

OKINAWA INSTITUTE OF SCIENCE AND TECHNOLOGY
GRADUATE UNIVERSITY

Thesis submitted for the degree

Doctor of Philosophy

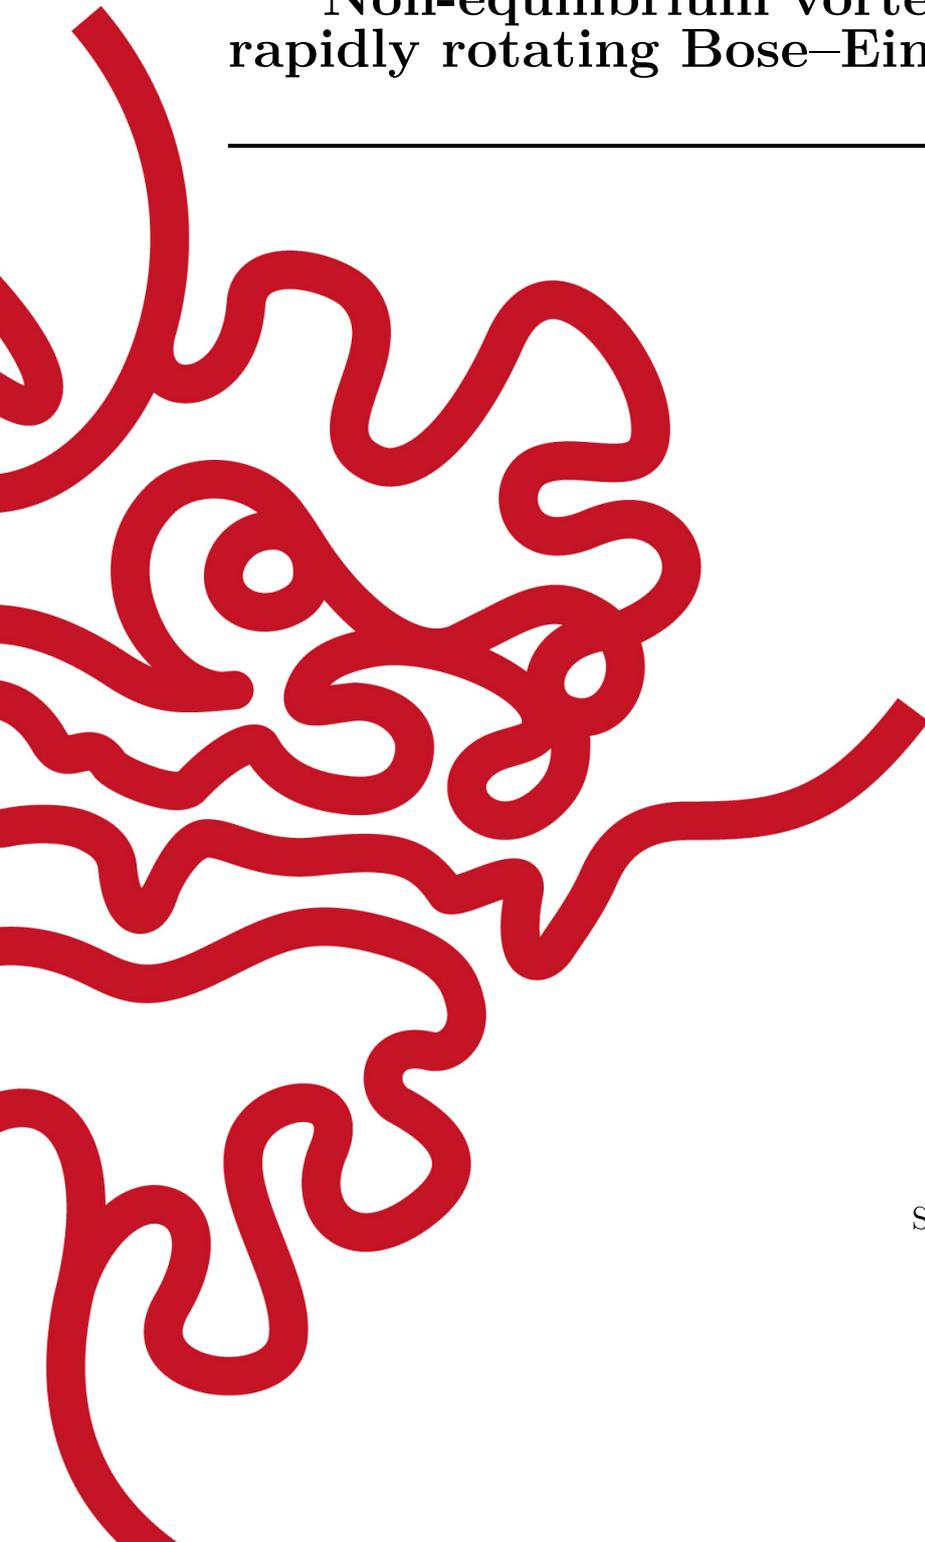
**Non-equilibrium vortex dynamics in
rapidly rotating Bose–Einstein condensates**

by

Lee James O’Riordan

Supervisor: **Prof. Thomas Busch**

February, 2017



Declaration of Original and Sole Authorship

I, Lee James O’Riordan, declare that this thesis entitled *Non-equilibrium vortex dynamics in rapidly rotating Bose–Einstein condensates* and the data presented in it are original and my own work.

I confirm that:

- This work was done solely while a candidate for the research degree at the Okinawa Institute of Science and Technology Graduate University, Japan and the University College Cork, Ireland.
- No part of this work has previously been submitted for a degree at this or any other university.
- References to the work of others have been clearly attributed. Quotations from the work of others have been clearly indicated, and attributed to them.
- In cases where others have contributed to part of this work, such contribution has been clearly acknowledged and distinguished from my own work.

Date: February, 2017

Signature:

Abstract

Non-equilibrium vortex dynamics in rapidly rotating Bose–Einstein condensates

This body of work examines the non-equilibrium dynamics of vortex lattice carrying Bose–Einstein condensates. We solve the mean-field Gross–Pitaevskii equation for a two-dimensional pancake geometry, in the co-rotating frame within the limit of high rotation frequencies. The condensate responds to this by creating a large periodic lattice of vortices with 6-fold triangular symmetry. By applying two distinct perturbations to this lattice, we examine the resulting effects on the vortices during time evolution. The first perturbation involves applying an optical potential with matching geometry to the vortex lattice. We observe the appearance of interference fringes, and we show that these can be described by moiré interference theory. This is backed up by a decomposition of the kinetic energy spectra of the condensate. The applied perturbation only modifies the condensate density, with the vortex positions largely unaffected. From this we conclude that the vortex lattice is very stable and robust against phononic disturbances.

Next, by removing vortices at predefined positions in the lattice using phase imprinting techniques, we examine the resulting order of the lattice. By performing this we generate stable topological defects in the crystal structure. The resulting lattice remains highly ordered in the presence of low numbers of these defects, where crystal structure and order of the lattice shows to be highly robust. By varying the type of imprinted phases we can create controllable degrees of disorder in the lattice. This disorder is analysed using orientational correlations, Delaunay triangulation, and Voronoi diagrams of the vortex lattice, and demonstrates a method for examining order and generating disorder in vortex lattices in Bose–Einstein condensates.

All work described makes extensive use of GPU computing techniques, and allows for the simulation of these systems to be realised in short times. The implementation of the calculations using GPU computing are also discussed, where the software is shown to be the fastest of its kind out of the independently tested software suites.

License and copyright information

The work contained in this thesis includes materials that have been published elsewhere. Notably, Chapters 3, 5 and 6 include figures and text that are published in adapted and original form in the American Physical Society's journal *Physical Review A*. Copyright permission has been obtained for use of these figures, and for brevity the copyright notice is given here.

Chapter 3. Numerical methods

Reprinted excerpts and figures with permission from Tadhg Morgan, Lee James O'Riordan, Neil Crowley, Brian O'Sullivan and Thomas Busch, *Physical Review A* **88**, 053618 (2013). Copyright (2013) by the American Physical Society.

Chapter 5. Moiré superlattice structures

Reprinted excerpts and figures with permission from Lee James O'Riordan, Angela White and Thomas Busch, *Physical Review A* **93**, 023609 (2016). Copyright (2016) by the American Physical Society.

Chapter 6. Defect engineering of the vortex lattice

Reprinted excerpts and figures with permission from Lee James O' Riordan, and Thomas Busch, *Physical Review A* **94**, 053603 (2016). Copyright (2016) by the American Physical Society.

Acknowledgments

/* Firstly, I would like to express my sincerest gratitude to my family back in the 'mel who have always been there for me, and provided amazing support from across the planet since I started my PhD (and before I started too) — most notably, to Mum, Nan, and Maria, who have helped so much over the past few years. Next, my sincerest gratitude and thanks goes to Prof. Thomas Busch, for freeing me from the working world and introducing me to the (ultra)cool world of atomic physics, as well as the guidance he has offered over the years. To past and present members of both the Ultracold Quantum Gases group at University College Cork, and Quantum Systems Unit at OIST, thanks for keeping me sane and helping out through the work and writing over the years — in no particular order: Mossy, Steve, Jeremie, Tara, Albert, Angela, Sawako, Rashi, James, Irina, the magnificent Dave Rea, and of course Tadhg. Next, I would like to thank the Graduate School for all of their help removing bureaucratic dealings from my everyday life here. A special thanks goes to IT & the Scientific Computing Section, for allowing me access to some very amazing toys over these past few years. A special thanks also to Annie, for supplying me with enough caffeine to get to the moon and back. Lastly, but not least-ly, to Christina for her compassion, understanding, and companionship over the past few years. */

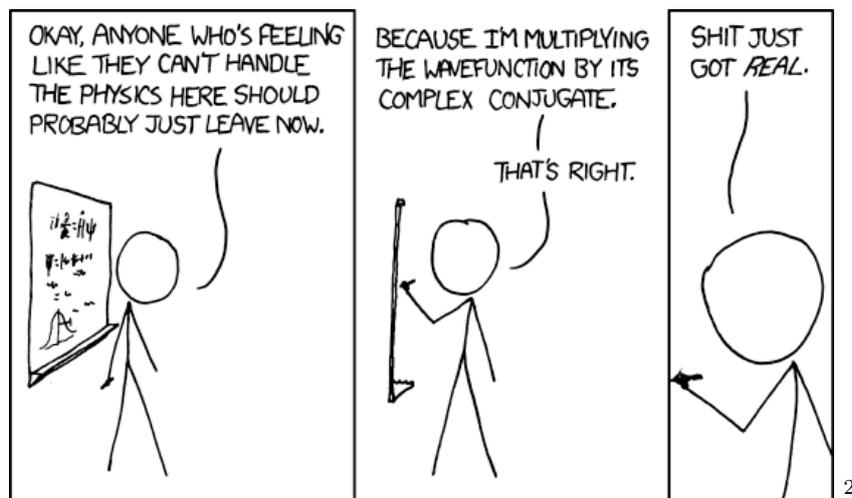
for all $i \in \{S | S \in \text{People to thank}\}$ **do**

$i \leftarrow$ ありがとうございます!

end for

Up-Goer Five abstract

I work at making really small things act like a big thing, and go around and around until they give an empty spot in the big thing. The more around and around it goes, the more spots appear in the big thing and become ordered. When I have lots of ordered spots I kick the big thing with a strong light, and watch the spots and big thing move. When the light is in the right position, cells appear in the big thing, because of how the light and the spots work together. I also take away spots from the big thing, which causes all other ordered spots to move and break the order. I do this using really fast computer cards.¹



¹<https://xkcd.com/1133/>, <http://splasho.com/upgoer5/>

²<https://xkcd.com/849/>

" MOVEMENT IS THE KEY TO VICTORIOUS COMBAT "
- DOOM (ZDTE) ³

³AmazDooM font, CC BY-NC 3.0

Contents

Declaration of Original and Sole Authorship	iii
Abstract	v
License and copyright	vii
Acknowledgment	ix
Up-Goer Five abstract	xi
Contents	xv
1 Introduction	1
2 Background	5
2.1 Ultracold atoms	5
2.1.1 Cooling of atomic gases	5
2.1.2 Introduction to Bose–Einstein condensation	5
2.1.3 Theoretical description of BECs: Gross–Pitaevskii equation	7
2.1.4 Bogoliubov-de Gennes equations	11
2.1.5 Lower dimensional condensates	12
2.2 Superfluidity	12
2.2.1 Introduction to superfluidity	12
2.2.2 Vortices in Bose–Einstein condensates	15
2.2.3 Vortex lattices	16
2.3 Recent progress and outlook	18
3 Numerical methods	19
3.1 Time evolution	20
3.2 Fourier split-operator method	21
3.2.1 Resolution considerations	24
3.3 Parallel computing	25

3.3.1	Parallel operations	26
3.4	CUDA programming model	27
3.5	GPU-enabled parallel Schrödinger simulations	31
3.5.1	Spatial adiabatic passage	31
3.5.2	Atom-chip model	34
3.5.3	3D Simulations	35
3.5.4	GPU computing performance	36
3.6	GPUE: GPU Gross–Pitaevskii equation solver	39
3.6.1	Angular momentum operators using Fourier split-operator method	41
3.6.2	Vortex tracking	44
4	Bose–Einstein condensate dynamics	49
4.1	Simulating Bose–Einstein condensate dynamics	49
4.1.1	Few vortex condensates	50
4.2	Rapidly rotating vortex lattice	51
4.2.1	Model system	51
4.3	Condensate perturbations	55
4.3.1	Trapping potential control	55
4.3.2	Direct phase manipulation	60
4.4	Condensate analysis	63
4.4.1	Kinetic energy decomposition	63
4.4.2	Delaunay triangulation and Voronoi tessellation	65
5	Moiré superlattice structures	69
5.1	Introduction	69
5.2	Model	70
5.3	Delta-kick dynamics	72
5.3.1	Non-rotating condensate	72
5.3.2	Rapidly rotating condensate	74
5.4	Vortex dynamics following a kick	80
5.5	Conclusions and outlook	83
6	Defect engineering in vortex lattices	85
6.1	Introduction	85
6.2	Model	86
6.3	Phase imprinting defects	87
6.3.1	Single vortex dynamics	88
6.3.2	Lattice dynamics	89
6.4	Discussion and conclusions	99
7	Conclusions and outlook	107
7.1	Conclusions	107

7.2 Outlook 108

Bibliography **111**

Chapter 1

Introduction

The purpose of this work is to understand the dynamics of rapidly rotating Bose–Einstein condensates subjected to perturbations, and to develop techniques to control and engineer specific non-equilibrium states. While it is possible to derive some analytical solutions for rapidly rotating condensates (e.g. lowest Landau level approach), such solutions are rare. This thesis concentrates on the numerical solutions of the Gross–Pitaevskii equation, and the resulting dynamics within this framework. It focuses on gaining an understanding of the dynamical behaviour of quantum vortices in an Abrikosov geometry following a perturbation. This body of work was carried out during my time as a Ph.D student at Okinawa Institute of Science and Technology Graduate University (OIST), and grew out of work and ideas I started to pursue at University College Cork (UCC), Ireland.

Understanding ultracold Abrikosov vortex systems can help with engineering quantum states for future technologies. Ideally, these systems can be used for long-term memory storage in computing applications as individual vortices are topologically protected and, therefore, very robust. They also allow the study of quantum mechanical effects on mesoscopic scales, and the inherent periodicity makes them a promising tool for simulating condensed matter physics. Furthermore, perturbed vortex lattices can be used to investigate turbulent, and possibly chaotic, quantum behaviour. While turbulent classical systems are notoriously hard to understand and control, quantum turbulence is thought to offer a more controllable route to understanding the nature of turbulence, due to the quantisation condition of the circulation. It is therefore of large interest to develop new tools for manipulating and engineering specific states of rotating condensates. In the following work I concentrate on two types of perturbations to the equilibrium state of a rotating condensate: i) the modification of the phonon spectrum of the condensate which does not influence the angular momentum, in particular through the use of a kicked optical potential; ii) the direct control of the topological excitations, and hence the angular momentum, which is performed with direct phase engineering of the condensate wavefunction. I examine both in the above order, and investigate their usefulness in controlling and manipulating condensate dynamics.

For investigating these perturbances I assume a system of a rapidly rotating BEC having a large number of vortices, arranged in a triangular Abrikosov lattice pattern. This requires the solution of a two-dimensional partial differential equation at high grid resolution with a variety of different initial conditions and controllable perturbations. The solution of the proposed system is a non-trivial numerical problem, and requires the use of advanced numerical computing techniques to allow for results in a reasonable time. For this I make use of graphics processing unit (GPU) computing, and I will discuss the development of such tools, my numerical contributions, and compare them against conventional simulation techniques.

The thesis is organised as follows:

Background

I will first give a brief introduction to the field of cold-atomic gases, and discuss the theoretical framework to describe Bose–Einstein condensation. Emphasis will be placed on material and works relevant to the studies I have performed in this thesis. I will present a derivation of the Gross–Pitaevskii equation, used to model Bose–Einstein condensates, as well as a discussion of the Bogoliubov–de Gennes equations. I will then discuss the hydrodynamic description of the condensate, and give the hydrodynamic form of the Gross–Pitaevskii equation. Here I introduce superfluidity, and the nature of quantised vortices in these systems. I will conclude with an outlook on the cutting edge work in the field in the context of condensate trapping and control.

Numerical methods

In this chapter I will discuss methods for numerically solving the Gross–Pitaevskii equation for simulating the dynamics of Bose–Einstein condensates. The Fourier split-operator method will be introduced, as well as the need for imaginary time evolution, and considerations required to effectively simulate the condensate. Graphics processing unit (GPU) computing will be introduced here, with the implementation of the Gross–Pitaevskii equation discussed. To demonstrate the power of GPU computing we will present and solve a difficult numerical problem, namely the solution of an experimentally realistic situation for a single, ultracold atom on an atomchip, for which the treatment of the fully three-dimensional Schrödinger equation is required. The use of GPU computing makes this problem tractable in realistic times. The work focuses on the area of adiabatic control techniques, and demonstrates the use of GPU computing to describe the long-time dynamics of a system for observing matter-wave spatial adiabatic passage. This work has been published in *Phys. Rev. A* **88**, 053618 (2013) [1].

Bose–Einstein condensate dynamics

In this chapter I will examine the dynamics of Bose–Einstein condensates under rotation, and discuss the methods used for perturbing the condensate system. I begin by introducing the dynamical behaviour of the condensate in the presence of vortices, and introduce a model system used for further discussions. I present the velocity profiles and discuss some of the dynamics a condensate with many vortices is expected to follow. This will be followed by an introduction to the two main perturbation methods for the condensate that I will later use: optical kicking, and phase imprinting. I will also discuss the techniques that I use to analyse the vortex dynamics, concentrating primarily on the kinetic energy spectrum, Delaunay triangulation and Voronoi tessellation.

Moiré superlattice structures

Here I investigate effects stemming from the optical kicking of a condensate carrying a vortex lattice. The dynamics of the condensate after a kick with an optical potential of the same geometry as the vortex lattice is demonstrated, and shows little to no deviation of ideal vortex positions. However, the resulting condensate density shows the appearance of a superlattice pattern. I analyse this system, and demonstrate that the resulting superlattice pattern stems from interference between the optical kicking potential and the present vortex lattice in reciprocal space. Moiré interference theory accurately predicts the observed behaviour, and is backed up by examining the kinetic energy spectrum of the condensate. To conclude, I discuss applications of this optical kicking technique and the resulting moiré interference. The results presented in this chapter have been published in *Phys. Rev. A* **93**, 023609 (2016) [2].

Defect engineering of the vortex lattice

To investigate the robustness of a vortex lattice in a rapidly rotating BEC, I will in this chapter discuss the effect of perturbations induced by adding or removing angular momentum through phase imprinting. This technique creates lattice imperfections, with stable topological lattice defects appearing during time evolution. The behaviour of these resulting defects is investigated over long times. I show that the vortex lattice demonstrates highly robust behaviour, even in the presence of such defects. I discuss the use of this method for creating varying degrees of disorder in the lattice, and propose it as a system for investigating transitions from ordered to disordered lattice geometries. The results presented in this chapter have been published in *Phys. Rev. A* **94**, 053603 (2016) [3].

Conclusions and outlook

In this chapter I conclude the work discussed in the thesis, and discuss extensions, and future ideas for the field.

Chapter 2

Background

2.1 Ultracold atoms

2.1.1 Cooling of atomic gases

One of the major advances in experimental physics towards creating matter in extreme situations has been the cooling of trapped atoms to temperatures near absolute zero. This feat resulted from the pioneering work of C. Cohen-Tannoudji, S. Chu and W. Phillips, and earned them the Nobel Prize in Physics, 1997 [4–6]. It relies on the use of counter-propagating detuned laser fields which act upon a trapped cloud of atoms. Due to Doppler shifting of the frequencies, atoms moving towards the respective beams see resonant photons, absorb them and slow down due to the momentum absorbed. This is followed by a spontaneous emission in a random direction, for which the recoil kicks average out to zero; hence, the atoms become cooler. This technique is known as “Doppler cooling”. As a result, the atoms eventually reach a velocity below that which no photons can be absorbed from the lasers due to the change in resonance frequency and the limit imposed by the resonance width of the atomic levels, leaving a narrower velocity distribution with a peak at a lower value. Although Doppler cooling allowed temperatures to reach micro-Kelvin regimes, additional techniques, such as evaporative cooling, must be used to obtain atoms deep in the nano-Kelvin temperature range. A further discussion of these cooling methods is presented in [7, 8]. These cooling techniques allow for the creation of Bose–Einstein condensates in dilute atomic gases.

2.1.2 Introduction to Bose–Einstein condensation

Upon learning of a work by S. N. Bose on the statistical behaviour of photons, A. Einstein translated and arranged for publication of his work, and generalised it to also describe systems of ideal bosons with mass [9]. This led to the Bose–Einstein distribution, which in the framework of the grand canonical ensemble and following the description given by Pitaevskii

and Stringari [10, chap. 2], can be written as

$$\bar{n}_i = \frac{1}{e^{\beta(\epsilon_i - \mu)} - 1}, \quad (2.1)$$

where \bar{n}_i is the average occupation number of the i -th energy state, $\beta = (k_B T)^{-1}$ with k_B being the Boltzmann constant and T the temperature, ϵ_i is the i -th energy eigenvalue, and μ is the chemical potential giving the energy required to add an atom to the system for a fixed volume and entropy. The total number of particles in the system, can be evaluated by summing over the individual occupation numbers, as

$$N = \sum_i \bar{n}_i. \quad (2.2)$$

This work predicted that non-interacting, indistinguishable bosonic particles would undergo a phase transition below a critical temperature into a new phase in which all particles would occupy the same lowest lying energy state of the system. The number of atoms occupying this lowest lying state, N_0 , at a given temperature is provided by

$$N_0 \equiv \bar{n}_0 = \frac{1}{e^{\beta(\epsilon_0 - \mu)} - 1}, \quad (2.3)$$

with ϵ_0 representing the lowest energy eigenvalue. Since negative occupation numbers would be a nonphysical result, the chemical potential is limited to values of $\mu < \epsilon_0$. As μ tends to ϵ_0 , the occupation of the lowest energy state grows large. Separating the total number of atoms into the lowest lying (condensed), N_0 , and higher lying (thermal), N_T , states as

$$N = N_0 + N_T = N_0 + \sum_{i \neq 0} \bar{n}_i(T, \mu), \quad (2.4)$$

allows for a relation for the onset of Bose–Einstein condensation to be given. For a finite temperature system ($T > 0$) this will happen when the temperature drops below a critical value, T_c , where μ will approach ϵ_0 . This results in the macroscopic occupation of the lowest lying state, yielding a Bose–Einstein condensate (BEC).

Given a cloud of identical bosons at temperatures higher than T_c , the particles behave classically and exhibit hard-core scattering interactions. As they are cooled, their thermal de Broglie wavelength increases as

$$\lambda_{\text{dB}} = \sqrt{\frac{\hbar^2}{2\pi m k_B T}}, \quad (2.5)$$

where m is the atom mass, and the wave-nature of the atoms becomes more prominent. The de Broglie waves start to overlap when $\lambda_{\text{dB}} \approx$ inter-particle separation. The quantity $\chi = \rho \lambda_{\text{dB}}^3$, where ρ is the density of the gas, is known as the *phase-space density*. Physically this denotes the number of particles present in a box with sides of λ_{dB} length. The phase transition to the condensate sets in at $\chi = \zeta(3/2) \approx 2.612$, where $\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}$ is the

Riemann-Zeta function. [11]. The value of T_c for this to occur in a uniform gas is given by

$$k_B T_c = \frac{2\pi}{\zeta(3/2)^{2/3}} \frac{\hbar^2 \rho^{2/3}}{m}. \quad (2.6)$$

It is worth noting that the above derivation assumes an ideal gas, in which there are no interactions between particles. As this is not the case for most physical systems, one must consider systems where interactions are weak. Fritz London, in 1938, following on from the body of work derived by Einstein and Bose drew the connection between superfluidity in liquid ^4He and Bose–Einstein condensation [10, Chap. 1]. However, due to the strongly interacting nature of liquid ^4He at low temperatures only approximately 10% of the atoms condense into a BEC [12]. In order to achieve the large occupation of the lowest energy state the system must be prepared so that the inter-particle interaction strength does not degrade the coherence, as is the case for liquid ^4He .

Due to their weak interactions, dilute atomic gases are closer to the ideal case discussed by Bose and Einstein. One negative impact, though, were the diluteness requirements and higher masses of most atoms compared to ^4He , which required reaching much lower transition temperatures. As such, the use of lighter elements was considered. Spin-polarised hydrogen was one of the first systems to be investigated to create a BEC [10, 13, 14] in the 1970’s. Using trapping and cooling techniques available at the time, this type of system came close, but did not quite reach the required temperatures and phase-space densities for Bose–Einstein condensation to occur until over twenty years later [14]. With the advent of laser cooling in the 1980’s, the use of alkali atoms was considered partly due to the ease of accessibility of their optical transition frequencies. It was not until 1995 that the first BECs were experimentally realised [15, 16].

For a finite number of atoms in a realistic experimental scenario, one can consider the case of a BEC trapped in a harmonic oscillator potential. For this finite-sized sample and inhomogeneous density profile, the transition temperature is given by [10]

$$k_B T_c = \frac{\hbar \bar{\omega} N^{1/3}}{\zeta(3)^{1/3}}, \quad (2.7)$$

where $\bar{\omega} = (\omega_x \omega_y \omega_z)^{1/3}$ is the geometric mean of the harmonic oscillator frequencies, and interactions between the particles have been neglected. Critical temperatures for harmonically trapped dilute gases are on the order of nano-Kelvin for experimentally realistic systems.

2.1.3 Theoretical description of BECs: Gross–Pitaevskii equation

We will, in the following section, outline the derivation of the mean-field Gross–Pitaevskii equation, which is widely used to study the behaviour of condensates in many works cited in this review. Following the *Les Houches 2013 Lecture Course* by J. Walraven [17] the

second quantized form of the many-body Hamiltonian for interacting particles in an external potential is given by

$$\hat{\mathcal{H}} = \hat{H}_1 + \hat{H}_2 = \int \hat{\Psi}^\dagger(\mathbf{r}) H_0(\mathbf{r}, \mathbf{p}) \hat{\Psi}(\mathbf{r}) d\mathbf{r} + \frac{1}{2} \iint \hat{\Psi}^\dagger(\mathbf{r}') \hat{\Psi}^\dagger(\mathbf{r}) V_{\text{int}}(\mathbf{r}', \mathbf{r}) \hat{\Psi}(\mathbf{r}) \hat{\Psi}(\mathbf{r}') d\mathbf{r}' d\mathbf{r}, \quad (2.8)$$

with $H_0(\mathbf{r}, \mathbf{p}) = \mathbf{p}^2/(2m) + V_{\text{ext}}(\mathbf{r}) = -(\hbar/2m)\nabla^2 + V_{\text{ext}}(\mathbf{r})$, and $\mathbf{r} = (x, y, z)$. The external potential, $V_{\text{ext}}(\mathbf{r})$ is taken as harmonic, of the form

$$V_{\text{ext}}(\mathbf{r}) = \frac{m}{2} \sum_i (\omega_i r_i)^2, \quad (2.9)$$

where ω_i represents the trapping frequency in the i -th spatial dimension. The interaction potential, V_{int} is assumed to be point-like as

$$V_{\text{int}}(\mathbf{r}, \mathbf{r}') = g\delta(\mathbf{r} - \mathbf{r}'), \quad (2.10)$$

where δ is the Dirac delta function and the mean-field interaction, g , is given by

$$g = \frac{4\pi\hbar^2 a_s}{m}, \quad (2.11)$$

with a_s being the s -wave scattering length. Inserting the contact potential Eq. (2.10) into the second quantised interaction Hamiltonian \hat{H}_2 from Eq. (2.8) above yields the relations

$$\hat{H}_2 = \frac{g}{2} \int \hat{\Psi}^\dagger(\mathbf{r}) \hat{\Psi}^\dagger(\mathbf{r}') \delta(\mathbf{r} - \mathbf{r}') \hat{\Psi}(\mathbf{r}') \hat{\Psi}(\mathbf{r}) d\mathbf{r} d\mathbf{r}' \quad (2.12a)$$

$$= \frac{g}{2} \int \hat{\Psi}^\dagger(\mathbf{r}) \hat{\Psi}^\dagger(\mathbf{r}) \hat{\Psi}(\mathbf{r}) \hat{\Psi}(\mathbf{r}) d\mathbf{r} \quad (2.12b)$$

$$= \frac{g}{2} \int \hat{\Psi}^\dagger(\mathbf{r}) \hat{n}(\mathbf{r}) \hat{\Psi}(\mathbf{r}) d\mathbf{r}. \quad (2.12c)$$

In the Heisenberg picture, the evolution of the system is governed by the equation

$$i\hbar \frac{d}{dt} \hat{\Psi}_H(\mathbf{r}, t) = \left[\hat{\Psi}_H(\mathbf{r}, t), \hat{\mathcal{H}} \right], \quad (2.13)$$

where the Heisenberg field annihilation operator, $\hat{\Psi}_H(\mathbf{r}, t)$, is given by

$$\hat{\Psi}_H(\mathbf{r}, t) = e^{i\hat{\mathcal{H}}t/\hbar} \hat{\Psi}(\mathbf{r}) e^{-i\hat{\mathcal{H}}t/\hbar}. \quad (2.14)$$

The operator, $\hat{\Psi}_H$, can be interpreted as the one removing an atom from a given state of the system. Therefore, if all N atoms in the system are in the ground state, $|0_N\rangle$, as would be the case in an ideal condensate, the following relationship holds

$$\hat{\Psi}_H(\mathbf{r}, t)|0_N\rangle = e^{\frac{iE_0(N-1)t}{\hbar}} \hat{\Psi}(\mathbf{r}) e^{-\frac{iE_0(N)t}{\hbar}} |0_N\rangle \quad (2.15a)$$

$$= \hat{\Psi}(\mathbf{r}) e^{\frac{i[E_0(N-1) - E_0(N)]t}{\hbar}} |0_N\rangle \quad (2.15b)$$

$$= \hat{\Psi}(\mathbf{r}) e^{-\frac{i\mu t}{\hbar}} |0_N\rangle, \quad (2.15c)$$

where the chemical potential is defined as $\mu = E_0(N) - E_0(N - 1)$. Using the bosonic commutation relations

$$\left[\hat{\Psi}(\mathbf{r}'), \hat{\Psi}^\dagger(\mathbf{r}) \right] = \delta(\mathbf{r}' - \mathbf{r}), \quad (2.16)$$

$$\left[\hat{\Psi}(\mathbf{r}'), \hat{\Psi}(\mathbf{r}) \right] = \left[\hat{\Psi}^\dagger(\mathbf{r}'), \hat{\Psi}^\dagger(\mathbf{r}) \right] = 0, \quad (2.17)$$

and noting the following relations

$$\left[\hat{\Psi}(\mathbf{r}), \hat{H}_1 \right] = \hat{H}_0(\mathbf{r}, \mathbf{p}) \hat{\Psi}(\mathbf{r}), \quad (2.18)$$

$$\left[\hat{\Psi}(\mathbf{r}), \hat{H}_2 \right] = g \hat{n}(\mathbf{r}) \hat{\Psi}(\mathbf{r}), \quad (2.19)$$

$$\left[\hat{\Psi}(\mathbf{r}), \hat{N} \right] = \hat{\Psi}(\mathbf{r}), \quad (2.20)$$

upon substitution of Eq. (2.15c) into Eq. (2.13), it can be rewritten as

$$i\hbar \frac{d}{dt} \left(\hat{\Psi}(\mathbf{r}) e^{-i\mu t/\hbar} \right) = H \hat{\Psi}(\mathbf{r}) e^{-i\mu t/\hbar} \quad (2.21)$$

with

$$H = -\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) + g \hat{n}(\mathbf{r}). \quad (2.22)$$

Due to the macroscopic occupation of the lowest lying single-particle state, $\hat{\Psi}$ can be treated as the sum of a condensed term and a quantum fluctuation (uncondensed) term, as

$$\hat{\Psi} = \Psi + \delta\hat{\Psi}, \quad (2.23)$$

where $\Psi = \langle \hat{\Psi} \rangle$ is known as the condensate wavefunction. Ψ takes the form of a classical field, and can be used to model the behaviour of the condensate, provided the number of atoms is sufficiently large ($N > 10^3$ for most experimental set-ups) and the correlations are not too strong i.e. the gas is sufficiently dilute that only two-body interactions occur. By substituting Eq. (2.23) into (2.8), and assuming a condensate at $T = 0$ K, the number of uncondensed atoms will be essentially zero, and thus we can safely ignore all terms in $\delta\hat{\Psi}$ and $\delta\hat{\Psi}^\dagger$. The above procedure gives the time dependent mean-field Gross-Pitaevskii equation (GPE) as

$$i\hbar \frac{\partial}{\partial t} \Psi(\mathbf{r}, t) = H_{\text{GP}} \Psi(\mathbf{r}, t), \quad (2.24)$$

with the nonlinear GPE Hamiltonian,

$$H_{\text{GP}} = \left[-\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) + g |\Psi(\mathbf{r}, t)|^2 \right], \quad (2.25)$$

where $\Psi(\mathbf{r}, t) = \Psi(\mathbf{r}) e^{-i\mu t/\hbar}$. The time independent form can be found by evaluating the left-hand side derivative in (2.24) and dividing across by $e^{-i\mu t/\hbar}$, yielding

$$\mu \Psi(\mathbf{r}) = \left[-\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) + g |\Psi(\mathbf{r})|^2 \right] \Psi(\mathbf{r}), \quad (2.26)$$

where the wavefunction is normalised to the particle number, N , as follows

$$\int_{-\infty}^{\infty} d\mathbf{r} |\Psi(\mathbf{r}, t)|^2 = N. \quad (2.27)$$

Assuming a large number of bosons in the condensate $N \gg 1$, the interaction term dominates over the kinetic term of the Hamiltonian which can therefore be neglected. This approximation turns finding the ground state of the system into a solvable, algebraic problem, and is known as the Thomas–Fermi approximation [11, p. 84]. The Hamiltonian can thus be reduced to a combination of the trapping potential and the mean-field interaction, where the ground state wavefunction, Ψ_{TF} , can be determined from the time independent GPE as

$$\Psi_{\text{TF}}(\mathbf{r}) = \sqrt{g^{-1}[\mu - V(\mathbf{r})]\Theta(\mu - V(\mathbf{r}))}, \quad (2.28)$$

where μ is the chemical potential, and Θ is the Heaviside step function, which ensures that the condensate density does not become negative. The boundary of the cloud is determined by the surface at which the density becomes zero, and corresponds to the point where the trapping potential and chemical potential are equal. This gives the Thomas–Fermi radius along the i -th direction of the cloud [18, p. 165] as

$$R_i = a \left(\frac{15Na_s}{a} \right)^{1/5} \frac{\bar{\omega}}{\omega_i}, \quad (2.29)$$

where the characteristic length of the harmonic oscillator is given by $a = \sqrt{\hbar/m\bar{\omega}}$, and

$$\bar{\omega} = \left(\prod_j \omega_j \right)^{1/3},$$

is the geometric mean of the harmonic oscillator trapping frequencies. This approximation holds valid for stationary condensate solutions.

To investigate condensate dynamics, it is convenient to consider rotation of the condensate about an axis. For such a rotating condensate, an additional term appears in the GPE Hamiltonian, $-\mathbf{\Omega} \cdot \mathbf{L}$, where $\mathbf{\Omega}$ is the angular rotation frequency, and \mathbf{L} is the angular momentum operator. Assuming rotation about a single axis, the longitudinal direction z , \mathbf{L} can be replaced with L_z , giving the form of the GPE in the co-rotating frame as

$$i\hbar \frac{\partial}{\partial t} \Psi(\mathbf{r}, t) = \left[-\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) + g|\Psi(\mathbf{r}, t)|^2 - \Omega L_z \right] \Psi(\mathbf{r}, t). \quad (2.30)$$

Within the Thomas-Fermi approximation, it is possible to determine analytical results for the profile of the BEC, and compare with exact results from numerically integrating the full GPE. However, in the case of a rotating condensate to solve the above form of the Gross–Pitaevskii equation will require numerical integration as the kinetic energy term can no longer be neglected.

2.1.4 Bogoliubov-de Gennes equations

While the Gross–Pitaevskii equation captures the rich array of dynamics exhibited by a condensate system, it is often necessary to examine the stability of its solutions. The previously neglected small quantum fluctuations (see Eq. (2.23)) can be included by writing them as a combination of counterpropagating waves as

$$\delta\hat{\Psi} = e^{-i\mu t/\hbar}[u(\mathbf{r})e^{-it\omega} + v^*(\mathbf{r})e^{it\omega}]. \quad (2.31)$$

From this expression, the wavefunction (2.23) can be written as

$$\Psi(\mathbf{r}, t) = e^{-i\mu t/\hbar} [\Psi_0(\mathbf{r}) + u(\mathbf{r})e^{-it\omega} + v^*(\mathbf{r})e^{it\omega}], \quad (2.32)$$

where $\Psi_0(\mathbf{r})$ is the stationary state solution. Firstly, we calculate the time-derivative of Eq. (2.32), which after simplification becomes

$$i\hbar\partial_t\Psi(t) = e^{-\frac{i\mu t}{\hbar}} [\mu\Psi_0 + (\mu + \hbar\omega)ue^{-it\omega} + (\mu - \hbar\omega)v^*e^{it\omega}], \quad (2.33)$$

where the dependence on \mathbf{r} is dropped for notational simplicity. Next, the nonlinear interaction term is given as

$$g|\Psi|^2\Psi = g\Psi^*\Psi\Psi \quad (2.34a)$$

$$= ge^{-\frac{i\mu t}{\hbar}} (\Psi_0^* + u^*e^{it\omega} + ve^{-it\omega}) \times (\Psi_0 + ue^{-it\omega} + v^*e^{it\omega})^2, \quad (2.34b)$$

$$\approx ge^{-\frac{i\mu t}{\hbar}} [|\Psi_0|^2(\Psi_0 + 2(ue^{-it\omega} + v^*e^{it\omega})) + \Psi_0^2(u^*e^{it\omega} + ve^{-it\omega})]. \quad (2.34c)$$

with the resulting equations linearised in terms of u and v . Plugging these terms into the GPE yields the Bogoliubov-de Gennes (BdG) equations,

$$\mu\Psi_0 = (H_0 - \Omega L_z + g|\Psi_0|^2)\Psi_0, \quad (2.35a)$$

$$(\mu + \hbar\omega)u = (H_0 - \Omega L_z + 2g|\Psi_0|^2)u + g\Psi_0^2v, \quad (2.35b)$$

$$(\mu - \hbar\omega)v^* = (H_0 - \Omega L_z + 2g|\Psi_0|^2)v^* + g\Psi_0^2u^*, \quad (2.35c)$$

where

$$H_0 = -\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{r}). \quad (2.36)$$

This can be written in matrix form as

$$\begin{pmatrix} H_0 + 2g|\Psi_0|^2 - \mu - \Omega L_z & g\Psi_0^2 \\ -g\Psi_0^{*2} & -H_0 - 2g|\Psi_0|^2 + \mu + \Omega L_z \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} = \hbar\omega \begin{pmatrix} u \\ v \end{pmatrix}, \quad (2.37)$$

where the eigenvalues of these modes can be used to determine the stability of the system. The norm of these functions is given by $N_{\text{BdG}} = \int d\mathbf{r}(|u|^2 - |v|^2)$. If the norm is positive, with

positive eigenvalues, the system is energetically stable. If the norm is positive with negative eigenvalues, the system will have an energetic instability, which will cause the system to move towards the lowest energy state only in the presence of dissipation. If, however, the norm is 0 with imaginary eigenvalues, the modes of the fluctuations are dynamically unstable. Due to the complex eigenvalues these modes will dominate exponentially over time and destroy the initial state of the condensate.

2.1.5 Lower dimensional condensates

Whilst the full three dimensional GPE describes the dynamics of a condensed cloud of cold atoms, often the physics of lower dimensional systems can also be quite interesting. For a BEC tightly confined along one axis, the system can be described as a pancake shaped condensate, or cigar shaped for tight confinement along two axes. These tightly confined systems allow for the examination of both two and one dimensional physics. For a BEC harmonically confined in a trap with a transverse frequency, ω_{\perp} , and tightly confined along z with frequency $\omega_z \gg \omega_{\perp}$, all dynamics can be frozen out along z , leaving the system in the ground state of the oscillator along z . This assumes that the energy to excite the system along z , $\hbar\omega_z$ is significantly greater than that to excite along the transverse plane $\hbar\omega_{\perp}$, and also the chemical potential μ . This assumption allows for the wavefunction to be written as $\Psi(\mathbf{r}, t) = \Psi(x, y, t)\phi(z)$ where $\Psi(x, y, t)$ is the transverse wavefunction, and $\phi(z) = (m\omega_z/(\pi\hbar)) \exp(-z^2 m\omega_z/(2\hbar))$ is the ground state along z . Substituting this separable wavefunction into the GPE and integrating over z modifies the nonlinear interaction strength as [18]

$$g_{2D} = g \sqrt{\frac{m\omega_z}{2\pi\hbar}}. \quad (2.38)$$

Though still technically a 3D system, this separation and integration allows one to consider a quasi-2D model, which can be modelled in two-dimensions and which reproduces the main aspects of the same vortex dynamics as a tightly confined three dimensional model. This is advantageous for numerical simulations, but also ensures that only the physics of interest, i.e. in this case the vortex-vortex interactions, play a role. Three dimensional effects, such as the possibility of excitations along the vortex line are removed. The above assumptions have been observed in many experimental systems [19–24], and therefore all work and models discussed later will be within the quasi-2D condensate regime, unless otherwise specified.

2.2 Superfluidity

2.2.1 Introduction to superfluidity

Superfluidity is a macroscopic quantum effect that is closely related to Bose–Einstein condensation. Liquid helium has been known for many years to exhibit superfluid behaviour [12].

One interesting property of superfluids is that they have quantized circulation which can lead to the appearance of quantum vortices under rotation. Traditionally, such excitations have been created in liquid ^4He by a “rotating bucket” type of experiment, where the container holding the superfluid is rotated about a single axis. When the fluid is initially above the λ -critical point, the temperature at which ^4He moves from a classical fluid to a superfluid, the atoms will undergo classical rotation with the container. On cooling the atoms below the λ -critical point, liquid ^4He will undergo a transition to the superfluid state. If the velocity is above a critical rotation frequency, Ω_c , vortices are nucleated in the rotating superfluid helium system. Due to the strongly interacting nature of liquid helium, the nucleated vortices are difficult to visualise as the healing length, ξ , of liquid ^4He is only on the order of Ångströms [25].

In order to visualise these vortices experimentally it was necessary to use an indirect means of visualisation in the form of tracer particles [26]. Limited success was had with this technique using solid hydrogen, and later plastic microspheres, as they tended to join together due to static charges. An improved technique, using charged particles, showed much greater success [27]. Ions, or electrons, were trapped inside the vortex lines, and an electric field along the direction of the lines allowed for acceleration of the charges towards a luminescent screen where they could be observed. Current experimental work in vortex visualisation with liquid ^4He has advanced significantly [28, 29], yet fine control over the behaviour of the liquid and the vortex dynamics remains difficult.

Superfluid behaviour is observed in liquid helium due partly to a portion of the atoms condensing into the ground state. Given that BECs show superfluidity, the use of a dilute gas where the majority of atoms are in the condensate state provides a more controlled means to investigate superfluidity and quantised vortices [11, 25, 28, 30]. In contrast with liquid ^4He , which has a healing length of the order of Ångströms, the healing length of a dilute gas of alkali atoms is on the order of microns [31]. This places condensates in a much more accessible regime for visualising vortices compared with liquid ^4He experiments. For many condensates visualisation of the vortices is provided by absorption imaging, following a time-of-flight expansion of the cloud to allow the vortex cores to expand and become more easily visible [32, 33]. One drawback of this method is that it is destructive, and requires multiple experimental realisations to acquire any dynamical information. To examine the dynamics of the vortices, successive imaging techniques are required. Single-shot experiments have been demonstrated, where a small percentage of the condensate is repeatedly transferred to an untrapped atomic state. This allows for the untrapped atoms to expand, and vortices have been imaged this way [34]. With this method the precession of the vortices in a trapped condensate can be directly observed, in a single-shot experiment. More recently, *in-situ* imaging of a condensate with a vortex lattice was demonstrated and allowed resolution of the cores without time-of-flight expansion with a minimally destructive optical refraction technique [35].

To fully understand the behaviour of these systems, it is necessary to have a framework for

modeling the condensate, in the absence and the presence of vortices. As described above, a dilute gas of condensed atoms close to absolute zero temperature can be readily modelled by the Gross–Pitaevskii equation (2.30), offering a direct means to examine condensate behaviour. It is, however, useful to obtain a hydrodynamic description for the condensate, performed by treating its wave-function as

$$\Psi(\mathbf{r}, t) = \sqrt{\rho(\mathbf{r}, t)}e^{i\theta(\mathbf{r}, t)}, \quad (2.39)$$

with

$$\rho(\mathbf{r}, t) = \Psi^*(\mathbf{r}, t)\Psi(\mathbf{r}, t) = |\Psi(\mathbf{r}, t)|^2. \quad (2.40)$$

where $\theta(\mathbf{r}, t)$ is the condensate phase [10]. Multiplying Eq. (2.24) by $\Psi^*(\mathbf{r}, t)$ then subtracting the complex conjugate of that expression allows one to obtain the continuity equation as [10, 18]

$$\frac{\partial}{\partial t}\rho(\mathbf{r}, t) + \nabla \cdot \mathbf{j}(\mathbf{r}, t) = 0, \quad (2.41)$$

where the current density of the condensate is

$$\mathbf{j}(\mathbf{r}, t) = \frac{-i\hbar}{2m} [\Psi^*(\mathbf{r}, t)\nabla\Psi(\mathbf{r}, t) - \Psi(\mathbf{r}, t)\nabla\Psi^*(\mathbf{r}, t)]. \quad (2.42)$$

Using ansatz (2.39) and substituting it into Eq. (2.42) then gives the form for the current density,

$$\mathbf{j}(\mathbf{r}, t) = |\Psi(\mathbf{r}, t)|^2 \frac{\hbar}{m} \nabla\theta(\mathbf{r}, t). \quad (2.43)$$

The velocity of the superfluid, $\mathbf{v}(\mathbf{r}, t)$, is defined as the ratio of the current density to the density, which is then given by

$$\mathbf{v}(\mathbf{r}, t) \equiv \frac{\mathbf{j}(\mathbf{r}, t)}{\rho(\mathbf{r}, t)} = \frac{\hbar}{m} \nabla\theta(\mathbf{r}, t). \quad (2.44)$$

The gradient of the phase therefore determines the velocity of the condensate atoms; this indicates that the superfluid behaviour in a condensate is irrotational ($\nabla \times (\nabla\theta) = 0$). Assuming a closed loop integral about a central point in the condensate, and recalling the single-valued nature of the wavefunction, yields the relationship

$$\oint \mathbf{v} \cdot d\mathbf{l} = \frac{\hbar}{m} 2\pi l. \quad (2.45)$$

This shows the quantised nature of circulation in a superfluid, with l representing the integer charge of the circulation. The phase winding around the central region is given by multiples of 2π , with the centre of the phase becoming ill-defined. To circumvent this problem the density at this point drops to zero, signalling the presence of a vortex in the condensate. This drop happens over the scale of the healing length, which, for repulsive interactions, is given by

$$\xi = \frac{1}{\sqrt{8\pi\rho_b a_s}}, \quad (2.46)$$

where ρ_b is the bulk density of the condensate, and a_s is the s -wave scattering length. For a cylindrically symmetric rotating condensate carrying a single vortex in its centre the wavefunction can be written as

$$\Psi = |\Psi|e^{il\theta}, \quad (2.47)$$

where l is the vortex charge. The velocity profile at a distance r_v away from the core can then be given as

$$v = \frac{\hbar l}{mr_v}. \quad (2.48)$$

To sustain a vortex the condensate must have sufficient angular momentum, which is imparted via the angular rotation frequency times the angular momentum operator $-\Omega L_z$ in the Hamiltonian given by Eq. (2.30). The rotation frequency, Ω , of the condensate has an upper-bound stability limit equivalent to the transverse trapping frequency, ω_\perp , of the harmonically trapped condensate.

2.2.2 Vortices in Bose–Einstein condensates

Given that the circulation of a vortex in a superfluid is quantised, one may assume that for faster rotation rates the circulation increases in integer multiples of h/m , beyond critical rotation thresholds to excite each higher multiple. To understand the effect of vortex charge on the condensate, it is necessary to examine the energy functional, as given by

$$E[\Psi] = \int \Psi^*(H_E)\Psi d\mathbf{r}, \quad (2.49)$$

where $H_E = (-\hbar^2/2m)\nabla^2 + V + g|\Psi|^2/2 - \Omega L_z$. By substituting (2.47) into Eq. (2.49) the energy functional is given as

$$E[\Psi] = \int \frac{\hbar^2}{2m} \left(|\nabla\Psi|^2 + \frac{|\Psi|^2 l^2 m \mathbf{v}^2}{2} \right) + V|\Psi|^2 + \frac{g}{2}|\Psi|^4 + \Omega\Psi^* L_z \Psi, \quad (2.50)$$

where \mathbf{v} is the superfluid velocity, as given by Eq. (2.44). The l^2 term shows that the energy scales quadratically with an increase in charge. This energy growth dependence indicates that it will be more favourable to allow for two singly charged vortices, than a single doubly charged vortex, which will decay due to the presence of complex eigenmodes [36]. For rates of rotation $\Omega_c < \Omega < \omega_\perp$, where Ω_c , is the threshold rotation rate to create a vortex, the condensate will favour many singly quantised vortices, rather than one or several multiply charged ones. For a superfluid condensate in the rotating frame $\Omega_c = E_v/L = E_v/(N\hbar)$, where E_v is the energy of the vortex, and L is the angular momentum component of the superfluid along z [10].

The stability of vortex states in a condensate is also a widely discussed topic [37, 38], as non-rotating traps show an instability for small displacements of the vortex from the trap centre. Since the ‘‘rotating bucket’’ technique used to generate vortices in ^4He cannot be used for

gaseous BECs, many other techniques have been developed [39–43]. For the work presented in this thesis, the method proposed by Dobrek *et al.* [42], is of particular interest, as it allows one to optically generate vortices by the use of what they term a “phase-imprinting” method. The authors describe a scheme where the phase of the condensate is directly controlled in such a way that the required topological charge to induce a vortex during evolution is provided by external lasers. Through use of an absorption plate whose absorption coefficient depends on the axis angle, the condensate can be imprinted with the required phase pattern. This method will form the basis of one work carried out in this thesis, and will be discussed in detail later in Sec. 4.3.2.

2.2.3 Vortex lattices

Although a large number of works exist which investigate systems with low numbers of vortices [44–54], such systems do not necessarily form periodic vortex lattices. We will therefore concentrate primarily on studies of systems containing many vortices, assuming large values of $\Omega \lesssim \omega_{\perp}$. For such systems, the vortices are arranged in a periodic triangular Abrikosov lattice, reminiscent of that which appears in type-II superconductors with magnetic flux lines [55]. This triangular lattice configuration is the most energetically favourable and stable arrangement [56]. Experimental setups have demonstrated upwards of over 130 vortices in a well ordered triangular formation. The resulting lattices are shown to be highly stable and are ideal setups for investigations of periodic systems [55, 57].

The Thomas–Fermi limit discussed earlier describes the case where the kinetic energy term of the Hamiltonian may be neglected in comparison to the interaction energy, as it offers little contribution to the condensate behaviour. This remains true for low rotation rates, however kinetic energy becomes important in the limit of fast rotation. In this case $\Omega/\omega_{\perp} \approx 1$, and the centrifugal force term, $m\Omega^2 r$, almost balances with the trapping force term, $-m\omega^2 r$. The condensate behaviour then closely resembles that of the two-dimensional quantum Hall regime, where the system is residing in the lowest Landau level, (LLL) ($n = 0$) [58]. As the rotation of the cloud approaches the trapping frequency, the system tends to the LLL, wherein the nonlinear interaction term becomes relatively weak due to the centrifugal forces on the atomic cloud. In this regime the Thomas–Fermi approximation becomes invalid, as the interaction term no longer dominates over the kinetic energy. The interaction energy is then much smaller than the gap between energy levels for single particle Landau levels. While there have been studies using harmonic-plus-quartic potentials to allow condensates to remain trapped beyond the harmonic trapping frequency limit [59, 60], we will in this thesis consider only systems in a harmonic confinement.

At high rotation rates very close to the trap frequency ($\Omega/\omega_{\perp} \approx 1$) the system enters a fractional quantum Hall (FQH) state [61], and will require treatment beyond that of mean-field methods. This state will require the use of a discrete boson model, which becomes next to impossible to simulate for a realistic number of atoms in a condensate. However,

for rotation rates just below this limit a mean-field approach can be used in the so-called “mean-field quantum Hall” (MFQH) regime [62]. The differences between these two regimes are characterised purely by the ratio of the kinetic energy to the interaction energy of the system, and in the MFQH regime we can assume that the LLL has been achieved [63, 64]. To identify these different regimes one may use the “filling factor”, $\nu = N/N_v$ [11, 58], which examines the ratio of atoms, N , to vortices in the system, N_v . This becomes an important characteristic at high rotation rates, and in the case where $1000 > \nu > 10$, the system may be accurately described to be in the MFQH regime. For values $\nu \leq 10$ the system is said to be strongly correlated [62], and the system enters the FQH regime. For an almost perfectly regular vortex lattice, the rotation rate must be sufficiently large so that the condensate width extends to large distances and a large number of vortices are generated, without entering the FQH regime [65]. To ensure the applicability of mean-field theory choosing rotational frequencies that guarantee this is important.

The distribution of the vortices in the MFQH regime forms a triangular lattice pattern that is almost regular [66]. As discussed earlier, the condensate has an irrotational flow profile due to velocity being defined as Eq. (2.44). This relation holds true provided that θ is well defined, which is not the case at the centre of a vortex. As the phase at the vortex core is ill-defined, this singular region creates a non-zero curl. A generalisation of the irrotational flow condition (i.e. $\nabla \times \mathbf{v} = 0$) to account for this can be given as as [10, 18]

$$\nabla \times \mathbf{v} = \frac{2\pi l\hbar}{m} \delta^{(2)}(\mathbf{r}_\perp) \hat{z}, \quad (2.51)$$

where δ^2 is a two-dimensional Dirac delta function, \hat{z} is the unit vector along z and $\mathbf{r}_\perp = (x, y)$. For large rotation frequencies this value becomes very similar to that of a solid-body rotation, as given by $\nabla \times \mathbf{v} = 2\Omega_z$. This result arises for a large regular vortex lattice, where the areal density of vortices can be specified by the Feynman relation [10]

$$n_v = \frac{m\Omega}{\pi\hbar}. \quad (2.52)$$

In the case of realistic condensates systems, where the densities are not uniform due to harmonic trapping, deviations exist from this value which have been calculated theoretically [67–70] and observed experimentally [71]. However, for the values chosen in all later discussed simulations these deviations tend to be small, and can in most cases be neglected.

Some key details for the classification of the regime of the condensate are [62]:

- Mean-field theory and experiment agree well for large atom numbers ($N > 10^3$) and rotational frequencies reaching approximately $0.995\omega_\perp$.
- The density profile of slowly rotating condensates differs very little from that of non-rotating condensates, except for the inclusion of the vortex cores.
- The Thomas–Fermi regime holds true for frequencies of $0.75 \leq \Omega/\omega_\perp \leq 0.99$, where kinetic energy remains negligible compared to flow velocity.

- The MFQH regime description, holding for $0.99 \leq \Omega/\omega_{\perp} \leq 0.999$, sees the vortex cores expanding and forming a densely packed lattice.

Given the above criteria, the rapidly rotating condensate system in our work can be described using mean-field Gross–Pitaevskii theory at a rotation rate of $\Omega = 0.995\omega_{\perp}$, assuming $N \approx 10^6$ atoms, as used in typical experimental setups.

2.3 Recent progress and outlook

Recently there have been many advances in the control of atomic BEC systems. A notable example is the work of Gauthier *et al.* [72], wherein they demonstrate arbitrary optical potential generation for Bose–Einstein condensates through use of digital micromirror devices (DMD). The authors demonstrate high resolution control and patterning of the condensate, and show near perfect control of the condensate atomic distribution. Given the current state of the art high performance imaging and control techniques available, these experimental systems can allow for high precision control and manipulations of the atoms, and therefore also vortices.

Another recent experimental work of note is that of Wigley *et al.* [73], with a completely automated approach to BEC generation and control. Such automation can allow for much higher throughput of experimental data collection, and allow for a much wider breadth of physics to be explored in condensate systems.

The current state of the art experimental systems can offer a very high degree of control of condensates, and their ensuing dynamics. In fact, all further methods discussed can be built using currently available state of the art systems, and hence are experimentally realisable.

Chapter 3

Numerical methods

Numerical solutions to problems in quantum physics are important, given the limited availability of exact solutions. Many such models and methods exist when dealing with many-body systems, and have been shown to provide good estimates of physical behaviour [74]. Techniques such as Monte Carlo methods, exact diagonalisation, and DMRG are used to solve a wide variety of many-body problems, but often fail to capture the full system dynamics of such problems [75]. For understanding the behaviour of systems such as Bose–Einstein condensates, the use of these methods is rather limited. For DMRG, the complexity of the problem grows quickly with increasing dimensionality and renders this technique unusable. Exact diagonalisation requires a linearised system to obtain realistic solutions, and also grows significantly in complexity with increased dimensionality. Monte Carlo methods generally do not allow for real-time dynamics, or allow one to calculate the underlying wavefunction.

To obtain solutions to BEC problems we make use of a mean-field approach, outlined previously in Sec. 2.1.3. Using the GPE and performing a numerical integration allows for almost all examinable dynamics that are valid in the mean-field limit. It can be noted though that the computational cost increases significantly with increased dimensions, and is already non-trivial for two-dimensions. The following chapter will introduce the necessary requirements to numerically solve quantum problems using state-of-the-art computational methods.

We will begin with an introduction to the time evolution of a quantum state. We then discuss the use of the time evolution approach to find the ground state of a quantum system using imaginary time evolution. After this necessary mathematical introduction we will discuss the implementation of both real and imaginary time evolution using the Fourier split-operator (split-step) algorithm. We give error bounds for the algorithm, and discuss its use in the context of solving for Hamiltonian dynamics. Though the discussed algorithm is well suited to solving quantum dynamics, the computational cost can be quite high, especially for systems with large grid sizes.

We next discuss ways to overcome this through the use of high performance computing meth-

ods, and introduce the concept of graphical processing unit (GPU) computing. We present many of the necessary considerations for mapping a computational problem onto GPUs. Making use of GPUs to numerically solve the Schrödinger equation with the Fourier split-operator method, we present the problem of coherent atomic transport. We introduce the “matter-wave spatial adiabatic passage” technique, with the goal of coherently transporting an atom between trapping potentials with high fidelity. The design of the model system is presented, and the results are shown.

We finish the chapter by introducing the developed algorithms for condensate systems. Performance metrics and considerations are given in the context of solving the Gross–Pitaevskii equation in the presence of vortices.

3.1 Time evolution

Given a quantum state, to examine the dynamics requires an understanding of how it evolves in time. Assuming a quantum state at time t_0 to be defined as $|\Psi(t_0)\rangle$, and a state at time t to be $|\Psi(t)\rangle$, the two states can be connected with a unitary evolution operator as

$$|\Psi(t)\rangle = \mathcal{U}(t, t_0)|\Psi(t_0)\rangle, \quad (3.1)$$

where we have assumed that $t > t_0$. One can write the wavefunction of a quantum system as the linear superposition of a set of basis states $|\Psi_m\rangle$ as

$$|\Psi\rangle = \sum_m C_m |\Psi_m\rangle, \quad (3.2)$$

with coefficients C_m . The unitary evolution operator can be written as

$$\mathcal{U}(t, t_0) = \exp\left(-\frac{i\mathcal{H}(t-t_0)}{\hbar}\right) \quad (3.3)$$

$$= \exp\left(\frac{-i\mathcal{H}\delta t}{\hbar}\right), \quad (3.4)$$

where \mathcal{H} is the Hamiltonian of the system, and $\delta t = t - t_0$.

The effect of $\mathcal{U}(t, t_0)$ on any arbitrary state can be examined by expanding it as

$$\mathcal{U}(t, t_0) = 1 - \frac{i\mathcal{H}\delta t}{\hbar} + \mathcal{O}(\delta t^2) \quad (3.5a)$$

$$= 1 - \frac{i\delta t}{\hbar} \sum_n E_n |\Psi_n\rangle \langle \Psi_n| + \mathcal{O}(\delta t^2), \quad (3.5b)$$

where the Hamiltonian operator has been written as a complete set of energy eigenkets, $\mathcal{H}|\Psi_n\rangle = E_n|\Psi_n\rangle$. Applying this to (3.2) gives

$$|\Psi(t)\rangle = \left(1 - \frac{i\delta t}{\hbar} \sum_n E_n |\Psi_n\rangle \langle \Psi_n|\right) \sum_m C_m |\Psi_m\rangle \quad (3.6a)$$

$$= \sum_m C_m |\Psi_m\rangle + \sum_n \left(-\frac{i\delta t}{\hbar} E_n\right) C_n |\Psi_n\rangle, \quad (3.6b)$$

where terms of order $\mathcal{O}(\delta t^2)$ and higher have been temporarily left out for notational simplicity, but are still included in the analysis. Time evolving the state from t_0 to a final time t by applying the evolution operator can then be written as

$$\mathcal{U}(t, t_0)|\Psi(t_0)\rangle = \sum_n C_n \exp\left(\frac{-iE_n \delta t}{\hbar}\right) |\Psi_n\rangle. \quad (3.7)$$

It follows from here that each state oscillates at a different rate, proportional to its eigenenergy; higher energy states will oscillate faster than those of lower energy. For a given set of states the dynamics and evolution of the quantum system can be fully determined for all times using the above evolution operation.

As systems will prefer to reside in the lowest energy state where they are most stable, it is often required to determine the ground state solution of a particular Hamiltonian. A common method for this is by evolving the system in imaginary time. Taking the evolution operator, and applying a Wick rotation [76] rotates the time component through $\pi/2$ into the imaginary plane, as $t \rightarrow -it$. This new evolution operator applied to the wavefunction gives

$$\mathcal{U}'(t, t_0)|\Psi\rangle = \sum_n C_n \exp\left(\frac{-E_n \delta t}{\hbar}\right) |\Psi_n\rangle. \quad (3.8)$$

This process removes the complex term in the operator, which now takes the form of sums of exponentially decaying states. When applied to the wavefunction the higher energy terms will decay at a rate faster than lower energy components increasing δt . This process also causes a loss of probability density, and so the wavefunction must be renormalised after each application. Through repeated application of this operator, and a renormalisation afterwards, the simulated quantum system converges to the ground state solution. To begin, however, we must make an initial guess for the wavefunction, which has some finite overlap with the lowest lying state. It should be noted that this method is a mathematical trick used to obtain a simulated ground state, with a real-world system only tending to the ground state in the presence of some form of dissipation. As effective as this technique is, the convergence to the lowest lying energy state becomes less effective as the computation approaches the expected value [77], and if many eigenstates are lying close to each other. To ensure the system converges to a sufficient degree the resulting energy can be checked after each iteration, and the evolution stopped only when the energy change fluctuates about a stable value.

With the time evolution method introduced, we will next discuss implementing this method. Although many such algorithms exist to implement time evolution, one that is well suited for this task is the Fourier split-operator method. This method works equally for real, as well as imaginary, time evolution.

3.2 Fourier split-operator method

The Gross–Pitaevskii equation is a second order nonlinear partial differential equation, and so very few exact solutions exist; the problem must often be tackled by a numerical approach.

Though there are many ways to solve such a system numerically, with the Crank–Nicolson and Trotter–Suzuki algorithms being notable examples, the method we have chosen is the pseudospectral Fourier split-operator method, described below [78].

If we consider a unitary evolution operator of the form

$$\Psi(\mathbf{x}, t + \tau) = \exp\left(-\frac{i\hat{H}\tau}{\hbar}\right)\Psi(\mathbf{x}, t), \quad (3.9)$$

where \hat{H} is the Hamiltonian, composed of momentum, potential, nonlinear interaction, and rotation terms defined in Eq. (2.24), we can solve for the wavefunction and its resulting dynamics over a specified timescale, assuming τ is a short time increment such that the formalism given in Sec. 3.1 is valid. Care must be taken during the implementation of such integration methods, as the loss of precision due to floating-point rounding, as well as the propagation of errors cannot be neglected. If we take \hat{H} in terms of its components as a combination of position and momentum space operators we obtain

$$\hat{H} = \hat{H}_{\mathbf{r}} + \hat{H}_{\mathbf{k}} + \hat{H}_{\mathbf{L}}, \quad (3.10)$$

where we first ignore the angular momentum operator, $\hat{H}_{\mathbf{L}}$, and consider only the two other non-commuting parts, $\hat{H}_{\mathbf{r}}$, containing the operators acting in position space, and $\hat{H}_{\mathbf{k}}$, containing the operators acting in momentum space only. The Baker–Campbell–Hausdorff formula [79] gives the relation for non-commuting operators as

$$\exp(\tau(A + B)) = \exp(\tau A) \exp(\tau B) \exp\left(-\frac{\tau^2}{2}[A, B] + \dots\right), \quad (3.11)$$

with \dots representing higher order commutators. This is directly mappable to the above Hamiltonian for time evolution. Due to the non-commutativity of $\hat{H}_{\mathbf{r}}$ and $\hat{H}_{\mathbf{k}}$, the above expression cannot be evaluated exactly, and so it is common to Taylor expand and truncate it. The resulting error can be determined as

$$\text{err} = \left\| \exp\left(-\frac{i\hat{H}_{\mathbf{k}}\tau}{\hbar}\right) \exp\left(-\frac{i\hat{H}_{\mathbf{r}}\tau}{\hbar}\right) - \exp\left(-\frac{i(\hat{H}_{\mathbf{k}} + \hat{H}_{\mathbf{r}})\tau}{\hbar}\right) \right\| \quad (3.12a)$$

$$= \left\| \left(1 + \left(\frac{-i\hat{H}_{\mathbf{k}}}{\hbar}\right)\tau + \left(\frac{-i\hat{H}_{\mathbf{k}}}{\hbar}\right)^2 \frac{\tau^2}{2}\right) \left(1 + \left(\frac{-i\hat{H}_{\mathbf{r}}}{\hbar}\right)\tau + \left(\frac{-i\hat{H}_{\mathbf{r}}}{\hbar}\right)^2 \frac{\tau^2}{2}\right) - \left(1 + \left(\frac{-i(\hat{H}_{\mathbf{r}} + \hat{H}_{\mathbf{k}})}{\hbar}\right)\tau + \left(\frac{-i(\hat{H}_{\mathbf{r}} + \hat{H}_{\mathbf{k}})}{\hbar}\right)^2 \frac{\tau^2}{2}\right) + \mathcal{O}(\tau^3) \right\|, \quad (3.12b)$$

which, upon simplification reduces to

$$\text{err} = \left\| \frac{\tau^2[\hat{H}_{\mathbf{r}}, \hat{H}_{\mathbf{k}}]}{2\hbar^2} + \mathcal{O}(\tau^3) \right\| = \mathcal{O}(\tau^2). \quad (3.13)$$

The error can be further reduced through the use of 2nd order Strang splitting [80], taking the error in the numerical integration scheme to $\mathcal{O}(\tau^3)$, with the resulting operator implementation given as

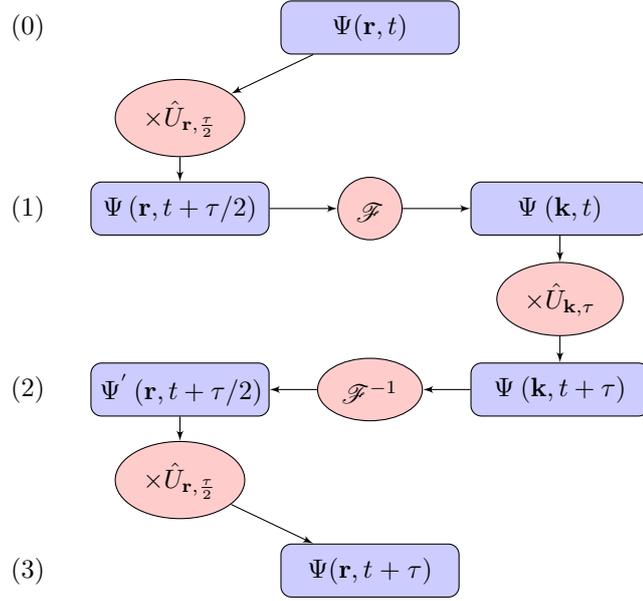


Figure 3.1: A single pass through the Fourier split-operator method.

$$\exp\left(-\frac{i(\hat{H}_{\mathbf{r}} + \hat{H}_{\mathbf{k}})\tau}{\hbar}\right) = \exp\left(-\frac{i\hat{H}_{\mathbf{r}}\tau}{2\hbar}\right) \exp\left(-\frac{i\hat{H}_{\mathbf{k}}\tau}{\hbar}\right) \exp\left(-\frac{i\hat{H}_{\mathbf{r}}\tau}{2\hbar}\right) + \mathcal{O}(\tau^3). \quad (3.14)$$

In the case of a nonlinear system, such as for solving the GPE, the above scheme attains a second-order error, resulting from the combination of the potential and nonlinear terms, with the respective mapping as [81]

$$\hat{H}_{\mathbf{r}} = V(\mathbf{r}) + g|\Psi(\mathbf{r}, t)|^2, \quad (3.15a)$$

$$\hat{H}_{\mathbf{k}} = \frac{-\hbar^2}{2m} \nabla^2. \quad (3.15b)$$

Following Bauke *et al.* [78], we can numerically solve this differential equation as

$$\Psi(\mathbf{r}, t + \tau) = \left[\hat{U}_{\mathbf{r}}\left(\frac{\tau}{2}\right) \mathcal{F}^{-1} \left[\hat{U}_{\mathbf{k}}(\tau) \mathcal{F} \left[\hat{U}_{\mathbf{r}}\left(\frac{\tau}{2}\right) \Psi(\mathbf{r}, t) \right] \right] \right] + \mathcal{O}(\tau^2), \quad (3.16)$$

where $\hat{U}_{\mathbf{r}}(\tau) = e^{-i\hat{H}_{\mathbf{r}}\tau/\hbar}$ is the time evolution operator in position space, $\hat{U}_{\mathbf{k}}(\tau) = e^{-i\hat{H}_{\mathbf{k}}\tau/\hbar}$ the time evolution operator in momentum space, and \mathcal{F} and \mathcal{F}^{-1} are the forward and inverse Fourier transform respectively. Taking the Fourier transform of the wavefunction allows the basis to be transformed between position and reciprocal space, wherein the time evolution operators are diagonal in each respective space. Figure 3.1 outlines a schematic representation of the method during a single pass of the algorithm.

The underlying theory of the Fourier split-operator method for the Gross–Pitaevskii equation is given by Javanainen *et al.* [81], showing how the choice of nonlinearity and operator splitting affects the outcome of the method. By taking the initial step as evolution in momentum

space, the choice of the most current wavefunction attains an error of third-order for the algorithm. However, this will require an additional two Fourier transform steps, and as such is rather costly in compute time for large systems. For an initial step in position space, the non-linear term is best calculated using a linear combination of all available wavefunctions through the algorithm as $\Psi = c_0\Psi_0 + c_1\Psi_1 + c_2\Psi_2$, where the subscripts denote the wavefunction at each stage of the evolution in position space, as indicated in Fig. 3.1, and the c_x are linear coefficients. This gives third order accuracy for the parameters $c_2 = \pm 1, c_1 = -c_0$. However, for simplicity and resource limitations we chose to work with the $\mathcal{O}(\tau^2)$ accurate scheme as depicted by Eq. (3.16), which was sufficient for the physics we aimed to describe.

An implementation of this method is a straight-forward process using MATLAB, and has been performed for the purpose of this study. However, due to the computational overhead required to time-evolve such a system, the procedure takes a long time to simulate at the required degree of accuracy for any dimension greater than one. Therefore, it is necessary to further develop the methods used, and to improve the implementation of this algorithm to leverage the recent advances in computational acceleration.

3.2.1 Resolution considerations

As the Fourier split-operator method requires special consideration of resolution in both position and momentum space, care must be taken while choosing numerical grids. The reciprocal relationship between position and momentum space is

$$k_{\max} = \frac{2\pi}{\Delta x}, \quad (3.17)$$

which follows directly from the uncertainty relation; better resolution in one space leads naturally to worse in the other. To allow for a condensate to be simulated efficiently in both spaces, it must fit within the grid on which it is defined, and resolve to at least half the size of the smallest structure. It is easy to estimate a radius for the position space wavefunction, following the Thomas–Fermi approximation. It is also rather easy to know that for a non-rotating condensate the wavefunction should occupy the lowest lying mode ($\mathbf{k} = 0$), and those close to it, assuming a harmonic trap. Rotating the condensate, however, has the effect of expanding the wavefunction in position space due to centrifugal forces. Additionally, the momentum space wavefunction also expands with increased angular momentum. With the addition of vortices to the system, there are now small scale structures to resolve. This leads to a system that is difficult to simulate; we have a simultaneously growing position space and momentum space wavefunction.

For a grid to effectively sample the wavefunction and capture all dynamics it will require a sampling rate of at least twice the smallest feature size following the Nyquist sampling theorem [82]. From this it is essential to have a large and finely sampled grid in order to resolve both position and momentum of the wavefunction, with all included features. For the simulations presented below a minimum grid size on the order of $2^8 = 256$ for low rotation

rates, to $2^{11} = 2048$ at high rotation rates in 2D for both X and Y dimensions is necessary to correctly resolve the system dynamics in both position and momentum space with vortices present. One such way of ensuring accurate resolution of the system is to define a sufficient smallest length scale on one such grid (such as position). By ensuring the position grid remains defined with the same lowest increment, it is possible to increase resolution in the reciprocal space with a larger grid. As vortex core sizes are on the order of μm , the above parameters allow between sub- μm (2^{10} and above) to few μm resolution. This also holds true for features in \mathbf{k} -space. Computationally, this can be costly, but quite effective when using compute accelerators (GPUs), which we will introduce next.

3.3 Parallel computing

As the dimensionality of a problem increases, so often too does the time required for performing simulations. One method for accelerating numerical solutions involves the use of multiple compute cores on a central processing unit (CPU) operating independently on different data elements in unison. This form of parallel computation can be achieved through the use of the OpenMP (Open Multi-Processing) application programming interface (API), which defines how a program may parallelise certain elements of code. It allows the developer to fully utilise the power of a multicore processor. However, the limit on how much performance can be gained by this method is set by the number of compute cores available to the system, as well as the limited support offered by compilers. It should also be noted that MATLAB has inherent support for such programming paradigms, and fully abstracts the implementation from the developer. In this instance writing a program from scratch in C/C++/Python/etc. for such a means of parallelisation may not be very beneficial due to the cost of diminishing returns; results would likely be obtained much faster from simply using a multicore supported package, provided that one includes the time to write, as well as simulate. However, this is not always the case, given that the size of problems can often require more cores than available on a single machine.

Another widely used programming paradigm that gets around this CPU core limitation is that of MPI (message passing interface). Where OpenMP allows a user to utilise all available processing cores on a single system, MPI allows the use of an (almost) unlimited number of networked computer systems operating in parallel together, each known as a node. This is the method generally preferred in programs written for compute clusters, where a large number of nodes are available to use. It is the preferable choice for distributed computing applications, with the best performance gains given if there is minimal dependence between data. A bottleneck may occur if data spread over multiple nodes is required for an operation, requiring continual transmission of data between individual nodes. At current data rates this would be limited to bus speeds (assuming an Infiniband optical connection) on the order of tens of gigabytes per second. Compared to a local calculation requiring few to no transfers,

the memory bandwidth (data quantity transferred between RAM and CPU per second) can be as high as 60 gigabytes per second [83]. It is important to note that transfers should be minimised to avoid bottlenecks, but transfers are often necessary to make use of the large number of processing cores. Therefore, to give a significant performance benefit, a large number of cores, a high memory bandwidth, a high-speed interconnect between cores (nodes), as well as sufficient space to store the problem in memory are required.

One means of achieving the required high performance is through the use of graphics processing units (GPUs). GPUs are signal processing devices created to offload from the central processing unit (CPU) much of the heavy computation required for displaying two and three-dimensional graphics. As a result, GPUs have been given the task of performing operations necessary to update a large number of pixels in a short amount of time, as well as complex 3D mathematics for image rendering. This has been achieved by giving the GPUs a large number of specialised compute cores for floating-point arithmetic, effectively operating in parallel. With the advent of general purpose GPU (GPGPU) computing, the ability to exploit these cores for the purpose of numerical computation has become possible. If a problem can be mapped effectively to the hardware of a GPU, it can reduce the overall compute time required for evaluating results, as well as reducing overall power usage for similar performance as cluster nodes. For the previous generation flagship industry standard GPUs used in computational acceleration (Nvidia M2090) the memory bandwidth for the device global memory (equivalent of RAM) is given as 288 gigabytes per second, with upwards of thousands of cores on demand, yielding a theoretical total of 1.41×10^{12} floating-point operations per second (FLOPS), following the formula

$$\text{FLOPS} = \text{cores} \times \text{clock frequency} \times \text{operations per clock cycle.} \quad (3.18)$$

For comparison, Intel Xeon CPU throughput values at best yield approximately 10^{11} FLOPS. As can be seen, performance of an order of magnitude greater can be gained by using a GPU for calculations, over high-performance (Xeon) CPUs. This has been shown to allow for effective implementation of the previously mentioned Fourier split-operator method [78]. We have also shown that it yields performance exceeding that of CPUs for a modest choice of GPU [1], of which we will discuss in detail in a later section.

3.3.1 Parallel operations

For a calculation to fully utilise all of the available throughput of a parallel-capable compute device, it is necessary to break down the problem into parallelised sub-problems, which can be easily achieved if data and required operations are uncoupled (*embarrassingly parallel*). Considering summation as an example, imagine that we have a large vector of floating point values that are to be summed together. The traditional way to solve this would be to iteratively add values to an accumulator, and return the final value at the end as the sum. This simple algorithm is $\mathcal{O}(n)$ complexity in time, as we iterate through each element at a

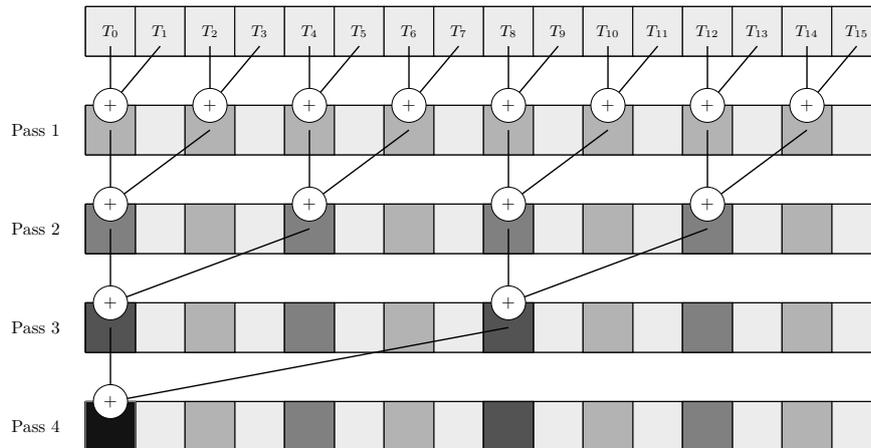


Figure 3.2: A simplified parallel summation algorithm. After each pass the number of elements to sum is halved, leading to a $\mathcal{O}(\log n)$ complexity in time. Since we require only the total sum, this algorithm chooses to read the value at T_0 at the end of the run.

time. Given that summation is associative, we easily parallelise this operation. By dividing our vector amongst a number of available processing cores recursively, we can reduce the computation to $\mathcal{O}(\log n)$ in time (see Fig. 3.2). This algorithm can give a significant benefit when a large number of summations are performed, such as for wavefunction normalisation, and can reduce the accumulated error resulting from floating point addition from $\mathcal{O}(n)$ to $\mathcal{O}(\log n)$ [84].

Another typical algorithm following this is the Hadamard product (element-wise multiplication) of two vectors or matrices. Although in parallel the number of operations and complexity remains the same, the advantage comes from the lack of interdependence between elements. The multiplications can be carried out asynchronously, allowing all freely available cores to work continuously until every element pair is multiplied.

Mapping the above problems onto GPU devices requires an understanding of some of the architecture and hierarchies of the programming model. We will next introduce this model, and discuss some optimal uses of the available memory and structures.

3.4 CUDA programming model

Although many multicore programming models exist for both CPUs and GPUs, with OpenCL and OpenACC being two such examples, we will concentrate on Nvidia's CUDA [85]. CUDA is a mature programming model and API for Nvidia GPUs, and has been well-received for high-performance parallel computing. CUDA operates on a single instruction multiple thread (SIMT) architecture, wherein a single operation is mapped to multiple compute threads independently operating on individual units of data. Writing a program using CUDA C/C++ is very similar to standard C/C++ programming, albeit with some minor differences to

account for control of GPU compute threads. CUDA manages calculations in a hierarchical structure with differing levels of fine-grained control over these threads. At the finest grained level (T) we have compute threads which operate directly on a single datum from memory. The maximum number of threads which can run simultaneously on the device in a single compute unit are known as a *warp*. Threads are grouped into blocks (\mathbf{B}), which is the next hierarchical level, which will optimally contain multiples of the maximum warp number for the device (32 elements is the standard warp size for Nvidia architecture).

At the coarsest level, the blocks are grouped together into a grid (\mathbf{G}), which encompasses the entire problem space. Hardware limits are specified limiting the upper-values of how many elements can appear in these units, and thus the hierarchy exists to allow for fine-grained control on memory usage, and hence performance optimisation on calculations. For a given problem it is necessary to find a mapping from data values in memory to execution threads. To ensure optimal use of GPU cores, the number of threads worked on simultaneously can be (for current hardware) up to 1024 threads per block, though keeping this value divisible by the warp size is recommended to allow an exact mapping to device cores. Optimal values can be found for balancing data computation and transfers, giving, depending on system size, necessary values for block size, as well as grid size. The dimensions of each hierarchical layer are independent of one another, and may be up to 3, with a hardware dependent limit only imposed on the maximum number of elements. Figure 3.3 gives a sample layout for a system of 2D threads and 2D blocks within a single grid.

Unlike programming for a CPU-based system, GPUs require explicit control of several different types of memory, and so it is necessary to be aware of the different aspects of managing this memory. Nvidia's CUDA programming model defines these different physical memories into *global*, *constant*, *shared* and *private*. Global memory is analogous to random access memory (RAM) on a CPU-based system, and is the location where we primarily store data for computation on the GPU. This memory block is accessible (readable and writable) to all threads in the computation. This is also the slowest memory on the GPU, with bandwidths of approximately 10^{11} bytes per second. Constantly reading and writing to global memory can hinder the performance of a computation. For memory that is statically defined at the start of a program and will only be read thereafter, the constant memory can be used. This is a special area of memory that can be used to store constants and other values that are often read during a computation. Once set, these values cannot be modified. The next memory level is shared memory, which is block accessible only i.e. threads within the same block have access. This allows threads to exchange information with close neighbours, and is the preferred method of inter-thread communication. Although higher performance than global, the size of shared memory is much smaller. An ideal use case is the parallel summation example given in Sec. 3.3.1. Performing this algorithm requires the copying of data from global to shared memory, performing the calculations, then copying the results back. Lastly, private memory which maps directly to device cache, is thread-accessible only i.e. each thread has access only to its own private values. This memory can be used for local variables defined

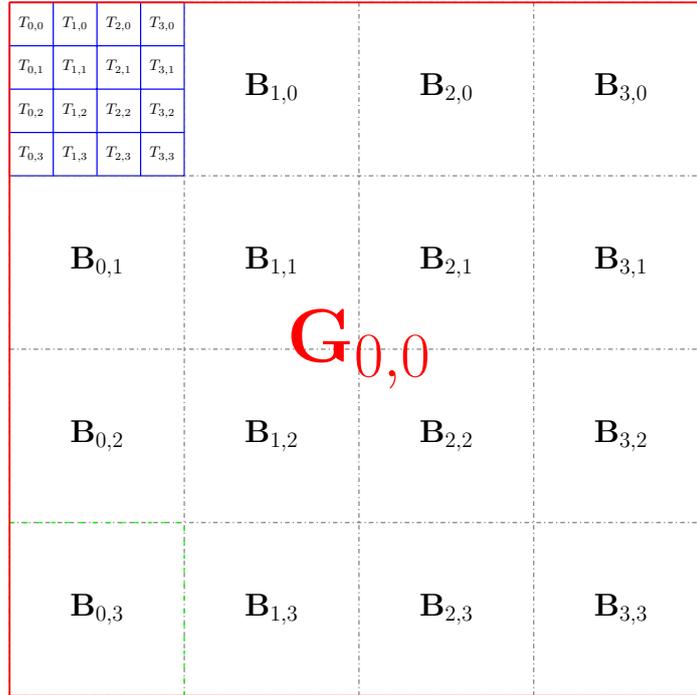


Figure 3.3: A sample CUDA data hierarchy model. Each thread ($T_{x,y}$) has use of private memory, which offers the highest performance overall, though does not allow transfer to other threads. Data within a block ($\mathbf{B}_{x,y}$) can be treated as adjacent and accessible to all other elements in the block using shared memory, and is next highest in terms of performance. Transfers between blocks require the use of global memory, which is slowest overall. Internally all data is stored and indexed lexicographically within the grid ($\mathbf{G}_{x,y}$).

in a function (known as a *kernel* for GPUs). Private memory has the highest performance, but the smallest available size. If possible, copying variables from global memory into private memory, performing all operations on private, then saving back to global can yield the highest performance, as global memory (slow) is read from and written to once each.

Limiting transfers between the CPU/RAM to the GPU/global memory are of high importance, as the slowest line of communication is the PCI-Express bus, connecting the GPU to the host system (≈ 16 GB/s max). Eliminating unnecessary communication will almost certainly allow for a gain in performance. Mapping a problem to the GPU requires parallelising the calculation and removing transfers where applicable. For maximum performance, a sample model of performing a GPU calculation is:

1. Define all variables and data on the host system.
2. Identify the optimal mapping onto the GPU thread execution model.
3. Send the data from RAM to the GPU global memory.
4. Perform computation, with necessary elements copied to shared/private memory.

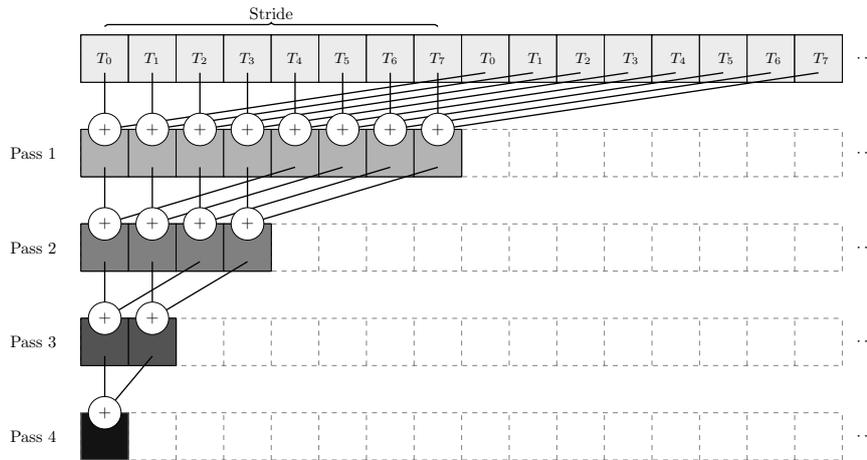


Figure 3.4: An optimised parallel summation, which allows for strided memory access. By halving the size of the stride and block size at every increment the number of evaluations carried out simultaneously will better utilise the available GPU cores.

5. Return final output of computation to host system when completed.

Although idealised and highly simplified, close adherence to such a model should yield significant performance gains compared to multicore CPU-based computation. An important point with memory access in GPUs is that to ensure optimal performance all access should be assumed to be in-order and adjacent. As an example, though the parallel sum method discussed in Sec. 3.3.1 is superior to an iterative summation, it is not an optimal example for GPU architectures [86]. With a small improvement, this can be highly optimised, as highlighted by Fig. 3.4. In this case, the memory access is carried out in strided linear chunks, and by ensuring that all memory accesses are performed this way as much as possible an overall higher degree of performance can be achieved [86]. As GPUs operate on blocks of memory simultaneously, minimising the amount of misaligned memory accesses can optimise the access performance. In this instance, these structures can be packed efficiently together and ensure a large number of copies are performed optimally, with coalesced memory access.

While details for the internal workings of the device are often abstracted, an awareness of these can prove useful. If we assume the number of data elements to be summed is larger than the available number of threads, this routine can perform several passes before the threads become sub-optimally used, as in the previous case. In this way, the summation ensures that the available threads are working mostly optimally.

Given that GPUs are first and foremost image processing devices, their ability to perform Fourier transforms rapidly is a well developed strength. Due to the large number of available cores, fast Fourier transforms (FFTs) can be significantly faster on a GPU than performing the same operation on many CPUs [1]. The CUFFT library allows for a seamless way to take advantage of this performance increase using GPUs. Next, we will discuss an example of GPU-enabled simulations for an experimentally realistic problem making use of the GPU

to perform the necessary operations discussed in Secs. 3.1 and 3.2, and offer some realistic performance measurements.

3.5 GPU-enabled parallel Schrödinger simulations

Here, we will introduce the problem of atomic transport in cold atomic systems, present the resulting model system, and give all the essential physics. Controlling the centre-of-mass movement of atoms has recently become a popular topic of investigation [87]. One family of techniques that aim to solve this are those of *spatial adiabatic passage*. In the following section we will describe the motivation, physical system, and the numerical implementation for the solution of the problem of transporting a single atom among three trapping potential wells. The fully three-dimensional numerical solution of this problem will then be provided using GPU computing methods, and the results discussed for both the physical and computing aspects.

For performance metrics, we will discuss the use of a GPU-enabled Schrödinger equation integrator developed by myself, based on and compared with the results of a multi-core MPI enabled version by T. Morgan and N. Crowley. The solution of the Schrödinger equation for a fully three-dimensional potential, will demonstrate the effectiveness and improved performance compared to standard HPC methods.

3.5.1 Spatial adiabatic passage

Controlling the internal degrees of freedom of atoms is well understood and many spectroscopic techniques exist. External degrees of freedom have, however, only recently become interesting due to the advancements in atom trapping, and techniques for controlling the centre of mass state of a single atom are still in development. One promising group of techniques for the generation of spatial superposition states or high fidelity transfer relies on ideas from spatial adiabatic passage (SAP) [87], which are analogous to STIRAP in optical systems [88]. These techniques are highly robust against variation in the system parameters [89], but suffer from being slow due to the adiabatic requirement. The use of SAP has recently been experimentally demonstrated in Lieb lattices [90], and many other accessible systems have been proposed [91–93].

To describe the method for SAP, we will first consider the case of a two-state system, which can be realised using two separated harmonic potential traps, with ground states $|L\rangle$ and $|R\rangle$. A reduction of the distance between the traps will increase the coupling, and hence tunneling rate, between them. This can be modelled with a two-level Hamiltonian as

$$H = -\frac{\hbar}{2} \begin{pmatrix} 0 & J_{LR} \\ J_{RL} & -2\Delta \end{pmatrix}, \quad (3.19)$$

where $J_{LR} = J_{RL}$ are the couplings between states, and Δ is the detuning of state $|R\rangle$, relative to $|L\rangle$. Assuming an atom initially localised in $|L\rangle$, and with an increase in coupling strength between the levels, the localised atom will tunnel from $|L\rangle$ to $|R\rangle$. However, this processes is difficult to control, as Rabi oscillations introduce an explicit time-dependence as

$$|c_L(t)|^2 \propto \sin^2 \frac{\omega t}{2}, \quad (3.20a)$$

$$|c_R(t)|^2 \propto \cos^2 \frac{\omega t}{2}, \quad (3.20b)$$

where the $|c_{L,R}(t)|^2$ are the populations of the respective states. This time dependence causes the atomic population to continuously tunnel between both traps. It will therefore require precise timing and control to ensure a full, robust transfer of population. From this we see that a double-well potential is a difficult system in which to realise coherent control, though methods such as rapid adiabatic passage (RAP) exist and can allow for this [94]. A more robust method, using three adjacent harmonic traps and the aforementioned matter-wave SAP process, can improve upon the standard double-well system. For this SAP technique, we model the trapping potentials arranged in a single line and coupled with their nearest neighbour only. For three equivalent potentials, L, M, R , with degenerate states $\{|L\rangle, |M\rangle, |R\rangle\}$, the system can be described by the Hamiltonian

$$H = -\frac{\hbar}{2} \begin{pmatrix} 0 & J_{LM} & 0 \\ J_{LM} & 0 & J_{MR} \\ 0 & J_{MR} & 0 \end{pmatrix}, \quad (3.21)$$

where J_{LM} , J_{MR} describe the left-middle and middle-right couplings respectively. Diagonalising this Hamiltonian gives three distinct eigenstates,

$$|\pm\rangle = \frac{J_{LM}|L\rangle \pm \sqrt{J_{LM}^2 + J_{MR}^2}|M\rangle + J_{MR}|R\rangle}{\sqrt{2(J_{LM}^2 + J_{MR}^2)}}, \quad (3.22a)$$

$$|D\rangle = \frac{J_{MR}|L\rangle}{\sqrt{J_{LM}^2 + J_{MR}^2}} - \frac{J_{LM}|R\rangle}{\sqrt{J_{LM}^2 + J_{MR}^2}}, \quad (3.22b)$$

with respective eigenvalues $E_{\pm} = \pm\sqrt{J_{LM}^2 + J_{MR}^2}$, $E_D = 0$. Here, only one eigenstate is of interest. For the zero-valued eigenstate of this Hamiltonian, $|D\rangle$, known as the *dark state*, the dependence on the middle potential vanishes. The state can be written as

$$|D\rangle = \cos \Theta |L\rangle - \sin \Theta |R\rangle, \quad (3.23)$$

where $\tan \Theta = J_{LM}/J_{MR}$ is the mixing angle, following directly from the trigonometric identities $\cos(\arctan(x)) = 1/\sqrt{1+x^2}$, $\sin(\arctan(x)) = x/\sqrt{1+x^2}$. From the adiabatic theorem of quantum mechanics, it is known that if a system has its Hamiltonian perturbed slowly

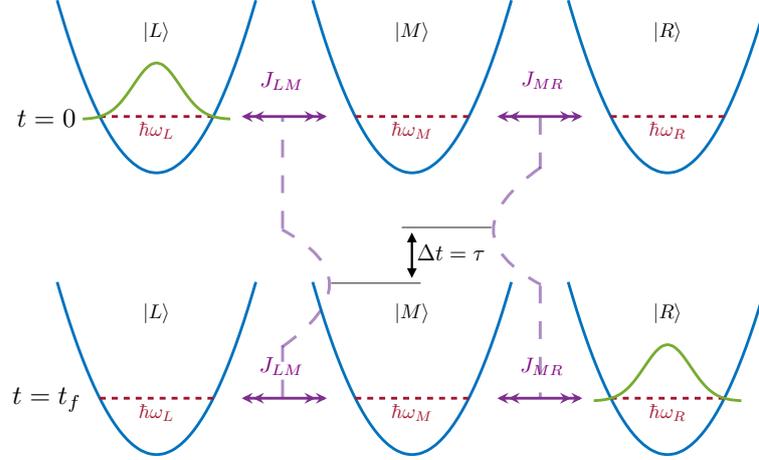


Figure 3.5: Three trapping potential model for matter-wave SAP. The atom (green) is initially localised in the leftmost potential, $|L\rangle$, at $t = 0$, with the couplings between adjacent traps controlled by varying the distance dependent parameters J_{LM} , J_{MR} . By varying the couplings between traps in the manner shown, with J_{MR} increasing initially, followed by J_{LM} after a delay τ , the atom is transferred completely from $|L\rangle$ to $|R\rangle$.

enough, then we can follow its evolution ensuring that it always remains in an eigenstate of the Hamiltonian. By preparing the system in state $|L\rangle$, and varying Θ slowly, we can shift the population from the leftmost harmonic potential to the rightmost, without populating the center. To ensure full transfer from $|L\rangle$ to $|R\rangle$, the mixing angle must change smoothly from $\Theta = 0 \rightarrow \pi/2$. From the properties of arctan, this can be achieved by applying the same spatial variation between $|M\rangle$ and $|R\rangle$, as between $|L\rangle$ and $|M\rangle$, with the latter coupling pulse following a delay, τ . A diagram of this is given by Fig. 3.5. Due to the robustness of this process a variety of coupling profiles can be used [88], as long as they have a smoothly varying background envelope p given by [95]

$$J_{LM}(t) = J_0 p(t - \tau), \quad (3.24a)$$

$$J_{MR}(t) = J_0 p(t). \quad (3.24b)$$

Varying the couplings between the trapping potentials, and hence controlling this mixing angle, can be achieved by either lowering the barrier height of adjacent potentials, or decreasing the distance between them. Here, we will discuss varying the spatial separation of the traps.

Typically, any method to adjust the couplings would be performed with time-dependent potentials. However, a static potential variant can be considered using parallel atomic waveguides, where the separation varies as a function of distance along the parallel axis. If we consider an atom that travels along such a waveguide, the couplings, and hence tunneling

rates, seen by the atom in the waveguide are altered as the atom propagates. Such work has been discussed and considered in a realistic system for two spatial dimensions [96]. Although a two dimensional model is effective at describing much of the relevant dynamics, the realism of the model is reduced by the lack of a third dimension. This is due to the lack of effects stemming from dispersion, curvature of the waveguides, as well as the absence of any such eigenstates along this dimension.

3.5.2 Atom-chip model

As discussed previously, to fully understand ultracold atom dynamics in appropriately shaped waveguide systems we must investigate the fully three-dimensional model. One method of creating the required potential landscape in experiments is through the use of atom-chips [97, 98]. These systems consist of micro-fabricated current-carrying wires, and can be used to create a variety of trapping potential shapes for controlled guidance of the atomic centre-of-mass [99]. The currents produce a magnetic field around individual wires, each of which has a minima at the wire core. Assuming the wire thickness to be negligible, this magnetic field at position \mathbf{r} can be calculated using the Biot–Savart law

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi} \oint I \frac{d\mathbf{l} \times \hat{\mathbf{r}}'}{|\mathbf{r}'|^2}, \quad (3.25)$$

where I is the current through the wire, μ_0 is the vacuum permeability, $d\mathbf{l}$ is the differential wire length, and $\hat{\mathbf{r}}'$ is the unit vector along $\mathbf{r}' = \mathbf{r} - \mathbf{l}$. To be able to trap atoms the minima must be raised to a position above the wire surface. This can be done using an orthogonally applied bias field, B_b , which raises the minima to a height of

$$\mathbf{r}_0 = \frac{\mu_0 I}{2\pi B_b}, \quad (3.26)$$

above the surface. Though, an issue still remains with the presence of the magnetic minima. If the field drops to zero at the centre of the trap, the atoms can be lost due to Majorana spin flips [100]. This can be prevented with the application of an additional field, B_{ip} , parallel to the wire direction, lifting the degeneracy of the atomic states, and ensuring they remain trapped. Spatial and temporal adjustments of the potentials are possible, with a fine degree of control, either during the production process, or by using time-dependent currents. These have been studied extensively in recent years for highly controllable trapping potentials [101, 102], and as atomic manipulators [103].

For this work, we model the system as three adjacent wires on the atom-chip surface. The direction of propagation is along z , and an additional harmonic oscillator potential, $V_z = m\omega_z^2(z - z_0)^2/2$, is applied in the same direction to impart motion to the atom, which is initially at the $z = 0$ position of the atom-chip. We set $z_0 = (\max z)/2$ to ensure the oscillator potential is symmetric around the centre of the atomchip. This potential also conveniently guarantees that the wavefunction refocuses after the transition at the opposite side of the

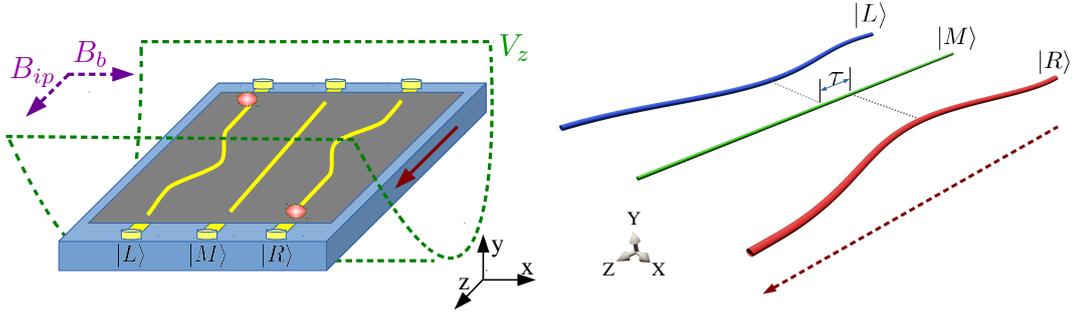


Figure 3.6: Schematic of the atom-chip and the resulting potentials. Reprinted from Morgan *et al.* [1].

atom-chip. A schematic of the atom-chip device and the respective potentials is given by Fig. 3.6.

The initial state is created by localising the atom with the help of a barrier at one end of the potential ($z = 0$), and in the leftmost waveguide using an additional barrier potential. After finding the ground state, the barriers are removed, and the atom is allowed to propagate along the length of the waveguide. The populations in each waveguide $|c_X|^2$ are tracked at each step of the process as

$$|c_X|^2 = \int_{\mathcal{V}_X} d\mathbf{r} \Psi^* \Psi \quad (3.27)$$

where \mathcal{V}_X is the volume encompassing each waveguide $X \in \{L, M, R\}$, and Ψ is the state of the system. The final populations were taken as the atom approached the other classical turning point of the harmonic oscillator along z . The fidelity of the process could then be calculated by comparing the initial populations in $|L\rangle$ and final ones in $|R\rangle$, as well as any ones left in $|M\rangle$.

Given that fully three-dimensional simulations of the Schrödinger equation are numerically expensive, the use of GPU computing methods were ideal for accelerating the simulation [78]. At the time of writing, as far as we had been aware, no other work using GPU computing to solve a three dimensional Schrödinger equation had been presented. We will now discuss the data and metadata of the simulations.

3.5.3 3D Simulations

Simulations of the proposed system assumed a single ${}^6\text{Li}$ atom localised in the left waveguide. A localised harmonic oscillator potential was added along z , with the transverse guiding potential, and the resulting ground state solution found numerically. The harmonic oscillator width was chosen to closely match the expected size in the transverse direction. The atom was then allowed to propagate along the waveguide potential (z).

As the shift in magnetic moment of the atoms is given by $\Delta E = -\boldsymbol{\mu} \cdot \mathbf{B}$, for regions with a

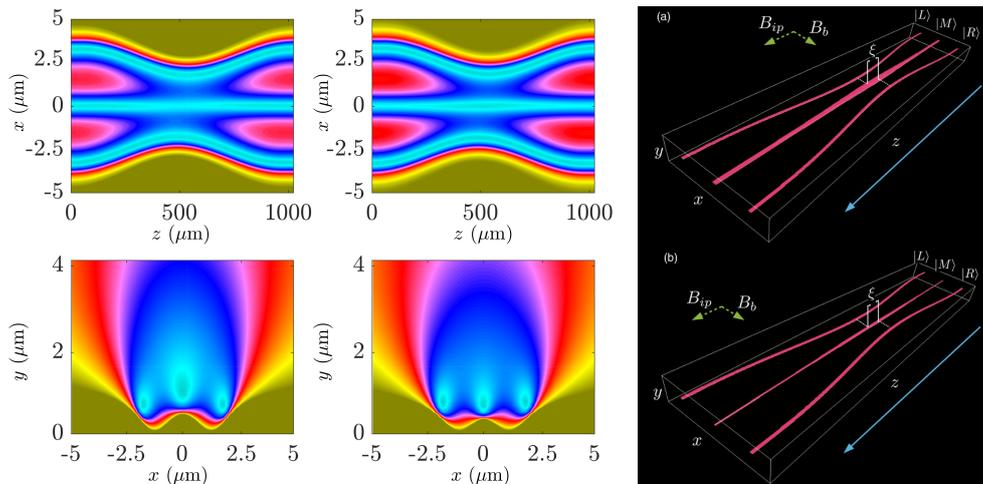


Figure 3.7: (Left) Two dimensional slices through the potential centre along $x - z$ (top) and tunneling region in $x - y$ (bottom). The out of resonance middle trapping potential can be seen where all three currents are the same. (Middle) Optimised currents which allow the same field value in the tunneling region ensure that the potentials are in resonance with each other. (Right) 3D isosurface plots of the numerically calculated magnetic minima potentials showing the out of resonance (a) and optimal (b) fields. Adapted from Morgan *et al.* [1].

larger magnetic field the atom will experience a greater energy shift, and thus the assumption of all traps being on resonance (degenerate) has to be carefully checked. In fact, the simulations showed that the addition of the magnetic fields stemming from the different wires at the center of the atom-chip leads to the central potential moving out of resonance with the outer two. To make SAP work, it therefore required adjusting the current in the central wire, such that the magnetic minima were in resonance within the tunneling region near the atom-chip centre. The resulting potentials for non-optimal (left) and optimal (middle) currents are shown in Fig. 3.7, which also depicts a three dimensional isosurface of the potential minima along the chip surface for both situations (right).

The populations for both the direct tunneling case, and the matter-wave SAP processes are shown in Fig. 3.8. The direct tunneling case can be seen to show Rabi-type oscillations between the waveguides, while the matter-wave SAP process shows a much cleaner transfer, and only a minor occupation of the central potential. The dependence of the transfer probability on the current in the central wire is shown in Fig. 3.9.

3.5.4 GPU computing performance

Given the large parameter space over which this system could be evaluated (e.g wire current, spatial separations, trap frequencies), a large number of simulations were required to determine optimal system behaviour. As discussed in Sec. 3.4, one example where GPU computing offers large performance gains are FFTs, which makes the Fourier split operator method an

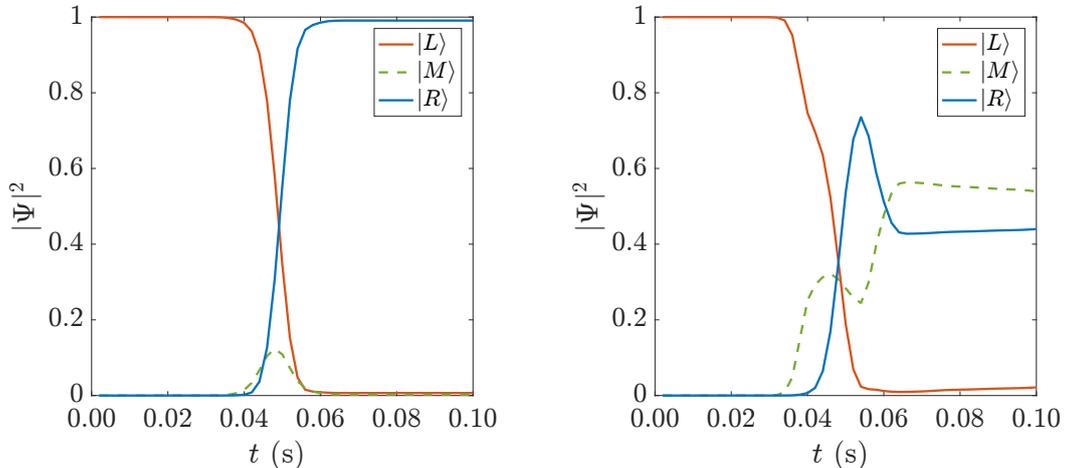


Figure 3.8: Transfer fidelities for the three trapping potentials are given over time for both the matter-wave SAP process (left, particle starting in $|L\rangle$) and direct tunneling (right, particle starting in $|L\rangle$). Matter-wave SAP clearly shows greater population transfer in comparison with a direct tunneling approach. Reprinted from Morgan *et al.* [1].

ideal candidate for GPU systems [78]. The body of work for implementing this algorithm was using C, CUDA and Nvidia’s CUFFT libraries for the Fourier transforms, whereas the MPI-enabled code was implemented using C.

To demonstrate the performance offered by GPU computing we compared it to using FFTW with MPI, a well used parallel programming library and paradigm. The MPI implementation allows code to be run across multiple machines, benefiting from the parallelism which may be offered by a supercomputing cluster. Although MPI-enabled FFTW is fast and supports extremely large grid sizes, it requires cluster access of a significant size to be a viable option for this type of system. The MPI work on this project was carried out on the Irish Center for High-End Computing (ICHEC) supercomputer system “Stoney” [104] over the period 2011 to 2012, with all performance metrics data calculated therefrom. This cluster system had 64 available compute nodes, each housing two 2.8GHz Intel Xeon X5560 processors with 4-cores each, and a total of 48GB of RAM per node with inter-node communication using double data rate Infiniband.

Due to the hardware limited memory on the GPU, and because the dynamics along the $x-z$ plane were of most importance, the grid-size of the simulations were scaled as $256 \times 64 \times 1024$ ($x \times y \times z$). Of next importance were the choice of timesteps for the simulations. To ensure minimal loss in accuracy, the timesteps were chosen as $\Delta t = 10^{-6}$ s. By approximating the waveguides as harmonic oscillators, the relevant timescales of the dynamics are of the

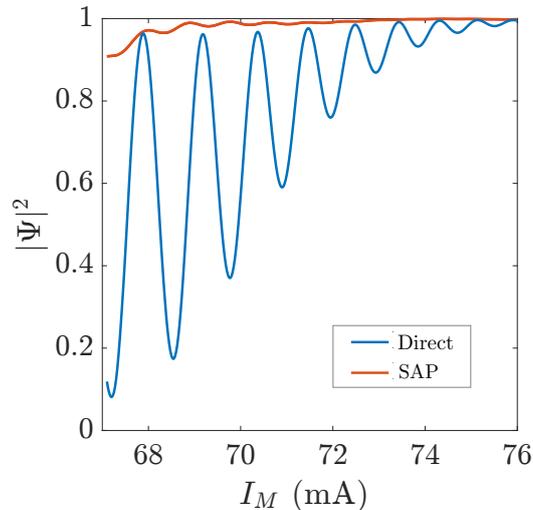


Figure 3.9: The final state population versus middle wire current for direct tunneling and matter-wave SAP. The robustness of the SAP technique can be seen, and gives a large range of currents with population transfer fidelity $\approx 99\%$. The oscillations in the direct tunneling regime are due to the presence of time-dependent Rabi couplings. Reprinted from Morgan *et al.* [1].

order $T_t \approx 10^{-4}$ s, and so we can accurately capture all relevant dynamics in the transverse direction. The timescales in the longitudinal direction, which requires a large oscillation period compared to the transverse plane to ensure the adiabaticity condition ($T_1 \approx 10^{-1}$ s), are also fulfilled. For the GPU simulations, the test system was an Intel Core i7 2600K CPU at stock frequency, 8GB DDR3 memory operating at 1600 MHz, 7200 RPM HDD, Nvidia GeForce GTX 580 with 3GB of onboard memory running at 783 MHz GPU core frequency, 1566 MHz shader processor frequency, and 2010 MHz memory frequency. For all simulations the desktop was running Ubuntu 11.10 64-bit operating system and all calculations were performed in double precision (64-bit floating point) where applicable.

Table 3.1 shows the approximate timings for the completion of runs using GPU and CPU codes. Not only does GPU computing offer a 6-fold improvement over a single CPU, it also allows us to achieve a performance level which is comparable to an 8-node 8-core (64 cores) core MPI enabled CPU calculation. Even for a modest choice of gaming GPU this offers substantial performance gains. Higher performance was achieved by using specific compute accelerators designed for double precision arithmetic. Making use of eight Nvidia M2090 GPUs available at OIST, terabytes of numerical results were generated, and allowed the problem to become tractable on a short timescale.

Device	Num. Devices	Timing	Rel. Improvement
CPU (MPI)	8	~6 Hr	1.0×
	16	~4 Hr	1.5×
	32	~1.5 Hr	4.0×
	64	~1 Hr	6.0×
GPU	1	~1 Hr	6.0×

Table 3.1: The approximate times taken to simulate the propagation of an atom through the atom chip system on both GPU and CPU. Reprinted from Morgan *et al.* [1].

3.6 GPUE: GPU Gross–Pitaevskii equation solver

Given the effectiveness of GPU computing in the simulation of the linear Schrödinger equation system for SAP, we next applied the newly-developed techniques to simulating Bose–Einstein condensates, which formed the bulk of work during my thesis. The body of software developed for this project has been released as the tool “GPUE”, available at <https://github.com/mlxd/gpue> [105]. Performance testing of this code was carried out by Peter Wittek, ICFO, Barcelona [106]. A comparison was performed between GPUE, the Trotter–Suzuki (TS) package developed by Wittek *et al.* [107], and the mature GPELab software suite for MATLAB [108, 109]. The sample results taken for time evolution are given in Fig. 3.10. GPUE and GPU-enabled TS clearly beat MATLAB, and CPU performance by a significant margin. Although TS is a more generalised suite for computing, as far as we are currently aware the GPU computation does not yet allow for Gross–Pitaevskii solutions with angular momentum. GPUE is currently the optimal choice for rotating condensate systems out of the examined software suites.

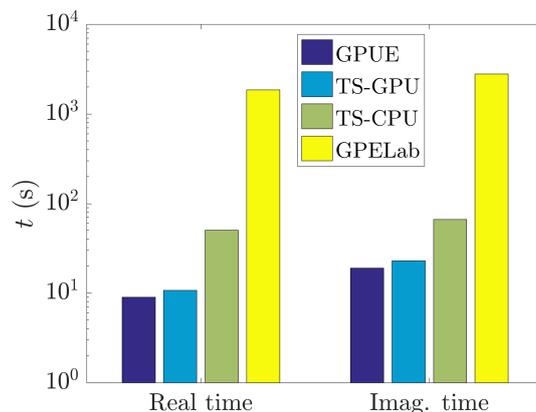


Figure 3.10: Performance benchmark of GPUE and other simulation packages for the evolution of a harmonically trapped atom in a superposition state between ground and first excited states. Lower numbers are better and give results in faster times. Data adapted from [106].

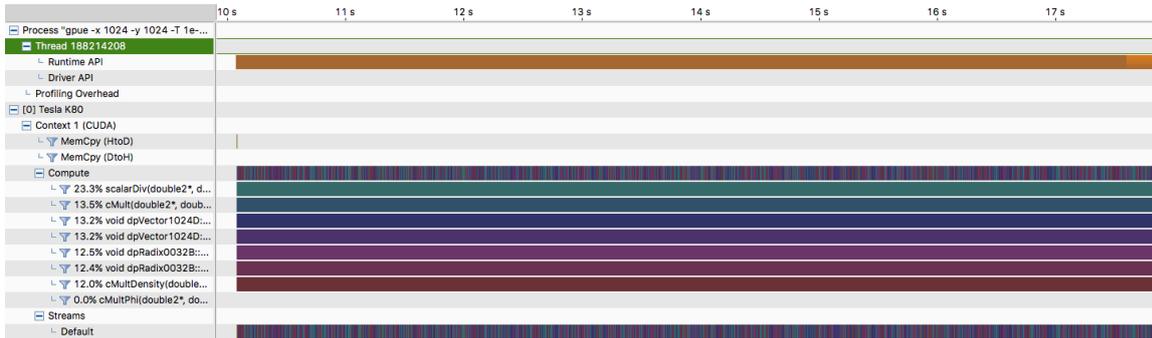


Figure 3.11: Nvidia Nsight performance analysis of GPUE for real time evolution simulation for 1010 steps at $2^{10} \times 2^{10}$ resolution. The respective kernel calls and total utilisation are listed on the left.

Kernel	Info	Avg. runtime	# Runs	Total time
Mem. copy [H2D]	Memory copy from host to GPU	2.312 ms	11	25.432 ms
Mem. copy [D2H]	Memory copy from GPU to host	2.18 ms	9	19.62 ms
cMult	Complex mult. in time ev.	0.342 ms	3030	1.0363 s
cMultDensity	Complex mult. in time ev. for nonlinear op.	0.456 ms	2020	0.921 s
scalarDiv	Renorm. of Ψ following FFT	0.2216 ms	8080	1.791 s
dpRadix0032B	Internal CUFFT operation	0.237 ms	8080	1.915 s
dpVector1024D	Internal CUFFT operation	0.252 ms	8080	2.0362 s

Table 3.2: Kernel usage and timings for 1010 steps of real time evolution with angular momentum.

Figures 3.11 and 3.12 demonstrate some of the resulting calls to different segments of the code, with timings given in Tables 3.2 and 3.3. The important data of the figures is both the kernel percentage utilisation, and that the operations are mostly saturating the available number of GPU cores. The data shows the average time spent in each individual kernel during both real and imaginary time evolution for 1010 steps at $2^{10} \times 2^{10}$ resolution with (real) and without (imaginary) angular momentum operators. As can be seen, the inclusion of angular momentum operators lead to a performance hit, compared with an imaginary time evolution for a static condensate. While further optimisations can almost always be provided for such simulations, the performance of the software as a whole is defined by its slowest component. In this case the routines are equally met in performance by the Fourier transforms, which are already fully optimised as an external library. As such, improving performance much beyond this with the other kernels will be wasteful in time and resources.

A simplified sequence and state diagram combination is given in Figs. 3.13 and 3.14 which describes the operating process for GPUE. A document listing all aspects of component dependencies and intercommunication is available at [105]. For brevity, we will refer the reader to this location for more information.¹

¹The documentation is built using Doxygen with the command “doxygen ./docs/gpue_doxy.conf”, and requires the *dot* package for figure generation. This may change during future releases.

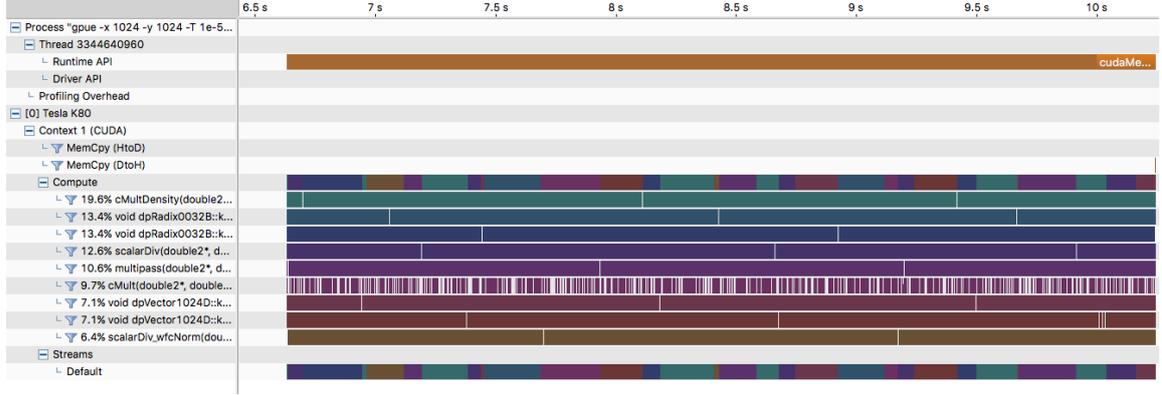


Figure 3.12: Nvidia Nsight performance analysis of GPUE for imaginary time simulation for 1010 steps at $2^{10} \times 2^{10}$ resolution with angular momentum. The respective kernel calls and total utilisation are listed on the left.

Kernel	Info	Avg. runtime	# Runs	Total time
Mem. copy [H2D]	Memory copy from host to GPU	2.311 ms	3	6.933 ms
Mem. copy [D2H]	Memory copy from GPU to host	1.85 ms	5	9.25 ms
cMult	Complex mult. in time ev.	0.342 ms	1010	0.345 s
cMultDensity	Complex mult. in time ev. for nonlinear op.	0.346 ms	2020	0.698 s
multipass	Optimised parallel summation	0.125 ms	3030	0.378 s
scalarDiv	Renorm. of Ψ following FFT	0.221 ms	2020	0.446 s
scalarDiv_wfcNorm	Normalisation of wavefunction during ev.	0.226 ms	1010	0.228 s
dpRadix0032B	Internal CUFFT operation	0.237 ms	4040	0.957 s
dpVector1024D	Internal CUFFT operation	0.251 ms	2020	0.507 s

Table 3.3: Kernel usage and timings for 1010 steps of imaginary time evolution for a non-rotating condensate.

3.6.1 Angular momentum operators using Fourier split-operator method

As discussed earlier, in the presence of large values of angular momentum, the condensate wavefunction will accommodate many vortices. To ensure a well ordered lattice, more consideration is required than to just directly numerically solve the GPE at the required rotation rate. Assuming an initial Gaussian guess, and using the imaginary time evolution algorithm to find the ground state, a large number of vortices will enter the condensate from the edge and compete for lattice sites to form the expected Abrikosov pattern. Due to the highly dense spectrum of the condensate close to the ground state in this regime, only minimal energy shifts will be given for deviations from the perfect Abrikosov geometry. As a result, it can take a significantly long time to reach an ordered state for rotation frequencies close to the transverse trapping frequency [110]. To overcome this issue, one can choose to follow the ground state of the condensate with a ramp of the rotation rate. This essentially mimics adiabatic evolution, and allows for the determination of the vortex lattice ground state for all rotation frequencies.

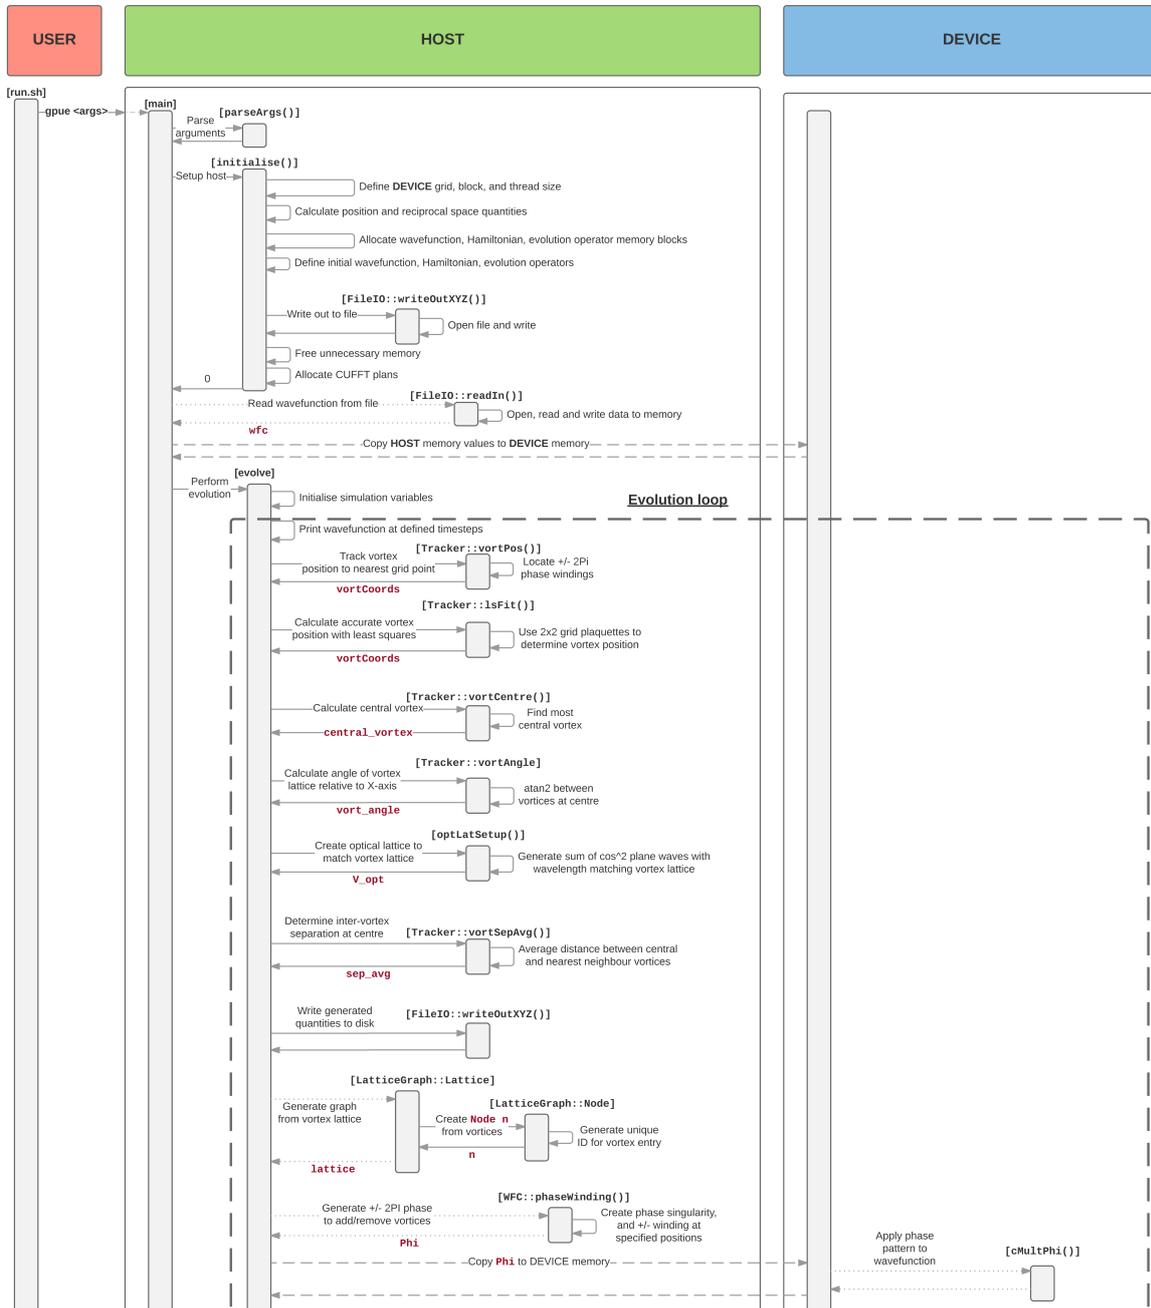


Figure 3.13: Simplified combined sequence and state diagram for GPUE operation (1 of 2). The operation procedure of GPUE is outlined in sequence from top-to-bottom. While much of the setup and analysis takes place on the host (CPU), the device (GPU) is used to offload all the time-evolution calculations. After setup, the wavefunction and all required operators are sent to the GPU.

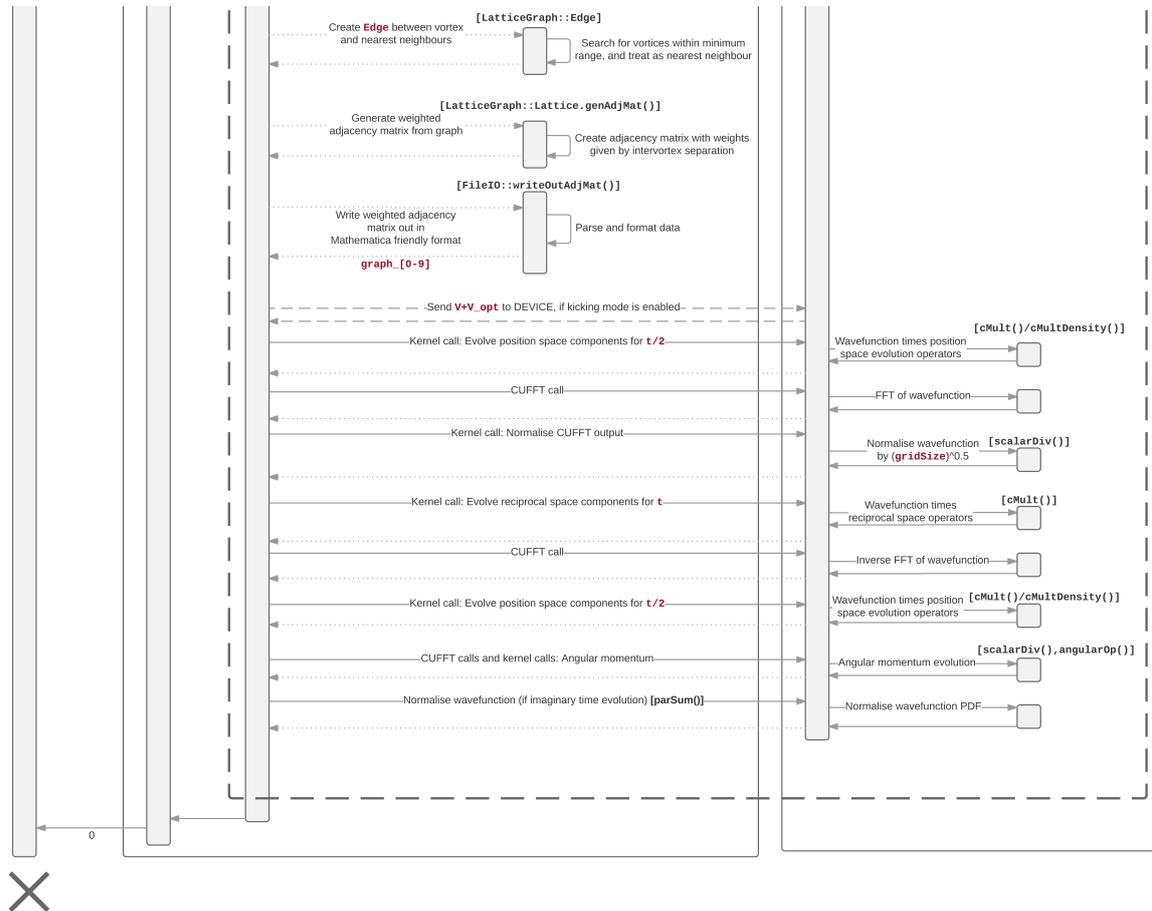


Figure 3.14: Simplified combined sequence and state diagram for GPUE operation (2 of 2). Following the completion of the time evolution after a predetermined number of steps, the wavefunction is unloaded from the GPU and returned to the CPU for output. Minimising this transfer allows for optimal performance from the device. Further details of dependencies and data flow are given by [105, docs/gpue.pdf].

The Fourier split-operator algorithm described earlier works well in handling cases where the individual operators live in position or momentum space respectively. However, the angular momentum operators are a combination of both spaces. Taking the angular momentum operator along the z -axis, $L_z = xp_y - yp_x$, and applying it to the wavefunction requires each basis element to be in a different space in the different directions. For applying this operator we must therefore Fourier transform along a single dimension, multiply by the respective \mathbf{k} -space component, take the inverse, multiply by the respective \mathbf{r} -space component in the other direction, and then perform this operation along the other dimensions, summing the results.

This accrues an error which is not encountered using methods that are solely in position or momentum space. Following the process given in Sec. 3.2 the error can be determined by checking the commutativity of the respective components of the angular momentum operator

as

$$\begin{aligned}
\alpha_1 &= [xp_y, -yp_x] = [xp_y, -y]p_x - y[xp_y, p_x] = -[-y, xp_y]p_x + y[p_x, xp_y] \\
&= -\left(\cancel{[-y, x]p_y}^0 + x[-y, p_y]\right)p_x + y\left([p_x, x]p_y + \cancel{x[p_x, p_y]}^0\right) \\
&= -x\cancel{[-y, p_y]p_x}^{-i\hbar} + y\cancel{[p_x, x]p_y}^{-i\hbar} \\
&= i\hbar(xp_x - yp_y). \tag{3.28}
\end{aligned}$$

The complex error term can be seen as, in the case of the above implemented evolution, allowing the angular momentum operator to change from imaginary time to real-time, and vice-versa in each respective case. To overcome this, we simply swap the application order of the operator components, between odd and even steps during the evolution. Starting with the alternate order we obtain a value of $\alpha_2 = [-yp_x, xp_y] = i\hbar(-xp_x + yp_y)$. Since we are applying these operators to the condensate we can overcome the error of one term by the application of the other, as

$$\exp i\alpha_1 \exp i\alpha_2 = 1. \tag{3.29}$$

Although alternating will provide a cancellation of this error, it can be assumed that for large timesteps the error will have a significant contribution to the overall dynamics, as the wavefunction evolves during each timestep. For greater accuracy of this method one can perform a decomposition following Eq. (3.14) for a third-order error, or using the above splitting for second-order.

An example of the density of the ground state and the associated wavefunction phase at a rotation frequency of $\Omega = 0.995\omega_x$ is given in Fig. 3.15 at a resolution of $2^{11} \times 2^{11}$ (2048×2048). Although aliasing may be apparent in the phase, this is due to the limited resolution of the computer monitor (printer). The presence of a well ordered Abrikosov lattice is clearly visible.

3.6.2 Vortex tracking

To efficiently follow the dynamics of individual vortices, a robust algorithm is needed to track their positions. One could track regions where the density drops to zero. However, this gives very little information on the topological excitation, and may miss many vortices, as the numerical wavefunction may never truly approach this value. A more effective way is to locate the $\pm 2\pi$ charge in the wavefunction phase, which is a signature of quantum vortices. For this we examine each 2×2 subgrid of the underlying lattice and check if the phase rotates from $-\pi$ to $+\pi$ (or vice versa). After an initial pass to identify vortex locations closest the nearest grid element, a least-squares fit is performed to more accurately determine the vortex core position [111]. Linear least squares is used generally for an overdetermined linear system

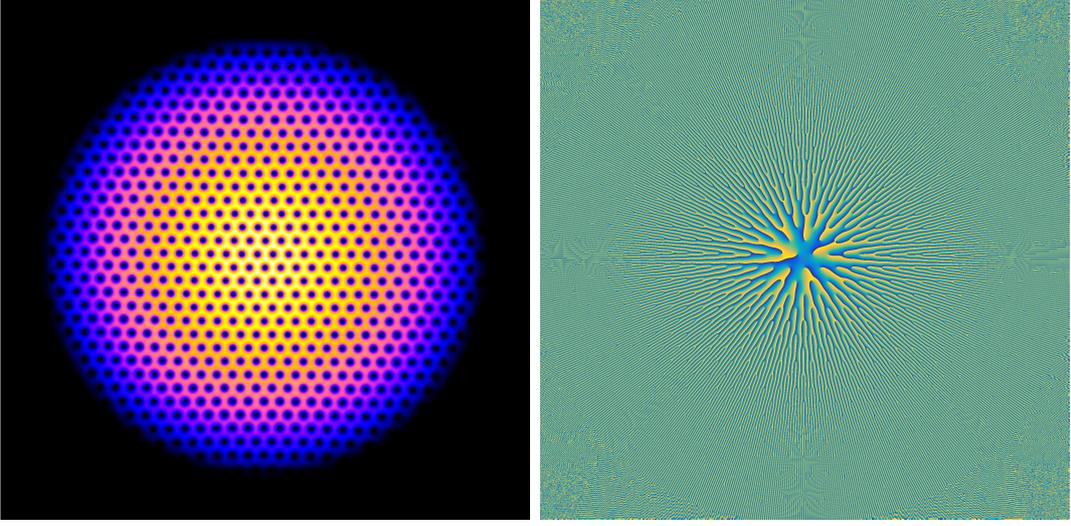


Figure 3.15: Condensate density (left) and phase (right) at a rotation rate of $\Omega = 0.995\omega_x$ for a $2^{11} \times 2^{11}$ grid showing approximately 600 vortices in the visible density regions. For both images the box size is $700 \mu\text{m} \times 700 \mu\text{m}$.

$\mathbf{A}\mathbf{r} = \mathbf{b}$, where unique solutions are unlikely to exist. Thus, for a solution, we seek the best fit plane that minimises the error, of the form

$$S(\mathbf{r}) = \sum |b_i - \sum A_{ij}r_j|^2 \quad (3.30)$$

where S is the objective function to be minimised, following $\mathbf{b} = \arg \min S(\mathbf{r})$. The solution of this minimisation problem is given by

$$\mathbf{A}^T \mathbf{A} \mathbf{r} = \mathbf{A}^T \mathbf{b}, \quad (3.31a)$$

$$\mathbf{r} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}. \quad (3.31b)$$

The best-fit plane is sought of the form $a_0c + \sum_i^m a_i r_i = f(\mathbf{r})$ which for a two-dimensional system, $\mathbf{r} = (x, y)$, is given by the matrix,

$$\mathbf{A} = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix}. \quad (3.32)$$

The above matrix is composed of all possible planes that can fit over a square 2×2 grid plaquette, and

$$\mathbf{b} = \left(\Psi(x_0, y_0) \quad \Psi(x_0, y_1) \quad \Psi(x_1, y_0) \quad \Psi(x_1, y_1) \right)^T, \quad (3.33)$$

are the wavefunction values around the sampled 2×2 grid. Upon evaluating the vector \mathbf{r} above, one can obtain the best fit plane solution as

$$\begin{pmatrix} x \\ y \\ c \end{pmatrix} = \begin{pmatrix} (-\Psi(x_0, y_0) + \Psi(x_0, y_1) - \Psi(x_1, y_0) + \Psi(x_1, y_1))/2 \\ (-\Psi(x_0, y_0) - \Psi(x_0, y_1) + \Psi(x_1, y_0) + \Psi(x_1, y_1))/2 \\ 3\Psi(x_0, y_0) + \Psi(x_0, y_1) - \Psi(x_1, y_0) - \Psi(x_1, y_1) \end{pmatrix}. \quad (3.34)$$

The goal is to find where both the real and imaginary components cross through zero, and thus we seek a solution of the form $x + y = -c$. Rearranging the above equations as

$$\begin{pmatrix} \Re(x) & \Re(y) \\ \Im(x) & \Im(y) \end{pmatrix} \begin{pmatrix} \delta x \\ \delta y \end{pmatrix} = - \begin{pmatrix} \Re(c) \\ \Im(c) \end{pmatrix}, \quad (3.35)$$

and again solving the linear system by inverting the left-hand matrix and multiplying across allows one to seek the corrections to the vortex position, $\delta \mathbf{r} = (\delta x, \delta y)$.

With this, we can accurately determine the position of the vortices with high precision. To track their motion during the evolution, we create an initial list of positions and give each vortex a unique identifier (UID). Assuming the vortex cores can travel a limited distance (some multiple of the grid resolution) between time steps, we can say at subsequent times which vortex has moved to the newly found positions.

This process is performed by representing the vortices as a graph, each with an assigned unique identifier, associated location, phase winding and on/off flag. Edges are created between vortices that are separated by at most root-two the average of the inter-vortex spacings. A finite boundary is chosen to examine only vortices in areas of significant condensate density, since vortices can easily appear and disappear close to the condensate boundary. Any vortex which appears without association to an initial vortex, or any tracked vortex that crosses the boundary, is switched off and remains so for all analysis. A graph for an initial ($t = 0$) vortex lattice shown in Fig. 3.16, with the identifiers indicated on each node, and the neighbouring distances indicated by the edge weights.

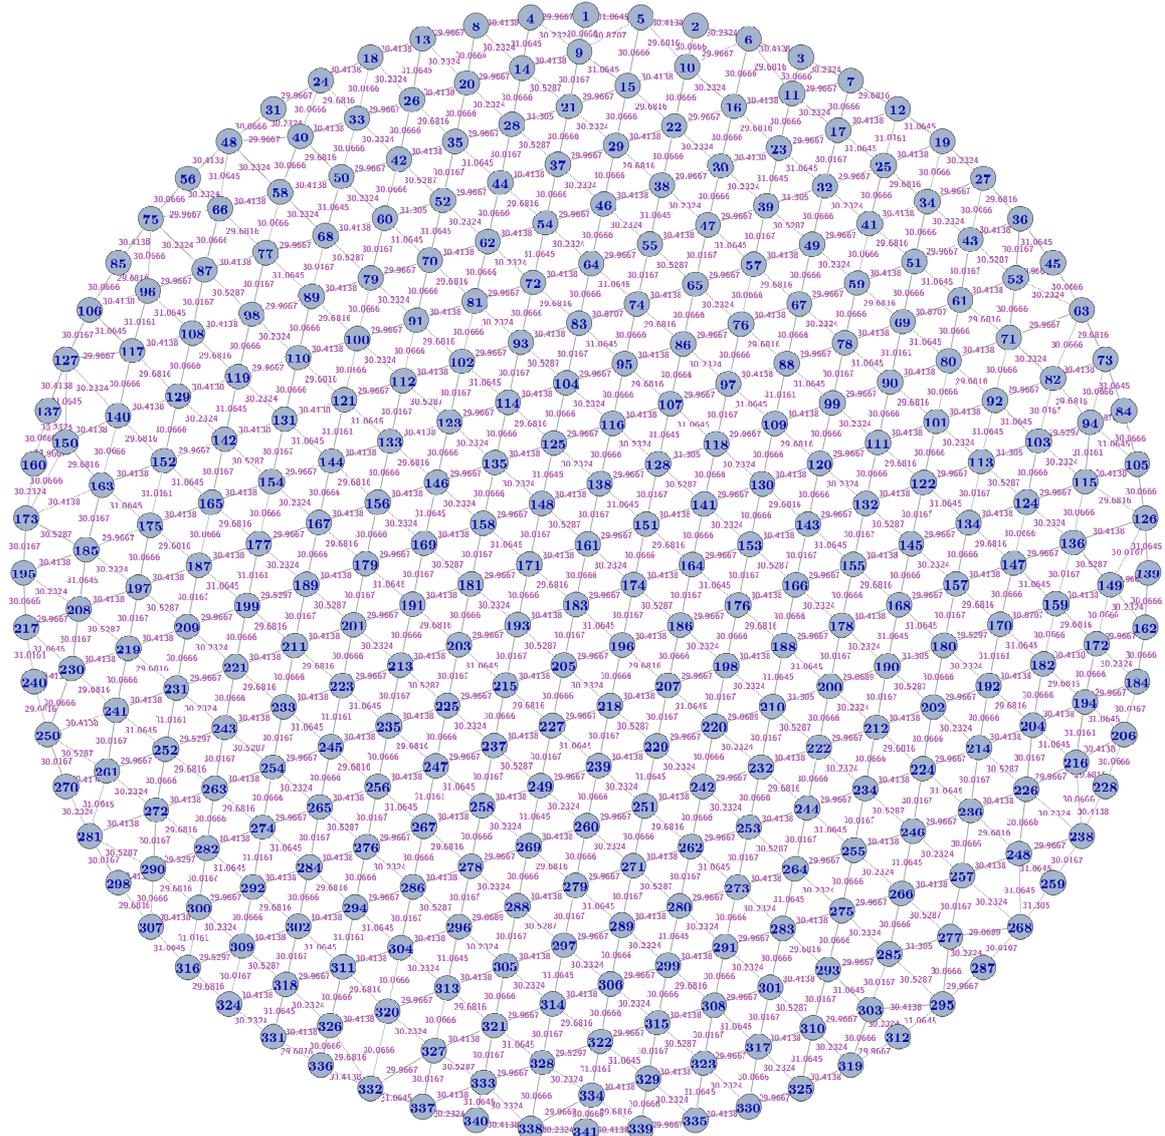


Figure 3.16: Graph of vortex lattice positions indicating the vortex identifier and the intervortex distances in units of grid-spacing $\Delta = \Delta x = \Delta y$. A hard-walled boundary is chosen such that the vortex distances remain almost uniform, and are shown here with a mean value of $\bar{r} = 30.2451$ and variance of $\sigma^2 = 0.16118$.

Chapter 4

Bose–Einstein condensate dynamics

The purpose of this chapter is to introduce the reader to vortex states of BECs, and eventually discuss manipulations of these states. In the following sections we will examine both the static and the dynamical solutions of the Gross–Pitaevskii equation (GPE), using the methods previously discussed. We will first briefly discuss few vortex solutions, with an analysis of their properties via an examination of the velocity field.

We will follow this by providing the details of the model system of a vortex lattice in a rapidly rotating BEC, using the theoretical framework developed in Chapter 2. Next, we will discuss the application of two distinct ways for manipulating and controlling the condensate: (1) applying an external potential to the trapped condensate for a short timescale, and (2) direct manipulation of the wavefunction phase. Both of these are experimentally realistic, and though they can be experimentally implemented by similar means, they serve two very different purposes, and they will be treated as such.

4.1 Simulating Bose–Einstein condensate dynamics

In Sections 2.1.3 and 3.2 we have outlined both the analytical Thomas–Fermi (TF) and numerical solutions of the GPE. It is instructive to compare the results from these two methods. For a stationary condensate in the ground state of a harmonic trapping potential the profile and width should be comparable, with the TF solution deviating only in the low-density regions. Fig. 4.1 shows a comparison of the two-dimensional profile for both methods, for condensates of ^{87}Rb with $N = 10^5$ atoms, and trapping frequencies of $\omega_{\mathbf{r}} = 2\pi \times (1, 1, 16)$ Hz. Both solutions show good agreement, deviating only where expected, and hence showing a well developed numerical procedure.

While the Thomas–Fermi solution closely agrees with the numerical solution of the Gross–



Figure 4.1: The numerical solution (left) and Thomas-Fermi (middle) solutions for a two-dimensional condensate of ^{87}Rb with $N = 1 \times 10^5$ atoms. The lineplot (right) is a central cut through both profiles showing the close match.

Pitaevskii equation, this solution is only applicable for stationary states with negligible kinetic energy, i.e. when the nonlinear interaction dominates over the kinetic energy term. For more complex problems involving dynamics, full numerical integration of the Gross–Pitaevskii equation is required.

One example where the TF approach fails is when investigating superfluid vortex dynamics. To generate vortices in the condensate angular momentum must be added to the system. Seeding a single vortex in the condensate requires that the frequency of rotation, Ω , must be higher than the critical rotation frequency Ω_c , as discussed in Sec. 2.2.2, and numerically, this is most easily simulated in the co-rotating frame as given by Eq. (2.30). Experimentally, there are many ways to create vortices in a condensate, such as stirring with a blue-detuned laser beam [32], carefully inverting the trap bias field potential [43, 112–115], or through the use of artificial gauge fields [116], to name but a few. In the next sections we will concentrate on solutions where the vortices are already present in the condensate and have settled into the lowest energy configuration.

4.1.1 Few vortex condensates

As discussed in Section 2.2, the study of quantum vortices remains an active area of research. As the condensate kinetic energy scales as $E_k \propto l^2$ (Sec. 2.2.2), increased angular momentum leads to the appearance of more singly charged vortices rather than one multiply charged vortex. For a single vortex in a rotating condensate the only stationary solution that exists is if the vortex resides at the exact centre of the trapping potential, making the system radially symmetric. For two vortices the radial symmetry of the system is broken, with the two vortices arranging themselves in the most favourable position to minimise the energy of the system, as with all higher vortex states. Placing vortices of the same sign into a condensate will allow them to behave as identically charged particles, and they will repel, and form ordered lattices with increasing vortex numbers (see Fig. 4.2). Similarly, if instead we pair a vortex and an antivortex, they will attract, and remain in constant motion, or disappear

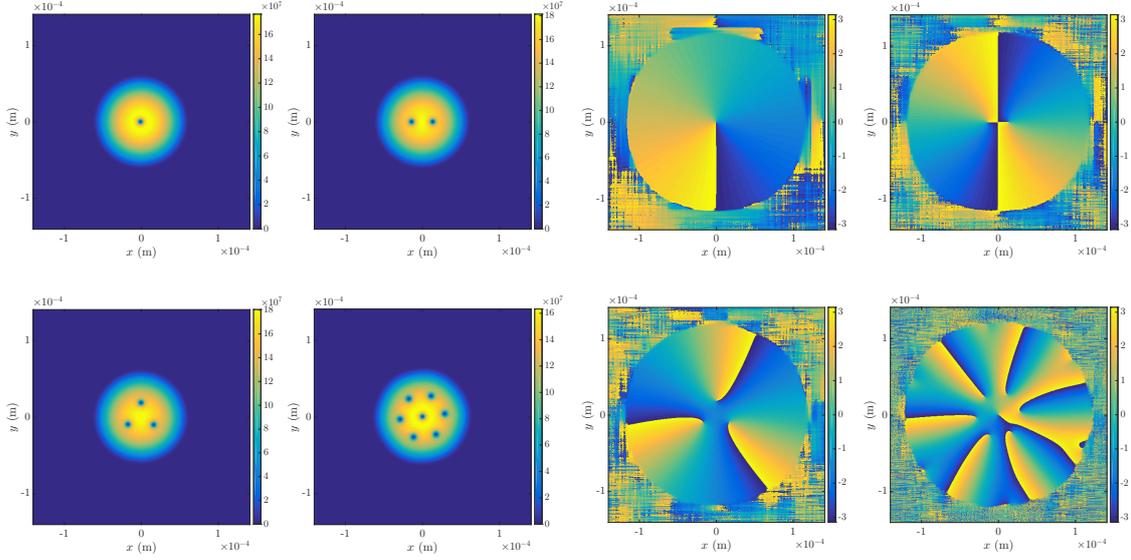


Figure 4.2: Density (left) and phase (right) for condensates carrying 1,2,3,and 7 vortices respectively.

altogether in the presence of dissipation.

This stability and dynamical behaviour can be understood by examining the velocity fields of the vortices, as shown in Fig. 4.3 for like-signed vortices, and Fig. 4.4 for a vortex-antivortex pair. The velocity fields of the vortices are additive, which for like-signed windings creates regions of zero flow, and for opposite-windings creates large flows between them respectively. From Eq. 2.48 the velocity field of the irrotational quantum vortex scales as $v \propto 1/r$, and hence becomes singular at the vortex centre, giving rise to the large velocities close to the singularity. Fig. 4.5 shows a clearer view of the velocity fields as depicted by Fig. 4.3, where the presence of maxima is easily observed at the cores, with field minima also observed due to the vortex-vortex interactions.

4.2 Rapidly rotating vortex lattice

4.2.1 Model system

For the work to follow we assume a standard single component Bose–Einstein condensate in a radially symmetric trap of frequency $\omega_{\perp} = 2\pi \times 1$ Hz. By tightly confining the condensate along the z -dimension, with trapping frequency $\omega_z = 2\pi \times 16$ Hz, such that ω_z is greater than ω_{\perp} , the condensate enters the desired pancake-shaped geometry. The system is then modelled using the mean-field GPE Hamiltonian as

$$H_{\text{GP}} = -\frac{\hbar^2}{2m}\nabla^2 + \frac{1}{2}m\omega_{\perp}^2\mathbf{r}^2 + g_{2\text{D}}|\Psi(\mathbf{r},t)|^2. \quad (4.1)$$

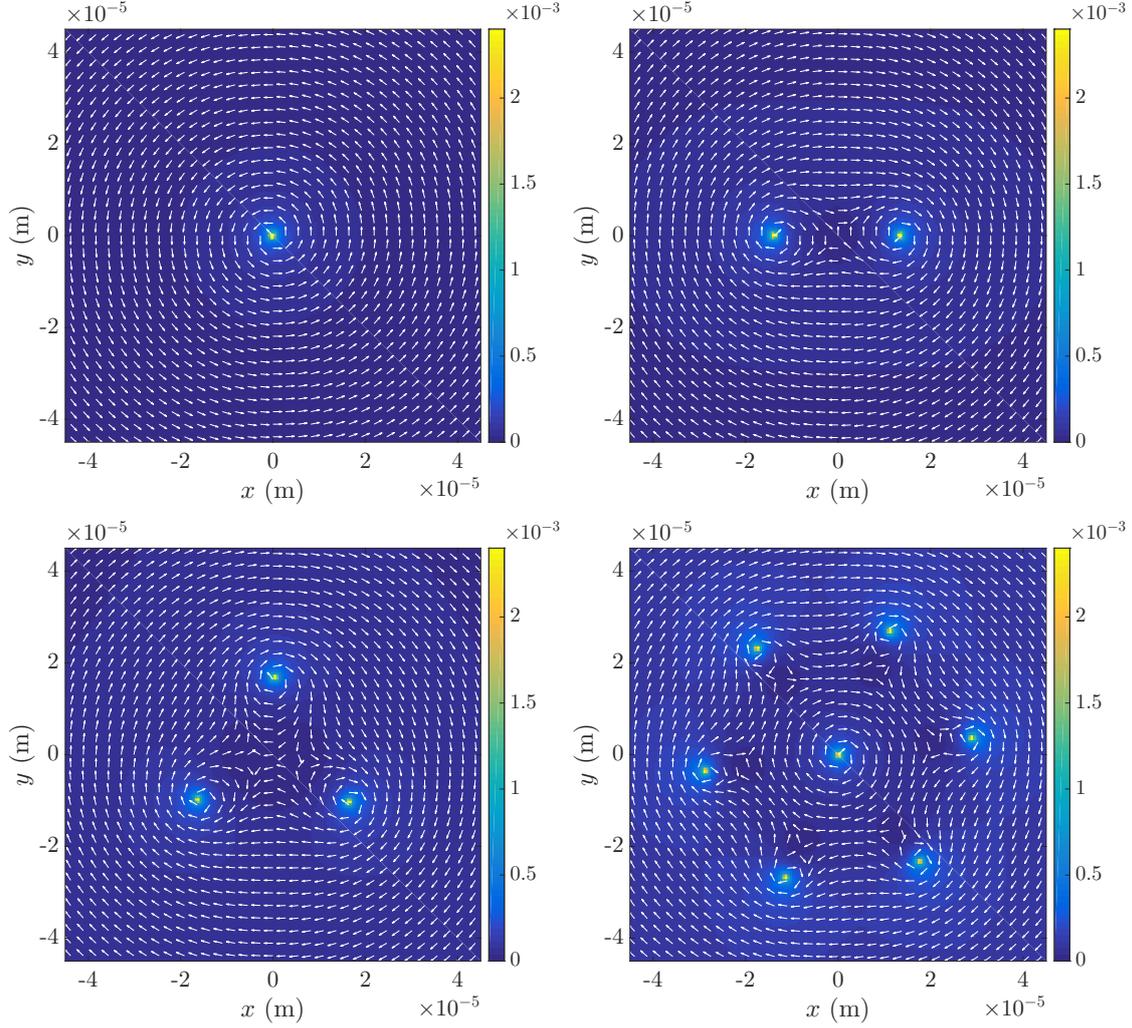


Figure 4.3: Magnitude of the velocity field and the direction of rotation of the field for clockwise circulating vortices. As the velocity follows a $1/r$ profile the magnitude becomes singular at the centre of the vortices, which is capped here at $v = 2.4 \times 10^{-3}$ (ms^{-1}).

Here, we define the two-dimensional effective interaction strength $g_{2\text{D}}$ given by Eq. (2.38), as

$$g_{2\text{D}} = g \sqrt{\frac{m\omega_z}{2\pi\hbar}} = 4g \sqrt{\frac{m}{\hbar}}. \quad (4.2)$$

To describe the system we enter the co-rotating frame, which is done by including the angular momentum operator L_z in the Hamiltonian. The time dependent dynamics of the system are then given by the GPE as

$$i\hbar\partial_t\Psi(\mathbf{r}, t) = [H_{\text{GP}} - \Omega L_z] \Psi(\mathbf{r}, t). \quad (4.3)$$

If the angular rotation frequency approaches the condensate trapping frequency, $\Omega \approx \omega_\perp$, the condensate gains a large triangular lattice of vortices. The effect of setting $\Omega = \omega_\perp$, can be

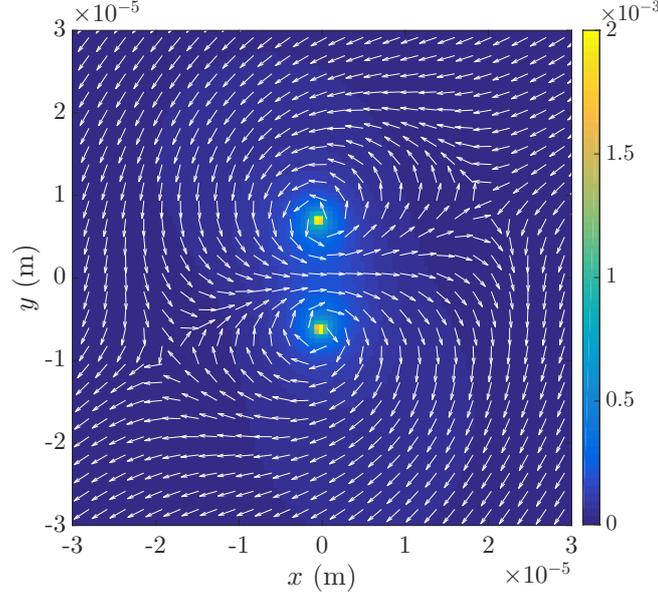


Figure 4.4: Velocity field direction and magnitude originating from vortex and antivortex present in a condensate. Colour scale is in units of (ms^{-1}).

partially understood in a mean-field setting by rewriting the GPE kinetic and rotation terms in the form

$$\frac{\mathbf{p}^2}{2m} + \frac{m\Omega\mathbf{r}^2}{2} - \Omega L_z = \frac{(\mathbf{p} - m\boldsymbol{\Omega} \times \mathbf{r})^2}{2m}. \quad (4.4)$$

The resulting changes to the GPE are then given by

$$i\hbar\partial_t\Psi = \left(\frac{1}{2m}(-i\hbar\nabla - m\boldsymbol{\Omega} \times \mathbf{r})^2 + \frac{m}{2}(\omega_{\perp}^2 - \Omega^2)r^2 + g_{2D}|\Psi|^2 \right) \Psi. \quad (4.5)$$

The above Hamiltonian demonstrates the correspondence between a rotating condensate and that of a non-relativistic charged particle in a magnetic field [110]. One can see that when the rotation and trapping frequencies are equal, the condensate no longer sees a confining potential, due to the centrifugal force experienced, given by the term $-m\Omega^2 r^2/2$.

As stated earlier in Sec. 2.2.3 as Ω approaches ω_{\perp} , the use of mean-field theory becomes less justified, and the system enters a strongly correlated regime. If the filling fraction ν is in the range of $10 \leq \nu \leq 1000$, the system enters the “mean-field quantum Hall” regime, where the Gross–Pitaevskii theory is still working well. The system examined here is well within this regime for a frequency of $\Omega = 0.995\omega_{\perp}$. This allows us to work with systems that have a very large vortex lattice, and are still described by mean-field theory [62].

In the following we will numerically solve Eq. (4.3), using the pseudospectral Fourier split operator method and making use of GPU computing as described in Sec. 3.6. For realistic experimental parameters we assume $N \approx 10^6$ atoms of ^{87}Rb , with an s -wave scattering length of $a_s = 4.76 \times 10^{-9}$ m [117]. The numerically evaluated ground-state for the given set of parameters is shown previously in Fig. 3.15 and has a radius of approximately 3.5×10^{-4} m.

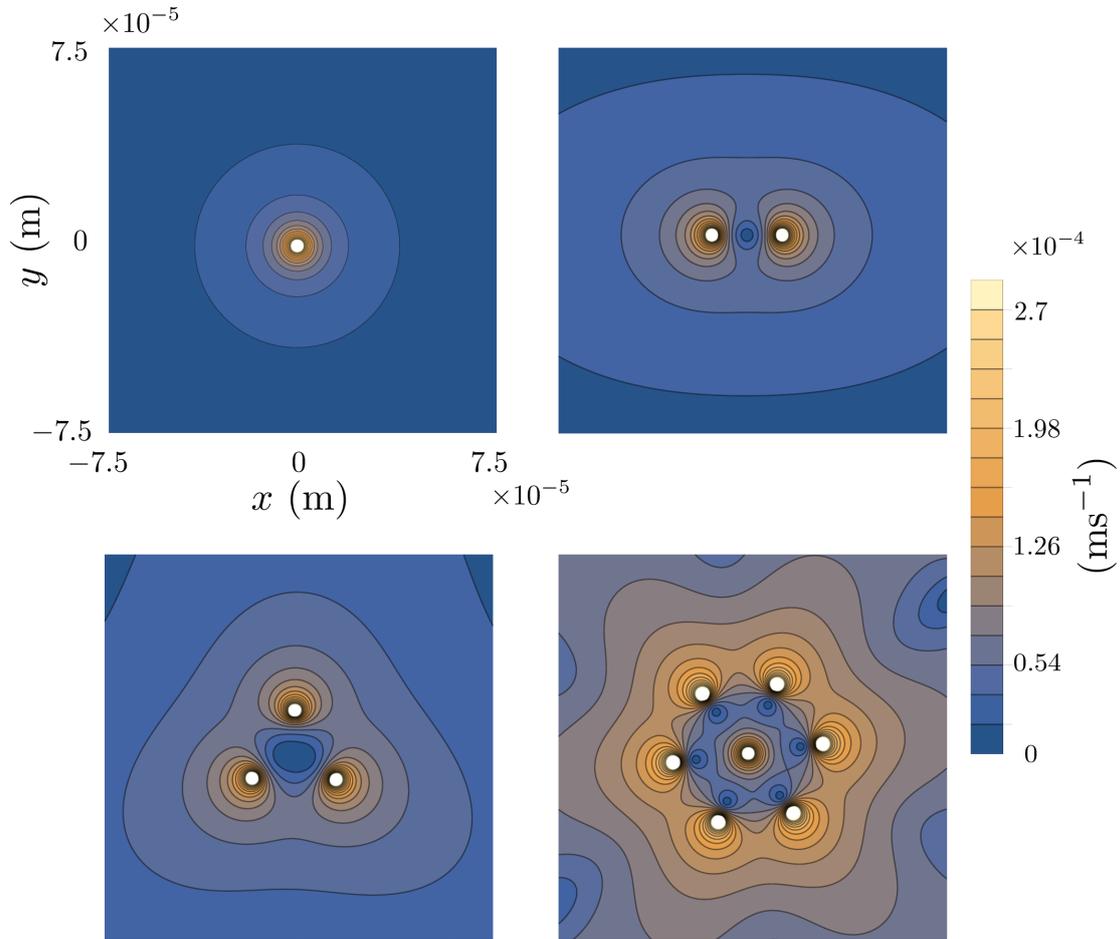


Figure 4.5: Magnitude of the velocity field for small numbers of like-signed vortices in a condensate. The velocity field can be seen to become singular at the centre, which here has its magnitude capped at a lower value than Figs. 4.3 and 4.4 to aid visibility. Regions of minimal velocity are seen, wherein the opposing field lines compensate one another.

For these parameters, the number of vortices within the visible density region is approximately 600, giving a filling factor of $\nu \approx 800$. This places the system within the mean-field quantum Hall regime, and therefore a description using Gross–Pitaevskii theory is adequate [66].

Following an imaginary time-evolution as outlined in Sec. 3.1, we first find the ground state of the condensate in a harmonic potential, starting with a condensate without vortices. We then linearly ramp the rotation frequency to avoid the lattice disordering issue discussed in Sec. 3.6.1. This allows us to follow the ground state solution at all times while the rotation frequency is increased, which has the added advantage of returning a ground state solution for any required rotation frequency. Examples of several states obtained during a single simulation are given in Fig. 4.6. The previously discussed resolution considerations become apparent as the rotation rate is increased for both the position and momentum space representations of the wavefunction. A movie of the wavefunction density is available at the

following URL [118], in which the frequency is ramped from $\Omega/\omega_{\perp} = 0.39 \rightarrow 0.995$.

The rapidly rotating vortex lattice is known to exhibit solid-body-like rotation [55]. This can be seen in Fig. 4.7, where the coarse-grained flow of the velocity field increases as a function of the distance from the lattice centre. This rigid-body behaviour allows us to treat the lattice as a solid object. Recalling the Feynman relation for vortex density Eq. (2.52), $n_v = m\Omega/(\pi\hbar)$, we choose an area over which the lattice spacing is almost constant and the vortex density closely matches this relation. In this rapidly rotating regime the vortices close to the centre will have an almost uniform profile [119]; for the above system parameters this is fulfilled by a hard-walled radial boundary of $r = 2 \times 10^{-4}$ m, in which the number of vortices are calculated to being $N_v \approx 342$. Employing the vortex detection and tracking methods described in Sec. 3.6.2, within the same region gives $N_v = 341$ vortices, with a lattice spacing of $a_v \approx 2.1 \times 10^{-5}$ m, and with a standard deviation of $\sigma \approx 2.7 \times 10^{-7}$ m. These values indicate that within this region the lattice is well ordered. While the condensate has more vortices outside this boundary that are initially ordered, during time evolution many of these move more easily due to the large velocity fields closer to the edges. The above boundary gives a well ordered lattice that remains well ordered for several seconds of time evolution.

If, however, the linear ramp is performed too quickly, or an initial state is chosen that already contains a large amount of angular momentum without being the eigenstate, the vortices tend to enter from the boundary all at once, and fail to converge to the well ordered ground state. A demonstration of this such issue is shown in Fig. 4.8, and indicates the need for a slow ramp of Ω that is essentially adiabatic in imaginary time. As higher rotation frequencies are reached, the rate at which vortices enter the condensate increases rapidly. While the rapid entry of vortices was a problem for reaching an ordered lattice without a ramp of rotation frequency, the presence of an existing lattice during ramping allows all newly entered vortices to order more easily. Therefore, a linear ramp is effective and does not require a more complex scaling.

In the following section we will discuss perturbations of the condensate in the presence and absence of vortices. For the above system, we will investigate the effect of perturbations to both the global and local condensate order due to changes in the lattice structure.

4.3 Condensate perturbations

4.3.1 Trapping potential control

With the model system outlined in Sec. 4.2.1, we will now imagine an abrupt change to the Hamiltonian, such that, $H(t) = H_{\text{GP}} + f(t)V_{\text{ext}}$, where V_{ext} is an external potential, and $f(t)$ is some function of time to control the application of V_{ext} . In this scenario the initial wavefunction, which is a stationary state of H_{GP} , will no longer remain so provided that

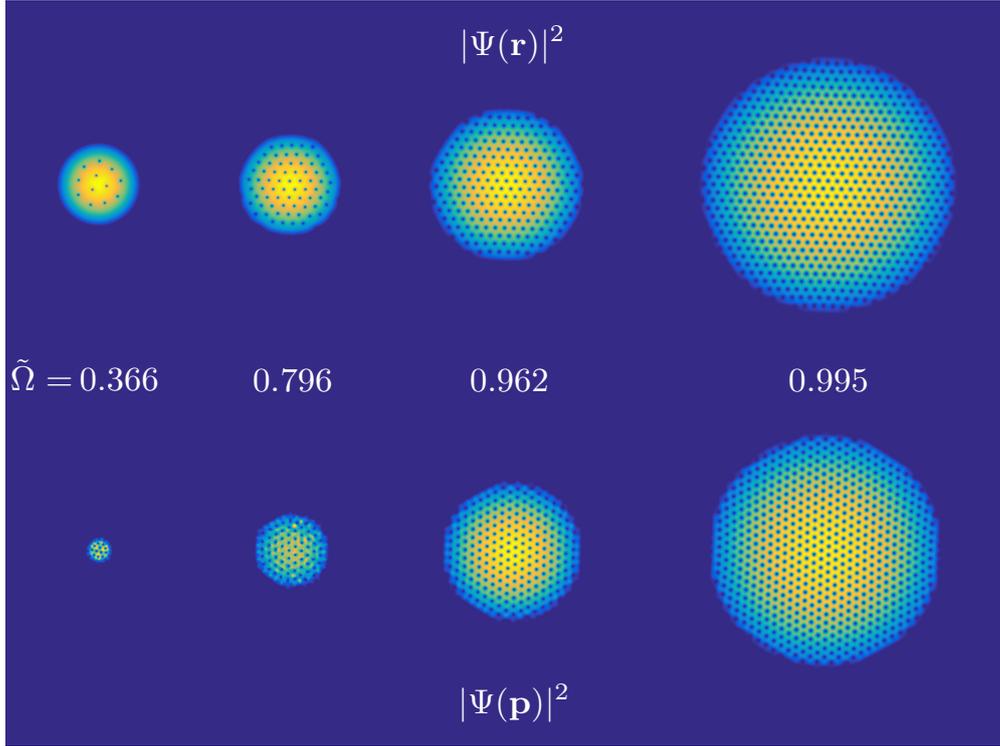


Figure 4.6: Density distributions of the wavefunction in position (top) and momentum (bottom) space for increasing rotation frequencies (from left to right). The color axis differs for each plot for visibility, as with a constant axis it is difficult to view densities across all magnitudes. The growth rate of the condensate radius in both position and momentum space becomes large when $\Omega_z \approx \omega_\perp$.

H_{GP} and $f(t)V_{\text{ext}}$ are non-commuting. Assuming the time of application of the additional term is much shorter than any other timescale of the condensate dynamics, any modification of the Hamiltonian in this way can be viewed as a method for changing the phase of the wavefunction. The resulting effect on the wavefunction can be given as

$$\Psi(t=0) = |\Psi(t=0)\rangle e^{i\theta_0} \quad (4.6a)$$

$$\begin{aligned} \Psi(t) &= \Psi(t=0) e^{-i\frac{V_{\text{ext}}\Delta t}{\hbar}} \\ &= |\Psi(t=0)\rangle e^{i\left(\theta_0 - \frac{V_{\text{ext}}\Delta t}{\hbar}\right)} \end{aligned} \quad (4.6b)$$

where we have made use of Eq. (2.39). After application of V_{ext} the wavefunction phase is given by

$$\theta' = \theta_0 - \frac{E_{\text{ext}}\Delta t}{\hbar}, \quad (4.7)$$

where E_{ext} is the perturbation energy. One commonly used method to manipulate the condensate is through the use of optical potentials, which offer a large degree of control over the respective system's Hamiltonian [120]. The electric field component of an arbitrary optical

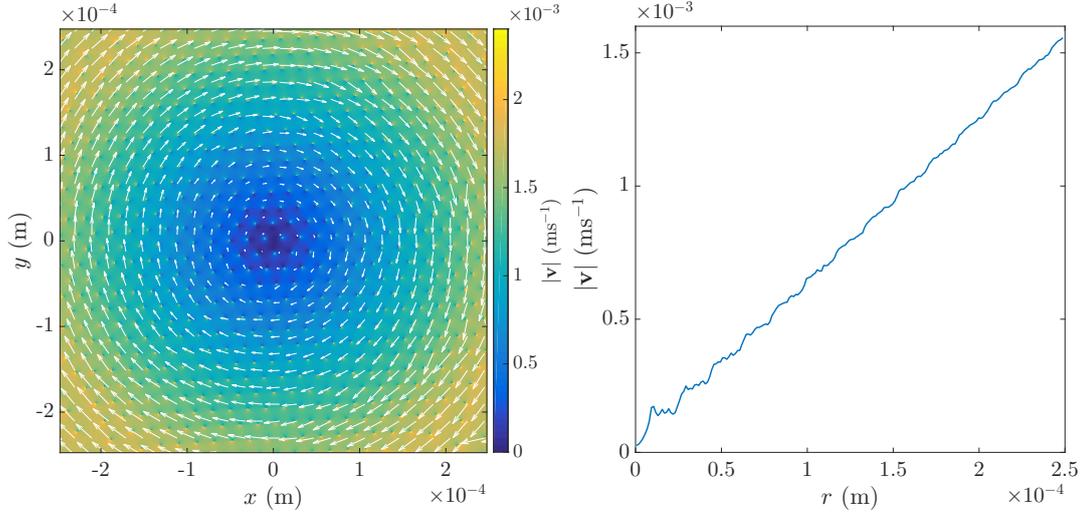


Figure 4.7: (Left) Magnitude and direction of flow of the velocity field for a rapidly rotating vortex lattice in the rapidly rotating frame. (Right) The angle-averaged velocity field from condensate centre. The solid-body rotation can be seen, where the coarse-grained condensate velocity scales as $v \propto r$ from the condensate centre. Close to the vortex cores the velocity field retains a $1/r_{\text{vtx}}$ profile, becoming singular at the core centres, with r_{vtx} as the distance from the respective cores.

field, described by a wavevector \mathbf{k} , and frequency, ω is given by

$$\mathbf{E}(\mathbf{r}, t) = \varepsilon_0 e^{i(\mathbf{k}\cdot\mathbf{r}-\omega t)} + \varepsilon_0^* e^{-i(\mathbf{k}\cdot\mathbf{r}-\omega t)}, \quad (4.8)$$

where ε_0 is the field amplitude. Using the dipole approximation, the interaction of an atom with a laser field is given by [7]

$$\mathcal{V} = -\mathbf{d} \cdot \mathbf{E}, \quad (4.9)$$

where \mathbf{d} is the electric dipole moment operator. For a two level atom with a ground state $|g\rangle$, and an excited state $|e\rangle$ with energy difference $\hbar\omega_0$, the dipole operator can be written as

$$\begin{aligned} \mathbf{d} &= \langle g|\mathbf{d}|e\rangle|g\rangle\langle e| + \langle e|\mathbf{d}|g\rangle|e\rangle\langle g| \\ &= \boldsymbol{\mu}_{eg}|g\rangle\langle e| + \boldsymbol{\mu}_{eg}^*|e\rangle\langle g|, \end{aligned} \quad (4.10)$$

where we have made use of $\langle g|\mathbf{d}|g\rangle = \langle e|\mathbf{d}|e\rangle = 0$, since the atoms have no permanent dipole moment. From Eq. (4.10) the Hamiltonian of the two-level system can then be written as

$$\begin{aligned} H &= H_0 + \mathcal{V} \\ &= \hbar\omega_0|e\rangle\langle e| - (\boldsymbol{\mu}_{eg}|g\rangle\langle e| + \boldsymbol{\mu}_{eg}^*|e\rangle\langle g|) \cdot (\varepsilon_0 e^{i(\mathbf{k}\cdot\mathbf{r}-\omega t)} + \varepsilon_0^* e^{-i(\mathbf{k}\cdot\mathbf{r}-\omega t)}). \end{aligned} \quad (4.11)$$

After expanding (4.11) we can then rewrite $\boldsymbol{\mu}_{eg} \cdot \varepsilon_0 = \hbar\Omega_r$, where Ω_r is the Rabi-oscillation frequency between the states. Assuming that the detuning $\Delta = \omega - \omega_0$ between the laser field

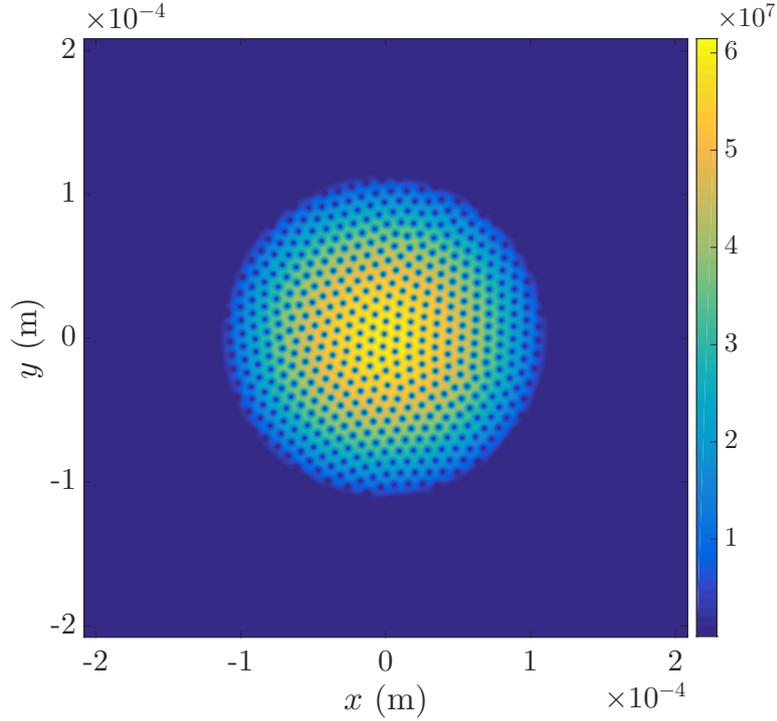


Figure 4.8: Disordered lattice resulting from starting in imaginary time evolution at the required rotation rate (here $\Omega = 0.99\omega_{\perp}$). If the rotation frequency is chosen too large without allowing the lattice to form and order, the resulting vortices all enter instantaneously and compete for their final positions. The system only converges one timescales that exceed reasonable computing times.

and transition frequency is small, $\Delta \ll \omega + \omega_0$, allows use of the rotating wave approximation. For this we first perform a unitary transformation of the system into the interaction picture rotating with H_0 using the operator $U = e^{-i\omega_0 t|e\rangle\langle e|}$ as

$$H_{\text{int}} = U^{\dagger} \mathcal{V} U. \quad (4.12)$$

All resulting terms featuring $\omega + \omega_0$ can be considered to be rapidly oscillating, and average out to zero. Following this approximation, the Hamiltonian can be transformed back into the Schrödinger picture, giving

$$H' = \hbar\omega_0 - \hbar\Omega (e^{i\omega t}|e\rangle\langle g| + e^{-i\omega t}|g\rangle\langle e|). \quad (4.13)$$

The final shift of the energies is then given as

$$\mathcal{V} = -\frac{\alpha}{2} \langle \mathbf{E}^2 \rangle_t, \quad (4.14)$$

where α is the real component of the atomic polarisability, and $\langle \cdot \rangle_t$ denotes the time average. If the electric field is spatially inhomogeneous, then this leads to a force of the form

$$\mathbf{F}_d = \frac{\alpha}{2} \nabla \langle \mathbf{E}^2 \rangle_t, \quad (4.15)$$

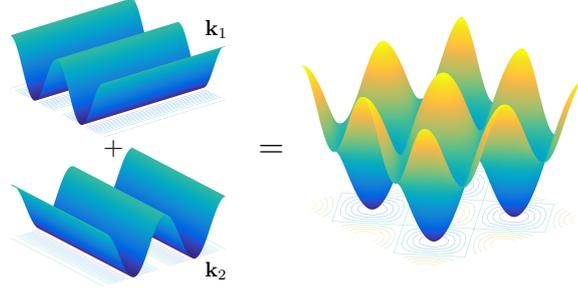


Figure 4.9: Square lattice generation using two orthogonal propagating laser fields with wavevectors \mathbf{k}_1 and \mathbf{k}_2 , as defined by Eq. (4.17).

which is known as the dipole force. Assuming counter propagating plane waves, we can then model a standing wave solution of the resulting optical potential as

$$V_{\text{ext}} \approx -\frac{\Omega_0^2(\mathbf{r})}{4\Delta} = V_0 \cos^2(\mathbf{k} \cdot \mathbf{r}), \quad (4.16)$$

where $V_0 = -\Omega_0^2/4\Delta$ is the field intensity, and $\Omega_0 \propto |\varepsilon_0|^2$ is the Rabi-frequency of the standing wave. The optical potential forms a highly periodic system given an appropriately chosen \mathbf{k} , and is known as an *optical lattice*. Optical lattices have become very common in BEC experiments as they allow for control of the kinetic energy term of the atoms to a very high degree [121–126].

Different geometric potentials can be formed with optical lattices by using laser fields with different \mathbf{k} vectors. Assuming standing waves with different polarisation axes or slightly different wavelengths, the interference effects between two or more fields can be ignored, with the resulting optical field given by the summation of lattice potentials with wavevectors $\mathbf{k}_{1,\dots,n}$. Creating a 2D lattice with n -fold rotational symmetry requires $n/2$ \mathbf{k} -vectors separated by $2\pi/n$, and with $n/2 \in \mathbb{Z}^+$. Taking a square lattice as an example, which has a 4-fold rotational symmetry, it can be created by two \mathbf{k} -vectors, separated by $\pi/2$, as

$$\mathbf{k}_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad \mathbf{k}_1 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \quad (4.17)$$

The resulting potential is shown in Fig. 4.9. The time the optical lattice is applied to the condensate can be controlled by choosing the function $f(t)$. Applying a lattice for a finite, but short, time only will lead to a modification of the wavefunction phase, which then subsequently, and on a much longer time-scale, will have an effect on the density distribution. Of particular interest to us is the use of an optical potential that is pulsed one or several times, which can be described by $f(t)$ as a periodic delta function. The condensate phase profile is the only quantity immediately modified, and any change in the density distribution appears only in the following evolution.

4.3.2 Direct phase manipulation

While ground state condensates will have a flat phase across the system, there are two interesting examples where a spatially dependent phase exists: dark solitons [127] and vortices [42]. We have previously discussed the 2π phase profile of a vortex that leads to the singularities in the wavefunction. In contrast, dark solitons feature a π phase jump profile. These excitations are unstable in dimensions higher than one, and will decay via the snake instability to paired vortices and antivortices [128–130]. Where so far we have only considered the short-term evolution of the wavefunction after being kicked by an optical potential, we will in the following consider what structures can be created in the condensate by careful phase manipulation techniques, besides dark solitons and vortices.

For this we will assume that the BEC allows for a short enough application of potentials so that only the phase is affected, and discuss direct manipulation of the wavefunction, as opposed to modification of the Hamiltonian. Following [10] and taking Eq. (2.39), the phase of the condensate can then be written as

$$\theta = \theta_0 + \theta_i, \quad (4.18)$$

where θ_0 is the unperturbed condensate phase, and θ_i is the phase pattern to be imprinted. Upon solving for the initial condensate ground state with deterministic phase, an additional phase pattern can be imprinted at any time by simply multiplying the wavefunction by $e^{i\theta_i}$. However, without careful choice of the phase terms their addition can lead to unwanted dynamics, so care must be taken to choose a well defined initial and imprinted phase pattern.

The advantage of the phase imprinting model is that for topological defects, one can imprint the required winding instantaneously, allowing them to appear at predefined positions. The density also needs to only adjust itself locally to the phase singularity, with the remaining condensate seeing an almost constant shift in phase. The creation of vortices through application of localised $\pm 2\pi$ phase winding defects in the condensate therefore allows for direct control of the angular momentum and vorticity within the BEC. While discussed in the literature for the creation of vortices, it is worth noting that the phase imprinting method can also be used to annihilate a vortex from the condensate by applying a phase profile of opposite winding, removing the singularity. This will leave the condensate with a density depletion at the prior location of the phase singularity. Without the phase singularity this depletion will fill in and excite phonon modes in the condensate during time evolution. This process will form the basis for further discussions and analysis of vortex carrying condensates.

Experimental realisation of arbitrary potential patterns to achieve the required phase is accessible through the use of spatial light modulators (SLM) [131]. These devices behave as digital displays, through which visual patterns can be expressed in a time dependent manner, allowing the application of a laser field in the required form. We will assume for all future

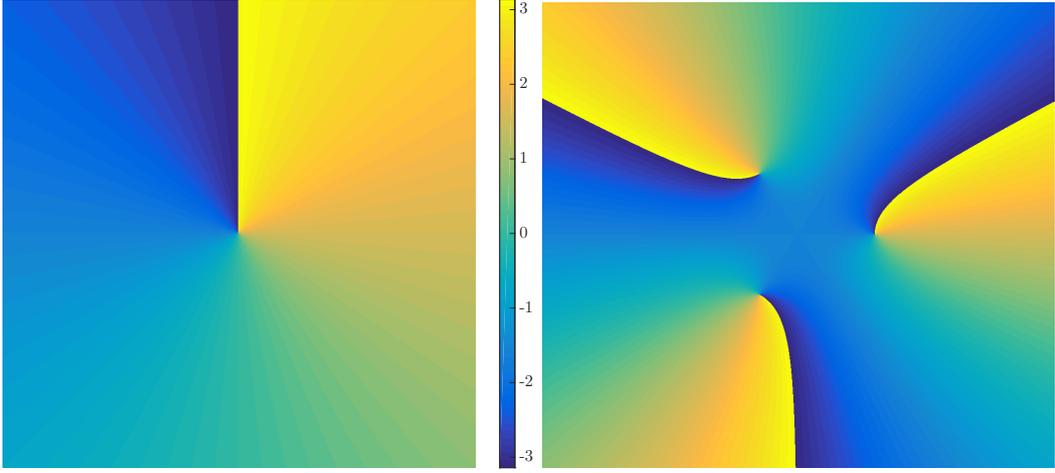


Figure 4.10: 2π phase winding shown for a single (left) and three separated (right) phase singularities. The application of separate phase singularities can be treated as summing each individual phase profile, $\left(\sum_i \theta_i\right) \bmod 2\pi$. Note that the resolution used in the numerical examples is much higher than can be obtained from currently available SLMs.

discussions that the potentials we require are experimentally realisable with sufficient resolution, and focus on the resulting effect on the condensate. For the creation of a single vortex the 2π phase winding pattern can be created spatially using the two-argument four-quadrant form of arctan as given by

$$\theta_i(\mathbf{x}, \mathbf{y}; x_0, y_0) = \arctan(\mathbf{y} - y_0, \mathbf{x} - x_0), \quad (4.19)$$

which locates the singularity at the position (x_0, y_0) . The resulting phase is shown in Fig. 4.10(left), and including the additional phase singularity term θ_i , the condensate wavefunction following an imprint is given as

$$\Psi'(\mathbf{r}, t) = |\Psi(\mathbf{r}, t)| e^{i(\theta_0(\mathbf{r}, t) + \theta_i(\mathbf{r}))}. \quad (4.20)$$

Following the imprint this process will create phonons in the condensate density that will radiate outwards from the singularity imprint. As imprinting is directly controlling the condensate phase, it can also be considered a direct manipulation of the kinetic energy since the superfluid velocity depends on the phase gradient (see Eq. (2.44)). This, in reverse, means that by applying spatially inhomogeneous phase profiles the atomic velocity can be adjusted to different values in different regions of the condensate. To demonstrate this we consider a simple example of a Gaussian phase profile applied to the condensate. The imprinted profile has the form

$$\theta_i(\mathbf{r}) = A \exp\left(-\frac{|\mathbf{r} - \mathbf{r}_0|^2}{2\sigma^2}\right) \bmod 2\pi, \quad (4.21)$$

where A is the phase profile amplitude, \mathbf{r}_0 is the centre of the Gaussian curve, and σ is adjusted to match the condensate width. The modulo 2π ensures that the phase wraps around for

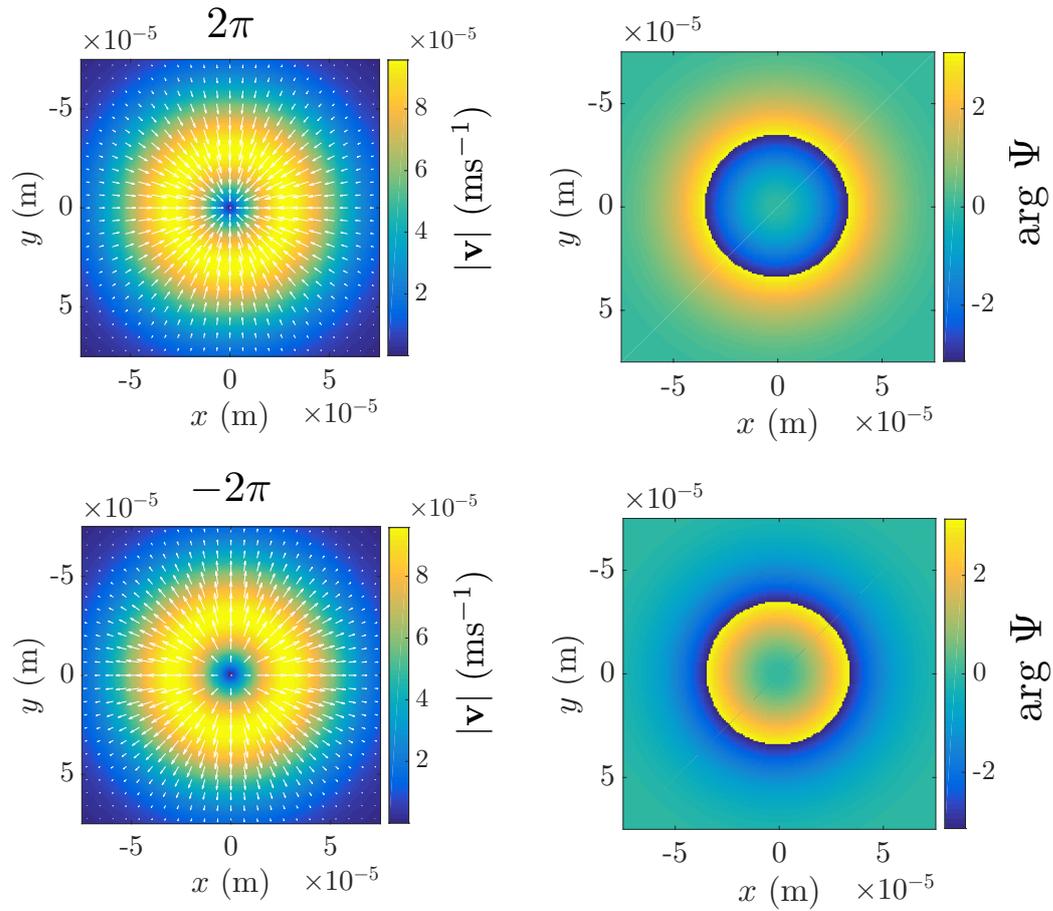


Figure 4.11: Velocity fields and magnitude (left) and phase (right) for a condensate directly following a Gaussian phase imprint. The lengths of the arrows give the magnitude of the respective velocity components, with the color map indicating this also for clarity. The sign of the imprint changes the direction of the respective kinetic components, with a positive imprint initially creating a density contraction (top), and negative creating an expansion (bottom), with the arrows indicating the direction of the flow.

amplitudes exceeding the $(0, 2\pi)$ range. The Gaussian profile has large radial gradients in two-dimensions, so that an imprint on the condensate should lead to radial velocities and therefore an expansion or contraction of the cloud (see Fig. 4.11). This can be expected to lead to interference fringes, as the faster moving atoms have the possibility to overtake the slower ones for sufficiently large amplitudes [132]. In Fig. 4.12, where a slice through the condensate centre is given for both position and momentum space, this can be clearly seen.

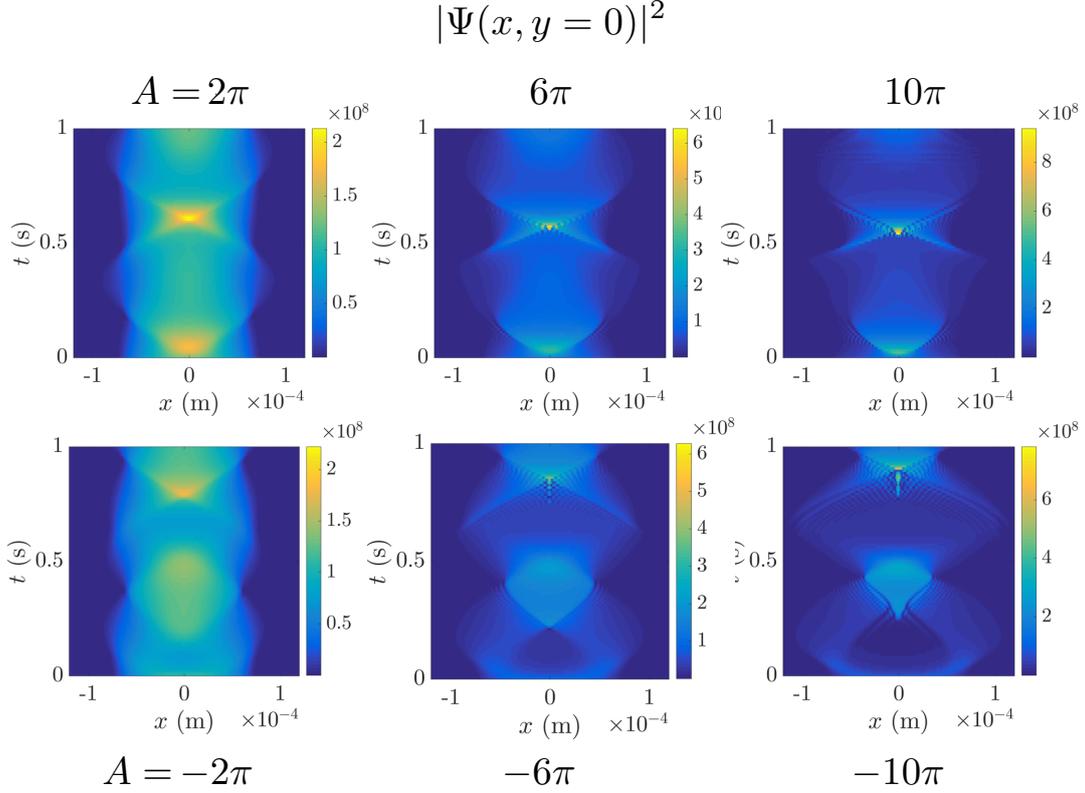


Figure 4.12: A cut through the condensate wavefunction density following the phase imprinting of a Gaussian with amplitudes $A = \pm(2\pi, 6\pi, 10\pi)$. The positive imprints create an initial contraction of the cloud (top), while the negative imprints lead to expansion (bottom). For the larger kicking strengths interference fringes can be observed during expansion and contractions.

4.4 Condensate analysis

To analyse the effects of the phase imprinting we will below introduce the decomposition of the kinetic energy to isolate the effect from phonons and vortices. Following this, for the vortex lattice we will introduce two closely linked methods to examine geometric structure — Delaunay triangulation and Voronoi tessellation. These methods are dual to one another, and can be used to easily identify order, structure and local parameters within systems of many particles.

4.4.1 Kinetic energy decomposition

Given that the phase engineering modifies the condensate kinetic energy profile, it is instructive to quantify this effect. One can apply a spectral decomposition of the kinetic energy of the condensate into contributions solely from the vortices (incompressible), and those from

the phonons (compressible) [133–135]. For this, the wavefunction is again written in terms of amplitude $\sqrt{\rho(\mathbf{r}, t)}$ and phase $S(\mathbf{r}, t)$, which allows the kinetic part of the Gross–Pitaevskii energy functional Eq. (2.50) to be calculated as

$$E_{\text{kqp}} = \int d\mathbf{r} \left(\frac{\hbar^2}{2m} |\nabla \sqrt{\rho(\mathbf{r}, t)}|^2 + \frac{m}{2} |\sqrt{\rho(\mathbf{r}, t)} \mathbf{v}(\mathbf{r}, t)|^2 \right). \quad (4.22)$$

One can then decompose this into the quantum pressure (first) and kinetic energy (second) terms. The kinetic energy term can be seen as a density-weighted velocity field, $\mathbf{u}(\mathbf{r}, t) = \sqrt{\rho(\mathbf{r}, t)} \mathbf{v}(\mathbf{r}, t)$, and it can be further decomposed into the sum of compressible and incompressible terms,

$$\mathbf{u}(\mathbf{r}, t) = \mathbf{u}^c(\mathbf{r}, t) + \mathbf{u}^i(\mathbf{r}, t). \quad (4.23)$$

The two terms can be calculated by performing a Helmholtz decomposition of the field \mathbf{u} , which separates terms that are longitudinal (\mathbf{u}^c) and transversal (\mathbf{u}^i) with

$$\nabla \times \mathbf{u}^c(\mathbf{r}, t) = 0, \quad (4.24a)$$

$$\nabla \cdot \mathbf{u}^i(\mathbf{r}, t) = 0. \quad (4.24b)$$

By introducing the vector potential, \mathbf{A} , and the scalar potential, B , such that

$$\mathbf{u}^c = \nabla B, \quad (4.25a)$$

$$\mathbf{u}^i = \nabla \times \mathbf{A}, \quad (4.25b)$$

we can rewrite Eq. (4.23) as

$$\nabla \times \mathbf{u}(\mathbf{r}, t) = -\nabla^2 \mathbf{A}, \quad (4.26)$$

$$\nabla \cdot \mathbf{u}(\mathbf{r}, t) = \nabla^2 B. \quad (4.27)$$

To solve the above equation we begin by seeking a solution for B by performing a spectral decomposition of the full density-weighted velocity field as

$$B = \sum_j \frac{k_j}{|\mathbf{k}|^2} \mathcal{F}[\mathbf{u}], \quad (4.28)$$

where k_j is the j -th component in \mathbf{k} space, and \mathcal{F} is the Fourier transform. The resulting solution for \mathbf{u}^c is then given by

$$\mathcal{F}[\mathbf{u}_i^c] = \sum_j \frac{k_i k_j}{|\mathbf{k}|^2} \mathcal{F}[\mathbf{u}], \quad (4.29)$$

which after taking note of Eq. (4.23) gives

$$\begin{aligned} \mathcal{F}[\mathbf{u}_i^i] &= \mathcal{F}[\mathbf{u}_i] - \mathcal{F}[\mathbf{u}_i^c], \\ &= \sum_j \left(\delta_{i,j} - \frac{k_i k_j}{|\mathbf{k}|^2} \right) \mathcal{F}[\mathbf{u}_i]. \end{aligned} \quad (4.30)$$

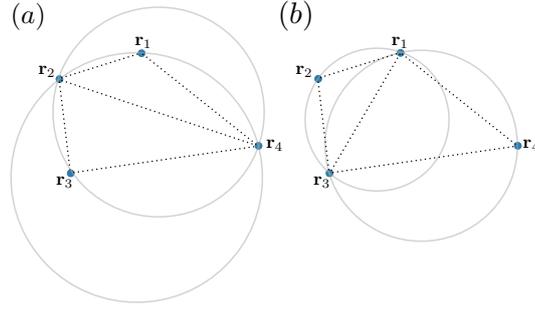


Figure 4.13: Non-Delaunay (a) and Delaunay (b) triangulation of 4 Euclidian points.

This decomposition separates the energy contribution from phonons and vortex cores, represented by compressible and incompressible terms respectively [136]. By averaging over binned shells in \mathbf{k} -space, the kinetic energy spectra, $E^{c,i}(k)$, are calculated as [135]

$$E^{c,i}(k) = \frac{mk}{2} \sum_{j \in \mathbf{r}} \int_0^{2\pi} d\phi_k \frac{|\mathcal{U}_j^{c,i}(\mathbf{k}, t)|^2}{s_k}, \quad (4.31)$$

where

$$\mathcal{U}_j^{c,i}(\mathbf{k}, t) = \int d^2\mathbf{r} e^{-i(\mathbf{k} \cdot \mathbf{r})} u_j^{c,i}(\mathbf{r}, t). \quad (4.32)$$

The terms $u_j^{c,i}(\mathbf{r}, t)$ represent the position-space density-weighted velocity components in the specified shell, where ϕ_k is the polar angle, and s_k is the number of values in the chosen shell.

4.4.2 Delaunay triangulation and Voronoi tessellation

A common method for examining the ordering and periodicity of large-scale crystalline structures is to generate a mesh with each vertex being the location of a particle. With this, one can easily observe ordered and disordered regions in a material; well defined straight lines indicate a perfect crystal, with any bends indicating the presence of imperfections. Some of the most widely used methods for this are the dual techniques from computational geometry of Delaunay triangulation and Voronoi tessellation.

The Delaunay triangulation of an arbitrary set of points in Euclidian space, \mathbf{R} , which we will denote as $D(\mathbf{R})$, is constructed in the following way:

1. No point will fall within the interior of any circumcircle of 3 points where $\mathbf{r}_{1..3} \subset \mathbf{R}$
2. The Delaunay triangulation will maximise the minimum angle between points.
3. If four points are on the same circumcircle, then both possible configurations give a Delaunay triangulation.

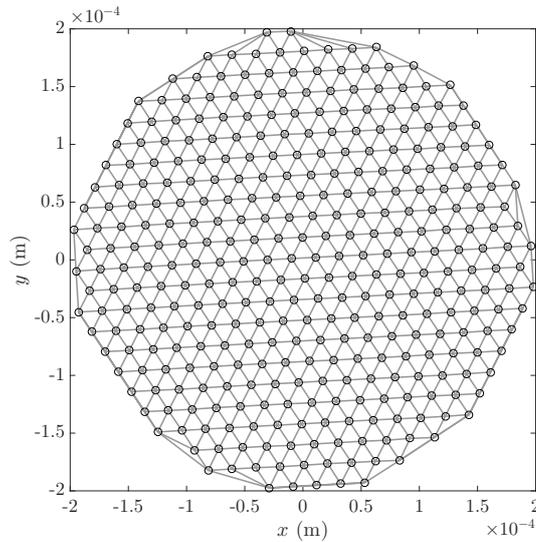


Figure 4.14: Delaunay triangulation of the vortex lattice ground state. The vertices away from the condensate boundary have the expected 6-edge structure.

This concept is more easily explained visually. Fig. 4.13 shows two different triangulations of four points; situation (a) is a non-Delaunay triangulation, as the points \mathbf{r}_1 and \mathbf{r}_3 each fall within the circumcircle of the other points. However, by simply flipping the central edge from $(\mathbf{r}_2, \mathbf{r}_4)$ to $(\mathbf{r}_1, \mathbf{r}_3)$ we can see in (b) that we now have a valid Delaunay triangulation. No point falls within the circumcircle of the other points, and the minimum angle formed is maximised relative to configuration (a). Following directly from this, one can see that Delaunay triangulation can be used to connect the closest vertices in a network. A nice side-effect of Delaunay triangulation is that one can examine when the number of edges from a vertex deviates from the expected value in the lattice, which is 6 for triangular lattices. This can be a useful means to locate defects in a crystal lattice, and we will make use of this during later discussions. This is performed using the built-in MATLAB function “delaunayTriangulation”, and counting the number of attachments to each individual vertex. A triangulation of the vortex lattice from Fig. 3.15 within the previously discussed radial boundary of $r = 2 \times 10^{-4}$ m is shown in Fig. 4.14.

An alternative representation, using the dual of the Delaunay triangulation, is that of the Voronoi tessellation (diagram). The characteristic of these diagrams is that they are composed of cells each encompassing an individual vertex, within which all enclosed points are closer to that particular vertex than any other. This representation can be generated from the Delaunay triangulation and vice-versa. Taking the centres of the circumcircles describing the Delaunay triangulations, and connecting these forms the boundaries of the Voronoi cells. A simple generation method can be seen as creating and expanding the radius of circles (or n -spheres in n -dimensions) centred on each vertex. Where the circles intersect with one another defines the boundary of each individual cell. An example of a Voronoi diagram compared with a Delaunay triangulation is given by Fig. 4.15. The area of each cell can be used as a

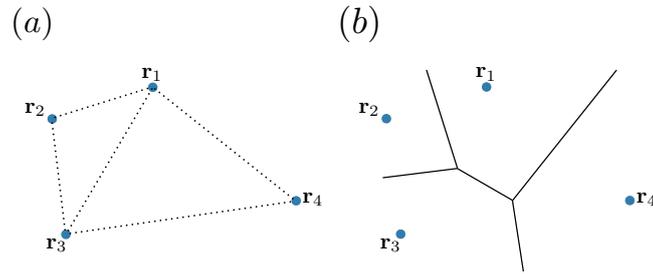


Figure 4.15: Comparison of Delaunay triangulation (a) with a Voronoi diagram (b). These graphs are duals, which means that one can be used to generate the other.

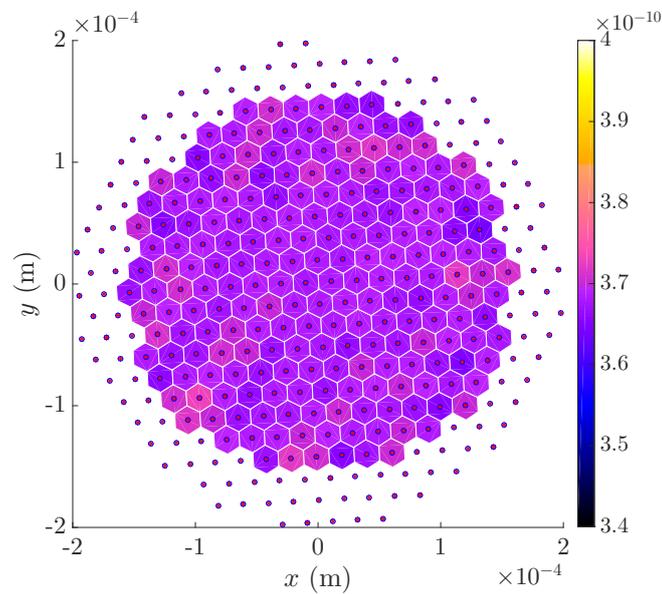


Figure 4.16: Voronoi tessellation of the vortex lattice. The area of each cell is represented by the color mapping. To avoid the tessellation tending to infinity, a buffer region of vortices is created close to the boundary.

metric of the strength of the interaction between particles in a many-body system, but also we may represent quantities local to each region in a system by the colour-scale of each cell. For the vortex lattice as given by Fig. 3.15, a sample Voronoi diagram is given in Fig. 4.16 with the color representing the area spanned by each Voronoi cell of the lattice.

Chapter 5

Moiré superlattice structures

5.1 Introduction

As mentioned in Sec. 4.3.1, the use of optical lattice potentials is ubiquitous in the field of ultracold atoms. Given their near perfect periodic structure, it is interesting to examine the effect these lattices have on already periodic systems. For a rapidly rotating condensate, the density profile will feature a large periodic Abrikosov lattice having 6-fold rotational symmetry i.e. a triangular lattice structure. The use of optical lattices stationary in the co-rotating frame on vortex lattices has previously been shown to allow for pinning of the vortex cores to the lattice maxima [122, 124]. Assuming blue detuned laser fields, the lattice maxima will drive atoms away, and make it favourable for vortex cores to sit in these positions. A transition between lattice geometries can be observed in these cases. This is, however, experimentally challenging as in the lab frame the optical lattice must rotate with the vortex lattice and match its rotation rate. Nevertheless, the creation of a rotating quasi two-dimensional optical lattice was achieved in [124], where the authors used a rotating mask with the holes placed in such a way as to allow for the required optical diffraction patterns. However, the system suffered from difficulties arising from the inclusion of a mechanically rotating stage, as well as the heating effects due to the presence of the always-on optical lattice. As the vortex lattice has 6-fold rotational symmetry, and assuming the chosen rotation rate of $\Omega = 0.995\omega_{\perp}$, after rotating for $t \approx 1/6$ s the system has returned to its initial orientation. It is therefore intriguing to see if the need for a rotating optical potential can be replaced by a periodically flashed optical lattice. Here, we will investigate the effect of a single application of an optical lattice to the vortex lattice.

For short application times, the potential will kick the wavefunction, and can be seen as leaving a phase imprint with minimal immediate density change. As discussed in Sec. 4.3.1, the effect of the kick can then be observed in the ensuing dynamics. Interestingly, one can observe the appearance of moiré interference patterns in the condensate density, while the positions of the vortex cores are only minimally affected. We will begin by discussing the

model, and present the explanation of the resulting interference pattern. We will show that it stems from the interference between the different \mathbf{k} -vectors, with signatures of the effect being visible in the condensate's kinetic energy spectrum. Finally, we will discuss the uses of this effect as a microscope for observing the presence of vortex lattices without time-of-flight.

As already discussed in detail above, given the large number of vortices in the condensate density, as well as the lowered density resulting from the large centrifugal forces, the vortex cores become large and form a highly periodic triangular lattice. The non-orthogonal lattice vectors are given in position space by $\mathbf{a}_1 = a_v[1, 0]$ and $\mathbf{a}_2 = a_v[-1/2, \sqrt{3}/2]$, where a_v is the distance between vortex cores. The lattice “constant” is truly only constant in areas of near uniform density. Thus, due to the inhomogeneous wavefunction density, the vortices near the edges are separated by a slightly larger distance than those at the centre. For the analysis we have therefore chosen to ignore all vortices near the condensate edge. The momentum space (\mathbf{k} -space) lattice vectors are reciprocal to the given \mathbf{r} -space vectors, and given by $\mathbf{b}_1 = 4\pi/(\sqrt{3}a_v)[\sqrt{3}/2, 1/2]$ and $\mathbf{b}_2 = 4\pi/(\sqrt{3}a_v)[0, 1]$.

5.2 Model

Using an optical potential one can create a perturbation in the condensate density by switching it on for a brief period of time. Such a kick must occur only for a time much shorter than both the rotation period of the vortex lattice, as well as the inter-vortex spacing divided by the speed of sound in the condensate, so that the effect is limited to a phase imprinting and not to a direct change in the density distribution. For the work carried out herein, the duration of the kick is $\tau_{\text{kick}} = 10^{-5}$ s, which is fast compared to the classical rotation period of the 6-fold symmetric lattice of $T_6 \approx 1.66 \times 10^{-1}$ s. This ensures that the vortex lattice is essentially stationary during the kick. To examine the effect that periodicity plays on the system, the optical potential was chosen to match the geometry of the Abrikosov vortex lattice, with the angle of relative alignment θ_Δ given as a free parameter. The optical potential was modelled as the sum of counter-propagating laser beams, as

$$V_{\text{opt}} = V_0 \sum_j \cos^2 [\mathbf{k}_j \cdot \mathbf{r}], \quad (5.1)$$

where V_0 is the optical lattice potential amplitude, and $j = [0, 1, 2, \dots]$ is the index of each respective laser with a differing \mathbf{k} -space wave-vector. By careful choice of the wave-vectors \mathbf{k}_j the density modulation given by the vortices is matched to the optical potential. This is achieved using the reciprocal lattice vectors defined by the vortex lattice, $\mathbf{b}_{1,2}$, and an additional third wave-vector, $\mathbf{k}_3 = 4\pi/(\sqrt{3}a_o)[\sqrt{3}/2, -1/2]$. For the resulting optical potential, we define the lattice sites as the maxima of intensity. These are in turn matched with the condensate density minima of the vortex lattice.

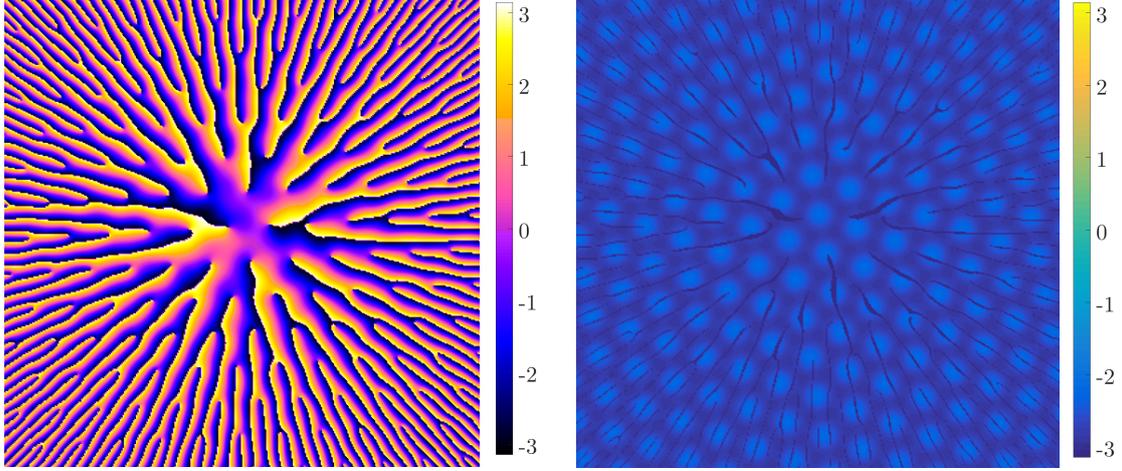


Figure 5.1: The condensate phase for the vortex lattice following a kick (left) and without the background lattice phase at $t = 0$ (right). The phase modulation through the optical lattice kick is clearly visible.

As the vortex lattice constant in a rapidly rotating atomic BEC can be quite large depending upon the system parameters, optical lattices with wavelengths on the order of tens of microns can be necessary [137, 138]. By ensuring a short kick with an amplitude on the order of $10^{-2}\mu$, with μ as the BEC's chemical potential, the result of the kick is an imprinted phase on the wavefunction [42]. As an example, following a single kick the vortex lattice phase can be observed to be modified, with localised phase gradients at the optical lattice maxima, as seen in Fig. 5.1. These localised gradients lead to the development of a flow originating from each optical lattice potential maxima location, and in turn create well-defined interference patterns resulting from phonons created by the kick. In the presence of a vortex lattice with a periodic core arrangement, this creates moiré superlattice structures [139] in the density of the condensate. Such patterns have been observed in many solid-state systems, for example in graphene on hexagonal boron nitride [140], though it is static in such cases. In our system, the structures are dynamical, and revive at well defined intervals, where the frequency of the oscillations can be related to the speed of sound

$$c(\mathbf{r}, t) = \sqrt{\frac{\rho(\mathbf{r}, t)g}{m}}, \quad (5.2)$$

divided by the lattice constant.

For state-of-the-art theoretical analysis of condensates, one of the most widely used methods to understand the effects of the kinetic energy is to perform a spectral decomposition via the method explained in Sec. 4.4.1. This separates the contributions from vortices and phonons in the kinetic energy, and is analogous to methods used in classical fluid systems. For vortices, in both liquid helium and atomic Bose–Einstein condensates, this decomposition also allows one to characterise the onset of quantum turbulence, and is used extensively in the literature [135, 141–144].

However, Reeves *et al.* [145] have recently discussed how the kinetic energy spectral decomposition as presented in Sec. 4.4.1 does not truly represent the kinetic energy of the system, as the spectral terms are not additive in \mathbf{k} -space. The classical interpretation is described as being “*obtained by applying the general correspondence between a two-point correlation function and its associated power spectrum to the velocity field*”, as stated by the authors. They propose a quantum version of the above description, which involves including the condensate phase, alongside the density weighted velocity field, as $\mathbf{u}(\mathbf{r}, t) = \sqrt{\rho(\mathbf{r}, t)}\mathbf{v}(\mathbf{r}, t) \exp(iS(\mathbf{r}, t))$. The resulting kinetic spectra then accurately describes the true spectrum of the condensate. For an examination of the vortex lattice, it is instructive to investigate the effectiveness of both approaches, and an example of the kinetic energy spectra is given in Fig. 5.2 for both the classical (left) and quantum (right) methods. For the classical spectrum, highly periodic structures are observed which appear at wavenumbers corresponding to the distances between vortices, which is not present in the quantum spectrum. Though the quantum spectrum may more accurately represent the kinetic energy of the condensate system, the periodic structure observed in the classical case are lost.

This is due to the included phase term breaking the correspondence between the power spectrum of a signal and its associated autocorrelation function [145], which are related by a Fourier transform according to the Wiener-Khinchin formula [146]. Though not truly the compressible and incompressible energies of a fully quantum system, due to the mean-field description the condensate system can be written in the form of a continuity equation (see Eq. (2.41)) with the probability current taking the form of the density weighted velocity field as used in Sec. 4.4.1. For such classical fluid descriptions, the classical method can be applied for isolating and describing both compressible and incompressible energies, and is used in the above mentioned studies. One can clearly see that this formalism allows us to identify the spatial structure of the observed pattern, which results from ensuring that the autocorrelation–power spectrum relation remains. As the periodicity is an important characteristic of the vortex lattice, the use of the classical kinetic energy spectral decomposition gives more information useful to this work, and will be used in the following sections.

5.3 Delta-kick dynamics

5.3.1 Non-rotating condensate

To fully characterise the effect kicking has on a rapidly rotating BEC carrying a vortex lattice it is instructive to first understand the dynamics following a kick in the absence of a vortex lattice. For an accurate comparison we need to take into account that in the absence of rotation no centrifugal forces appear, and we therefore adjust the trapping frequency of the non-rotating condensate such that the background densities match that of the rotating condensate. This is achieved from Eq. (4.5) by assuming the centrifugal term competing

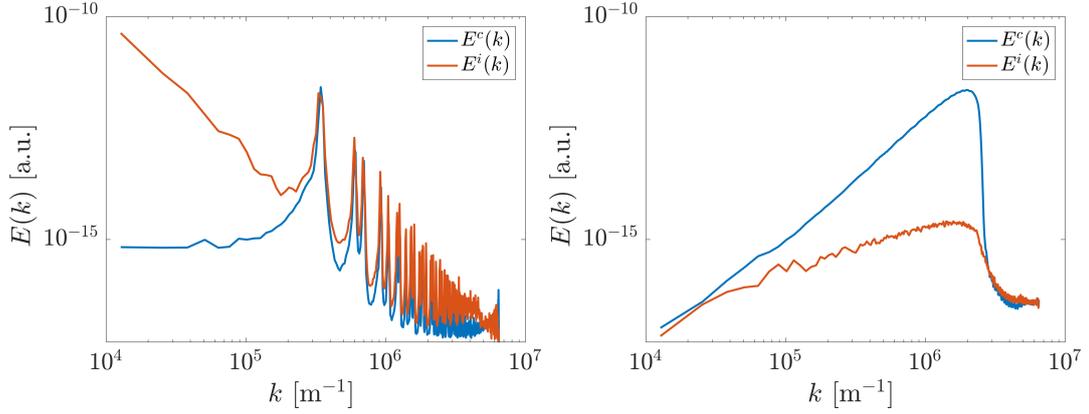


Figure 5.2: The compressible and incompressible kinetic energy spectra comparing the classical (left) and the quantum version (right).

against the harmonic oscillator term as $V_{\text{opt}} = (1/(2m))(\omega_{\perp}^2 - \Omega^2)r^2$, where $\Omega = 0.995\omega_{\perp}$ is the value chosen for the rapidly rotating case, as given by Eq. (4.5).

For a stationary (non-kicked) condensate the kinetic energy spectrum will remain constant during time-evolution, with a kick leading to the appearance of new, time varying components. To observe this we numerically evolve the system by setting $V(\mathbf{r}, t) = V_{\text{ext}}(\mathbf{r}) + V_{\text{opt}}(\mathbf{r}, t)$, where the time dependent optical potential is only active for $\tau = 10^{-5}$ s of the simulation time. As discussed earlier, this short kicking duration is sufficient to allow a phase imprint onto the condensate wavefunction, without any change in the density on the timescales of the kick itself. Next, we examine the compressible and incompressible kinetic energy spectrum following the kick (see Fig. 5.3). Unsurprisingly, the spectrum is dominated by a peak corresponding to the wavenumber associated with the optical potential at $k = 4\pi/(\sqrt{3}a_o)$ (indicated by the dashed line), and several smaller ones corresponding to higher harmonics of n -th next nearest neighbour components of the lattice. This makes sense as the nearest neighbour lattice spacings should be the dominant structure with higher order spacings dropping in intensity due to the finite system size. As no rotation is added by the imparted phonon modes, the incompressible energy is an order of magnitude smaller compared to the compressible energy. Although the incompressible energy should technically be zero for a condensate with no vortices, given the difference in magnitudes ($10^2 \rightarrow 10^4$) this can be taken as a k -space resolution issue following the decomposition, and can be assumed as zero. Following from this we will therefore restrict the analysis to the compressible part of the spectrum.

The evolution of the nearest neighbour peak in the compressible kinetic energy spectrum during the first 250 ms after the kick is shown in Fig. 5.4(a). It initially oscillates in and out of existence and eventually disperses over a wide range of wavenumbers. Snapshots of the density evolution are given in Fig. 5.4(b), which clearly show that the oscillations correspond to the existence of a transient lattice pattern with several revivals, having the same underlying

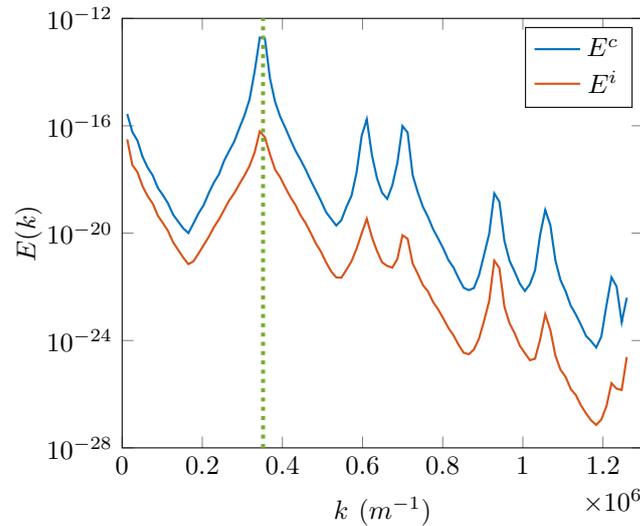


Figure 5.3: Compressible and incompressible energy spectra of a non-rotating condensate directly following a kick. A peak at $k = 4\pi/(\sqrt{3}a_o)$ can be seen, which corresponds to the lattice spacing, a_o (indicated by the dashed line), and the smaller, higher energy peaks can be attributed to higher harmonics between nearest and next-nearest neighbours. The incompressible spectrum is much smaller than the compressible spectrum, and thus is neglected for all further analysis. Reprinted from O’Riordan *et al.* [2].

structure as the optical potential. In fact, the lattice pattern is best formed whenever the main peak in the kinetic energy spectrum goes to zero, i.e. when the imprinted kinetic energy has been converted into density modulations.

5.3.2 Rapidly rotating condensate

Kicking a condensate carrying an Abrikosov vortex lattice with the above optical lattice gives an additional parameter, θ_Δ , which describes the orientation of the imprinted phonon lattice angle relative to the vortex lattice. We initially assume that the vortex and optical potential lattices have the same lattice constant, $a_v = a_o = a$, which means that symmetry allows for a restriction of the angle to $\theta_\Delta \in [0, \pi/3)$. In the following it can be seen that adjusting θ_Δ leads to the appearance of different, transient super-structures in the condensate density. If $\theta_\Delta = 0$ (see Fig. 5.5(a)) the kicking imparts kinetic energy at wavenumbers that are already well defined in the lattice. No significant change to the compressible kinetic energy spectrum is observed in this case, apart from small amplitude modulations on the well defined peaks.

However, if the angle between both lattices is finite, and not an integer multiple of $\pi/3$, superlattice structures appear after a short time (see Fig. 5.5(b)-(d)), which have a structure cell size that decreases for increasing values of $\theta_\Delta \in [0, \pi/6]$ and beyond which increases for larger values until the misalignment angle reaches the lattice symmetry point again at

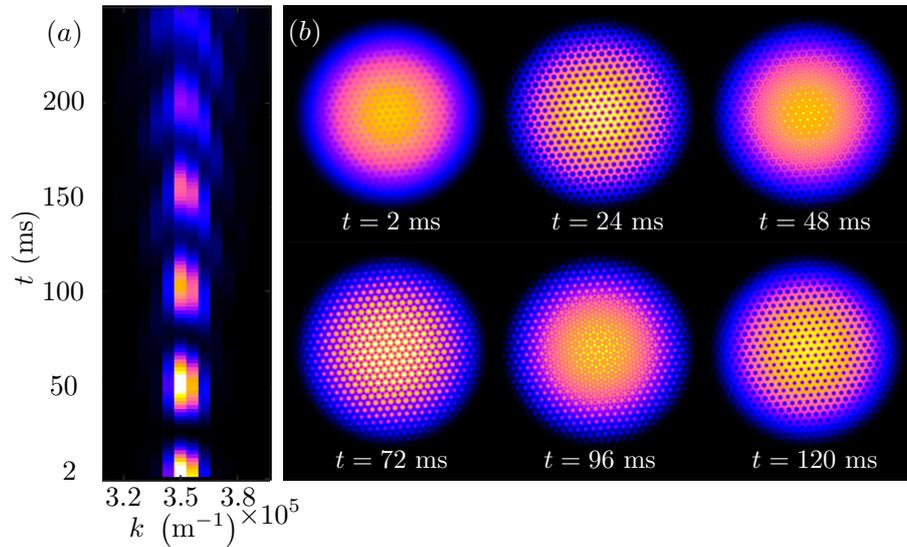


Figure 5.4: (a) Main peak of the compressible kinetic energy spectrum for a kicking strength of $V_0 \approx 1.35 \times 10^{-2} \mu$. It can be seen to revive, and eventually disperse over a wide range of wavenumbers. (b) Condensate densities at several times during the evolution. A pattern matching the optical potential can be observed to appear and disappear several times over the course of the evolution. Reprinted from O’Riordan *et al.* [2].

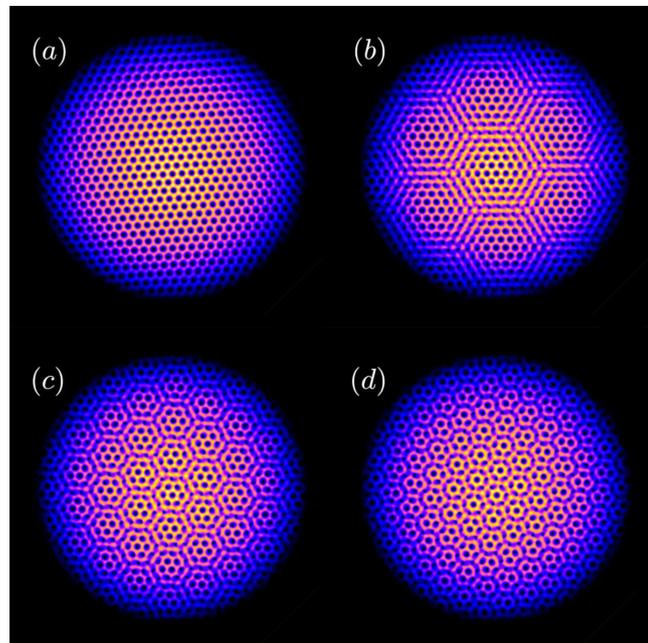


Figure 5.5: Condensate density at $t = 1.4 \times 10^{-2}$ s for several optical lattice rotation angles. The cell size of the super-lattice structures can be seen to shrink as the angle is increased. The angles for the examples shown are (a) $\theta_\Delta = 0$, (b) $\theta_\Delta = 2\pi/45$, (c) $\theta_\Delta = 4\pi/45$, (d) $\theta_\Delta = 2\pi/15$. Reprinted from O’Riordan *et al.* [2].

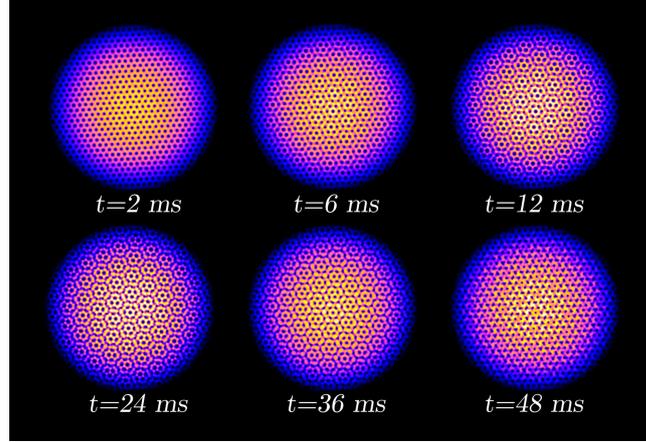


Figure 5.6: Condensate density shows visible moiré structures upon receiving a kick with $\theta_{\Delta} = \pi/9$. The appearance and disappearance of a moiré structure with wavelength $\lambda_M \approx 2.9a$ over a timescale of about 50 ms can be seen. Reprinted from O’Riordan *et al.* [2].

$\theta_{\Delta} = \pi/3$. These structures are transient, and several revivals can be observed before the condensate settles back into the vortex lattice structure with an increase in the background wavenumber spread, as expected based on kicking the non-rotating condensate. This settling does not imply that the condensate has dissipated the energy, but merely that the energy has spread over all possible wavenumbers, with the underlying lattice structure remaining dominant. An example of this for a fixed angle is shown in Fig. 5.6.

To explain the interference patterns observed for misaligning the optical and the vortex lattice, we employ moiré interference theory [147]. Moiré patterns are known to appear when two periodic structures are overlaid while slightly misaligned to each other, and can be calculated from the reciprocal lattice vectors. In all generality, any choice of equidistantly separated reciprocal lattice vectors can be parameterised as

$$\mathbf{g}_l = g_0 \left[\sin \left(\frac{2\pi l}{v} + \theta \right), \cos \left(\frac{2\pi l}{v} + \theta \right) \right], \quad (5.3)$$

where v describes the rotational symmetry of the lattice, l labels the vector direction on the unit circle, θ is the angle with respect to a chosen coordinate system and g_0 is the reciprocal lattice constant. For commensurate and triangular lattices we get $g_0 = 4\pi/(\sqrt{3}a)$, $v = 6$ and the vector directions are $l = [0 \dots v - 1]$. As only the relative mis-alignment between the vortex and the phonon lattice matters, we choose $\theta = 0$ for the vortex lattice and $\theta = \theta_{\Delta}$ for the optical potential alignment. All possible wavelengths that can appear in an interference pattern between two such lattices in position space are then given by

$$\lambda_{ll'} = \frac{\lambda_0}{|\mathbf{g}_{ll'}|}, \quad (5.4)$$

where $\mathbf{g}_{ll'} = \mathbf{g}_l^{\text{vtx}} - \mathbf{g}_{l'}^{\text{opt}}$, and $\lambda_0 = 4\pi/\sqrt{3}$ for the commensurate triangular lattices. This yields 6 interfering wavevectors, and as a result 6 moiré wavelengths. Figure 5.7 shows these

resulting interferences in both position and reciprocal space, with the colour indicating the corresponding wavenumbers and wavelengths. The band-like structure of the interferences have contact points at the angle of maximal and minimal alignment of the two lattices, $\theta_\Delta = (j\pi/3, j\pi/6)$ respectively. For a chosen rotation angle between these limiting values, we obtain several wavenumbers, and hence wavelengths. One can see from Fig. 5.6 for $\theta_\Delta = \pi/9$ that a pattern matching the longest wavelength, $\lambda_M = \max[\lambda_{ll'}] \approx 2.9a$, appears around $t = 24$ ms and is clearly the most visible one for the given angle. Shorter wavelengths, while are expected to be present, are hard to discern in this system, and therefore we will concentrate on the lowest wavenumber for the following analysis.

In \mathbf{k} -space the shortest $|\mathbf{g}_{ll'}|$ corresponds to adjacent wave-vectors with the smallest θ_Δ between them (see inset in Fig. 5.8). Due to the symmetry of the lattices the most visible structures are therefore given by $\lambda_M = \lambda_{00}$ for $\theta_\Delta \in [0, \pi/6]$ and $\lambda_M = \lambda_{01}$ for $\theta_\Delta \in [\pi/6, \pi/3]$ (see inset of Fig. 5.8). While this symmetry assumption no longer holds strictly true after the system has been kicked, it is still fulfilled to a very good approximation during the initial dynamics. One can then obtain the wavelength of the dominating moiré structure as [140, 148]

$$\lambda_M = \frac{a}{2 \sin(\eta/2)}, \quad (5.5)$$

where $\eta = \min(\theta_\Delta, \frac{\pi}{3} - \theta_\Delta)$ (see Fig. 5.8). These super-structures become observable when the wavelength becomes smaller than the radius of the condensate, which for the chosen parameters is $\lambda_M \approx 11a$ and which corresponds to an angle $\theta_\Delta \approx \pi/36$. One can see from Fig. 5.8 that once the relative angle is increased beyond this value the structure sizes shrink to a minimum value at the point of complete misalignment, $\theta_\Delta = \pi/6$, giving $\lambda_M \approx 1.93a$, and increase again up to the point of symmetry. Beyond this point the behaviour starts over, due to the symmetry of the lattice. Note that in principle the above procedure can be carried out for square or other optical lattice geometries.

The appearance of the moiré vector in \mathbf{k} -space can be confirmed from the numerical simulations by looking at the compressible kinetic energy spectra which is given in Fig. 5.9. Apart from the dominant peaks corresponding to the underlying triangular geometry of the Abrikosov lattice, which are independent of θ_Δ (straight lines in Fig. 5.9), a number of additional peaks appear. Their position is a function of the misalignment angle and the lowest wavenumber that appears increases its value with increasing θ_Δ . This is consistent with the moiré model and the appearance of density structures of differing size. Furthermore, a symmetric repeat of this structure about the $\theta_\Delta = \pi/6$ point is also visible, which corresponds to the $\pi/3 - \theta_\Delta$ lattice vector component. The minimum wavelength observed agrees with the theoretically determined minimum value of $\lambda_M \approx 1.93a$ and all other values over the range of observed angles. Note that for the higher harmonics at larger wavenumbers similar behaviour exists and is also covered by the moiré model.

To better observe the wavenumbers in the spectra, one can remove the background lattice spectrum. Corresponding to the systems shown in Fig. 5.9 this is given in Fig 5.10, with

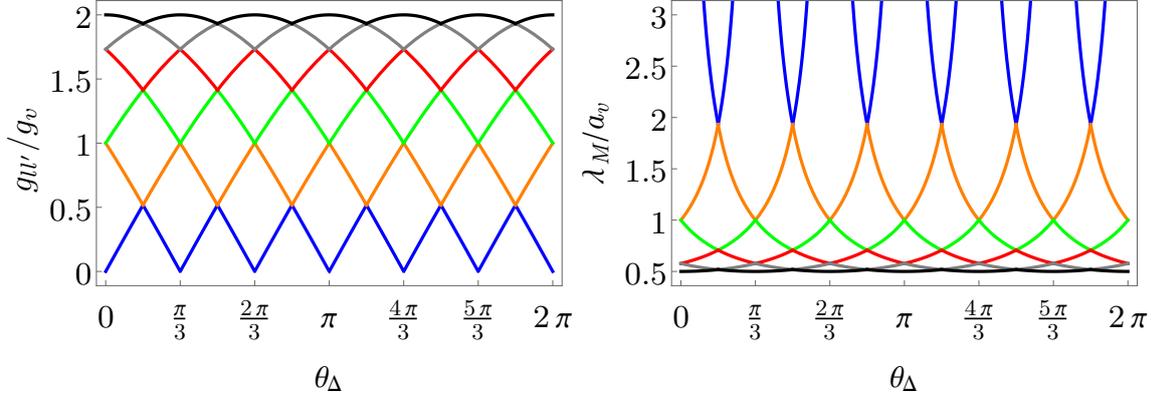


Figure 5.7: The moiré interference pattern in terms of wavenumbers (left) and corresponding wavelengths (right). The resulting wavenumber $g_{ll'}$ is normalised by the vortex lattice spacing in reciprocal space g_v , and the moiré wavelength λ_M is normalised by the vortex spacing in position space a_v . The lowest lying band in reciprocal space is the main contributor to the visible moiré interference patterns, corresponding to the largest wavelength (both in blue). Higher orders are shown for the reciprocal and the corresponding position space, with the higher modes only visible between the beats of the lowest wavenumber moiré interferences. Colours correspond between wavelengths and wavenumbers.

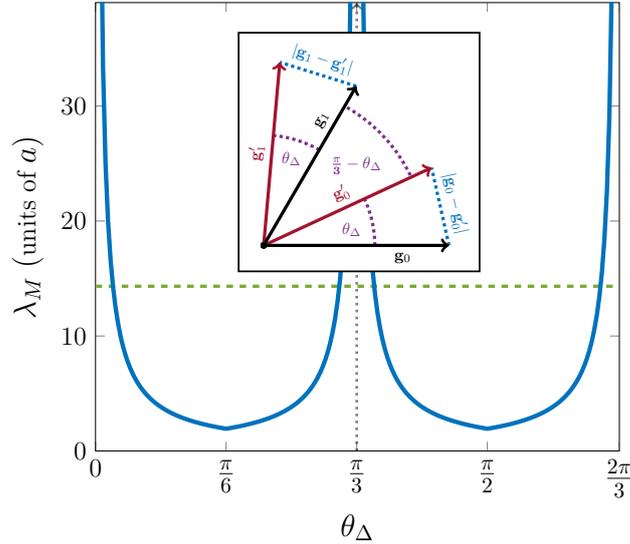


Figure 5.8: The dashed green line indicates the condensate radius. Inset: The different vectors in \mathbf{k} -space of the two lattices, with the optical lattice rotated by an angle θ_Δ . The $\mathbf{g}_{ll'} = |\mathbf{g}_l - \mathbf{g}_l'|$ vectors defining the dominant moiré wavelength are those for which the enclosed angle is smallest. Reprinted from O’Riordan *et al.* [2].

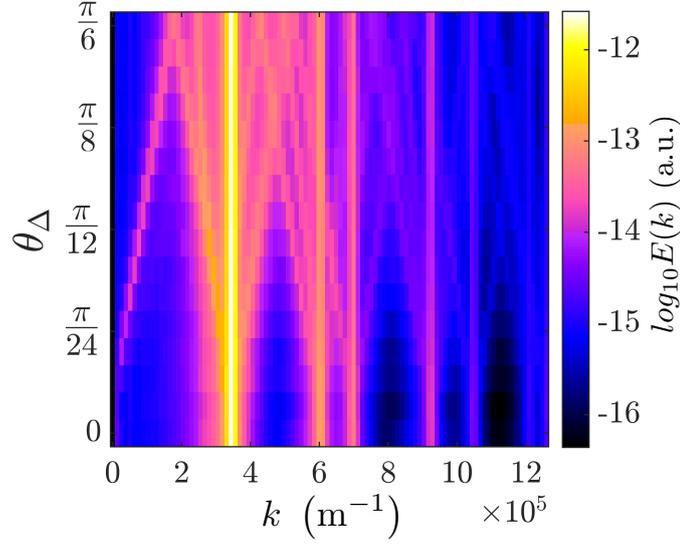


Figure 5.9: Compressible kinetic energy spectrum as a function of θ_{Δ} . All values are time-averaged over an interval $t = 0$ s to $t = 1$ s. The moiré peak corresponding to the lowest wavenumber can be seen shifting to larger values for increasing angles and similar behavior is visible for the higher order components. Reprinted from O’Riordan *et al.* [2].

the peaks at $t = 0$ removed. In this instance, lower wavenumber interference patterns are more easily visualised, with a slight enhancement of higher orders. We can also observe the appearance of secondary interference patterns across low to high wavenumbers, though visibility quickly diminishes approaching higher wavenumbers. It should be noted that given sufficient time the higher order wavenumbers can influence the density dynamics, but the system remains dominated by the lowest wavenumber. Due to the coupling between adjacent modes in the system, all interference patterns become hard to discern after some time.

In the following we will briefly discuss what happens for stronger kicking, or when the two lattices are non-commensurate. In the above the strength of the kicking pulse was chosen such that its perturbation only leads to a phase imprinting [42, 127], with minimal change to the initial density. If one increases the kicking intensity the situation becomes quite different and one can see from Fig. 5.11(a) that higher order wavenumbers become more strongly excited. This, in turn, leads to modulations of the condensate density at shorter wavelength and an example is shown in Figs. 5.11(b)-(e), with the numerically calculated averaged spectra given in Fig. 5.12. For fully realistic experimental situations it is necessary to also consider the heating of the condensate once the kicking becomes stronger. However, we do not extend this direction towards stronger kicking strengths in this work.

A situation where the optical and the vortex lattice have different lattice constants can be imagined to appear naturally due to experimental uncertainties. Defining $a_o = a_v(1 + \epsilon)$ the

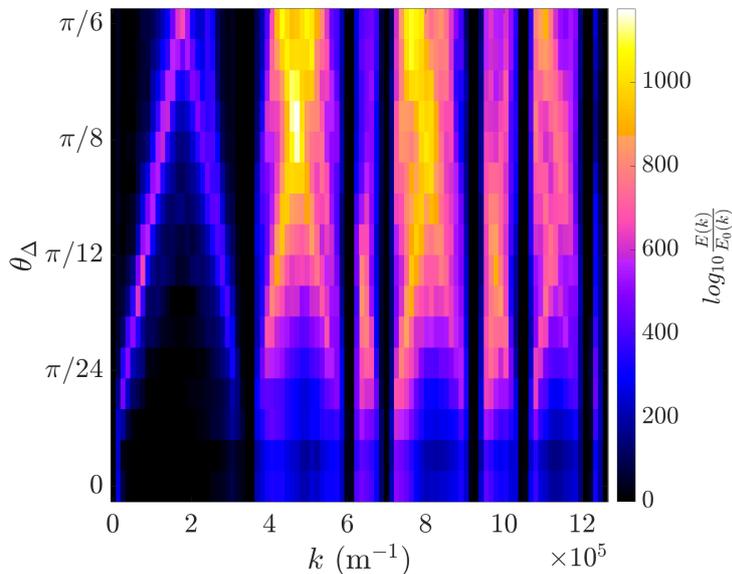


Figure 5.10: Compressible kinetic energy spectrum with the structures stemming from the static background removed. The effect of the moiré interference pattern for lower and higher modes is more clearly visible compared to Fig. 5.9.

expression in Eq. (5.4) can be calculated to be

$$\lambda_M = \frac{a_v(1 + \epsilon)}{\sqrt{2(1 + \epsilon)(1 - \cos \theta) + \epsilon^2}}, \quad (5.6)$$

which reduces to Eq. (5.5) for $\epsilon = 0$. Evaluating this expression for $\epsilon = (-0.1, 0, 0.1)$ shows that the largest moiré wavelength changes slightly for small values of ϵ , but it remains distinct enough from the higher order wavelengths to stay visible in the evolution (see Fig. 5.13). This ensures that the system examined here is experimentally realistic. For large deviations of the lattice constant, the assumption of the largest moiré wavelength playing the dominant role is no longer valid, as many of the interfering wavevectors approach comparable scales.

5.4 Vortex dynamics following a kick

While we have concentrated mostly on the phonon modes of the condensate density it is necessary to also examine the result of kicking on the vortex lattice. Figure 5.14 shows the densities and trajectories in the co-rotating frame for kicks of $\theta_\Delta = \pi(1/60, 1/10)$. Some deviation from solid-body rotation is observed at the edge of the examined vortex region, but the central vortex regions show almost no variation from their ideal lattice positions. Both cases demonstrate the same overall dynamics, with the vortex lattice showing extreme robustness against the phononic perturbances that build following a kick.

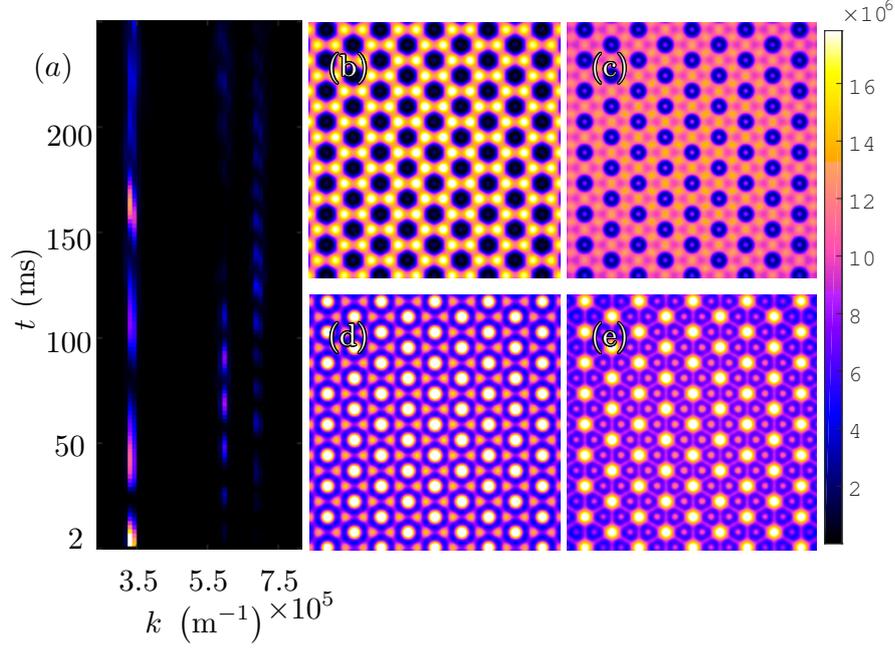


Figure 5.11: (a) For a kicking strength of $V_0 = 5.4 \times 10^{-2} \mu$ for a non-rotating condensate higher order modes become non-negligible contributors to the compressible kinetic energy spectrum. This leads to the appearance of higher-order peaks and a different behaviour of the condensate density. Close-ups of the density structures are shown for (b) 24 ms, (c) 36 ms, (d) 56 ms, and (e) 88 ms. Note that the structures with higher amplitudes in these plots are given by the optical lattice constant, a_o , which is set to the mean inter-vortex distance for the rapidly rotating condensate. In the presence of vortices this is anticipated to create many additional interferences. Reprinted from O’Riordan *et al.* [2].

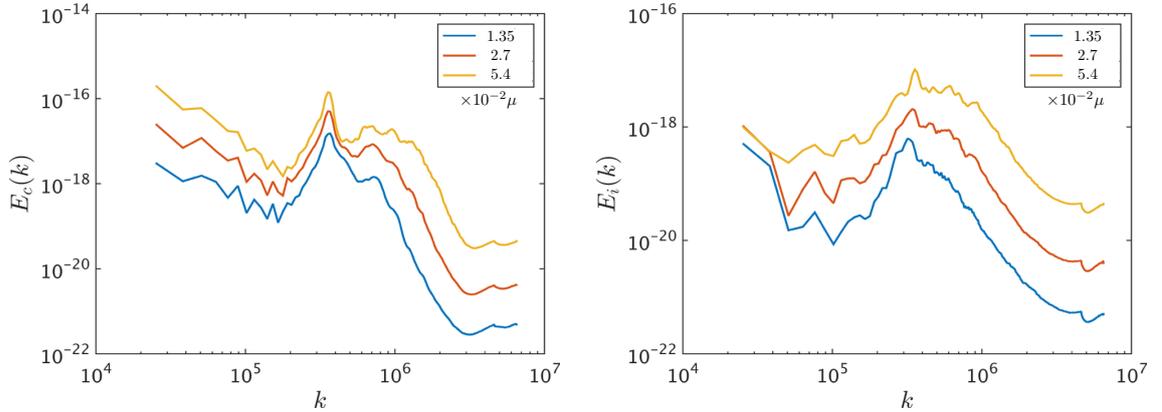


Figure 5.12: For kicking pulses of $V_0 = (1.35, 2.7, 5.4) \times 10^{-2} \mu$ the kinetic energy spectra show a clear increase in higher order modes being excited for both compressible (left) and incompressible (right) cases with increased kicking strength.

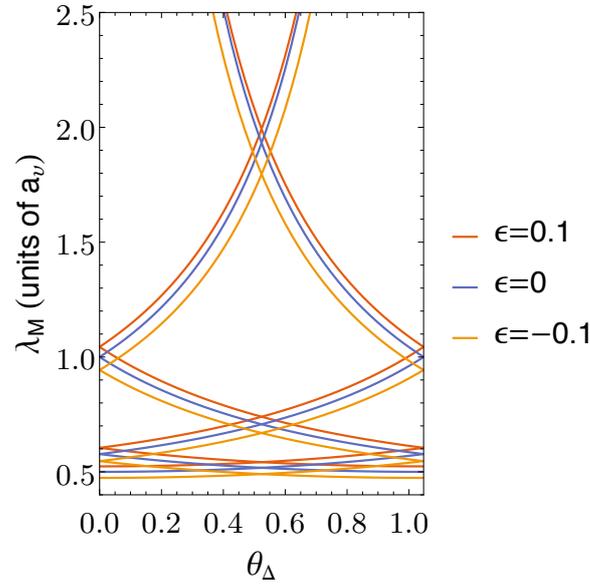


Figure 5.13: Small deviations for ϵ can be seen to only lead to minor changes of the induced moiré wavelength.

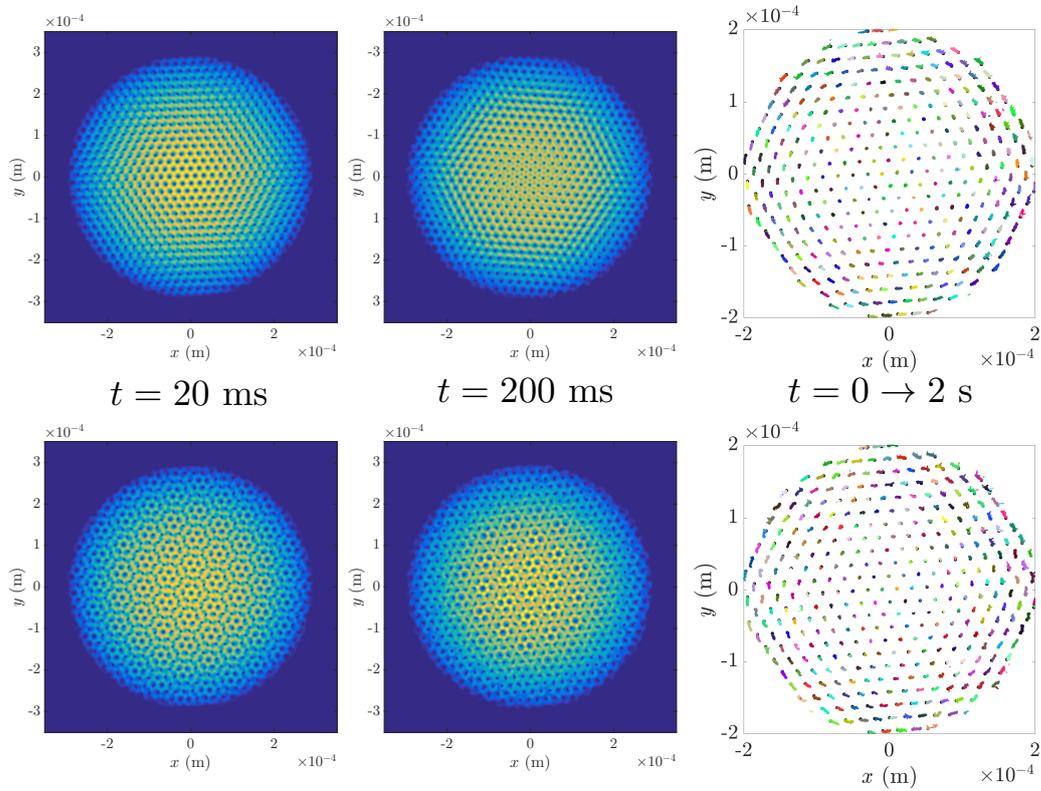


Figure 5.14: The condensate densities are shown for times $t = (20, 200)$ ms, and for $\theta_\Delta = (\pi/60, \pi/10)$ (top and bottom). The resulting trajectories of the vortices show minimal deviation from their initial lattice positions over 2 seconds of time evolution.

5.5 Conclusions and outlook

The purpose of the previous work was to investigate the effect of a kicked optical lattice potential with a well defined periodicity on a vortex lattice having its own periodic structure. As observed, the kicking had a negligible effect on the vortex positions, with the lattice remaining well defined with near constant spacings, even in the long time limit. The optical lattice kicks solely modified the background density around the vortex cores. One could observe that the vortex lattice remains incredibly stable and strongly resilient to perturbations and modulations in density. Given the stability of the lattice, it can be expected that even for a finite number of kicks, the lattice will remain mostly unbroken. It can also be expected that heating effects would play a role in destabilising the condensate, and to determine any long-term realistic behavior from this system, including the treatment of a non-negligible thermal cloud contribution would be necessary.

The kinetic energy spectra provided a useful tool to investigate the order of the imparted wavenumbers following a kick, and were very closely matched with the results from moiré interference theory. The existence of this moiré interference effect has some interesting consequences. It can, for example, be thought of as a tool to test the periodicity of a structure that is otherwise too small to be resolved without time of flight. In general, this technique may be used on other types of periodic structures, for example crossed solitons in 2D condensates [149]. As optical lattices can offer a large number of free parameters (lattice constant, amplitude, misalignment, geometry, ...), a full toolbox can be developed from this technique. This could potentially be used for probing condensate systems, applying well-defined structures onto stationary condensates, and for generating the aforementioned interferences in the presence of an underlying periodicity.

While in my thesis proposal the system discussed in this chapter was suggested to investigate delta-kicked chaotic dynamics [150], it became clear very quickly that the robustness of the vortex lattice with respect to realistic kicking strengths did not allow for significant vortex dynamics. Nevertheless, an investigation into periodic kicking is still an interesting (and numerically even more challenging) endeavor. As vortex lattices feature a well-defined rotation rate and 6-fold symmetry, the system rotates to a symmetric position every $\pi/3$. Providing a periodic kick at this rate, would therefore allow the lattice to always have the same alignment angle, at least for an initial number of kick repetitions.

Chapter 6

Defect engineering in vortex lattices

6.1 Introduction

In the previous chapter we have discussed the application of kicked optical potentials to vortex lattice carrying condensates. The vortex lattice proved to be very resilient, with the individual vortex positions largely unaffected by the perturbation, apart from some minor oscillations about their equilibrium positions. The global structure of the lattice was preserved, and remained well ordered. It begs the question of how to create a disordered arrangement from this well ordered structure. Given the similarity of vortex lattices to some crystal structures in condensed matter, it is interesting to examine lattice ordering, and investigate the behaviour of the vortex lattice subjected to a defect or vacancy. True crystalline solid-state materials have feature sizes on the order of Ångstroms, and can be difficult to observe experimentally, especially the appearance of lattice defects. Even more difficult can be the creation of well defined defects within the crystal structures. Cold atomic systems allow for the same fundamental effects to be investigated on mesoscopic scales, and in currently experimentally realisable systems. Here, we will demonstrate the creation of crystal defects in a BEC vortex lattice. With this, an interesting question to ask is if the lattice were to lose a vortex at a specified location, how would the overall system respond?

In Section 4.3.2 we discussed the use of phase imprinting techniques to modify the condensate phase. While it was discussed that a vortex can be created with the required $\pm 2\pi$ phase winding, phase imprinting can also be used to annihilate a pre-existing vortex. Assuming a vortex with clockwise phase winding, the direct application of a phase profile with counter-clockwise winding can remove the topological charge of the vortex and annihilate it. In this chapter we will examine the use of this technique, and investigate the resulting effects on the vortex lattice order.

This work has been accepted for publication in Physical Review A, with the content referenced therefrom.

6.2 Model

To investigate the evolution of a vortex lattice subjected to the phase imprinting technique we once again numerically solve the Gross–Pitaevskii equation in two dimensions, assuming a strong confinement along the third axis, following the model given in Sec. 4.2.1. This allows us to restrict the dynamics to the x - y plane and focus fully on the Abrikosov lattice geometry. In the frame co-rotating with the condensate the nonlinear mean-field equation governing the BEC wave-function is once again given by

$$i\hbar\partial_t\Psi(\mathbf{r},t) = \left[-\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{r}) + g_{2D}|\Psi(\mathbf{r},t)|^2 - \Omega L_z \right] \Psi(\mathbf{r},t). \quad (6.1)$$

While the Abrikosov ground state is perfectly ordered, removing or adding vortices will lead to disorder in the lattice. Recently, quantifying the disorder of vortex lattices has become an active topic of interest [33, 151]. Here, we suggest to quantify the order of the vortex lattice by first determining the position of each vortex by summing the wavefunction phase over adjacent grid sites and locating $\pm 2\pi$ windings. This gives a vortex position estimated to the nearest numerical grid point. A linear least-squares fit is then performed to more accurately determine the vortex core location to sub-grid resolution, as described by Sec. 3.6.2. This allows us to determine the wavefunction zeroes within the region, and obtain a continuous rather than discrete range of values for the vortex positions. Since tracking many-body dynamics is a difficult problem, we make use of the Delaunay triangulation and Voronoi tessellation techniques from computational geometry to examine the ordering of the the vortex lattice, as introduced in Sec. 4.4.2. In the ideal triangular Abrikosov lattice, every vortex has 6 nearest neighbours, $l = (0, \dots, 5)$, located at $\theta_l = l\pi/3$ around the polar angle. Any perturbation that breaks this symmetric arrangement can be easily observed using Delaunay triangulation. This technique generates a mesh from the vortex positions, which makes it easy to check for the presence of non 6-fold neighbouring vortices. These vortices are termed as n -fold topological lattice defects, where n is the number of connected edges, and dislocation defects can form when, for example, a 5-fold and a 7-fold defect pairs. Since all these structures are easily countable, this method allows us to characterize the effect that a well-defined perturbation has on the lattice. Alternatively, Voronoi tessellations, which can be generated from the Delaunay triangulation and vice versa, make it easy to observe structural changes in the lattice, and are useful for visualising quantities local to the regions of interest.

As unperturbed Abrikosov lattices in BECs are well ordered everywhere in the bulk region [152] we define the previously mentioned radial boundary at approximately 2/3 of the maximum density, which corresponds to $r = 2 \times 10^{-4}$ m from the BEC centre and restrict our

analysis to vortices inside it. This leaves an edge boundary of approximately $4a_v$ wide where vortices are not counted. As discussed in Sec. 4.2.1, this gives approximately $N_v = 341$ vortices in the region of interest, which is a sufficient number for the following analysis. Given the coordinate locations for each vortex within the boundary, it is possible to calculate statistical quantities that characterised the degree to which the lattice is ordered. While the Delaunay and Voronoi representations of the lattice are useful methods for the reasons described above, having a metric to measure the deviation from a symmetric arrangement is also useful. As our system is of finite size, the commonly used translational correlations have only limited value, and we will focus in the following on orientational correlations which quantify how the rotational symmetry of the vortices correlates across length-scales. The orientational correlation function for a two-dimensional lattice with 6-fold rotational symmetry is defined as

$$g_6(r) = \frac{1}{N(r)} \sum_{j,k}^{N(r)} \zeta_6(\mathbf{r}_j) \zeta_6^*(\mathbf{r}_k), \quad (6.2)$$

with

$$\zeta_6(\mathbf{r}_j) = \frac{1}{n_j} \sum_l^{n_j} \exp(i6\theta_{jl}), \quad (6.3)$$

where $N(r)$ is the number of paired vortices at locations \mathbf{r}_j and \mathbf{r}_k separated by $r = |\mathbf{r}_j - \mathbf{r}_k|$, ζ_6 is the orientational order parameter, l runs over the nearest neighbouring vortices, n_j is the number of elements in the respective bin-range ($n_j = 6$ for a perfect triangular lattice), and θ_{jl} is the angle a paired vortex and nearest neighbour makes relative to a reference axis [153]. We examine the orientational correlation function as a measure of the order of a “vortex unit cell”, defined by the angle made by nearest neighbours to an individual vortex. For a perfectly ordered triangular lattice this value will tend to 1 at $r = a_v$, nearest neighbour and higher order crystal spacings, and 0 elsewhere.

6.3 Phase imprinting defects

The phase imprinting methods discussed in Sec. 4.3.2 can be used to annihilate a vortex from the lattice by applying a phase profile of opposite winding to remove the vortex phase singularity. This will leave the condensate with a density depletion at the prior location of the singularity, which will consequently fill in and excite phonon modes in the condensate. Alternatively, one can also change the direction of rotation of a vortex in the lattice by applying a 4π magnitude phase in opposition to the present direction of rotation. Many proposed methods and resulting implementations of this set of techniques have been demonstrated for vortex generation [42, 43, 112–114, 154], and so it is assumed that the methods and discussions provided are realisable. In the following we will introduce the effect of removing a vortex on the lattice system, and discuss the resulting dynamics.

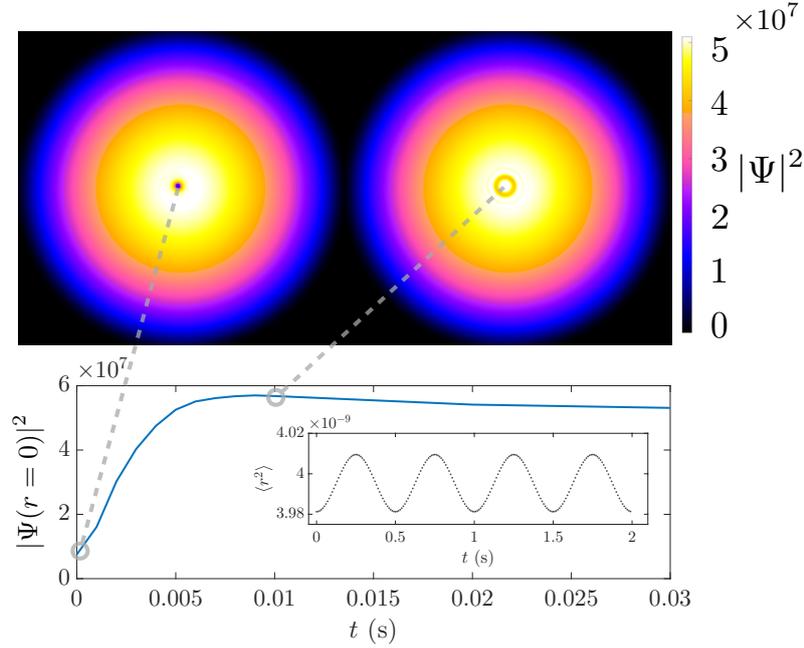


Figure 6.1: The evolution of the condensate density is shown for the initial state, and after 10 ms of evolution. The removal of the phase singularity at $t = 0$ leads to a filling in of the density dip, which can also be seen from the line-plot (the singularity is initially located at $r = 0$). The process excites the monopole mode at frequency $2\omega_{\perp}$ (see inset).

6.3.1 Single vortex dynamics

To fully understand the effects of removing a vortex from the lattice system, we will first investigate removing the angular momentum from a single vortex-carrying condensate. For this we apply a phase pattern that cancels the pre-existing 2π phase winding, and time evolve the system to examine the resulting dynamics. Figure 6.1 demonstrates the resulting dynamics, and as expected, the density depletion in the condensate fills in following the phase singularity removal. Since the system is rotationally symmetric, we also plot the expectation value of the squared radius, $\langle r^2 \rangle$, where $r^2 = x^2 + y^2$, which clearly shows that the annihilation process excites the breathing mode at the expected frequency of $2\omega_{\perp}$ for a two-dimensional system [155, 156]. The energy shift due to the phase removal can also be meaningfully characterised via the ratio of compressible (phonon) and incompressible (vortex) kinetic energy as given by Sec. 4.4.1, with the spectra shown in Fig. 6.2. The kinetic energies are determined from the density weighted velocity field, $\mathbf{u} = |\Psi| \frac{\hbar}{m} \nabla \theta$, where θ is the phase of the condensate, by making use of Eq. (4.31). As one can see from Fig. 6.2 after the vortex is annihilated and phonons are created, the energy ratio drops to favour lower incompressible-to-compressible values, in particular for higher wavenumbers. The latter is likely due to the removal of larger kinetic energies from the atoms closer to the vortex core.

While the above example suggests that erasing vortices is a straightforward and controllable

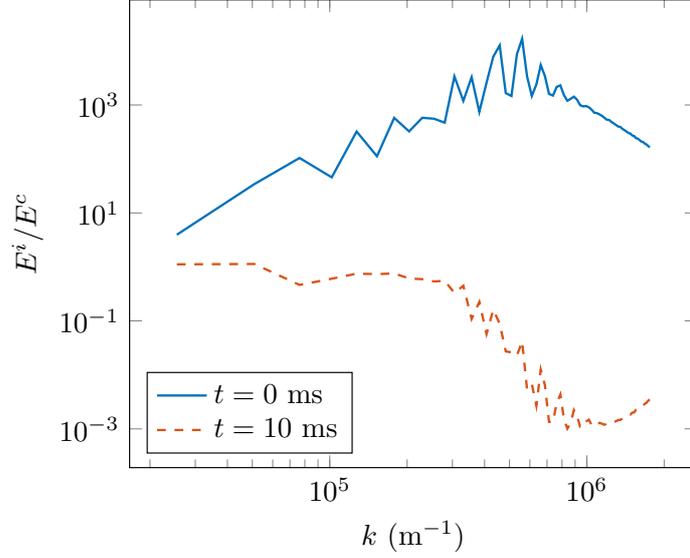


Figure 6.2: Ratio of incompressible to compressible energy at $t = 0$ (solid) and $t = 10$ ms (dashed). Initially the incompressible energy is greater than the compressible due to the presence of the vortex, giving values greater than unity for all k . After application of the phase profile, the vortex is annihilated, with the energy released as phonons, indicated by a decrease in incompressible energy for all k values.

process, this assumption must be checked for the situation where the imprinted phase singularity and the existing phase are not perfectly centered on each other. This is shown in Fig. 6.3, and one finds that cases where the imprinted profile is sufficiently close to the core (i.e. within twice the healing length, $\xi \approx 1.06 \times 10^{-6}$ m) the existing vortex gets erased as before. However, beyond this distance a separate antivortex gets created and the vortex-antivortex pair travels to the edge of the condensate system and begins to circulate around [157], as is shown in Sec. 4.1.1. For a densely packed lattice of vortices, however, this is not a problem as the typical distance between vortices is of comparable size with the healing length.

6.3.2 Lattice dynamics

The removal of a single vortex from the vortex lattice by phase erasing initially affects only the nearest neighbours, as the phase gradient is only significant over the length scale of a healing length close to the erased singularity. The altered velocity profile will lead to the remaining vortices leaving their initial equilibrium positions in the Abrikosov lattice and the excitation of phonon modes. However, in the lattice areas away from the impurity, these phonon modes have only minimal impact on the geometry, as was observed with the optical lattice kicking described in Sec. 5.4. To characterize the vortex dynamics following the application of the phase profile, we will in the following track each individual vortex throughout the full time-evolution and use the resulting trajectories, Delaunay triangulations,

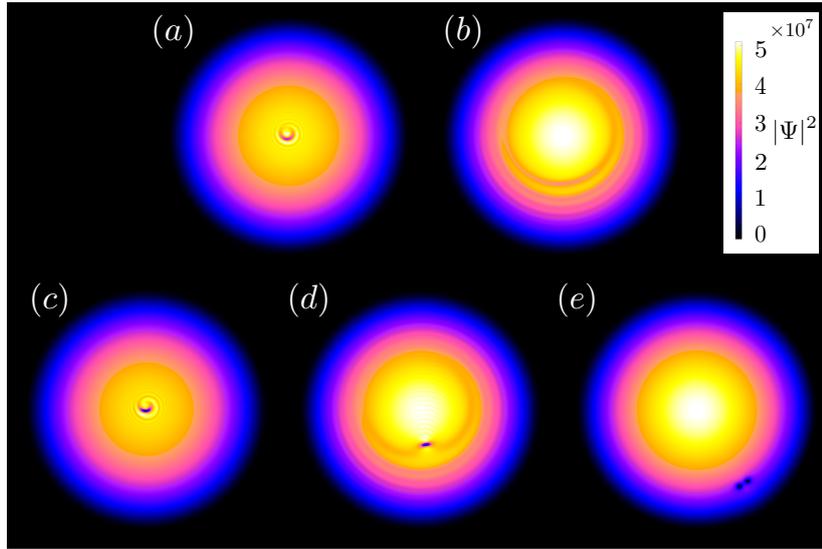


Figure 6.3: The condensate evolution following an uncentered phase imprint. For an imprint where the singularities of the vortex and the imprinted phase are less than a healing length away from each other, the existing vortex is annihilated and phonon modes are excited (*a, b*). However, beyond this distance an antivortex is created, which travels with the pre-existing vortex and circulates the condensate (*c, d, e*). The distance for cases (*a, b*) and (*c, d, e*) are $r = 1.36 \times 10^{-6}$ m, and $r = 2.73 \times 10^{-6}$ m respectively.

and Voronoi tessellations for analysis.

Let us first consider the situation where a single vortex is erased within the central area of the vortex lattice. In Fig. 6.4 we show the trajectories of the remaining vortices over a time-scale of 4 seconds, and the corresponding wavefunction densities at $t = (0.1, 1, 4)$ s in Fig. 6.5. One can see that a long-lived vacancy is maintained close to the centre with the adjacent vortices rotating faster than the lattice due to the loss of the local velocity field. The honeycomb-like vacancy region eventually decays and the system settles into a new local geometry. The velocity field magnitude and direction of the condensate local to the removal site are shown in Fig. 6.6 for $t = (0.01, 0.1)$ s. Almost directly following the imprint, the removal site shows a phononic flow that fills in the now unstable dip in the density, and which eventually disperses across the lattice. The small-scale flow magnitude that exists in the remaining region is sufficient to allow for the decay of the honeycomb-like structure during the subsequent time evolution.

Similar behavior can be observed if the erased vortex is not within the central area, as long as it is within a region of constant areal vortex density. However, being closer to the edge of the lattice reduces the stability of the perturbed region, which is likely due in part to the solid-body velocity field of the vortex lattice away from the centre. The overall lattice remains well structured after a vortex removal, as can be seen from the orientational correlation function shown in Fig. 6.7 for different times. Although the gaps between the peaks that exist at $t = 0$ disappear during the evolution due to the presence of the phonon excitation, the overall

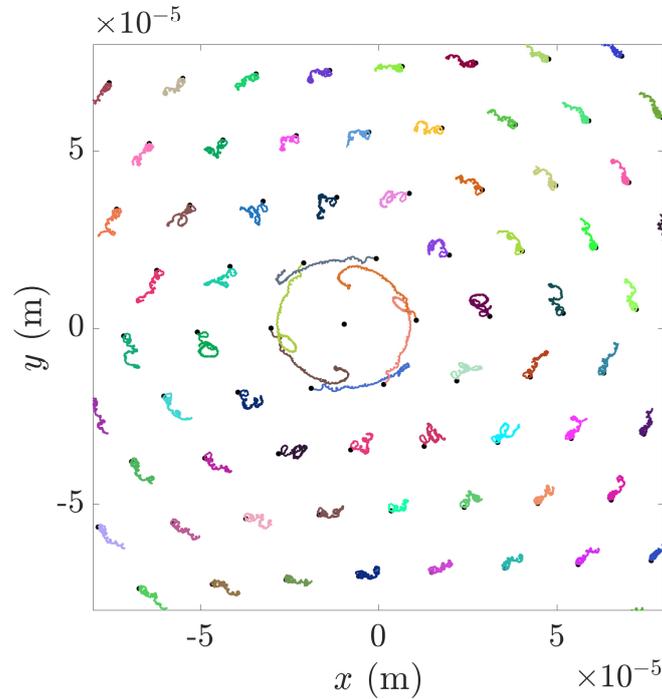


Figure 6.4: The trajectories of the vortices over 4 seconds ($dt=10^{-5}$, $\text{steps}=4 \times 10^5$) following the removal of the vortex closest to the center, where each color represents a unique trajectory path, and black dots indicate positions at $t = 0$. The vortices can be seen to move counter-clockwise in the co-rotating frame due to the loss of the local velocity field. However, the effect of the removal decreases quickly with increased radial distance.

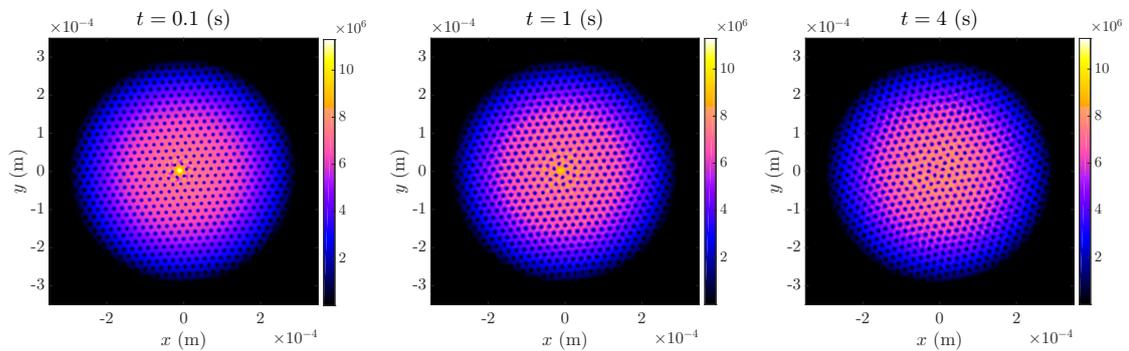


Figure 6.5: Wavefunction densities following a vortex removal close to the centre of the condensate. The vacancy region exists for some time after the removal, before the honeycomb-like structure decays and locally disorders the vortex lattice.

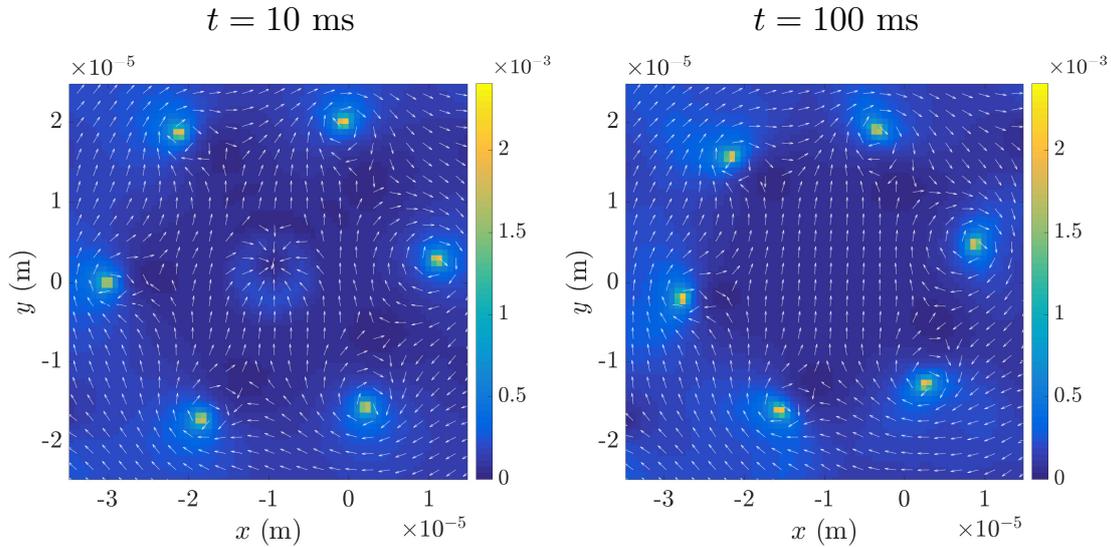


Figure 6.6: The velocity field magnitude and direction following the removal of a centrally located vortex from the lattice. At $t = 10$ ms a phononic flow can be observed, resulting from the vortex annihilation. For $t = 100$ ms, a net flow across the previous position of the erased vortex indicates an instability in the region, eventually leading to decay of the vacancy region.

correlations remain high for long times and near constant across all length scales. The slight peak softening arises from the vortices no longer being aligned to a perfect triangular lattice position, which is indicative of a weak disordering or distortion of the lattice structure.

As described above, the Delaunay triangulation of the lattice can give a graphical overview of how connected the different vortices are, and therefore what changes to the lattice structure have occurred [153]. We show the case where a single vortex was removed from the lattice centre in Fig. 6.8. One can see that a pair of (5,7)-fold connected lattice defects have formed after the removal (at 10 ms), which slightly adjusts and becomes stable for longer times. Removing vortices at different positions in the lattice shows similar behavior, with a localization of the disordered region not far from the site of the vortex removal.

If the phase imprinting is not directly aligned with the vortex singularity, other n -fold dislocations can be found in the Delaunay triangulation. This is due in part to the vortex core size becoming comparable to the average spacing between the vortices in rapidly rotating condensates, and therefore the imprinted changes to the velocity field affect more nearby vortices. To investigate this effect at the lattice centre, a series of simulations were performed by imprinting at different locations as a function of r and θ . Due to the lattice symmetry, the resulting dynamics can be completely obtained by examining the region of $r = [0, a_v/2)$ and $\theta = [0, \pi/6)$. A schematic of this shown in Fig. 6.9.

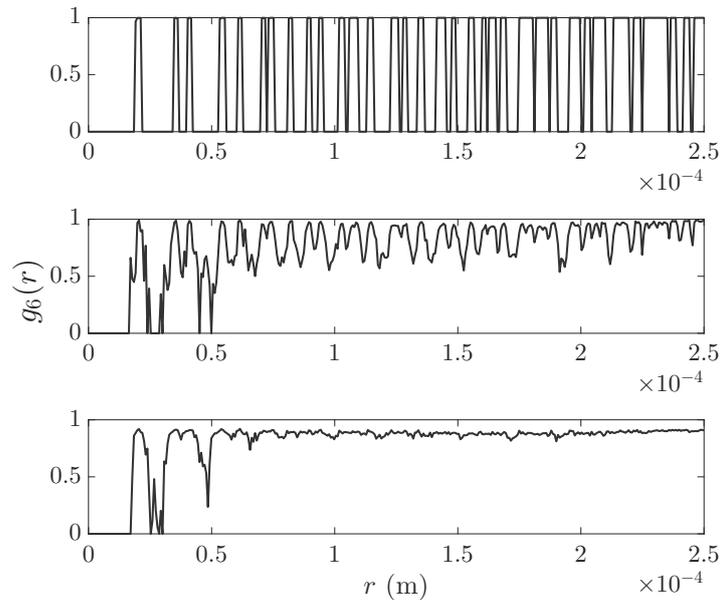


Figure 6.7: The orientational correlation function for the vortex lattice after removing the central vortex is given for $t = (0, 1, 6)$ seconds (top to bottom). The peaks at $t = 0$ appear at nearest neighbour, next nearest, and higher order distances. Due to finite binning of the lengths, the peaks become grouped to 1 at higher length scales. For times greater than $t = 0$ the peak correlations drop, however, the large value at long times indicates a well ordered lattice as high correlations are observed across all length scales.

Figure 6.10 shows the time-averaged number of lattice defects following the imprint within this examined region. One can see that if the displacement is still within the core of the vortex, on average 1 or 2 defects are created of the 5-fold (*a*) and the 7-fold (*b*) kind. At $r \approx a_v/4$ from the core centre, the imprint tends to create upwards of 3 to 4 defects, which again tends back to the average of 2 beyond this region. This shows that the previously discussed issue resulting from the creation of antivortices through imperfect alignment does not exist in Abrikosov lattices, and we will concentrate on the perfect imprint of the phase in the following discussions.

To further demonstrate the localized nature of the defects, let us briefly discuss the situation where two vortices are erased in separate regions away from the lattice centre. The Delaunay triangulation for this case is shown in Fig. 6.11, and the independence of the two localized regions is clearly visible, with each showing behavior similar to the case discussed above. Since we are limiting ourselves here to perfect imprinting, we also show the number of edges formed between vortices as a function of time for 5, 6 and 7 nearest neighbours respectively (N_x) in Fig. 6.12. One can see that the initial perturbation settles quickly to values similar to the ones above.

In addition to simply erasing vorticity, we can also use phase imprinting to create varying degrees of disorder. By, for example, applying an appropriate 4π magnitude phase imprint

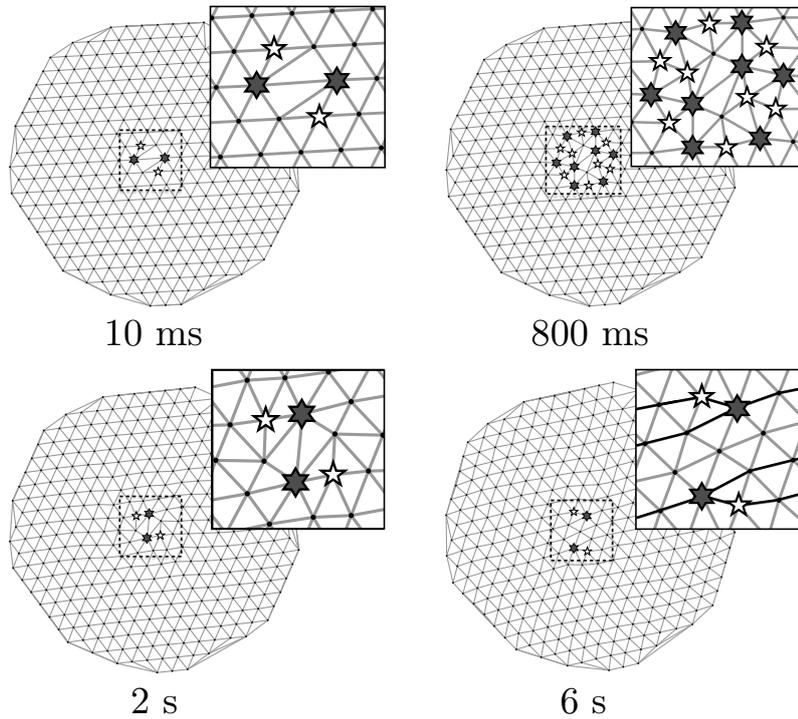


Figure 6.8: Delaunay triangulation of the vortex lattice after removing one vortex, shown at $t = (0.01, 0.8, 2, 6)$ s. The resulting lattice defects are indicated by white and gray stars for 5-fold and 7-fold defects respectively. One can see that two (5,7) dislocations are formed quickly, which settle and persist in the lattice for long times. Lattice dislocation lines are indicated as black lines for inset $t = 6$ s.

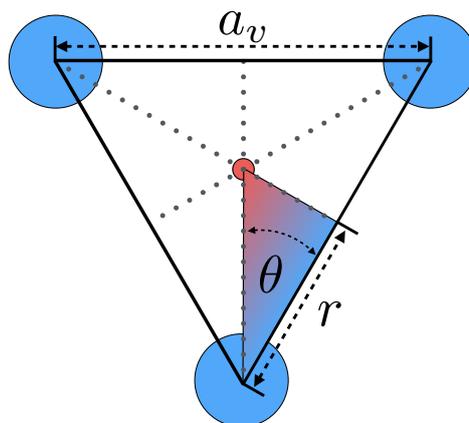


Figure 6.9: Schematic of triangular cell of the vortices. Due to the symmetry of the cell, an imprint within the shaded region can allow for all of the resulting dynamics to be determined.

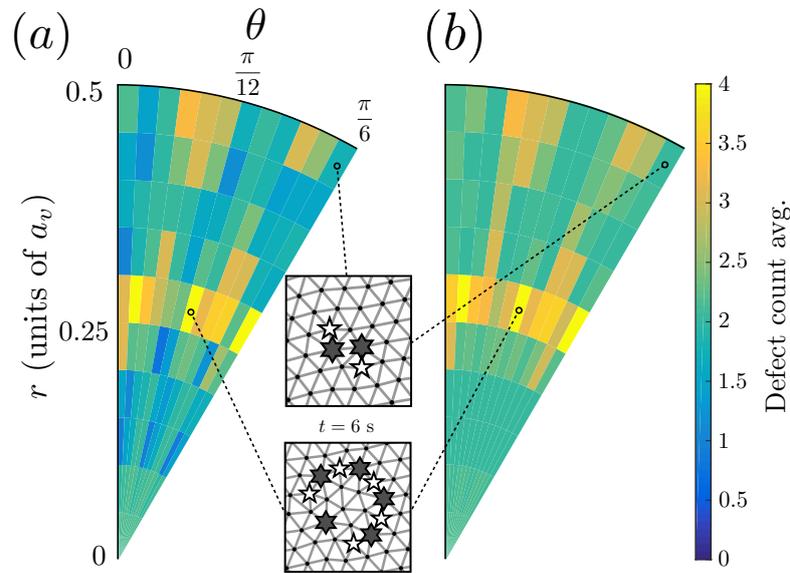


Figure 6.10: The time-averaged number of defects appearing over a range of imprint positions, relative to a central vortex from $t = 1 \rightarrow 10$ s, and allowing 1 s of settling time. Both 5-fold (a) and 7-fold (b) defects are shown. The insets show a snapshot of two different parameter regions at $t = 6$ s. A high simultaneity is observed between their appearance, where a paired (5,7) defect indicates a lattice dislocation. However, not all 5 and 7-fold defects pair, as some can exist individually, or pair with other n -fold defects.

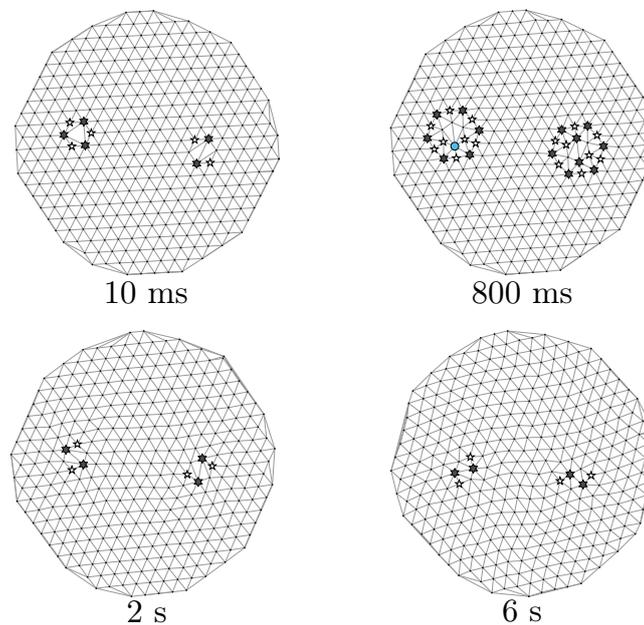


Figure 6.11: Delaunay triangulation of the vortex lattice upon removal of two vortices at either sides of the lattice for $t = (0.01, 0.8, 2, 6)$ s. The resulting defects remain localized for long times and can therefore be considered independent. The lattice largely remains ordered, similar to the case of removing the central vortex discussed above. The appearance of an 8-fold connected vortex can also be observed at $t = 0.8$ s (blue dot).

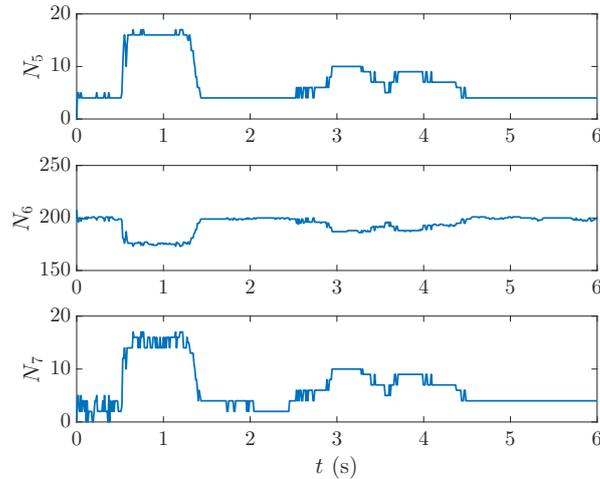


Figure 6.12: The defect (N_5 , N_7) and triangular lattice (N_6) nearest neighbour count taken from a Delaunay triangulation of the vortex lattice following the removal of two vortices on opposite sides as a function of time. After a brief settling time, the lattice attains an almost constant defect count.

we can replace a vortex with an antivortex at a given position. Since this does not require a change in the local density, all resulting perturbations stem from the adjusted velocity field of the vortex that has been flipped [158], with snapshots at $t = (10, 100)$ ms shown in Fig. 6.13. Here we can clearly see the change in velocity field direction around the central vortex relative to the surrounding vortices. This in turn causes the surrounding vortices to rotate at a different rate than the solid-body rate of the overall lattice, creating a loss of triangular symmetry in this region on a faster scale than removing a single vortex. It is immediately obvious that such a situation is unstable, which can be confirmed by observing the creation of a large number of defects during the evolution, as shown in Fig. 6.14. An increase in the number of defects can be seen up to approximately $t = 3$ s, during which the antivortex causes local disordering of the lattice, annihilates with a nearby vortex, and gives rise to the creation of a large number of (5,7) defect pairs. After this the number of defects no longer grows, but instead fluctuates about a stable value which is greater than that of the previously examined cases.

A final class of possible perturbations is the removal of a cluster of neighbouring vortices from the lattice, and in Fig. 6.15 we show the results from erasing an entire seven vortex unit cell from the condensate. As expected, one can see that the number of lattice defects rises considerably and does not settle during the time over which we can simulate the condensate. In this case, the disordered regions occupy a large area of the lattice and the number of 6-fold connected vortices becomes very low.

Comparing the orientational correlation function for the three cases discussed in this section (removing two distant vortices, creating an antivortex, and removing 7 vortices) also demonstrates the different degree of disorder they produce (see Fig. 6.16). The removal of

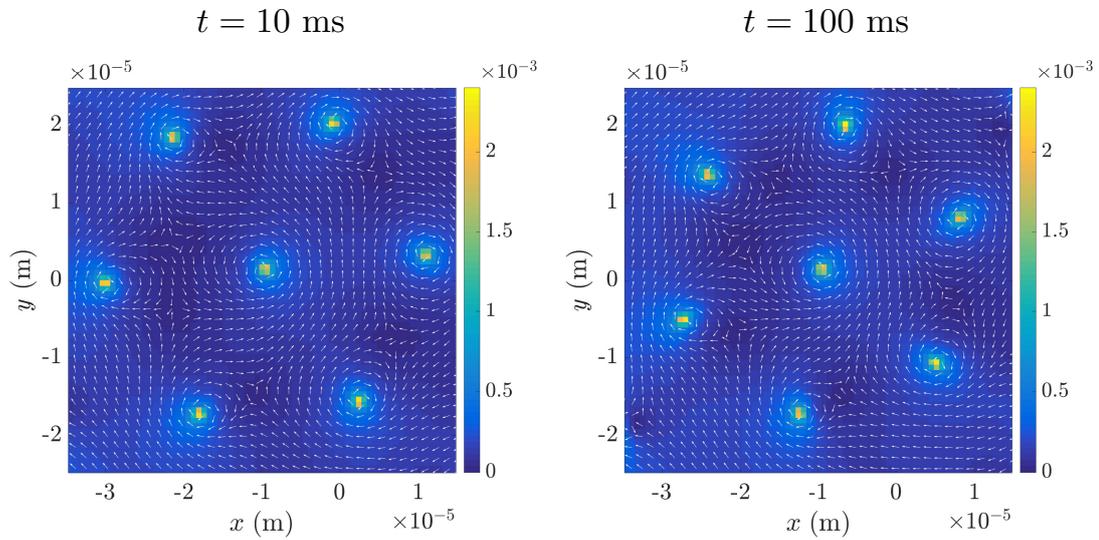


Figure 6.13: Velocity field of the vortex lattice following the flipping of a vortex to an antivortex for $t = (10, 100)$ ms. The change in the velocity field adjacent to the vortex causes the surrounding vortices to rotate at a much slower rate than the solid-body rotation of the vortex lattice. This in turn creates a large change in the resulting ordering of the lattice. The unbalanced local field will eventually lead to the antivortex moving and annihilating with a nearby vortex.

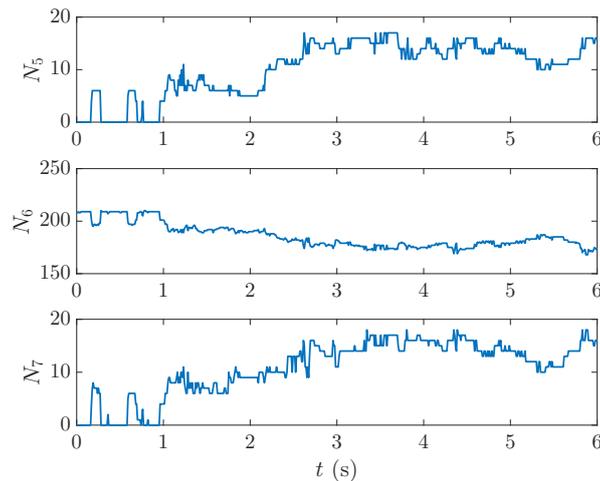


Figure 6.14: The defect (N_5 , N_7) and triangular lattice (N_6) nearest neighbour count taken from a Delaunay triangulation of the vortex lattice following the insertion of an antivortex. The number of defects increases as the local structure decays, and eventually gives rise to a quasi-constant state.

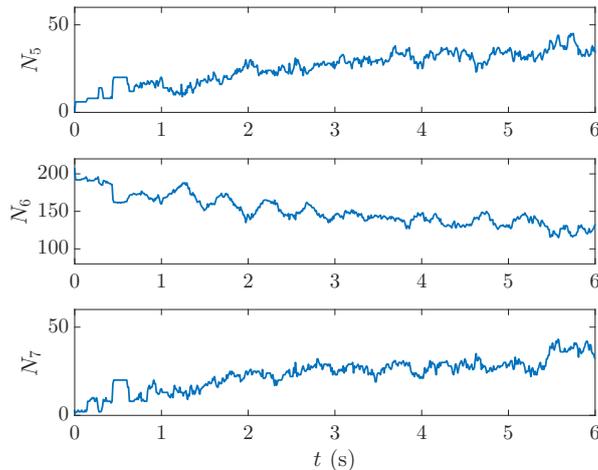


Figure 6.15: The defect (N_5 , N_7) and triangular lattice (N_6) nearest neighbour count from a Delaunay triangulation of the vortex lattice following the removal of 7 vortices from the centre of the lattice.

the two vortices at opposite sides of the condensate still yields reasonably high correlations ($\langle g_6(r) \rangle \approx 0.8$) at all times and length scales, indicating a well ordered lattice. Creating an antivortex in the lattice leads to lower correlations across all length scales, especially in the long time limit ($\langle g_6(r) \rangle \approx 0.7$), but still tends to the same long-ranged value as the previous case. This indicates an ordered lattice outside the region of the localized defects. Lastly, the removal of seven vortices shows a significant drop in correlations at all length scales and across both times ($\langle g_6(r) \rangle \approx 0.5$), indicating a global disordering of the vortices, which is consistent with the large number of defects identified earlier.

An alternative view of the above dynamics can be given from the Voronoi tessellation of the lattice. For the above cases of (a) removing a vortex from the centre, (b) removing 2 away from the centre, (c) flipping the rotation direction, and (d) removing 7, the ensuing evolutions are shown in Fig. 6.17, with the colours representing the vortex cell area. Snapshots are taken at times $t = (0.01, 0.1, 1, 6)$ s, with the mean area \bar{A} and standard deviation σ over time also shown. We can use this to observe the perturbation on the lattice, and see how it disturbs positions across the entire condensate. One can see that the perturbation initially affects only the vortices close to the defect. The Voronoi cells within this region have a much larger area than those of the surrounding regions. As the system evolves the defect disturbs the system outward from the initial imprint, with the strength of this disturbance being visible in the change of cell areas. As described earlier, the defect region maintains its honeycomb-like structure for up to $t = 1$ s. Rapid oscillations in \bar{A} and a peak in σ persist up to approximately $t = 1.5$ s, after which they settle into an oscillation at the breathing mode frequency $2\omega_\perp$ and a constant value respectively for cases (a), (b) and (c). The settling of these values is indicative of the lattice vacancy decaying and the lattice settling into a more disordered arrangement. For case (d) the removal affects the entire lattice, with the

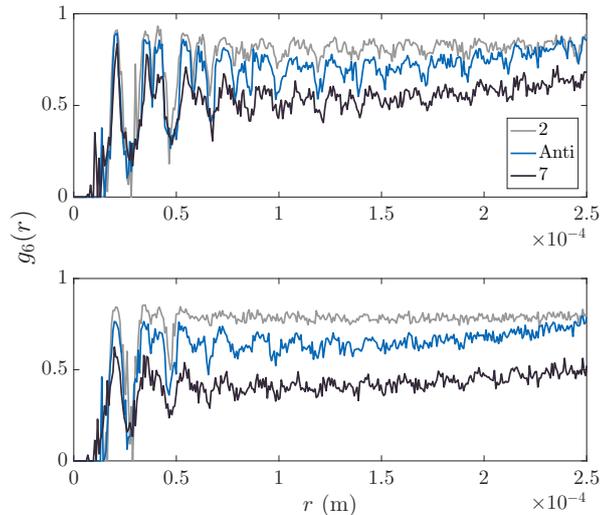


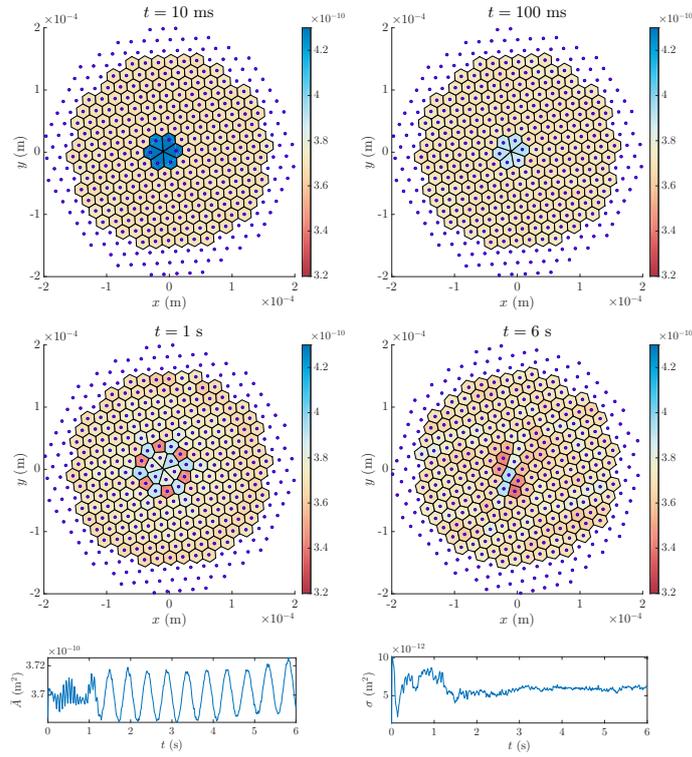
Figure 6.16: The orientational correlation function is given for moderate ($t = 3$ seconds, top) and long ($t = 6$ seconds, bottom) times after the phase imprint for removing 2, creating an antivortex, and removing 7 vortices respectively. The general behaviour at short and long ranges is similar for all three scenarios, but the correlations are significantly reduced, especially for the situation where seven vortices are removed.

finer details of settling into a new arrangement now completely overtaken by the large-scale disordering.

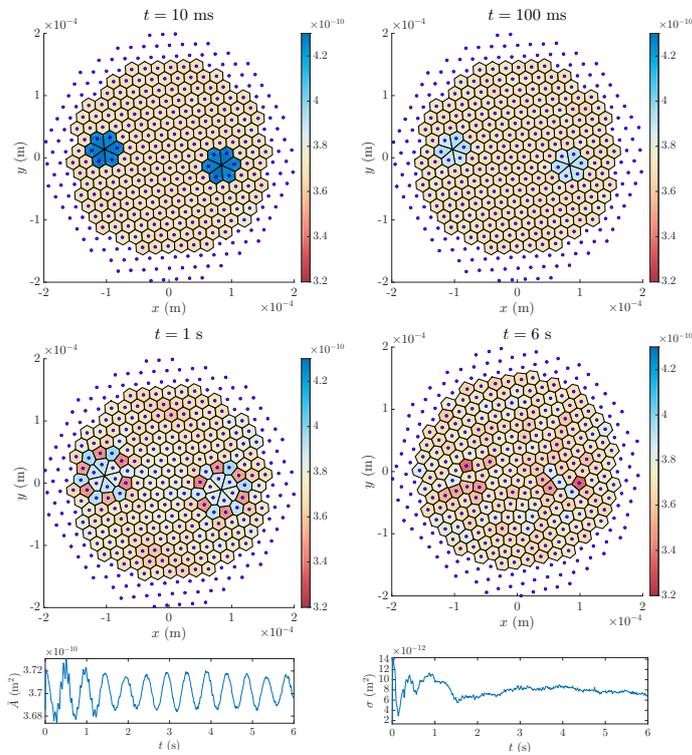
We can use the local orientational order parameter $g_6(0) = |\zeta_6(\mathbf{r}_0)|^2$ as the colour scale for the above examples. This is shown in Fig. 6.18, again with the mean-value and standard deviations over time given. As with the Delaunay triangulation, and the orientational correlations one can see that the effects of the perturbations are confined to the region close to the affected vortex site. The loss of triangular symmetry is observed for the cells surrounding the affected sites for (a) and (b), which show an almost identical profile for mean-value and standard deviation. For well separated sites the resulting effects on the lattice can be considered as being independent of one another, due to the local nature of the perturbation. For cases (c) and (d), the local disordering of the regions is much higher, with significant deviation from triangular lattice symmetry for $t = 1$ s and beyond. These above methods allow us to clearly identify the effects of the perturbations to the lattice, and are a way towards an understanding of the resulting non-equilibrium dynamics.

6.4 Discussion and conclusions

We have discussed the situation where controllable amounts of disorder can be created in a BEC vortex lattice through vortex removal or rotational direction flipping at predetermined positions. A single vortex removed from an Abrikosov vortex lattice via phase imprinting creates a quasi-stable honeycomb-like vacancy site. The removal of the associated velocity

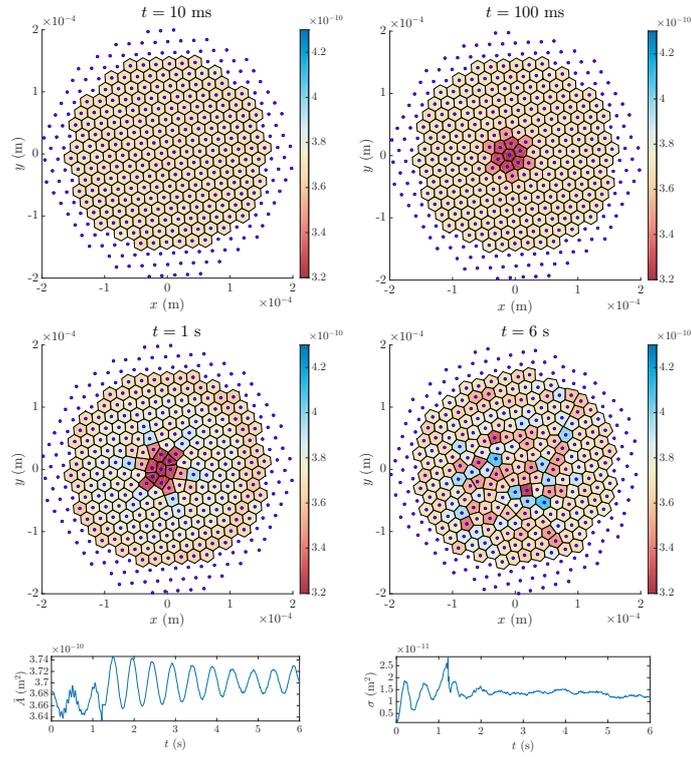


(a) Central vortex removed.

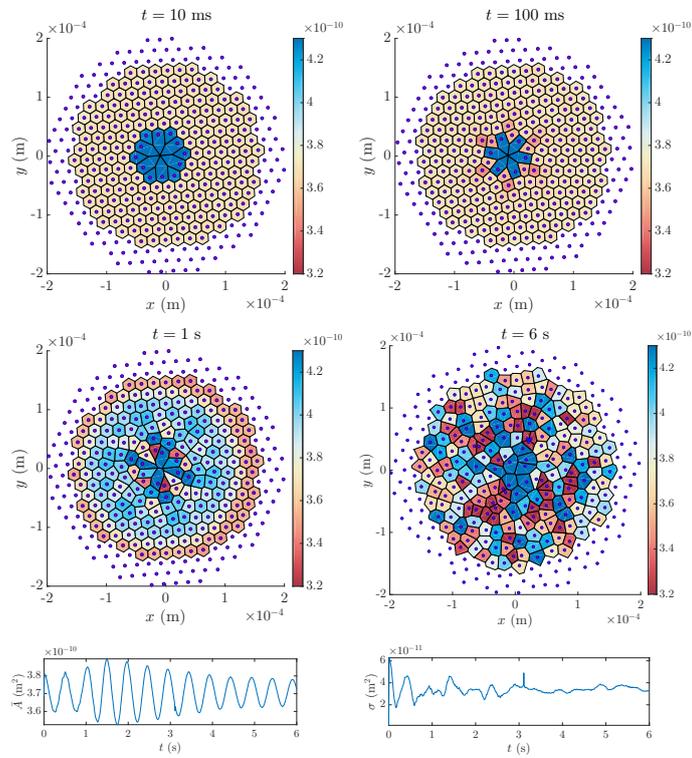


(b) Two vortices away from centre removed.

Figure 6.17: Voronoi diagram of a perturbed vortex lattice, with mean cell area \bar{A} and standard deviation σ following a phase imprint. The cell colour is indicative of the area spanned by the Voronoi cell.

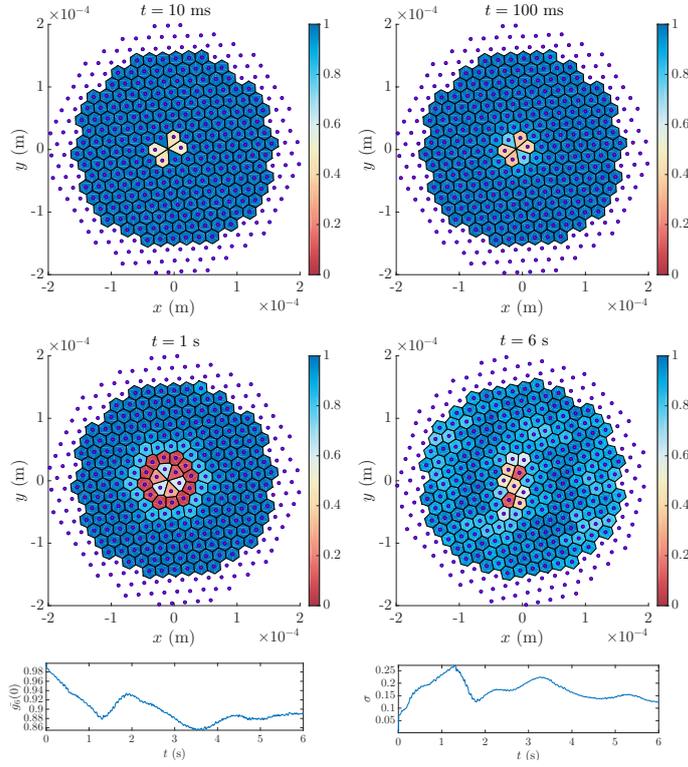


(c) Central vortex rotation direction flipped.

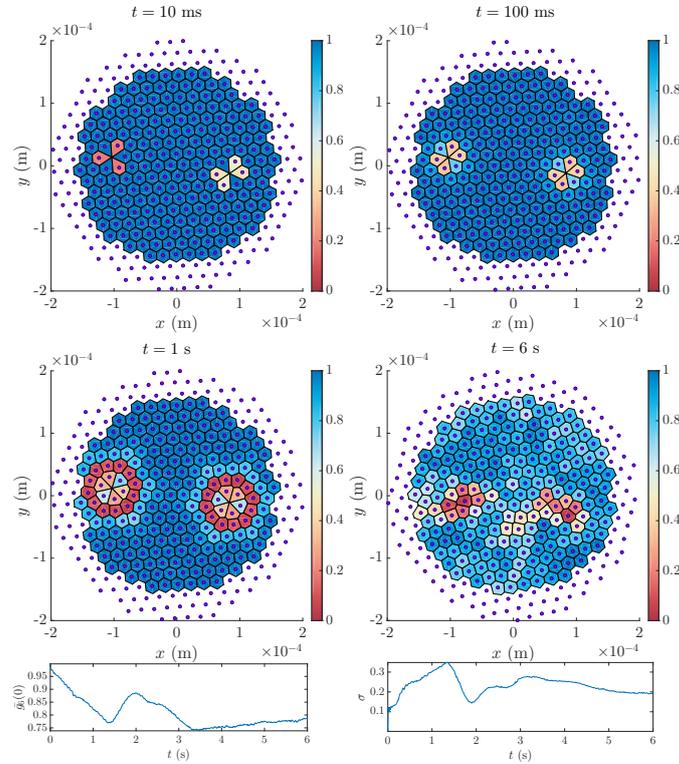


(d) Central 7 vortex cell removed.

Figure 6.17: Voronoi diagram of a perturbed vortex lattice, with mean cell area \bar{A} and standard deviation σ following a phase imprint. The cell colour is indicative of the area spanned by the Voronoi cell.

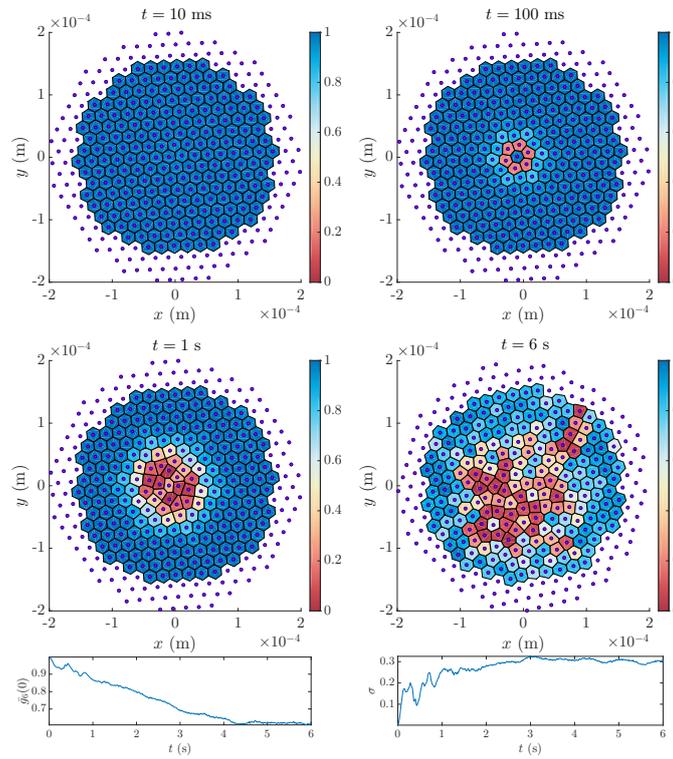


(a) Central vortex removed.

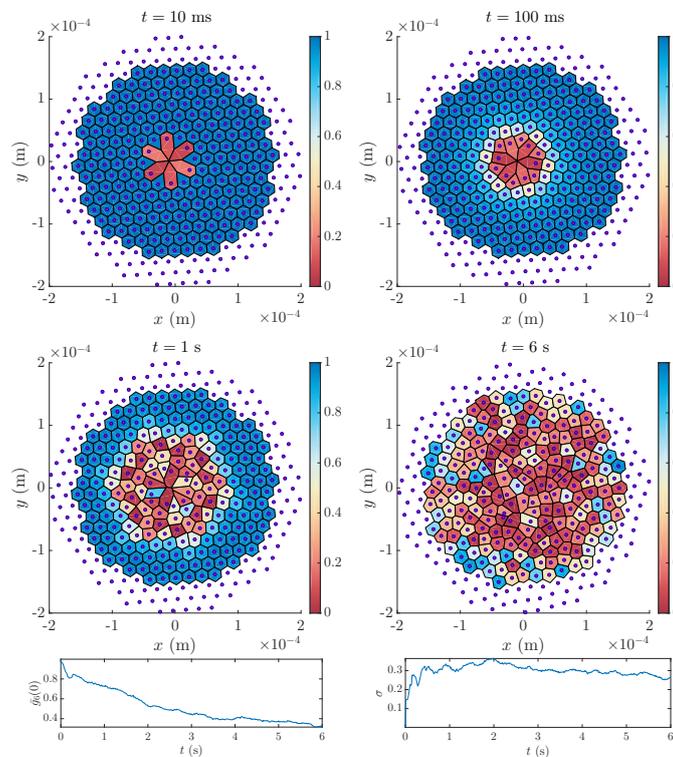


(b) Two vortices away from centre removed.

Figure 6.18: Voronoi diagram of a perturbed vortex lattice following a phase imprint. The cell colour is indicative of local orientational ordering of the vortex lattice $g_6(0) = |\zeta_6(\mathbf{r}_0)|^2$. The mean value of the local orientational order $\bar{g}_6(0)$ and standard deviation σ over time area also given.



(c) Central vortex rotation direction flipped.



(d) Central 7 vortex cell removed.

Figure 6.18: Voronoi diagram of a perturbed vortex lattice following a phase imprint. The cell colour is indicative of local orientational ordering of the vortex lattice $g_6(0) = |\zeta_6(\mathbf{r}_0)|^2$. The mean value of the local orientational order $\bar{g}_6(0)$ and standard deviation σ over time area also given.

field near the vacancy, however, disturbs the local solid-body behavior and the vacancy region rotates slower than the surrounding vortex lattice. It eventually decays, creating highly stable topological lattice defects due to the local rearrangement of the vortices, which persist for long times.

In fact, the resulting defects can be seen to pair, with (5,7) lattice defects being the most prominent, and manifesting themselves as dislocation defects in the lattice. Similar behaviour was observed for removing two separated vortices on the lattice, with the resulting defects being independent of one another. This lowered the correlation value, which indicated a drop in the ordering of the vortex lattice. Next, we examined the effect of introducing an antivortex into the lattice by flipping the phase profile of a pre-existing vortex. The effect of this was to further reduce the order of the vortex lattice, and settle into a lower correlated state than the above cases. Finally, we examined the removal of a large number of adjacent vortices in the lattice, which showed a significant drop in correlations, and hence ordering, over the course of evolution.

The characterization of perturbed lattices put forward by us complements the recent work of Rakonjac *et al.* [33], where the authors determine the disorder present in a vortex lattice in a BEC by comparing the ratio of the standard deviation of nearest neighbor distances to the mean distance. Here we extend the available tools by using orientational correlations, Delaunay triangulation for topological defect detection, Voronoi tessellation for identifying regions of modified orientational order, and by introducing a method to controllably engineer lattice defects through phase imprinting.

One possible use for the techniques we have discussed is to create vortex turbulence in low-dimensional condensate systems. Contrary to all currently existing discussions, this would make use of the phase erasing technique in Abrikosov lattices to examine turbulence starting from highly ordered systems [22, 159, 160]. While the use of vortex flipping has been considered previously [158], doing this in a highly controllable manner would be advantageous, as one could potentially also investigate phase transitions in the vortex system.

KTHNY (Kosterlitz, Thouless, Halperin, Nelson, Young) theory predicts the melting of two-dimensional systems as a two-step process, where the paired (5,7) defects dissociate, causing a loss of translational correlations, while maintaining quasi-long range orientational order [161–165]. This resulting phase is known as “hexatic”, and exists between the solid and isotropic liquid phases in two-dimensions. This phase may potentially be observed from the decay profile of correlations in this system. However, given the finite size of this system, observing the hexatic phase may be challenging, as one normally uses the combination of translational correlations as well as orientational correlations to identify this behaviour. In this instance, given the lack of applicability of translational correlations, one might instead opt to examine the structure factor

$$S(\mathbf{q}) = \frac{1}{N_v} \sum_{j,k=1}^{N_v} e^{-i\mathbf{q}\cdot(\mathbf{r}_j - \mathbf{r}_k)},$$

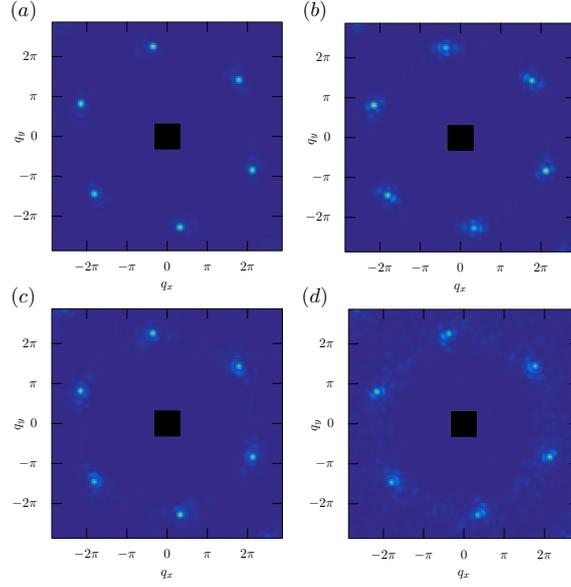


Figure 6.19: Structure factor at $t = 4$ s for (a) removing central vortex, (b) removing 2 separated vortices, (c) flipping vortex rotation, and (d) removing 7 adjacent vortices. The appearance of a smeared reciprocal peak between that of an amorphous fluid (ring) and a solid (peaks) at the lattice constant can be indicative of a hexatic phase in 2D materials. A continuous smearing of the structure factor peaks cannot be seen, making the determination of a hexatic phase inconclusive. The existence of this phase will require alternative methods to identify the phase, such as those proposed by [168, 169].

where N_v is the number of vortices, and $\mathbf{r}_{i,j}$ are the respective vortex positions. The structure factor can be examined for an arc-like profile, which is indicative of the presence of a hexatic phase [166, 167]. While this was briefly examined (see Fig. 6.19), conclusive results of such a phase were not initially found. This is likely due to the finite-size of the examined system.

A potential criteria for the onset of the hexatic phase can be given between the stretching and orientational order of nearest neighbouring vortex positions and a Lindemann parameter which quantifies the loss of both translational and orientational order [168]. Additionally, one can examine the statistics of the endpoints of defect strings, formed between pairs of dislocation (5,7) defects, which can be determined from small particle number systems [169]. While not examined in the preceding work, these methods can form the basis of a future investigation for the existence of the hexatic phase in the vortex lattice system.

Chapter 7

Conclusions and outlook

7.1 Conclusions

In this thesis we have presented our work on the dynamical behaviour of non-equilibrium Bose–Einstein condensates, where we have examined the behaviour of vortex lattices subjected to two distinct perturbations. We began by modelling the condensate using the mean-field Gross–Pitaevskii equation. Using this formalism we then discussed the superfluid properties of the condensate, and concentrated primarily on states with vorticity in two-dimensions. We primarily discussed high rotation rates of the condensate, where the condensate attains a large number of singly charged vortices arranged in a triangular Abrikosov vortex lattice. We restricted ourselves to a rotation rate of $\Omega = 0.995\omega_{\perp}$, and sought a numerical solution of the system with $N \approx 10^6$ atoms of ^{87}Rb .

We next introduced the algorithmic framework to numerically solve this system, making use of the Fourier split-operator method. With this we discussed imaginary time evolution to determine the vortex lattice ground state, and real time evolution for all subsequent dynamics. The simulation of this system was computationally challenging due to the finely sampled numerical grid required to resolve all features of the condensate in position and momentum space. To overcome this challenge, we introduced GPU computing methods, which were demonstrated for the problem of coherent atomic transport. For this we investigated a system using SAP methods with magnetic waveguides on atom-chips. These advanced computational techniques allowed for a fully three-dimensional simulation of the Schrödinger equation to be solved in significantly less time than a standard CPU implementation. The resulting GPU code was compared with a traditional MPI-enabled code, and showed equivalent performance to an 8-core 8-node cluster for the same system parameters. This led to the development of a software suite for the numerical solution of the Gross–Pitaevskii equation titled “GPUE”. An independently operated performance test of this suite was found to outperform other numerical softwares for the same class of problems.

Using the developed numerical suite, we performed a series of simulations for stationary and rotating condensates with a low number of vortices. We mentioned the necessary criteria for ensuring a well-ordered vortex lattice, and examined a condition where these criteria were unfulfilled. Following this, we introduced two distinctive perturbation techniques to disturb the the condensate, and allow non-equilibrium dynamics to be observed. The first method used a kicked optical potential, which modified the condensate phase. By matching the structure and lattice constants of both the optical and vortex lattices, the kick allowed for the generation of transient, time-varying superlattice structures in the density. These superlattice structures were observed during the subsequent dynamical evolution following the kick. By varying the alignment angle of the optical lattice relative to the vortices, we showed that the wavelength of the structures could be changed. The change in the structures were explained using moiré interference theory, and arose from the interference between the reciprocal lattice vectors of both the optical and vortex lattices. This was confirmed by examining the compressible kinetic energy spectrum of the condensate. The kicking perturbation showed how robust the vortex lattice was to density variations, with the phonons generated by the kick having little to no effect on the vortex positions.

As the vortex lattice proved to be very robust following the kicked potential, we next investigated methods to controllably create disorder in the vortex lattice. By directly phase imprinting topological excitations (phase singularities), we demonstrated that this was possible. From the well ordered vortex lattice ground state we annihilated or flipped the rotation direction of vortices at predefined positions in the lattice. As a vacancy was created in the vortex lattice following an annihilation, the remaining vortices attempted to redistribute and reorder to the most favourable position. Through extensive simulations, this was shown to create localised topological defects in the lattice, with the overall lattice still maintaining a large degree of order. Varying degrees of disorder were then created by removing additional vortices, or by flipping a vortex rotation profile. The use of Delaunay triangulation allowed us to easily identify the defect types, and largely showed the appearance of (5,7) topological lattice defects. By examining the orientational correlations of the lattice we observed that different imprints created varying degrees of lattice disordering. We then made use of Voronoi tessellations to allow local variations in lattice area and orientational correlations respectively to be identified following an annihilation, and demonstrated the effect the phase imprinting had on the vortex lattice on different timescales.

7.2 Outlook

Given the current state-of-the-art experimental control of condensate systems through use of SLMs, the perturbation methods discussed within this thesis are expected to be realisable. These perturbations represent two very useful techniques for quantum state control and engineering. For the kicked optical lattice, the creation of moiré interference patterns with

wavelengths much greater than the lattice spacing opens the possibility for detecting vortices without time-of-flight expansion in a lattice. We consider this technique to be a unique method for examining the periodicity of a lattice system, where the evolving pattern can also potentially be observed through the *in-situ* imaging techniques, as discussed in 2.2.1. Further extensions of this work can involve investigating the periodicity of large-scale soliton trains in quasi-1D condensates.

Some preliminary work in small-scale zig-zag and linear vortex crystals was carried out in conjunction with A. Barahmi and Th. Busch. This showed that with little periodicity in the system there were negligible peaks in the compressible energy spectrum. As a result, there were no discernible moiré superlattice patterns in the condensate density. It is expected that for these structures to be observed that highly periodic systems with a well defined reciprocal lattice are required. However, given a highly periodic system, any disordering of the system will affect the visibility of the peaks. As a result, this method could potentially allow for an examination of lattice disorder, and can form the basis of a future investigation.

The vortex annihilation/flipping through phase imprinting appears to be a very good candidate to create varying degrees of disorder in a vortex lattice system. The analysis methods discussed and used for this work can easily be applied to real experimental data. A potential use for this is to create controllable routes towards quantum turbulence from a well-ordered system. While the examination presented focussed primarily on the use of phase profiles opposite to that of the lattice, the imprinting of like-signed vortices also remains an interesting choice. Forcing vortices into different locations in the lattice is potentially an additional method to create lattice dislocations, and hence, topological lattice defects. One might consider erasing and adding vortices at different locations to both create and remove topological lattice defects. This can form the basis for a memory storage technique in a quantum computing system. The applicability of this method can potentially be examined in a future work.

Additionally, one can also create multi-charged vortices in the condensate. The effect of the surrounding lattice on the resulting multi-charged vortex would be an interesting problem. One might expect the l -charge vortex decay to be suppressed if the energy to move the surrounding lattice vortices is greater than the energy to maintain the l -charged vortex. This was briefly investigated by examining the Bogoliubov-de Gennes solutions of the imprinted vortex lattice system, with the aim of observing if the resulting excitation modes were complex. These modes were, however, not found due to the numerical complexity of the problem, and it remains an open question if this suppression exists. This will be investigated in a future work.

While we briefly mentioned the search for a KTHNY hexatic phase transition in this system, this will require further examination. Future work can include an investigation for the existence of this transition, and examine whether dislocation mediated melting of the vortex lattice can occur as a result of the phase imprinting techniques. Though we consider

the framework developed and examined for all the above methods to be valid, the consideration of finite temperature effects would ensure that the investigated methods are truly physically realistic. For such finite temperature condensates, one might consider use of the Zaremba–Nikuni–Griffin (ZNG) formalism [170, 171], or the formalism of Billam *et al.* [172]. An extension of the above works can examine this.

The use of GPU computing for simulating quantum dynamics is currently an under-utilised paradigm. The potential for a significant performance gain exists, given an effective mapping of a numerical algorithm to the GPU hardware. While the code developed and utilised for all the above simulations offers a clear performance advantage, it should be noted that further development and maintenance of such code can be challenging. Rapid changes to the CUDA programming models have introduced many new features to the standard which could potentially be used for solving more complex problems of both linear and nonlinear Schrödinger-type problems. However, such changes often require training, software rewrites, or newer hardware to take advantage of these. An extension of the GPUE codebase to cover one and three dimensional Gross–Pitaevskii systems will allow for this suite to be as feature rich as the currently most capable suites available [107, 108], whilst still holding the current edge in performance. Solutions using arbitrary gauge fields for these problems will also offer a distinctive advantage. Additionally, the inclusion of a numerical BdG solver for the resulting numerical solutions will allow for this software to become a very general suite for BEC problems.

The methods and works examined in this thesis offer interesting answers, questions and possibilities for the future of controllable quantum systems and technologies.

Bibliography

- [1] T. Morgan, L. J. O’Riordan, N. Crowley, B. O’Sullivan, and T. Busch. *Coherent transport by adiabatic passage on atom chips*. Physical Review A, 88(5) (2013).
- [2] L. J. O’Riordan, A. C. White, and T. Busch. *Moire superlattice structures in kicked Bose-Einstein condensates*. Physical Review A, 93(2) (2016).
- [3] Lee James O’Riordan and Thomas Busch. *Topological defect dynamics of vortex lattices in Bose-Einstein condensates*. Phys. Rev. A, 94:053603 (2016). doi:10.1103/PhysRevA.94.053603. URL <http://link.aps.org/doi/10.1103/PhysRevA.94.053603>.
- [4] S. Chu. *The manipulation of neutral particles*. Reviews of Modern Physics, 70(3):685–706 (1998).
- [5] C. N. Cohen-Tannoudji. *Manipulating atoms with photons*. Reviews of Modern Physics, 70(3):707–719 (1998).
- [6] W. D. Phillips. *Laser cooling and trapping of neutral atoms*. Reviews of Modern Physics, 70(3):721–741 (1998).
- [7] C. J. Foot. *Atomic physics*. Oxford master series in physics. Oxford University Press, Oxford ; New York (2005).
- [8] Harold J. Metcalf and Peter Van der Straten. *Laser cooling and trapping*. Graduate texts in contemporary physics. Springer, New York (1999).
- [9] A Einstein. *Quantentheorie des einatomigen idealen Gases*. Sitzungsberichte der Preussischen Akademie der Wissenschaften, 1 (1925).
- [10] L. P. Pitaevskii and S. Stringari. *Bose-Einstein Condensation* (2003).
- [11] Masahito Ueda. *Fundamentals and New Frontiers of Bose-Einstein Condensation* (2010).
- [12] Oliver Penrose and Lars Onsager. *Bose-Einstein Condensation and Liquid Helium*. Physical Review, 104(3):576–584 (1956).

-
- [13] Daniel Kleppner, Thomas C. Killian Thomas J. Greytak, Dale G. Fried, Lorenz Willmann, David Landhuis, and Stephen C. Moss. *Bose-Einstein Condensation of Atomic Hydrogen*. Proceedings of the International School of Physics “Enrico Fermi”, Course 140, 1998 (1998).
- [14] D. G. Fried, T. C. Killian, L. Willmann, D. Landhuis, S. C. Moss, D. Kleppner, and T. J. Greytak. *Bose-Einstein condensation of atomic hydrogen*. Physical Review Letters, 81(18):3811–3814 (1998).
- [15] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. *Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor*. Science, 269(5221):198–201 (1995).
- [16] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. Vandrueten, D. S. Durfee, D. M. Kurn, and W. Ketterle. *Bose-Einstein Condensation in a Gas of Sodium Atoms*. Physical Review Letters, 75(22):3969–3973 (1995).
- [17] J. Walraven. *Elements of Quantum Gases Collisions and Statistics - Master course*. Les Houches Cold-Atoms PreDoc School 2013: Manipulation of quantum degenerate gases (2013).
- [18] H. Pethick, C. J.; Smith. *Bose-Einstein Condensation in Dilute Gases, 2E*. (2008).
- [19] A. Gorlitz, J. M. Vogels, A. E. Leanhardt, C. Raman, T. L. Gustavson, J. R. Abo-Shaeer, A. P. Chikkatur, S. Gupta, S. Inouye, T. Rosenband, and W. Ketterle. *Realization of Bose-Einstein condensates in lower dimensions*. Physical Review Letters, 87(13) (2001).
- [20] S. Stock, Z. Hadzibabic, B. Battelier, M. Cheneau, and J. Dalibard. *Observation of phase defects in quasi-two-dimensional Bose-Einstein condensates*. Physical Review Letters, 95(19) (2005).
- [21] T. W. Neely, E. C. Samson, A. S. Bradley, M. J. Davis, and B. P. Anderson. *Observation of Vortex Dipoles in an Oblate Bose-Einstein Condensate*. Physical Review Letters, 104(16) (2010).
- [22] W. J. Kwon, G. Moon, J. Y. Choi, S. W. Seo, and Y. I. Shin. *Relaxation of superfluid turbulence in highly oblate Bose-Einstein condensates*. Physical Review A, 90(6) (2014).
- [23] S. W. Seo, J. Y. Choi, and Y. I. Shin. *Free expansion of quasi-2D Bose-Einstein condensates with quantized vortices*. Journal of the Korean Physical Society, 64(1):53–57 (2014).
- [24] L. Chomaz, L. Corman, T. Bienaime, R. Desbuquois, C. Weitenberg, S. Nascimbene, J. Beugnon, and J. Dalibard. *Emergence of coherence via transverse condensation in a uniform quasi-two-dimensional Bose gas*. Nature Communications, 6 (2015).

-
- [25] R. Srinivasan. *Vortices in Bose-Einstein condensates: A review of the experimental results*. Pramana-Journal of Physics, 66(1):3–30 (2006).
- [26] R. E. Packard. *Vortex Photography in Liquid-Helium*. Physica B & C, 109(1-3):1474–1484 (1982).
- [27] R. E. Packard and T. M. Sanders. *Detection of Single Quantized Vortex Lines in Rotating He2*. Physical Review Letters, 22(16):823– (1969).
- [28] M. Tsubota, K. Kasamatsu, and M. Kobayashi. *Quantized vortices in superfluid helium and atomic Bose-Einstein condensates*. arXiv, 1004.5458 (2010).
- [29] Wei Guo, Marco La Mantia, Daniel P. Lathrop, and Steven W. Van Sciver. *Visualization of two-fluid flows of superfluid helium-4*. Proceedings of the National Academy of Sciences (2014).
- [30] Makoto Tsubota. *Quantum Turbulence*. Journal of the Physical Society of Japan, 77(11):111006 (2008).
- [31] T. Isoshima and K. Machida. *Vortex stabilization in Bose-Einstein condensate of alkali-metal atom gas*. Physical Review A, 59(3):2203–2212 (1999).
- [32] C. Raman, J. R. Abo-Shaeer, J. M. Vogels, K. Xu, and W. Ketterle. *Vortex Nucleation in a Stirred Bose-Einstein Condensate*. Physical Review Letters, 87(21):210402 (2001).
- [33] A. Rakonjac, A. L. Marchant, T. P. Billam, J. L. Helm, M. M. H. Yu, S. A. Gardiner, and S. L. Cornish. *Measuring the disorder of vortex lattices in a Bose-Einstein condensate*. Physical Review A, 93(1) (2016).
- [34] D. V. Freilich, D. M. Bianchi, A. M. Kaufman, T. K. Langin, and D. S. Hall. *Real-Time Dynamics of Single Vortex Lines and Vortex Dipoles in a Bose-Einstein Condensate*. Science, 329(5996):1182–1185 (2010).
- [35] K. E. Wilson, Z. L. Newman, J. D. Lowney, and B. P. Anderson. *In situ imaging of vortices in Bose-Einstein condensates*. Physical Review A, 91(2) (2015).
- [36] Y. Kawaguchi and T. Ohmi. *Splitting instability of a multiply charged vortex in a Bose-Einstein condensate*. Physical Review A, 70(4) (2004).
- [37] P. O. Fedichev and G. V. Shlyapnikov. *Dissipative dynamics of a vortex state in a trapped Bose-condensed gas*. Physical Review A, 60(3):R1779–R1782 (1999).
- [38] D. L. Feder, C. W. Clark, and B. I. Schneider. *Vortex stability of interacting Bose-Einstein condensates confined in anisotropic harmonic traps*. Physical Review Letters, 82(25):4956–4959 (1999).
- [39] J. R. Anglin and W. H. Zurek. *Vortices in the wake of rapid Bose-Einstein condensation*. Physical Review Letters, 83(9):1707–1710 (1999).

- [40] B. M. Caradoc-Davies, R. J. Ballagh, and K. Burnett. *Coherent dynamics of vortex formation in trapped Bose-Einstein condensates*. Physical Review Letters, 83(5):895–898 (1999).
- [41] R. J. Marshall, G. H. C. New, K. Burnett, and S. Choi. *Exciting, cooling, and vortex trapping in a Bose-condensed gas*. Physical Review A, 59(3):2085–2093 (1999).
- [42] Ł Dobrek, M. Gajda, M. Lewenstein, K. Sengstock, G. Birkl, and W. Ertmer. *Optical generation of vortices in trapped Bose-Einstein condensates*. Physical Review A, 60(5):R3381–R3384 (1999).
- [43] M. Nakahara, T. Isoshima, K. Machida, S. Ogawa, and T. Ohmi. *A simple method to create a vortex in Bose-Einstein condensate of alkali atoms*. Physica B, 284:17–18 (2000).
- [44] B. M. Caradoc-Davies. *Vortex Dynamics in Bose-Einstein Condensates*. Ph.D. thesis, University of Otago (2000).
- [45] F. Chevy, K. W. Madison, and J. Dalibard. *Measurement of the angular momentum of a rotating Bose-Einstein condensate*. Physical Review Letters, 85(11):2223–2227 (2000).
- [46] N. R. Cooper, N. K. Wilkin, and J. M. F. Gunn. *Quantum phases of vortices in rotating Bose-Einstein condensates*. Physical Review Letters, 87(12) (2001).
- [47] P. Rosenbusch, V. Bretin, and J. Dalibard. *Dynamics of a single vortex line in a Bose-Einstein condensate*. Physical Review Letters, 89(20) (2002).
- [48] S. I. Ogawa, M. Mottonen, M. Nakahara, T. Ohmi, and H. Shimada. *Method to create a vortex in a Bose-Einstein condensate*. Physical Review A, 66(1) (2002).
- [49] V. Bretin, P. Rosenbusch, and J. Dalibard. *Dynamics of a single vortex line in a Bose-Einstein condensate*. Journal of Optics B-Quantum and Semiclassical Optics, 5(2):S23–S28 (2003).
- [50] K. W. Madison, F. Chevy, W. Wohlleben, and J. Dalibard. *Vortex Formation in a Stirred Bose-Einstein Condensate*. Physical Review Letters, 84(5):806–809 (2000).
- [51] K. W. Madison, F. Chevy, W. Wohlleben, and J. Dalibard. *Vortices in a stirred Bose-Einstein condensate*. Journal of Modern Optics, 47(14-15):2715–2723 (2000).
- [52] F. Chevy, K. Madison, V. Bretin, and J. Dalibard. *Nucleation of quantized vortices in a gaseous Bose-Einstein condensate*. Comptes Rendus De L Academie Des Sciences Serie Iv Physique Astrophysique, 2(4):663–669 (2001).
- [53] K. W. Madison, F. Chevy, V. Bretin, and J. Dalibard. *Stationary states of a rotating Bose-Einstein condensate: Routes to vortex nucleation*. Physical Review Letters, 86(20):4443–4446 (2001).

-
- [54] M. Mottonen, N. Matsumoto, M. Nakahara, and T. Ohmi. *Continuous creation of a vortex in a Bose-Einstein condensate with hyperfine spin $F=2$* . Journal of Physics-Condensed Matter, 14(49):13481–13491 (2002).
- [55] J. R. Abo-Shaeer, C. Raman, J. M. Vogels, and W. Ketterle. *Observation of Vortex Lattices in Bose-Einstein Condensates*. Science, 292(5516):476–479 (2001).
- [56] V. Tkachenko. *Stability of Vortex Lattices*. Soviet Physics JETP-USSR, 23(6):1049– (1966).
- [57] P. Engels, I. Coddington, P. C. Haljan, and E. A. Cornell. *Nonequilibrium effects of anisotropic compression applied to vortex lattices in Bose-Einstein condensates*. Physical Review Letters, 89(10) (2002).
- [58] T. L. Ho. *Bose-Einstein condensates with large number of vortices*. Physical Review Letters, 87(6) (2001).
- [59] V. Bretin, S. Stock, Y. Seurin, and J. Dalibard. *Fast rotation of a Bose-Einstein condensate*. Physical Review Letters, 92(5) (2004).
- [60] T. K. Ghosh. *Vortex formation in a fast rotating Bose-Einstein condensate*. Physical Review A, 69(4) (2004).
- [61] N. Regnault and T. Jolicoeur. *Quantum Hall fractions in rotating Bose-Einstein condensates*. Physical Review Letters, 91(3) (2003).
- [62] Alexander L. Fetter. *Rotating trapped Bose-Einstein condensates*. Reviews of Modern Physics, 81(2):647–691 (2009).
- [63] H. Zhai, Q. Zhou, R. Lu, and L. Chang. *Double-layer Bose-Einstein condensates with a large number of vortices*. Physical Review A, 69(6) (2004).
- [64] S. Stock, B. Battelier, V. Bretin, Z. Hadzibabic, and J. Dalibard. *Bose-Einstein condensates in fast rotation*. Laser Physics Letters, 2(6):275–284 (2005).
- [65] A. Aftalion, X. Blanc, and J. Dalibard. *Vortex patterns in a fast rotating Bose-Einstein condensate*. Physical Review A, 71(2) (2005).
- [66] V. Schweikhard, I. Coddington, P. Engels, V. P. Mogendorff, and E. A. Cornell. *Rapidly Rotating Bose-Einstein Condensates in and near the Lowest Landau Level*. Physical Review Letters, 92(4):040404 (2004).
- [67] Gentaro Watanabe, Gordon Baym, and C. J. Pethick. *Landau Levels and the Thomas-Fermi Structure of Rapidly Rotating Bose-Einstein Condensates*. Phys. Rev. Lett., 93:190401 (2004). doi:10.1103/PhysRevLett.93.190401. URL <http://link.aps.org/doi/10.1103/PhysRevLett.93.190401>.
- [68] D. E. Sheehy and L. Radzihovsky. *Vortex lattice inhomogeneity in spatially inhomogeneous superfluids*. Physical Review A, 70(5) (2004).

- [69] D. E. Sheehy and L. Radzihovsky. *Vortices in spatially inhomogeneous superfluids*. Physical Review A, 70(6) (2004).
- [70] N. R. Cooper, S. Komineas, and N. Read. *Vortex lattices in the lowest Landau level for confined Bose-Einstein condensates*. Phys. Rev. A, 70:033604 (2004). doi: 10.1103/PhysRevA.70.033604. URL <http://link.aps.org/doi/10.1103/PhysRevA.70.033604>.
- [71] I. Coddington, P. C. Haljan, P. Engels, V. Schweikhard, S. Tung, and E. A. Cornell. *Experimental studies of equilibrium vortex properties in a Bose-condensed gas*. Physical Review A, 70(6) (2004).
- [72] I.; McKay Parry N.; Baker M.; Davis M. J.; Rubinsztein-Dunlop H.; Neely T. W. Gauthier, G.; Lenton. *Direct imaging of a digital-micromirror device for configurable microscopic optical potentials* (2016).
- [73] P. B. Wigley, P. J. Everitt, A. van den Hengel, J. W. Bastian, M. A. Sooriyabandara, G. D. McDonald, K. S. Hardman, C. D. Quinlivan, P. Manju, C. C. N. Kuhn, I. R. Petersen, A. N. Luiten, J. J. Hope, N. P. Robins, and M. R. Hush. *Fast machine-learning online optimization of ultra-cold-atom experiments*. Scientific Reports, 6 (2016).
- [74] Werner Krauth. *Statistical Mechanics: Algorithms and computations*. Oxford Master Series in Physics (2006).
- [75] U. Schollwock. *The density-matrix renormalization group*. Reviews of Modern Physics, 77(1):259–315 (2005).
- [76] P. Bader, S. Blanes, and F. Casas. *Solving the Schrodinger eigenvalue problem by the imaginary time propagation technique using splitting methods with complex coefficients*. Journal of Chemical Physics, 139(12) (2013).
- [77] Ionut Danaila. *Three-dimensional vortex structure of a fast rotating Bose-Einstein condensate with harmonic-plus-quartic confinement*. Physical Review A, 72(1) (2005).
- [78] Heiko Bauke and Christoph H. Keitel. *Accelerating the Fourier split operator method via graphics processing units*. Computer Physics Communications, 182(12):2454–2463 (2011).
- [79] M. Weyrauch and D. Scholz. *Computing the Baker-Campbell-Hausdorff series and the Zassenhaus product*. Computer Physics Communications, 180(9):1558–1565 (2009).
- [80] V. Gradinaru. *Strang Splitting for the Time-Dependent Schrodinger Equation on Sparse Grids*. Siam Journal on Numerical Analysis, 46(1):103–123 (2007).
- [81] Juha Javanainen and Janne Ruostekoski. *Symbolic calculation in development of algorithms: split-step methods for the Gross-Pitaevskii equation*. Journal of Physics A: Mathematical and General, 39(12):L179–L184 (2006).

-
- [82] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press (2007).
- [83] Intel. *Intel® Xeon® Processor E5-2697 v2 Product Brief*. URL <http://ark.intel.com/products/75283/>.
- [84] N. J. Higham. *The Accuracy of Floating-Point Summation*. Siam Journal on Scientific Computing, 14(4):783–799 (1993).
- [85] John Nickolls, Ian Buck, Michael Garland, and Kevin Skadron. *Scalable Parallel Programming with CUDA*. Queue, 6(2):40–53 (2008).
- [86] David Kirk and Wen-mei Hwu. *Programming massively parallel processors hands-on with CUDA*. Morgan Kaufmann Publishers, Burlington, MA (2010).
- [87] R. Menchon-Enrich, A. Benseny, V. Ahufinger, A. D. Greentree, Busch Th, and J. Mompert. *Spatial adiabatic passage: a review of recent progress*. Reports on Progress in Physics, 79(7):074401 (2016).
- [88] K. Bergmann, N. V. Vitanov, and B. W. Shore. *Perspective: Stimulated Raman adiabatic passage: The status after 25 years*. Journal of Chemical Physics, 142(17) (2015).
- [89] K. Eckert, M. Lewenstein, R. Corbalan, G. Birkl, W. Ertmer, and J. Mompert. *Three-level atom optics via the tunneling interaction*. Physical Review A, 70(2) (2004).
- [90] Y. Takahashi S. Taie. *Conference talk*. In *Mini Symposium: Spatial Adiabatic Passage*. Okinawa Institute of Science and Technology Graduate University, Japan (2016).
- [91] K. Eckert, J. Mompert, R. Corbalan, M. Lewenstein, and G. Birkl. *Three level atom optics in dipole traps and waveguides*. Optics Communications, 264(2):264–270 (2006).
- [92] T. Morgan, B. O’Sullivan, and T. Busch. *Coherent adiabatic transport of atoms in radio-frequency traps*. Physical Review A, 83(5) (2011).
- [93] J. Huneke, G. Platero, and S. Kohler. *Steady-State Coherent Transfer by Adiabatic Passage*. Physical Review Letters, 110(3) (2013).
- [94] N. V. Vitanov, T. Halfmann, B. W. Shore, and K. Bergmann. *Laser-induced population transfer by adiabatic passage techniques*. Annual Review of Physical Chemistry, 52:763–809 (2001).
- [95] M. P. Fewell, B. W. Shore, and K. Bergmann. *Coherent population transfer among three states: Full algebraic solutions and the relevance of non adiabatic processes to transfer by delayed pulses*. Australian Journal of Physics, 50(2):281–308 (1997).
- [96] B. O’Sullivan, P. Morrissey, T. Morgan, and Th Busch. *Using adiabatic coupling techniques in atom-chip waveguide structures*. Physica Scripta, T140 (2010).

- [97] M. Bartenstein, D. Cassettari, T. Calarco, A. Chenet, R. Folman, K. Brugger, A. Haase, E. Hartungen, B. Hessmo, A. Kasper, P. Kruger, T. Maier, F. Payr, S. Schneider, and J. Schmiedmayer. *Atoms and wires: Toward atom chips*. Ieee Journal of Quantum Electronics, 36(12):1364–1377 (2000).
- [98] R. Folman, P. Kruger, D. Cassettari, B. Hessmo, T. Maier, and J. Schmiedmayer. *Controlling cold atoms using nanofabricated surfaces: Atom chips*. Physical Review Letters, 84(20):4749–4752 (2000).
- [99] J. Denschlag, D. Cassettari, and J. Schmiedmayer. *Guiding neutral atoms with a wire*. Physical Review Letters, 82(10):2014–2017 (1999).
- [100] D. M. Brink and C. V. Sukumar. *Majorana spin-flip transitions in a magnetic trap*. Physical Review A, 74(3) (2006).
- [101] M. Yun and J. P. Yin. *Practical scheme to realize 2D array of BECs on an atom chip: novel 2D magneto-optical and magnetic lattices*. Optics Express, 14(7):2539–2551 (2006).
- [102] D. Gallego, S. Hofferberth, T. Schumm, P. Kruger, and J. Schmiedmayer. *Optical lattice on an atom chip*. Optics Letters, 34(22):3463–3465 (2009).
- [103] G. Bensky, R. Amsuss, J. Majer, D. Petrosyan, J. Schmiedmayer, and G. Kurizki. *Controlling quantum information processing in hybrid systems on chips*. Quantum Information Processing, 10(6):1037–1060 (2011).
- [104] Irish Center for High End Computing. *Decomissioned systems: Stoney, July 2009 - October, 2014*. URL <https://www.ichec.ie/infrastructure/previous>.
- [105] Lee James O’Riordan. *GPUE: GPU Gross Pitaevskii Equation solver*. URL <https://github.com/mlxd/gpue>.
- [106] Peter Wittek. *Comparing three numerical solvers of the Gross-Pitaevskii equation* (2016). URL <http://peterwittek.com/gpe-comparison.html>.
- [107] P. Wittek and F. M. Cucchietti. *A second-order distributed Trotter-Suzuki solver with a hybrid CPU-GPU kernel*. Computer Physics Communications, 184(4):1165–1171 (2013).
- [108] X. Antoine and R. Duboscq. *GPELab, a Matlab toolbox to solve Gross-Pitaevskii equations I: Computation of stationary solutions*. Computer Physics Communications, 185(11):2969–2991 (2014).
- [109] X. Antoine and R. Duboscq. *GPELab, a Matlab toolbox to solve Gross-Pitaevskii equations II: Dynamics and stochastic simulations*. Computer Physics Communications, 193:95–117 (2015).

-
- [110] Erich J. Mueller and Tin-Lun Ho. *Two-Component Bose-Einstein Condensates with a Large Number of Vortices*. Physical Review Letters, 88(18):180403 (2002).
- [111] Chris Foster. *vortutils: Utilities for finding vortices in complex fields*. GitHub (2013). URL <https://github.com/c42f/vortutils>.
- [112] A. E. Leanhardt, A. Gorlitz, A. P. Chikkatur, D. Kielpinski, Y. Shin, D. E. Pritchard, and W. Ketterle. *Imprinting vortices in a Bose-Einstein condensate using topological phases*. Physical Review Letters, 89(19) (2002).
- [113] Y. Kawaguchi, M. Nakahara, and T. Ohmi. *Topological vortex formation in a Bose-Einstein condensate under gravitational field*. Physical Review A, 70(4) (2004).
- [114] Takeshi Kuwamoto, Hayato Usuda, Satoshi Tojo, and Takuya Hirano. *Dynamics of Quadruply Quantized Vortices in 87Rb Bose-Einstein Condensates Confined in Magnetic and Optical Traps*. Journal of the Physical Society of Japan, 79(3):034004 (2010).
- [115] S. Masuda, U. Gungordu, X. Chen, T. Ohmi, and M. Nakahara. *Fast control of topological vortex formation in Bose-Einstein condensates by counterdiabatic driving*. Physical Review A, 93(1) (2016).
- [116] Jean Dalibard, Fabrice Gerbier, Gediminas Juzeliūnas, and Patrik Öhberg. *Colloquium : Artificial gauge potentials for neutral atoms*. Reviews of Modern Physics, 83(4):1523–1543 (2011).
- [117] J. L. Roberts, N. R. Claussen, James P. Burke, Chris H. Greene, E. A. Cornell, and C. E. Wieman. *Resonant Magnetic Field Control of Elastic Scattering in Cold ^{85}Rb* . Physical Review Letters, 81(23):5109–5112 (1998).
- [118] Lee James O’Riordan. *BEC vortex lattice generation* (2014). URL <http://goo.gl/30Mj1q>.
- [119] G. Watanabe, S. A. Gifford, G. Baym, and C. J. Pethick. *Structure of vortices in rotating Bose-Einstein condensates*. Physical Review A, 74(6) (2006).
- [120] I. Bloch, J. Dalibard, and W. Zwerger. *Many-body physics with ultracold gases*. Reviews of Modern Physics, 80(3):885–964 (2008).
- [121] M. Greiner, O. Mandel, T. Esslinger, T. W. Hansch, and I. Bloch. *Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms*. Nature, 415(6867):39–44 (2002).
- [122] J. W. Reijnders and R. A. Duine. *Pinning of vortices in a Bose-Einstein condensate by an optical lattice*. Physical Review Letters, 93(6) (2004).
- [123] A. S. Sorensen, E. Demler, and M. D. Lukin. *Fractional quantum hall states of atoms in optical lattices*. Physical Review Letters, 94(8) (2005).

- [124] S. Tung, V. Schweikhard, and E. A. Cornell. *Observation of Vortex Pinning in Bose-Einstein Condensates*. Physical Review Letters, 97(24) (2006).
- [125] P. Vignolo, R. Fazio, and M. P. Tosi. *Quantum vortices in optical lattices*. Physical Review A, 76(2) (2007).
- [126] G. Watanabe, B. P. Venkatesh, and R. Dasgupta. *Nonlinear Phenomena of Ultracold Atomic Gases in Optical Lattices: Emergence of Novel Features in Extended States*. Entropy, 18(4) (2016).
- [127] J. Denschlag, J. E. Simsarian, D. L. Feder, C. W. Clark, L. A. Collins, J. Cubizolles, L. Deng, E. W. Hagley, K. Helmerson, W. P. Reinhardt, S. L. Rolston, B. I. Schneider, and W. D. Phillips. *Generating solitons by phase engineering of a Bose-Einstein condensate*. Science, 287(5450):97–101 (2000).
- [128] Vladimir Tikhonenko, Jason Christou, Barry Luther-Davies, and Yuri S. Kivshar. *Observation of vortex solitons created by the instability of dark soliton stripes*. Opt. Lett., 21(15):1129–1131 (1996). doi:10.1364/OL.21.001129. URL <http://ol.osa.org/abstract.cfm?URI=ol-21-15-1129>.
- [129] B. P. Anderson, P. C. Haljan, C. A. Regal, D. L. Feder, L. A. Collins, C. W. Clark, and E. A. Cornell. *Watching Dark Solitons Decay into Vortex Rings in a Bose-Einstein Condensate*. Phys. Rev. Lett., 86:2926–2929 (2001). doi:10.1103/PhysRevLett.86.2926. URL <http://link.aps.org/doi/10.1103/PhysRevLett.86.2926>.
- [130] J. Brand and W. P. Reinhardt. *Solitonic vortices and the fundamental modes of the "snake instability": Possibility of observation in the gaseous Bose-Einstein condensate*. Physical Review A, 65(4) (2002).
- [131] S. Moulder, S. Beattie, R. P. Smith, N. Tammuz, and Z. Hadzibabic. *Quantized supercurrent decay in an annular Bose-Einstein condensate*. Physical Review A, 86(1) (2012).
- [132] T. Busch, M. Kohl, T. Esslinger, and K. Molmer. *Transverse mode of an atom laser*. Physical Review A, 65(4) (2002).
- [133] C. Nore, M. Abid, and M. E. Brachet. *Kolmogorov Turbulence in Low-Temperature Superflows*. Physical Review Letters, 78(20):3896–3899 (1997).
- [134] C. Nore, M. Abid, and M. E. Brachet. *Decaying Kolmogorov turbulence in a model of superflow*. Physics of Fluids, 9(9):2644–2669 (1997).
- [135] Ashton S. Bradley and Brian P. Anderson. *Energy Spectra of Vortex Distributions in Two-Dimensional Quantum Turbulence*. Physical Review X, 2(4) (2012).

-
- [136] T. L. Horng, C. H. Hsueh, S. W. Su, Y. M. Kao, and S. C. Gou. *Two-dimensional quantum turbulence in a nonuniform Bose-Einstein condensate*. Physical Review A, 80(2) (2009).
- [137] L. Fallani, C. Fort, J. E. Lye, and M. Inguscio. *Bose-Einstein condensate in an optical lattice with tunable spacing: transport and static properties*. Optics Express, 13(11):4303–4313 (2005).
- [138] R. A. Williams, J. D. Pillet, S. Al-Assam, B. Fletcher, M. Shotter, and C. J. Foot. *Dynamic optical lattices: two-dimensional rotating and accordion lattices for ultracold atoms*. Optics Express, 16(21):16977–16983 (2008).
- [139] Y. Murata, V. Petrova, B. B. Kappes, A. Ebnonnasir, I. Petrov, Y. H. Xie, C. V. Ciobanu, and S. Kodambaka. *Moire Superstructures of Graphene on Faceted Nickel Islands*. Acs Nano, 4(11):6509–6514 (2010).
- [140] M. Yankowitz, J. M. Xue, D. Cormode, J. D. Sanchez-Yamagishi, K. Watanabe, T. Taniguchi, P. Jarillo-Herrero, P. Jacquod, and B. J. LeRoy. *Emergence of superlattice Dirac points in graphene on hexagonal boron nitride*. Nature Physics, 8(5):382–386 (2012).
- [141] M. Kobayashi and M. Tsubota. *Kolmogorov spectrum of superfluid turbulence: Numerical analysis of the Gross-Pitaevskii equation with a small-scale dissipation*. Physical Review Letters, 94(6) (2005).
- [142] M. Tsubota. *Quantum turbulence—from superfluid helium to atomic Bose-Einstein condensates*. Journal of Physics-Condensed Matter, 21(16) (2009).
- [143] A. C. White, N. P. Proukakis, and C. F. Barenghi. *Topological stirring of two-dimensional atomic Bose-Einstein condensates*. Journal of Physics: Conference Series, 544(1):012021 (2014).
- [144] A. Skaugen and L. Angheluta. *Vortex clustering and universal scaling laws in two-dimensional quantum turbulence*. Physical Review E, 93(3) (2016).
- [145] M. T. Reeves, T. P. Billam, B. P. Anderson, and A. S. Bradley. *Signatures of coherent vortex structures in a disordered two-dimensional quantum fluid*. Physical Review A, 89(5) (2014).
- [146] U. Frisch and A. N. Kolmogorov. *Turbulence : the legacy of A.N. Kolmogorov*. Cambridge University Press, Cambridge, Eng. ; New York (1995).
- [147] K. Hermann. *Periodic overlays and moire patterns: theoretical studies of geometric properties*. Journal of Physics-Condensed Matter, 24(31) (2012).

- [148] H. T. Blair, A. C. Welday, and K. C. Zhang. *Scale-invariant memory representations emerge from moiré interference between grid fields that produce theta oscillations: A computational model*. *Journal of Neuroscience*, 27(12):3211–3229 (2007).
- [149] T. Morgan and T. Busch. *Stability and dynamics of cross solitons in harmonically confined Bose-Einstein condensates*. *Physical Review A*, 88(6) (2013).
- [150] S. A. Gardiner, D. Jaksch, R. Dum, J. I. Cirac, and P. Zoller. *Nonlinear matter wave dynamics with a chaotic potential*. *Physical Review A*, 62(2) (2000).
- [151] T. Mithun, K. Porsezian, and B. Dey. *Disorder-induced vortex lattice melting in a Bose-Einstein condensate*. *Physical Review A*, 93(1) (2016).
- [152] J. R. Anglin and M. Crescimanno. *Inhomogeneous vortex matter*. arXiv (2002).
- [153] I. Guillamon, R. Cordoba, J. Sese, J. M. De Teresa, M. R. Ibarra, S. Vieira, and H. Suderow. *Enhancement of long-range correlations in a 2D vortex lattice by an incommensurate 1D disorder potential*. *Nature Physics*, 10(11):851–856 (2014).
- [154] J. F. S. Brachmann, W. S. Bakr, J. Gillen, A. Peng, and M. Greiner. *Inducing vortices in a Bose-Einstein condensate using holographically produced light beams*. *Optics Express*, 19(14):12984–12991 (2011).
- [155] L. P. Pitaevskii and A. Rosch. *Breathing modes and hidden symmetry of trapped-atoms in two dimensions*. *Physical Review A*, 55(2):R853–R856 (1997).
- [156] G. Watanabe. *Axial breathing mode in rapidly rotating Bose-Einstein condensates and uncertainty of the rotation velocity*. *Physical Review A*, 76(3) (2007).
- [157] J. P. Martikainen, K. A. Suominen, L. Santos, T. Schulte, and A. Sanpera. *Generation and evolution of vortex-antivortex pairs in Bose-Einstein condensates*. *Physical Review A*, 64(6) (2001).
- [158] E. J. M. Madarassy and C. F. Barenghi. *Disordered vortex arrays in a two-dimensional condensate*. *Geophysical and Astrophysical Fluid Dynamics*, 103(2-3):269–278 (2009).
- [159] T. W. Neely, A. S. Bradley, E. C. Samson, S. J. Rooney, E. M. Wright, K. J. H. Law, R. Carretero-Gonzalez, P. G. Kevrekidis, M. J. Davis, and B. P. Anderson. *Characteristics of Two-Dimensional Quantum Turbulence in a Compressible Superfluid*. *Physical Review Letters*, 111(23) (2013).
- [160] A. J. Groszek, T. P. Simula, D. M. Paganin, and K. Helmerson. *Onsager vortex formation in Bose-Einstein condensates in two-dimensional power-law traps*. *Physical Review A*, 93(4) (2016).
- [161] J. M. Kosterlitz and D. J. Thouless. *Ordering, metastability and phase transitions in two-dimensional systems*. *Journal of Physics C: Solid State Physics*, 6(7):1181 (1973). URL <http://stacks.iop.org/0022-3719/6/i=7/a=010>.

-
- [162] B. I. Halperin and David R. Nelson. *Theory of Two-Dimensional Melting*. Phys. Rev. Lett., 41:121–124 (1978). doi:10.1103/PhysRevLett.41.121. URL <http://link.aps.org/doi/10.1103/PhysRevLett.41.121>.
- [163] David R. Nelson and B. I. Halperin. *Dislocation-mediated melting in two dimensions*. Phys. Rev. B, 19:2457–2484 (1979). doi:10.1103/PhysRevB.19.2457. URL <http://link.aps.org/doi/10.1103/PhysRevB.19.2457>.
- [164] A. P. Young. *Melting and the vector Coulomb gas in two dimensions*. Phys. Rev. B, 19:1855–1866 (1979). doi:10.1103/PhysRevB.19.1855. URL <http://link.aps.org/doi/10.1103/PhysRevB.19.1855>.
- [165] Sebastian C. Kapfer and Werner Krauth. *Two-Dimensional Melting: From Liquid-Hexatic Coexistence to Continuous Transitions*. Phys. Rev. Lett., 114:035702 (2015). doi:10.1103/PhysRevLett.114.035702. URL <http://link.aps.org/doi/10.1103/PhysRevLett.114.035702>.
- [166] A. Brodin, A. Nych, U. Ognysta, B. Lev, V. Nazarenko, M. Skarabot, and I. Musevic. *Melting of 2D liquid crystal colloidal structure*. Condensed Matter Physics, 13(3) (2010).
- [167] X. Y. Sun, Y. Li, Y. Q. Ma, and Z. X. Zhang. *Direct observation of melting in a two-dimensional driven granular system*. Scientific Reports, 6 (2016).
- [168] G. M. Bruun and D. R. Nelson. *Quantum hexatic order in two-dimensional dipolar and charged fluids*. Physical Review B, 89(9) (2014).
- [169] W. Lechner, H. P. Buchler, and P. Zoller. *Role of Quantum Fluctuations in the Hexatic Phase of Cold Polar Molecules*. Physical Review Letters, 112(25) (2014).
- [170] E. Zaremba, T. Nikuni, and A. Griffin. *Dynamics of Trapped Bose Gases at Finite Temperatures*. A. Journal of Low Temperature Physics, 116(3):277 ((1999)).
- [171] Nick Proukakis, S. A. Gardiner, Matthew J. Davis, and Marzena Szymanska. *Quantum gases : finite temperatures and non-equilibrium dynamics*. Cold atoms. Imperial College Press ; Distributed by World Scientific, London Singapore ; Hackensack, NJ (2013).
- [172] T. P. Billam, P. Mason, and S. A. Gardiner. *Second-order number-conserving description of nonequilibrium dynamics in finite-temperature Bose-Einstein condensates*. Physical Review A, 87(3) (2013).