require real-time updates. These include the PID gains, update rates or the setpoints (without external modulation). Such parameters are typically adjusted at the beginning of an experimental sequence before the high-precision RTIO timeline is restarted. This motivated the development of an interface which is similar to how the cameras are controlled through Ethernet. A Raspberry Pi 4B single-board computer was made rack-compatible and its SPI buses were broken out to 8 ribbon-cable connectors which are used to provide links to three PID controller units. This is showed in Figure 4.19. In the remaining part of this section, this device will be referred to as the interface PC. In total, 8 units and 16 channels can be controlled by a single interface PC, this is due to the limited number of SPI buses. The interface PC acts as a router between the connected PID controllers and reads from and writes into the PID controller registers. Furthermore, it hosts an RPC server powered by sipycos which is the primary way to interface network devices to the ARTIQ hard-

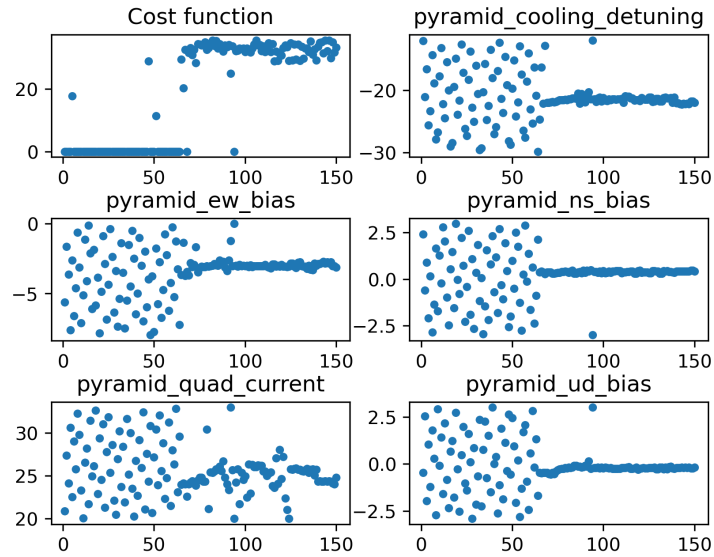


Figure 4.17: Evolution of cost function and parameters during an ML optimisation.

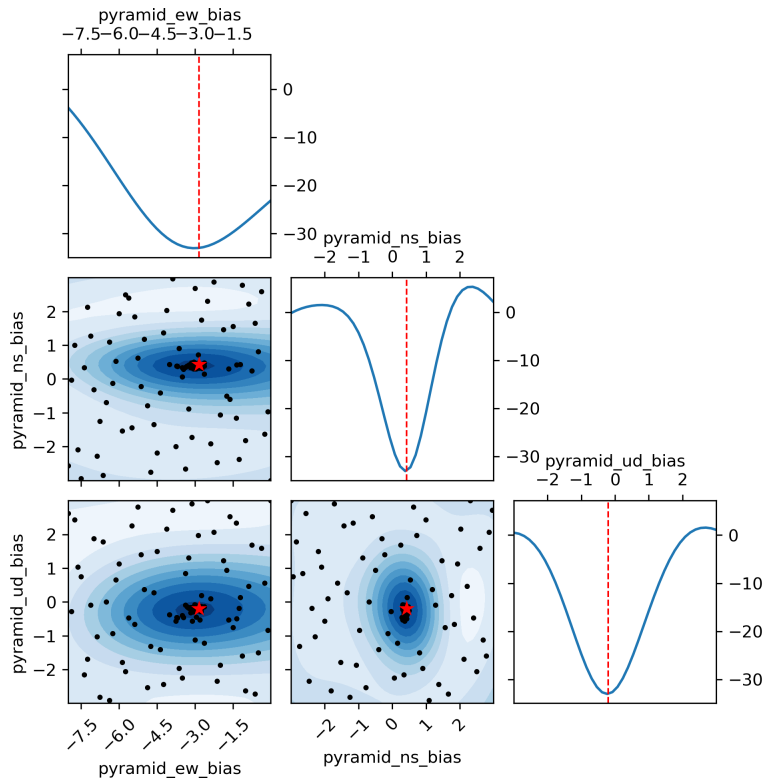


Figure 4.18: Mutual dependence plots of selected variables after ML optimisation.

ware. The control software running on the interface PC interprets the commands received through Ethernet and transmits the corresponding instructions over SPI to the PID controllers. Besides the control interface, a monitor GUI was developed which is used to visualise the input, output and error signal of the six PID controller channels.

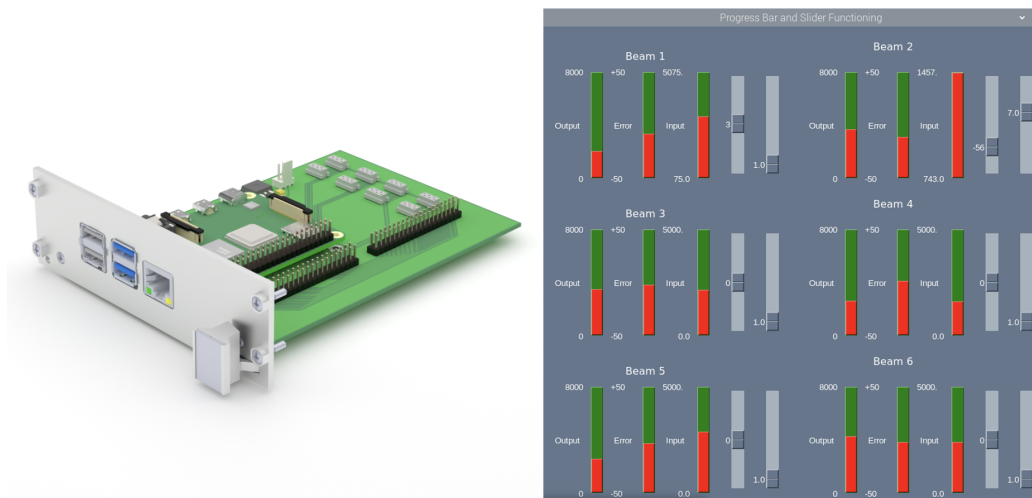


Figure 4.19: The Raspberry Pi based interface PC and monitor GUI for the PIDs.

4.8 Future upgrades of the new experiment

4.8.1 Laser system

While the current laser system performs well, it consists of numerous independent setups which are all connected by optical fibres. This results in a significant optical loss and the need for three ILDLs and two TAs. The future version of the setup will involve frequency tuning the cooling light of the pyramid source and the glass cell MOT together by serrodyne modulation [100] which offers fast and broad frequency shifting without a decrease in optical power as with AOMs. The light from the serrodyne modulation setup will be amplified by a single ILDL,

combined with repumper light with a fibre splitter and amplified by a fibre-coupled TA with free space output. This beam will be split between the pyramid source and glass cell MOT. This system is expected to offer better frequency tuning while reducing the number of required lasers and free-space optics.

4.8.2 PCB RF coils

A straightforward approach of eliminating cross-coupling between the RF coils entirely is using a PCB-coil which produces RF magnetic field along one of the directions (such as x) can be used. This configuration restricts the RF polarisation to linear only, which has a higher loss-rate at the same Rabi-frequency at the bottom of the shell trap compared to circular polarisation. Additionally, a linear shell trap would also capture atoms in the $|F = 2\rangle$ state, which must be expelled for efficiently BEC production of the $|F = 1\rangle$ species [73]. Since the completion of experimental work written up in this thesis, Erik Rydow has successfully optimised the experiment to create a BEC using linear RF polarisation only, which will be described in detail in his thesis [101]. The planned upgrade of the RF coil also involves an improved version of the RF pickups, mounted on the PCB itself. Instead of active temperature control, a low temperature coefficient (2 ppm/K), low-resistance (100 mOhm) current sense resistor (VCS1265 series) will be used to provide passive stability. Furthermore, main RF signal chain and the pickup signal chain will be isolated by an RF transformer (ADT1-1WT+) which is identical to the one used in the Urukul DDS synthesiser. This upgrade will reduce the number of components, shorten the coaxial cables and reduce the number of connectors used in the main RF signal chain, therefore increasing the passive stability of the experiment.

4.8.3 RF-dressed dimple trap

While BEC production is possible with the new apparatus, we were not able to achieve a significant improvement in repetition rate due to the difficulty of evaporative cooling in the shell trap. To increase the radial confinement for efficient evaporative cooling in the RF-dressed potential, a possible future direction is adding a vertical red-detuned dipole-trapping beam. The resulting combined magnetic and optical potential configuration is often called a dimple trap and has been used to produce rubidium BECs at a rate of 1 Hz using a Z-wire magnetic trap [77]. In this trap, there are two effects which contribute to fast evaporative cooling. Firstly, the increased radial trapping frequency at the bottom of the shell increases the collision and thermalisation rates, allowing faster evaporation. Secondly, the atomic population far from the dipole trap serves as a reservoir. After the end of the evaporation process, the dipole trap will be turned off adiabatically to minimise the excitation of collective modes in the condensate. Since the radial trapping frequency of the final potential is ≈ 10 Hz, this can be done safely on the timescale of a few seconds, leading to significantly shorter experimental sequences.

5 | Direct Loading and Bose-Einstein condensation in the shell trap

In this chapter, we start with a quick review of various existing loading schemes of RF dressed traps. We introduce a novel shell trap loading scheme which involves significantly fewer stages and control parameters compared to the method used in the previous generation BEC apparatus in Oxford. Next, we investigate evaporative cooling in the shell trap and show that a degenerate quantum gas can be created by conventional RF induced evaporation, without employing any additional modifications of the trapping potential for example by optical dipole forces. Finally, we characterise the evaporative cooling ramp and compare its efficiency to other experiments.

5.1 RF dressed trap loading schemes

Trapping atoms in RF-dressed adiabatic potentials was first realised in 2006 by Colombe *et al.* in a dressed quadrupole Ioffe-configuration (QUIC) trap [70]. The trap was loaded by ramping the frequency of the dressing RF field from below the Larmor frequency of atoms at the centre of the QUIC trap (1 MHz starting frequency) to a desired detuning (e.g. 2 MHz). This ensured that the atoms remained in the trapped state when the dressing RF field was turned on. The same loading procedure can be applied to load an RF-dressed quadrupole trap by ramping down the current in the Ioffe coils of the QUIC trap [69]. Different loading schemes of the dressed quadrupole trap (without Ioffe coils) involve using hybrid magnetic and optical potentials such as a plugged quadrupole trap [68] or a dimple trap [102]. However, in these schemes, the magnetic potential is significantly changed by optical dipole forces which makes them sensitive to control

parameters and the alignment of laser beams. Another scheme developed in Oxford involves only magnetic potentials which are deformed slowly until the final, RF-dressed quadrupole trap is loaded [71]. Initially, atoms are trapped in a time-orbiting potential (TOP) trap from which they are loaded into a time-averaged adiabatic potential (TAAP). Ramping down the rotating bias field loads the atoms into the RF-dressed quadrupole trap. This scheme is still being used by the previous generation experiment in Oxford. As this method does not rely on ramping the frequency of the RF field, it allows for the usage of a resonant matching circuit which reduces reflections from the RF antenna improving amplitude stability, and rejects low-frequency noise which would result in atom loss. The main disadvantage of this scheme is the number of control parameters needed as well as the complexity of the hardware, especially the need for an extra set of coils used to generate the rotating bias field for the TOP trap. The coupling of these extra coils to the RF antennae can drift over time which results in Rabi-frequency variations. Experiments such as splitting a quantum gas in a double-well potential require exceptionally stable Rabi-frequencies as the balance of atom number in the two wells is strongly dependent on them. An order of magnitude estimate for the required stability is 1000 ppm. In all of the previous loading schemes, either a BEC or a very cold gas close to the BEC critical point is loaded into the final RF-dressed trap because of the finite trap depth and the anisotropy of these traps. Therefore, it is very important to either adiabatically deform the potentials or to mode-match them. Otherwise, collective modes such as monopole or dipole oscillations are present in the cloud which has detrimental effects for matter-wave interferometry due to linear and quadratic phase gradients induced by these excitations. This requires careful and frequent optimisation of the experimental parameters in the previous generation apparatus in Oxford.

5.2 Loading a shell trap from optical molasses

The high loading rates of the second MOT in the UHV glass cell of the new apparatus allowed us to investigate a new scheme for loading the shell trap directly from a cloud of laser-cooled atoms. Even though the typical trapping volume of shell traps is small and highly anisotropic, we successfully loaded 80 million atoms at a temperature of $30 \mu\text{K}$ from optical molasses. We argue that this one-step scheme is conceptually simpler than those presented in the previous section. We load atoms at a phase-space density approximately 5 orders of magnitude below the BEC critical point, therefore evaporative cooling helps remove any collective mode present in the trapped gas, making this scheme ideal for matter-wave interferometry. The ability to load a large, hot cloud into the shell trap is also advantageous when the system is uncalibrated (this happens when a new radio-frequency is used). The new scheme allows for immediate RF spectroscopy, giving accurate measurements of RF amplitudes, while the previous generation experiment requires a straightforward but long and iterative approach in this scenario. We also demonstrate later in this chapter that the initial phase-space density is sufficient to achieve Bose-Einstein condensation by RF induced evaporative cooling. This section presents the different stages of the shell trap loading scheme, providing a detailed record of the parameters used, which could be replicated in other experiments.

5.2.1 Cell MOT loading

We operate the pyramidal MOT source with a cooling laser detuning of -22 MHz and with a quadrupole gradient of 13 G/cm (weaker directions). The single laser beam of the source has a peak intensity of 19 mW/cm^2 . The second MOT is operated with a cooling laser detuning of -23.7 MHz and a quadrupole gradient of

23 G/cm. The optical power in each beam is 5 mW. We typically load 4 billion atoms in 5 seconds into the glass cell MOT, which results significantly higher phase-space density compared to the 1 billion atoms at a temperature of 2 mK after transporting the atoms into the UHV glass cell in the previous generation Oxford experiment. The high loading rates are achieved by careful optimisation of the source through adjustments in the bias fields, beam alignment and polarisation. The loading is monitored with a photodiode onto which we image the cloud of fluorescing atoms in the MOT with a single lens. The photodiode reading is used to servo control the MOT loading time. After the loading is terminated we shutter the laser beam of the pyramid MOT source as a small fraction of it leaks through the low-conductance pipe used for differential pumping. We ramp the cooling detuning, quadrupole current and bias fields of the second MOT to different values in 100 ms to adjust for the absence of the scattering force of the pyramid leakage beam and to prepare for the next stage. The optimal parameters for this intermediate stage were found by machine learning optimisations. Only the change in the quadrupole gradient is significant, the best value was found to be 41 G/cm. This higher value in the intermediate stage helps by reducing the rate of change of current that the power supply has to achieve at the beginning of the next stage, where the required quadrupole gradient is much higher. The higher gradient also makes the MOT less sensitive to imbalances in the scattering force and stabilises the cloud against the lack of the pyramid MOT leakage beam.

5.2.2 Compressed MOT and molasses

To increase the phase space density of the laser-cooled atoms, we perform a compressed MOT (cMOT) stage [103] followed by optical molasses to achieve sub-Doppler cooling [104], as in the previous work by our research group [73]. We linearly ramp the quadrupole gradient and the laser detuning in 5 steps, each last-

ing 7.5 ms. The quadrupole gradients are (99, 100, 176, 69, 29) G/cm and the corresponding cooling laser detunings are (-19.6, -32.8, -37.7, -20.0, -20.0) MHz. At the end of the cMOT ramp, we trap ~ 1 billion atoms which have a temperature of $\sim 700 \mu\text{K}$ ¹. The quadrupole gradient and detuning values were found by machine-learning optimisations due to the high dimensionality of the parameter space. The reason for the non-monotonous increase in the quadrupole gradients is the finite response time of the power supply used to drive the quadrupole coils. While in an ideal cMOT, the quadrupole gradient increases monotonously, performing the molasses stage right after the cMOT stage requires quickly ramping the quadrupole gradient to zero. This is more efficient if the quadrupole gradient is relaxed before the molasses ramp. During the five-stage cMOT ramp we also perform a single stage linear ramp of the bias fields. We hold the atoms at the end of the cMOT stage for 5 ms. In the next 5 ms, we ramp down the quadrupole gradient to zero and the cooling laser detuning to -70 MHz, further cooling the atoms in optical molasses. At the end of the molasses ramp, we typically have half a billion atoms at a temperature of $\sim 100 \mu\text{K}$. The high detuning from resonance is necessary as the optical molasses contains a large number of atoms which can be heated by rescattering. During this ramp, we also change the bias fields of the cell MOT to values which cancel the stray magnetic fields created by the ion pump magnets and other sources. These values largely influence the efficiency of cooling during the optical molasses stage, hence loading the atoms into the shell trap. At the end of the control parameter ramps of this stage, the atoms are held in optical molasses for a brief period of 100 μs .

¹These numbers are approximate as the magnification of the imaging system is too high to measure a large, quickly expanding cloud in a time-of-flight measurement

5.2.3 Shell trap loading

The shell trap loading procedure starts with optically pumping the atoms into the $|F = 1\rangle$ hyperfine ground level by shuttering the repumping light in the cell MOT beams; 250 μs after turning off the repumping light, we shutter the cell MOT beams. Note that the shutter has a 1 ms delay in response so the actual optical pumping duration is 1.25 ms. After this, only a small population remains in the $|F = 2\rangle$ level. At this point we rapidly ramp up the quadrupole gradient to 37 G/cm and use the vertical and horizontal bias fields to displace the quadrupole node such that the centre of mass of the atoms coincides with the minimum of the RF-dressed trap to be loaded. At the end of the ramp we suddenly turn on a circularly polarised dressing RF field of frequency 4.2 MHz and Rabi frequency of 200 kHz which projects the atoms into the dressed eigenstates. The two hyperfine ground levels have g-factors of different sign, therefore only the lower $|F = 1\rangle$ level is trapped at the bottom of the shell if correct handedness of circular polarisation is used. This eliminates the need for removing them optically to avoid sympathetic cooling and collisional losses [73, 82]. A small fraction of the atoms remains trapped around the quadrupole node which are in the dressed state of opposite sign to the one trapped at the bottom of the shell. The ratio of the number of atoms in the shell trap and in the quadrupole trap depends on the combination of RF frequency, vertical bias field strength, and quadrupole gradient. The number of atoms in the quadrupole trap always remains small as the strong dressing-RF field acts as an RF-knife just as in the case of evaporative cooling in the quadrupole trap. The atoms trapped in the quadrupole trap decay slowly with an estimated lifetime of 25 seconds [59] due to Majorana spin-flip transitions, but they are quickly removed when the atoms in the shell trap are cooled with RF induced evaporation. An example of an in-situ absorption image of atoms trapped

in shell trap shortly after loading is shown in Figure 5.1.




Figure 5.1: In-situ absorption image of atoms trapped in the shell trap (crescent shape) along with a faint elliptical cloud trapped around the node of the quadrupole field.

When loading the shell trap from a cloud of laser cooled atoms, many factors influence the efficiency of transfer. Firstly, it is important to use a high Rabi frequency as this suppresses the non-adiabatic losses and increases the trapping volume. Secondly, the RF frequency should be chosen appropriately. For too low RF frequencies such as 1 MHz, the trapping volume is too small. The Landau-Zener losses are also more significant in small shells as atoms have enough kinetic energy to explore the upper half of the resonant spheroid where the local coupling is weak. Using a too high RF frequency is also problematic since it is difficult to produce strong RF fields at high frequencies due to stray capacitances of the coils. A too high RF frequency also results in the need to displace the trap a higher amount vertically which can be a technical limitation. The resonant spheroid may be compressed by increasing the quadrupole gradient but this results in a smaller trap volume due to the increased trapping frequency (see 3.20). Moreover, sud-

denly increasing the quadrupole gradient to high values is difficult in practice because of the inductance of the coils generating the quadrupole field. In our work, dressing frequencies ranging from 3 MHz to 6 MHz have worked best.

5.2.4 RF spectroscopy of the shell trap

RF spectroscopy is an important tool in characterising RF-dressed traps [92]. The method relies on using a weak RF-field to drive transitions between trapped and untrapped dressed states. Since the energy splitting between the dressed states varies spatially, it is possible to eject atoms which are above a certain potential energy in the trap. Equivalently, the probe RF-field is a dressing field itself creating its own avoided crossing with a low Rabi frequency, which acts as a controlled hole in the trap. There are a large number of transitions which can be driven, the simplest being the first-order transitions corresponding to single-photon processes. For a circularly polarised RF dressing field, there are 3 such transitions resonant at Ω_R and $\omega_{RF} \pm \Omega_R$ [105]. The $\omega_{RF} + \Omega_R$ transition is illustrated in Figure 5.2 a and is resonant for atoms at the bottom of the RF-dressed potential. If the probe field of this frequency is applied to the atoms for a long-enough duration they all escape, providing a method to measure the Rabi frequency with high accuracy. The polarisation of the probe field (along the x -direction) and the vanishing gain of the RF amplifier used at low frequencies mean that it is the most convenient to use the $\omega_{RF} + \Omega_R$ transition. For this transition, any probe frequency above the resonant frequency ejects a certain fraction of the atoms depending on the energy distribution in the trap. However, for probe frequencies below the resonant frequency, we expect no atom loss as there exists no energy levels in the RF-dressed trap below Ω_R (relative to the energy of the magnetically insensitive state). This is illustrated with experimental results in Figure 5.2 b, where the orange line corresponds to a Boltzmann energy distribution at 7.5 μK which

was proposed as an ansatz and experimentally verified for similar temperatures in [92]. The exact shape of the spectral lines depends on many factors such as power or Fourier broadening, the presence of a condensate in the trap or the anharmonic nature of RF-dressed potentials. Nevertheless RF spectroscopy provides accurate measurements of Rabi frequencies which is an important step towards manipulating atoms in MRF potentials.

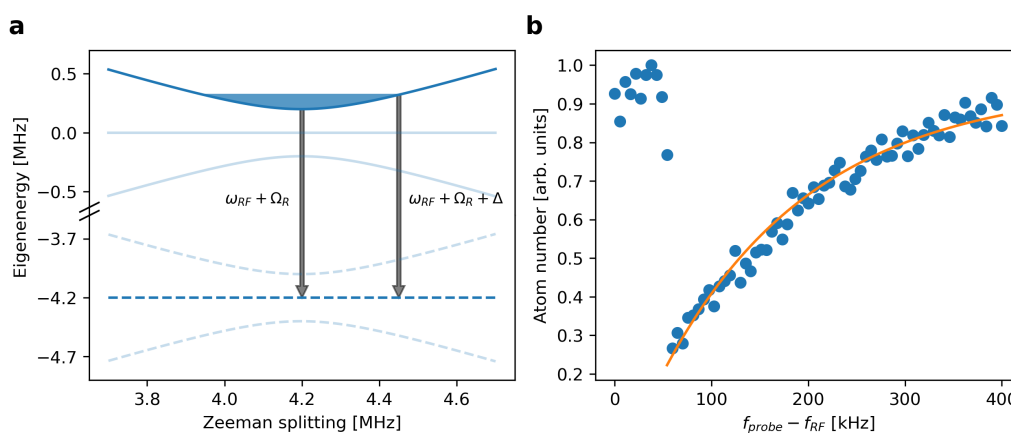


Figure 5.2: **a**: Illustration of RF spectroscopy considering the $\omega_{RF} + \Omega_R$ transition driven resonantly and off-resonantly. For blue frequency detuning, a fraction of the atoms remain in the trap as shown by the blue filled zone. **b**: Experimental results for RF spectroscopy. Orange line serves as a guide to the eye and it corresponds to a Boltzmann distribution at $7.5 \mu\text{K}$, assuming a harmonic trap.

5.2.5 Calibration of the RF amplitude and polarisation

The RF amplitudes are calibrated using the RF spectroscopy method outlined above. After measuring the Rabi frequency of linear shells produced by the x and y coils separately, we calculate the conversion factor between Rabi frequency and the value in the amplitude control register of the DDS, assuming a linear relationship². This enables us to supply meaningful control parameters such as Rabi

²Note that amplifier saturation and other non-linear effects may modify this trend at high RF amplitudes.

frequency instead of a machine unit numbers. Once the RF amplitudes are calibrated, the phase difference between the x and y dressing fields is swept linearly and the number of atoms in the shell is measured after a long, e.g. 10 seconds hold time. For polarisations other than the circular polarisation of the correct handedness, the atom loss rate is higher and therefore we measure a lower number of atoms following a long hold time. Once the calibration procedure is completed, signals from the RF pickups are recorded and kept as reference. Signals measured in later experiments can be compared to these and deviation from the reference signals indicates the need for repeating the calibration procedure.

5.3 Evaporative cooling and Bose-Einstein condensation

Evaporative cooling of trapped atoms relies on the truncation of the long tail of the energy distribution and re-thermalisation to a lower temperature via collisions. For this scheme to be effective, the thermalisation rate must be significantly faster than the heating rate in the trap. In the shell trap, two energy-dependent loss mechanisms can be used to achieve this. Firstly, atoms with high kinetic energy can escape via non-adiabatic losses as they traverse through the avoided crossing. Secondly, a weak RF field can be used to perform RF-induced evaporative cooling in a conventional way as in DC magnetic traps. This exploits the same loss mechanism as RF spectroscopy, but requires ramping the probe frequency.

Evaporative cooling in the shell trap presents several challenges. Firstly, to the lowest order, the trap is harmonic along all 3 spatial dimensions. The scaling of density and collision rate with atom number makes harmonic potentials inferior to linear potentials realised in the quadrupole trap [106] from the point of view of

evaporative cooling. The efficiency of evaporative cooling is further reduced by the high trap anisotropy, as the atoms predominantly cross the evaporation surface along the vertical direction. This gives an effective evaporation dimension of one [107]. For these reasons, Bose-Einstein condensation has not been achieved in the dressed quadrupole trap prior this work [69].

The efficiency of evaporative cooling is typically characterised two parameters that relate logarithmic changes in temperature T , number of atoms N , and phase space density \mathcal{D} :

$$\alpha_{\text{evap}} = \frac{d(\log T)}{d(\log N)}, \quad (5.1)$$

and

$$\gamma_{\text{evap}} = -\frac{d(\log \mathcal{D})}{d(\log N)}. \quad (5.2)$$

Considering a constant α_{evap} and a harmonic trap, Bose-Einstein condensation is only possible if $\alpha_{\text{evap}} > 1/3$ since the transition temperature scales with atom number as $N^{1/3}$. This scaling also demonstrates the effect of various evaporation efficiencies. Given an initial state with a certain atom number N_0 and temperature T_0 , different evaporation efficiencies change the intercept between the evaporation trajectory and the BEC critical point and provide an upper bound on the number of atoms in the condensate. γ_{evap} gives a more direct figure of merit of evaporative cooling as higher values of this parameter result in a larger condensate at a given phase space density. Note that α_{evap} or γ_{evap} may not be constant during evaporation. Depending on the trapping scheme and the temperature, the effective dimension of the evaporation can change, affecting its efficiency. The elastic collision rate (per particle) at the centre of the trap is given by

$$\Gamma_{el} = n_0 \sigma \sqrt{\frac{16k_B T}{\pi m}} \quad (5.3)$$

where n_0 is the density at the centre of the trap and σ is the s-wave scattering cross section which is $5.4 \times 10^{-16} \text{ m}^2$ for Rubidium 87 in the $|F = 1\rangle$ level [108]. For high collision rates, evaporation can enter the runaway regime, where the elastic collision rate increases as time progresses, thus accelerating BEC production.

To perform evaporative cooling, we use the $\omega_{RF} + \Omega_R$ transition for the same reasons as in the case of RF spectroscopy. The evaporation cooling ramp is executed in six steps, each characterised by the frequency and amplitude of the probe RF field, quadrupole gradient, Rabi frequency and duration. Between the steps, we interpolate the respective parameters linearly. Initially, we employed a 1 ms update rate and a varying number of steps, hence controlling the duration of each evaporation segment. This scheme was later improved by using the Digital Ramp Generator feature of the DDS synthesiser which can create high update rate and high-resolution ramps of frequency, amplitude or phase. Using this feature, the evaporative cooling DDS generates a frequency ramp and the two dressing RF DDS amplitudes are ramped similarly. This approach offloads the simple operation of linear interpolation of RF amplitudes and the evaporation frequency from the control system to the DDS chips, reducing the data transfer and the failure rate. Since the ramp generator can only control a single parameter, the amplitude of the evaporative cooling field is controlled by successive updates at a slower rate. This is not a limitation as the amplitude affects the efficiency of evaporative cooling weakly [85]. To enable full DRG functionality, the CPLD gateway of the Urukul DDS synthesiser was customised. The parameters of a typical evaporation ramp are displayed in Table 5.1: f_{evap} is the frequency of the evaporative cooling field measured relative to the frequency which empties the trap completely ($f_{RF} + \Omega/2\pi$). For instance, at step 5 the actual frequency is 4350 kHz. The amplitude of the evaporative cooling field is controlled logarithmically, such that 0 is the maximum amplitude and -9 is the amplitude that would correspond to 1

LSB of the output DAC of the DDS synthesiser. Logarithmic control is useful as we have to change the DDS output amplitude of the evaporation signal significantly because of the resonant match. The parameters in Table 5.1 were found by machine learning optimisation and it can be seen that they are not necessarily monotonous. For instance, the very low starting amplitude can be explained by a third-order transition at a frequency of $2\omega_{RF} - \Omega_R$ [105]. While this transition is in general weak, high amplitudes are preferred in the initial phase of evaporation to remove a sufficient number of atoms from the trap. Therefore, we attribute the low starting amplitude to a compromise between a high starting frequency and driving the third-order transition weakly enough. The small change of f_{evap} between steps 1 and 2 has a similar explanation. Step 1 is likely to cut too deep into the cloud but it is necessary to avoid the higher-order transitions. Therefore, keeping the frequency around 2 MHz for 2.5 seconds compensates for this. After this step, the amplitude is dropping almost monotonously. The quadrupole gradient and Rabi frequency decrease around steps 3 and 4 compresses the cloud hence increases the collision rate. However, with high collision rates, the cloud may enter the hydrodynamic regime where the efficiency of evaporative cooling is reduced [109]. This occurs when the collision rate exceeds the oscillation frequency and we are close to this in the radial direction.

	f_{evap} (kHz)	$\ln(A_{evap})$	b' (G/cm)	$\Omega/2\pi$ (kHz)	Time (s)
Step 0	2800	-9.00	36.6	98	0
Step 1	2000	0.00	44.1	96	2.5
Step 2	1981	-0.77	53.6	85	5.0
Step 3	705	-0.66	61.8	49	8.5
Step 4	175	-2.52	47.6	29	13.0
Step 5	52	-3.91	47.0	98	19.5
Step 6	22	-3.59	49.6	98	26.5

Table 5.1: Parameters along the evaporative cooling ramp

We characterise the evaporative cooling ramp in the following way. We carried out

the evaporation ramp truncated at various percentages of the full duration of 26.5 s and held the atoms for 500 ms to allow them to reach thermal equilibrium after the end of the ramp. We imaged the cloud after 15 ms time-of-flight and fit the resulting column density with polylogarithmic and bimodal (sum of polylogarithmic and Thomas-Fermi) distributions to extract quantities such as atom number, temperature or phase-space density from the time-of-flight distribution of atoms [110]. The results of the characterisation sweep can be seen in Figure 5.3. On some panels, the x -axis shows the control parameter "Evap percentage" which is the ratio between the length of the truncated and full ramps. For instance a value of 20 percent means that the ramp is terminated at 5.3 seconds, shortly after the start of Stage 2. In panels a and b, the atom number and phase space density are shown as a function of the evap percentage. Due to the anisotropy of the shell trap, $t_{TOF} \gg \omega_z^{-1}$ but $t_{TOF} \approx \omega_r^{-1}$ therefore T_z provides a better estimate to the temperature of the cloud. Nevertheless, there is a good agreement between the two temperatures when plotted on a logarithmic scale. Panel c shows the calculated collision rate per particle showing that the experiment operates in the runaway regime initially. Approximately, 60 percent along the evaporation ramp, the collision rate starts dropping which is compensated with longer stages as shown in Table 5.1. Panel d shows the temperature (along the z -direction) as a function of atom number, both on logarithmic scales. The data shows a linear trend which implies an approximately constant value of $\alpha_{\text{evap}} = 1.00(1)$. This value is larger than the efficiency of the previous generation Oxford experiment (0.93) [73], despite being faster, and is comparable to recent strategy employed in hybrid traps which give an efficiency of 1.1 [102]. This result is surprising as it was previously thought in our research group that evaporative cooling is more effective in the TOP or TAAP traps because of their stronger radial confinement. Furthermore, it demonstrates that the high initial atom number increases the efficiency

of evaporative cooling. The linear trend also indicates that the effective dimensionality of evaporation stays one, which is expected due to the anisotropy of the trap. In panel e we show the phase space density as a function of the atom number, both on logarithmic scales. Similar to the earlier panel, the results show an approximately constant $\gamma_{\text{evap}} = -1.99(2)$. This is smaller compared to -3.5 in hybrid traps reported in [102]. However, in that work, high efficiency evaporation was only sustained until a phase space density of the order of 10^{-1} and the number of condensed atoms at the end of evaporation was similar to here. Finally, in panel f the image of a Bose-Einstein condensate is shown with integrated profiles along both directions fitted with Gaussian and bimodal distributions. For the sake of clarity, this image shows a cloud that has a significant condensate fraction and it does not correspond to the 100 percent evaporation datapoint where the system just crossed the critical point. The bimodal nature of the distribution is easier to resolve along the horizontal profile because of the anisotropy of the shell trap as excitations predominantly occur along the radial direction. The two fit models are compared by sum of squares (SSQ) tests. In case of the horizontal profile, the bimodal fit yields a factor of 5 lower SSQ. In case of the vertical profile, the bimodal result still yields a lower SSQ than a polylogarithmic but only by 20 percent.

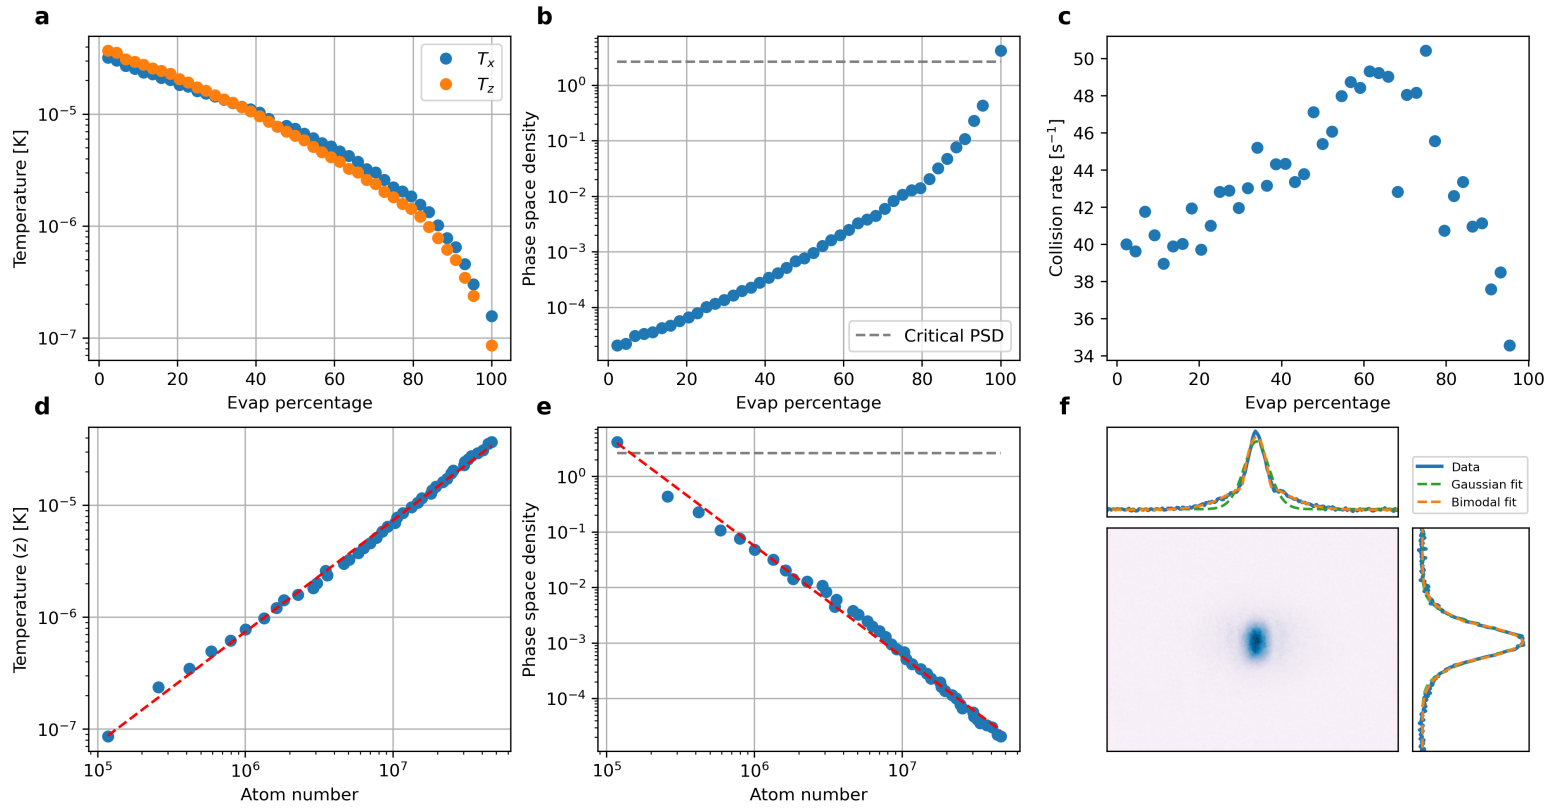


Figure 5.3: Characterisation of the evaporative cooling scheme. **a**: Temperature measured along the horizontal (x) and vertical (z) directions as a function of evap percentage. Due to the high anisotropy of the shell trap measurements of T_x are less reliable compared to T_z . **b**: Phase space density vs. evap percentage. The BEC critical point is crossed at the end of the ramp. **c**: Collision rate per particle vs. evap percentage. The last data point was excluded due to the presence of a condensate. **d**: Temperature vs. atom number with linear fit indicating a nearly constant α_{evap} . **e**: Phase space density vs. atom number with linear fit indicating a nearly constant γ_{evap} . **f**: Absorption image of the Bose-Einstein condensate at with integrated density profiles along the horizontal and vertical directions.

5.4 Conclusion and outlook

In this chapter, we reviewed different loading schemes of RF-dressed adiabatic traps, outlining their advantages and disadvantages. We presented a new scheme of loading a shell trap directly from a cloud of laser-cooled atoms. We showed how RF spectroscopy is utilised to calibrate the amplitudes of RF fields in the shell trap. It was shown that despite the disadvantages of the highly anisotropic shell trap being far from the optimal shape for evaporative cooling it is possible to create a Bose-Einstein condensate in this trap. We characterised the evaporative cooling ramp and compared its efficiency to the previous generation Oxford experiment as well as recent results using a combination of optical and magnetic potentials.

The experimental procedures presented in this chapter substantially simplify the experimental hardware required to realise a shell trap. Since the trapped gas is cooled over orders of magnitude of temperature in the final trapping potential, collective excitations are strongly suppressed which makes this system ideal to perform precision matter-wave interferometry. Therefore, this loading scheme is an important advance towards commercialisation of atom interferometers utilising RF-dressed technology.

6 | Full counting statistics of interference contrast of bilayer 2D Bose gases

In this chapter, we present the results of experiments investigating the full counting statistics of interference contrast in bilayer 2D Bose gases. We start with a brief review of theoretical and experimental results which use full counting statistics of observables to explore many-body quantum states. Following this, we introduce our experimental protocol which is based on matter-wave interferometry. We discuss the analysis of matter-wave interference fringes of two identical 2D Bose gases. From the local phase data, we analyse the resulting correlation functions and identify the BKT critical point. We observe that the shape of the measured contrast distribution changes across the BKT critical point and compare the data with numerical simulations and our findings demonstrate that it is important to account for additional noise sources in such comparisons. We discuss the universal properties of contrast statistics and we find good agreement between theory and experiment in the low temperature limit. Finally, we analyse the central moments of the measured distributions of the interference contrast and show that as the critical point is crossed, the higher order moments saturate, which provides an alternative way to identify the BKT critical point in bilayer 2D Bose gases.

6.1 Related work

Characterising quantum systems experimentally is a difficult task because of the projective nature of quantum measurements. Quantum tomography, first introduced in the context of measuring the quantum state of light [111] offers a solution but in practice it is limited to a few particles [112]. For quantum many-body systems, to obtain the most information, one has to measure the full counting statistics

of an observable such as individual electrons in superconductors [113], or current noise in a quantum dot [114]. These solid-state systems, however, do not have reproducible sample preparation and suffer from noise, which result in the need for averaging over many realisations of the system or over a long time. In contrast, ultracold atomic systems offer extreme levels of isolation from the environment and repeatable sample preparation. Experimental measurements of the full counting statistics of interference contrast yielded very interesting results for 1D quantum gases. It has been shown that the mean value of the interference contrast can be used as a thermometer and that from its distribution functions it is possible to infer whether the fluctuations of this observable are of classical or quantum origin [115]. The same method has also been used to probe out-of-equilibrium physics in 1D systems [116, 117]. The full counting statistics of interference contrast has been predicted to exhibit similar characteristics for 2D systems in equilibrium, as these can be thought of as being coupled 1D systems [14, 118]. Approaching the BKT transition, it has been predicted that the distribution of interference contrast changes from a peaked Gumbel to a monotonically decaying Poisson distribution in the case of 2D systems. However, the non-equilibrium generalisation of the distribution functions is an open theoretical problem, which our experiments can illuminate in the future.

6.2 Experimental protocol

We prepare a bilayer 2D Bose gas system using the MRF-dressed potential splitting procedure described in detail in [74]. In this experiment, we use radio-frequencies of 7.0, 7.1, and 7.2 MHz¹ and a quadrupole gradient of 170 G/cm.

¹The exact frequencies are about 1 percent different for technical reasons, but the sum of the lower and upper frequencies precisely equals to twice the middle frequency, as generated digitally with a common clock signal.

While the gas is still in a single-well potential, a 532 nm, ring-shaped laser beam propagating along the vertical direction creates an additional optical dipole potential. The optical system is based on direct imaging of a DMD onto the trapped atoms and it is described in detail in the thesis of David Garrick [76]. Since that work, the telescope system before the DMD has been changed to use best-form singlet lenses because earlier doublet lenses in the previous optical system were damaged by the high optical intensity. The new telescope system has a different magnification, therefore the optical potential strength calibrations in [76] had to be updated. The ring of light is set to have a 25 μm inner and 50 μm outer radius at the location of the atoms and the potential is turned on slowly over 0.8 seconds to avoid exciting collective modes. The resulting trap with hard walls realises a near-uniform system. The optimal strength of the dipole potential was found by measuring the phase correlation function at different green laser powers for cold systems in the superfluid regime. As the system becomes more uniform, longer range correlations are observed resulting in a lower value of the algebraic exponent. The in-situ density distribution captured with imaging in the vertical direction, after averaging over the angular coordinate, is shown in Figure 6.1 for an earlier experiment with a ring potential of inner diameter 35 μm .

The cloud that is initially trapped in a single-well potential is split slowly, over 0.45 seconds, to avoid exciting collective modes. For the last 0.15 seconds of the splitting, the barrier height is higher than 1.5 kHz, therefore we can consider the cloud to be fully split. Since the RF amplitudes, phases and quadrupole gradient are calibrated precisely to the 1 percent level, we use numerical methods based on Floquet theory to calculate the trapping potential [73]. The accuracy of the numerical predictions has been tested extensively in the research group by comparing to theoretical expressions of trapping frequencies and splitting experiments which are particularly sensitive to the shape of the potential [73, 74, 119]. At the




Figure 6.1: In-situ density distribution, averaged over the angular coordinate, showing a near-uniform system over a width of $\approx 32 \mu\text{m}$ for a ring potential of inner diameter $35 \mu\text{m}$.

end of the splitting procedure, the wells are separated by $4.9 \mu\text{m}$ by a barrier of height 4 kHz . The axial trapping frequencies are $900(60) \text{ Hz}$ and $830(40) \text{ Hz}$ in the higher and lower wells respectively, therefore the gases are in the quasi 2D regime in both wells for the typical temperature of 40 nK ². The radial trapping frequencies are $10.6(2) \text{ Hz}$ and $11.0(2) \text{ Hz}$ respectively. The difference in axial trapping frequencies arises from a technical compromise between adequate barrier height and small manifold losses (see Section 3.4), resulting in a small, 4 percent difference between the 2D interactions strengths in the two clouds. Such a small difference results in a BKT critical temperature difference of 0.4 percent between the two subsystems. For typical experimental conditions, this broadens the BKT crossover by 200 pK which is small relative to the typical critical temperature of 40 nK and well within the temperature measurement uncertainties and the width

²For these parameters, $\hbar\omega_z/k_B T \approx 1$ therefore the system is not deep in the quasi 2D regime. However, we expect a smooth shift between the 2D BKT and 3D BEC critical temperatures across the dimensional crossover [120]. A more detailed calculation for our system [74] based on numerical solutions of the equation of state of a 2D Bose gas [121] shows that the quasi 2D condition is satisfied for $\hbar\omega_z/k_B T \gtrsim 0.5$.

of the BKT crossover in our finite-sized system. After completing the splitting procedure the RF amplitudes are kept constant and we let the system equilibrate for 0.3 seconds before releasing the atoms from the trap to undergo time-of-flight expansion for 16.5 ms. We rapidly turn off the optical potential 2 ms before the release to ensure that it does not affect the atoms during time-of-flight. Before imaging, we optically pump the atoms to the $|F = 2\rangle$ level under 150 μs by a light sheet of thickness 8 μm , enabling us to image only atoms in this thin slice. Note that the thickness of the light sheet is a factor of two larger than what is used in previous works [12, 75] where the emphasis was on probing the local phase. Here, it is practical to extend the region of interest such that the measured contrast is proportional to the product of the number of condensed atoms in the two clouds. Moreover, we reduce the signal to noise ratio on the images by 2, since we image approximately twice as many atoms.

To control phase space density and tune the system across the BKT critical point, we aim to keep the temperature constant and change the number of atoms before the splitting procedure. This is achieved by changing the temperature of the cloud at the end of evaporation in the quadrupole trap. Loading a hot cloud into the RF-dressed potential results in the rapid loss of energetic atoms from the finite trap depth, leading to a lower final number after the second stage of evaporation. However, the temperature stays approximately constant as it is determined predominantly by the final frequency of the second evaporation stage. To verify this, we repeat the measurements with two modifications. Firstly, we repump the entire cloud to ensure accurate atom number measurement. Secondly, we adiabatically ramp down the hard-wall potential before time of flight expansion to load the atoms back to the harmonic trap. This way, we can accurately measure the temperature and condensate fraction through bimodal fits. In Figure 6.2 we show that the temperature stays constant within measurement error while the atom

number decreases approximately linearly with the "Quadrupole cut" control parameter which defines the lowest frequency reached by the evaporative cooling RF field. At each value of the control parameter, we repeat the experiment 10 times for these measurements, while we perform 200 repeats with selective repumping to obtain reliable statistics of interference contrast. The temperature and atom number measurements suffer from two systematic errors due to ramping down the hard-wall potential before time-of-flight. Firstly, this method does not distinguish between atoms trapped inside and outside the ring, giving a higher atom number count. Secondly, ramping down the hard-wall potential decompresses the trapped cloud adiabatically which results in a decrease in temperature. This effect is investigated in detail in Chapter 7 for a wide range of laser powers and is shown to not result in a significant shift in temperature. To eliminate these systematic effects in the future, thermometry and atom number measurements should be performed without adiabatically loading the atoms back into the all-magnetic harmonic trap.

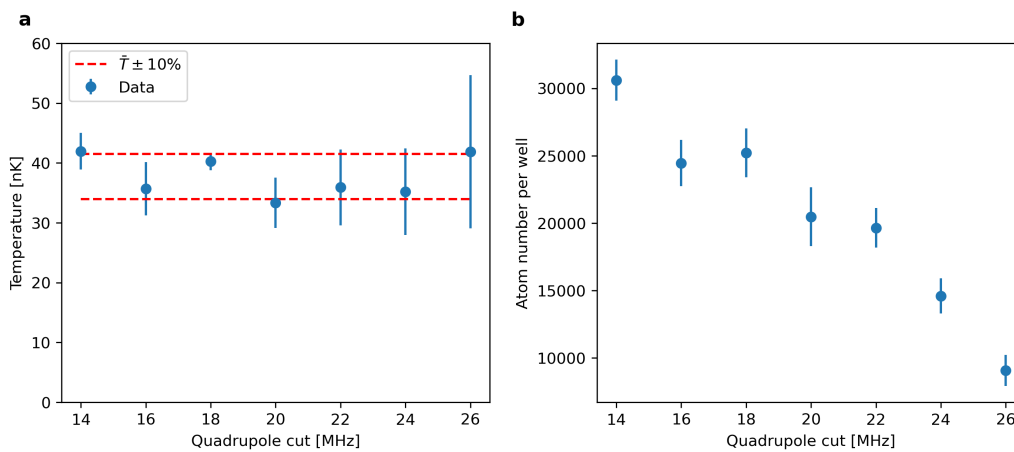


Figure 6.2: Thermometry results for the contrast statistics experiment. **a**: Temperature as a function of quadrupole cut plotted with lines indicating 10 percent deviation around the mean. **b**: Atom number in each potential well (assuming equal split).

It is important to ensure equal populations in the two wells and prevent relative

motion between the clouds. Therefore we optimise the splitting procedure by changing RF amplitudes and timings at various stages of the splitting by maximising the interference contrast. We detect this as the Fourier amplitude of the fully repumped system, integrated over the horizontal coordinate, calculated at the expected spatial frequency of the fringes. This favours, high contrast, flat fringes. Based on our experience, the population difference is best controlled by the amplitudes of the two RF fields which create the two wells when the barrier is approximately 1 kHz high. The most common source of relative velocity between the clouds is sloshing which can occur if there is a non-circular RF polarisation when the RF-dressed trap is loaded. With accurate calibration of the RF field strength and setting the correct phase shift between the signals it is possible to eliminate sloshing, as demonstrated in our previous work [12, 13].

6.3 Correlation function analysis

The data analysis begins with extracting the density distribution of the atoms from the raw images. To remove spurious interference fringes which arise from the use of coherent probe light, we reconstruct the optimal basis images utilising the eigenface algorithm [122]. In the following, we refer the extracted atom number distribution as the image. Initially, we sum the images over the repetitions as well as along the vertical z direction. From the integrated 1D profile, we extract the centres of the clouds on each image along the x direction as well as the size of the quasicondensate, which we refer to as the system size. We then analyse the images column by column in a range given by 85 percent of the system size. This is chosen as reasonable compromise between utilising as much of the image as possible and avoiding edge effects and has been used in our previous work [12, 13]. The Discrete Fourier Transform (DFT) of each column is computed and

the absolute values of the resulting spectra are averaged for all images and all columns. The resulting average spectrum consists of a Gaussian centered at zero spatial frequency (representing the Gaussian envelope of the cloud), and a higher momentum peak which corresponds to the matter-wave interference fringes, thus we can identify the relevant spatial frequency k_0 for the fringe analysis. Instead of extracting the phase from a single spatial frequency, the phases are averaged over several points in the neighbourhood of k_0 . The range of averaging is determined by the full width at half maximum of the peak in the average spectrum. The advantages of this method are outlined in Appendix C where we demonstrate using computer-generated data that such an averaging process reduces the uncertainty in the measured phase. Compared to the earlier method of finding the phase by fitting a sinusoidal curve with a Gaussian envelope [12, 74], this DFT-based scheme provides fast, less noisy and deterministic results. An example for phase extraction is shown in Figure 6.3 a. After the extraction of phases, we calculate the correlation matrix $C_{i,j}^k$ for each image labelled by k that we define as

$$C_{m,n}^k = e^{i(\phi_m^k - \phi_n^k)}, \quad (6.1)$$

where ϕ_m^k is the extracted phase at column m (relative to centre) of image k . The correlation matrices are then averaged for all runs with identical experimental parameters to obtain $\bar{C}_{m,n}$. One such averaged correlation matrix is shown in Figure 6.3 b. Finally, we calculate the correlation function $C(\bar{x}_j)$ by averaging $Re(\bar{C}_{m,m+j})$, i.e. the diagonal of $Re(\bar{C}_{m,n})$ with offset j in a range ± 5 around the centre of the averaged correlation matrix. This final averaging process is illustrated in Figure 6.3 c.

The measured correlation functions are fitted with the theoretical algebraic and exponential models. The fit is constrained to distances less than half the system




Figure 6.3: Illustration of the correlation function analysis. **a**: Phase extraction from an absorption image. **b**: Average correlation matrix. **c**: Correlation function extracted from correlation matrix. The three datapoints highlighted with different colours correspond to the three dashed lines indicated on the correlation matrix.

size since these data points correspond to image columns where the density is far from the value in the central, near-uniform region. The correlation length for the exponential fit is bounded from above by the system size as this quantity is ill-defined when larger than the system size. To compare the validity of the two fit models, we perform the reduced chi-squared test and identify the critical point as being where the two models describe the data equally well. From the thermometry and atom number measurement results, we calculate the mean phase space density \mathcal{D}_{avg} across the region of interest assuming a uniform system of radius $25 \mu\text{m}$ and use this dimensionless quantity as our primary independent variable. Note that there are two systematic errors which affect this calculation. Firstly, the system is not truly uniform and the decaying density near the edge of the ring potential results in a lower calculated mean phase space density at the centre of the trap. Secondly, the measured atom number also contains atoms which, in an interferometry experiment, are outside the ring potential, thus increasing the calculated mean phase space density. The results of the fits for selected mean atom numbers across the BKT transition are shown in Figure 6.4. For $\mathcal{D}_{avg} = 12.5$ the algebraic model clearly fits the observations better. At larger distances (above $\approx 12 \mu\text{m}$), the

correlation function changes behaviour and starts decaying at a faster rate because of the residual non-uniformity of the gas in our trapped system. Approaching the crossover region ($\mathcal{D}_{avg} = 11.7$ and 11.1), the measured correlation function is decaying faster, closer to the decay curve of the exponential model with the longest possible correlation length given by the system size. For the lowest mean phase space density, the data clearly follows the exponential model.




Figure 6.4: Correlation functions fitted with algebraic and exponential models for selected values of the mean phase space density across the BKT transition.

The fit results are summarised in Figure 6.5. Firstly, we performed a reduced chi-squared test to investigate whether the algebraic model describes the correlation function better at each \mathcal{D}_{avg} . The reduced chi-squared values for the two models cross around 9.4. We estimate an uncertainty of 0.25 to the critical phase space density as illustrated by the shaded region. This value is in good agreement with experiments performed in harmonic traps [74] and close to the theoretical value in true uniform systems (8.5) [40]. In Figure 6.5 b we plot the exponent of the algebraic model η as a function of \mathcal{D}_{avg} . Using the previously found crossover region, we calculate the lower and upper bounds of the critical exponent by approximating the data with a cubic interpolant. This gives a critical exponent of 0.18(2) which is in good agreement with earlier results in trapped, finite-size systems [12, 74]. While there exist theoretical predictions for $\eta(\mathcal{D})$ in true uniform systems, they are either valid only deep in the superfluid regime [1], or below the critical point

[31] as opposed to our experiment which measured the algebraic exponent across the critical point. Therefore, we performed MC simulations, which incorporate effects from the residual inhomogeneity and the finite size of the system and find good agreement between the experiment and simulation results. Finally, we plot the correlation length r_0 which characterises the decay of the exponential model. As the crossover region is passed, the correlation length drops rapidly, which is a signature of the BKT transition. Once again, the experimental data shows good agreement with the MC simulations.

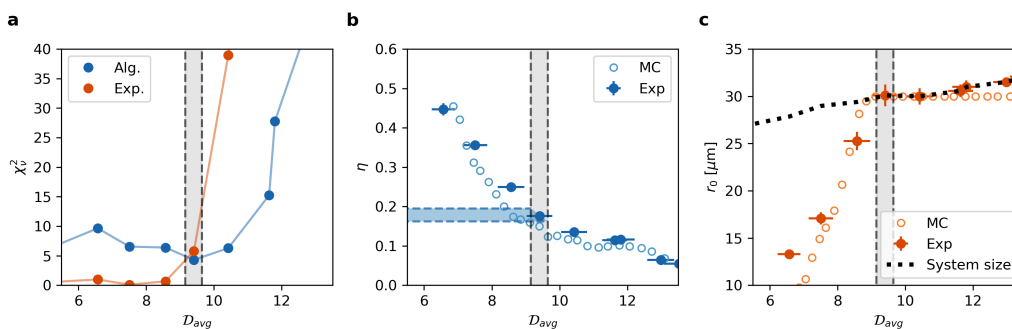


Figure 6.5: Results of the correlation function analysis. **a**: Reduced chi-square test of the fits using the two models. **b**: Measurements of the algebraic exponent as a function of the mean phase space density. The observed values are in good agreement with MC simulations (open circles). **c**: Measurements of the correlation length as a function of the mean phase space density, with the system size plotted as a dashed line. Open circles show results of MC simulations.

6.4 Contrast analysis

The same dataset was analysed to extract the full counting statistics of interference contrast. Similar to the phase analysis, a DFT-based method was used to obtain the contrast. Also, we take the mean of the absolute values of the complex Fourier amplitudes around k_0 for reasons explained in Appendix C. After contrast extraction, we post-select the data by ignoring datapoints with contrast above 0.5. This

is to avoid the inclusion of datapoints with unrealistically high contrasts, often reaching near 1 which is due to low signal to noise ratio. This cutoff is determined from the highest contrast measured with very cold clouds and using the thinnest possible imaging region, which corresponds to the upper bound given by the finite imaging resolution [75]. We further post-select the data by imposing a lower limit of contrast equal to the RMS contrast fluctuations which we estimate from the standard deviation of the real and imaginary parts of the Fourier amplitudes above $2.5k_0$ and averaged over all columns and repetitions. This is to remove data points with low signal to noise ratios. Finally, we only analyse columns which lie within half of the system size (central 50 percent of the cloud). For significantly larger analysis regions, the contrast is reduced because of the non-uniform density of the condensate, while for significantly smaller regions, the amount of data is not sufficient to determine the probability distributions accurately. We then define uniform bins of the normalised squared contrast variable $\tilde{v} = C^2 / \langle C^2 \rangle$ on the $0 \rightarrow 3.5$ range and count the observation frequencies in each bin. We approximate the uncertainties as \sqrt{N} where N is the count in each bin.

6.5 Simulating contrast full counting statistics

To compare experiment with theory, one has to rely on numerical simulations as both the quantum operator corresponding to contrast and the quantum state of the system are complicated mathematical objects. To simulate the full counting statistics of interference contrast of a pair of 2D gases, we use the approach developed in the theoretical work of Rath *et al.* [14]. For a pair of 2D gases described by classical fields $\psi_1(\mathbf{r})$ and $\psi_2(\mathbf{r})$ before time-of-flight expansion, the observed

squared contrast is given by

$$C^2 = \frac{1}{N^2} \int_D d\mathbf{r}_1 \int_D d\mathbf{r}_2 \psi_1^*(\mathbf{r}_1) \psi_1(\mathbf{r}_2) \psi_2(\mathbf{r}_1) \psi_2^*(\mathbf{r}_2), \quad (6.2)$$

where it is assumed that both gases contain N atoms within the integration domain D which corresponds to the region of selective imaging. The two classical fields are generated by MC methods as described in Appendix D.8 and the expression 6.2 is evaluated by numerical integration. If the density fluctuations of the gases are negligible, the above equation can be simplified, giving

$$C^2 \approx \left| \frac{1}{N} \int_D d\mathbf{r} \bar{n}(\mathbf{r}) e^{i(\phi_1(\mathbf{r}) - \phi_2(\mathbf{r}))} \right|^2, \quad (6.3)$$

where $\bar{n}(\mathbf{r})$ is the mean density which we estimate for a given potential using the Thomas-Fermi approximation or by numerical GPE simulations. The phase fluctuations can be generated by the transfer-matrix approach and sampling an Ornstein-Uhlenbeck (OU) process [123]; this is significantly faster than the MC approach. Strictly speaking, this method is only applicable in the superfluid regime since the correlation functions it generates always decay algebraically (within the local correlation approximation). For our finite size system the algebraic model fits the measured correlation function with reasonable accuracy even for conditions above the critical point where the normal phase is expected; in this regime the algebraic exponent is greater than the value at the critical point. This method takes into account the increase in phase fluctuations due to slow density variations, but the contrast statistics calculated from such fluctuations ignores density fluctuations which are present in 2D systems. To match the simulation parameters to the experimental conditions, we generate Thomas-Fermi profiles corresponding to the measured values of atom number in the condensate. We then adjust the temperature such that the phase fluctuations give the correct superfluid

density, thus the correct algebraic exponent and use the generated phase fields to compute the distribution of interference contrast using 6.3. The difference between the predictions of the MC and stochastic simulations is expected to grow as the system moves towards the normal phase. The computation work to implement the stochastic algorithm was mainly carried out by En Chang, and will be written up in detail in his thesis [124].

6.6 Comparison of simulation and experiment

The simulated probability distributions at selected points across the BKT transition characterised by the measured algebraic exponent are displayed in Figure 6.6. For the OU simulation, we generate 1000 realisations of the fluctuating phase and evaluate the integral in 6.3 over a region which corresponds to the thickness of the repumping light sheet along the direction of the imaging beam and the effective pixel size perpendicular to the imaging beam. We employ a similar procedure for the MC simulation, but we evaluate the integral in 6.2 to account for the effects of density fluctuations. It is necessary to set the region of interest to match the experimental one as the distributions are expected to change with the aspect ratio of the integration region [14]. For example, integrating through longer distance along the imaging direction results in more peaked contrast distributions as the fluctuations are averaged out. The experimentally measured distribution at the lowest algebraic exponent displays a peaked distribution centered around $\tilde{v} \approx 0.75$. As the system exhibits stronger fluctuations, the observed distribution becomes flatter for normalised contrasts under unity. At and above the BKT critical point, the distribution decays approximately monotonically. The decay at lower values is due to the difficulty of resolving low contrasts from a noisy background. The simulated distributions show similar qualitative features of the distribution transforming into

a monotonically decaying, Poisson-like distribution which is a characteristic of thermal noise. However, the two simulations only agree well with each other deep in the superfluid regime, where the density fluctuations are strongly suppressed. At $\eta \approx 0.11$ and at the critical point, the MC simulation results in a broader distribution which is possibly due to the density fluctuations. Deep in the normal phase, we again find reasonable agreement between the two simulations. However, this is most likely to be accidental since the OU simulations are not expected to give accurate results in this regime since they always result in algebraic correlation functions while the actual correlation function is exponential.




Figure 6.6: Comparison of observed contrast distributions and results of MC and OU simulations. The experimental data shows a transformation from a peaked to a monotonously decaying distribution (apart from very low contrasts), which is qualitatively similar to the simulation results. At cold temperatures, the two simulations agree as expected due to the suppression of density fluctuations.

There are qualitative similarities between theory and experiment, but the comparison indicates that the theoretical distributions tend to underestimate low normalised contrasts and give more peaked distributions deep in the superfluid regime. A possible reason for this could be additional effects which reduce the contrast which were ignored in the simulations. These include the finite imaging resolu-

tion, the finite condensate fraction or experimental noise sources, such as imaging noise or imperfect splitting. The high number of repetitions gives us accurate noise statistics and we can predict their influence on the contrast distribution. To appropriately account for these effects, we have to ensure that the signal to noise ratio is identical in simulation and experiment. Therefore, given that the measured contrast noise is characterised by the standard deviation σ and the ratio between the simulated and measured mean contrast is r , we add a Gaussian random variable of standard deviation σr to the simulated contrast values. Since the white noise level σ was chosen as the lower bound for postselection based on contrast, we discard simulated contrasts which are below σr . We compare the four selected measured distributions to the MC result in Figure 6.7, where the simulated data including the effect of the rescaled noise is plotted with solid lines. For comparison, we plot the simulated distributions analysed the same way but without rescaling the noise with dashed lines. Deep in the superfluid regime ($\eta = 0.05$), the agreement between experimental data appears to agree well with the simulation that includes rescaled noise. As the system is closer to the critical point at $\eta = 0.12$, the simulation reproduces the broad "peak" of the measured distribution. At and above the critical point, the simulation matches well the long tail of the measured distributions but as before, the counts in the first few bins are often underestimated. This discrepancy is possibly because of the increased signal to noise ratio of the experimental data, as indicated above the subplots. Deep in the normal regime, the measured white noise level is slightly more than half the mean contrast, therefore the measured distribution is noise-dominated. In this case, the simple scaling of noise to match the signal to noise ratios between experiment and simulation fails.

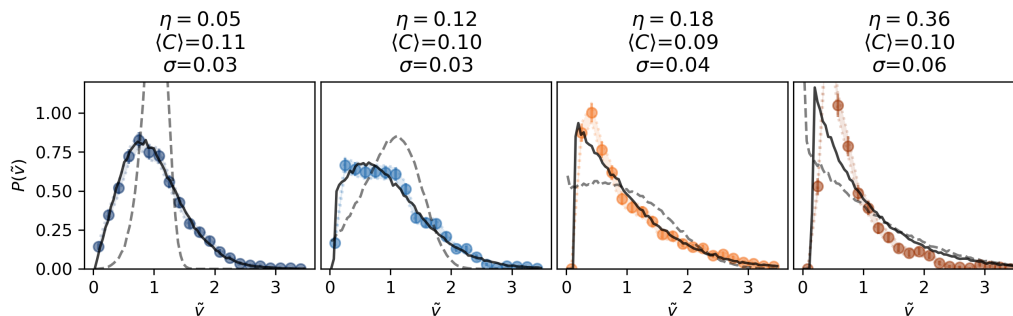


Figure 6.7: The observed contrast distributions (data points) and the simulated contrast distributions including noise with (solid lines) and without (dashed lines) rescaling.

6.7 Universal properties of contrast statistics in 2D

In the limit of large regions of interest (significantly larger than the interparticle spacing), the measured interference contrast is directly related to the number of condensed atoms. For three-dimensional Bose-Einstein condensates at the critical point, it has been shown that the full counting statistics of condensate fraction obeys a universal non-Gaussian distribution [125]. This can be extended to the critical regime for finite size systems using scaling relations [126]. For 2D systems, while it is not possible to use such scaling relations, one expects a universal distribution at the critical point. Furthermore, numerical simulations in the superfluid regime revealed that the contrast distribution converges to the convolution of two identical Gumbel functions as the temperature is reduced [14]. To compare this prediction to the experimental results, it is convenient to introduce an auxiliary variable

$$u = \frac{\pi L}{\eta}(1 - C^2), \quad (6.4)$$

where L is the system size and η is the algebraic exponent which we measure simultaneously through the analysis of correlation functions. We fit the observed

distributions of the auxiliary variable u with the model

$$y(x) = Ag_{\alpha,\beta}(x) \otimes g_{\alpha,\beta}(x) \quad (6.5)$$

where

$$g_{\alpha,\beta}(x) = e^{-(x-\alpha)/\beta - e^{-(x-\alpha)/\beta}} \quad (6.6)$$

is the Gumbel function and A is a normalisation factor. In Figure 6.8 a we plot the sum of squared residuals as a function of the measured algebraic exponents. While the sum of squares drops in the normal phase with decreasing η , we observe a much sharper drop as the critical point is crossed, indicating rapid convergence to the Gumbel convolution model. Since the change in the behaviour of the sum of squares is significant, this method could provide alternative means in identifying the BKT critical point. In Figure 6.8 b, we show the distribution $p(u)$ for the dataset with the lowest η as well as the fitted curve, showing good agreement, as expected from the low sum of squares. At each value of the algebraic exponent, we show the fit parameters α and β in Figure 6.8 c.

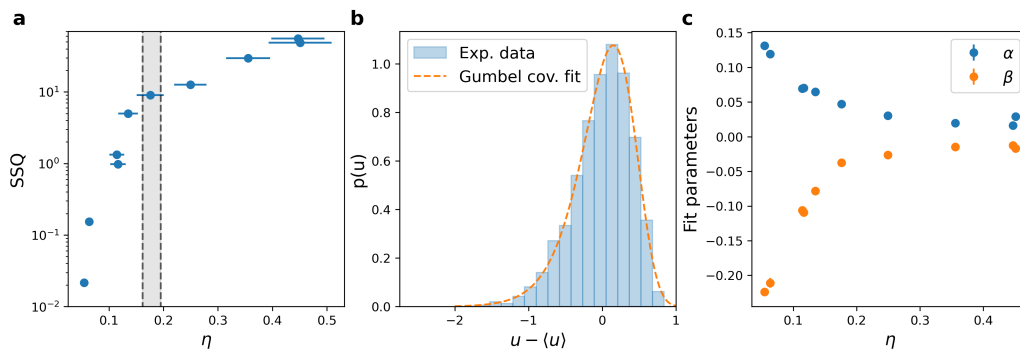


Figure 6.8: **a**: Sum of squared residuals (SSQ) for the Gumbel convolution fit of the observed distributions for datasets at various values of η . **b**: Observed distribution with Gumbel convolution fit for the lowest value of η . **c**: Fitted parameters of the Gumbel convolution model as a function of η .

6.8 Moments of contrast distribution

We further analysed the measured contrast statistics data by calculating the central moments up to the fourth order for the observable \tilde{v} . The central moment of order n for variable \tilde{v} for N samples is defined by

$$m_n = \frac{1}{N} \sum_{i=1}^N (\tilde{v}_i - \langle \tilde{v} \rangle)^n \quad (6.7)$$

where $\langle \tilde{v} \rangle$ is the arithmetic mean of the N samples of \tilde{v} . On a bounded interval, the collection of all moments uniquely defines a probability distribution. Therefore, they can capture the functional change of the contrast distribution across the BKT transition and have the ability to identify the critical point. Moreover, it can be shown that the moment of order n is related to the $2n$ -point correlation function [44]. This analysis method is advantageous as it does not rely on assumptions such as the choice on binning when calculating the discrete probability distributions or comparing with computationally expensive simulations. While it has been confirmed that simulated data fits well the convolution of two Gumbel distributions [14], there exists no theoretical or simulation result of how the fit parameters depend on the system parameters such as η or r_0 . The long tails of the distributions contribute with more weight to the higher-order moments, thus it is important to set a meaningful cutoff to the values of \tilde{v} . High values of this quantity occur rarely and they are more affected by Poissonian shot noise. Therefore, we set the upper limit of \tilde{v} to be 5.0 for both the experimental and simulated data, which includes all the meaningful information as can be seen from the measured distributions (see Figures 6.6 and 6.7). The moments of order 2-4 are plotted as a function of the algebraic exponent from the correlation analysis in Figure 6.9 a. As the gas crosses to the normal phase, the values of the moments increase until the critical

point is reached and then remain approximately constant. This is not surprising as the functional form of the distribution of \tilde{v} is expected to stay Poissonian above the critical point which is characteristic to thermal noise [115, 118]. The same moments are calculated for the contrast distribution from the MC simulations, without adding additional noise, and shown in Figure 6.9b. These simulation results also show the same increase and saturation of moments, however, the jump upwards is sharper compared to the experimental values which we attribute to the lack of technical noise.

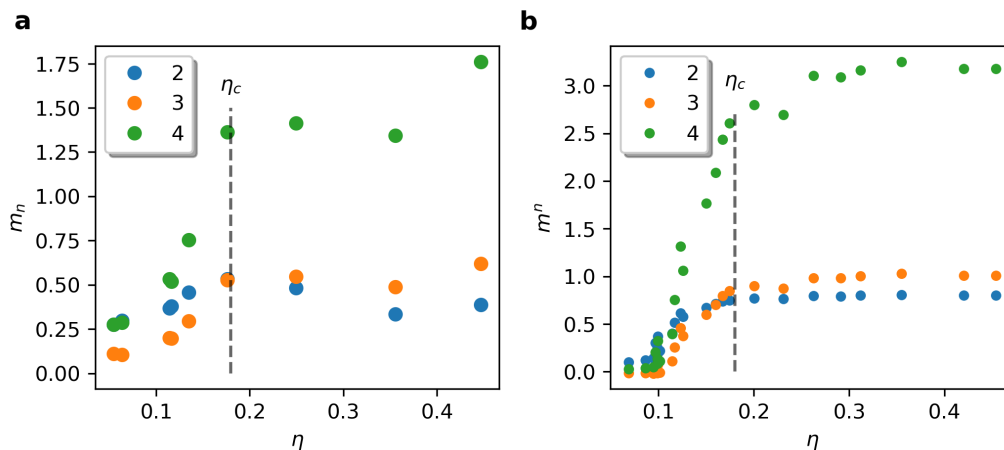


Figure 6.9: Analysis of moments of contrast statistics. **a**: Experimental results. **b**: MC simulation results without adding noise. The dashed line indicates the critical exponent from correlation function analysis. The moments on the simulated data show a rapid increase as the critical point is approached from the superfluid phase, indicating a change in shape of the contrast distribution. The experimental data shows qualitatively similar features.

6.9 Conclusion and outlook

In conclusion, we have performed comprehensive measurements of correlation functions and full counting statistics of interference contrast in a bilayer 2D Bose gas system across the BKT critical point. Our findings revealed that the transfor-

mation of the shape of the distribution function of interference contrast is qualitatively similar to theoretical predictions. We modelled the contrast statistics with MC and stochastic simulations and highlighted the effect of additional noise sources which affect the experimental data. The resulting simulated distributions were in good agreement with the observed distributions. We found that with increasing superfluid phase space density, the experimental distributions converged towards the convolution of two identical Gumbel functions as predicted by theory. Finally, we analysed the moments of the experimental distributions and found that the higher order moments saturate as the system crosses the BKT critical point from the superfluid to the normal phase. This is an important step to understand non-equilibrium quantum dynamics of quantum fields in 2D that is not necessarily characterised fully by simple two-point correlations that we access by the standard phase correlation analysis [127].

7 | Superfluid transition in disordered systems

In this chapter, we investigate the effect of speckle disorder on the superfluid transition in bilayer 2D Bose gases. We use matter-wave interferometry to measure the coherence of the relative phase at various strengths of the disorder potential. The transition between a quasi-ordered and a normal phase is identified by analysing the functional form of the measured correlation functions. We compare the critical disorder strength to predictions of numerical simulations based on classical and quantum field theories. The experimentally measured value of the critical disorder strength agrees well with the quantum field prediction, suggesting strong non-classical effects. We introduce an effective coupling model to compare the scaling of measured quantities at different disorder strengths to theoretical results in the thermodynamic limit using RG theory.

7.1 Generation of disordered optical potential

There are two main experimental schemes being used in ultracold atom experiments to realise a disordered optical potential. The first method involves directly imaging a DMD with pixels turned "on" at random locations [76]. This method creates point-like disorder which has the advantage of independent control over the strength and the shape of the potential. However, it requires accurate optical alignment and is fundamentally limited by the imaging resolution. The second common method is shining laser light through an optical diffuser which generates a random optical speckle pattern [128]. The case of 2D gases is particularly interesting since the lack of the third dimension results in a significant effect even in the presence of weak speckle disorder. Speckle patterns also result from imperfections in optical interfaces for example reflections from a mirror or transmis-

sion through a lens. Therefore speckle patterns with strength proportional to laser power are unavoidable when creating an optical box trap. While these inhomogeneities have no significant effect in typical cold-atom experiments, by increasing the laser power beyond what is required for realising the trap, it is possible to subject the trapped atoms to a disordered potential. This has been exploited in [129] to investigate far-from-equilibrium systems in the presence of weak disorder. These inhomogeneities in the optical potential were also present in other experiments as in Chapter 6 or in [130]. However, those experiments only require the hard wall potential to be at most a few times higher than the chemical potential, such that the potential coming from the speckle pattern inside the box is significantly weaker than any relevant energy scale. We have enough laser power to realise a disordered potential without having to make any changes to the optical setup.

7.2 Experimental protocol

The experimental protocol used is mostly identical to the one used in the contrast statistics experiment described in Section 6.2. We use radio-frequencies of 7.1, 7.15, and 7.2 MHz and a quadrupole gradient of 179 G/cm (along the weaker directions). At the final state, the wells are separated by $2.6 \mu\text{m}$ with a barrier height of 4 kHz. Under these conditions, the tunnelling rate between the two-wells is not negligible, we estimate that this causes a shift in the BKT critical temperature of approximately 30 percent based on experimental results in the same system and MC simulations [130]. We use numerical methods to calculate the MRF potential as in the previous chapter. The axial trapping frequencies are 1400(100) Hz and 1260(90) Hz in the upper and lower wells respectively, therefore the gases are in the 2D regime in both wells for the typical temperature of 32 nK used in the ex-

periment [121, 120]. The difference in the axial trapping frequencies results in a small, 5 percent difference in the 2D interaction strengths between the two gases which does not significantly affect the results. The radial trapping frequencies are 11.6(2) Hz and 11.3(2) Hz respectively. Similar to the experimental methods in the previous chapter, the splitting is performed over 0.45 seconds and the system is left to equilibrate for 0.3 seconds before time-of-flight expansion and detection. To ensure that the system is in equilibrium in the disordered potential, the green light is turned on adiabatically over 0.8 seconds to the desired laser power when the atoms are in a single-well RF trap such that the strength of the optical potential remains constant for 0.25 seconds before the splitting procedure starts. We quickly turn off the optical potential 2 ms before the release to ensure that it does not affect the atoms during time-of-flight. Before imaging, we repump the atoms to the $|F = 2\rangle$ level by a light sheet of thickness $4 \mu\text{m}$ so that we only image atoms in this thin slice to probe the relative phase fluctuations locally.

7.3 Estimating the strength of disorder

To estimate the strength of the disorder potential, we first measure the height of the ring potential by the methods used earlier in the research group [76]. This relies on splitting a single pancake-shaped cloud of known chemical potential into two halves using the optical potential and observing the decay of atomic density between the two daughter clouds at various laser powers. The strength of disorder relative to the hard walls has been reported to be on the few percent level and it is expected to scale inversely with box size [131, 132]. This disorder strength and dependence of box size was verified on a conceptually similar DMD setup which contains a single telescope. However, the disorder strength used in the experiment is expected to be higher because of the larger number of imperfect optical sur-

faces used compared to the simplified setup¹. To obtain a more reliable value for the strength of the disorder, we prepared a single pancake-shaped 3D gas close to the dimensional crossover. This results in a thin sample with strongly suppressed density fluctuations. We then set the green laser power to give wall heights of 2μ and 16μ and recorded in-situ images with the vertical imaging system. The effect of the hard walls is not expected to change significantly between these two values while the strength of the speckle potential changes by a factor of 8. We compared the mean density profiles averaged over 30 realisations of the two different potential heights and found that the mean speckle potential is 10 percent of the height of the hard walls. This measurement suffers from a systematic error due to finite imaging resolution and the healing length of the Bose gas as both effects reduce the magnitude of density fluctuations, hence the strength of the measured potential. The healing length is $1 \mu\text{m}$, while width of the point-spread function of the imaging system is $1.6 \mu\text{m}$ [133], therefore we expect the contribution of the imaging system to be dominant. To estimate the magnitude of this systematic error, we convolved 1000 realisations of the disorder potential with correlation length of $5 \mu\text{m}$ with a Gaussian function of width 1.6μ , and observed 17 percent decrease in strength.

7.4 Effects of optical potential on bulk parameters

Despite being far-detuned, the 532 nm green laser light used to realise the box potential may cause heating or atom loss due to photon scattering and imperfect pointing stability. These effects would systematically shift the experimental conditions and the observables. To investigate the magnitude of these shifts, we performed thermometry and atom number measurements at various intensities of the

¹There are 10 lenses and 9 mirrors along the beam path

green light, with repumping of the entire system before imaging. The measured temperatures and atom numbers are shown in Figure 7.1 as a function of the disorder strength. Both the temperature and atom number stay within 10 percent of the mean values of 32 nK and 21000 respectively, a repeatability which is typical of cold-atom experiments [77]. The corresponding mean phase space density of 11.5(2), while the critical phase space density is expected to be $\mathcal{D}_c^0 \approx 7.5$ in this tunnel-coupled bilayer system [130]. This gives a critical temperature of 49.8(9) nK and a critical atom number of 13700(200) as shown on the plot. However, similar to the experimental protocol in Chapter 6, the thermometry measurement was performed after a slow rampdown of the optical potential such that before time-of-flight, the atoms are confined in a well-known harmonic potential and therefore they suffer from the same systematic effects as discussed earlier. In this case, not only the ring potential but also the disorder potential results in an adiabatic compression of the trapped gas which can alter the measured temperature. To investigate this effect, the temperature of the system was calculated for images from matter-wave interferometry with selective repumping as well. For these experiments, the optical potential was only switched off 2 ms before time-of-flight, therefore these measurements do not suffer from adiabatic compression effects. In this case, only the wings of the measured atomic density were fitted with a polylogarithmic function (the spatial cutoff for the fit was 1.5-times the Thomas-Fermi radius about the peak density). The temperature extracted from the fit result is plotted in Figure 7.1 as green circles, showing good agreement with the thermometry results without selective imaging (blue circles) and no significant heating effect due to compression. The results from selective imaging consistently predict a lower temperature compared to the thermometry results. We attribute this to the fact that the region of interest for imaging might have been slightly misaligned from the centre of the cloud due to the discretised nature

of the available positions on the DMD used to create the repumping light sheet. Despite the systematic errors of these measurements, we can conclude that the temperature and atom number stay sufficiently constant for range of green laser power employed and stay sufficiently far from the superfluid transition. While fluctuations in these quantities might affect the quantities of interest such as the algebraic exponent, we expect them to be small compared to the measured trends.

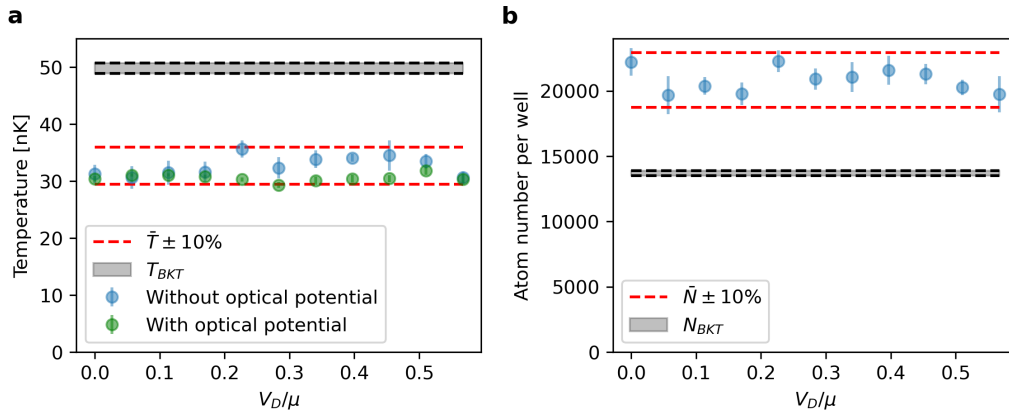


Figure 7.1: Results of thermometry and atom number measurements at various strengths of the disorder potential. **a**: Temperature measurements without the optical potential and imaging the entire system (blue circles) and with the optical potential and selective imaging (green circles). **b**: Atom number in each potential well (assuming equal split) without the optical potential, including the atoms which are outside the ring potential in experiments. Shaded regions show the expected critical values of these parameters.

7.5 Correlation function analysis

We used the same phase extraction and correlation function analysis procedure as in Chapter 6 to extract the algebraic exponent η and correlation length r_0 from the experimental images. We identified the critical point by performing reduced chi-squared analysis. The results are shown in Figure 7.2. The chi-squared analysis predicts a crossover between the algebraic and exponential models at a disorder

strength $V_D/\mu = 0.29(2)$. At this value, the algebraic exponent's interpolated value is $0.15(3)$ which is 15 percent lower than the critical exponent measured in Chapter 6. Similar to the earlier results, the correlation length r_0 deviates from the system size as the critical point is crossed from the superfluid to the normal phase. As the behaviour of the correlation function is qualitatively similar to other measurements (Chapter 6 or [12]), the results suggest that our disordered system is still within the BKT universality class. However, as the critical point is crossed, although there is an increase in the rate of change of the algebraic exponent, the trend appears to saturate at higher disorder strengths instead of diverging. The saturation is also present on the measured correlation length. There are multiple effects which could cause saturation. Firstly, the saturation occurs around $V_D/\mu \approx 0.45$. While the mean value of the disorder is less than the chemical potential, there must be regions where the disorder potential is comparable to or higher than the chemical potential and can no longer be considered weak, which may modify the scaling of observables. Secondly, strong disorder results in regions in the trap where the density is increased significantly, limiting the divergence of the algebraic exponent and the decay of the correlation length.

Using the QMC simulation results of Carleo *et al.* [15], the scaling of T_c with disorder strength (2.22) gives a critical value of

$$\frac{V_D^{crit}}{\mu} = 1 - \frac{T_c}{T_c^0} = 0.36, \quad (7.1)$$

where we used the estimated critical temperature and the system temperature from the thermometry measurement. This value is close to the observed value of $0.29(2)$. The lower value can be explained by the systematic error in the calibration procedure of the disorder strength as discussed earlier. A similar calculation based on classical field theories (2.21), yields a critical disorder strength of 0.69




Figure 7.2: Identification of the BKT critical point by correlation function analysis. **a:** Comparison of reduced chi-squared values for the two fit models. **b:** Algebraic exponent as a function of disorder strength, shaded regions indicate crossover. **c:** Correlation length of the exponential model as a function of disorder strength. Dashed line indicates system size and its variations are within the typical measurement uncertainty.

for our system. Therefore, the experimental results suggest classical theories are not sufficient to explain the effect of disorder on a 2D Bose gas, and quantum fluctuations play a significant role.

7.6 Vortex number analysis

As vortices play a critical role in the BKT phase transition, detecting these topological defects provides complementary information to the correlation function analysis. We use the methods described in [74] for vortex detection which is based on counting sharp discontinuities corresponding to $\approx \pi$ phase shifts in the matter-wave interference pattern. Figure 7.3 a shows an absorption image displaying a vortex near the centre of the cloud, marked by the red dashed line. It has been shown by RG and MC methods [32], and confirmed by experiments [12] that the number of vortices N_v scales exponentially with phase-space density as $N_v \propto \exp(-\gamma\mathcal{D})$ where the decay constant γ is approximately 0.5. Since matter-wave interferometry probes the relative phase between the clouds, the de-

cay constant is expected to decrease with tunnel coupling strength J_{\perp} in bilayer systems, as large values of the relative phase carry a high energy cost as shown by recent experiments [130]. To obtain the decay constant from the results in a disordered system, we introduce the effective phase space density based on 2.22 as

$$\mathcal{D}_{\text{eff}} = \mathcal{D}_{\text{exp}} \left(1 - \frac{V_D}{\mu} \right), \quad (7.2)$$

giving consistent results at zero disorder strength and at the critical point. In this picture, the critical point is not shifted by disorder but other quantities such as temperature or phase space density take effective values. We normalise the number of detected vortices by the number of columns the vortex counting was evaluated at, giving the vortex detection probability P_v which we plot as a function of the effective phase-space density in Figure 7.3 b. We fit the observed vortex density with the exponential decay as in [130]. We restrict the fit for $\mathcal{D}_{\text{eff}} > 6.5$, corresponding to $V/\mu < 0.43$, where we observe exponential scaling. This coincides with the regime where we have observed the expected scaling of the algebraic exponent and correlation length. Since $D_c^0 \approx 7.5$, the fit range covers a wide range of effective phase-space densities across the critical point. We find that the exponential model fits the data well with a decay constant $\gamma = 0.9(1)$. Experiments with similar parameters, but in absence of disorder resulted in $\gamma \approx 1.1$ for a well separation of $3.0 \mu\text{m}$ [130]. Direct comparison is not possible for a number of factors, the primary being the presence of disorder and the introduction of \mathcal{D}_{eff} from the linear scaling of the critical temperature. Others include the temperature, 2D interaction strength or the shape of the trapping potential being different in the two experiments. However, the functional form and decay constant of the detected vortex number with effective phase space density provides further evidence that the superfluid transition in this system is in the BKT universality class for weak disorder.

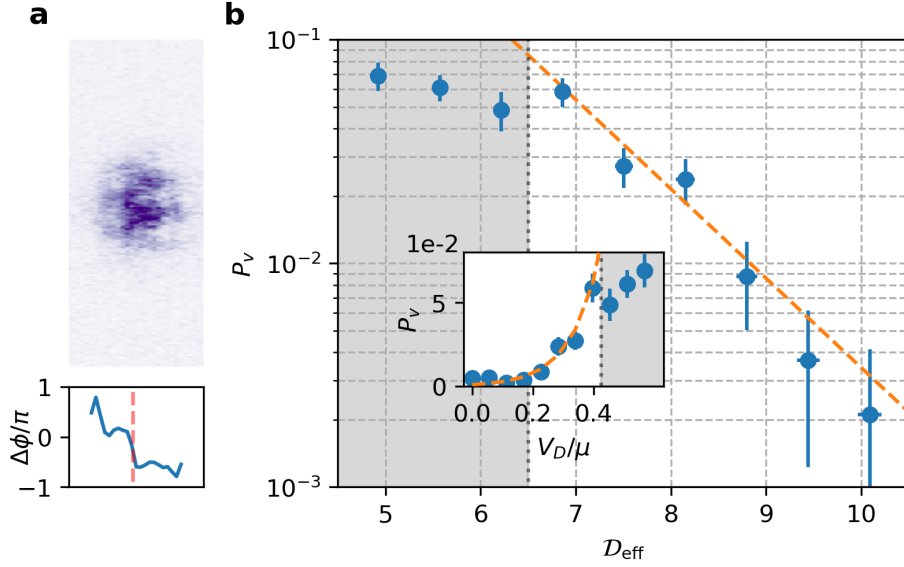


Figure 7.3: **a**: Absorption image of interference pattern with a vortex near the middle of the cloud (top). Extracted phase profile of the image showing a sharp discontinuity around the red dashed line, indicating the presence of a vortex (bottom) **b**: Vortex detection probability as a function of the effective phase-space density with fitted exponential decay. Inset shows the same quantity on a linear scale as a function of disorder strength.

7.7 The effective coupling model

While theoretical models, based on either classical or quantum field theories, predict the shift of the superfluid transition critical temperature, their results do not include the scaling of experimental observables such as η across the critical point. Inspired by the linear scaling of the critical temperature with disorder strength in 2.22, we heuristically absorb all effects of disorder into the superfluid coupling constant and refer to this as the *effective coupling model*. According to the RG treatment of the 2D XY-model, as well as free-energy considerations, the critical temperature is proportional to the spin-spin coupling strength J . Therefore, by introducing an effective value J_{eff} we modify the critical temperature such that

$T_c/T_c^0 = J_{eff}/J$. This is similar to the effective value of J used in initial points for RG flows as the superfluid coupling is temperature-dependent for a 2D Bose gas through the renormalisation of superfluid density in the effective low energy Hamiltonian (2.12). Therefore, we keep temperature constant and model the effect of disorder as the renormalisation of superfluid coupling, and assume that the long-range behaviour of the system is not affected by other mechanisms.

This model allows us to extend our analysis to the scaling of measured algebraic exponent across the critical point. We numerically integrate the RG equations that describe the coupled bilayer system (2.28 - 2.32) and extract the algebraic exponent from the final value of the superfluid coupling constant of the relative mode. We set the initial symmetric and antisymmetric superfluid couplings to $J_{eff} = \frac{\hbar^2 n}{2m} (1 - V_D/\mu)$. We use the vortex core energy expression from [52] and set all three vortex fugacities accordingly, allowing for a small common scaling factor as these quantities are not strictly defined in the case of a 2D Bose gas. The approximate value of the tunnelling energy J_\perp is found by calculating the algebraic exponent for small disorder strength where the system is still in the superfluid phase. In this regime, the results depend weakly on the choice of the vortex fugacity scaling factor as the RG flow drives the system towards high vortex core energies. We allow another scaling factor of order unity to change the tunnelling energy. We then find the optimal values (within reasonable bounds) of these two scaling parameters which describe the experimental data. To ensure that classical trapping does not significantly influence the measured algebraic exponent, we only consider disorder strengths up to $V_D/\mu = 0.35$. We compare the experimental data and the RG results in Figure 7.4, showing good agreement between theory and experiment for weak disorder. Apart from the effects of strong disorder, the increase in η could be affected by residual non-uniformity of the system which increases the width of the phase transition. The RG calculation was repeated with

identical scaling parameters but with effective coupling values which correspond to the classical field result (2.21). For the investigated range of disorder strength, the system remained far from the BKT critical point without a significant increase in the algebraic exponent.

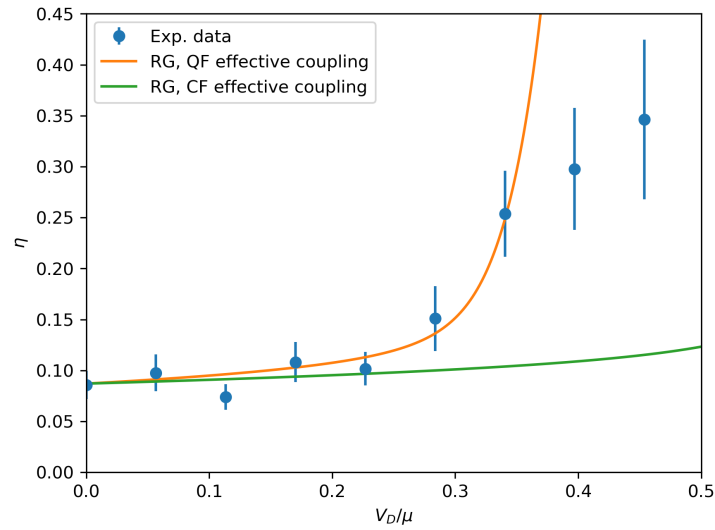


Figure 7.4: Comparison of measured algebraic exponent and RG results using the quantum field and classical field effective coupling models.

7.8 Simulating the effect of disorder on Bose gases

To guide the interpretation of the experimental results, we performed various classical field MC simulations to model the increase in the algebraic exponent. Firstly, we simulated a gas in a circular box trap of radius $25 \mu\text{m}$ in the presence of a disordered potential. These simulations were performed by Dr. Vijay Singh (TII, Abu Dhabi). We model the disorder potential $V_d(\mathbf{r})$ by a zero-mean isotropic speckle pattern with Gaussian autocorrelation

$$\overline{V_d(\mathbf{r})V_d(\mathbf{r}')} = V_D^2 e^{-|\mathbf{r}-\mathbf{r}'|^2/2l_c^2} \quad (7.3)$$

as often assumed in theoretical works [134], where l_c is the correlation length of the disorder and we refer to V_D as the disorder strength. This expression for the disorder correlation function holds if the speckle pattern arises from a large number of random phase shifts [135], which is true in the experiment. We generate $V_d(\mathbf{r})$ by multiplying white noise with the Fourier transform of the autocorrelation function (7.3), inverse Fourier transform the product and finally take the real part. Since the gas is confined in a finite-size box trap, the mean of V_d is not necessarily zero. This results in an atom number fluctuation between different realisations of the disorder. To correct this, we calculate the mean of the disorder potential inside the box and subtract this value from V_d . Furthermore, we carefully adjust the chemical potential to ensure that the number of atoms is the same at different disorder strengths. We calculate the phase correlation function from the simulation results, and estimate the uncertainty of the algebraic exponents by bootstrapping over the raw simulation data. We repeated the simulations for a range of different disorder correlation lengths. For a disorder correlation length of $5\mu\text{m}$ and above, we obtained identical $\eta(V_D)$ trends (within the error bars). For significantly lower disorder correlation lengths such as $1\mu\text{m}$, the disorder potential had no significant effect on the system which demonstrates the breakdown of the LDA. We compare the results of the simulation with a disorder correlation length of $10\mu\text{m}$ with the effective coupling model in Figure 7.5. The MC results are in qualitative agreement with the RG prediction, showing no significant increase in the algebraic exponent over the investigated range of disorder strengths ($0-0.35\mu$).

To benchmark the bilayer RG effective coupling model result, we performed simulations in disorder-free systems within the effective coupling model by introducing a scaling factor J_s to the kinetic energy term of the Hamiltonian. The advantage of this approach is that it does not rely on non-universal scaling constants. Since the RG results include coupling between the layers, we simulate a system of two

classical fields ψ_1 and ψ_2 which are described by the 2D Bose gas Hamiltonian independently and interact through the coupling Hamiltonian

$$\mathcal{H}_{int} = -t \sum_{\mathbf{r}} \psi_{1,\mathbf{r}}^* \psi_{2,\mathbf{r}} + c.c. \quad (7.4)$$

where \mathbf{r} denotes lattice sites. This coupling term is equivalent to 2.25 but it is more practical for numerical simulations. We set the temperature and chemical potential to match the experimental conditions at $J_s = 1$. Furthermore, we set the interlayer hopping t such that the simulated algebraic exponent matches the experimental value at zero disorder strength. We let the system equilibrate for 250000 updates per site per layer and successively sample the fields after 1000 updates per site per layer afterwards. The value of superfluid coupling weakly affects the number of atoms in the system, therefore we adjust the chemical potential at each value of J_s to ensure the accuracy of the simulation. We analyse the results using the same methods as the experimental data following phase extraction and we compare the simulation to the RG and experimental results in Figure 7.5 (orange circles). The MC results follow the RG trend qualitatively for the range of investigated disorder strengths and demonstrate a sharp increase of the algebraic exponent as the critical disorder strength is crossed. These results demonstrate that the non-trivial scaling parameters used in the RG calculations were chosen within reasonable bounds.

7.9 Conclusion and future plans

We have performed measurements on the effect of speckle disorder on bilayer 2D Bose gases and observed a crossover from the superfluid to the normal phase with increasing disorder strength. The critical disorder strength was compared with classical and quantum field predictions and we found better agreement with the latter. This result was further corroborated with the RG theory of bilayer gases




Figure 7.5: Comparison of MC simulations of a 2D Bose gas in a disordered environment and the experimental results.

within an effective coupling model as well as classical field MC simulations.

To understand the effect of speckle disorder on 2D Bose gases in more detail, the following changes will be made to the experimental methods in the future. Firstly, it is desirable to perform the measurements with larger separation between the layers to eliminate the interplay of coupling and disorder. Secondly, a two-dimensional sweep of disorder strength and temperature would allow precise measurements of the shift in the critical temperature. Finally, the control over the speckle disorder could be improved by creating it independently of the optical ring potential, by diffusive plates.

The results obtained so far, as described in this chapter, show that our interferometric method gives a powerful way to investigate the effect of disorder on BKT physics in 2D quantum gases, and we have obtained clear evidence for a significant shift in the critical temperature. This testing of the experimental methodology and its validation by comparison with theoretical modelling is a key step towards

gathering a larger set of data to provide an even more comprehensive understanding of this system.

8 | Dynamical control of bilayers across the BKT transition

In this chapter, we show the results of three proof of concept experiments in which we quench bilayer 2D Bose gases across the BKT critical point in three different ways. In the first experiment, we realise a quench from the superfluid phase into the normal phase by reducing the tunnel coupling between the clouds. We observe the decoherence dynamics and show that the evolution over intermediate timescales is consistent with the universal scaling supported by theory and our previous measurement using the density quench with harmonically-trapped gases. Next, we show experimental results where we realise a quench by changing the strength of the speckle potential which is present in the trap. For a quench into the normal phase we show similar results to earlier experiments. For a quench into the superfluid phase, we observe a linear decay of the algebraic exponent until equilibration.

8.1 Tunnel coupling strength quench

As the BKT critical temperature in coherently coupled bilayers depends on the interlayer coupling (see Section 2.9), this can be exploited to quench the system across the critical point from the superfluid to the normal phase by increasing the separation between the clouds. This scheme is advantageous as it provides independent control of the initial and final state of the system (instead of halving the phase space density as in the case of quench by splitting). As 2D Bose gases are quenched through the BKT transition in this direction, the relaxation dynamics of the system exhibits reverse Kibble-Zurek phenomena. Such a process can be modelled by real-time RG theory [16], which showed that the relaxation is a two-

step process. On short time-scales, phonons equilibrate through a light-cone-like process, limited by the finite speed of sound in the gas. This is followed by the slow, dynamical proliferation of free vortices. The first experiment in which out-of-equilibrium dynamics of 2D Bose gases following a quench across the BKT critical point was performed in Oxford and realised the quench by splitting a single 2D system into two parallel subsystems, thereby halving the phase-space density [13].

8.1.1 Experimental methods

To realise the interlayer coupling quench experimentally, we prepare a bilayer system in an MRF dressed potential with radio-frequencies 7.1, 7.15, and 7.2 MHz and a quadrupole gradient of 179 G/cm (along the weaker directions). Similar to the experimental methods described in Chapter 6, we project a ring-shaped laser beam to realise a near-uniform system, which is turned on adiabatically, over 0.8 seconds. We perform splitting of a single cloud slowly, over 0.47 seconds, to avoid collective excitations. At the final stage of splitting, which we consider to be the pre-quench state, the well separation is $2.4 \mu\text{m}$ with a barrier height of 2 kHz. After finishing splitting, we let the cloud equilibrate for 0.3 seconds. We then rapidly (over 20 ms) reduce the interlayer coupling by changing the amplitudes of the dressing RF fields to increase the well separation to $3.1 \mu\text{m}$ and the barrier height to 4.5 kHz. We let the system evolve for various times. Similar to the experimental scheme described in earlier chapters, we turn off the ring potential 2 ms before releasing the atoms from the trap. We image the matter-wave interference pattern after 17 ms time-of-flight. Before imaging, we repump the atoms to the $|F = 2\rangle$ level by a light sheet of thickness $4 \mu\text{m}$ to ensure probing the local relative phase between the two clouds.

8.1.2 Relaxation dynamics after quench

We use the same relative phase and correlation function extraction methods as in the previous chapters to analyse the images. The correlation functions at each hold time are fitted with the algebraic and exponential functions characteristic to the superfluid and normal phases, and we use a reduced chi-squared test to compare the two models. The results of the statistical test are plotted in Figure 8.1 a, showing a crossover between the two models at a dynamical critical time t_c of 0.11(3) s. Despite the crossover, the reduced chi-square values for the algebraic model stay low even at longer hold times, therefore the algebraic exponent remains a meaningful parameter for characterising the phase fluctuations. From the time evolution of the algebraic exponent shown in Figure 8.1 b, we identify the critical exponent to be 0.12(2) by approximating the time-dependence of η with a cubic interpolant to smoothly connect the datapoints. This is in agreement with previous results for similar experiments with harmonic confinement in which the critical point was crossed by spitting a single cloud into two [13]. At subsequent times, the algebraic exponent increases monotonically until it reaches saturation (within measurement error) after 0.7 s. For a dynamic crossing of the BKT critical point from the superfluid to the normal phase, the system is expected to exhibit self-similar dynamics [16], which results in a universal linear scaling of the algebraic exponent with time, as indicated by the dotted line. At intermediate times (between 0.3 and 0.7 seconds), the data fits this behaviour well. At short times, the deviation can be explained by starting with a split system, with a relative phase correlation function characterised by a non-zero algebraic exponent. This is in contrast with our earlier experiments in which the pre-quench state was a single cloud, giving long range coherence of the relative phase at short times after the quench [13]. Finally, we show the time evolution of the correlation length which

characterises the exponential decay of the correlation function in Figure 8.1 c, displaying similar behaviour to the algebraic exponent. As in the case of the equilibrium measurements in Chapter 6, the BKT crossover can be observed by the correlation length decaying below the system size.




Figure 8.1: Relaxation dynamics following a reduction of interlayer coupling. **a:** Statistical test of the fitting with algebraic and exponential models. **b:** Time evolution of the algebraic exponent. **c:** Time evolution of the correlation length.

The self-similar relaxation dynamics results in a universal correlation function which depends only the rescaled spatial coordinate $x' = \bar{x} \times t/t_c$. In Figure 8.2, we plot the measured correlation functions up to 0.7 s as a function of x' and observe a collapse onto a universal function for $0.3 \text{ s} \leq t \leq 0.7 \text{ s}$, the range of times we expect self-similar dynamics.

8.1.3 Conclusion and outlook

In this section, we presented experimental results for the relaxation dynamics of a bilayer 2D system after a sudden reduction of the tunnel coupling strength. We observed a temporal change of the functional form of the correlation function from algebraic to exponential, demonstrating a quench into the normal phase. We showed that at intermediate times, the correlation function obeys a universal scaling relation, as in similar quench experiments. To further test the universality of relaxation dynamics, future experiments will perform the same quench at different




Figure 8.2: Correlation functions at various hold times as a function of the rescaled time, displaying a collapse onto a universal curve for $0.3 \text{ s} \leq t \leq 0.7 \text{ s}$.

initial conditions. This is an important step towards exploiting the tunnel coupling to tune or modulate bilayers across the superfluid transition.

8.2 Disorder strength quench

The disorder-induced loss of coherence of a 2D Bose gas, as shown in Chapter 7, can be exploited to make out-of-equilibrium measurements by rapidly changing the strength of disorder. We expect this to be helpful in understanding the scaling behaviour and universality class of the disorder-induced transition, as well as potential localisation properties by the looking for signatures of non-ergodic dynamics. Similar methods were used to probe the dynamical break-down and establishment long-range phase coherence of a cigar-shaped 3D Bose-Einstein condensate following a rapid change in the strength of an optical speckle potential [136]. In our system, such experiments are straightforward to realise as the disorder potential is independent from the double-well potential.

8.2.1 Experimental methods

The preparation of the bilayer system is identical to the equilibrium measurements in Chapter 7. We set the power of the green laser to maximise the phase coherence of the system, which happens at $V_D/\mu \approx 0.1$ for which the equilibrium measurement in Chapter 7 resulted in an algebraic exponent of $\eta = 0.07(1)$. After splitting the cloud and equilibration for 0.3 seconds, we ramp the laser power linearly to a higher value over 25 ms such that the strength of the disorder potential is $V_D/\mu \approx 0.40$. At this value, the equilibrium system was in the normal phase, characterised by an algebraic exponent of 0.29(5). We refer to this scenario as a "reverse quench" since the expected dynamics is the reverse of the Kibble-Zurek type. In a second set of measurements, we prepare the system to be in the normal phase with an algebraic exponent of 0.35(7) by setting the green laser power to correspond to $V_D/\mu \approx 0.60$. Similarly, we wait 0.3 seconds after splitting and decrease the laser power linearly to $V_D/\mu \approx 0.1$ to attempt to quench the system from the normal to the superfluid phase. In this case, the ramp is performed over 125 ms to avoid heating that would drive the system back to the normal phase. We refer to this experiment as "forward quench". In both experiments, after the quench, we let the bilayer system evolve for a variable amount of time. Similar to the other quench experiment, we turn off the optical dipole potential 2 ms before the 17 ms long time-of-flight and we use a light sheet of thickness $4 \mu\text{m}$ to selectively repump the atoms before imaging. We analyse the data using the same methods as described in the previous section.

8.2.2 Relaxation dynamics after a reverse quench

We show the results of the statistical test comparing the two fit models in Figure 8.3 a. The result displays a crossover between the algebraic and exponential

models at a dynamical critical time of 0.1(1) s. In Figure 8.3 b, we plot the time evolution of the algebraic exponent, which shows three characteristic regimes. Firstly, at short times (until 60 ms), the algebraic exponent stays constant. At intermediate times (between 60 ms and 150 ms), η increases approximately linearly. The superfluid transition happens in this regime at a critical exponent of 0.12(2), which is in agreement with the tunnel coupling quench result in the previous section. Finally, at long times, the algebraic exponent shows slow evolution and stays above the critical value. In Figure 8.3 c, we show the correlation length, which drops under the system size as the critical point is crossed.




Figure 8.3: Relaxation dynamics after a reverse quench. **a**: Statistical test of the algebraic and exponential models. **b**: Time evolution of the algebraic exponent. **c**: Time evolution of the correlation length. Dashed line indicates system size.

8.2.3 Relaxation dynamics after a forward quench

We show the results of the statistical test comparing the two fit models in Figure 8.4 a. The test shows a crossover between the reduced chi-squared values of the two fit models after 48(5) ms. After this time, the algebraic model describes the data better, except at a hold time of 80 ms. The time evolution of the algebraic exponent (Figure 8.4 b) shows an approximately linear decay over approximately 60 ms from 0.22(3) to 0.11(1). At later times, the algebraic exponent stays approximately constant, which suggests reaching the new equilibrium state. A slow

increase is present after 100 ms which we attribute to the finite heating rate in the trap which would eventually drive the system back into the normal phase. The time evolution of the correlation length (Figure 8.4 c) shows an initial linear increase from $20(1) \mu\text{m}$ to the system size, over 30 ms. After this time, the correlation length stays close to the system size. Using a cubic interpolation to smoothly connect between the data points, the dynamical critical time yields a critical exponent of $0.096(4)$, which is lower than the expected value from equilibrium measurements and the other, superfluid to normal quench results. Since the measured critical exponent is only slightly higher than the values measured at later times, this could explain the anomalous chi-squared values at a hold time of 80 ms. A possible explanation for the low critical exponent could be that to reach the superfluid phase in a quench experiments, phase fluctuations must become weaker from both phonon and vortex excitations. Since vortices are topological excitations, they either have to form closely bound pairs or leave the central, high-density region. Therefore, their relaxation is expected to happen on a longer timescale compared to phonons, as demonstrated in [13] for superfluid to normal quench experiments. To investigate this, further experiments are required to verify the relaxation of vortex excitations and the mechanisms behind this process.




Figure 8.4: Relaxation dynamics after a forward quench. **a:** Statistical test of the algebraic and exponential models. **b:** Time evolution of the algebraic exponent. **c:** Time evolution of the correlation length. Dashed line indicates system size.

8.2.4 Conclusion and outlook

In this section, we presented results of experiments in which we probed the out of equilibrium dynamics of bilayer 2D Bose gases following a disorder strength quench. In the reverse quench scenario, we observed a dynamical crossover to the normal phase and equilibration at long timescales. The measured critical exponent agrees well with our previous quench experiments. In the forward quench scenario, there were some unexpected observations. The time evolution of the measured algebraic exponent and correlation length demonstrated a dynamic crossover to the superfluid phase and the system reached a new equilibrium state. However, the critical exponent, identified by the statistical test of the fit models, was unexpectedly low compared to equilibrium and other quench experiments.

To investigate the evolution of our system following disorder strength quenches, we will carry out future experiments with more separated wells to eliminate the interplay between coupling and disorder. Also, performing the same experiments with various initial points across the BKT crossover would be advantageous for probing the universality of relaxation dynamics. These experiments can provide important insights into the role of disorder in 2D Bose gases both in and out of equilibrium that will guide the development of theories describing such systems.

9 | Creating uniform two-dimensional RF-dressed potentials

Uniform traps are beneficial for theoretical interpretation of cold atom experiments. However, most RF-dressing based trapping methods inherently result in harmonic confinement. In this chapter we propose an experimental scheme to realise uniform two-dimensional trapping potentials for cold atoms using the RF-dressing technique.

9.1 The advantages of uniform potentials

Uniform potentials have numerous advantages for confining ultracold atoms. Firstly, the Hamiltonian of uniform systems is inherently simpler, therefore theorists often assume quantum gases in a box, simplifying the interpretation of experimental data. In homogeneous systems, the phonon modes are cleaner and easier to observe as demonstrated in [137] and [138]. Phase transitions, such as the BKT transition, are easier to observe unlike in harmonic traps where significant fractions of the system can be above and below the critical point. Bragg spectroscopy in uniform traps yields high resolution measurement of the momentum distribution due to the lack of inhomogeneous broadening [139]. Box traps are also advantageous for exploring non-equilibrium phenomena such as the Kibble-Zurek mechanism, ensuring that the different parts of system contribute with equal weight to the choice of the order parameter [140]. Finally, in quasi 2D systems, a uniform in-plane potential which is independent from the transverse potential makes it possible to perform a two-dimensional time-of-flight measurement in which the superfluid component can be identified [141].

9.2 Realising uniform 2D potentials in a magnetic dipole trap

From equation 3.9 it is clear that the geometry of the RF-dressed trap depends on two main factors: the spatial dependence of the detuning from magnetic resonance, and the Rabi frequency. For common RF-dressed traps, the radial trapping geometry is determined by the shape of isomagnetic surfaces of the DC magnetic field. Therefore, by engineering DC magnetic field configurations that create a uniform magnetic field strength in the radial direction of the RF-dressed potential, it is possible to create a trap with a very high uniformity in the radial direction. One such configuration is the magnetic dipole field, which can be created by a single wire loop topology. In the next sections, we explore and model various coil geometries for this trap.

9.2.1 Magnetic field of a single wire loop

Figure 9.1 shows selected contours of constant magnetic field strength for a single wire loop while the background illustrates the magnetic field strength on a logarithmic scale. It is clear that with increasing distance away from the coil along its axis (downwards), the radius of curvature changes from negative to positive. Therefore, there must exist a contour with zero radius of curvature which is shown by the solid contour in Figure 9.1.

To find the distance along z where the contours are flat, we consider the magnetic field in the x - z plane

$$\vec{B}(x, y = 0, z) = B_x \vec{e}_x + B_z \vec{e}_z. \quad (9.1)$$




Figure 9.1: Isomagnetic surfaces of a single wire loop. The curvature of contours marked with dashed lines changes sign at a certain distance away from the plane of the coil, highlighted with the solid line, resulting in flat contours. Colour scale displays magnetic field strength on a logarithmic scale with arbitrary reference level.

B_z has a well-known expression along the axis of the current loop

$$B_z(z) = \frac{\mu_0 I}{2} \frac{R^2}{(z^2 + R^2)^{\frac{3}{2}}}, \quad (9.2)$$

where I is the current in the loop and R is the radius of the loop. For flat contours, we require that

$$\frac{\partial \vec{B} \cdot \vec{B}}{\partial x} = \frac{\partial^2 \vec{B} \cdot \vec{B}}{\partial x^2} = 0. \quad (9.3)$$

The gradient of the contours is zero due to the reflection symmetry of the system.

The curvature of contours on axis is given by

$$\frac{\partial^2 \vec{B} \cdot \vec{B}}{\partial x^2} = 2 \left(\frac{\partial B_x}{\partial x} \right)^2 + 2B_x \frac{\partial^2 B_x}{\partial x^2} + 2 \left(\frac{\partial B_z}{\partial x} \right)^2 + 2B_z \frac{\partial^2 B_z}{\partial x^2}, \quad (9.4)$$

where the second and third terms vanish. Using $\vec{\nabla} \cdot \vec{B} = 0$ and $\vec{\nabla} \times \vec{B} = 0$ we obtain

$$\frac{\partial B_x}{\partial x} = -\frac{1}{2} \frac{\partial B_z}{\partial z} \quad (9.5)$$

and

$$\frac{\partial^2 B_z}{\partial x^2} = -\frac{1}{2} \frac{\partial^2 B_z}{\partial z^2}. \quad (9.6)$$

Therefore for flat contours,

$$\frac{1}{2} \left(\frac{\partial B_z}{\partial z} \right)^2 - B_z \frac{\partial^2 B_z}{\partial z^2} = 0. \quad (9.7)$$

Using 9.2 we find that this happens at

$$z_0 = \sqrt{\frac{3}{7.5}} R. \quad (9.8)$$

While this result gives the correct distance from the plane of a wire loop where the contours are flat, calculating the RF-dressed potential requires the off-axis magnetic field, therefore we rely on numerical calculations and use the Magpylib package¹ [91]. To verify the analytical result, the magnetic field strength is calculated in a 10 mm wide region along the x -axis, in a 5 mm to 40 mm range along the z -axis, measured from the plane of the loop. For each z -coordinate, the magnetic field strength as a function of x is fitted with $B(x) = c_0 + c_2 x^2 + c_4 x^4 + c_6 x^6$. We consider the z -coordinate to be optimal if the quadratic coefficient c_2 vanishes and denote it as z_0 . This calculation is repeated for a range of coil diameters and z_0 and $|B(z_0)|$ are recorded. From $|B(z_0)|$ we calculate the RF frequency to create a trap at this location. The results are shown in Figure 9.2 b/c. The optimal z distance scales linearly with coil diameter with a gradient of 0.32, which is in agreement with 9.8. This is expected as the phenomenon is purely geometric since an ideal current loop is considered. This correlation remains valid as long as the conductor is thin compared to the loop diameter.

¹There is a known analytical solution for the off-axis magnetic field for a wire loop [142], however it involves elliptic integrals which are difficult to manipulate.

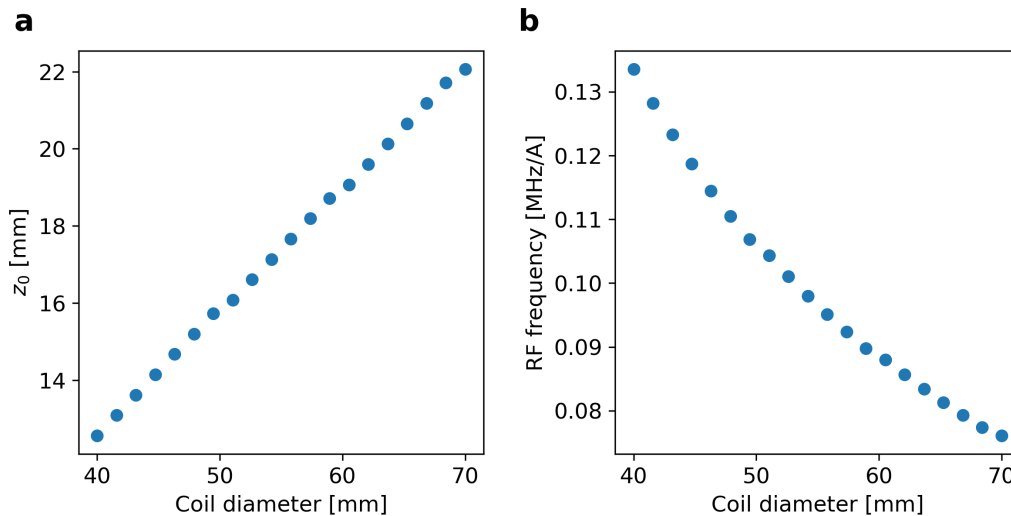


Figure 9.2: Numerical results for RF-dressed potential by a single wire loop. **a**: Optimal distance to realise uniform potential measured from coil plane as a function of coil diameter. **b**: RF frequency per unit current as a function of coil diameter to realise quasi-uniform potential.

Besides the isomagnetic surfaces, it is important to consider the confinement and gravitational sag along the z -direction. For a single wire loop the gravitational sag is more significant, and the trapping frequency along the z -direction is weaker due to the lower magnetic field gradient in this direction compared to conventional anti-Helmholtz coils. Realising a quasi 2D gas requires strong confinement along the z -direction to freeze out the excitations. This happens when the characteristic thermal energy $k_B T$ is less than both the vibrational energy spacing and the chemical potential [143]. For experimentally feasible temperatures, the trapping frequency must be at least 200 Hz which gives a temperature of 10 nK to be in the quasi 2D regime. This sets a lower bound for the magnetic field gradient which affects the current as well as the RF frequency required. With a 200 Hz axial trapping frequency and considering a system size of 50 μm , the chemical potential is 100 Hz for 100000 atoms, therefore the proposed trap is capable of producing quasi-2D gases considering the characteristics of the trap only. The magnetic field

strength and gradient at z_0 can be calculated from the on-axis magnetic field 9.2, giving

$$B_z(z = z_0) \approx \frac{0.30\mu_0 I}{R}, \quad (9.9)$$

and

$$\left. \frac{dB_z}{dz} \right|_{z=z_0} \approx -\frac{0.41\mu_0 I}{R^2}. \quad (9.10)$$

9.2.2 Quasi-uniform trap in absence of gravity

Firstly, we consider the case of a microgravity cold-atom experiment as there is no need to compensate gravity by a magnetic force. Examples of such an experiment include the Cold Atom Laboratory on the International Space Station [144] or the ICE experiment which can operate on earth or on a reduced-gravity aircraft [145]. In this case, from Eq. 3.9 the trapping frequency along z for the trapped dressed state is given by

$$\omega_z^2 = \frac{\hbar}{m} \frac{\partial^2}{\partial z^2} \sqrt{\Delta^2 + \Omega^2}. \quad (9.11)$$

Ignoring spatial variations in the Rabi frequency and ignoring the second derivative in Δ , the trapping frequency is

$$\omega_z^2 = \frac{\hbar}{m} \frac{(\Delta')^2}{\Omega}. \quad (9.12)$$

Combining Eqs. 9.10 and 9.12 yields

$$\omega_z = \sqrt{\frac{\hbar}{m\Omega} \mu_B g_F m_F \left(\frac{0.41\mu_0 I}{R^2} \right)}. \quad (9.13)$$

Considering $\Omega/2\pi = 20$ kHz which is a typical experimental lower bound due to non-adiabatic losses, we obtain $\omega_z/2\pi = 200$ Hz for a coil diameter of 2 cm current of 72 A, producing a uniform potential 6 mm away from the plane of the

coil². Furthermore, the required dressing frequency to achieve a uniform trap can be calculated from Eq. 9.9, giving 20.2 MHz. These values lie within the typical parameters of RF-dressed setups. Note that the trapping frequency can be improved further by a factor of $\sqrt{2}$ by using the $|F = 2\rangle$ hyperfine level which has $m_F = \pm 2$. There are two effects which modify slightly the axial trapping frequency. Firstly, with a dressing frequency of 20.2 MHz, the non-linear Zeeman effect must be taken into account which results in a weaker vertical confinement for $|F = 2\rangle$ and stronger vertical confinement for $|F = 1\rangle$ [146]. Secondly, we have neglected second derivative of detuning in 9.12, but this increases the trapping frequency, irrespective of which hyperfine level is being used.

9.2.3 PCB-based uniform trap in absence of gravity

The relatively low current needed in a single wire loop makes it possible to create a multi-turn PCB-based coil to produce required DC magnetic field. We consider a design with 8 turns, thus expecting a current around 10 A. Following the IPC-2221A standard, with an ambient temperature of 20 degrees and a maximum operating temperature of 60 degrees, using 5-oz copper thickness, the minimum trace width for the required current is 0.62 mm. For increased robustness, we opt for 0.9 mm wide traces separated by 0.2 mm which is about 2.5-times larger than the smallest recommended trace spacing. We model the coil as concentric circular current lines. With an innermost diameter of 15 mm the uniform region is produced 6 mm below the coil. A numerical trap frequency calculation yields 200 Hz for a current of 10.5 A, Rabi frequency of 20 kHz and dressing RF frequency of 22 MHz. An alternative is to use direct bonded copper substrates which are used

²The 20 kHz lower bound is for the old apparatus in a quasi-2d harmonic trap with $\omega_z/2\pi = 1\text{kHz}$. The non-adiabatic losses are expected to be smaller in the proposed trap as the characteristic velocity of atoms traversing the non-adiabatic crossing is smaller due to the lower axial trapping frequency.

in ex vacuo atom chips [147]. These have significantly higher current carrying capabilities compared to conventional traces on PCBs³, hence a lower number of turns can be used. For both cases, we can conclude that the proposed design is experimentally feasible.

So far, we have considered the Rabi frequency as a constant. However, the RF coupling between the Zeeman sublevels generally depends on the angle between the RF polarisation vector and the direction of the local DC magnetic field. For circular polarisation about the z -axis, the Rabi frequency is given by [65]

$$\Omega(r) = \frac{g_F \mu_B |B(r)|}{\sqrt{2} \hbar} (1 + u_z(r)), \quad (9.14)$$

where $u_z(r)$ is the z -projection of the unit vector which points along the DC magnetic field. This means that the RF-dressed potential will be slightly repulsive despite the flat isomagnetic surface near the trap centre due to the decrease in coupling strength. After calculating the RF-dressed potential including spatial modulations of the Rabi frequency, the RF frequency was tuned to minimise the quadratic coefficient of the radial potential. The necessary adjustment was small, the optimal value was 23 MHz which is about 5 percent different from the earlier value. For this coil geometry, it was not possible to achieve a vanishing quadratic coefficient in the potential. The lowest quadratic coefficient (in modulus) gave a 84 mHz radial trapping frequency, which is 2 orders of magnitude lower compared to the harmonic trap used on the old experimental apparatus. The resulting potential is displayed in Figure 9.3 as a function of x and z (left) on a logarithmic scale and as a function of x at the trap minimum. The variations in the radial potential are negligible compared to the chemical potential of a typical BEC. Furthermore, this is valid in an 1 mm wide region which is significantly larger than

³200 A for a wire of cross section 100 μm x 100 μm [148].

typical clouds which have a diameter of $50 \mu\text{m}$. This makes it possible to combine this technique with a repulsive optical dipole potential to provide hard-walls of the required shape and size. Finally, the reduction of the local coupling strength due to the non-uniform nature of the DC magnetic field was investigated. The Rabi-frequency was calculated at the trap minimum along z up to 1 mm radially from the trap centre. It was found that the Rabi frequency only reduced by 600 Hz which means that the spatial variation of non-adiabatic losses are suppressed in this trap.




Figure 9.3: RF-dressed potential created by the 8-turn PCB-based coil. On the left, the potential is displayed as a colour map on a base-10 logarithmic scale with a reference level of 1 Hz along the x - z plane. On the right, the potential is plotted along the x -direction around the global minimum on logarithmic (solid) and linear (dashed) scales, demonstrating good uniformity.

9.2.4 Quasi-uniform trap on Earth

A necessary condition for forming a dressed trap on Earth is that the magnetic force on the atoms must exceed the gravitational force, as represented by the following equation for the single-turn design:

$$\left| m_F g_F \mu_B \frac{dB_z}{dz} \right|_{z=z_0} \geq mg. \quad (9.15)$$

A practical lower limit for the ratio of these two forces is 1.5 to ensure stable trapping. For the previously considered single wire loop this ratio is only 1.2. The required ratio can be achieved by increasing the current by a factor of 1.25 to 90 A. The required RF frequency grows by the same factor, to 25.3 MHz. Operating at this frequency would require an RF coil with low self-capacitance which is achieved by reducing the number of turns. This results in a weaker RF-field, making it challenging to initially trap a cloud of atoms at a temperature of tens of microkelvin. Although a factor of 2 in the magnetic force can be gained by using the $|F = 2\rangle$ hyperfine ground states of ^{87}Rb which would make this scheme feasible, we prefer to work in the lower hyperfine level where the collisional loss is smaller. Therefore, we propose a different coil geometry to create quasi-uniform magnetic traps on Earth at lower RF frequencies. To reduce the magnetic field strength of the single wire loop while keeping the quadrupole gradient, one can use a uniform magnetic field created by coils in the Helmholtz configuration. The coil array consists of two coils of 6 cm diameter in the Helmholtz configuration and a single wire loop of diameter 2 cm as shown on Figure 9.4 a. Contours of constant magnetic field strength are displayed on Figure 9.4 b indicating an approximately flat contour at 17 G. Red markers show the directions of currents into and out of the page for clarity. For the magnetic field depicted in this figure a current of 150 A in the single coil and 93 A in the Helmholtz coil were used. The ratio of magnetic and gravitational forces is 1.75 for the $|F = 1\rangle$ hyperfine ground state, making this configuration practical. Finally, Figure 9.4 c shows the radial trapping frequency as a function of RF frequency at 3 different values of currents in the Helmholtz coil. Negative values indicate a repulsive potential. Unlike a single wire loop in absence of gravity, with this coil geometry, it is possible to create a trap with a vanishing radial quadratic coefficient if the RF frequency and Helmholtz current are chosen appropriately. Therefore, we can conclude that such

a trap is feasible on Earth to create a quasi-uniform Bose gas.

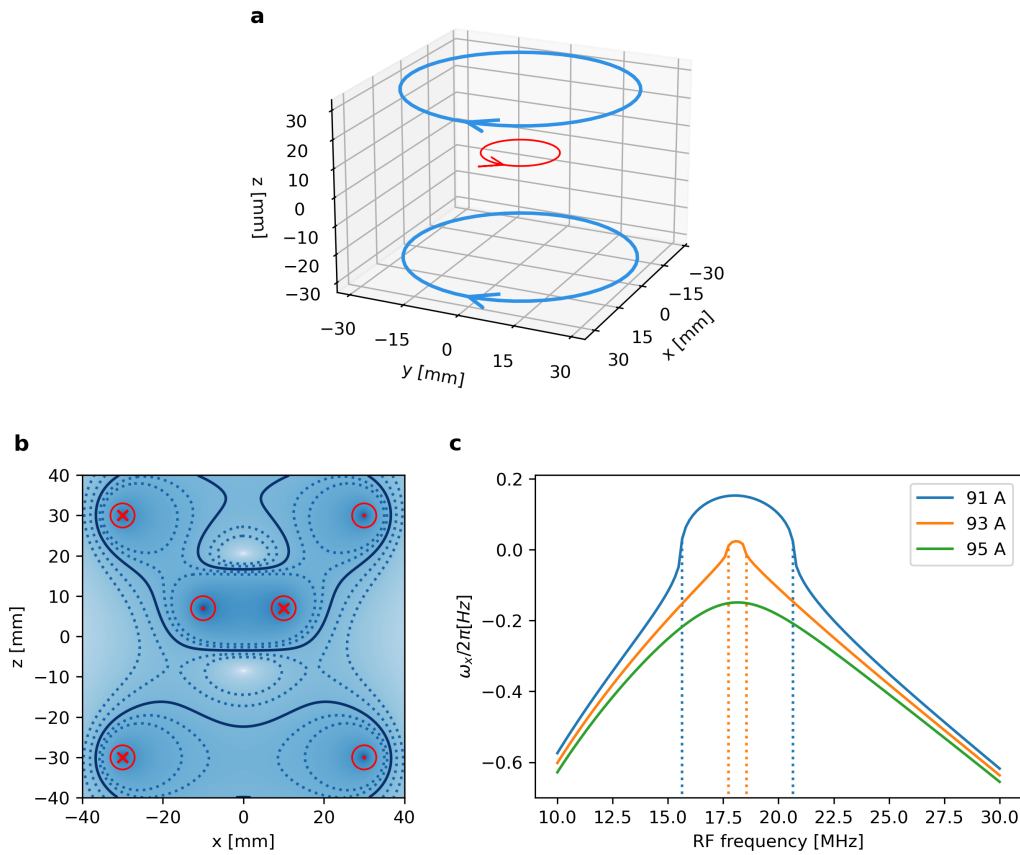


Figure 9.4: **a**: Coil geometry to realise quasi-uniform RF-dressed traps in the presence of gravity. **b**: Isomagnetic contours of a single wire loop and a pair of Helmholtz coils. **c**: Radial trapping frequency as a function of RF frequency at various Helmholtz currents. Dotted lines indicate vanishing harmonic term in the radial potential.

10 | Conclusion and outlook

10.1 Conclusion

This thesis describes experimental work on probing universality of 2D quantum systems with 2D bilayer Bose gases. We have built a new experimental apparatus which utilises a high-flux source of cold-atoms, thereby eliminating the transport of cold atoms to the science chamber. This allowed the development of a new loading scheme for loading atoms into the RF-dressed quadrupole trap which is a simple, one-step process compared to existing methods. We demonstrated the production of a degenerate quantum gas by evaporative cooling in the shell trap. We have performed experiments to explore the full counting statistics of interference contrast of 2D Bose gases. We outlined the importance of additional experimental noise sources in these experiments and found good agreement between theory and experiment. We showed that the experimental distributions converge to the functional form of the universal theoretical distribution as the superfluid phase space density is increased. We investigated the effect of a disordered external potential on the superfluid transition of 2D Bose gases. We found that with increasing disorder strength, the system lost phase coherence and crossed into the normal phase. A comparison of the experimental results with theoretical results based on classical and quantum field theories suggests that quantum fluctuations play an important role in the loss of coherence. We also introduced an effective coupling model to explore the scaling of observables with increasing disorder strength. We have performed various quench experiments across the BKT critical point. By suddenly reducing the interlayer coupling strength, we drove the system into the normal phase, and observed that the time evolution obeyed universal scaling laws at intermediate times. In a separate experiment, we utilised the

strength of disorder to quench the system across the superfluid transition in both directions. For the superfluid phase to normal phase quench, we observed the expected crossover into the normal phase. For the normal phase to superfluid phase quench, we observed unexpected results which could indicate the presence of an out-of-equilibrium state that will be investigated further. Finally, we proposed an RF-dressed experimental scheme to realise uniform magnetic potentials.

10.2 Future experimental directions

10.2.1 Matter-wave interferometry with the next-generation experiment

A straightforward way to continue the work on the next-generation experiment is to use MRF-dressing to split a quantum gas and perform matter-wave interferometry. Since the cloud in the trap is expected to have collective excitations strongly suppressed, the experiment is ideal for performing matter-wave interferometry. To benchmark the performance of the experiment, it is desirable to use clouds in the 3D regime to ensure a spatially uniform condensate phase. From the interference patterns, fluctuations in the spatial phase profile as well as the stability of phase difference between clouds can be deduced. These measurements are an important step towards the deployment of RF-dressed technologies for precision measurements and possible applications in commercial devices.

10.2.2 Phase diagram of disordered 2D Bose gases

While this work is an important step towards understanding the effect of disorder on 2D Bose gases, further measurements are needed to comprehensively map out the phase diagram as a function of disorder strength and phase-space density.

For these measurements, a larger well separation between the wells is desirable to eliminate the possible interplay between interlayer coupling and disorder. Furthermore, it would be advantageous to create a disordered potential by a diffusive plate which allows control over the disorder correlation length. These measurements are also key to better understand the out of equilibrium dynamics of 2D Bose gases following a disorder strength quench. Currently the new apparatus does not have as much optical access to the cold atoms as the previous generation apparatus¹ but this can be addressed by rearranging the MOT beams to be at 45 degrees to the horizontal, all lying in a vertical plane rather than a horizontal plane as in the current setup. This configuration would permit the positioning of high NA objective lenses both above and below the UHV glass cell, allowing for high resolution for both imaging and projection of the light for the optical dipole potentials.

¹Since the previous generation apparatus has no MOT beams propagating through the glass cell.

A | XY model correlation function integral

We are interested in evaluating the integral

$$I(x, y) = \int \frac{1 - \exp(i\mathbf{k} \cdot \mathbf{r})}{k^2} d\mathbf{k}. \quad (\text{A.1})$$

We start with calculating the Laplacian of $I(x, y)$ over spatial coordinates, giving

$$\nabla_{\mathbf{r}}^2 I(x, y) = \int \exp(i\mathbf{k} \cdot \mathbf{r}) d\mathbf{k}, \quad (\text{A.2})$$

where we performed the derivatives inside the integral as they do not act on the variables of integration. By definition, the integral results in $\delta(\mathbf{r})/(2\pi)^2$. Therefore, $I(x, y)$ is proportional to Green's function of the 2D Laplacian operator, giving

$$I(r) = 2\pi \ln(r), \quad (\text{A.3})$$

where $r = \sqrt{x^2 + y^2}$.

B | PID controller firmware and control interface

The firmware of the PID controller consists of two distinct units: the PID algorithm and the control interface which is used to set parameters such as the setpoint as well as read out the error signal. Parallel programming is difficult on MCUs therefore these two processes run concurrently. Priority is given to the control interface which activates as a hardware interrupt. The PID controller halts execution while the control interface is active and resumes when all control tasks are finished. The duration of performing the control tasks is typically 2 ms, and considering the usual 500 Hz update rate, at most 2 PID cycles are affected when a parameter is being updated. This typically happens before an experimental run, therefore it has no negative influence. The PID controller has 8 writable control registers and 3 readable status registers which are initialised when the device is turned on. These registers can be classified into four distinct categories with different functions. Firstly, binary registers have a value of either zero or one and they are used for turning the PID controller on and off, to start the sample and hold mode or to set the gain sign of the control loop. Secondly, gain registers store the proportional, integral, and derivative gains of the controller in an encoded format. The special feature of this register type is that it can efficiently apply the encoded gain on values such as the error signal or its cumulative sum (the integral term). All processing is done using integer arithmetic until applying the gains without any discretisation error, resulting in fast computation. Therefore, it is desirable to restrict applying the gains to integer arithmetic. This does not directly allow a gain to be set to for example 1.5. This is an issue as the control loop might be self-oscillating at a gain of 2 but a gain of 1 might not provide optimal stability.

To circumvent this issue, the gain is represented by 2 numbers A and B such that the output x , for a given input y (which can be the error signal or its integral) is given by

$$x = A2^B y \quad (\text{B.1})$$

where A and B are both integers. This is advantageous as multiplying the error signal by A can be performed using integer arithmetic and multiplying or dividing by powers of 2 can be efficiently executed by bit shifting. For example, an error signal value of 1235 and a gain of 1.5 represented with $A = 1536$ and $B = -10$ yields 1852 while the floating-point calculation gives 1852.5, indicating that the errors are only due to rounding. This representation of the gains is not unique and values for A and B are calculated such that A is the largest possible 8-bit integer which can represent the desired gain. Registers of the third type contain integer values up to 16 bits and they are designated for the setpoint and the update rate. Finally, read-only registers contain the input, error signal and output values and these can be read back to a computer for monitoring purposes. At each PID update cycle, the ADC takes 16 samples of the input signal and we average the samples by bit shifting. The value of the error signal is updated and it is added to the integrator buffer. To prevent unnecessary accumulation of the integral of the error signal, updating the integrator buffer is suspended when the input or output of the control loop is outside the normal range (maximum range with 5 percent margins) for more than 50 update cycles. This is useful for example when the MOT beams are shuttered. Without the selective update of the integrator buffer, the time until the integrator would return to a normal value would be comparable to the time the MOT beams were off which is far from ideal behaviour. While deactivating the control loop off along with closing mechanical shutter would address this, the simple rule for updating the integrator buffer eliminates the need for real-time adjustments of the control loop. Following the update of the integrator buffer, the

control loop's output value is calculated and the duty cycle of the output signal is adjusted accordingly. At the end of a single update cycle, the elapsed time is measured and the PID algorithm pauses for a period ensuring evenly spaced updates in time. The output indicator LEDs and the status registers are updated at every 5 PID update cycles.

Remote control of the PID controller functions through manipulating the previously described registers via an SPI interface. The communication is conducted in three 8-bit words. The first word addresses the PID controller, selecting a register and an action. The first 2 bits specify which the channel being addressed, bits 3 and 4 denote the action to be taken (such as read or write) and the remaining 4 bits determine the target register. The last 2 words contain the values which are transferred to the PID controller's addressed register. When a register's value is read, the last 2 words are ignored and the PID controller sends the values stored in the register.

C | FFT-based analysis of interference phase and contrast

C.1 FFT analysis for clean data

We create noiseless data computationally which is of the form

$$f(x) = A \exp\left(-\frac{x^2}{2w^2}\right) (1 + C \cos(k_0x + \phi)). \quad (\text{C.1})$$

We choose the length unit to be 1 pixel as in images and use experimentally relevant parameters: $|x| < 75 px$, $w = 10 px$, $2\pi/k_0 = 10 px$, $A = 20$, and $C = 0.3$. Such a fringe pattern is shown in Figure C.1 a where a phase of $\pi/4$ was chosen. While there is an analytic expression for the Fourier transform of $f(x)$, it is a convoluted formula, therefore we rely on discrete Fourier transforms in this investigation. We plot the absolute value and phase its discrete Fourier transform in Figure C.1 b, where we normalised the absolute value by the sum of $f(x)$. Note that $|\tilde{F}(k_0)|$ is half the contrast. In the experiment, it is not guaranteed that the Fourier peak of the matter-wave interference fringes coincides one of the discrete spatial frequencies of the FFT, which reduces the measured contrast at k_0 . However, this is a small effect, we expect an error of at most 2.5 percent due to the discretisation. By taking the average of $|\tilde{F}(k)|$ a few points around k_0 , the measured contrast is reduced. This is not an issue since, the decay of $|\tilde{F}(k)|$ is deterministic, therefore the correct contrast can be found by a simple scaling factor. In the analysis of the experimental data, the contrasts were always normalised to a mean value therefore there was no need for such a correction. The phase of the $\tilde{F}(k_0)$ is $\pi/4$ as expected and exhibits jumps of π around the peak. Since these jumps are also deterministic, averaging the phase a few points around

k_0 gives a deterministic offset, therefore this method is able to correctly identify phase profile of a set of columns in an image.




Figure C.1: a: Computer generated fringe pattern with typical experimental parameters. b: Discrete Fourier transform of the generated pattern.

C.2 FFT analysis of noisy data

Noise on the images due to CCD readout or any other sources affects the real and imaginary parts of $\tilde{F}(k)$ and can be well modelled with a Gaussian white noise. Therefore we repeat the FFT-based analysis by adding Gaussian white noise to the generated real-space data. We chose a standard deviation of 0.05 which corresponds to the observed noise level. The effect of this noise for $\tilde{F}(k_0)$ represented as phase and contrast noise is illustrated in Figure C.2a. To reduce the noise, we take the averages of the phases and absolute values at k_0 and $k_0 \pm 1$. We repeat this for 10000 realisations of the noise and compare the spread of the resulting phase and contrast data for analysis by fit, FFT analysis from a single complex amplitude at k_0 , and FFT analysis from 3 complex amplitudes in Figure C.2b/c. For both contrast and phase, the fitting analysis is in agreement with FFT analysis from a single complex amplitude. However, the spread of data for

analysis from 3 complex amplitudes narrows the distribution by approximately 1.7 as expected from the scaling of standard error of the mean with sample size. By analysing 5 complex amplitudes, the spread can be further narrowed. Using more complex amplitudes is not practical as the signal to noise ratio drops quickly as we move away from the Fourier peak at k_0 .

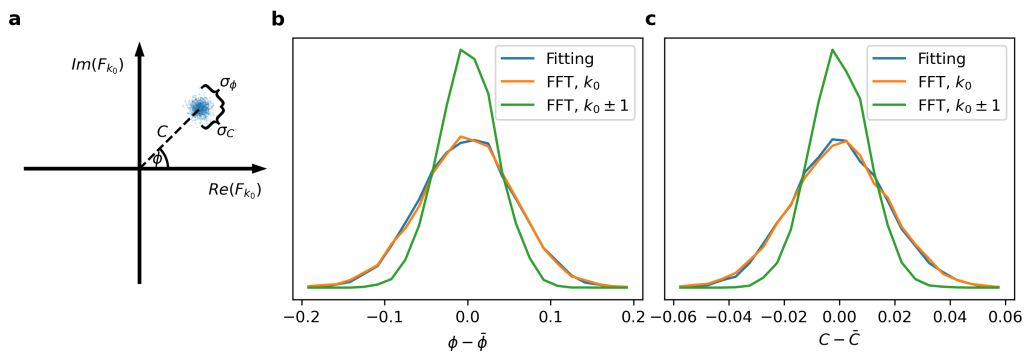


Figure C.2: a: Phase noise and contrast noise due to Gaussian white noise of the real-space data. b: Distribution of phase error. c: Distribution of contrast error.

D | Simulating quantum gases

Simulations play an important role in exploring the static and dynamic properties of quantum gases and help with the interpretations of experiments. In this chapter, we introduce pyGPE, a simulation toolbox which can be used to simulate quantum gases at zero and finite temperatures. Furthermore, we will present two case studies in which pyGPE was used to help with analysis of experimental data. Finally, another simulation based on Monte Carlo methods is introduced to simulate a Bose gas at a finite temperature.

D.1 Coupled Gross-Pitaevskii equations

The Hamiltonian of a Bose gas interacting via s-wave scattering, written in a second quantised formalism is [149]

$$\hat{\mathcal{H}} = \int \Psi^\dagger(\mathbf{r}) \left(-\frac{\hbar^2 \nabla^2}{2m} + V(\mathbf{r}) \right) \Psi(\mathbf{r}) d\mathbf{r} + \frac{g}{2} \int \Psi^\dagger(\mathbf{r}') \Psi^\dagger(\mathbf{r}) \Psi(\mathbf{r}) \Psi(\mathbf{r}') d\mathbf{r} d\mathbf{r}', \quad (\text{D.1})$$

where $\Psi^\dagger(\mathbf{r})$ and $\Psi(\mathbf{r})$ are Bosonic creation and annihilation operators at location \mathbf{r} . Within the Hartree-Fock approximation, a Bose gas with a finite condensate fraction is described by functions obeying the equations [110]

$$\left(-\frac{\hbar^2 \nabla^2}{2m} + V(\vec{r}) + g(n_0(\vec{r}) + 2n_T(\vec{r})) \right) \psi_0 = \mu \psi_0 \quad (\text{D.2})$$

and

$$\left(-\frac{\hbar^2 \nabla^2}{2m} + V(\vec{r}) + 2g(n_0(\vec{r}) + n_T(\vec{r})) \right) \phi_i = \epsilon_i \phi_i, \quad (\text{D.3})$$

where ψ_0 is the wavefunction of the condensate with corresponding chemical potential μ , ϕ_i are the single-particle wavefunctions of the excited states with cor-

responding energies ϵ_i , V is the external potential, g is the interaction strength, n_0 is the condensate density, and n_T is the density of the thermal component. By convention, ψ_0 is normalised to the number of atoms in the condensate such that $n_0 = |\psi_0|^2$ and ϕ_i are normalised to unity, therefore $n_T = \sum_{i>0} n_i |\phi_i|^2$ where n_i is the occupation number given by the Bose-Einstein statistics. At zero temperature, the occupations of all the excited states are zero and we recover the normal Gross-Pitaevskii equation. Note that the Hartree-Fock approximation involves replacing the operators in [D.1](#) with c-fields (the condensate and excited state wavefunctions). This has significant implications, for example the Gross-Pitaevskii equations cannot explain quantum phase transitions which arise from the commutation relations of the Bosonic field operators, such as transition to a Bose-glass phase [[150](#)].

D.2 Finding the ground state at zero temperature

To find the ground state at zero temperature, the imaginary time propagation method is used [[151](#)]. The key idea is to make the $t \rightarrow -it$ transformation in the time-dependent Schrödinger equation. This results in the non-unitary time evolution

$$|\psi(t)\rangle = \sum_i a_i(0) e^{-E_i t/\hbar} |\phi_i\rangle, \quad (\text{D.4})$$

where $a_i(0) = \langle \phi_i | \psi(0) \rangle$. From this it is clear that the higher the energy of a given eigenstate is, the faster its amplitude decays. Therefore, starting from an arbitrary initial state, by successive propagation in imaginary time and renormalisation, the state converges to the lowest energy state. To apply the operator $e^{-Ht/\hbar}$ we use a second order symmetrised Suzuki-Trotter expansion [[152](#)]

$$e^{-Ht/\hbar} \approx e^{-Vt/2\hbar} e^{-Tt/\hbar} e^{-Vt/2\hbar}, \quad (\text{D.5})$$

where V is the position dependent part and T is the momentum-dependent part of the Hamiltonian. This approximation is only valid for small t , therefore one often needs to adjust this parameter. Too large t might not result in a converged solution while a too small t slows down the convergence. The various terms in the approximate evolution operator are diagonal either in position and momentum spaces. Therefore, Fourier transforms can be employed to efficiently change the basis. After every iteration, the chemical potential is calculated at the location of the trap minimum. The algorithm is stopped when the change in chemical potential between two successive iterations is less than a given threshold value. In the software implementation of this method, the FFTW library [153] and the Numba just-in-time compiler [154] is being used to ensure fast execution speeds. Furthermore, a GPU-accelerated version of pyGPE was developed which uses the CuPy library [155]. An example of executing such a calculation using the high-level API of pyGPE is found below.

```
import numpy as np
from pygpe_oop.util.grid import Grid
from pygpe_oop.util.potentials import Potential
from pygpe_oop.nonthermal.groundstate import Groundstate

gld = Grid(dimension=1, dx=3e-8, Mx=1024)
wx = 10*2*np.pi
wr = 500*2*np.pi
V1d_harmonic = Potential('harmonic', gld, wx=wx, wy=wr, wz=wr)
Gnd1d = Groundstate(V1d_harmonic, N=100, dt=1e-7)
psi0, mu0 = Gnd1d.get_groundstate(tolerance=1e-8)
```

We create a Grid object which defines a simulation grid in 1 dimension with 30 nm spacing and 1024 points in the positive x -direction (2048 points in total). It is in general advantageous to use a grid size of a power of 2 as it makes the FFTs computationally efficient. We then create a Potential object based on the grid

which defines a harmonic potential with 10 Hz trapping frequency. Providing the other two trapping frequencies is necessary as the interaction strength depends on these within the quasi 1D approximation. Finally, we create a Groundstate object based on the potential and we set the atom number to 100 and the time step for imaginary time propagation to be 100 ns. The final line performs the imaginary time propagation algorithm until the relative change in chemical potential between two steps is 10^{-8} and returns the ground state wavefunction and the corresponding chemical potential. The resulting density is plotted in Figure D.1 a. To verify the results, this calculation is repeated with atom numbers in the 1 to 10000 range and the chemical potential is plotted as a function of atom number in Figure D.1 b. For high atom numbers, the chemical potential asymptotically approaches the Thomas-Fermi limit, while for low atom numbers it converges to the zero-point energy of the quantum harmonic oscillator.

D.2.1 The case of rotational symmetry

In many scenarios the quantum gas has rotational symmetry due to the trapping geometry, for example, in the dressed quadrupole trap with circular RF polarisation. In this case, the wavefunction is a function of the radial coordinate only in two dimensions or the radial and axial coordinates in three dimensions. This significantly reduces the memory requirements and speeds up simulations. However, one must replace the Fourier transforms in the simulations with Hankel transforms [156]. This is implemented by the standard way of matrix-vector products using the discrete Hankel transform (DHT) matrix [157]. Since the transformation matrix does not have the same symmetry properties as a Vandermonde matrix, the DHT is an $O(N^2)$ operation and it is slower than FFT for large enough N. Another disadvantage of the DHT is that the spatial grid is non-uniform. Nevertheless, for realistic simulation grid sizes, it was found that DHT approach in one or two di-

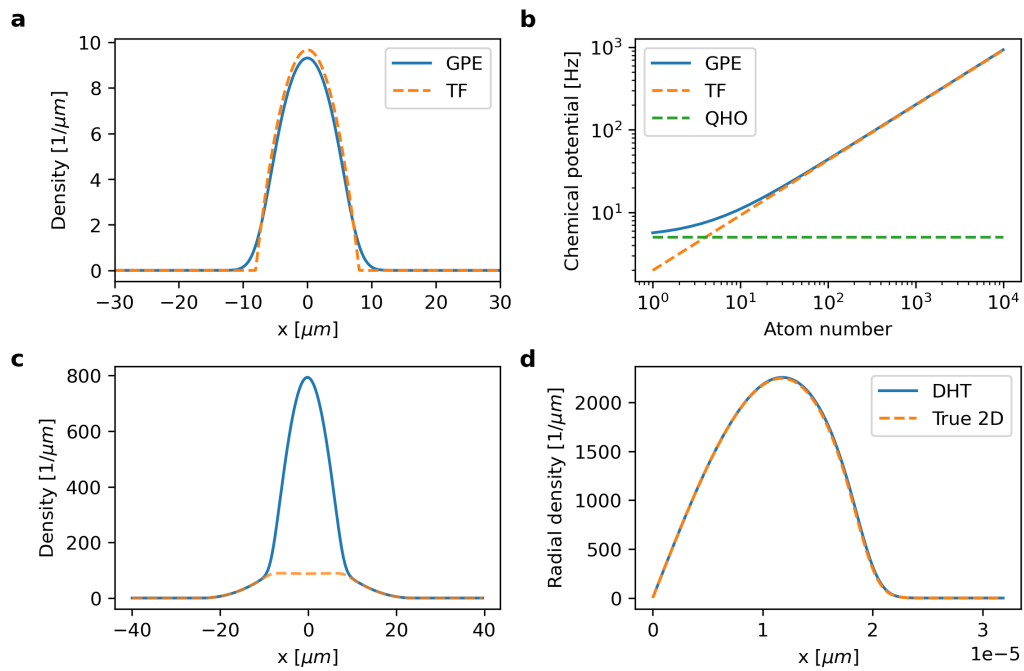


Figure D.1: Ground state calculations with pyGPE. **a:** Quasi 1D calculation with 10 Hz trapping frequency and 100 atoms. Dashed line shows the density in the Thomas-Fermi limit with identical parameters. **b:** Scaling of chemical potential with atom number. Dashed lines show the Thomas-Fermi and the non-interacting quantum harmonic oscillator limits. **c:** Finite temperature simulation with total density (solid blue line) and excited state density (orange dashed line). **d:** Radial density for a rotationally symmetric simulation via DHT and for a true 2D simulation.

mensions is significantly faster than the FFT approach for the corresponding two or three-dimensional problem. Radial densities calculated with this method are compared with true 2D results in Figure D.1 d, demonstrating excellent agreement. The ground states were calculated with $\omega_r/2\pi = 10$ Hz, $\omega_z/2\pi = 1000$ Hz and for 30000 atoms. The resulting chemical potentials were 0.2 percent different which further demonstrates the feasibility of the DHT approach. The execution time was 260 s for the true 2D simulation while only 1 s when the radial symmetry was exploited.

D.3 Finding the initial state at a finite temperature

To find the initial state of a Bose gas at a finite temperature, one must solve the coupled GP equations. We start by calculating the condensate fraction from a given total atom number and temperature, assuming a certain trapping potential such as harmonic or box potential which has a known expression for this. Then the condensate wavefunction and its chemical potential is calculated, ignoring the thermal states. After this step, the Hamiltonian matrix for the thermal states is constructed using a second-order central finite difference method. Since the density of thermal states ¹ is unknown, the interaction part of the Hamiltonian only considers the condensate. Using ARPACK's sparse solver [158], the first N eigenvalues and eigenvectors are calculated. From the eigenvalues and eigenvectors, occupation numbers and the density of thermal states is calculated. After this we iterate between imaginary time propagation for the condensate wavefunction until a desired tolerance, and partial diagonalisation of the Hamiltonian for the thermal states. Since the initial guess for the thermal states ignored interactions with themselves, this procedure usually results in a slowly converging, oscillatory

¹as particles per volume

behaviour of the energy eigenvalues. To solve this, the self-interaction term is turned on slowly by scaling it with $1 - e^{-6j/N}$ where j labels the iteration step and N is the total number of steps. It was found that this method provides converging solutions for 15 total steps, which limits the number of times the Hamiltonian is partially diagonalised. For large grid sizes, partial diagonalisation is computationally expensive. In this case, imaginary time propagation can be used for finding the excited single particle states. To prohibit the wavefunctions converging to the lowest energy eigenstate, we perform a Gram-Schmidt process after every iteration to force the eigenstates to be orthogonal.

D.4 Finding ground state with a fluctuating phase

The ground state found by imaginary time propagation always has a uniform phase since it is at zero temperature. In quasi-2d systems that we investigate in experiments, phase fluctuations and vortices destroy long range order at any finite temperature. To simulate these systems, the phase of the wavefunction must be provided externally and imprinted on the wavefunction after every iteration of imaginary time propagation. This way, the algorithm converges to the lowest energy state with a given phase fluctuation which results in density fluctuations. Ground states calculated by this method are shown in Figure D.2, considering 30000 atoms in a harmonic trap with 10 Hz radial and 1 kHz axial trapping frequencies. The phase fluctuations for the first four images are results of finite temperature classical field Monte Carlo simulations of a harmonically trapped Bose gas. Details of this simulation are found in [12]. For the last image, phase profiles of 10 vortices at random positions are imprinted which have zero net angular momentum.




Figure D.2: Ground state densities with various phase profiles. The first four images were generated using phase profiles from Monte-Carlo simulations at various temperatures. The last image was generated using the phase profiles of 10 vortices.

D.5 Simulating dynamics

Dynamics is simulated by the same method as finding the ground state but without the imaginary time transformation. In this case, there is no need for renormalising the condensate wavefunction as the evolution is unitary. We differentiate between three different types of evolution, trap, hold, and time-of-flight. In the trap phase, the potential is an arbitrary time-dependent potential. In the hold phase, the potential is constant. Finally in the time-of-flight phase the external potential is turned off. In practice we use the last stage the most to compare theory, simulation and experiment. Therefore, two improvements in speed were implemented. Firstly, if interactions can be ignored as in the case of ballistic time-of-flight expansion, the Hamiltonian has no spatial dependency. In this case, there is no need to iteratively propagate the wavefunctions. When interactions cannot be ignored, the speed can be improved by adaptively adjusting the time step. The time steps dt are chosen to keep $g|\psi_0|^2 dt$ (or the corresponding term for the thermal states) constant. As the cloud expands, dt grows rapidly as larger values of dt result in more expansion. Examples of such simulations are shown in Figure D.3 in the quasi-1d case. In the images, the amplitude of the wavefunction is plotted (for better visibility) at various times along the simulation. In Figure D.3 a, a time-of-flight expansion is

simulated for a condensate in a 50 Hz 1-d harmonic trap (with the other two trapping frequencies being 500 Hz). The simulation ignores expansion along the other dimensions. The dashed line indicates the width of the cloud at half maximum, indicating the expected linear expansion in time for $\omega_x t \gg 1$. In Figure D.3 b, a wavefunction identical to the one as in the previous case is displaced left and right and summed. As the clouds undergo time-of-flight expansion and overlap, matter-wave interference fringes are created. Since the phase of the two wavefunctions is identical, a bright fringe is created in the middle. Finally, in Figure D.3 c, we start with a single cloud as in scenario a, but we displace the trap centre by 25 μm as well as reduce the trapping frequency to 25 Hz in 75 ms and hold in the final trap for 150 ms. Since the change in the potential happened on a timescale comparable with the trap frequency, sloshing is observed.




Figure D.3: Examples of dynamical simulations. **a**: Time-of-flight. **b**: Matter-wave interference. **c**: Displacing and decompressing harmonic trap.

D.6 Case study 1: Simulating matter-wave interferometry

The main experimental protocol used in this work is assumed to measure the in-situ local phase difference between two pancake-shaped Bose-gases via matter-wave interferometry. However, there are many factors which influence whether

this assumption is true, for example the effect of time of flight, the effect of interactions during time of flight or the finite imaging resolution. To quantify these effects, we performed simulations using pyGPE. The zero-temperature 3D ground state of a single cloud was calculated initially with parameters corresponding to the experiment. We used a cylindrically symmetric harmonic trap with 11 Hz radial and 1 kHz axial trapping frequencies and the atom number was set to 35000. For such a system, the ideal gas BEC critical temperature is 77 nK while the BKT critical temperature is 70 nK [121]. However, due to the nonuniform, trapped gas, the crossover regime is significantly broadened. Phase fluctuations from classical field Monte Carlo simulations were imprinted on the ground state for temperatures deep in the superfluid regime (10 nK) and in the crossover regime (37.5 nK). This procedure was repeated for an identical copy of the ground state but with a different realisation of the phase fluctuations. The initial state was constructed by summing these two wavefunctions, displaced by the typical well separation as shown in Figure D.4 a/b. Time-of-flight simulations were then performed, simulating both ballistic and non-ballistic expansions. After 15 ms, the density was calculated in a region of width $3.5 \mu\text{m}$ which corresponds to the thickness of the selective imaging region in the experiment (Figure D.4 c).

The simulation protocol was repeated for 100 different pairs of phase profiles at the two temperatures. The resulting images were analysed using the same phase and correlation function extraction methods as experimental images. To account for the finite imaging resolution, we approximated the point-spread function of the imaging system with a gaussian of width $2 \mu\text{m}$ and took the convolution of the raw images and the point spread function. The correlation functions for the two temperatures corresponding to the superfluid and crossover regime are shown in Figure D.5 a/b with and without interactions. The correlation function in the case of ballistic expansion demonstrates better coherence at large distances when com-

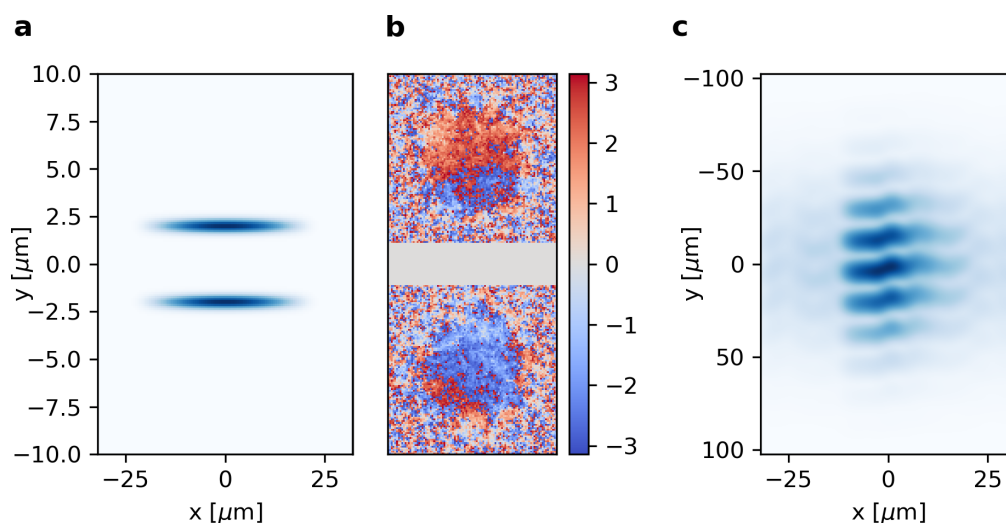


Figure D.4: Illustration of the protocol used for simulating matter-wave interference. **a:** The integrated density of the initial state showing two clouds separated by $5 \mu\text{m}$. **b:** Phase profiles of the two clouds, from classical field MC simulations. **c:** Image of the system after 15 ms of time-of-flight.

pared to the interacting expansion, the effects are small however. Interestingly, the in-situ phase correlation function lies between the two GPE results. We attribute this to the averaging effects of calculating the mean density in the selective imaging region and the finite imaging resolution. Overall, there is a good qualitative agreement between the two GPE and the MC simulations. In the crossover regime, a similar effect can be observed when comparing the two GPE results, the non-interacting correlation function decays slower than its interacting counterpart. However, there is a more significant qualitative difference between the GPE and the MC results. There are multiple effects which contribute to this. The integration effects of imaging resolution as well as the selective imaging region increases the measured coherence at low distances. As the phase fluctuations on short length scales are stronger in this case, the integration effects result in a larger relative increase in coherence. This is not a significant limitation of the experimental scheme as we are primarily interested in the intermediate-range behaviour

of the correlation function. The effect of imaging resolution is further investigated in Figure D.5 c/d in the superfluid and crossover regimes respectively for the interacting GPE simulations by changing the Gaussian width of the point-spread function. Due to the integration effect, larger point-spread function widths result in a correlation function decaying slower. In the superfluid regime, the decay of correlations is faster than expected when infinite resolution is assumed. This highlights the decay of coherence due to time of flight expansion in the radial direction as well as the role of interactions. As a fortunate coincidence, the finite imaging resolution brings the correlation function closer to the one in-situ. In the crossover regime, a similar effect is present, but due to heavy phase fluctuations, the correlation function is overcorrected at short distances. As a general observation, the shape of the correlation functions does not change at distances greater than the imaging resolution. From these results we can conclude that the experimental matter-wave interferometry scheme reproduces qualitatively equivalent correlation functions to the in-situ correlation function. However, it is important to consider the effects of imaging resolution, time of flight expansion and interactions when comparing theoretical and experimental results.

D.7 Case study 2: Matter-wave lensing

Releasing atoms from RF-dressed traps can be done by two methods. The first method, commonly referred to as normal turn-off, involves turning off the DC magnetic field at first and then turning off the RF field. This turn-off scheme imparts momentum on the atoms since the resonant spheroid is swept outwards as the DC magnetic field turns off. When releasing atoms from the double-well trap, normal turn-off imparts a different momentum on the two clouds which can influence matter-wave interferometry. To avoid this, a second method called RF

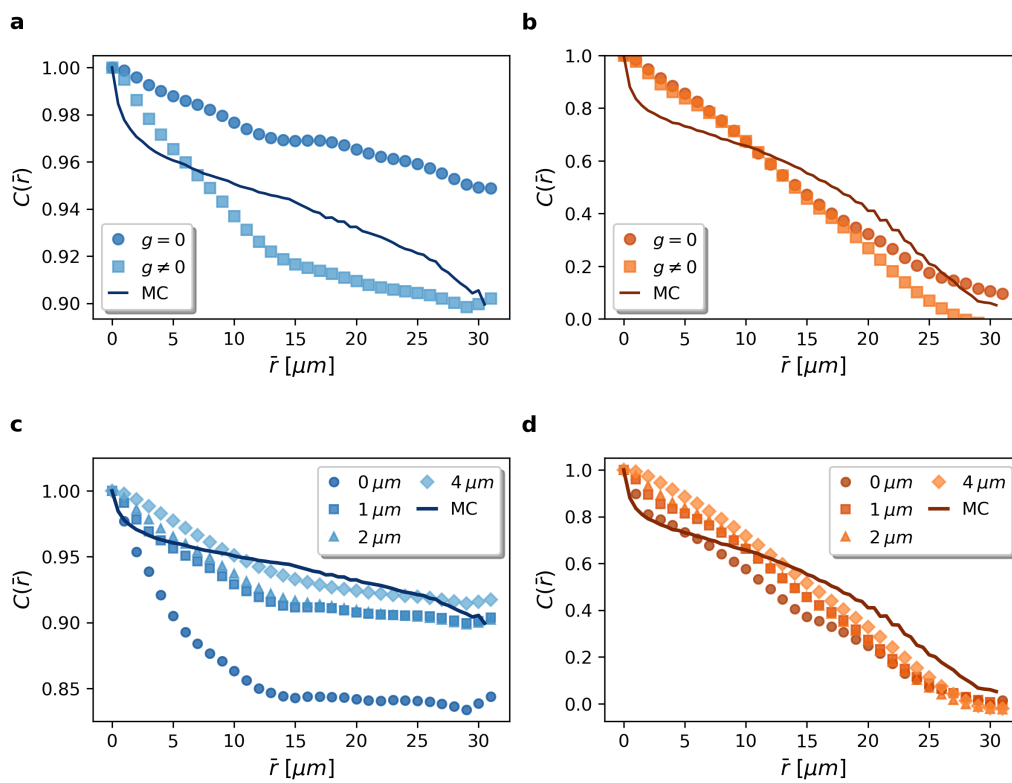


Figure D.5: Effects of interactions and imaging resolution on measuring correlation functions via matter-wave interferometry.

turn-off can be employed. In this case, the RF field is turned off at first, projecting the atoms into the Zeeman states. The magnetically insensitive state is only affected by gravity while atoms in the other Zeeman substates experience a magnetic force while the DC magnetic field is on. Typically, the $|m_F = 0\rangle$ state is used in measurements as it is unaffected by the residual magnetic field. In typical matter-wave interferometry measurements, fringes can be observed when the other two states are imaged as shown in Figure D.6 a. The low-field seeking state is compressed by the DC magnetic field while the high-field seeking state is stretched in the radial direction. The latter is advantageous as this magnification effect can result in more precise measurement of the correlation function at low distances. However, it is not known whether such a measurement is accurate as dynamics in the magnetic field can alter the correlation function. Note that the three clouds are displaced horizontally a small amount with respect to each other. This could be either due to stray magnetic fields created by Eddy currents during turn-off or a small angular misalignment of the quadrupole field with respect to gravity. Since the displacement is small compared to the size of the cloud size or the difference in cloud size between the $|m_F = 0\rangle$ and the high-field seeking states, we ignore this misalignment. To simulate this effect, we firstly calculate the ground state wavefunction of a cigar-shaped condensate in a harmonic trap with trap frequencies 5 Hz, 50 Hz and 1000 Hz along the x , y , and z directions respectively within the quasi-2D approximation. Conceptually, the interference of two such cigar shaped clouds in the quasi-2D approximation is equivalent to the interference of two pancake-shaped clouds in 3D as in the experiment, however the latter requires longer execution times for the simulation. The choice of the trapping frequency along x was motivated by the radial trapping frequency of the typical RF-dressed trap and the system size while the trapping frequency along y was chosen such that the interfering clouds expand slower to reduce the size

of the simulation grid. The atom number in the condensate was chosen to be 50 which gives a chemical potential of 300 Hz, which is typical in the performed 2D experiments. After this, phase profiles with exponential correlation functions of different correlation length² were generated by shaping the power spectral density of white noise in k-space and were imprinted on the condensate. Similar to the previous case study, the clouds with different realisations of phase profiles were displaced and added. The density in the initial states as well as examples for the imprinted phases is shown in Figure D.6 b/c. The initial wavefunction was then evolved in the quadrupole magnetic potential whose magnitude was scaled temporally in a sigmoid-like fashion as

$$B_0(t) \propto 1 - \frac{e^{(t-t_0)/5t_0}}{1 + e^{(t-t_0)/5t_0}}, \quad (\text{D.6})$$

with $t_0 = 0.5$ ms. This corresponds to the typical turn-off of the current in the quadrupole coils over 1 ms. After 1 ms evolution, the wavefunction was multiplied by a phase factor of form e^{ik_0x} to ensure that the momentum expectation value of the wavefunction is near zero. This is equivalent to simulating the cloud in its rest frame, as there are no external forces acting on the cloud after the initial 1 ms period, allowing for smaller simulation grid sizes. The wavefunction is then further evolved for 15 ms in absence of an external potential. An example of the fringe pattern after evolution is shown in Figure D.6 d.

The simulation was repeated for 50 different realisations at each correlation length as well as for clouds without an imprinted phase. The interference patterns were analysed with the same phase and correlation function extraction methods as experimental images after correcting for the finite resolution of the imaging system. The fringe pattern of the wavefunction amplitude in absence of imprinted phases

²Our primary interest is how the detected correlation length relates to the real correlation length. In the superfluid regime, we expect the algebraic exponent to be scale invariant.




Figure D.6: Experimental image and simulation protocol for matter-wave lensing. Note that as opposed to the usual convention in the thesis, the vertical axis shows the y direction. This is because pyGPE automatically considers x and y to be the two spatial directions for quasi 2D simulations.

is shown in Figure D.7 a. As expected due to the shape of isomagnetic surfaces of the quadrupole field, a quadratic phase profile appears. However, the quadratic coefficient is very small, $125 \mu\text{rad}/\mu\text{m}^2$. Between the centre and the edge of a typical cloud, this amounts to a phase shift of 0.18 rad. The corresponding correlation function decays by only 0.4×10^{-4} over a distance of $25 \mu\text{m}$ which is orders of magnitude less than the typical decay in experiments. Therefore, the change in the detected phase profile due to evolution in the quadrupole potential is negligible. Figure D.7 b shows the measured correlation length as a function of the in-situ correlation length. For the $|m_F = 0\rangle$ state, the two are in good agreement as indicated by the dashed line of slope 1. The high-field seeking state shows a similar linear scaling but with a slope of approximately 2. This equals the ratio of the Thomas Fermi radii of the two states after time of flight expansion. For large in-situ correlation lengths, the measured correlation lengths are consistently

lower than expected due to finite size effects. In experiments, the typical measured correlation lengths are in the 10 to 30 μm range, therefore this only introduces a small error. We can conclude that the evolution in the quadrupole potential does not introduce a significant change in the correlation lengths and therefore in the correlation functions for the high-field seeking state. Therefore, this scheme is experimentally feasible and can be used to measure the same phase profile twice, reducing the uncertainties. Moreover, due to the stretching effect, the resulting correlation functions are expected to be more accurate at short distances.

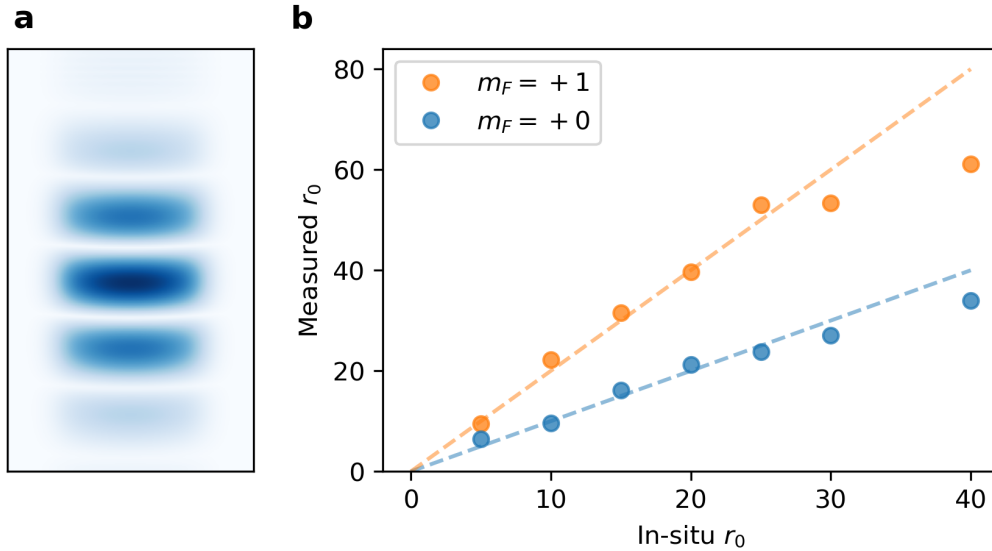


Figure D.7: Results of the matter-wave lensing simulations. **a**: Fringe pattern of the wavefunction amplitude without phase imprinting. **b**: Measured vs. in-situ correlation length for the two Zeeman states of interest.

D.8 Classical field Monte Carlo simulations

Calculating initial states with GPE alone will always result in the zero temperature state. To simulate thermal properties of Bose gases, another approach based on classical field Monte-Carlo methods and the Metropolis algorithm can be em-

ployed. Firstly, we make the classical field approximation by replacing the field operators in D.1 with c-numbers. This is a good approximation for dilute gases and modes occupied by a large number of bosons [159]. We then discretise the spatial coordinate, ensuring that the lattice spacing is comparable or smaller than the healing length [160]. An field configuration is proposed by setting all of the field values to zero. Then, a random lattice site is picked where a new field value is proposed. The corresponding change in energy and particle number is calculated and the associated grand canonical weight $\exp(-\beta(\Delta E - \mu))$ is calculated³. A random number x is then generated from a uniform distribution in the 0 to 1 range. The new field configuration is accepted only if the weight is larger than x . This process is repeated until equilibrium in energy and particle number is reached, which takes of the order of 100000 updates per site. Such simulations are useful in generating phase and density fluctuations for dynamical GPE simulations or in understanding how correlation functions change when working with trapped systems.

D.9 Case study 3: Effect of box potential on correlation functions

Equilibrium BKT measurements in harmonic traps suffer from both finite size and local density effects. As it was shown in Section D.6 and in [12], careful data is needed as these effects modify correlation functions. To reduce the local density effects, a circular box potential can be added to the harmonic potential. Experimentally, this can be done by a ring-shaped blue-detuned beam. By adiabatically ramping up the intensity of this potential, a fraction of the thermal atoms can be

³We found that it is important to evaluate the kinetic term by a forward or backward difference method. Employing a symmetric difference method does not change the kinetic energy locally, leading to spurious results.

pushed inside the trap, increasing the density. Moreover, a sharp cutoff in density is introduced at the edge of the box. These two effects are expected to reduce the density variations in the trap and improve the behaviour of correlation functions at long distances. To investigate this effect, we use classical field Monte Carlo simulations to generate phase fluctuations of a 2D Bose gas in a trap which is the sum of a harmonic and finite size box potentials of variable height. We chose a radial trap frequency of 11 Hz, dimensionless interactions strength of 0.076 (corresponding to 1 kHz axial trapping frequency) and a circular box potential of radius $22.5 \mu\text{m}$ and height of $20/2\pi$ kHz. We chose the chemical potential to correspond to typical experimental atom numbers and the temperature to be 40 nK. The lattice spacing was chosen to be $0.5 \mu\text{m}$ such that the lattice system approximates well the corresponding continuous system. From a random initial state, the system was left to equilibrate for 500000 updates per site. After this, we sampled the field after 1000 updates per site 500 times. The simulation was repeated for a range of chemical potentials with and without the box potential. The mean density was calculated from the fields at each chemical potential, which are shown in Figure D.8 a, matched approximately by atom number. Since the hard walls repel the atoms, the peak density is higher with the hybrid potential, as one would expect from the Thomas-Fermi approximation. Due to the sharp drop in density at the edge of the box potential, the gas is more uniform. The correlation functions calculated from the phase of the fields is shown in Figure D.8 b in the superfluid regime (blue, $N \approx 25000$) and in the crossover regime (orange, $N \approx 15000$). The solid lines show the correlation function with the hybrid potential, showing a slower decay compared to the harmonic potential only with similar atom numbers. However, when we compare identical chemical potentials, the correlation functions are nearly identical in the two different potentials. At large distances, the gas trapped in the hybrid potential demonstrates a slightly higher coherence.

Therefore, we can conclude that by adding a hard-wall potential to the harmonic trap, at a constant atom number, the density at the centre of the trap can be increased which results in longer range correlations.

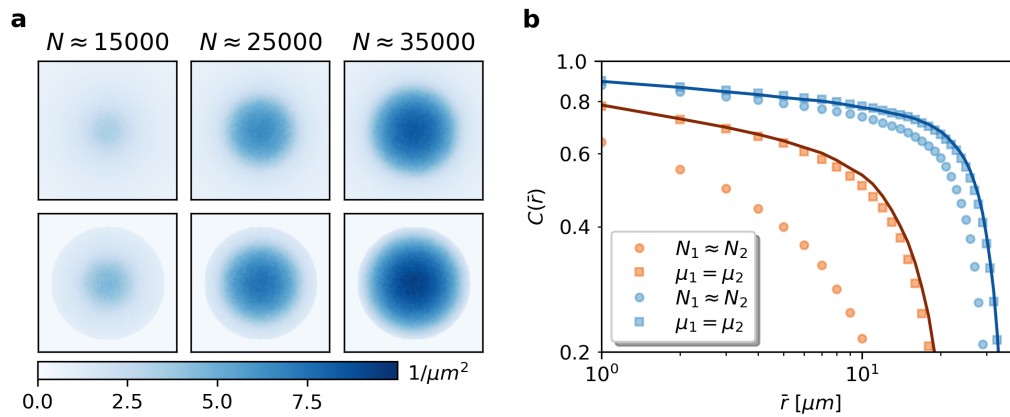


Figure D.8: Results of the classical field MC simulations in a harmonic trap with hard walls. **a**: Mean density at various atom numbers. **b**: Phase correlation functions extracted from the classical field MC simulations. The solid blue and red lines correspond to the harmonic potential with hard walls in the superfluid and crossover regimes respectively. The circle and square markers show the correlation functions without the hard wall potential for matched atom numbers and matched chemical potentials respectively (compared to the system with hard walls).

Bibliography

- [1] D. R. Nelson and J. M. Kosterlitz. Universal Jump in the Superfluid Density of Two-Dimensional Superfluids. *Phys. Rev. Lett.*, 39(19):1201–1205 (1977).
- [2] M. R. Beasley, J. E. Mooij, and T. P. Orlando. Possibility of Vortex-Antivortex Pair Dissociation in Two-Dimensional Superconductors. *Phys. Rev. Lett.*, 42(17):1165–1168 (1979).
- [3] D. J. Resnick, J. C. Garland, J. T. Boyd, et al. Kosterlitz-Thouless Transition in Proximity-Coupled Superconducting Arrays. *Phys. Rev. Lett.*, 47(21):1542–1545 (1981).
- [4] A. Posazhennikova. Colloquium: Weakly interacting, dilute Bose gases in 2D. *Rev. Mod. Phys.*, 78(4):1111–1134 (2006).
- [5] V. L. Berezinskii. Destruction of Long-range Order in One-dimensional and Two-dimensional Systems Possessing a Continuous Symmetry Group. *J. Exp. Theor. Phys.*, 61(3):1144 (1972).
- [6] J. M. Kosterlitz and D. J. Thouless. Ordering, metastability and phase transitions in two-dimensional systems. *J. Phys. C Solid State Phys.*, 6(7):1181–1203 (1973).
- [7] M. H. Anderson, J. R. Ensher, M. R. Matthews, et al. Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor. *Science*, 269(5221):198 LP – 201 (1995).
- [8] I. Bloch, J. Dalibard, and W. Zwerger. Many-body physics with ultracold gases. *Rev. Mod. Phys.*, 80(3):885–964 (2008).
- [9] I. Bloch, J. Dalibard, and S. Nascimbène. Quantum simulations with ultracold quantum gases. *Nat. Phys.*, 8(4):267–276 (2012).
- [10] M. R. Andrews, C. G. Townsend, H.-J. Miesner, et al. Observation of Interference Between Two Bose Condensates. *Science*, 275(5300):637 LP – 641 (1997).
- [11] Z. Hadzibabic, P. Krüger, M. Cheneau, et al. Berezinskii–Kosterlitz–Thouless crossover in a trapped atomic gas. *Nature*, 441(7097):1118–1121 (2006).
- [12] S. Sunami, V. P. Singh, D. Garrick, et al. Observation of the Berezinskii–Kosterlitz–Thouless Transition in a Two-Dimensional Bose Gas via Matter-Wave Interferometry. *Phys. Rev. Lett.*, 128:250402 (2022).

-
- [13] S. Sunami, V. P. Singh, D. Garrick, et al. Universal scaling of the dynamic BKT transition in quenched 2D Bose gases. *Science*, 382(6669):443–447 (2023).
- [14] S. P. Rath and W. Zwerger. Full counting statistics of the interference contrast from independent Bose-Einstein condensates. *Phys. Rev. A*, 82(5):053622 (2010).
- [15] G. Carleo, G. Bo  ris, M. Holzmann, and L. Sanchez-Palencia. Universal Superfluid Transition and Transport Properties of Two-Dimensional Dirty Bosons. *Phys. Rev. Lett.*, 111(5):050406 (2013).
- [16] L. Mathey and A. Polkovnikov. Light cone dynamics and reverse Kibble-Zurek mechanism in two-dimensional superfluids following a quantum quench. *Phys. Rev. A*, 81(3):33605 (2010).
- [17] N. D. Mermin and H. Wagner. Absence of Ferromagnetism or Antiferromagnetism in One- or Two-Dimensional Isotropic Heisenberg Models. *Phys. Rev. Lett.*, 17(22):1133–1136 (1966).
- [18] J. V. Jose. *40 years of Berezinskii-Kosterlitz-Thouless theory*. World Scientific, Singapore (2013).
- [19] I. Boettcher and M. Holzmann. Quasi-long-range order in trapped two-dimensional Bose gases. *Phys. Rev. A*, 94(1):11602 (2016).
- [20] E. P. Gross. Structure of a quantized vortex in boson systems. *Il Nuovo Cimento Ser.10*, 20(3):454–477 (1961).
- [21] L. P. Pitaevskii. Vortex Lines in an imperfect Bose gas. *J. Exp. Theor. Phys.*, 13(2):451–454 (1961).
- [22] Y. Kagan, V. A. Kashurnikov, A. V. Krasavin, et al. Quasicondensation in a two-dimensional interacting Bose gas. *Phys. Rev. A*, 61(4):43608 (2000).
- [23] A. I. Safonov, S. A. Vasilyev, I. S. Yasnikov, et al. Observation of Quasicondensate in Two-Dimensional Atomic Hydrogen. *Phys. Rev. Lett.*, 81(21):4545–4548 (1998).
- [24] C. Pethick and H. Smith. *Bose-Einstein condensation in dilute gases*. Cambridge University Press, Cambridge, 2nd edition (2008).
- [25] A. L. Fetter. Ground state and excited states of a confined condensed Bose gas. *Phys. Rev. A*, 53(6):4245–4249 (1996).
- [26] P. Minnhagen. The two-dimensional Coulomb gas, vortex unbinding, and superfluid-superconducting films. *Rev. Mod. Phys.*, 59:1001 (1987).

- [27] Y. Shin, M. Saba, M. Vengalattore, et al. Dynamical Instability of a Doubly Quantized Vortex in a Bose-Einstein Condensate. *Phys. Rev. Lett.*, 93(16):160406 (2004).
- [28] J. M. Kosterlitz. The critical properties of the two-dimensional XY model. *J. Phys. C Solid State Phys.*, 7(6):1046 (1974).
- [29] Kenneth G. Wilson. Renormalization Group and Critical Phenomena. I. Renormalization Group and the Kadanoff Scaling Picture. *Phys. Rev. B*, 4:3174–3183 (1971).
- [30] Kenneth G. Wilson. Renormalization Group and Critical Phenomena. II. Phase-Space Cell Analysis of Critical Behavior. *Phys. Rev. B*, 4:3184–3205 (1971).
- [31] N. Prokof'ev and B. Svistunov. Two-dimensional weakly interacting Bose gas in the fluctuation region. *Phys. Rev. A*, 66(4):43608 (2002).
- [32] I. Maccari, N. Defenu, L. Benfatto, et al. Interplay of spin waves and vortices in the two-dimensional XY model at small vortex-core energy. *Phys. Rev. B*, 102(10):104505 (2020).
- [33] J. Kondo. Resistance Minimum in Dilute Magnetic Alloys. *Prog. Theor. Phys.*, 32(1):37–49 (1964).
- [34] P. W. Anderson. Absence of diffusion in certain random lattices. *Phys. Rev.*, 109(5):1492–1505 (1958).
- [35] M. Imada, A. Fujimori, and Y. Tokura. Metal-insulator transitions. *Rev. Mod. Phys.*, 70(4):1039–1263 (1998).
- [36] S. H. Pan, J. P. O'Neal, R. L. Badzey, et al. Microscopic electronic inhomogeneity in the high-T_c superconductor Bi₂Sr₂CaCu₂O_{8+x}. *Nature*, 413(6853):282–285 (2001).
- [37] A. Weinrib. Percolation threshold of a two-dimensional continuum system. *Phys. Rev. B*, 26(3):1352–1361 (1982).
- [38] T. Bourdel. Phase diagrams of two-dimensional and three-dimensional disordered Bose gases in the local density approximation. *Phys. Rev. A*, 86(6):063626 (2012).
- [39] B. Allard, T. Plisson, M. Holzmann, et al. Effect of disorder close to the superfluid transition in a two-dimensional Bose gas. *Phys. Rev. A*, 85(3):1–5 (2012).

- [40] N. Prokof'ev, O. Ruebenacker, and B. Svistunov. Critical Point of a Weakly Interacting Two-Dimensional Bose Gas. *Phys. Rev. Lett.*, 87(27):270402 (2001).
- [41] G. Bertoli, V. P. Michal, B. L. Altshuler, and G. V. Shlyapnikov. Finite-Temperature Disordered Bosons in Two Dimensions. *Phys. Rev. Lett.*, 121(3):30403 (2018).
- [42] P. A. Crowell, F. W. Van Keuls, and J. D. Reppy. Onset of superfluidity in ^4He films adsorbed on disordered substrates. *Phys. Rev. B*, 55(18):12620–12634 (1997).
- [43] M. C. Beeler, M. E.W. Reed, T. Hong, and S. L. Rolston. Disorder-driven loss of phase coherence in a quasi-2D cold atom system. *New J. Phys.*, 14 (2012).
- [44] A. Polkovnikov, E. Altman, and E. Demler. Interference between independent fluctuating condensates. *Proc. Natl Acad. Sci.*, 103(16):6125 LP – 6129 (2006).
- [45] Y. D. van Nieuwkerk, J. Schmiedmayer, and F. H. L. Essler. Projective phase measurements in one-dimensional Bose gases. *SciPost Phys.*, 5:046 (2018).
- [46] M. Gring, M. Kuhnert, T. Langen, et al. Relaxation and prethermalization in an isolated quantum system. *Science*, 337(6100):1318 (2012).
- [47] T. Langen, R. Geiger, M. Kuhnert, et al. Local emergence of thermal correlations in an isolated quantum many-body system. *Nat. Phys.*, 9(10):640–643 (2013).
- [48] L. Mathey, A. Polkovnikov, and A. H. C. Neto. Phase-locking transition of coupled low-dimensional superfluids. *Europhys. Lett.*, 81(1):10008 (2007).
- [49] G. Bighin, N. Defenu, I. Nándori, et al. Berezinskii-Kosterlitz-Thouless Paired Phase in Coupled XY Models. *Phys. Rev. Lett.*, 123(10):100601 (2019).
- [50] L. Mathey, K. J. Günter, J. Dalibard, and A. Polkovnikov. Dynamic Kosterlitz-Thouless transition in two-dimensional Bose mixtures of ultracold atoms. *Phys. Rev. A*, 95(5):1–7 (2017).
- [51] K. Brown, T. Bland, P. Comaron, and N. P. Proukakis. Periodic quenches across the Berezinskii-Kosterlitz-Thouless phase transition. *Phys. Rev. Res.*, 3(1):13097 (2021).

- [52] L. Benfatto, C. Castellani, and T. Giamarchi. Kosterlitz-Thouless Behavior in Layered Superconductors: The Role of the Vortex Core Energy. *Phys. Rev. Lett.*, 98(11):117008 (2007).
- [53] H. A. Bethe and E. E. Salpeter. *Quantum mechanics of one- and two-electron atoms*. Plenum, New York (1977).
- [54] L. Labzowsky, I. Goidenko, and P. Pyykkö. Estimates of the bound-state QED contributions to the g-factor of valence ns electrons in alkali metal atoms. *Phys. Lett. Sect. Gen. At. Solid State Phys.*, 258(1):31–37 (1999).
- [55] W. H. Wing. On neutral particle trapping in quasistatic electromagnetic fields. *Prog. Quantum Electron.*, 8(3):181–199 (1984).
- [56] K. B. Davis, M. O. Mewes, M. A. Joffe, et al. Evaporative cooling of sodium atoms. *Phys. Rev. Lett.*, 75(26):2909 (1995).
- [57] D. M. Brink and C. V. Sukumar. Majorana spin-flip transitions in a magnetic trap. *Phys. Rev. A*, 74(3):35401 (2006).
- [58] D. E. Pritchard. Cooling neutral atoms in a magnetic trap for precision spectroscopy. *Phys. Rev. Lett.*, 51:1336–1339 (1983).
- [59] W. Petrich, M. H. Anderson, J. R. Ensher, and E. A. Cornell. Stable, tightly confining magnetic trap for evaporative cooling of neutral atoms. *Phys. Rev. Lett.*, 74:3352–3355 (1995).
- [60] K. B. Davis, M. O. Mewes, M. R. Andrews, et al. Bose-Einstein Condensation in a Gas of Sodium Atoms. *Phys. Rev. Lett.*, 75:3969–3973 (1995).
- [61] C. Cohen-Tannoudji and S. Reynaud. Dressed-atom description of resonance fluorescence and absorption spectra of a multi-level atom in an intense laser beam. *J. Phys. B At. Mol. Phys.*, 10(3):345–363 (1977).
- [62] T. Yabuzaki, S. Nakayama, Y. Murakami, and T. Ogawa. Interaction between a spin-1/2 atom and a strong RF field. *Phys. Rev. A*, 10(6):1955–1963 (1974).
- [63] C. Zener. Non-adiabatic crossing of energy levels. *Proc. R. Soc. Lond. Ser. A*, 137(833):696–702 (1932).
- [64] M. Born and V. Fock. Beweis des adiabatensatzes. *Z. Phys.*, 51(3):165–180 (1928).
- [65] H. Perrin and B. M. Garraway. Chapter four - trapping atoms with radio frequency adiabatic potentials. volume 66 of *Advances In Atomic, Molecular, and Optical Physics*, pages 181–262. Academic Press, (2017).

- [66] F. Bloch and A. Siegert. Magnetic resonance for nonrotating fields. *Phys. Rev.*, 57(6):522–527 (1940).
- [67] J. H. Shirley. Solution of the Schrödinger Equation with a Hamiltonian Periodic in Time. *Phys. Rev.*, 138:B979–B987 (1965).
- [68] K. Merloti, R. Dubessy, L. Longchambon, et al. A two-dimensional quantum gas in a magnetic trap. *New J. Phys.*, 15(3):33007 (2013).
- [69] O. Morizot, C. L. Garrido Alzar, P-E. Pottie, et al. Trapping and cooling of rf-dressed atoms in a quadrupole magnetic field. *J. Phys. B At. Mol. Opt. Phys.*, 40(20):4013 (2007).
- [70] Y. Colombe, E. Knyazchyan, O. Morizot, et al. Ultracold atoms confined in rf-induced two-dimensional trapping potentials. *Europhysics Letters*, 67(4):593 (2004).
- [71] M. Gildemeister, E. Nugent, B. E. Sherlock, et al. Trapping ultracold atoms in a time-averaged adiabatic potential. *Phys. Rev. A*, 81:031402 (2010).
- [72] Ph. W. Courteille, B. Deh, J. Fortágh, et al. Highly versatile atomic micro traps generated by multifrequency magnetic field modulation. *J. Phys. B At. Mol. Phys.*, 39(5):1055–1064 (2006).
- [73] E. Bentine. *Atomic mixtures in radiofrequency dressed potentials*. PhD thesis, University of Oxford (2018).
- [74] S. Sunami. *Non-Equilibrium Dynamics in Two-Dimensional Quantum Systems*. PhD thesis, University of Oxford (2021).
- [75] A. J. Barker, S. Sunami, D. Garrick, et al. Coherent splitting of two-dimensional Bose gases in magnetic potentials. *New J. Phys.*, 22(10):103040 (2020).
- [76] D. Garrick. *Two-dimensional Bose gases in systems with pointlike disorder*. PhD thesis, University of Oxford (2022).
- [77] D. M. Farkas, E. A. Salim, and J. Ramirez-Serrano. Production of Rubidium Bose-Einstein Condensates at a 1 Hz Rate. *arXiv:1403.4641* (2014).
- [78] A. J. Barker, H. Style, K. Luksch, et al. Applying machine learning optimization methods to the production of a quantum gas. *Mach. Learn. Sci. Technol.*, 1(1):015007 (2020).
- [79] S. Ravenhall. *A compact, high-flux source of cold atoms*. PhD thesis, University of Oxford (2018).

- [80] H. A. Steinherz. *Handbook of high vacuum engineering*. Reinhold Pub. Corp., New York (State) (1963).
- [81] A. N. Nesmeianov and R. Gary. *Vapor pressure of the chemical elements*. Elsevier Pub. Co, Amsterdam; London (1963).
- [82] E. Bentine, A. J. Barker, K. Luksch, et al. Inelastic collisions in radiofrequency-dressed mixtures of ultracold atoms. *Phys. Rev. Res.*, 2:033163 (2020).
- [83] D. J. McCarron, S. A. King, and S. L. Cornish. Modulation transfer spectroscopy in atomic rubidium. *Measurement Science and Technology*, 19(10):105601 (2008).
- [84] R. H. Pantell. The laser oscillator with an external signal. *Proc. IEEE*, 53(5):474–477 (1965).
- [85] H. J. Lewandowski, D. M. Harber, D. L. Whitaker, and E. A. Cornell. Simplified system for creating a Bose-Einstein condensate. *Journal of low temperature physics*, 132(5-6):309–367 (2003).
- [86] D. P. George. The Micro Python Board, (2015).
- [87] D. P. George. The micropython project. <https://github.com/micropython>, (2013).
- [88] Y. Colombe, E. Knyazchyan, O. Morizot, et al. Ultracold atoms confined in rf-induced two-dimensional trapping potentials. *Europhys. Lett.*, 67(4):593–599 (2004).
- [89] H. Ott, J. Fortagh, G. Schlotterbeck, et al. Bose-Einstein Condensation in a Surface Microtrap. *Phys. Rev. Lett.*, 87(23):230401 (2001).
- [90] M. Keil, O. Amit, S. Zhou, et al. Fifteen years of cold matter on the atom chip: promise, realizations, and prospects. *J. mod. opt.*, 63(18):1840–1885 (2016).
- [91] Silicon Labs. magpylib. <https://github.com/magpylib/magpylib>, (2019).
- [92] R. K. Easwaran, L. Longchambon, P-E. Pottie, et al. RF spectroscopy in a resonant RF-dressed trap. *J. Phys. B At. Mol. Phys.*, 43(6):65302 (2010).
- [93] S. Bourdeauducq et al. ARTIQ 1.0., (2016).
- [94] F. Nogueira. Bayesian Optimization: Open source constrained global optimization tool for Python, (2014–).

- [95] P. Virtanen et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Method*, 17:261–272 (2020).
- [96] R. Storn and K. Price. Differential Evolution - A Simple and Efficient Heuristic for global Optimization over Continuous Spaces. *J. Glob. Optim.*, 11(4):341–359 (1997).
- [97] T. Head, M. Kumar, H. Nahrstaedt, et al. scikit-optimize/scikit-optimize, (2021).
- [98] F. Pedregosa, G. Varoquaux, A. Gramfort, et al. Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.*, 12:2825–2830 (2011).
- [99] I. M. Sobol. On the distribution of points in a cube and the approximate evaluation of integrals. *Ussr Comput. Math. Math. Phys.*, 7:86–112 (1967).
- [100] R. Houtz, C. Chan, and H. Müller. Wideband, Efficient Optical Serrodyne Frequency Shifting with a Phase Modulator and a Nonlinear Transmission Line. *Opt. Express*, 17(21):19235–19240 (2009).
- [101] E. Rydow. *In preparation*. PhD thesis, University of Oxford.
- [102] D. Rey, S. Thomas, R. Sharma, et al. Loading a quantum gas from a hybrid dimple trap to a shell trap. *J. Appl. Phys.*, 132(21):214401 (2022).
- [103] W. Petrich, M. H. Anderson, J. R. Ensher, and E. A. Cornell. Behavior of atoms in a compressed magneto-optical trap. *J. Opt. Soc. Am. B*, 11(8):1332–1335 (1994).
- [104] P. D. Lett, R. N. Watts, C. I. Westbrook, et al. Observation of atoms laser cooled below the doppler limit. *Phys. Rev. Lett.*, 61:169–172 (1988).
- [105] K. Luksch. *Investigating the spectrum of atoms in multiple-radiofrequency dressed potentials for coherent splitting of quantum gases*. PhD thesis, University of Oxford (2019).
- [106] K. B. Davis, M-O. Mewes, and W. Ketterle. An analytical model for evaporative cooling of atoms. *Appl. phys. B Laser opt.*, 60(2-3):155–159 (1995).
- [107] E. L. Surkov, J. T. M. Walraven, and G. V. Shlyapnikov. Collisionless motion and evaporative cooling of atoms in magnetic traps. *Phys. Rev. A*, 53:3403–3408 (1996).
- [108] N. R. Newbury, C. J. Myatt, and C. E. Wieman. s-wave elastic collisions between cold ground-state ^{87}Rb atoms. *Phys. Rev. A*, 51:R2680–R2683 (1995).

- [109] Z.-Y. Ma, A. M. Thomas, C. J. Foot, and S. L. Cornish. The evaporative cooling of a gas of caesium atoms in the hydrodynamic regime. *J. Phys. B At. Mol. Opt. Phys.*, 36(16):3533 (2003).
- [110] V. V. Goldman, I. F. Silvera, and A. J. Leggett. Atomic hydrogen in an inhomogeneous magnetic field: Density profile and Bose-Einstein condensation. *Phys. Rev. B*, 24(5):2870–2873 (1981).
- [111] D. T. Smithey, M. Beck, M. G. Raymer, and A. Faridani. Measurement of the Wigner distribution and the density matrix of a light mode using optical homodyne tomography: Application to squeezed states and the vacuum. *Phys. Rev. Lett.*, 70(9):1244–1247 (1993).
- [112] C. Schwemmer, G. Tóth, A. Niggebaum, et al. Experimental Comparison of Efficient Tomography Schemes for a Six-Qubit State. *Phys. Rev. Lett.*, 113(4):40503 (2014).
- [113] V. F. Maisi, D. Kambly, C. Flindt, and J. P. Pekola. Full Counting Statistics of Andreev Tunneling. *Phys. Rev. Lett.*, 112(3):36801 (2014).
- [114] S. Gustavsson, R. Leturcq, B. Simovic, et al. Counting Statistics of Single Electron Transport in a Quantum Dot. *Phys. Rev. Lett.*, 96(7):76605 (2006).
- [115] S. Hofferberth, I. Lesanovsky, T. Schumm, et al. Probing quantum and thermal noise in an interacting many-body system. *Nat. Phys.*, 4(6):489–495 (2008).
- [116] T. Kitagawa, A. Imambekov, J. Schmiedmayer, and E. Demler. The dynamics and prethermalization of one-dimensional quantum systems probed through the full distributions of quantum noise. *New J. Phys.*, 13 (2011).
- [117] M. Gring, M. Kuhnert, T. Langen, et al. Relaxation and Prethermalization in an Isolated Quantum System. *Science*, 337(6100):1318–1322 (2012).
- [118] A. Imambekov, V. Gritsev, and E. Demler. Mapping of Coulomb gases and sine-Gordon models to statistics of random surfaces. *Phys. Rev. A*, 77(6):1–5 (2008).
- [119] A. J. Barker. *Bose gases in double-well potentials*. PhD thesis, University of Oxford (2020).
- [120] N. A. Keeper, I-K. Liu, F. Dalfovo, and N. P. Proukakis. Phase transition dimensionality crossover from two to three dimensions in a trapped ultracold atomic Bose gas. *Phys. Rev. Res.*, 4(3):33130 (2022).

- [121] M. Holzmann, M. Chevallier, and W. Krauth. Semiclassical theory of the quasi-two-dimensional trapped Bose gas. *Europhys. Lett.*, 82(3):30001 (2008).
- [122] X. Li, M. Ke, B. Yan, and Y. Wang. Reduction of interference fringes in absorption imaging of cold atom cloud using eigenface method. *Chin. Opt. Lett.*, 5(3):128–130 (2007).
- [123] I. E. Mazets. Two-dimensional dynamics of expansion of a degenerate Bose gas. *Phys. Rev. A*, 86(5):55603 (2012).
- [124] E. Chang. *In preparation*. PhD thesis, University of Oxford.
- [125] X. S. Chen, V. Dohm, and N. Schultka. Order-Parameter Distribution Function of Finite $O(n)$ Symmetric Systems. *Phys. Rev. Lett.*, 77(17):3641–3644 (1996).
- [126] K. Binder. Finite size scaling analysis of Ising model block distribution functions. *Z. Phys. B Condens. Matter*, 43(2):119–140 (1981).
- [127] T. Langen, S. Erne, R. Geiger, et al. Experimental observation of a generalized Gibbs ensemble. *Science*, 348(6231):207–211 (2015).
- [128] J. E. Lye, L. Fallani, M. Modugno, et al. Bose-Einstein Condensate in a Random Potential. *Phys. Rev. Lett.*, 95(7):70401 (2005).
- [129] G. Martirosyan, C. J. Ho, J. Etrych, et al. Observation of subdiffusive dynamic scaling in a driven and disordered box-trapped Bose gas. *arXiv:2304.06697* (2023).
- [130] S. Sunami, V. P. Singh, A. Beregi, et al. Observation of a bilayer superfluid with interlayer coherence. *In preparation* (2023).
- [131] P. Christodoulou. *Superfluidity in a uniform two-dimensional Bose gas*. PhD thesis, University of Cambridge (2022).
- [132] L. Chomaz, L. Corman, T. Bienaimé, et al. Emergence of coherence via transverse condensation in a uniform quasi-two-dimensional Bose gas. *Nat. Commun.*, 6(1):6162 (2015).
- [133] T. Harte. *Ultracold atoms in dressed potentials*. PhD thesis, University of Oxford (2017).
- [134] T. Haga and M. Ueda. Effective temperature of a superfluid flowing in a random potential. *Phys. Rev. Res.*, 2(4):43316 (2020).

- [135] M. Piraud, L. Pezzé, and L. Sanchez-Palencia. Quantum transport of atomic matter waves in anisotropic two-dimensional and three-dimensional disorder. *New J. Phys.*, 15(7):75007 (2013).
- [136] B. Nagler, S. Barbosa, J. Koch, et al. Observing the loss and revival of long-range phase coherence through disorder quenches. *Proc. National Acad. Sci.*, 119(1):1– (2022).
- [137] J. L. Ville, R. Saint-Jalm, É. Le Cerf, et al. Sound Propagation in a Uniform Superfluid Two-Dimensional Bose Gas. *Phys. Rev. Lett.*, 121:145301 (2018).
- [138] P. Christodoulou, M. Gałka, N. Dogra, et al. Observation of first and second sound in a BKT superfluid. *Nature*, 594(7862):191–194 (2021).
- [139] J. Stenger, S. Inouye, A. P. Chikkatur, et al. Bragg spectroscopy of a Bose-Einstein condensate. *Phys. Rev. Lett.*, 82(23):4569–4573 (1999).
- [140] J. Beugnon and N. Navon. Exploring the Kibble-Zurek mechanism with homogeneous Bose gases. *J. Phys. B At. Mol. Phys.*, 50(2) (2017).
- [141] S. P. Rath, T. Yefsah, K. J. Günter, et al. Equilibrium state of a trapped two-dimensional Bose gas. *Phys. Rev. A*, 82(1):13609 (2010).
- [142] J. D. Jackson. *Classical electrodynamics*. Wiley, New York ; Chichester, 3rd ed. edition (1999).
- [143] D. S. Petrov, M. Holzmann, and G. V. Shlyapnikov. Bose-Einstein Condensation in Quasi-2D Trapped Gases. *Phys. Rev. Lett.*, 84(12):2551–2555 (2000).
- [144] K. Frye, S. Abend, W. Bartosch, et al. The Bose-Einstein Condensate and Cold Atom Laboratory. *EPJ Quantum Technol.*, 8(1):1 (2021).
- [145] R. A. Nyman, G. Varoquaux, F. Lienhart, et al. I.C.E.: a transportable atomic inertial sensor for test in microgravity. *Appl. Phys. B*, 84(4):673–681 (2006).
- [146] G. Sinuco-León and B. M. Garraway. Radio-frequency dressed atoms beyond the linear Zeeman effect. *New J. Phys.*, 14(12):123008 (2012).
- [147] M. B. Squires, S. E. Olson, B. Kasch, et al. Ex vacuo atom chip Bose-Einstein condensate. *Appl. Phys. Lett.*, 109(26) (2016).
- [148] S. E. Olson. Private communication.
- [149] A. L. Fetter. Nonuniform states of an imperfect Bose gas. *Ann. Phys.*, 70(1):67–101 (1972).

- [150] M. P. A. Fisher, P. B. Weichman, G. Grinstein, and D. S. Fisher. Boson localization and the superfluid-insulator transition. *Phys. Rev. B*, 40(1):546–570 (1989).
- [151] M. L. Chiofalo, S. Succi, and M. P. Tosi. Ground state of trapped interacting Bose-Einstein condensates by an explicit imaginary-time algorithm. *Phys. Rev. E*, 62(5):7438–7444 (2000).
- [152] M. Suzuki. Generalized Trotter’s formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems. *Commun. Math. Phys.*, 51(2):183–190 (1976).
- [153] M. Frigo and S. G. Johnson. The design and implementation of FFTW3. *Proc. IEEE*, 93(2):216–231 (2005).
- [154] S. K. Lam, S. Archibald, A. Pitrou, et al. numba/numba: Version 0.57.1, (2023).
- [155] R. Okuta, Y. Unno, D. Nishino, et al. CuPy: A NumPy-Compatible Library for NVIDIA GPU Calculations. In *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*, (2017).
- [156] George B. Arfken, Hans J. Weber, and Frank E. Harris. Chapter 20 - integral transforms. In George B. Arfken, Hans J. Weber, and Frank E. Harris, editors, *Mathematical Methods for Physicists (Seventh Edition)*, pages 963–1046. Academic Press, Boston, seventh edition, (2013).
- [157] N. Baddour and U. Chouinard. Theory and operational rules for the discrete Hankel transform. *J. Opt. Soc. Am. A*, 32(4):611–622 (2015).
- [158] D.C. Sorensen. 1 - implicitly restarted arnoldi/lanczos methods and large scale svd applications. In *SVD and Signal Processing, III*, pages 21–31. Elsevier Science, (1995).
- [159] B. Svistunov, E. Babaev, and N. Prokof’ev. *Superfluid states of matter*. CRC Press, Boca Raton, 1st edition.
- [160] C. Mora and Y. Castin. Extension of Bogoliubov theory to quasicondensates. *Phys. Rev. A*, 67(5):53615 (2003).