Ripples on a Quantum Interface

Waves, Instabilities, and Prethermalization in a Binary Bose–Einstein Condensate

by

Yanda Geng

Submitted to the Faculty of the Graduate School of the
University of Maryland, College Park
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

Abstract

Ultracold atomic Bose-Einstein condensates (BECs) provide us a versatile platform for investigating superfluid hydrodynamics and quantum statistical mechanics. In this thesis, I present experimental research on the non-equilibrium phenomena at the interface in immiscible binary ^{23}Na condensates within a quasi-2D potential: ripplons, the Rayleigh-Taylor instability (RTI), and the prethermalization of the interface following a quench.

Our experimental system features immiscible Bose-Einstein condensates, whose spin domains are separated by a nearly 1D interface. By modulating the magnetic field gradient, we parametrically excited ripplons and measured their dispersion relation. The RTI was triggered by a rapid inversion of magnetic gradient; we observed its evolution and quantified its growth rate. Additionally, we demonstrated a matter-wave interferometry technique to visualize the counterflow velocity at the interface.

Separately, we investigated the relaxation of the 1D quantum fluid interface following a sudden quench. By quantifying the interfacial fluctuations, we observed that the quench leads to a long-lived prethermal state characterized by a persistent non-equipartition of energies among long-wavelength ripplon modes. We attributed this prethermal state to kinematic isolation: the absence of energy- and momentum-conserving relaxation in our system at ultracold temperature.

Dissertation Committee:

Professor Steven Rolston (Chair)

Adjunct Professor Ian B. Spielman (Co-Chair/Advisor)

Professor Gretchen K. Campbell (Co-chair/Advisor)

Assistant Professor Nathan Schine

Professor Christopher Jarzynski (Dean's Representative)

To the curious warriors
at the frontiers of
Knowledge

Acknowledgements

Completing my Ph.D. in the Department of Physics at the University of Maryland, College Park over the last five years has been a tremendous pleasure.

First, I thank my advisors, Ian Spielman and Gretchen Campbell, for welcoming me into their research groups. Ian taught me physics and the attitude towards physics research; from him, I also mastered the art of talking fast and skipping lunch. Throughout the years, I have been continually amazed by his profound intuition and boundless enthusiasm for new physics. Countless times, I walked into his office with a tangled mess in my head and walked out with a clear action plan. His incredible ability to spot interesting questions amid fragmented lab observations was the driving force behind several of my research tangents. Gretchen cared not just about my research progress, but also about my holistic growth as a scientist. I will always treasure her patience and willingness to lend an ear during the most difficult and emotionally taxing moments.

My research benefited immensely from the supportive, collaborative environment at the Joint Quantum Institute (JQI). I am appreciative of the kindness and openness shown by all the faculty members I have worked and brainstormed with. Alessandro Restelli was the one who led me into the wonderful world of electronics; together, we brought several research electronics projects to life. I cannot thank him enough for our many hands-on lab sessions and all the circuit design tricks he passed on to me. Trey Porto has a contagious love for atomic physics; I took great pleasure in the many conversations we had and the insights into condensate physics he kindly provided. Despite their busy schedules as junior faculty, both Alicia Kóllar and Nathan Schine generously offered me impromptu (sometimes late-night) lectures and guidance—not only on research questions, but also on navigating the human side of science. I also had the privilege of interacting frequently with William D. Phillips; I will always cherish our freewheeling conversations spanning physics, history, career development, and beyond.

I am also deeply grateful to the fellow students and postdocs I was lucky enough to work with over the years. I owe a huge debt to the previous generation of the sodium lab—Swarnav Banik, Monica Gutiérrez Galán, Hector Sosa Martínez, and Madi Anderson; my research projects would have been impossible without their pioneering work in building the experimental platform from the ground up. Shouvik Mukherjee accompanied me throughout my early days in the sodium lab. Running an ultracold atomic lab is both exhausting and stressful, and I cannot imagine having built and troubleshot the apparatus during my first few years without him. I am also greatly indebted to Junheng Tao and Mingshu Zhao, who bootstrapped my research initiative during our year-long collaboration. From Mingshu, I received both scientific ideas and moral support. Working with

Junheng was both exhilarating and productive—it was all too easy to lose track of time and forget to head home when running experiments together. I truly appreciate his contributions to my research and all the philosophical exchanges we had together. I also thank our long-time collaborator Steve Eckel for his remote guidance on metrology and vacuum technologies, and his careful reading of all my manuscripts.

Beyond research, I would like to recognize the members of various student organizations and administrative staff in the Department of Physics. Yonatan Gazit and Anna Fitzmaurice served alongside me the longest term on the Graduate Student Committee; I will fondly remember the fun days of planning events or painting rooms. I also thank Patrick Banner and Kellen O'Brien—my basement friends and fellow leaders of our sister organization, the Mental Health Task Force—for their assistance and shared efforts to make the department a better place. My sincere appreciation goes to Josiland Chambers, Donna Hammer, and Angel Torres in the Office of Student and Education Services, not only for supporting our student events, but also for guiding me through a maze of forms and administrative hurdles throughout graduate school. I am likewise grateful to the Graduate Program Director Tom Cohen for his wit and timely help during some highly unexpected situations. In my final year of graduate school, I spent a significant amount of time at the UMD Sandbox Makerspace. I am especially thankful to Director Gordon F. Crago for not kicking me out and allowing me to burn through 3D printing filament on my weird ideas. Kudos as well to the other student volunteers who eagerly shared their passion, taught me CNC skills, and injected a welcome dose of undergraduate vigor back into my life.

Life would have been immeasurably harder without the support from my housemates and graduate school friends. I am endlessly appreciative of my fellow physics students and housemates—Emily Jiang, Quincy Webb, and Bihan Shen—for all our late-night kitchen banter and the silly jokes we all laughed hard at, for the food we offered each other, for dragging and driving me to events and parties, and simply for their company to keep me sane. The years spent living with them will forever remain a cherished highlight of my life. I thank Zhiyin Tu for the guitar skills he taught me, the weekend jam sessions, and the funny music videos we shot together. Sarthak Subhankar was always willing to share inspiration and gave me numerous crash courses on how to build solid optics. I value his candor and patience, as well as his generosity in involving me in his projects, which eventually turned into a joint publication. My gratitude likewise goes to my second-floor office friends—Rahul Shrestha, Khoi Hoang, and Huy Nguyen—for all the spontaneous chats on photonics and the Lindblad form, for the PC-buying advice and the photography tutorials, and more. I also thank my other fellow basement dwellers in the Physical Sciences Complex—Martin Ritter, Yaxin Li, Deniz Kurdak, Kevin Weber, Marissa McMaster, Tim Li,

Jeff Lee, Scott Eustice, and Shukai Liu—for generously sharing lab equipment and sheltering me whenever I needed a quick break from the lab. I am grateful for the camaraderie of the other students in the Spielman Group: I had a brief but wonderful time sharing my lab with Davis Garwood, and I also appreciated Dario D'Amato and Michael Doris reliably showing up at my talks. As I depart from the sodium lab, I am glad to see Qi and Lynnx picking up the baton and working towards integrating the long-promised erbium into our apparatus—I wish them all the best.

Finally, I would like to thank my family and friends back in China. I thank my parents Yan and Yuexing for their unwavering support of every goal I pursued in my life, even when those paths did not look apparently promising to others. Due to immigration hurdles, I was not able to return home and spend time with them over the last five years; I am extremely grateful for their patience and understanding of my absence from family duties. Lastly, I thank Xiaowei Zhang for enlightening me in my youth and bringing me into the wonderful world of physical science.

Contents

1 Introduction **1**

1.1 Historical background: 1925–2026 . . . 2
1.2 Bose-Einstein condensate: a closer look . . . 4
1.3 Making a Bose-Einstein condensate . . . 5
1.4 Bose-Einstein condensate and superfluidity . . . 8
1.5 Interface phenomena in quantum fluids . . . 10
1.6 Structure of this thesis . . . 12

2 Theory of Bose-Einstein Condensates **13**

2.1 Quantum statistics of bosons . . . 13
2.2 Non-interacting Bose-Einstein condensate: a thermodynamic point of view 17
2.3 Weakly interacting Bose-Einstein condensate at zero temperature: static properties . . . 23
2.4 Weakly interacting Bose-Einstein condensate at zero temperature: dynamical properties . . . 29
2.5 Binary Bose-Einstein condensate . . . 33
2.6 Excitation spectrum of a weakly interacting Bose-Einstein condensate . . . 36
2.7 Hydrodynamic formalism . . . 44

3 Experimental Realization of a Quasi-2D Bose-Einstein Condensate **46**

3.1 Properties of sodium-23 . . . 46
3.1.1 Atomic structure and light-atom interaction . . . 47
3.1.2 Low-energy scattering properties . . . 48
3.2 Path to quantum degeneracy . . . 50
3.2.1 Laser systems . . . 51
3.2.2 Magnetic coils and antennas . . . 52
3.2.3 Atomic source . . . 53
3.2.4 Multistage cooling sequence . . . 54
3.3 Diagnostics and optimization of BEC production . . . 57
3.3.1 Absorption imaging of an atomic cloud . . . 57

3.3.2 Time-of-flight measurement of cloud expansion . . . 58
3.3.3 Optimization of condensate production . . . 60
3.4 Programmable quasi-2D potential . . . 63
3.4.1 Vertical confinement: sheet trap . . . 63
3.4.2 Horizontal confinement: DMD trap . . . 66

4 Imaging, Control, and Characterization of a Quasi-2D Bose-Einstein Condensate 70
4.1 Imaging a BEC . . . 70
4.1.1 Partial transfer absorption imaging (PTAI) technique . . . 71
4.1.2 Imaging theory: a quick tour . . . 72
4.1.3 Aberrations in the imaging system . . . 75
4.1.4 How optical density (OD) is modified by the imaging system . . . 76
4.2 Characterizing the imaging system . . . 79
4.2.1 Calibrating the magnification M . . . 79
4.2.2 Characterizing aberrations of the imaging system . . . 80
4.3 Control of magnetic field gradients in optical traps . . . 82
4.4 Calibrating the atom number by phonon excitations . . . 85

5 Study of Ripplons and the Rayleigh-Taylor Instability in a Binary Condensate 90
5.1 Dispersion relation of a quantum fluid interface . . . 91
5.2 Experimental measurement of the ripplon dispersion in a binary condensate 93
5.3 Experimental observation of the Rayleigh-Taylor instability in a binary condensate . . . 95
5.4 Obtaining the interfacial tension . . . 98
5.5 Visualizing counterflow velocity at the interface . . . 99
5.6 Conclusion . . . 101

6 Study of Prethermal Ripplons on a One-Dimensional Quantum Fluid Interface 103
6.1 Fluctuations on a quenched quantum interface . . . 104
6.2 Theory of thermal ripplons . . . 106
6.3 Quantifying prethermal ripplon modes . . . 107
6.4 Proposed prethermal mechanism: kinematic isolation . . . 109
6.5 Measuring the bulk temperature of the condensate . . . 110
6.6 Quantum statistics of thermal ripplons . . . 113

6.7 Conclusion 114

7 Conclusion and Outlook **116**

Appendix **117**

A Proper Definition of Interfacial Tension **118**

B Numerical Computation of BEC Dynamics and Quasi-Particle Spectrum **122**

B.1 Further reduction of dimension by symmetry 122

B.2 Defining the numerical units 123

B.3 Solving for the condensate ground state 124

B.4 Survey of numerical methods for solving GPE 125

B.5 Computing the quasi-particle spectrum 134

Acronyms **137**

Bibliography **139**

1 Introduction

> Too far into the early morning hours, three young men sit in a cold room, darkened save for the phosphor glow of computer screens. Behind them lurks a disorderly ton of steel and glass, layered with hours of care like a wedding cake, but interlaced as a chaotic tower of Babel. [...] The three listen intently, not so much to a physical sound, but rather to an abstract music which accompanies an unobserved dance of light and atoms. In this unimaginably cold spot, a new form of matter is born and reborn every minute.
>
> *— Dan M. Stamper-Kurn describes the operation of his BEC experiment at MIT (2000) [1]*

A century has passed since Albert Einstein first predicted the existence of Bose-Einstein condensates [2]. More than 30 years ago, researchers first observed a Bose-Einstein condensate (BEC) in weakly interacting ultracold atomic gases [3–6], which has since ignited a remarkable progression in the study of quantum gases—a center pillar of modern physics. In the first quarter of the 21st century, ultracold atomic Bose-Einstein condensates have served as pristine platforms for exploring a broad range of physical phenomena: superfluidity and quantum hydrodynamics, analogue gravity, thermalization in quantum systems, just to highlight a few. This field has witnessed rapid technological leaps that continually push the boundaries of experimental possibility, enabling the study of new physics at unprecedented scales and under extreme conditions. Yet amid a constantly shifting scientific landscape, the fundamental protocol for achieving quantum degeneracy—laser cooling and evaporative cooling—remains indispensable. Likewise unchanged is the daunting complexity of the experimental contraptions in which these studies take place. The typical laboratory setup differs little from the labyrinth of hardware described by Stamper-Kurn in the epigraph—if anything, these setups have become even more intricate and convoluted after thirty years of evolution!

The research work I present in this thesis branches from the ever-growing tree of atomic Bose-Einstein condensate and quantum fluid studies—specifically, the interface phenomena of an immiscible binary quantum fluid composed of a pair of ferromagnetic Bose-

Einstein condensates of ultracold sodium (^{23}Na) atoms. In this introductory chapter, I feel obliged to trace the series of historic milestones that led to this work, and meanwhile introduce some fundamental concepts that will lay the groundwork for the chapters ahead in this thesis.

1.1 Historical background: 1925–2026

At its core, experimental physics relies on the dual pillars of *control* and *measurement*. It is through precise control of the physical systems that their responses can be conditioned and studied, and subsequently quantified through accurate measurements. The facts collected then become the fertile soil in which theories are conceived, and these theories, in turn, guide the development of control and measurement techniques, enabling the study of more complex and delicate physical systems. The pursuit of the Bose-Einstein condensate (BEC) in ultracold atomic gases stands as one of the most striking manifestations of this paradigm.

The concept of BEC dates back to 1925, when Albert Einstein [2] first proposed it. His intention was to address a quite subtle issue in Bose statistics, specifically, that the chemical potential of an ideal Bose gas reaches zero at finite temperature, resulting in an unphysical divergence in the population of the lowest energy state. His proposal—that this divergence signifies a macroscopic fraction of atoms occupying the lowest energy state—might seem quite counterintuitive to researchers of that era, since most matter solidifies as the temperature decreases rather than transforming into an exotic new state like a BEC. Yet Einstein's theory soon found an unexpected experimental realization: in 1937, Pyotr Kapitsa, along with John F. Allen, and Don Misener discovered the superfluidity of liquid helium-4 [7, 8]—a phenomenon Fritz London [9] quickly linked to BEC. It is fair to say that the development of technologies in low-temperature physics enabled the cooling of a substantial amount of liquid helium-4, making such striking discoveries possible. Kapitsa later was awarded the Nobel Prize for this contribution. But superfluid helium-4 is not strictly a condensate, or at least not a "good" one: neutron scattering experiments later revealed that the actual population fraction in the zero-momentum state (which is the lowest energy single-particle state) is a mere 10% of the total population at saturated vapor pressure [10–12]. The connection between BEC and superfluidity, as it turns out, is a subtle one and warrants a careful discussion.

The problem of helium-4 is its strong interaction between atoms, which causes very severe depletion of the zero-momentum state population even at absolute zero temperature. As atoms in solids and liquids (like superfluid helium-4) interact strongly, researchers

were motivated to search for weakly interacting systems that remain gaseous at low temperature. Spin-polarized hydrogen [13, 14] emerged as the next front-runner in the race to achieve BEC. Under a strong magnetic field, the interaction between polarized hydrogen atoms is so weak that they will never liquify down to arbitrarily low temperature at ambient pressure. Furthermore, hydrogen can be prepared at large density by cryogenic technology, permitting a relatively high transition temperature. Significant progress was made in cooling spin-polarized hydrogen gases: they were magnetically trapped, and the concept of forced evaporative cooling—an important step in all modern atomic BEC experiments—was first demonstrated on them [15]. However, this momentum was ultimately slowed down by rapid dipolar spin-relaxation and inelastic three-body recombination in a dense cloud [5]. As a result, the spin-polarized hydrogen was eventually surpassed by alkali atoms in the BEC search, which benefited immensely from the maturing of laser cooling and trapping techniques. In 1995, the first realization of BEC took place in a cold cloud of ^{87}Rb atoms at JILA [3], two years before hydrogen was finally condensed [16]. This spectacular journey is recognized by two Nobel Prizes: the first in 1997 was awarded to Steven Chu, Claude Cohen-Tannoudji, and William D. Phillips for the development of laser cooling and trapping, and the second in 2001 was awarded to Eric Cornell, Wolfgang Ketterle, and Carl Wieman for achieving BEC. The techniques pioneered during this era have turned into the bread and butter for all ultracold atomic experiments in the 21st century.

Arguably, BEC in weakly interacting atomic gases is the first pristine isolated quantum system ever produced by mankind, sitting idly in vacuum, awaiting detection by laser light. It has become a unique and compelling platform for studying quantum effects on a macroscopic scale (the size of a condensate is $\approx 10^{-5}$ m, compared to typical atomic size 10^{-10} m). The initial successful preparation and observation of this remarkable physical system set off a global frenzy of condensate research.

In hindsight, alkali atoms proved to be better candidates for BEC than hydrogen—not only due to their laser-accessible energy structure, but also because their hyperfine structures give rise to a plethora of fascinating physics. These hyperfine structures, combined with their Zeeman splitting under magnetic fields, provide an atomic analog to electron spin, hence the term "spinor condensate" [17, 18]. In the first generation of BEC experiments, atoms were confined in magnetic potentials where evaporative cooling took place [3, 4]. However, this magnetic trap often locks atoms to a single spin state (the "low-field seeking" state), expelling other spin states with incompatible magnetic dipole moments and effectively freezing out the spin degree of freedom. The situation changed in 1997, when optical trapping techniques were demonstrated by pioneering work in Ket-

terle's group [19], led by graduate student Dan M. Stamper-Kurn, who later became a faculty member at the University of California, Berkeley. Together with Nicholas Bigelow at the University of Rochester, Yong-il Shin at Seoul National University, Paul Lett at the University of Maryland (UMD), and several other experimental groups, alongside key theoretical groundwork from Masahito Ueda's group, they helped usher in two decades of fruitful spinor condensate research.

In the 2020s, a new generation of researchers around the world took up the baton—including the KAIST group led by Jae-Yoon Choi, the Trento group led by Gabriele Ferrari, and the group I am part of at the Joint Quantum Institute (JQI), a collaboration between the University of Maryland and the National Institute of Standards and Technology (NIST). In this thesis, I present my research at the frontier of spinor condensate studies, focusing on the interface phenomena at the domain wall between two ferromagnetic condensates. This is where the two spins meet—or even clash—stirring up chaotic, yet sometimes spectacular and beautiful, patterns.

Embarking on the study of spinor condensate research was neither a journey I had planned nor one I initially envisioned. It arose from a serendipitous combination of experimental capabilities, fortunate events, and the collective know-how readily available in the group [1], to which I will be forever grateful.

In the following sections of this chapter, I will provide a heuristic introduction to the very core of this thesis—Bose-Einstein condensates and superfluidity—as well as some historical context for the interface phenomenon.

1.2 Bose-Einstein condensate: a closer look

Bose-Einstein condensation is a phase transition that occurs in a bosonic system at temperatures near absolute zero. All macroscopic objects are composed of microscopic building blocks: atoms and molecules. In quantum mechanics, these microscopic constituents are "indistinguishable". Unlike classical macroscopic objects, which can be differentiated by their compositions or defects, the atoms and molecules are truly identical and cannot be "tagged". There are two categories of indistinguishable particles: bosons and fermions, named after Satyendra Nath Bose and Enrico Fermi—two giants in quantum statistical mechanics. Fermions cannot occupy the same quantum state, whereas bosons can, allowing them to condense *en masse* into the ground state. This particular configuration—where a macroscopic number of bosons occupy the ground state—is peculiar. It is the equilibrium configuration of the bosonic system at temperatures near absolute zero, maximizing the

[1] I'm deeply indebted to Mingshu and Junheng, whose insights and curiosity made this work possible.

system's entropy. At first glance, it may seem implausible that placing most bosons in the ground state would yield a higher entropy than distributing them among different excited states. Two clarifications are in order: first, atoms in the ground state do not contribute entropy, but atoms in the excited states do; second, the energy constraint limits the number of atoms in excited states, forcing the majority into the ground state.

The process of Bose-Einstein condensation can be intuitively understood through a semiclassical picture [5]. Heisenberg's uncertainty principle states that the exact coordinates and momentum of a particle cannot be determined simultaneously. This means atoms are not point particles but wavepackets, whose extent is on the order of the *thermal* de Broglie wavelength $\lambda_{\text{th}} = (2\pi\hbar^2/mk_BT)^{1/2}$ (not to be confused with the de Broglie wavelength), where T is the temperature and m is the atomic mass. At high temperature, the broad momentum distribution results in a narrow wavepacket; λ_{th} is small, rendering the wavepacket "point"-like. As temperature decreases, λ_{th} becomes larger and larger, until the point that the wavepackets start to overlap, merging hundreds of thousands of atoms into a macroscopic quantum object—a Bose-Einstein condensate. For an atomic cloud that has a peak density of n, this phase transition happens at $n\lambda_{\text{th}}^3 \approx 2.6$. [2]

This phase transition is unique because it occurs without any interaction between bosons. Usually, interactions drive phase transitions in matter; for example, chemical bonds are formed, broken, or modified during this process. In contrast, Bose-Einstein condensation is purely a result of quantum statistics. Nevertheless, real atoms do interact, albeit weakly under our typical experimental conditions. This interaction turns out to be a fundamental element in ultracold atomic BECs and provides a crucial bridge between condensation and superfluidity. In a non-interacting BEC, each atom is separately described by the single-particle Schrödinger equation; the condensate state is just a product state of these identical single-particle ground states. With interactions, the condensate state becomes a mixture of single-particle states, though under weak interactions—as in our experiments—it can still be well described by a product state of individual atoms' single-particle states. More intuitively, the interaction creates pressure that "flattens" the density peak, producing a more uniform density profile. A detailed description of the thermal and quantum mechanical properties, along with the effect of interactions, is the topic of Chapter 2.

1.3 Making a Bose-Einstein condensate

In our experiment (pictured in Fig. 1.1), the condensate is prepared in an ultra-high vacuum (UHV) system at a pressure of $\sim 10^{-10}$ mbar. Figure 1.2 illustrates the BEC creation process.

[2] For a rigorous derivation of this argument, see steps leading to Eq. (2.22)

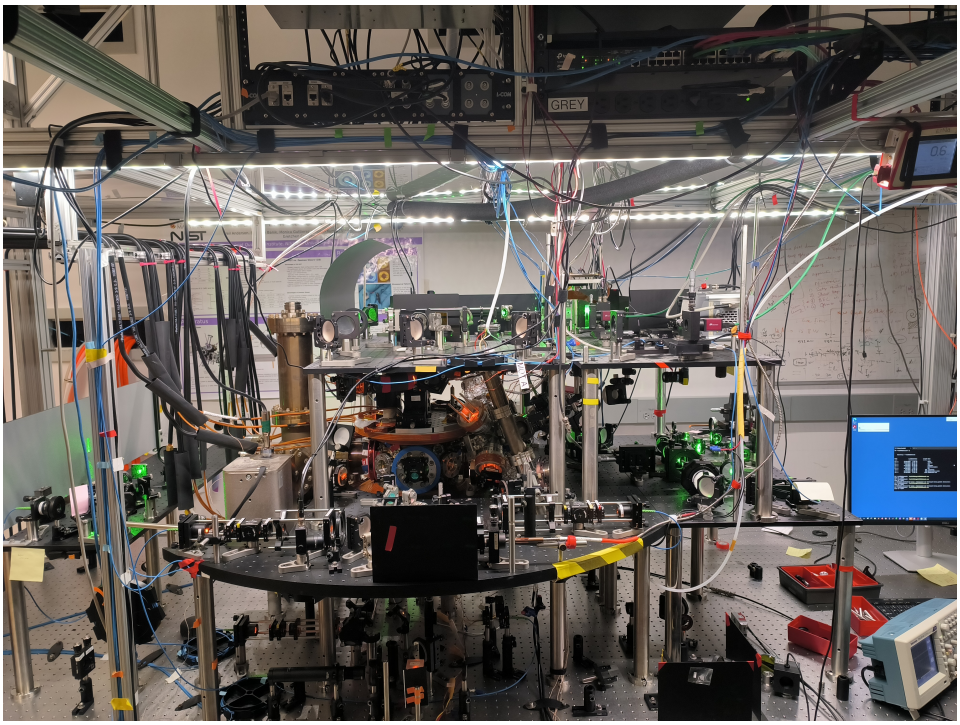

Figure 1.1: A photo of the main vacuum system and the optics surrounding it.

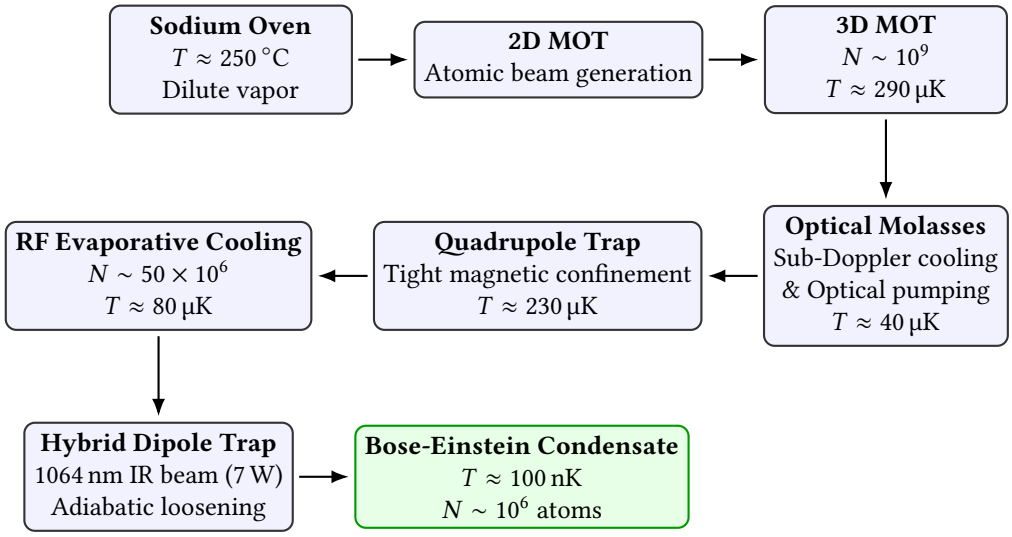

Figure 1.2: Experimental sequence for generating a ^{23}Na Bose-Einstein condensate in a hybrid trap.

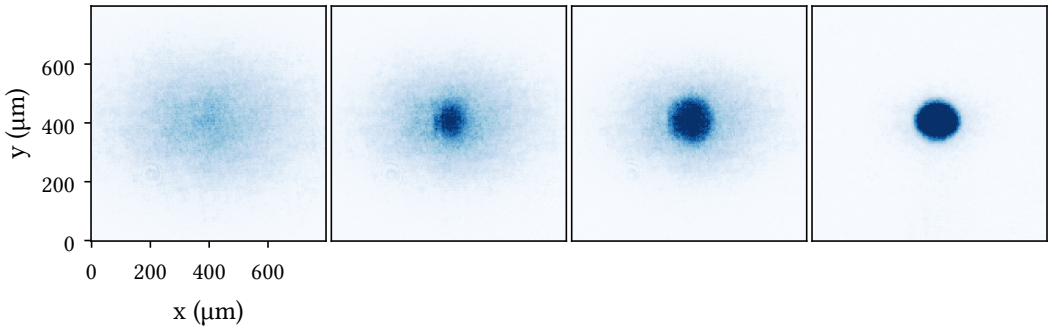

Figure 1.3: From left to right: an ultracold cloud of sodium atoms before and after crossing the condensation threshold during evaporative cooling. The images are taken after 7 ms of time-of-flight.

In the oven section, a chunk of metallic sodium is vaporized into a hot, dilute gas at a temperature of 250 °C. The hot atomic gas flows into a two-dimensional magneto-optical trap (2D MOT) that provides slowing and trapping in two dimensions, effectively confining cold atoms into a short "string". This string of cold atoms is pushed into a 3-dimensional MOT by a blue-detuned laser beam, where around 10^9 atoms reach a phase space density of $\sim 10^{-6}$ at around 290 μK. After this loading process is finished, we turn off the magnetic field and let atoms further cool through optical molasses, before optically pumping them into the low-field-seeking state and turning on the quadrupole magnetic trap with a field gradient of 190 G cm^{-1} (and 170 A of DC current inside the coils). Now, atoms are tightly trapped in the magnetic trap, yielding a high rate of collisions. We turn on an RF field to selectively flip the magnetic dipole of energetic atoms into the untrapable state, "evaporating" them out of the trap to bring down the temperature of atoms remaining inside the trap. At the end of this evaporative cooling process, the magnetic trap is loosened adiabatically. Now, atoms start to gather near the center of the trap, where the low magnetic field exposes them to loss and heating induced by the uncontrolled spin-flip (Majorana spin-flip). To mitigate, we focus an infrared 1064 nm laser with 7 W of power down to a diameter of 40 μm near the center of the magnetic trap, forming an optical dipole trap. This optical confinement traps atoms before they reach the low-field region. Evaporation continues by reducing the optical trap depth; once the temperature crosses ~1 μK, a BEC with $\sim 10^6$ atoms finally forms (Fig. 1.3) within the combined optical and magnetic trap, termed the "hybrid trap". The details of the experimental setup and protocol are discussed in Chapter 3.

1.4 Bose-Einstein condensate and superfluidity

The confusing connection between Bose-Einstein condensate and superfluidity has been a subject of decades of debate and remains a conceptual hurdle for graduate students even today. It is worth noting that the concept of superfluidity was developed in the context of helium-4 research, and historically, the same term has been used to describe a range of different phenomena. Here we shall focus on the most common definition of superfluidity: the ability to flow without dissipation. This property is found in helium-4 at temperatures below its superfluid phase transition (which forms a "poor" BEC, see above), in most (if not all) atomic BECs, and also in superconductors (in which electron pairs flow without resistance). The two-fluid model, developed by Landau in 1941 [20], phenomenologically captures many key properties of superfluids and superconductors. In particular, Landau introduced an elegant picture explaining how a fluid can flow without dissipation.

A normal fluid experiences viscosity and friction with the wall when flowing in a pipe. The kinetic energy of the flow is converted into heat—random density fluctuations inside the fluid. Landau argued that dissipating the macroscopic flow into density fluctuations (or phonon excitations in quantum terms) is not always energetically favorable.

Landau considered a uniform liquid at zero temperature, where the entire system rests in its ground state with energy E_0. Now, imagine this liquid is in an infinitely long pipe moving with velocity $\mathbf{u}$. The question is: Will the liquid move with the pipe or remain stationary? In the reference frame where the pipe is stationary, but the fluid is moving, the energy of the fluid, after a Galilean transformation, is

$$E(\mathbf{u}) = E_0 - \mathbf{p}_0 \cdot \mathbf{u} + \frac{1}{2} M \mathbf{u}^2, \tag{1.1}$$

where $\mathbf{p}_0$ is the momentum of the fluid in the initial rest frame (hence $\mathbf{p}_0 = 0$) and M is the total mass.

Now, if the fluid in the original frame has one single massless excitation of momentum $\mathbf{p}$, the momentum and energy of the system become

$$\mathbf{p}_{\text{ex}} = \mathbf{p}, \tag{1.2}$$

$$E_{\text{ex}} = E_0 + \epsilon_p. \tag{1.3}$$

In the moving frame, the energy is

$$E_{\text{ex}}(\mathbf{u}) = E_{\text{ex}} - \mathbf{p}_{\text{ex}} \cdot \mathbf{u} + \frac{1}{2} M \mathbf{u}^2 \tag{1.4}$$

$$= E_0 + \epsilon_p - \mathbf{p} \cdot \mathbf{u} + \frac{1}{2} M \mathbf{u}^2. \tag{1.5}$$

Now, comparing the energy of the system before and after adding the excitation

$$E_{\text{ex}}(\mathbf{u}) - E(\mathbf{u}) = \epsilon_p - \mathbf{p} \cdot \mathbf{u}. \tag{1.6}$$

If the difference is greater than zero, $\epsilon_p > \mathbf{p} \cdot \mathbf{u}$, then creating an excitation actually *increases* the total energy in the moving frame. Consequently, it is energetically unfavorable to spontaneously create these excitations. The fluid will show superfluidity.

For an excitation with energy $\epsilon_p = cp$, there exists a maximal flowing velocity $u_c = \epsilon_p / p = c$ below which excitations cannot be generated. This threshold is termed the Landau critical velocity. For a system with a more general dispersion relation, the Landau critical velocity becomes

$$u_c = \min\left(\frac{\epsilon_p}{p}\right). \tag{1.7}$$

From this picture, it becomes clear that it is the nature of excitations of a quantum system that determines its superfluidity. This explains why, at $T = 0$, superfluid helium-4 achieves a superfluid density of $\rho_s / \rho = 1$, despite having a condensate fraction of only ~10%. Its ground state is a highly entangled many-body state, yet the excitations are well described by phonons and rotons, which yield a Landau critical velocity of $\approx 60\ \text{m s}^{-1}$ [21]. Experimentally, the measured critical velocity is generally lower than this number, due to the creation of other excitations (such as vortices).

Landau's theory also explains why the ideal, non-interacting BEC lacks superfluidity, whereas a weakly interacting BEC exhibits it. As will be discussed in the later chapters, the spectrum of a non-interacting bosonic system is $\epsilon_p = p^2 / 2m$ where m is the particle mass. From Eq. (1.7), we find that its $u_c = 0$, implying any non-zero flow velocity will lead to dissipation, preventing a non-interacting gas from sustaining superfluidity. On the other hand, a weakly interacting BEC is superfluid: it has a phonon spectrum $\epsilon_p = cp$, resulting in a finite critical velocity $u_c = c > 0$ that permits the existence of a flow without dissipation.

For binary condensates separated by an interface—the main focus of this thesis—a distinct type of excitation emerges: ripplons, which are capillary wave-like interfacial excita-

tions with a dispersion relation $\epsilon_p \sim p^{3/2}$. Similar to a non-interacting BEC, this dispersion yields $u_c = 0$, meaning that "flow" in a binary condensate generally excites ripplons along the interface, dissipating energy. This presents an interesting case: a system where both individual components are superfluids by themselves, yet the system as a whole is not. It illustrates another pitfall: superfluidity is a global property of the entire system and cannot always be inferred by studying a local portion of it. Conversely, applying a stabilizing force on two sides of the interface will modify the low-$\mathbf{p}$ portion of the dispersion relation into $\epsilon_p \sim p^{1/2}$, thus restoring a non-zero u_c that permits superfluidity.

So far, we have only considered systems at $T = 0$. At finite temperature ($T \neq 0$), thermal excitations are present in the fluid. Scattering processes between them enable energy transfer; therefore, the thermal excitations form a fraction of fluid that actually dissipates on top of the superfluid fraction that does not. This dissipative component, termed the normal component, together with the superfluid component, completes Landau's two-fluid theory.

A term frequently used alongside superfluid is *quantum fluid*, although there is no single consensus on its definition, and it is often used to refer to a broader class of quantum many-body systems. In this thesis, I will restrict the term quantum fluid to quantum systems with non-negligible inter-atomic interactions—systems that typically exhibit superfluidity precisely because of those interactions.

1.5 Interface phenomena in quantum fluids

The main focus of this thesis is the interfacial capillary excitations of a binary Bose-Einstein condensate. Capillary waves are commonplace in the classical world; we see them as ripples when dropping a pebble into a pond or gently blowing across hot coffee. *Ripplons* are the quantized version of capillary excitations in quantum fluids.

The study of capillary waves dates back to Lord Rayleigh in the 19th century [22]. Their quantum counterpart, ripplons, gained prominence during the study of trapped electrons above the surface of superfluid helium. In 1979, Grimes and Adams [23] from Bell Laboratories first confirmed that surface-trapped electrons form a crystalline structure by using RF fields to excite standing capillary waves at the liquid-helium interface. Their experimental report was complemented by a detailed theoretical analysis of this "phonon-ripplon coupling" by their Bell Laboratories colleagues Fisher, Halperin, and Platzman [24], who subsequently popularized the term "ripplon".

The context of their work was a search for the Wigner crystal [25, 26], a phase of electrons that occurs at low temperatures when the Coulomb interaction between localized

electrons overpowers their kinetic energy. The electrons trapped near the helium fluid surface actually create "dimples" on the surface, making the strong coupling to capillary standing waves not a surprising one. The electron lattice moves under oscillating RF fields. When the drive frequency resonates with a ripplon whose wave vector $\vec{q}$ is an integer multiple of a reciprocal lattice vector $\vec{G}$ (or the lattice spacing is a multiple of the ripplon wavelength), that ripplon is excited. Experimentally, this resonance is reflected as a sharp change in the drive impedance. It is worth noting that although the term "ripplon" carries a quantum connotation, both the theoretical treatments and experimental measurements in these studies were completely classical. Nevertheless, the electron-ripplon coupling can exhibit genuine quantum effects, such as the shift in the electron energy level [27] (a condensed matter analogue of the Lamb shift in quantum electrodynamics). The study of electrons trapped near the liquid helium interface and their interaction with ripplons proved to be a fruitful direction, with research on this system continuing to this day.

The investigation of the interface phenomena in BECs dates back to 1997, when the Wieman and Cornell group at JILA observed the first binary condensate of $|F = 1, m_F = -1\rangle$ and $|2, 2\rangle$ hyperfine states of ^{87}Rb [28]. In the following year, they observed spatial separation of BECs in $|1, -1\rangle$ and $|2, 1\rangle$ states [29], though their "interface" was not quite sharp—arguably due to the unfavorable scattering properties of ^{87}Rb and the low density of their condensates. In a seminal work published in the same year, Ao and Chui [30] analyzed the condition for spatial segregation and evaluated the interfacial tension between two immiscible BECs (a value later shown to be off by a factor of 4 [31]). Capillary waves at the interface of immiscible BECs were first analyzed theoretically in 2001 by Mazets [32], who derived the approximate dispersion relations of these interfacial modes.

Interestingly, despite clear theoretical interest, direct experimental observation of ripplons in binary condensates did not materialize for nearly two decades. During this interim, theoretical works continued to grow, exploring phenomena such as the quantum analogue of the Rayleigh-Taylor instability (RTI) at the quantum fluid interface.

The experimental hiatus ended in 2024. One month before the 2024 DAMOP conference, our group observed and systematically studied the RTI and ripplon dynamics in an immiscible binary mixture of spinor ^{23}Na BECs (which grew out of a fun project with Junheng on a Sunday afternoon) [33] [3]. Our observation was complemented by a study of the Kelvin-Helmholtz instability in ^{6}Li BECs carried out by Jae-Yoon Choi's group at KAIST,

[3] The scientific writer at JQI, Bailey Bedford, did a really good popular write-up of our work. See https://web.archive.org/web/20260117105243/https://jqi.umd.edu/news/when-superfluids-collide-physicists-find-mix-old-and-new.

published the same year [34]. The detail of our investigations into ripplons and RTI will be presented in Chapter 5.

It is worth noting that RTI—defined by the iconic mushroom-shaped patterns formed when two fluids are pressed into each other—actually originates from a very simple physical mechanism, which makes it ubiquitous across diverse physical systems. As an example, it even plays a (notoriously negative) role during the laser ignition process of inertial confinement fusion [35]. Our observation of RTI in binary BECs is yet another manifestation of this universal phenomenon. It serves as a reminder that, despite the seemingly vast differences in scale and energy, analogous dynamics naturally emerge if the governing principles coincide. This universality is one of the most enchanting traits of physics—and one that has continually inspired me throughout my research journey.

1.6 Structure of this thesis

In the following chapters, before I dive into my original research, I will first introduce the basic theory of atomic physics and Bose-Einstein condensation. I will also outline the experimental apparatus, though the discussion will be brief—as most of it is detailed in the dissertations of the former lab members Swarnav Banik and Monica Gutierrez Galan. I will focus primarily on the apparatus upgrades and technical developments I implemented. The main research chapters will then follow, covering my two core original projects: the investigation of ripplons and the RTI in the binary quantum fluid, and the prethermalization of ripplons at immiscible BEC interfaces.

This thesis will be logically divided into three parts: the background and introduction (Chapter 1–2), technical implementation (Chapter 3–4), and primary scientific results (Chapter 5–6).

While this chapter is written in an accessible style for readers with general scientific training, the remainder of this thesis assumes physical knowledge at the advanced undergraduate or junior graduate level.

2 Theory of Bose-Einstein Condensates

The theory of BEC in dilute cold atomic gases is well established and thoroughly covered in standard textbooks [21, 36]. The goal of this chapter is to present an experimentalist's perspective on condensate theory that is directly relevant to interpreting experiments and serving as a solid basis for numerical calculations of the work presented in this thesis.

This chapter largely follows the conventional pedagogical structure. Alongside technical derivations, I emphasize the underlying physical principles and conceptual caveats. I start with a brief review of quantum statistical mechanics of bosons before examining the thermodynamic properties of non-interacting Bose gases in three-dimensional free space and harmonic traps. Next, I introduce the mean-field description of a condensate at zero temperature: the ground state, the equation of motion—the renowned Gross-Pitaevskii equation (GPE)—and the reduction of dimension for analyzing two-dimensional systems. Finally, I expand the discussion to binary condensates, examine elementary excitations in a condensate, and quickly review the hydrodynamics formalism which describes density excitations in a slightly different flavor.

2.1 Quantum statistics of bosons

In thermal equilibrium, quantum and classical systems alike arrive at configurations that maximize entropy, subject to relevant constraints. The classical entropy for a system in contact with a bath is defined as

$$S = -k_B \sum_i P_i \ln P_i, \tag{2.1}$$

where P_i is the probability of finding the system in microstate i. Similarly, one can define the entropy for a quantum system as (following the notation of Sakurai [37])

$$S/k_B \equiv \sigma = -\mathrm{tr}(\hat{\rho} \ln \hat{\rho}), \tag{2.2}$$

where the summation over probabilities is replaced by the trace of the density operator $\hat{\rho}$. Its connection to the classical entropy can be easily verified by looking at this expression in the basis where the density matrix is diagonalized.

The density matrix itself can describe any quantum system, be that a single particle or a complicated many-body system. To find the thermal equilibrium configuration of a quantum system and obtain its macroscopic properties, we maximize Eq. (2.2) for that system while constraining the average energy $\langle H \rangle$ and average particle number $\langle N \rangle$, similar to the treatment in classical statistical mechanics. The maximization requires

$$\delta\sigma = -\sum_j [\delta\rho_{jj}(\ln \rho_{jj} + 1)] = 0. \tag{2.3}$$

The two constraints, along with the normalization $\mathrm{tr}\rho = 1$, demand that

$$\delta\langle H \rangle = \delta \mathrm{tr}(\hat{\rho}\hat{H}) = 0, \tag{2.4}$$

$$\delta\langle N \rangle = \delta \mathrm{tr}(\hat{\rho}\hat{N}) = 0, \tag{2.5}$$

$$\delta(\mathrm{tr}\rho) = \sum_j \delta\rho_{jj} = 0. \tag{2.6}$$

If $\hat{\rho}$ is diagonal in the energy and particle number basis,

$$\delta\,\mathrm{tr}(\hat{\rho}\hat{H}) = \sum_j \delta\rho_{jj}E_j, \tag{2.7}$$

$$\delta\,\mathrm{tr}(\hat{\rho}\hat{N}) = \sum_j \delta\rho_{jj}N_j, \tag{2.8}$$

applying the Lagrange multiplier technique yields the variational equation

$$\sum_j \delta\rho_{jj}(\ln \rho_{jj} + 1 + \beta E_j + \xi N_j + \gamma) = 0 \tag{2.9}$$

that must hold true for all possible small variations $\delta\rho_{jj}$, which requires each term in the summation to vanish, yielding

$$\rho_{jj} = \exp\big(-\beta E_j - \xi N_j - \gamma\big), \tag{2.10}$$

which is exactly the form of the Boltzmann factor. Note that arriving at Eq. (2.10) relies solely on maximizing the entropy of the quantum system within the density matrix for-

malism. We can readily identify the Lagrange multiplier $\beta \equiv 1/k_B T$ and $\xi \equiv -\mu/k_B T$ for a grand canonical ensemble, where μ is the chemical potential. The normalization

$$Z \equiv \exp(\gamma) = \sum_j \exp\big(-\beta(E_j - \mu N_j)\big) = \mathrm{tr}\Big(e^{-\beta(\hat{H}-\mu\hat{N})}\Big) \tag{2.11}$$

is defined as the partition function. The density operator can therefore be rewritten as

$$\hat{\rho} = \frac{e^{-\beta(\hat{H}-\mu\hat{N})}}{Z}. \tag{2.12}$$

Now we apply this formalism to a non-interacting bosonic system. Assuming each boson has its energy eigenstates labeled by index k, and the many-body energy eigenstates are labeled by $|\cdots, n_k, n_{k+1}, \cdots\rangle$, where n_k is the occupation number of state k (for fermions, n_k must be 0 or 1; for bosons, it can be any non-negative integer), we have

$$\hat{H}|\cdots, n_k, n_{k+1}, \cdots\rangle = \left(\sum_k n_k \epsilon_k\right)|\cdots, n_k, n_{k+1}, \cdots\rangle, \tag{2.13}$$

$$\hat{N}|\cdots, n_k, n_{k+1}, \cdots\rangle = \left(\sum_k n_k\right)|\cdots, n_k, n_{k+1}, \cdots\rangle. \tag{2.14}$$

This permits us to evaluate the mean occupation number $\bar{n}_i$ of state i by summing over all possible configurations of $\{n_k\}_{k=1,2,\cdots}$:

$$\bar{n}_i = \mathrm{tr}(\hat{\rho}\hat{n}_i) = \frac{\sum_{\{n_k\}} n_i \prod_k e^{-\beta(\epsilon_k-\mu)n_k}}{\sum_{\{n_k\}} \prod_k e^{-\beta(\epsilon_k-\mu)n_k}} \tag{2.15}$$

$$= \frac{\sum_{n_i} n_i e^{-\beta(\epsilon_i-\mu)n_i}}{\sum_{n_i} e^{-\beta(\epsilon_i-\mu)n_i}} \tag{2.16}$$

$$= \frac{1}{e^{\beta(\epsilon_i-\mu)} - 1}. \tag{2.17}$$

This is the celebrated Bose distribution function, which forms the basis for discussions throughout the rest of this chapter. We can quickly verify that in the high energy limit ($\epsilon_i - \mu \gg k_B T$), the state has very low occupation, and Eq. (2.17) reduces to the Maxwell-Boltzmann distribution $\bar{n}_i = e^{-\beta(\epsilon_i-\mu)}$ as expected. For any bosonic system, the chemical potential μ can never exceed the single-particle ground state energy ϵ_0; otherwise, $\bar{n}_0$ becomes unphysical. As $T \to 0$, μ approaches ϵ_0, driving $\bar{n}_0$ to a macroscopic fraction of the total particle number and signaling the formation of a BEC. The conservation of total particle number plays a central role: bosons that do not obey number conservation (such

as photons or phonons) have $\mu = 0$. When T decreases, their total population decreases continuously, preventing them from condensing under usual circumstances[1].

[1] However, a BEC of photons has been demonstrated in a dye-filled microcavity, where the long decay time of photons in the cavity effectively enforces a temporary photon number conservation on the thermalization timescale [38]. The prethermalization of ripplons, which will be covered in later chapters, shares the same conceptual thread.

2.2 Non-interacting Bose-Einstein condensate: a thermodynamic point of view

Having derived the Bose distribution function, we now have the foundation to study the thermodynamics of bosonic gases at low temperatures. The usual presentation of this topic (e.g., [36, Ch. 2], [21, Ch. 3], or [39, Ch. 12]) involves the calculation of the total atom number of a bosonic system in the continuum limit, showing that it saturates at low temperatures—suggesting that condensation must occur, largely in line with Einstein's original argument in 1925. Before we follow suit, several key nuances warrant clarification.

First, most theoretical work on BECs is formulated in the grand canonical ensemble, which assumes the system of interest is embedded in a bath that can exchange both energy and particles. In actual experiments, the atoms are trapped in vacuum. Certainly, the isolated cloud has no path to exchange energy or particles with the outside world (indeed, minimizing such interactions is precisely why an ultra-high vacuum is a requirement in the first place). Technically, atoms in a static, lossless trap are described by the microcanonical ensemble, corresponding to an isolated system with fixed total energy and particle number. One reason for the widespread usage of the grand canonical ensemble, despite this apparent mismatch with experimental reality, is the mathematical simplicity it offers [2]. The microcanonical ensemble is notoriously difficult to work with! Think about it: the key question in working with the microcanonical ensemble is to find the number of ways to distribute a fixed amount of energy among a fixed number of atoms—the number of possible partitions $\Gamma(E, N)$. This is a hard combinatorial problem. Canonical and grand canonical ensembles are far more tractable because they allow some "looseness" in the exact energy and atom numbers. When considering the mean thermodynamic quantities, the choice of ensemble does not matter in the thermodynamical limit $N \to \infty$—the "looseness", or the relative variation from the fixed total atom number N, is infinitesimally small ($\sim N^{-1/2}$). While the grand canonical ensemble empirically yields decent results for mean quantities even at moderate N, when the looseness itself—fluctuations of condensate atom number—is concerned, the grand canonical ensemble predicts "unphysically large" condensate population fluctuations $\delta N_0 \sim N_0$, which persist however large the system is—an artifact known as the *grand canonical catastrophe*. That's when one must revert to the microcanonical ensemble [40, 41].

[2] Much of the literature treats the thermal cloud as the condensate's reservoir. This is technically wrong: the thermal fraction is itself predicted by the Bose distribution, so invoking it to justify the grand canonical treatment is tautological. Furthermore, as the system temperature drops below the transition temperature, the population of the thermal fraction is quickly reduced compared to the condensate mode, failing the basic requirement that the reservoir must remain overwhelmingly larger.

The second point to clarify is the meaning of "chemical potential". Though colloquially referred to as "the energy increase due to the addition of a single particle", this imprecise language is misleading and routinely generates pseudo-problems. Taking the above statement at face value, a puzzle immediately arises: why does μ in Eq. (2.17) have a negative sign if it indicates an increase in energy? Recall how the chemical potential μ was introduced in the last section: I steered clear of mentioning "energy" altogether—it is fundamentally a tool to regulate the average atom number, exactly as temperature regulates the system's total energy. Statistical mechanics revolves around maximizing entropy. The effect of injecting an extra atom is two-fold: the system energy is increased by at least the single-particle ground-state energy—the zero-point energy in a harmonic trap; yet, the new atom provides additional degrees of freedom among which the total energy can redistribute, yielding a gain in entropy. The thermodynamic machinery ties μ to other thermodynamic quantities, including energy but with a crucial caveat: μ represents the change in *free energy* upon particle addition,

$$\mu = \left(\frac{\partial F}{\partial N}\right)_{T,V}, \tag{2.18}$$

not the change in the total internal energy[3]. This explains why μ must always remain strictly below the ground-state energy ϵ_0: the entropy payoff of adding a particle offsets its raw energy cost.

This picture becomes even simpler for trapped gases. Here, the grand canonical ensemble is merely a convenient approximation; the atom number is fixed by whatever we, the experimentalists, load into the trap. The chemical potential μ acts simply as an adjustable knob: we tune it so that summing Eq. (2.17) over all available states yields the desired total atom number N set by experimental reality.

A crucial scenario emerges in the continuum limit (as will be shown below): for a gas with a fixed average atom number N, lowering the temperature causes μ to rise toward ϵ_0. Beyond this threshold, adding further particles only contributes to the system's total energy without altering its configuration (entropy). This can only occur when a Bose-Einstein condensate forms: any newly added atoms fall directly into the ground state, contributing nothing to the entropy.

[3] However, when a constant-volume system enters the condensate phase, added particles go into the ground state because all the excited states have saturated, leaving the system's entropy unchanged. In this case, μ directly tracks the internal energy via $\mu = (\partial U/\partial N)_{S,V}$.

With all these nuances in mind, let us calculate the total atom number of a bosonic gas in free space and determine its connection to μ, aiming to uncover useful thermodynamic properties along the way.

The standard approach is to convert the summation into an integral in the continuum limit: the (non-condensate) population

$$N_T = \sum_i \frac{1}{e^{\beta(\epsilon_i-\mu)}-1} = V\int \frac{\mathrm{d}^3\mathbf{p}}{(2\pi\hbar)^3}\frac{1}{e^{\beta(\epsilon(\mathbf{p})-\mu)}-1}, \tag{2.19}$$

where V is the total volume of the system.

For convenience, the integral over $\mathbf{p}$ is usually converted to an integral over energy by expressing $\mathbf{p}$ in terms of ϵ: $\epsilon(p) = p^2/2m$. The Jacobian factors occurring during this substitution are absorbed into the "density of states" $g(\epsilon)$, yielding

$$N_T = \int \mathrm{d}\epsilon\, g(\epsilon)\frac{1}{e^{\beta(\epsilon-\mu)}-1}. \tag{2.20}$$

For a non-relativistic gas in 3D free space, the density of states

$$g(\epsilon) = \frac{Vm^{3/2}}{\sqrt{2}\pi^2\hbar^3}\epsilon^{1/2}. \tag{2.21}$$

This is a Bose integral, and its result can be expressed in terms of polylogarithm functions:

$$N_T = \frac{V}{\lambda_T^3} g_{3/2}\left(e^{\beta(\mu-\epsilon_0)}\right), \tag{2.22}$$

where ϵ_0 is the ground-state energy, $\lambda_T \equiv \sqrt{2\pi\hbar^2/mk_BT}$ is known as the thermal de Broglie wavelength, and the polylogarithm functions are defined as

$$g_\nu(z) = \frac{1}{\Gamma(\nu)}\int_0^\infty \mathrm{d}x\, \frac{x^{\nu-1}}{z^{-1}e^x-1}. \tag{2.23}$$

Up to this point, the adoption of the continuum limit leading to Eq. (2.22) remains unjustified. In fact, taking the continuum limit introduces a pitfall: when $\mu \to \epsilon_0$, the summation in Eq. (2.19) diverges due to the singular contribution of the ground-state term—a mathematical indication of the formation of a condensate. The integral, however, does not diverge because the density of states $g(\epsilon) \propto \sqrt{\epsilon}$ wipes out this divergence by assigning to it zero weight. Strictly speaking, taking the continuum limit requires the states to be "dense" within a small energy interval and the distribution function to remain well-behaved—an

assumption that breaks down precisely at the condensation threshold. The singularity at the ground state must be isolated and treated separately, yielding the total atom number

$$N = N_0 + N_T, \tag{2.24}$$

where $N_0(T,\mu) = 1/[\exp\beta(\epsilon_0 - \mu) - 1]$ represents the condensate population and N_T accounts for all excited states. When $N < N_T(T, \mu = \epsilon_0)$, no condensate has formed yet and N_0 remains negligible. Once N exceeds the maximum capacity of the excited states, $N > N_T(T, \mu = \epsilon_0)$, the extra atoms spill macroscopically into N_0. The critical temperature marking this onset is

$$k_B T_c = \frac{2\pi\hbar^2}{m}\left[n/g_{3/2}(1)\right]^{2/3}, \tag{2.25}$$

where $g_{3/2}(1) = \zeta(3/2) = 2.612\ldots$, and $n = N/V$ is the 3D density. Below T_c, the condensate population N_0 scales with the total atom number N according to

$$N_0(T,\mu) = N - N_T(T, \mu = \epsilon_0) = N\left[1 - \left(\frac{T}{T_c}\right)^{3/2}\right]. \tag{2.26}$$

To derive other thermodynamic quantities, a more systematic approach starts from the partition function Z defined in Eq. (2.11). Evaluating Z is greatly simplified in the grand canonical ensemble, which allows us to swap the summation and product operations because the occupation number n_i for each state runs from 0 to ∞. As a result,

$$Z = \prod_i \frac{1}{1 - e^{\beta(\mu - \epsilon_i)}} = \prod_i \frac{1}{1 - z\, e^{-\beta\epsilon_i}}, \tag{2.27}$$

where $z = e^{\beta\mu}$ is known as the fugacity. Invoking the standard results of the grand canonical ensemble yields [39, Ch. 8]

$$P = \frac{k_B T}{V} \ln Z(V, z, \beta) \tag{2.28}$$

$$= \frac{k_B T}{\lambda_T^3} g_{5/2}\left(z\, e^{-\beta\epsilon_0}\right), \tag{2.29}$$

$$\frac{E}{V} = -\frac{1}{V}\frac{\partial}{\partial\beta}\bigg|_z \ln Z(V, z, \beta) \tag{2.30}$$

$$= \frac{3}{2}\frac{k_B T}{\lambda_T^3} g_{5/2}\left(z\, e^{-\beta\epsilon_0}\right). \tag{2.31}$$

Note that in the condensate phase, the chemical potential saturates at $\mu = \epsilon_0$; hence, $ze^{-\beta\epsilon_0} = 1$. We immediately see that P is volume-independent, implying that a non-interacting BEC is infinitely compressible. Furthermore, as the temperature approaches zero, the pressure P likewise vanishes. These behaviors stand in stark contrast to those of weakly interacting Bose gases, which always possess finite compressibility and whose pressure remains non-zero at zero temperature due to the interactions, as will be shown in the next section.

All of the preceding discussion assumes the non-interacting Bose gas is homogeneous. In actual experiments, atoms are always trapped, which inherently introduces spatial inhomogeneity. Here, I briefly outline how the preceding results are modified by a harmonic trapping potential [36, Ch. 2].

In the harmonic trap with trap frequencies $\omega_x, \omega_y, \omega_z$, the potential and the resulting energy levels are

$$V(x, y, z) = \frac{1}{2}m(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2), \tag{2.32}$$

$$\epsilon_{n_x,n_y,n_z} = \left(n_x + \frac{1}{2}\right)\hbar\omega_x + \left(n_y + \frac{1}{2}\right)\hbar\omega_y + \left(n_z + \frac{1}{2}\right)\hbar\omega_z. \tag{2.33}$$

Similarly, we can take the continuum limit under the assumption that the thermal energy is much larger than the trap level spacings $k_B T \gg \hbar\omega_i$. Offsetting the zero-point energy, we obtain the density of states

$$g(\epsilon) = \frac{\epsilon^2}{2\hbar^3\omega_x\omega_y\omega_z}. \tag{2.34}$$

Consequently, the transition temperature becomes

$$k_B T_c = \hbar\left(\frac{\omega_x\omega_y\omega_z N}{g_3(1)}\right)^{1/3}, \tag{2.35}$$

where $g_3(1) = \zeta(3) = 1.202\ldots$. The condensate atom number below T_c scales as

$$N_0(T, \mu) = N\left[1 - \left(\frac{T}{T_c}\right)^3\right]. \tag{2.36}$$

Importantly, our argument for the onset of BEC hinges on the saturation of the integral in Eq. (2.22), whose convergence places a strict upper bound on the excited-state population. In lower dimensions—such as a homogeneous 2D gas—the corresponding integral diverges logarithmically at low energy, allowing excited states to accommodate an infi-

nite number of particles at any finite temperature. This absence of BEC in homogeneous 2D systems is formally a consequence of the Mermin-Wagner theorem, which forbids the spontaneous breaking of continuous symmetries with short-range interactions in any dimension $d \leq 2$ at finite temperature. Nonetheless, the presence of a harmonic trap in 2D modifies the density of states to $g(\epsilon) \propto \epsilon$, restoring the convergence of the integral in Eq. (2.22), thereby permitting a true BEC transition. Even though a true BEC is genuinely absent in homogeneous 2D systems, interactions can still enable superfluidity via the Berezinskii-Kosterlitz-Thouless (BKT) mechanism. In a trapped 2D gas, the fragile BEC transition is absent again at finite interactions, but superfluidity is robustly achieved through the BKT process [42].

The experiments described later in this thesis operate in the quasi-2D regime, where tight vertical harmonic confinement restricts atoms to the xy-plane; though, unlike in Ref. [43], it is not sufficiently tight to completely inhibit thermal excitations along the z-direction. Therefore, we expect a genuine BEC transition in our experiments, accompanied by robust superfluidity.

2.3 Weakly interacting Bose-Einstein condensate at zero temperature: static properties

In the last section, we developed the thermodynamics of non-interacting Bose gases, building up important intuition for experimentally cooling toward condensation. Real alkali atoms, however, weakly interact at ultracold temperatures; their interactions are governed by s-wave scattering properties at low energies. In a harmonic trap, repulsive interactions slightly suppress the transition temperature T_c—a simple consequence of reduced peak density—whereas in a homogeneous gas, the shift runs in the opposite direction [44]. We leave these nuances to the literature and turn now to the condensate itself at zero temperature.

The core focus of this thesis is the low-temperature regime (typically $\sim 0.05T_c$), where a weakly interacting Bose gas forms a nearly pure condensate. At these temperatures, the total non-condensate population—from both thermal excitations and quantum fluctuations—remains negligible $\sim 10^{-3}$. Elementary excitations can thus be treated as perturbations to the condensate ground state, a formalism we develop in subsequent sections. Here, we begin with the ground-state properties of a weakly interacting BEC at zero temperature. As shown below, the problem simplifies significantly because (i) weak interactions allow the ground state to be approximated as a product state of single-particle states, justifying a mean-field treatment; (ii) low temperatures suppress thermal depletion, rendering the non-condensate fraction negligible.

Our starting point is the Hamiltonian of a weakly interacting bosonic gas in second-quantized form:

$$\hat{H} = \int \mathrm{d}\mathbf{r}\, \hat{\Psi}^\dagger(\mathbf{r}) \left[-\frac{\hbar^2 \nabla^2}{2m} + U(\mathbf{r}) + \frac{g}{2} \hat{\Psi}^\dagger(\mathbf{r}) \hat{\Psi}(\mathbf{r}) \right] \hat{\Psi}(\mathbf{r}). \tag{2.37}$$

Here, $\hat{\Psi}(\mathbf{r})$ is the bosonic field operator. The first two terms represent the single-particle kinetic energy and the external trapping potential, while the third accounts for contact interactions via the pseudopotential (whose microscopic origin will be detailed in Section 3.1):

$$V(\mathbf{r}_1 - \mathbf{r}_2) = g\delta(\mathbf{r}_1 - \mathbf{r}_2). \tag{2.38}$$

The total atom number operator is defined as $\hat{N} = \int \mathrm{d}\mathbf{r}\, \hat{\Psi}^\dagger(\mathbf{r}) \hat{\Psi}(\mathbf{r})$. One readily verifies that $[\hat{H}, \hat{N}] = 0$, conserving total particle number.

For a Bose-Einstein condensate, the immediate question is: what is its many-body wavefunction? Since atoms interact weakly, we adopt the Hartree-Fock or mean-field approximation: all N atoms occupy the same single-particle state, yielding a symmetrized product state of the single-particle wavefunction $\phi(\mathbf{r})$:

$$\Psi(\mathbf{r}_1, \mathbf{r}_2, \cdots, \mathbf{r}_N) = \prod_{i=1}^{N} \phi(\mathbf{r}_i) \tag{2.39}$$

or, in the language of second quantization [45]:

$$|\Psi\rangle = \frac{1}{\sqrt{N!}} \left(\int \mathrm{d}\mathbf{r} \phi(\mathbf{r}) \hat{\Psi}^\dagger(\mathbf{r}) \right)^N |0\rangle \tag{2.40}$$

where $|0\rangle$ is the vacuum state and ϕ is normalized: $\int \mathrm{d}\mathbf{r} |\phi(\mathbf{r})|^2 = 1$. This formalism is equivalent to the single-particle density matrix $\langle \mathbf{r}' | \rho_1 | \mathbf{r} \rangle = \langle \hat{\Psi}^\dagger(\mathbf{r}) \hat{\Psi}(\mathbf{r}') \rangle$ possessing a single dominant macroscopic eigenvalue N corresponding to eigenvector $|\phi\rangle$, with all other eigenvalues vanishingly small—the precise criterion for complete condensation defined by Penrose and Onsager [10]. While this picture generally fails in strongly interacting quantum systems, such as superfluid helium-4—where strong quantum depletion populates other single-particle modes—it provides an excellent description for the weakly interacting gases in our experiments.

The energy of this product state is

$$E = \langle \Psi | \hat{H} | \Psi \rangle = N \int \mathrm{d}\mathbf{r} \left(\frac{\hbar^2}{2m} |\nabla \phi(\mathbf{r})|^2 + U(\mathbf{r}) |\phi(\mathbf{r})|^2 \right) + \frac{N(N-1)}{2} g \int \mathrm{d}\mathbf{r} |\phi(\mathbf{r})|^4. \tag{2.41}$$

Here, an integration by parts was performed on the kinetic energy term, assuming ϕ vanishes at infinity. The factor $N(N-1)/2$ simply counts the total number of interacting atom pairs.

For a non-interacting gas ($g = 0$), $\phi(\mathbf{r})$ is simply the single-particle ground state: a constant in a homogeneous system or a half-period sinusoidal function in a box potential. Interactions provide a source of *pressure*, flattening the density peak in $\phi(\mathbf{r})$. We determine the optimal $\phi(\mathbf{r})$ by minimizing the energy functional Eq. (2.41) subject to the normalization constraint $\int |\phi(\mathbf{r})|^2 = 1$. This yields the variational problem:

$$\delta \left[E - \mu N \int \mathrm{d}\mathbf{r} |\phi(\mathbf{r})|^2 \right] = 0 \tag{2.42}$$

where μ is the Lagrange multiplier. In the limit $N \gg 1$ ($N - 1 \approx N$), varying against $\delta\phi^*$ yields the celebrated Gross-Pitaevskii equation (GPE)[4] [46]:

$$-\frac{\hbar^2}{2m}\nabla^2\phi(\mathbf{r}) + [U(\mathbf{r}) + gN|\phi(\mathbf{r})|^2]\phi(\mathbf{r}) = \mu\phi(\mathbf{r}). \tag{2.43}$$

To identify the physical meaning of μ, consider the energy change associated with adding a single particle ($N \to N + 1$), noting that the variations in ϕ contribute only at second order near the energy minimum:

$$\Delta E = \int \mathrm{d}\mathbf{r}\left(\frac{\hbar^2}{2m}|\nabla\phi(\mathbf{r})|^2 + U(\mathbf{r})|\phi(\mathbf{r})|^2\right) + gN\int \mathrm{d}\mathbf{r}|\phi(\mathbf{r})|^4 \tag{2.44}$$

$$= \mu. \tag{2.45}$$

We immediately recognize that $\mu = \partial E/\partial N$ is the chemical potential—the energy required to add one particle to the condensate at zero temperature, as discussed in Section 2.2.

Thomas-Fermi limit

The Gross-Pitaevskii equation (2.43) rarely admits analytical solutions, generally necessitating numerical methods. Analytical treatment, however, is possible in two limiting regimes. One trivial limit occurs as the density $n(\mathbf{r}) = N|\phi(\mathbf{r})|^2 \to 0$, where Eq. (2.43) reduces to the single-particle Schrödinger equation. A far more experimentally relevant regime arises when the repulsive interaction energy ($g > 0$) dominates the kinetic energy:

$$gN|\phi(\mathbf{r})|^4 \gg \frac{\hbar^2|\nabla\phi(\mathbf{r})|^2}{2m}. \tag{2.46}$$

This is the Thomas-Fermi (TF) limit[5], in which the kinetic term is neglected [46]:

$$U(\mathbf{r}) + gN|\phi(\mathbf{r})|^2 = \mu, \tag{2.47}$$

[4] Careful readers will notice the kinetic term reverts to its standard form before the integration by parts. This is a standard result of functional differentiation that involves ∇:

$$\frac{\delta\int|\nabla'\phi(\mathbf{r}')|^2}{\delta\phi^*(\mathbf{r})} = \int \mathrm{d}\mathbf{r}'\frac{|\nabla'\phi(\mathbf{r}') + \nabla'\delta\phi(\mathbf{r})\delta(\mathbf{r}-\mathbf{r}')|^2 - |\nabla'\phi(\mathbf{r}')|^2}{\delta\phi^*(\mathbf{r})}$$
$$= \int \mathrm{d}\mathbf{r}'\nabla'\delta(\mathbf{r}-\mathbf{r}')\cdot\nabla'\phi(\mathbf{r}') = -\nabla^2\phi(\mathbf{r}).$$

yielding the wavefunction:

$$\phi_{\mathrm{TF}}(\mathbf{r}) = \begin{cases} \left(\frac{\mu - U(\mathbf{r})}{gN}\right)^{1/2} & \text{if } \mu \geq U(\mathbf{r}), \\ 0 & \text{if } \mu < U(\mathbf{r}), \end{cases} \tag{2.48}$$

known as the Thomas-Fermi wavefunction. Importantly, the Thomas-Fermi approximation breaks down at the condensate boundary where $\mu = U(\mathbf{r})$. There, Eq. (2.48) predicts a sharp drop in the condensate density to zero, instead of the actual smooth transition produced by the balancing between the zero-point kinetic energy and the trapping potential. We can, in fact, extract the characteristic length scale ξ over which the wavefunction smoothly heals from zero near a hard boundary (such as a square well) by examining on what scale the Thomas-Fermi approximation fails:

$$gn_0 \sim \frac{\hbar^2}{2m\xi^2}, \tag{2.49}$$

where n_0 is the bulk condensate density far from the walls. This yields the *healing length*:

$$\xi^2 = \frac{\hbar^2}{2mgn_0}. \tag{2.50}$$

This sets the minimal scale over which the condensate wavefunction can respond to a localized perturbation in the trapping potential; variations below ξ are smoothed out by kinetic pressure.

Reduction of dimension

Consider a common experimental geometry: a condensate tightly confined along one direction (z) by a strong harmonic potential, restricting the low-energy dynamics—such as sound propagation—to the two-dimensional (xy) plane. This separation of the energy scales largely decouples low-energy in-plane dynamics from the excitations along the tightly confined direction. In this case, constructing an effective 2D description dramatically reduces the numerical complexity of the full 3D theory.

[5] This name is borrowed from the classic "Thomas-Fermi model" of the electrons in atoms, though the connection is largely historical. While the original theory uses a semiclassical expression for the kinetic energy, the Thomas-Fermi limit in BECs neglects the kinetic energy term altogether.

First, we assume the condensate wavefunction is separable: $\phi(x,y,z) = \phi_\perp(x,y)\phi_z(z)$. Substituting this ansatz into Eq. (2.41) (with $N-1 \approx N$, $\nabla_\perp \equiv \hat{\mathbf{x}}\partial_x + \hat{\mathbf{y}}\partial_y$) produces:

$$E = N\int \mathrm{d}\mathbf{r}\left(\frac{\hbar^2}{2m}|\phi_z\nabla_\perp\phi_\perp + \hat{\mathbf{z}}\,\phi_\perp\partial_z\phi_z|^2 + [V_\perp(x,y)+V_z(z)]|\phi_\perp\phi_z|^2\right) \tag{2.51}$$

$$+\frac{N^2}{2}g\int \mathrm{d}\mathbf{r}\,|\phi_\perp\phi_z|^4. \tag{2.52}$$

Here, the external potential separates into a strong vertical trap $V_z(z)$ and a weak in-plane potential $V_\perp(x,y)$, so that $U(\mathbf{r}) = V_\perp(x,y) + V_z(z)$. Expanding the kinetic term—where the cross term vanishes by orthogonality—yields:

$$E = N\int \mathrm{d}\mathbf{r}\left(\frac{\hbar^2}{2m}(|\phi_z\nabla_\perp\phi_\perp|^2 + |\phi_\perp\partial_z\phi_z|^2) + (V_\perp+V_z)|\phi_\perp\phi_z|^2\right) + \frac{N^2}{2}g\int \mathrm{d}\mathbf{r}\,|\phi_\perp\phi_z|^4 \tag{2.53}$$

$$= N\int \mathrm{d}\mathbf{r}\left[\phi_z^2\left(\frac{\hbar^2}{2m}|\nabla_\perp\phi_\perp|^2 + V_\perp|\phi_\perp|^2 + \frac{N}{2}g|\phi_z|^2|\phi_\perp|^4\right)\right. \tag{2.54}$$

$$\left.+\phi_\perp^2\left(\frac{\hbar^2}{2m}|\partial_z\phi_z|^2 + V_z|\phi_z|^2\right)\right].$$

Since V_z is tight, the second term dominates the total energy. Yet the interaction term has been arbitrarily grouped alongside the weak in-plane terms. Assuming this grouping is valid, as in cases where V_z is much larger than the interaction energy, we minimize the z-direction energy contribution (the second term) first:

$$-\frac{\hbar^2}{2m}\frac{\partial^2}{\partial z^2}\phi_z(z) + V_z(z)\phi_z(z) = \epsilon_z\phi_z(z), \tag{2.55}$$

which is simply the one-dimensional Schrödinger equation, with the ground-state energy ϵ_z serving as the Lagrange multiplier that enforces normalization of ϕ_z. Now the system energy becomes:

$$E = N\int \mathrm{d}\mathbf{r}\,\phi_z^2\left(\frac{\hbar^2}{2m}|\nabla_\perp\phi_\perp|^2 + V_\perp|\phi_\perp|^2 + \epsilon_z|\phi_\perp|^2\right) + \int \mathrm{d}z\,\phi_z^4\frac{N^2}{2}g\int \mathrm{d}x\mathrm{d}y|\phi_\perp|^4. \tag{2.56}$$

Minimizing this functional with respect to $\phi_\perp(x,y)$ yields an effective two-dimensional Gross-Pitaevskii equation:

$$-\frac{\hbar^2}{2m}\nabla_\perp^2\phi_\perp(x,y) + [V_\perp(x,y) + \alpha g N|\phi_\perp(x,y)|^2]\phi_\perp(x,y) = (\mu - \epsilon_z)\phi_\perp(x,y) \tag{2.57}$$

where the effective 2D interaction parameter is rescaled by the factor $\alpha = \int \mathrm{d}z\ \phi_z^4$. In typical experiments, V_z is harmonic with trap frequency ω_z. The z-direction wavefunction ϕ_z then reduces to the harmonic-oscillator ground state, yielding $\alpha = 1/(\sqrt{2\pi}\, a_z)$ with oscillator length $a_z = \sqrt{\hbar/m\omega_z}$.

It is now clear that the initial grouping of the interaction term with the in-plane dynamics is therefore valid provided the interaction energy is much smaller than $N\epsilon_z$, or equivalently, $\mu - \epsilon_z \ll \epsilon_z$.

In the opposite limit ($\mu \gg \epsilon_z$), the strong interaction regime requires another dimensional reduction approach, where ϕ_z transitions from a Gaussian to a Thomas-Fermi profile [21, Chapter 23]. The experiments presented in this thesis operate in a crossover regime where $\mu \sim \epsilon_z$, precluding either limit. Nevertheless, because the relevant excitation energies satisfy $\epsilon \ll \mu \sim \epsilon_z$, the effective 2D GPE (2.57) remains robust. In this regime, the prefactor α requires modification, and the eigenvalue on the right-hand side loses its clean interpretation as $\mu - \epsilon_z$.

2.4 Weakly interacting Bose-Einstein condensate at zero temperature: dynamical properties

In the previous section, we derived the ground state $|\Psi\rangle$ (2.40)—a pure N-particle Fock state—for a weakly interacting Bose gas. Tracking the time evolution of a Fock state containing 10^6 atoms is a highly non-trivial many-body problem. Studying the dynamics of such systems requires a suitable approximation.

Before addressing the problem of dynamics, let us consider a subtle question: what is the order parameter of a Bose-Einstein condensate? In the classical Landau paradigm, continuous phase transitions signal spontaneous symmetry breaking of the underlying Hamiltonian[6]. Ever since Bogoliubov's treatment of the weakly interacting Bose gas and Anderson's broken-symmetry formulation of superfluidity, the condensate phase transition has been interpreted as the spontaneous breaking of the global $U(1)$ symmetry of the Hamiltonian, Eq. (2.37). Within this framework, the condensate state is not a fixed-N Fock state, but a coherent state $|\phi\rangle$—a superposition of Fock states with atom numbers extending to infinity—which acts as an eigenstate of the annihilation operator $\hat{\Psi}(\mathbf{r}, t)$[7]:

$$\hat{\Psi}(\mathbf{r}, t)|\phi\rangle = \sqrt{N}\phi(\mathbf{r}, t)|\phi\rangle, \tag{2.58}$$

The broken "global $U(1)$ symmetry" refers to the phase θ of the single-particle wavefunction $\phi(\mathbf{r}) = |\phi(\mathbf{r})|e^{i\theta}$, and this θ is believed to have been spontaneously picked by the condensate when crossing the condensate transition. The complex field $\sqrt{N}\phi(\mathbf{r})$ plays the role of the macroscopic order parameter of the condensate.

This picture, though mathematically convenient, is problematic and has sparked extensive debate within the BEC community [47]—a point we shall clarify later in this section. For the moment, however, we adopt this symmetry-breaking picture to demonstrate how it remarkably streamlines the calculation of condensate dynamics.

To this end, we write the Heisenberg equation of motion for the field operator $\hat{\Psi}(\mathbf{r}, t)$:

$$i\hbar\frac{\partial}{\partial t}\hat{\Psi}(\mathbf{r}, t) = [\hat{\Psi}(\mathbf{r}, t), \hat{H}]. \tag{2.59}$$

[6] However, this statement is not always true—the BKT transition serves as a notable counterexample.

[7] Following our convention set in the last section, ϕ is normalized to unity; the atom number expectation value N of the state is explicitly expressed by the factor $\sqrt{N}$.

Expanding the commutator with the Hamiltonian (2.37) via standard bosonic commutation relations yields:

$$i\hbar\frac{\partial}{\partial t}\hat{\Psi}(\mathbf{r},t) = \left(-\frac{\hbar^2}{2m}\nabla^2 + U(\mathbf{r},t)\right)\hat{\Psi}(\mathbf{r},t) + g\hat{\Psi}^\dagger(\mathbf{r},t)\hat{\Psi}(\mathbf{r},t)\hat{\Psi}(\mathbf{r},t). \tag{2.60}$$

Taking the expectation value of both sides with respect to $|\phi\rangle$ and assuming $|\phi\rangle$ remains an eigenstate of $\hat{\Psi}(\mathbf{r},t)$ with eigenvalue $\sqrt{N}\phi(\mathbf{r},t)$ throughout the evolution, we obtain:

$$i\hbar\frac{\partial}{\partial t}\phi(\mathbf{r},t) = \left(-\frac{\hbar^2}{2m}\nabla^2 + U(\mathbf{r},t)\right)\phi(\mathbf{r},t) + gN|\phi(\mathbf{r},t)|^2\phi(\mathbf{r},t). \tag{2.61}$$

This is commonly known as the *time-dependent Gross-Pitaevskii equation*, governing the mean-field evolution of the condensate order parameter. For a stationary ground state $\phi(\mathbf{r},0) = \phi_0(\mathbf{r})$ obtained from Eq. (2.43), the time dependence simplifies to a pure phase rotation

$$\phi_0(\mathbf{r},t) = e^{-i\mu t/\hbar}\phi_0(\mathbf{r}). \tag{2.62}$$

It is now time to critically examine the conceptual flaws of this $U(1)$ symmetry breaking formalism. First and foremost, condensates prepared in real experiments are isolated in vacuum with a fixed atom number; assuming the condensate state overlaps even infinitesimally with states of infinite particle number is flatly incorrect (and will be vehemently rejected by all experimentalists). Second, because a coherent state lacks a definite particle number, it cannot possibly represent a stationary state of a number-conserving Hamiltonian (2.37). Unphysical pathologies immediately emerge. For example, quantum phase diffusion causes phase fluctuations to grow rapidly in time [48], which invalidates the linearized Bogoliubov expansion on the ground state (discussed below) on timescales as short as ~0.2 s under typical experimental conditions [49].

These pathologies do not imply that $\sqrt{N}\phi(\mathbf{r})$ fails as an order parameter. Rather, it *is* a physical order parameter in Penrose and Onsager's sense [10]—they rigorously define the condensate order parameter as the eigenfunction of the single-particle density matrix $\langle\mathbf{r}'|\hat{\rho}_1|\mathbf{r}\rangle = \langle\hat{\Psi}^\dagger(\mathbf{r})\hat{\Psi}(\mathbf{r}')\rangle$, which is exactly $\sqrt{N}\phi(\mathbf{r})$, corresponding to the eigenvalue N. While $\phi(\mathbf{r})$ possesses an overall phase θ, the absolute value of this phase in a condensate is fundamentally unobservable, and thus carries no physical meaning.

The primary argument historically invoked in favor of the $U(1)$ symmetry-breaking picture—beyond the historical precedent and the simplicity it offers for calculations—came from the landmark interference experiment of two isolated condensates [50] performed

in the early days of atomic BEC. In that experiment, interference fringes were observed when two independently prepared condensates expanded and overlapped in space. Initially, proponents of the $U(1)$ symmetry-breaking theory interpreted these fringes as direct evidence that each condensate possesses a well-defined phase associated with a coherent state, which sets the spatial phase of the interference fringes. However, the fringe offset in Ref. [50] varied randomly from shot to shot, demonstrating that if phase symmetry was "spontaneously broken", it was "picked" anew in every single realization. Crucially, Castin and Dalibard [51] showed that a coherent state with an unknown phase is statistically indistinguishable from a Poissonian mixture of Fock states. More profoundly, they demonstrated that even for two initially pure Fock states with no well-defined relative phase, the process of particle detection itself establishes a well-defined relative phase in every shot. Consequently, not only do pure Fock states exhibit identical interference fringes with random shot-to-shot phase offsets, but no measurement whatsoever can distinguish between the two descriptions. In a similar vein, Castin and Dum [49] constructed a fully number-conserving formalism of condensate dynamics. By employing a rigorous expansion in the powers of the square root of non-condensate atoms, they rederived Eq. (2.61) and the Bogoliubov-de Gennes equations (discussed below) without ever violating atom-number conservation. All of the above seems to assure us: no matter the nature of the condensate state being a Fock state or a coherent state with an unknown phase, the experimental observations should remain identical—at least if only macroscopic hydrodynamics is concerned.

In hindsight, this equivalence should come as no surprise. From the very moment we adopt mean-field theory, we effectively reduce the intricate quantum many-body system to a single-particle description. As a result, all multi-particle quantum correlations (and the associated theoretical difficulties) are systematically discarded. Indeed, setting aside the second-quantized formalism entirely, the time-dependent GPE (2.61) can also be derived directly as the stationary condition (Euler-Lagrange equation) of a classical field action [21, Ch. 5], [52, Ch. 1]:

$$\delta\left[-i\hbar N\int \mathrm{d}\mathbf{r}\,\mathrm{d}t\phi^*\partial_t\phi+\int \mathrm{d}tE\right]=0. \tag{2.63}$$

This classical field perspective, however, should not be mistaken for a rejection of the quantum nature of Bose-Einstein condensates. True quantum many-body effects become prominent under *extreme* experimental conditions—such as systems with very low atom numbers ($\sim 10^3$), strong interactions, or when probing observables that directly reflect quantum entanglement and fluctuations (though the difficulty and feasibility of measur-

ing those quantities are another matter). All experimental work presented in this thesis focuses on the low-energy hydrodynamics of the condensate—a regime where many-body corrections are vanishingly small and the GPE framework remains remarkably accurate.

2.5 Binary Bose-Einstein condensate

The preceding section treated the condensates with a single component only, which are governed by a scalar order parameter $\phi(\mathbf{r})$. Real atoms, however, usually possess internal hyperfine structures—an internal degree of freedom we refer to as *(pseudo-)spin*. The experiment platform central to this thesis is a binary BEC composed of ^{23}Na atoms occupying two immiscible hyperfine states. Consequently, the order parameter expands from a scalar to a two-component vector. In this section, we formulate the mean-field framework adopted to describe the ground states and the phase separation of the binary quantum fluids [30, 31].

We begin by constructing a general Hamiltonian for a dual-species bosonic system:

$$\hat{H} = \int \mathrm{d}^3\mathbf{r}\Bigg\{\hat{\Psi}_\uparrow^\dagger(\mathbf{r})\left[-\frac{\hbar^2}{2m}\nabla^2 + U_\uparrow(\mathbf{r})\right]\hat{\Psi}_\uparrow(\mathbf{r}) + \hat{\Psi}_\downarrow^\dagger(\mathbf{r})\left[-\frac{\hbar^2}{2m}\nabla^2 + U_\downarrow(\mathbf{r})\right]\hat{\Psi}_\downarrow(\mathbf{r}) \tag{2.64}$$

$$+ \frac{g_{\uparrow\uparrow}}{2}\hat{\Psi}_\uparrow^\dagger(\mathbf{r})\hat{\Psi}_\uparrow^\dagger(\mathbf{r})\hat{\Psi}_\uparrow(\mathbf{r})\hat{\Psi}_\uparrow(\mathbf{r}) + \frac{g_{\downarrow\downarrow}}{2}\hat{\Psi}_\downarrow^\dagger(\mathbf{r})\hat{\Psi}_\downarrow^\dagger(\mathbf{r})\hat{\Psi}_\downarrow(\mathbf{r})\hat{\Psi}_\downarrow(\mathbf{r}) \tag{2.65}$$

$$+ g_{\uparrow\downarrow}\hat{\Psi}_\uparrow^\dagger(\mathbf{r})\hat{\Psi}_\downarrow^\dagger(\mathbf{r})\hat{\Psi}_\downarrow(\mathbf{r})\hat{\Psi}_\uparrow(\mathbf{r})\Bigg\}. \tag{2.66}$$

Here, the labels $\uparrow$ and $\downarrow$ distinguish the two spins, which we conveniently refer to as "spin-$\uparrow$" and "spin-$\downarrow$", bearing in mind that this pseudo-spin reflects hyperfine levels rather than electronic spin. Interactions within identical species are governed by the "intra-spin" interaction strengths $g_{\uparrow\uparrow}, g_{\downarrow\downarrow}$, whereas cross-species interactions are quantified by the "inter-spin" interaction strength $g_{\uparrow\downarrow}$.

In the absence of population exchange between spins, the particle number of each spin is independently conserved, which is enforced by introducing two chemical potentials, $\mu_\uparrow$ and $\mu_\downarrow$. Following the identical variational approach from the last section, we obtain the time-independent Gross-Pitaevskii equations:

$$\mu_\uparrow\phi_\uparrow = \left[-\frac{\hbar^2}{2m}\nabla^2 + U_\uparrow(\mathbf{r})\right]\phi_\uparrow + g_{\uparrow\uparrow}N_\uparrow|\phi_\uparrow|^2\phi_\uparrow + g_{\uparrow\downarrow}N_\downarrow|\phi_\downarrow|^2\phi_\uparrow \tag{2.67}$$

$$\mu_\downarrow\phi_\downarrow = \left[-\frac{\hbar^2}{2m}\nabla^2 + U_\downarrow(\mathbf{r})\right]\phi_\downarrow + g_{\downarrow\downarrow}N_\downarrow|\phi_\downarrow|^2\phi_\downarrow + g_{\uparrow\downarrow}N_\uparrow|\phi_\uparrow|^2\phi_\downarrow, \tag{2.68}$$

where the individual wavefunctions[8] obey the normalization $\int \mathrm{d}^3\mathbf{r}\ |\phi_{\uparrow,\downarrow}(\mathbf{r})|^2 = 1$. The corresponding time-dependent GPEs follow by replacing $\mu_{\uparrow,\downarrow}\phi_{\uparrow,\downarrow}$ with $i\hbar\partial_t\phi_{\uparrow,\downarrow}$.

To study the role of interactions, let us first consider a spatially homogeneous system ($U_{\uparrow,\downarrow} = 0$). The central question is: how does the inter-species interaction $g_{\uparrow\downarrow}$ shape the

[8] The emergence of two order parameters $\phi_{\uparrow,\downarrow}$ is commonly associated with "the broken $U(1) \times U(1)$ symmetry".

ground state? If $g_{\uparrow\downarrow} = 0$, the two spins act as two independent condensates, both settling into a homogeneous configuration. A moderate inter-spin attraction $g_{\uparrow\downarrow} < 0$ preserves the uniform overlap, until attraction overpowers the intra-spin repulsion and destabilizes the mixture[9]. Conversely, as the inter-spin repulsion ($g_{\uparrow\downarrow} > 0$) is dialed up, at some point, the energy penalty of mutual overlap exceeds the energy cost of spatial displacement, forcing the system to phase separate.

To place this instability on quantitative footing, we evaluate the mean-field energy for the homogeneous configuration:

$$E_{\mathrm{PM}} = \frac{1}{2V}\Big(g_{\uparrow\uparrow}N_{\uparrow}^2 + g_{\downarrow\downarrow}N_{\downarrow}^2 + 2g_{\uparrow\downarrow}N_{\uparrow}N_{\downarrow}\Big), \tag{2.69}$$

where the subscript PM stands for *paramagnetic*. We examine the stability of the paramagnetic configuration by introducing a small density displacement within two spins $n_{\uparrow,\downarrow}(\mathbf{r}) = n_{\uparrow 0,\downarrow 0} + \delta n_{\uparrow,\downarrow}(\mathbf{r})$. Enforcing $\int \mathrm{d}^3\mathbf{r}\delta n_{\uparrow,\downarrow} = 0$ causes all first-order terms to vanish, leaving the second-order energy shift:

$$\Delta E_{\mathrm{PM}} = \int \mathrm{d}^3\mathbf{r}\frac{1}{2}\Big(g_{\uparrow\uparrow}\delta n_{\uparrow}^2 + g_{\downarrow\downarrow}\delta n_{\downarrow}^2 + 2g_{\uparrow\downarrow}\delta n_{\uparrow}\delta n_{\downarrow}\Big). \tag{2.70}$$

The paramagnetic configuration is energetically stable if and only if the quadratic form ΔE_{PM} is positive-definite, corresponding to the conditions $g_{\uparrow\uparrow}, g_{\downarrow\downarrow} > 0$ and $g_{\uparrow\uparrow}g_{\downarrow\downarrow} > g_{\uparrow\downarrow}^2$. While the first condition ensures the stability of individual components, the second dictates the miscibility: falling below $g_{\uparrow\downarrow} < -\sqrt{g_{\uparrow\uparrow}g_{\downarrow\downarrow}}$ leads to the condensate collapse, whereas exceeding $g_{\uparrow\downarrow} > \sqrt{g_{\uparrow\uparrow}g_{\downarrow\downarrow}}$ drives the spatial separation of two spins.

To study the ground state configuration of the phase-separated case, we neglect any spatial overlap at the interface. [10] The mean-field energy of the phase-separated—or ferromagnetic (FM)—state reads $E_{\mathrm{FM}} = g_{\uparrow\uparrow}N_{\uparrow}^2/2V_{\uparrow} + g_{\downarrow\downarrow}N_{\downarrow}^2/2V_{\downarrow}$. Minimizing it with respect to the volume ratio $V_{\uparrow}/V_{\downarrow}$ determines the optimal volume partitioning:

$$V_{\uparrow} = \frac{1}{1 + \sqrt{g_{\downarrow\downarrow}}N_{\downarrow}/\sqrt{g_{\uparrow\uparrow}}N_{\uparrow}}V, \tag{2.71}$$

$$V_{\downarrow} = \frac{1}{1 + \sqrt{g_{\uparrow\uparrow}}N_{\uparrow}/\sqrt{g_{\downarrow\downarrow}}N_{\downarrow}}V, \tag{2.72}$$

[9] However, the collapse can be prevented within a narrow window: the Lee-Huang-Yang correction to the mean-field condensate energy provides a repulsive pressure that balances the inter-spin attraction, resulting in self-bound "quantum droplets" [53].

[10] The energy associated with this overlap along the $(d-1)$-dimensional interface is almost always tiny compared to the d-dimensional bulk energy for this global volume calculation. However, it is crucial when determining the structure of the interface.

or, in terms of particle densities $n_{\uparrow,\downarrow} = N_{\uparrow,\downarrow}|\phi_{\uparrow,\downarrow}|^2 = N_{\uparrow,\downarrow}/V_{\uparrow,\downarrow}$:

$$\sqrt{g_{\uparrow\uparrow}}n_{\uparrow} = \sqrt{g_{\downarrow\downarrow}}n_{\downarrow}. \tag{2.73}$$

Substituting these equilibrium volumes back into the energy yields:

$$E_{\mathrm{FM}} = \frac{1}{2V}\left(g_{\uparrow\uparrow}N_{\uparrow}^2 + g_{\downarrow\downarrow}N_{\downarrow}^2 + 2\sqrt{g_{\uparrow\uparrow}g_{\downarrow\downarrow}}N_{\uparrow}N_{\downarrow}\right). \tag{2.74}$$

By examining the phase-separation condition $E_{\mathrm{PM}} - E_{\mathrm{FM}} > 0$, we rederive:

$$g_{\uparrow\downarrow} > \sqrt{g_{\uparrow\uparrow}g_{\downarrow\downarrow}}. \tag{2.75}$$

The density balance condition (2.73) can be understood equally well through a mechanical approach: equilibrium across the interface requires equal pressure $P = -(\partial U/\partial V)_S$. For a uniform condensate at $T = 0$,

$$P = -\frac{\partial E}{\partial V} = \frac{1}{2}gn^2 = \frac{\mu^2}{2g}. \tag{2.76}$$

Equating the pressures of the two spin domains immediately recovers Eq. (2.73).

The spatial separation in the ferromagnetic state yields domain walls that mimic fluid-fluid interfaces in immiscible liquids. The dynamics and energy associated with these quantum interfaces form the primary focus of Chapter 5.

2.6 Excitation spectrum of a weakly interacting Bose-Einstein condensate

In Section 2.4, we established the mean-field condensate ground state by approximating $\hat{\Psi}(\mathbf{r})$ with a scalar wavefunction $\phi(\mathbf{r})$ and minimizing the system energy. The logical next step is to examine the small perturbations about this mean-field background. We will show that after expanding to second order, the quadratic Hamiltonian is diagonalized into quasi-particle modes, whose spatial amplitudes are governed by the Bogoliubov-de Gennes equations.

We begin by analyzing a uniform condensate where momentum $\mathbf{p}$ remains a good quantum number; here, the low-energy Bogoliubov excitations manifest unambiguously as phonons. After establishing the basic formalism, we will extend the framework to inhomogeneous (trapped) gases and, ultimately, to binary condensate systems.

The homogeneous case

Following Bogoliubov's treatment, we first separate the condensate mode $\hat{a}_0 = \int \mathrm{d}^3\mathbf{r}\phi_0^*(\mathbf{r})\hat{\Psi}(\mathbf{r})$—the zeroth-order contribution—from the full field operator $\hat{\Psi}(\mathbf{r})$. In the plane-wave basis $\langle \mathbf{r}|\mathbf{p}\rangle = \exp(i\mathbf{p}\cdot\mathbf{r}/\hbar)/\sqrt{V}$:

$$\hat{\Psi}(\mathbf{r}) = \frac{1}{\sqrt{V}}\sum_{\mathbf{p}} e^{i\mathbf{p}\cdot\mathbf{r}/\hbar}\hat{a}_{\mathbf{p}}, \quad \hat{a}_{\mathbf{p}} = \frac{1}{\sqrt{V}}\int \mathrm{d}^3\mathbf{r} e^{-i\mathbf{p}\cdot\mathbf{r}/\hbar}\hat{\Psi}(\mathbf{r}). \tag{2.77}$$

The condensate mode corresponds to $\hat{a}_0 \sim \sqrt{N}$, and all other $\hat{a}_{\mathbf{p}\neq 0}$ are small perturbations around it.

Truncating the Hamiltonian (2.37) to second order in $\hat{a}_{\mathbf{p}\neq 0}$ yields:

$$\hat{H} = \frac{g}{2V}\hat{a}_0^\dagger\hat{a}_0^\dagger\hat{a}_0\hat{a}_0 + \sum_{\mathbf{p}\neq 0}\left[\frac{p^2}{2m}\hat{a}_{\mathbf{p}}^\dagger\hat{a}_{\mathbf{p}} + \frac{g}{2V}(4\hat{a}_0^\dagger\hat{a}_{\mathbf{p}}^\dagger\hat{a}_0\hat{a}_{\mathbf{p}} + \hat{a}_{\mathbf{p}}^\dagger\hat{a}_{-\mathbf{p}}^\dagger\hat{a}_0\hat{a}_0 + \hat{a}_0^\dagger\hat{a}_0^\dagger\hat{a}_{\mathbf{p}}\hat{a}_{-\mathbf{p}})\right]. \tag{2.78}$$

Following the same spirit, we approximate $\hat{a}_0$ and $\hat{a}_0^\dagger$ by $\sqrt{N}$ in the second term. Importantly, we carefully expand the first term with $\hat{a}_0^\dagger\hat{a}_0 + \sum_{\mathbf{p}\neq 0}\hat{a}_{\mathbf{p}}^\dagger\hat{a}_{\mathbf{p}} = \hat{N}$, retaining terms up to the quadratic order:

$$\hat{a}_0^\dagger\hat{a}_0^\dagger\hat{a}_0\hat{a}_0 = \hat{a}_0^\dagger\hat{a}_0(\hat{a}_0^\dagger\hat{a}_0 - 1) \approx N^2 - 2N\sum_{\mathbf{p}\neq 0}\hat{a}_{\mathbf{p}}^\dagger\hat{a}_{\mathbf{p}}. \tag{2.79}$$

The Hamiltonian now reduces to a quadratic form:

$$\hat{H} = \frac{gN^2}{2V} + \sum_{\mathbf{p}\neq 0}\left[\frac{p^2}{2m}\hat{a}^\dagger_{\mathbf{p}}\hat{a}_{\mathbf{p}} + \frac{gn}{2}(2\hat{a}^\dagger_{\mathbf{p}}\hat{a}_{\mathbf{p}} + \hat{a}^\dagger_{\mathbf{p}}\hat{a}^\dagger_{-\mathbf{p}} + \hat{a}_{\mathbf{p}}\hat{a}_{-\mathbf{p}})\right], \tag{2.80}$$

where $n \equiv N/V$ is the particle density. The actual diagonalization is performed by introducing quasi-particle creation and annihilation operators:

$$\hat{a}_{\mathbf{p}} = u_p\hat{b}_{\mathbf{p}} - v_p^*\hat{b}^\dagger_{-\mathbf{p}}, \quad \hat{a}^\dagger_{\mathbf{p}} = u_p^*\hat{b}^\dagger_{\mathbf{p}} - v_p\hat{b}_{-\mathbf{p}}. \tag{2.81}$$

This mapping is the well-known *Bogoliubov transformation.* Isotropy of the system ensures u_p and v_p depend solely on the magnitude of the momentum $p = |\mathbf{p}|$. To preserve the bosonic commutation relation $[\hat{b}_{\mathbf{p}}, \hat{b}^\dagger_{\mathbf{p}'}] = \delta_{\mathbf{p},\mathbf{p}'}$, the coefficients u_p and v_p must satisfy

$$|u_p|^2 - |v_p|^2 = 1. \tag{2.82}$$

Conversely, the *inverse Bogoliubov transformation* reads:[11]

$$\hat{b}_{\mathbf{p}} = u_p^*\hat{a}_{\mathbf{p}} + v_p^*\hat{a}^\dagger_{-\mathbf{p}}, \quad \hat{b}^\dagger_{\mathbf{p}} = u_p\hat{a}^\dagger_{\mathbf{p}} + v_p\hat{a}_{-\mathbf{p}}. \tag{2.83}$$

Requiring all the off-diagonal terms to vanish diagonalizes the Hamiltonian into a set of non-interacting quasi-particles:

$$\hat{H} = E_0 + \sum_{\mathbf{p}\neq 0}\epsilon_{\mathbf{p}}\hat{b}^\dagger_{\mathbf{p}}\hat{b}_{\mathbf{p}}, \tag{2.84}$$

which yields both the amplitudes u_p, v_p and the energy spectrum:

$$u_p = \sqrt{\frac{p^2/2m + gn}{2\epsilon_p} + \frac{1}{2}}, \quad v_p = \sqrt{\frac{p^2/2m + gn}{2\epsilon_p} - \frac{1}{2}}, \tag{2.85}$$

$$\epsilon_p = \sqrt{2gn\frac{p^2}{2m} + \left(\frac{p^2}{2m}\right)^2}. \tag{2.86}$$

This represents the celebrated Bogoliubov dispersion relation for elementary excitations in a homogeneous, weakly interacting BEC.

Physically, $\hat{b}^\dagger_{\mathbf{p}}$ and $\hat{b}_{\mathbf{p}}$ create and annihilate quasi-particles atop the condensate background. The spectrum ϵ_p dictates the characteristic energy scale of the condensate dynam-

[11] There isn't a consensus on which one of the two transformations is designated as the "inverse" in the literature.

ics and is of great experimental relevance. At high momentum ($p^2/2m \gg gn$), $u_p \to 1$ and $v_p \to 0$, causing the quasi-particle operator to reduce to a bare particle operator $\hat{b}_{\mathbf{p}} \to \hat{a}_{\mathbf{p}}$. The excitation energy asymptotically approaches that of a free particle $p^2/2m + gn$, where the chemical potential $\mu = gn$ represents the cost of adding a particle to the condensate. These high-momentum quasi-particles are therefore identified as free-particle-like excitations.

On the other hand, at low momentum ($p^2/2m \ll gn$), both u_p and v_p diverge while $u_p/v_p \to 1$. Now, applying $\hat{b}_p^\dagger$ simultaneously demotes a particle with momentum $-\mathbf{p}$ and promotes a particle with $\mathbf{p}$ from the condensate background; the resulting excitation carries a total momentum of $\mathbf{p}$ without changing the net particle number overall. In real space, this amounts to creating a density wave on the condensate ground state—a *phonon* traveling wave. The dispersion relation becomes linear in momentum: $\epsilon_p \approx c_s p$, where the speed of sound c_s reads

$$c_s = \sqrt{gn/m}. \tag{2.87}$$

Sound propagation was first observed in a harmonically trapped condensate [54], and phonon excitations in homogeneous ultracold Bose gases have since been directly probed in experiments [55, 56]; they will serve as an important calibration tool in Section 4.4.

It is instructive to compare the ground state of an interacting Bose gas with its non-interacting counterpart. Particularly, in a uniform system, both cases feature massive occupation of the zero-momentum state at zeroth order. Interactions, however, mingle opposite momentum modes, scattering pairs of particles with momenta $\mathbf{p}$ and $-\mathbf{p}$ out of the condensate mode. This is evident in the structure of the quasi-particle vacuum $|0\rangle_{\mathrm{QP}}$, defined by

$$\hat{b}_{\mathbf{p}}|0\rangle_{\mathrm{QP}} = 0. \tag{2.88}$$

Expressing $\hat{b}_{\mathbf{p}}$ in terms of $\hat{a}_{\mathbf{p}}$ and solving the resulting linear system [39, Section 13.8] yields the quasi-particle vacuum in the Fock basis:

$$|0\rangle_{\mathrm{QP}} = \prod_{\mathbf{p}\in\text{half space}}^{\otimes} \frac{1}{\sqrt{|u_p|^2}} \sum_j \left(-\frac{v_p^*}{u_p^*}\right)^j \left|N_{\mathbf{p}} = j, N_{-\mathbf{p}} = j\right\rangle \tag{2.89}$$

where "$\otimes$" denotes the direct product between the Hilbert spaces of different momentum modes.

The occupation of each non-zero momentum state is readily evaluated as

$$
\begin{aligned}
{}_{\mathrm{QP}}\langle 0|\hat{a}_{\mathbf{p}}^{\dagger}\hat{a}_{\mathbf{p}}|0\rangle_{\mathrm{QP}} &= {}_{\mathrm{QP}}\langle 0|(u_p^*\hat{b}_{\mathbf{p}}^{\dagger} - v_p\hat{b}_{-\mathbf{p}})(u_p\hat{b}_{\mathbf{p}} - v_p^*\hat{b}_{-\mathbf{p}}^{\dagger})|0\rangle_{\mathrm{QP}} && (2.90)\\
&= v_p^2. && (2.91)
\end{aligned}
$$

Summing over all $\mathbf{p} \neq 0$ (by evaluating an interesting integral) yields the total out-of-condensate atom number at $T = 0$

$$
N_{\mathrm{ex}} = \sum_{\mathbf{p}\neq 0} v_p^2 = V\int \frac{\mathrm{d}^3\mathbf{p}}{(2\pi\hbar)^3} v_p^2 = \frac{V}{3\pi^2}\left(\frac{mgn}{\hbar^2}\right)^{3/2}. \tag{2.92}
$$

This interaction-driven phenomenon—known as *quantum depletion*—persists at zero temperature and has been experimentally verified in homogeneous systems [57].

The E_0 in Eq. (2.84) contains terms from normal-ordering $\hat{b}_{\mathbf{p}}, \hat{b}_{\mathbf{p}}^{\dagger}$ during the Bogoliubov transformation, incorporating the zero-point energy summed over all $\hat{b}_{\mathbf{p}}$ modes. The summation exhibits an ultraviolet divergence—an artifact of replacing the true finite-range interatomic potential with a contact interaction, Eq. (2.38). Approximating the interaction to higher order in the scattering length yields a finite result—the famous Lee-Huang-Yang correction [21, 39, 58], which lies beyond the scope of our present discussion.

The inhomogeneous case

In the presence of an external trapping potential $U(\mathbf{r})$—the case in virtually all experiments—broken translational symmetry implies that $\mathbf{p}$ is no longer a good quantum number. Unless the trap possesses a specific symmetry (such as a periodic lattice potential or an axisymmetric trap), the non-interacting single-particle states and the quasi-particle excitations do not generally share a common set of quantum numbers. Consequently, each quasi-particle operator $\hat{b}$ becomes a linear combination of single-particle operators $\hat{a}$. In practice, it is both conceptually and numerically more straightforward to formulate the Bogoliubov transformation directly in real space.

Offsetting the Hamiltonian (2.37) by $-\mu N$[12] and expanding to second order in the field perturbation about the condensate wavefunction ϕ_0 yields

$$\hat{H}^{(2)} = \int \mathrm{d}^3\mathbf{r}\left[\delta\hat{\Psi}^\dagger(\mathbf{r})\left(-\frac{\hbar^2}{2m}\nabla^2 - \mu + U(\mathbf{r}) + 2gN|\phi_0(\mathbf{r})|^2\right)\delta\hat{\Psi}(\mathbf{r})\right. \tag{2.93}$$

$$\left. + \frac{1}{2}gN(\phi_0^*(\mathbf{r})^2\delta\hat{\Psi}(\mathbf{r})^2 + \text{h.c.})\right]. \tag{2.94}$$

We seek Bogoliubov modes $u_i(\mathbf{r})$ and $v_i(\mathbf{r})$ such that

$$\delta\hat{\Psi}(\mathbf{r}) = \sum_i \left[u_i(\mathbf{r})\hat{b}_i - v_i^*(\mathbf{r})\hat{b}_i^\dagger\right], \tag{2.95}$$

$$\delta\hat{\Psi}^\dagger(\mathbf{r}) = \sum_i \left[u_i^*(\mathbf{r})\hat{b}_i^\dagger - v_i(\mathbf{r})\hat{b}_i\right] \tag{2.96}$$

$$\tag{2.97}$$

diagonalizes the Hamiltonian (2.94).

The two commutator relations of $\hat{\Psi}(\mathbf{r})$ impose:

$$[\delta\hat{\Psi}(\mathbf{r}), \delta\hat{\Psi}^\dagger(\mathbf{r}')] = \delta(\mathbf{r}-\mathbf{r}') \quad \Longrightarrow \quad \sum_i (u_i(\mathbf{r})u_i^*(\mathbf{r}') - v_i^*(\mathbf{r})v_i(\mathbf{r}')) = \delta(\mathbf{r}-\mathbf{r}'), \tag{2.98}$$

$$[\delta\hat{\Psi}(\mathbf{r}), \delta\hat{\Psi}(\mathbf{r}')] = 0 \quad \Longrightarrow \quad \sum_i (u_i(\mathbf{r})v_i^*(\mathbf{r}') - u_i(\mathbf{r}')v_i^*(\mathbf{r})) = 0. \tag{2.99}$$

Conversely, requiring the quasi-particle operators to satisfy the commutator relations $[\hat{b}_i, \hat{b}_j^\dagger] = \delta_{ij}$ and $[\hat{b}_i, \hat{b}_j] = 0$ enforces:

$$\int \mathrm{d}^3\mathbf{r}\ [u_i^*(\mathbf{r})u_j(\mathbf{r}) - v_i^*(\mathbf{r})v_j(\mathbf{r})] = \delta_{ij}, \tag{2.100}$$

$$\int \mathrm{d}^3\mathbf{r}\ [u_i^*(\mathbf{r})v_j^*(\mathbf{r}) - v_i^*(\mathbf{r})u_j^*(\mathbf{r})] = 0. \tag{2.101}$$

The inverse Bogoliubov transformation is given by

$$\hat{b}_i = \int \mathrm{d}^3\mathbf{r}[u_i^*(\mathbf{r})\delta\hat{\Psi}(\mathbf{r}) + v_i^*(\mathbf{r})\delta\hat{\Psi}^\dagger(\mathbf{r})] \tag{2.102}$$

$$\hat{b}_i^\dagger = \int \mathrm{d}^3\mathbf{r}[u_i(\mathbf{r})\delta\hat{\Psi}^\dagger(\mathbf{r}) + v_i(\mathbf{r})\delta\hat{\Psi}(\mathbf{r})], \tag{2.103}$$

[12] Working in the frame rotating at frequency $\mu/\hbar$ (or equivalently shifting the Hamiltonian by $-\mu N$) ensures that the condensate ground state wavefunction is stationary, preventing the time-dependent phase rotation in the quasi-particle basis, and also streamlining numerical calculations.

as can be easily verified using the above commutator relations.

The eigenvalue problem can be cast into matrix form. Defining the Bogoliubov matrix

$$M(\mathbf{r}) = \begin{pmatrix} -\frac{\hbar^2}{2m}\nabla^2 - \mu + U(\mathbf{r}) + 2gN|\phi_0(\mathbf{r})|^2 & -gN\phi_0(\mathbf{r})^2 \\ -gN\phi_0^*(\mathbf{r})^2 & -\frac{\hbar^2}{2m}\nabla^2 - \mu + U(\mathbf{r}) + 2gN|\phi_0(\mathbf{r})|^2 \end{pmatrix}, \tag{2.104}$$

$u_i(\mathbf{r})$ and $v_i(\mathbf{r})$ solve the eigenvalue problem

$$\sigma_z M(\mathbf{r}) \begin{pmatrix} u_i(\mathbf{r}) \\ v_i(\mathbf{r}) \end{pmatrix} = \epsilon_i \begin{pmatrix} u_i(\mathbf{r}) \\ v_i(\mathbf{r}) \end{pmatrix}, \tag{2.105}$$

where the Pauli matrix $\sigma_z = \mathrm{diag}(1, -1)$. This constitutes the *Bogoliubov–de Gennes (BdG) equation*, and $\sigma_z M(\mathbf{r})$ is the non-Hermitian BdG operator[13].

Notably, $(u, v) = (\phi_0, \phi_0^*)$ is always an exact solution to Eq. (2.105) with eigenvalue $\epsilon = 0$, known as the *zero mode*. Physically, this zero mode represents an overall rotation of the condensate phase—the Nambu–Goldstone mode associated with the spontaneously broken $U(1)$ symmetry. Because its norm $\int \mathrm{d}^3\mathbf{r}\ (|u|^2 - |v|^2) = 0$, the zero mode cannot be normalized to unity; it falls outside the $\hat{b}_i$ basis and must be treated separately.

Furthermore, substituting Eq. (2.96) into the Hamiltonian (2.94) and invoking the orthogonality relations (2.100) and (2.101) diagonalizes the quadratic Hamiltonian:

$$\hat{H}^{(2)} = \sum_{i,j} \hat{b}_i^\dagger \hat{b}_j \int \mathrm{d}^3\mathbf{r} \Big(u_i^*(\mathbf{r}), \quad v_i^*(\mathbf{r}) \Big) M(\mathbf{r}) \begin{pmatrix} u_j(\mathbf{r}) \\ v_j(\mathbf{r}) \end{pmatrix} \tag{2.106}$$

$$= \sum_i \epsilon_i \hat{b}_i^\dagger \hat{b}_i. \tag{2.107}$$

All the relations in this section provide a complete mathematical foundation for numerically computing the excitation spectra and quasi-particle modes in arbitrary inhomogeneous geometries.

The two-component case

We are now in a position to generalize the real-space BdG framework to binary condensates. Similar to the single-component case, expressing the field operators as $\hat{\Psi}_{\uparrow,\downarrow}(\mathbf{r}) =$

[13] Since $\sigma_z M$ is non-Hermitian, it is not immediately clear that its eigenvalues are real. In fact, it can be shown that provided the ground state is energetically stable (M is semi-positive definite), all eigenvalues of $\sigma_z M$ are real, and the corresponding eigenvectors are normalizable by Eq. (2.100) (with the exception of the zero modes). A detailed proof can be found in Ref. [59, Chapter 3].

$\sqrt{N_{\uparrow,\downarrow}}\phi_{\uparrow 0,\downarrow 0}(\mathbf{r}) + \delta\hat{\Psi}_{\uparrow,\downarrow}(\mathbf{r})$ and expanding the two-component Hamiltonian (2.66) to second order in the field perturbations $\delta\hat{\Psi}_{\uparrow,\downarrow}(\mathbf{r})$ yields:

$$\hat{H}^{(2)} = \int \mathrm{d}^3\mathbf{r}\Bigg[\delta\hat{\Psi}_\uparrow^\dagger(\mathbf{r})\left(-\frac{\hbar^2}{2m}\nabla^2 - \mu_\uparrow + U_\uparrow(\mathbf{r}) + 2g_{\uparrow\uparrow}N_\uparrow|\phi_{\uparrow 0}(\mathbf{r})|^2 + g_{\uparrow\downarrow}N_\downarrow|\phi_{\downarrow 0}(\mathbf{r})|^2\right)\delta\hat{\Psi}_\uparrow(\mathbf{r}) \tag{2.108}$$

$$+ \delta\hat{\Psi}_\downarrow^\dagger(\mathbf{r})\left(-\frac{\hbar^2}{2m}\nabla^2 - \mu_\downarrow + U_\downarrow(\mathbf{r}) + 2g_{\downarrow\downarrow}N_\downarrow|\phi_{\downarrow 0}(\mathbf{r})|^2 + g_{\uparrow\downarrow}N_\uparrow|\phi_{\uparrow 0}(\mathbf{r})|^2\right)\delta\hat{\Psi}_\downarrow(\mathbf{r}) \tag{2.109}$$

$$+ \frac{1}{2}g_{\uparrow\uparrow}N_\uparrow(\phi_{\uparrow 0}^*(\mathbf{r})^2\delta\hat{\Psi}_\uparrow(\mathbf{r})^2 + \text{h.c.}) + \frac{1}{2}g_{\downarrow\downarrow}N_\downarrow(\phi_{\downarrow 0}^*(\mathbf{r})^2\delta\hat{\Psi}_\downarrow(\mathbf{r})^2 + \text{h.c.}) \tag{2.110}$$

$$+ g_{\uparrow\downarrow}\sqrt{N_\uparrow N_\downarrow}(\phi_{\uparrow 0}^*(\mathbf{r})\phi_{\downarrow 0}^*(\mathbf{r})\delta\hat{\Psi}_\uparrow(\mathbf{r})\delta\hat{\Psi}_\downarrow(\mathbf{r}) + \text{h.c.}) \tag{2.111}$$

$$+ g_{\uparrow\downarrow}\sqrt{N_\uparrow N_\downarrow}(\phi_{\uparrow 0}(\mathbf{r})\phi_{\downarrow 0}^*(\mathbf{r})\delta\hat{\Psi}_\uparrow^\dagger(\mathbf{r})\delta\hat{\Psi}_\downarrow(\mathbf{r}) + \text{h.c.})\Bigg]. \tag{2.112}$$

The Bogoliubov transform reads

$$\delta\hat{\Psi}_\uparrow(\mathbf{r}) = \sum_i\left[u_{\uparrow,i}(\mathbf{r})\hat{b}_i - v_{\uparrow,i}^*(\mathbf{r})\hat{b}_i^\dagger\right], \tag{2.113}$$

$$\delta\hat{\Psi}_\uparrow^\dagger(\mathbf{r}) = \sum_i\left[u_{\uparrow,i}^*(\mathbf{r})\hat{b}_i^\dagger - v_{\uparrow,i}(\mathbf{r})\hat{b}_i\right], \tag{2.114}$$

$$\delta\hat{\Psi}_\downarrow(\mathbf{r}) = \sum_i\left[u_{\downarrow,i}(\mathbf{r})\hat{b}_i - v_{\downarrow,i}^*(\mathbf{r})\hat{b}_i^\dagger\right], \tag{2.115}$$

$$\delta\hat{\Psi}_\downarrow^\dagger(\mathbf{r}) = \sum_i\left[u_{\downarrow,i}^*(\mathbf{r})\hat{b}_i^\dagger - v_{\downarrow,i}(\mathbf{r})\hat{b}_i\right]. \tag{2.116}$$

I now develop the excitations on the real ground state of the system $\phi_{\uparrow 0,\downarrow 0} = \phi_{\uparrow 0,\downarrow 0}^*$. Defining the shorthand symbols:

$$\blacksquare_\uparrow = -\frac{\hbar^2}{2m}\nabla^2 - \mu_\uparrow + U_\uparrow(\mathbf{r}) + 2g_{\uparrow\uparrow}N_\uparrow|\phi_{\uparrow 0}(\mathbf{r})|^2 + g_{\uparrow\downarrow}N_\downarrow|\phi_{\downarrow 0}(\mathbf{r})|^2, \tag{2.117}$$

$$\blacksquare_\downarrow = -\frac{\hbar^2}{2m}\nabla^2 - \mu_\downarrow + U_\downarrow(\mathbf{r}) + 2g_{\downarrow\downarrow}N_\downarrow|\phi_{\downarrow 0}(\mathbf{r})|^2 + g_{\uparrow\downarrow}N_\uparrow|\phi_{\uparrow 0}(\mathbf{r})|^2, \tag{2.118}$$

$$\diamond_\sigma = g_{\sigma\sigma}N_\sigma\phi_{\sigma,0}^2, \tag{2.119}$$

$$\blacktriangle = g_{\uparrow\downarrow}\sqrt{N_\uparrow N_\downarrow}\phi_{\uparrow,0}\phi_{\downarrow,0}, \tag{2.120}$$

then the Bogoliubov–de Gennes equations for the binary condensate can be written as

$$\sigma_z M(\mathbf{r}) \begin{pmatrix} u_{\uparrow,i}(\mathbf{r}) \\ v_{\uparrow,i}(\mathbf{r}) \\ u_{\downarrow,i}(\mathbf{r}) \\ v_{\downarrow,i}(\mathbf{r}) \end{pmatrix} = \epsilon_i \begin{pmatrix} u_{\uparrow,i}(\mathbf{r}) \\ v_{\uparrow,i}(\mathbf{r}) \\ u_{\downarrow,i}(\mathbf{r}) \\ v_{\downarrow,i}(\mathbf{r}) \end{pmatrix} \tag{2.121}$$

where $\sigma_z = \mathrm{diag}(1, -1, 1, -1)$, and the Hermitian Bogoliubov matrix is:

$$M(\mathbf{r}) = \begin{pmatrix} \blacksquare_\uparrow & -\diamond_\uparrow & \blacktriangle & -\blacktriangle \\ -\diamond_\uparrow & \blacksquare_\uparrow & -\blacktriangle & \blacktriangle \\ \blacktriangle & -\blacktriangle & \blacksquare_\downarrow & -\diamond_\downarrow \\ -\blacktriangle & \blacktriangle & -\diamond_\downarrow & \blacksquare_\downarrow \end{pmatrix}. \tag{2.122}$$

At first glance, it might seem surprising that only a single set of Bogoliubov quasiparticle operators $\hat{b}_i$ diagonalizes the two-component Hamiltonian (2.112) containing two distinct field operators $\hat{\Psi}_\uparrow(\mathbf{r})$ and $\hat{\Psi}_\downarrow(\mathbf{r})$, whereas in the limit $g_{\uparrow\downarrow} \to 0$ one expects two sets of clearly independent phonon modes residing in the uncoupled two components. In fact, any $g_{\uparrow\downarrow} \neq 0$ hybridizes these excitations into coupled density and spin modes in miscible systems, or into the bulk and interface modes in immiscible systems, all under the same index i. In numerical practice, isolating modes associated with different degrees of freedom from a single eigenvalue spectrum is always a challenge.

Furthermore, unlike in the single-component case, the binary BdG system Eq. (2.121) possesses multiple zero modes. Two of them correspond to independent phase rotations of each component, $(\phi_{\uparrow 0}, \phi_{\uparrow 0}, 0, 0)^T$ and $(0, 0, \phi_{\downarrow 0}, \phi_{\downarrow 0})^T$, one for each broken $U(1)$ symmetry. In a phase-separated system, spontaneous breaking of translational symmetry by the interface along the x-direction introduces a third zero mode: the rigid displacement of the interface along the x-direction $(\partial_x \phi_{\uparrow 0}, -\partial_x \phi_{\uparrow 0}, \partial_x \phi_{\downarrow 0}, -\partial_x \phi_{\downarrow 0})^T$ [60]. At finite transverse wavenumber k_y along the interface, these modes become *ripplons*—the capillary wave of the quantum fluid interface [60], which we study in detail in Chapter 5.

2.7 Hydrodynamic formalism

Another way to view the condensate dynamics, apart from GPE [Eq. (2.61)], is to use a set of equations analogous to hydrodynamic equations [36, Chapter 7]. This formalism clarifies the local velocity of a condensate and directly highlights the fluid-like behavior of interacting condensates during dynamic evolution—thereby establishing its connection to superfluidity. Furthermore, linearizing these hydrodynamic equations around equilibrium configurations offers an alternative to Eq. (2.105), which is more analytically elucidating in many cases.

We begin with the continuity equation derived from Eq. (2.61). Similar to the derivation from the single-particle Schrödinger equation, it reads

$$\frac{\partial|\phi|^2}{\partial t}+\nabla\cdot\left(\frac{\hbar}{2im}(\phi^*\nabla\phi-\phi\nabla\phi^*)\right)=0, \tag{2.123}$$

which can be written as

$$\frac{\partial n}{\partial t}+\nabla\cdot(n\mathbf{v})=0 \tag{2.124}$$

where $n=|\phi|^2$ is the particle density and $\mathbf{v}$ is the velocity associated with the particle transport:

$$\mathbf{v}=\frac{\hbar}{2im}(\phi^*\nabla\phi-\phi\nabla\phi^*). \tag{2.125}$$

We further separate ϕ into the amplitude f and phase φ: $\phi=fe^{i\varphi}$. This yields

$$n=f^2,\quad \mathbf{v}=\frac{\hbar}{m}\nabla\varphi. \tag{2.126}$$

Substituting ϕ in Eq. (2.61) with $fe^{i\varphi}$, separating the real part, and applying ∇ to both sides yields

$$m\frac{\partial\mathbf{v}}{\partial t}=-\nabla\left(\tilde{\mu}+\frac{1}{2}m\mathbf{v}^2\right) \tag{2.127}$$

which is very similar to Euler's equation. In the above equation, $\tilde{\mu}\equiv U+gn-\frac{\hbar^2}{2m\sqrt{n}}\nabla^2\sqrt{n}$, whose final term is commonly known as the "quantum pressure".

To determine the linearized excitations around the equilibrium configuration, we set $n=n_0+\delta n$. As a premise for the hydrodynamic formalism to be applicable, the interaction energy dominates over the kinetic energy, allowing us to neglect the quantum pressure

and the velocity term in Eq. (2.127). Substituting this into Eq. (2.124) and linearizing with respect to δn yields the equation of motion for elementary excitations

$$m\frac{\partial^2 \delta n}{\partial t^2} = g\nabla \cdot (n\nabla \delta n). \tag{2.128}$$

We are interested in the oscillatory solution of the form $\delta n \sim e^{i\omega t}$. Under the Thomas-Fermi limit [Eq. (2.47)], the above equation becomes

$$\omega^2 \delta n = \frac{1}{m}\left[\nabla U \cdot \nabla \delta n - (\mu - U)\nabla^2 \delta n\right], \tag{2.129}$$

which is the secular equation of condensate excitations in the hydrodynamic limit. Eq. (2.129) will be applied to analyze phonon excitations in Section 4.4.

3 Experimental Realization of a Quasi-2D Bose-Einstein Condensate

In the 1990s, creating a BEC with ultracold atoms was a Nobel Prize-winning scientific feat. Three decades later, condensate production is a daily routine in hundreds of laboratories worldwide—extending even to the International Space Station [61]. In this chapter, I present our experimental apparatus and various techniques used to prepare a BEC of ^{23}Na atoms in a quasi-2D potential.

I begin with a concise summary of the atomic properties of ^{23}Na pertinent to our experiments. Next, I briefly outline the condensate production sequence in a hybrid configuration of magnetic and optical dipole traps, alongside the diagnostic and optimization methods. For technical details on the apparatus, the reader is referred to the theses of Swarnav Banik [62] and Monica Gutiérrez Galán [63], as well as our publication [64]. Instead, the emphasis here is on the design and operation of the final traps: the "sheet" trap providing optical quasi-2D confinement in the z-direction, and the DMD trap enabling dynamic, arbitrary potential shaping during the experimental run.

Portions of this chapter have been published in Ref. [64].

3.1 Properties of sodium-23

Sodium-23 (^{23}Na)—one of the first atomic species cooled to quantum degeneracy—is the workhorse of our experiment. Solid at room temperature, sodium must first be vaporized in vacuum. Heating up metallic sodium produces a hot vapor, which we cool into a cold, dilute gas via laser cooling—a process critically shaped by the atom's energy structure. Next comes evaporative cooling: by selectively removing the most energetic atoms, collisions re-thermalize the remaining atoms to lower temperatures. The process continues until the condensate phase transition occurs, when the atomic cloud turns into a weakly interacting Bose-Einstein condensate. Both during evaporation and in the degenerate regime, low-energy scattering properties play a prominent role.

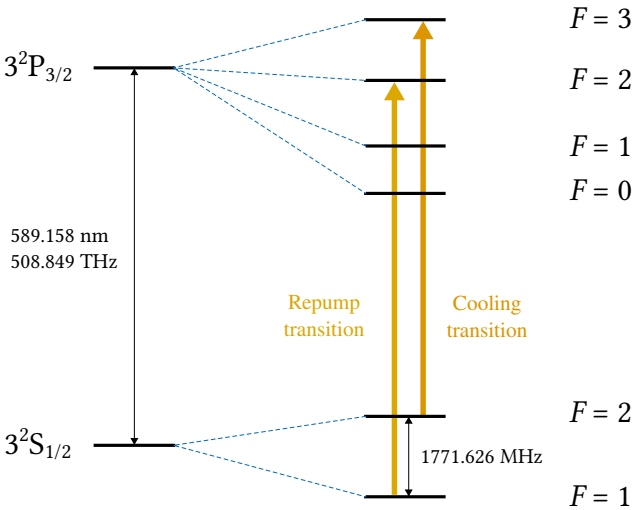

Figure 3.1: Structure of ^{23}Na D_2 line, not to scale.

3.1.1 Atomic structure and light-atom interaction

As the third element in the first column of the periodic table, ^{23}Na [65] belongs to the family of alkali atoms. Possessing only one electron in the outermost shell, ^{23}Na has a relatively simple energy structure.

In the experiments, we work with the ground state (with orbital angular momentum quantum number $L = 0$) and the first excited state ($L = 1$). The coupling between electronic spin $\mathbf{S}$ and the angular momentum $\mathbf{L}$ gives rise to the *fine structure* characterized by the total electronic angular momentum $\mathbf{J} = \mathbf{L} + \mathbf{S}$, splitting the first excited state into a doublet with $J = 1/2$ or $3/2$. The transitions from the ground state ($3^2S_{1/2}$ in compact notation) to this excited state doublet ($3^2P_{1/2}$, $3^2P_{3/2}$) are commonly referred to as the D_1 and D_2 lines, respectively. Each of the two fine-structure levels splits further under the coupling between the electron and the nuclear spin $\mathbf{I}$ ($I = 3/2$ for ^{23}Na). The resulting states, called *hyperfine states*, are labeled by the total atomic angular momentum $\mathbf{F} = \mathbf{J} + \mathbf{I}$, with the level scheme summarized in Fig. 3.1.

At the heart of ultracold atomic experiments is the interaction between atoms and light. In the literature, this interaction is described through various quantities: light absorption rate, photon scattering rate, radiation pressure, steady-state excited population, etc. These seemingly different metrics, in fact, stem from the same physics: coherent atomic excitation balanced by incoherent spontaneous emission. Central to all these observables

is the steady-state excited population ρ_{ee} (from which all other quantities differ only by numerical factors) [66, Chapter 2]:

$$\rho_{ee} = \frac{s_0/2}{1 + s_0 + (2\delta/\gamma)^2}, \tag{3.1}$$

where $s_0 \equiv I/I_{\text{sat}}$ is the on-resonance saturation parameter defined as the quotient of the light intensity I and the saturation intensity I_{sat}, $\delta = \omega_l - \omega_a$ is the laser detuning from the atomic resonance, and γ is the spontaneous decay rate (or the natural linewidth).

Notice that in the high-intensity limit, $s_0 \to \infty$, ρ_{ee} is capped at $1/2$. Under a coherent drive, atoms can only reside in the excited state for half the time. The scale of laser power for reaching this saturated regime is set by the saturation intensity

$$I_{\text{sat}} = \frac{\pi\hbar\omega\gamma}{3\lambda^2} = \frac{\hbar\omega\gamma}{2\sigma_0}. \tag{3.2}$$

Here, we define the *resonant scattering cross-section* σ_0 as

$$\sigma_0 \equiv 3\lambda^2/2\pi. \tag{3.3}$$

The saturation intensity depends intrinsically on the excited-state spontaneous decay rate γ, which is proportional to the electric dipole matrix element $|\langle g|e\hat{r}|e\rangle|^2$ [66, Chapter 2] [37, Section 5.9]. Typically, tabulated values of γ are specified for the broad fine-structure transitions. To compute the transition strength for specific hyperfine sublevels $|F, m_F\rangle \to \left|F', m_F'\right\rangle$, the bare matrix element must be weighted by appropriate Clebsch-Gordan coefficients [66, Appendix D]. The cycling transition $|F = 2, m_F = 2\rangle \to |F' = 3, m_{F'} = 3\rangle$ in the D_2 line features the largest Clebsch-Gordan coefficient ($= 1$), hence is preferred in laser cooling of alkali atomic species.

In a typical laser cooling scheme, atoms are first pumped from the "dark" $F = 1$ ground state to the $F' = 2$ excited state by the "repump" laser before spontaneously decaying into the $F = 1$ or $F = 2$ ground state. Atoms that end up in the $F = 2$ ground state enter the cyclic cooling transition driven by the "cooling" laser while the rest repeat the pumping process.

3.1.2 Low-energy scattering properties

At ultracold temperatures, interatomic interactions are governed by the s-wave scattering process in quantum scattering theory. At its heart, quantum scattering theory considers an incident plane wave deformed by a central potential into outgoing spherical waves. By

decomposing the outgoing wavefunction into partial waves, the angular dependence is isolated into Legendre polynomials $P_l(\cos\theta)$, while the radial wavefunctions encode the specific details of the scattering potential. In the limit $r \to \infty$ far outside the range of the central potential, the radial wavefunction approaches a free spherical wave with a phase shift [37, Section 6.4]. Because the interatomic van der Waals potential decays rapidly as $1/r^6$, low-energy scattering ($k \to 0$) cannot resolve its short-range details; the phase shift simplifies to $\delta_0 = -ka$. The entire scattering process reduces to a single parameter a—the *s-wave scattering length*.

We simplify the scattering problem by replacing the true atomic potential with a weak contact form that shares the s-wave scattering length with the true potential $U(\mathbf{r})$. Under the Born approximation, a is given by:

$$a = \frac{m_r}{2\pi\hbar^2} \int \mathrm{d}^3\mathbf{r}\, U_{\mathrm{eff}}(\mathbf{r}), \tag{3.4}$$

where $m_r = m/2$ is the reduced mass. The resulting fictitious contact potential is:

$$U_{\mathrm{eff}}(\mathbf{r}) = \frac{4\pi\hbar^2 a}{m}\delta(\mathbf{r}). \tag{3.5}$$

This matches the contact interaction used in the mean-field treatment of Chapter 2.

3.2 Path to quantum degeneracy

The journey towards quantum degeneracy consists of multiple cooling stages that bridge nine orders of magnitude in temperature—from vapor of ^{23}Na at ∼240 °C down to ∼100 nK inside an ultra-high vacuum (UHV) chamber—before loading the ultracold atoms into the final optical trapping potential.

In our experiments, the entire process takes place in interconnected UHV chambers situated on the primary optical table, enveloped by three layers of optics and various sets of custom coil assemblies, further supported by three (not very compactly arranged) auxiliary optical tables hosting four distinct laser systems and hundreds of optical elements for laser beam preparation and beam delivery via optical fibers, alongside an array of power supplies and current servos stably powering all the coils, plus a dedicated cooling process loop extracting heat produced by the ∼170 A current required for forced RF evaporative cooling; on top of them, a collection of software, timing devices, DACs, and cameras faithfully executes programmed experiments at sub-microsecond timing precision and record experimental data in every single experiment run. Within this environment, the graduate student's primary mandate is, in an unexaggerated way: to design and construct the aforementioned infrastructure, verify that every subsystem functions within spec, and ensure all subsystems operate well every single day, all at the same time. In ultracold atomic physics, establishing this operation baseline is merely the start, long before any meaningful scientific inquiry can commence. What should bring comfort to graduate student readers is that this apparatus is by far not the most intricate one in modern atomic physics research[1].

The layout of our vacuum system is illustrated in Fig. 3.2. It comprises two zones: the science chamber and the atomic source—a 2-dimensional magneto-optical trap (2D MOT). Isolated within the ultra-high vacuum throughout the entire experimental cycle, the atoms interact with the outside world exclusively through optical and magnetic fields. Rather than detailing every component, this section outlines the core experiment sequence leading to condensation and briefly highlights the key subsystems responsible for each stage. For comprehensive descriptions of hardware, beam paths, and laser parameters, etc., the reader is referred to the thesis of Swarnav and Monica [62, 63], as well as Ref. [64].

[1] This summary is intended not as a warning to incoming graduate student researchers, but rather to highlight the astonishing human capacity to repeatedly accomplish the mission impossible.

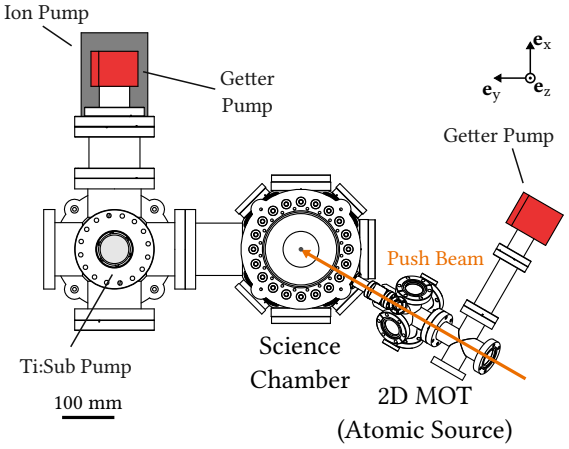

Figure 3.2: Schematic drawing of the experiment vacuum chambers and vacuum pumps.

3.2.1 Laser systems

Our laser systems divide naturally into two distinct functional classes: (i) near-resonant systems driving atomic transitions for laser cooling, optical pumping, and imaging; and (ii) far-off-resonant lasers generating optical dipole potentials.

Resonant light at the sodium D_2 transition is produced by two broadly tunable, narrow-linewidth external-cavity diode lasers:

- A Toptica RFA-SHG pro laser system at 589 nm tuned near the $|F = 2\rangle \longrightarrow |F' = 3\rangle$ transition ("cooling transition" in Fig. 3.1) for cooling and imaging. Its frequency is stabilized against a Doppler-free saturated absorption spectroscopy (SatAbs) setup. Its nominal output power is 1.1 W.

- A Toptica TA-SHG pro laser system at 589 nm tuned near the $|F = 1\rangle \longrightarrow |F' = 2\rangle$ transition ("repump transition" in Fig. 3.1) for optical pumping, stabilized against the cooling laser via an optical beat-note lock. Its nominal output power is 0.8 W[2].

For dipole trapping, far-off-resonant light is provided by:

- An IPG YLR-20-1064-LP-SF single-mode ytterbium fiber laser at 1064 nm that nominally outputs 25 W for the red-detuned ODT, delivering ~7 W of power to the experiment after intensity stabilization through a single-mode polarization-maintaining photonic crystal fiber (NKT Photonics aeroGUIDE PM-10).

- An IPG GLR-30 single-mode frequency-doubled ytterbium fiber laser at 532 nm nominally outputting 7 W for the blue-detuned quasi-2D traps.

[2] A large fraction of the repump power is dedicated to the 2D MOT repump beam.

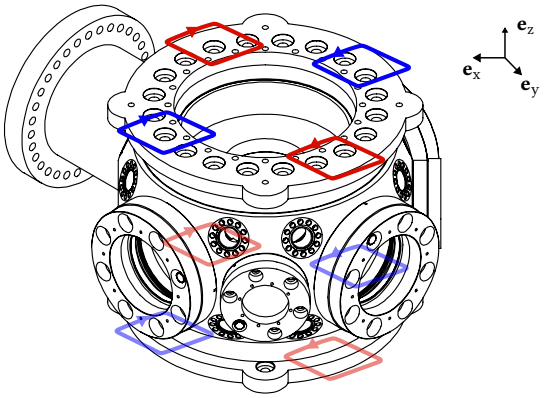

Figure 3.3: Schematic drawing of the configuration of the gradient compensation coils.

3.2.2 Magnetic coils and antennas

Laser cooling, magnetic trapping, and coherent spin manipulation require well-controlled magnetic fields.

The transverse quadrupole magnetic field for the 2D MOT is supplied by a compact set of permanent magnets (Fig. 3.4). Two orthogonal pairs of bias coils secured on the 2D MOT viewports provide fine control over the quadrupole center, optimizing the alignment of the atoms in the 2D MOT.

At the science chamber, a pair of water-cooled, hollow-core coils is mounted on the top and bottom recessed viewports to produce the quadrupole magnetic field required for the 3D MOT and magnetic trapping during RF forced evaporation. Driven by a Lambda EMI ESS 20-750 power supply, the current is regulated by a bank of six insulated-gate bipolar transistors (IGBTs) for fast control. During magnetic trapping, high-pressure distilled water (~180 psi or 1200 kPa) is forced through the coils at 0.16 GPM (0.61 L min^{-1}) to dissipate the thermal load generated by the ~170 A current during RF forced evaporation. Three additional pairs of bias coils installed on all six MOT beam viewports provide homogeneous bias magnetic fields, establishing the quantization axis and controlling the Zeeman splitting of the atoms throughout the experimental cycle.

To exert a spin-dependent magnetic force inside the final quasi-2D trap, two sets of four small coils (colored in red and blue in Fig. 3.3) produce in-plane magnetic field gradients $\partial_x B_z, \partial_y B_z$. Each coil set is arranged into two anti-parallel, top-bottom pairs. By driving anti-symmetric currents in these pairs (indicated by arrows in Fig. 3.3), the field ideally vanishes at the trap center while maintaining a field gradient across the $x-y$ plane. These coils are controlled by two low-noise current servos designed by Mingshu [67].

Two RF/microwave loop antennas are situated next to the bottom recessed window. The first antenna emits RF fields at ~10^1 MHz to drive spin-flip transitions between Zee-

man sublevels during forced evaporative cooling. The second antenna—a full-wave loop antenna matched to 50 Ω with a stub tuner—is driven by a 10 W amplifier. It radiates a 1.772 GHz microwave field tuned to the ground-state hyperfine transition $|F = 1\rangle \rightarrow |F = 2\rangle$ with a maximal Rabi frequency of 5 kHz. In the experiment, the microwave pulse is always triggered at the zero-crossing of the 60 Hz line noise to ensure the resonance stability across experimental runs.

3.2.3 Atomic source

Serving as the atomic source of our apparatus, a two-dimensional magneto-optical trap (2D MOT, Fig. 3.4) slows and transversely captures atoms exiting the resistively heated sodium oven (at ∼240 °C, well above the melting point of sodium of ∼100 °C) into an elongated cloud. During the loading phase of the science chamber, a slightly blue-detuned axial "push" beam pushes this elongated cold cloud into the ultra-high vacuum science chamber.

Our 2D MOT design closely resembles the setup developed by the Trento group [68]. Compared to traditional Zeeman slowers (affectionately dubbed "atomic fire hoses" by Gretchen Campbell), a 2D MOT offers a modern, compact alternative[3]—replacing water-cooled coils with permanent magnets reduces size and simplifies heat management. Loading into the science chamber is easily controlled by optically switching the 2D MOT laser beams, eliminating the need for a mechanical atomic shutter inside vacuum. From the oven section to the 2D MOT center, the fringe magnetic field created by the permanent magnets traces a profile similar to that of a Zeeman slower. Although transverse rather than axial (the "correct" direction for a Zeeman slower), this fringe field produces slowing effects when combined with counter-propagating slowing and repump light [68]. Nevertheless, Zeeman slowers retain advantages in their simple principle and sheer atomic flux ($\sim 10^{10}$–10^{11} atoms/s), operating robustly without fine optimization effort. Conversely, a 2D MOT demands multiple beams, greater optical power and meticulous alignment to achieve a fraction of the flux ($\sim 10^{9}$ atoms/s) of a Zeeman slower.

Our implementation differs from the original design [68] in a few notable ways. First, a 1/4-inch copper pinch-off installed directly downstream of the oven serves as a collimation tube. The resulting differential pumping drastically improved the vacuum quality, dropping the 2D MOT chamber pressure to $\sim 10^{-10}$ mbar compared to the original design of $\sim 10^{-6}$ mbar. Consequently, sodium deposition onto the MOT beam viewports is dramatically reduced. Residual ^{23}Na coated onto the viewports is sufficiently cleared overnight via light-induced atomic desorption (LIAD) using 365 nm UV LED lamps, enabling the contin-

[3] Note that modern, well-designed Zeeman slowers might also have a compact footprint.

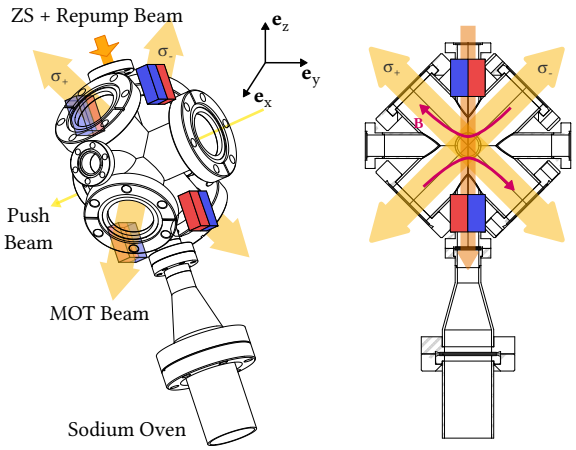

Figure 3.4: Schematic drawing of 2D MOT. Left: Front view. Right: Section view towards the push beam direction.

ued 2D MOT operation without periodic vacuum baking. However, this collimation tube restricts the solid angle of the atomic beam and partially obstructs access of the axial slowing beam. In our setup, we send the slowing and repump beams through the viewport right along the line of sight of the oven; this viewport is actively heated to 130 °C to prevent the large atomic flux from quickly coating it into a mirror. Empirically, we find the captured atomic flux is strongly limited by repump power, suggesting that the repump beam plays a crucial role in the slowing process, while reflecting the poor transmission of that viewport. Detailed operating parameters including laser detunings and optical powers for the 2D MOT are detailed in Table 1 of Ref. [64].

3.2.4 Multistage cooling sequence

After loading atoms into the science chamber, a lengthy multistage cooling process ensues. Separated from the 2D MOT source by a differential pumping tube, the science chamber maintains an ultra-high vacuum of $\sim 10^{-10}$ mbar, suppressing background collisions and ensuring long trap lifetimes (~25 s of $1/e$-lifetime in the magnetic trap) throughout the cooling sequence.

The sequence begins with capturing atoms in a standard six-beam magneto-optical trap (termed "3D MOT" in our experimental nomenclature). The cooling beams address the closed $|F = 2\rangle \longrightarrow |F' = 3\rangle$ cycling transition on the sodium D_2 line (F and F' denote the ground- and excited-state hyperfine quantum numbers). Simultaneously, a repump beam drives the $|F = 1\rangle \longrightarrow |F' = 2\rangle$ transition that recycles decayed atoms from the $F = 1$ dark state. With a ~9 mm-diameter dark center in the beam [69], this "hollow" repump

configuration reduces the photon rescattering of the already-cooled atoms, thus improving cloud density[4].

Loading for 5 s to 7 s yields 1.7×10^9 to 2.0×10^9 atoms in the MOT at a temperature of ~290 μK. Next, the quadrupole magnetic field is abruptly switched off to commence the sub-Doppler cooling in the polarization lattice formed by the six MOT beams (termed "polarization gradient cooling" or PGC [66, 70]), further reducing the cloud temperature to ~40 μK. This concludes the laser-cooling stage.

Next, a brief optical pumping pulse (designated "spin polarization", or SP in experimental terms) resonant with the $|F = 1\rangle \longrightarrow |F' = 1\rangle$ transition optically pumps 50 % to 60 % of the population into the low-field-seeking trappable state $|F = 1, m_F = -1\rangle$. Immediately following optical pumping, the quadrupole magnetic field gradient is rapidly snapped up to 190 $\mathrm{G\,cm^{-1}}$ (~ 170 A current through the water-cooled coils), compressing the cloud and initiating RF forced evaporation. Over ~8 s, the RF frequency is swept from 50 MHz to 5 MHz. This "RF knife" removes atoms with large kinetic energy while the remaining atoms re-thermalize via rapid elastic collisions in the tight magnetic trap. Meanwhile, a 6 W tightly focused IR laser beam at 1064 nm with a 22 μm beam waist is switched on, imposing an attractive optical dipole trap (ODT) of depth ~200 μK. As cooling proceeds and atoms concentrate in the low-field region, the adiabatic spin-following condition fails, leading to non-adiabatic spin flips into untrappable states ("Majorana loss") that induce atom loss and heating.

To inhibit Majorana loss, the quadrupole magnetic field is ramped down, and the atoms are captured into the ODT, positioned near the zero-field center, before settling into the low-field region. Over the subsequent ~ 4 s, the ODT power is ramped down from 6 W to 50 mW, releasing high-energy atoms to continue the evaporation. The small, residual quadrupole magnetic gradient maintains axial confinement along the ODT beam where the potential is loose, completing the "hybrid trap", with an effective potential of [71]

$$U(x, \rho, \theta) = -U_0 \exp\left(-\frac{2\rho^2}{w_0^2}\right) \tag{3.6}$$

$$+ \frac{\mu_B B'_z}{2}\sqrt{\frac{x^2 + \rho^2\cos^2\theta}{4} + (\rho\sin\theta - z_0)^2} \tag{3.7}$$

$$+ mg\rho\sin\theta, \tag{3.8}$$

[4] Conventionally termed "dark SPOT MOT" in modern literature, "SPOT" stands for "dark spontaneous-force optical trap" [69]. Technically, the "MOT" part in the name is redundant. The historical role of "dark SPOT MOT" in BEC research is reviewed in Ref. [5].

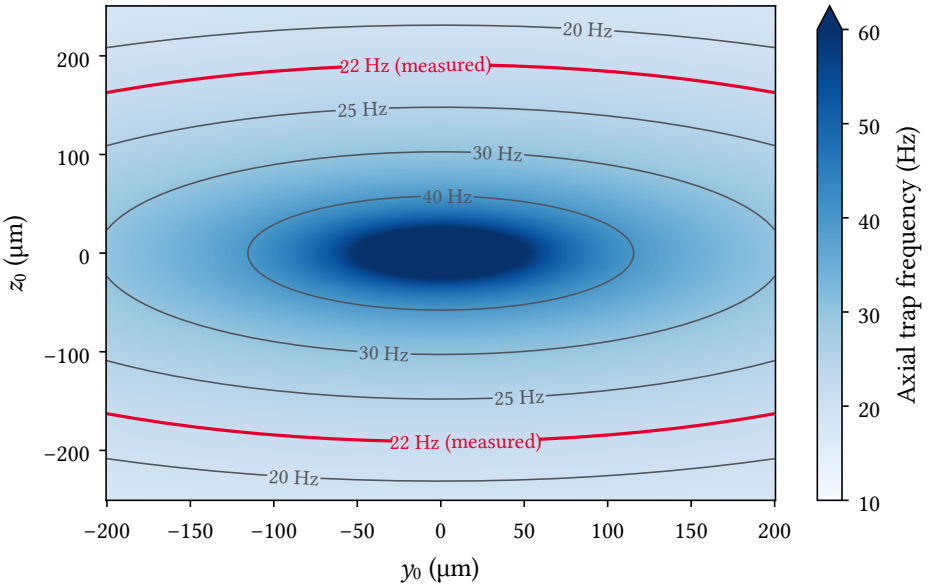

Figure 3.5: Axial trap frequency $\omega_x/2\pi$ of the hybrid trap as a function of the offset (y_0, z_0) of the magnetic field zero from the ODT beam center, from Eq. (3.9) with the residual gradient $B'_z \approx 12\ \mathrm{G\,cm^{-1}}$. Gray lines are iso-frequency contours; the red contour marks the measured axial frequency of 22 Hz, which places the field zero $r_0 \approx 190\,\mu\mathrm{m}$ from the beam.

experienced by $|F = 1, m_F = -1\rangle$ atoms. Here, (ρ, θ) are polar coordinates in the transverse yz-plane centered on the beam axis propagating in the x-direction, with the magnetic field zero a distance z_0 above the beam. $U_0 \propto P$ is the peak optical depth, B'_z is the quadrupole field gradient, and $\mu_B/2 = g_F m_F \mu_B$ is the magnetic moment of the trapped state ($g_F = -1/2$). The axial confinement is adjustable by moving the relative distance between the magnetic field center and the ODT beam waist, with the axial trap frequency approximated by

$$\omega_x = \sqrt{\frac{\mu_B B'_z}{8 m r_0}}, \tag{3.9}$$

where $r_0^2 = y_0^2/4 + z_0^2$, and y_0 is the horizontal offset of the field zero (Fig. 3.5).

Eventually, approximately 12 s after initiating evaporation, the cloud crosses the critical phase-space density threshold, yielding a BEC with $\sim 4 \times 10^6$ atoms. A detailed characterization of this cooling sequence is published in Ref. [64].

3.3 Diagnostics and optimization of BEC production

A reliable set of diagnostics is paramount both during the initial push toward condensation and throughout daily operation of a working BEC apparatus. These signals provide a quantitative, unambiguous gauge of the apparatus's status, offering a solid basis for systematic optimization. This section details these diagnostic tools and the optimization protocols we adopted in our daily routine.

3.3.1 Absorption imaging of an atomic cloud

Throughout the path to condensation, two central quantities govern the state of the cloud: the atom number and the temperature. These are extracted from the absorption of resonant light and the ballistic expansion speed of the cloud upon release from the trap. When a resonant "probe" beam passes through the atomic cloud, photon absorption casts a shadow across the beam. The remaining probe beam, carrying the dark shadow and propagating out of the vacuum chamber, is imaged onto a CCD camera, forming the *absorption image* of the atomic ensemble.

Attenuation of the probe beam propagating along the z-direction through a non-interacting cloud of atoms is governed by the Beer-Lambert law:

$$\frac{\mathrm{d}I(x,y,z)}{I(x,y,z)} = -n(x,y,z)\sigma(I)\mathrm{d}z, \tag{3.10}$$

where $n(x,y,z)$ is the atomic density, $I(x,y,z)$ is the intensity of the light along the z-direction, and $\sigma(I)$ is the intensity-dependent scattering cross-section:

$$\sigma = \frac{\sigma_0}{1 + I/I_{\mathrm{sat}} + (2\delta/\gamma)^2}. \tag{3.11}$$

For an on-resonant probe ($\delta = 0$), integrating Eq. (3.10) yields the column density n_{2D}:

$$\sigma_0 n_{\mathrm{2D}}(x,y) = -\ln\left(\frac{I(x,y)}{I_0(x,y)}\right) + \frac{I_0(x,y) - I(x,y)}{I_{\mathrm{sat}}}. \tag{3.12}$$

In the weak-intensity limit ($I_0 \ll I_{\mathrm{sat}}$), this reduces to:

$$\sigma_0 n_{\mathrm{2D}}(x,y) = -\ln\left(\frac{I(x,y)}{I_0(x,y)}\right). \tag{3.13}$$

The right-hand side—the negative logarithmic ratio of transmitted intensity I to incident intensity I_0—defines the *optical density (OD)*.

As a standard technique, absorption imaging carries three principal caveats:

1. It assumes the medium is optically thin, where photon re-absorption and re-emission is negligible. This assumption fails in dense ensembles, such as clouds in tightly confined dipole traps or condensates. In these scenarios, light scattering becomes a complex many-body problem not yet fully resolved [72, 73].

2. Extracting column density from OD requires precise knowledge of the scattering cross section σ. This is straightforward for a closed cycling transition where the system is restricted to two states only. For non-cycling transitions, optical pumping effects reduce the effective cross section σ, and sometimes even introduce time dependence dictated by the optical pumping timescale toward the steady state.

3. The linear relation in Eq. (3.13) holds only in the weak-probe limit ($I_0 \ll I_{\text{sat}}$), beyond which the saturation effect (I_0/I_{sat} term) emerges. To extract consistent density across varying probe intensities, I_{sat} must be calibrated empirically. While a weak probe can be completely absorbed by a dense cloud, increasing the probe intensity or the pulse duration risks kicking the cloud out of focus during exposure.

The ideal approach is clear: probe a cycling transition at low intensity. In practice, naive application of Eq. (3.13) can easily introduce systematic errors as large as a factor of two in atom numbers—tolerable for diagnosing a MOT, but unacceptable when characterizing condensate physics.

In our imaging scheme, we probe on the closed cooling transition $|F = 2\rangle \rightarrow |F' = 3\rangle$. Prior to the probe pulse, a repump beam optically pumps atoms from $|F = 1\rangle$ to $|F = 2\rangle$. For resolving features on a condensate, we avoid the optical pumping and employ a coherent microwave pulse to transfer a small fraction of atoms from $|F = 1\rangle$ to $|F = 2\rangle$—a technique known as the *partial transfer absorption imaging (PTAI)*, detailed in Section 4.1.1. To calibrate the imaged density, we usually first characterize the atom number in a condensate through collective dynamics, then tabulate the OD against the transferred population across varying microwave pulse durations.

3.3.2 Time-of-flight measurement of cloud expansion

A pillar of quantum gas diagnostics is the *time-of-flight* experiment. In this experiment, the confining potential is switched off to let the atomic ensemble expand ballistically. For a thermal cloud with negligible initial size, the ballistic expansion yields a spatial Gaussian density profile $\propto \exp\left(-r^2/2(\bar{v}t)^2\right)$, from which the RMS velocity $\bar{v} = \sqrt{k_B T/m}$ gives the ensemble temperature.

Historically, time-of-flight expansion provided the hallmark evidence for Bose-Einstein condensation. Driven by the mean-field interaction pressure, a condensate expands most

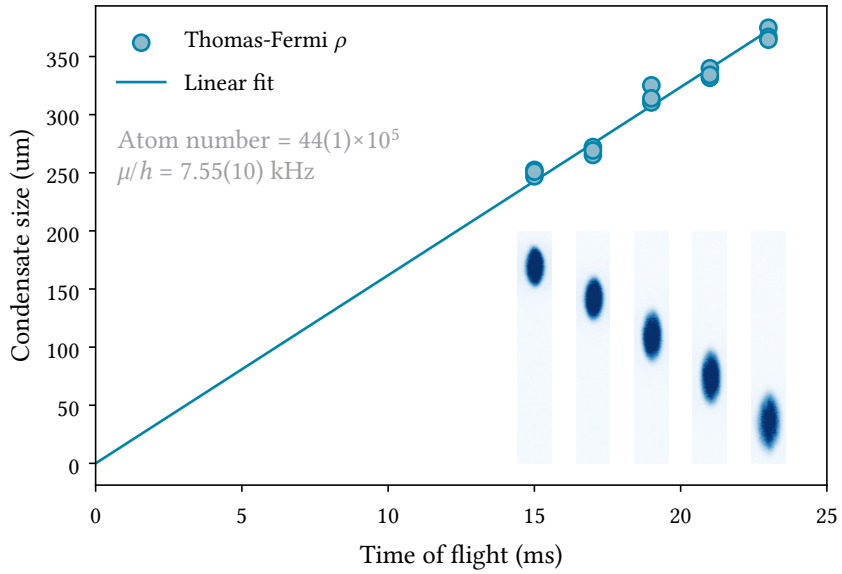

Figure 3.6: Condensate expansion during the time-of-flight. The dots are the radial size R of the condensate acquired from fitting the Thomas-Fermi condensate profile to the absorption images (bottom right) taken from the horizontal imaging system. The solid line is a linear fit of R against time-of-flight t, valid for $t \gg 1/\omega_\perp$. The annotated chemical potential and atom number are derived from Eq. (3.15).

rapidly in the direction of tightest confinement—strikingly different from the isotropic expansion of thermal atoms. Slightly below the transition temperature T_c, a sizable thermal fraction persists, resulting in a bimodal density distribution: a dense, Thomas-Fermi condensate core surrounded by a diffuse Gaussian thermal wing. A bimodal fit isolates these contributions, providing a direct measurement of the condensate fraction N_0/N.

Furthermore, for a BEC released from a harmonic trap with tight radial confinement ($\omega_\perp \gg \omega_z$), the expansion of its radial Thomas-Fermi radius $R(t)$, to leading order in $\omega_z/\omega_\perp$, obeys [74]:

$$R(t) = R_0\sqrt{1+(\omega_\perp t)^2}, \tag{3.14}$$

where the initial in-trap radius R_0 directly maps to the chemical potential μ and atom number N_0 by:

$$\mu = \frac{1}{2}m\omega_\perp^2 R_0^2 = \frac{15^{2/5}}{2}(N_0 a_{\mathrm{Na}})^{2/5}\left(\frac{m\bar{\omega}}{\hbar}\right)^{1/5}\hbar\bar{\omega}. \tag{3.15}$$

Here, $\omega_\perp$ denotes the radial trap frequency, and $\bar{\omega} = (\omega_\perp^2\omega_z)^{1/3}$ is the geometrically averaged trap frequency. This measurement, as shown in Fig. 3.6, is remarkably robust: the

Thomas-Fermi radius of the expanding condensate is adequately determined by the density profile at its edge, where the OD thins out enough to be reliably imaged, while trap frequencies and imaging magnification can be calibrated with high precision. Directly extracting absolute densities from absorption imaging is notoriously prone to systematic errors. Leveraging the collective condensate dynamics to infer densities will be a recurring theme throughout this thesis.

3.3.3 Optimization of condensate production

The primary goal of a BEC production sequence is to produce a large BEC—a critical asset for hydrodynamic studies. As outlined in Ref. [64], optimizing such a complicated multi-stage experiment sequence across dozens of parameters poses a formidable challenge. In this case, empirical pragmatism often outpaces idealized models—not because physics fails; it just works in a complicated way.

Yet, treating BEC production as a black box to be enumerated by brute-force search is both inefficient and unilluminating. For the initial investigation and optimization, physical intuition must take the lead. Our experimental sequence is partitioned into stages defined by the underlying physical processes. Because these stages are largely decoupled, we can isolate and optimize a small subset of parameters at a time.

Systematic optimization requires a quantitative metric—or, in the language of optimization theory, a *cost function*—a single numerical value that distills the quality of the experimental outcome. Prior to achieving degeneracy, the choice of cost function evolves with the primary objective of each stage. For instance, the MOT stage targets the atom number at a fixed loading time, whereas the sub-Doppler cooling stage optimizes peak OD after time-of-flight—a convenient proxy for phase-space density. Practically, the cost function must be straightforward to measure (preferably in a single experimental shot) and resilient against shot-to-shot noise. A rigorous physical metric like true phase-space density requires numerous measurements, rendering it unsuitable in practice. Conversely, a quick proxy like peak OD can be misleading if the magnetic trap geometry varies, making direct density comparison across shots meaningless.

The central limitation of the stage-by-stage strategy lies in its omission of inter-stage correlations. A prime example is the interplay between the MOT and sub-Doppler cooling. Unexpected from theoretical models, the actual magnetic field switching lags behind the experimental control, primarily due to coil inductance and eddy currents. To achieve the zero-field environment required for sub-Doppler cooling, the bias field must overshoot the decay tail—a value heavily dependent on the preceding MOT field settings (see Ref. [64], Fig. 5).

To overcome these cross-stage dependencies, once initial condensation is achieved, we shift from stage-by-stage tuning to end-to-end optimization of the entire sequence. At this point, the primary cost function is unambiguously the atom number of the BEC N_0. Alternatively, a more sensitive choice is the condensate fraction N_0/N just below the critical temperature T_c. Subtle variations in phase-space density near T_c produce dramatic changes in peak optical density—offering high sensitivity, albeit at the cost of increased shot-to-shot fluctuations. With this unequivocal metric, parameters across all stages can be co-optimized to maximize this single cost function—the one that ultimately matters.

Gaussian-process-based optimizer

This is also the place where automatic optimization algorithms shine.

An automatic optimization algorithm runs in a closed loop: it executes an experimental shot, evaluates the cost function from the raw data, proposes a new set of parameters for the next trial, and ultimately converges to the optimum. Importantly, it must tolerate the noise intrinsic to real-world experiments.

In our setup, we implement Bayesian optimization using a Gaussian process-based optimizer provided by the open-source M-LOOP package [64, 75]. Conceptually, a Gaussian process (GP) [76] is a regression machine: for an input parameter vector $\mathbf{x}$, it predicts the Gaussian distribution of the cost function $C(\mathbf{x})$ through its mean $\mu(\mathbf{x})$ and uncertainty $\sigma(\mathbf{x})$. Upon completing each cycle, our `labscripts`-based experimental control framework feeds the evaluated cost function $C(\mathbf{x}_n)$ back to the optimizer to update the GP model. The next trial parameter set $\mathbf{x}_{n+1}$ is selected by minimizing an acquisition function $\mathbf{x}_{n+1} = \mathrm{argmin}[(1-\alpha)\mu(\mathbf{x}) - \alpha\sigma(\mathbf{x})]$, where α is swept between 0 and 1 during the optimization, balancing exploitation of predicted minima with exploration of unknown regions[5].

Time reallocation strategy

To maximize condensate yield under a cycle-time budget, we developed a time-reallocation optimization strategy [64].

Throughout the BEC production process, the single most time-consuming phase is evaporative cooling. In our sequence, this is a 12-second routine comprising 7 sequential ramps that perform RF forced evaporation in the magnetic trap, transfer into the optical dipole

[5] Remarkably, oftentimes, the GP models from optimization runs seem bizarrely overfitted to the data—producing a highly unphysical cost landscape—but the optimization does bring the cost function to a better place.

trap, and additional evaporative cooling in the combined magnetic and dipole trap ("hybrid trap"). The underlying physical intuition is straightforward: not all ramps are equally important in the evaporation sequence. By reallocating time from non-critical ramps to efficiency bottlenecks, the total condensate yield can improve without the need to increase the total sequence duration.

Remarkably, this intuitive rebalancing strategy coincides with constrained gradient descent in optimization theory. Under this framework, the modified cost function is $C_{\mathrm{LM}}(\mathbf{x}) = C(\mathbf{x}) + \lambda f(\mathbf{x})$, where $f(\mathbf{x}) = 0$ is the constraint and λ is a Lagrange multiplier (LM). Extrema of $C_{\mathrm{LM}}(\mathbf{x})$ satisfy

$$\nabla_{\mathbf{x}} C_{\mathrm{LM}} = \nabla_{\mathbf{x}} C(\mathbf{x}) + \lambda \nabla_{\mathbf{x}} f(\mathbf{x}) = 0. \tag{3.16}$$

For time reallocation, the constraint takes the linear form $f(\mathbf{x}) = \sum_i t_i - T$ where t_i represents the duration of the i-th ramp and T is the fixed time budget. Evaluating the local gradient $\nabla_{\mathbf{x}} C(\mathbf{x})$ is achieved efficiently via local stochastic sampling. Starting from $\mathbf{x} = \mathbf{x}_N$, the updated parameter vector $\mathbf{x}_{N+1}$ is given by

$$\mathbf{x}_{N+1} = \mathbf{x}_N - \beta \mathbf{P}_f \nabla_{\mathbf{x}} C(\mathbf{x}_N), \tag{3.17}$$

which modifies the standard gradient descent update rule by projecting the increment onto the hyperplane defined by $f(\mathbf{x}) = 0$ using the projection operator $\mathbf{P}_f$, thereby explicitly conserving the total sequence time at every step. The learning rate β that regulates the step size is tuned empirically. This method proved highly effective in practice, yielding a 10 % to 20 % increase in the BEC atom number without extending the cycle time.

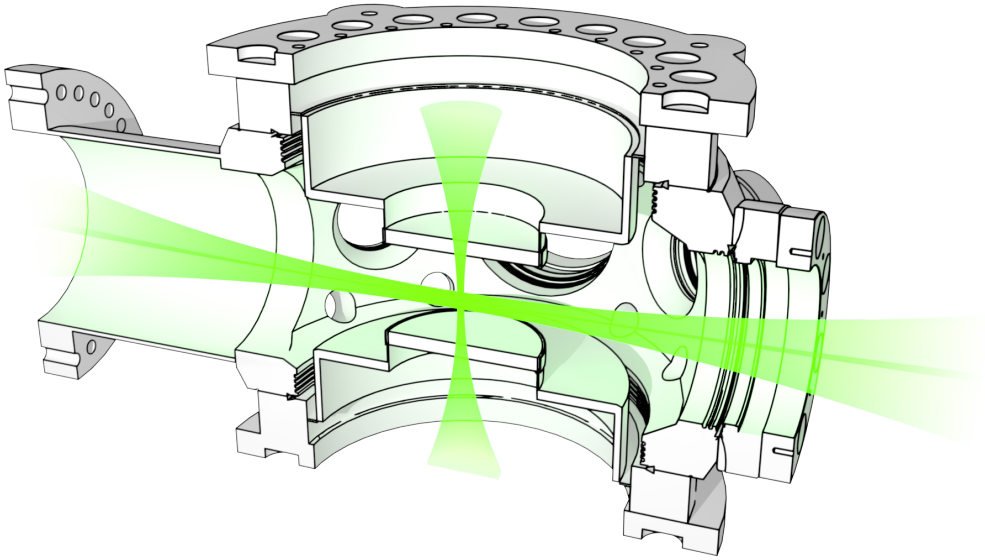

Figure 3.7: Schematic of the final optical dipole trap, consisting of the quasi-2D trap (sheet trap, in the horizontal direction) and the programmable trap created by a DMD (in the vertical direction).

3.4 Programmable quasi-2D potential

During the final stage of the BEC production sequence, the atoms are transferred into a programmable quasi-2D trapping potential made of blue-detuned laser light. This is where all scientific experiments presented in later chapters are carried out. As illustrated in Fig. 3.7, a vertically elongated 532 nm laser beam enters the science chamber from the side, focusing at the location of atoms. A sharp, dark gap is created at the center of this beam—the repulsive nature of 532 nm light (blue-detuned from the 589 nm ^{23}Na transition) pushes atoms into this dark gap, trapping atoms vertically into a "sheet", hence the name "sheet trap". The confinement in the horizontal direction is created by another 532 nm laser beam; the shape of this beam is modulated by a digital micromirror device (DMD), making the potential programmable during the experiment run. The light diffracted by the DMD is subsequently imaged onto the atoms through a microscope objective (designed by Monica [63]), which is shared with the vertical imaging subsystem.

3.4.1 Vertical confinement: sheet trap

The beam path of the sheet trap is illustrated in Fig. 3.8. A collimated 532 nm beam with a $1/e^2$-diameter of 1.3 mm delivered via a single-mode fiber, first enters a telescope of a concave cylindrical lens and a plano-convex lens that vertically elongates the collimated

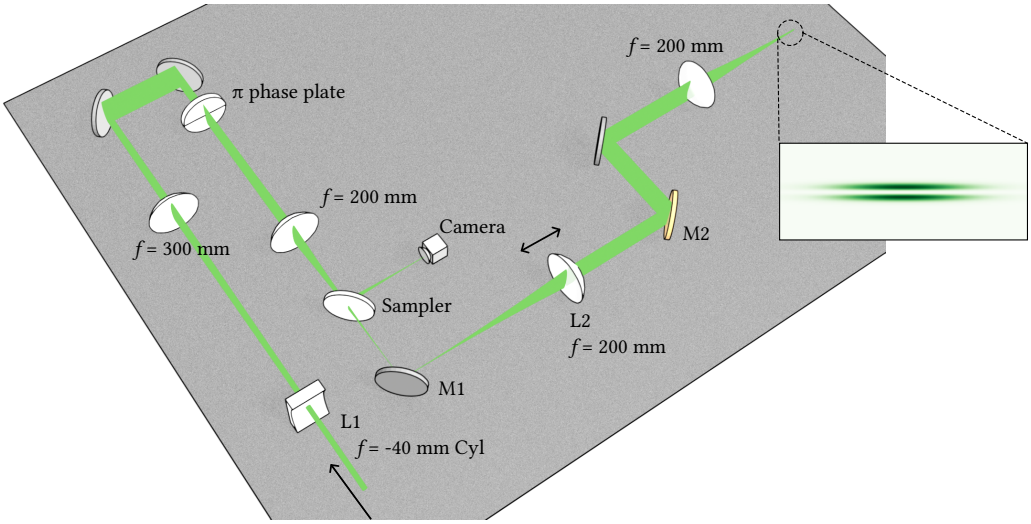

Figure 3.8: Schematic of the beam path of the sheet trap. All optics have a diameter of 2 inches.

beam while focusing it in the horizontal direction. The beam then passes the π-phase plate at the horizontal beam waist. A π-phase plate is a glass plate with a step-like thickness difference between the top and bottom halves. This thickness difference is fabricated to be exactly half a wavelength in optical path length, so that a laser beam, upon impinging on the center of the π-phase plate, acquires a π phase difference between the top and bottom halves. As a result, a vertically elongated Gaussian beam (TEM00 mode) becomes dominantly TEM01 (with 80 % spatial overlap in the electric field), with a distinct dark gap in the middle after propagation. The placement of the π-phase plate on the horizontal beam waist is intentional—it helps minimize the exposure to the manufacturing imperfection along the "cut" line. The next lens in the beam path focuses the beam in the vertical direction while collimates the beam in the horizontal direction, forming a real image of the sheet trap, which is then relayed onto atoms in the vacuum via a telescope. This image of the sheet beam is also sampled onto a camera to aid the alignment of the π-phase plate and check the quality of the sheet beam.

Five degrees of freedom are pertinent to the sheet trap alignment: the horizontal and vertical location of the focus (for intercepting atoms at its focus); the pitch and roll angles (for leveling against gravity). L1, L2, M1, and M2 labeled in Fig. 3.8 are the primary adjustable components in the system.

- L1 is installed on a rotational mount. It can be used for limited adjustment of the roll angle of the sheet before compromising the sheet structure horizontally.
- L2 can move forward and backward to focus the sheet trap onto the atoms.

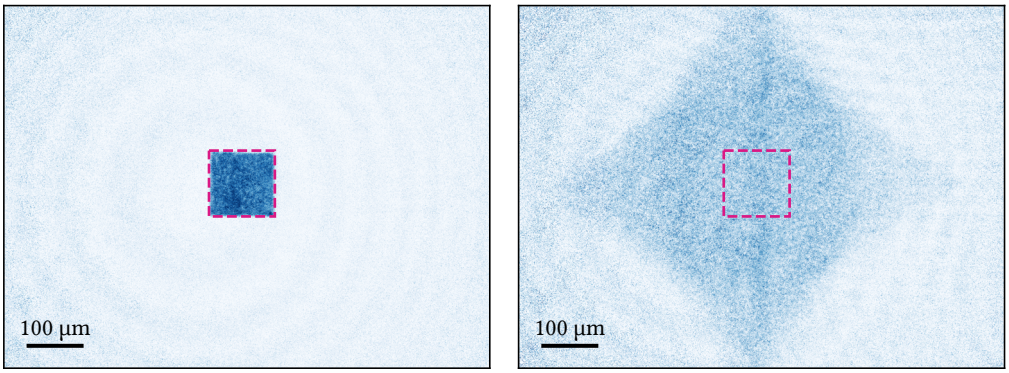

Figure 3.9: The free-expansion of a condensate in the sheet trap, initially confined in a square box potential with a size of 110 µm. Left: before turning off the box confinement. Right: 25 ms after the box confinement is turned off. The expansion is even in both $\mathbf{e}_x$ and $\mathbf{e}_y$ directions, indicating the sheet trap is free of pitch and roll misalignment.

- M1 is positioned very close to the first image plane of the sheet [6]. Adjusting its vertical pointing angle will change the pitch of the sheet beam at the atoms without moving the focus much.
- M2 is a dichroic mirror that reflects green but passes yellow (for the horizontal imaging system). It is very close to the chamber so that it is most sensitive to the location of focus but does not change the pitch as much as M1 does.

In practice, the separation of the pitch control from other degrees of freedom turned out to be crucial for quick convergence to ideal alignment. After roughly leveling and focusing the sheet trap, we adjust M2 to first intercept the atoms with the sheet trap and level it with M1. Usually, the result is evaluated by letting a fully optically trapped atom cloud expand in the sheet trap, as shown in Fig. 3.9. The focusing of the sheet trap is finalized by repeatedly moving L2 and maximizing the harmonic trap frequency, measured by quickly turning off and on the sheet trap for 3 ms to 5 ms to excite the dipole mode of the condensate (Fig. 3.10). Viewing from the side, the atoms should flow evenly in both directions after turning off the horizontal confinement (residue B-field has to be zeroed out for precise measurement); in the vertical image, the center-of-mass should remain the same as the cloud expands.

Under our typical experimental conditions, we measure a beam power at M1 to be around 300 mW, yielding a harmonic trap frequency of $f_z = 1.2$ kHz. We didn't observe a noticeable sign of anti-trapping in our system.

[6] Since it is close to the image of the sheet trap, contamination on M1 is detrimental to the quality of the trap.

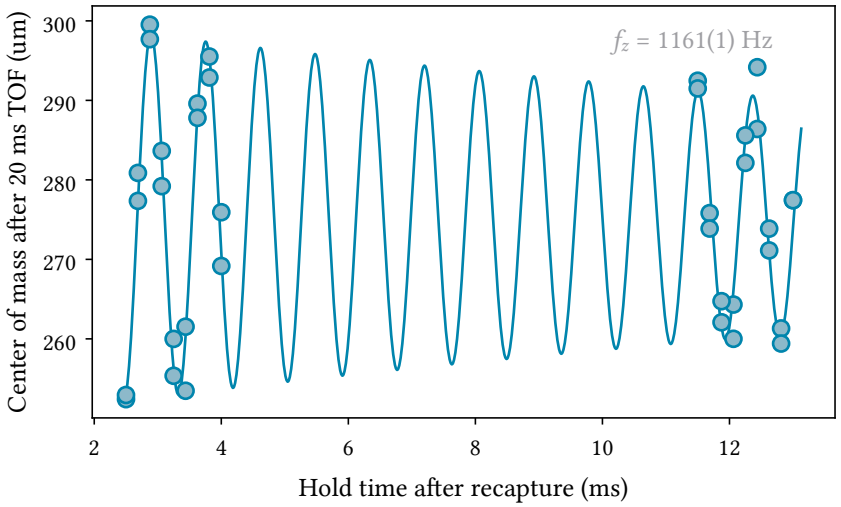

Figure 3.10: Measurement of the harmonic trap frequency of the sheet trap. The dipole mode is excited by quickly turning the sheet trap off and on, then measuring the center-of-mass (points) after a 20 ms time-of-flight against the hold time in the trap before release. The solid line is a fit of a damped sinusoid to the points.

3.4.2 Horizontal confinement: DMD trap

Apart from the sheet trap that confines atoms in the vertical direction, the other element of the final 2D traps is the confinement in the horizontal plane.

We illuminate a DMD (Texas Instrument DLP7000) with another 532 nm laser beam[7] and image its zero-th order diffraction onto the plane of the sheet trap. A DMD is an array of micro-fabricated mirrors (with a size of 13.7 µm) that can be individually controlled between the ON and OFF state. The array acts as a reflective grating; the diffracted light carries the shape of the pattern displayed on the DMD. This projection system acts as the inverse of the absorption imaging—the former projects a dark pattern into the vacuum while the latter images the dark pattern cast by the atoms onto the probe beam.

Compared to other light modulating devices like spatial light modulators (SLM), a DMD has the benefit of a fast update rate of, in our case, up to 32 kHz. But it also comes with some drawbacks: (i) The power efficiency is low (≈ 20 % in our case) since most of the power is lost to higher diffraction orders. (ii) It can only modulate light in a binary way; halftoning technique (see below) is necessary for mimicking the grayscale unless a DMD is used as a holographic device at the Fourier plane. (iii) Together with the imperfection in the imaging system, the actual projection at the plane of atoms suffers from a low-frequency wavefront

[7] The sheet beam and the DMD beam are sourced from the same laser but are taken from the $+1^{st}$ and -1st order of two AOMs, separating their frequencies by 160 MHz–necessary for avoiding spatial interference.

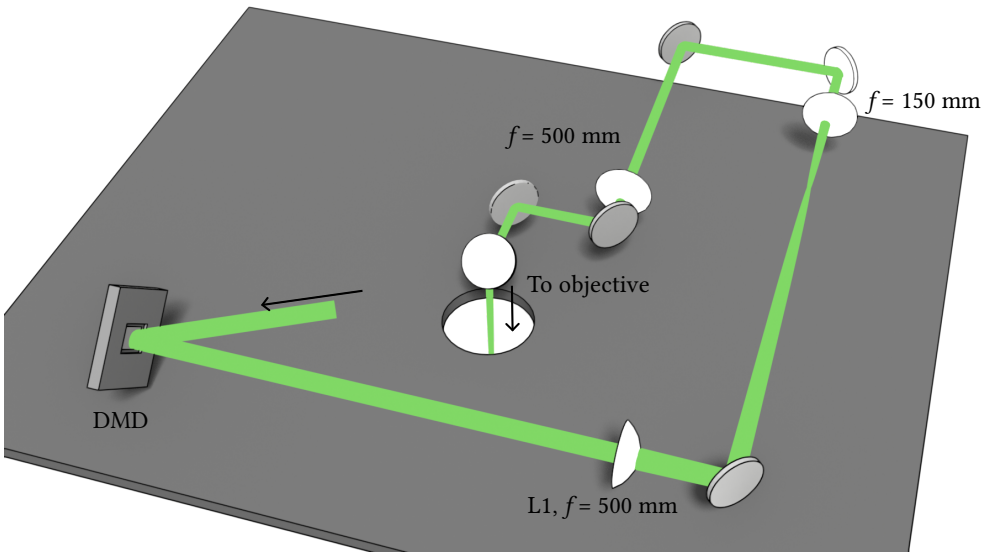

Figure 3.11: Schematic of the beam path of the DMD trap. All optics have a diameter of 2 inches. Our DMD (Texas Instrument DLP7000) has a mirror pitch of 13.68 µm.

error well-known to the community [77, 78]. This wavefront error causes aberration that limits the finest detail the projection can create; the quality of the projection beam and the reflections between optical surfaces inside the microscope objective compromise the overall homogeneity of the beam, rendering the creation of a clean potential difficult if not impossible, especially over the ~100 µm span of the typical box trap we usually use [8]. Fortunately, in our experiments, we mostly use our DMD trap to make box potentials—a bright pattern with a completely dark box region at the center. The atoms reside in the region without light, therefore do not experience the inhomogeneity and aberrations (except for the edge region).

The beam path of the DMD trap is illustrated in Fig 3.11. It is installed on a breadboard above the vacuum chamber. An input beam with $1/e^2$-diameter of 6 mm is aligned at a 24° incident angle with respect to the normal of the DMD [9] so that the main diffraction order from the DMD is perpendicular to the DMD surface, preserving the aspect ratio of the displayed pattern. The diffracted beam with the programmed pattern goes through a 4-f system that exposes the Fourier plane, before being handed over to the final telescope formed by a 150 mm lens and the microscope objective (f = 68 mm). We measured the

[8] I usually use the ugliness of the probe beam images as a reference for the look of the potentially compromised DMD potential. Now it is common to see the usage of incoherent light to eliminate those disorders [79].

[9] This angle is not the most power efficient.

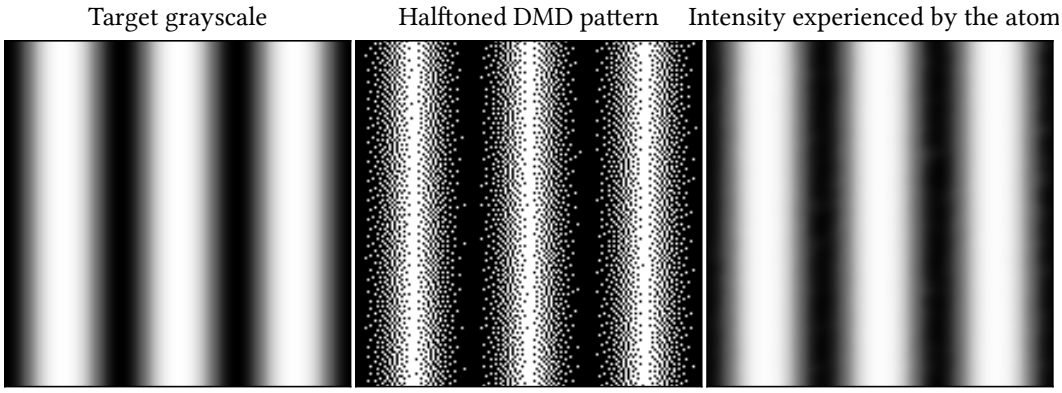

Figure 3.12: Numerical demonstration of halftoning. **Left**: the desired intensity pattern to project. **Middle**: the binary ON-OFF pattern displayed on DMD, obtained by dithering with the Floyd-Steinberg algorithm. **Right**: the intensity at the plane of atoms, after the finite aperture of the projection system has blurred away the dithering features.

power ratio of the trapping light right before entering the chamber and the total input, yielding an overall efficiency of ~20 %.

Creating grayscale from binary patterns: halftoning

Besides projecting a box-like potential with a completely dark center region, we often need to modulate the interior potential—a crucial capability necessary for exciting sound waves (detailed in later chapters) and creating vortices by phase imprinting. For this purpose, we employ the halftoning technique. Halftoning has a rich history in the printing industry. By dithering grayscale regions with tiny dots of varying densities, binary printing equipment can produce continuous-tone images for the human eye. Similarly, to project a continuously varying potential, we dither the patterns displayed on the DMD with densely interlaced ON-OFF pixel arrays. At the plane of atoms, those high-frequency dithering features are blurred by the finite resolution of the projecting system and are thus perceived by the condensate as continuously varying potential. However, when the original pattern contains large near-dark or near-bright regions, the spatial distribution of the ON/OFF pixels after dithering becomes too sparse, leaving the dithering features not sufficiently blurred at the plane of atoms.

In our experiment, we have chosen the Floyd-Steinberg algorithm[10] as our dithering method. Fig. 3.12 illustrates how halftoning works on a sinusoidal pattern for exciting a high-order sound mode.

[10] We use the Floyd-Steinberg algorithm implemented in the Pillow python package.

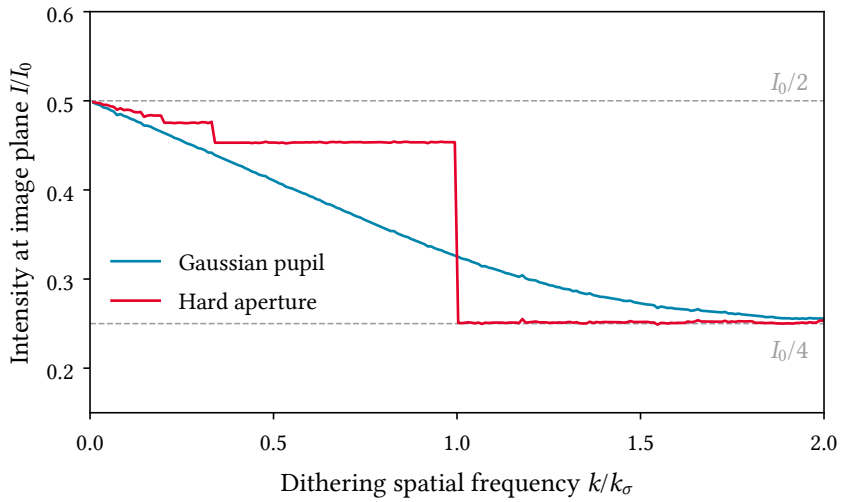

Figure 3.13: Total light intensity at the plane of atoms of a square-wave grid at 50 % duty-cycle, varying the spatial frequency. Two pupils are compared: the Gaussian (same as Fig. 3.12), and a hard circular aperture. In both cases, the intensity runs from $I_0/2$, while the grid is resolved, to $I_0/4$ once it is not.

A subtle question arises: How does the actual light intensity at the atoms scale with the gray level of the original pattern? Consider an image filled with a uniform gray tone containing an equal amount of black and white. The resulting dithered pattern will contain exactly 50 % OFF pixels, blocking half of the power; thus, one might expect the light intensity I at the atoms to be $0.5I_0$. Alternatively, one could argue that the DMD modulates the electric field, which is reduced by 50 % at the plane of atoms, yielding $I = 0.25I_0$. Both statements can not be true at the same time. The key insight is that the low-pass filtering of the imaging system comes at an energy cost: electric field components with higher spatial frequencies k are clipped by the aperture, and thus their energy is permanently lost, which is why the latter statement is correct. In fact, measuring the total power of square-wave grids with varying spatial periods at the imaging plane provides a way to characterize the resolution of the imaging system. When the imaging system fully resolves the square-wave grid, the transmitted power at the imaging plane is half of the initial power. As the grid period drops below the resolution limit, all high-frequency components are lost, leaving the transmitted power at 25 % of the initial power. This process is simulated in Fig. 3.13.

In our experiment, we safely assume the spatial frequency of the dithering features is way below the resolution limit, so that I and the gray level α follow $I \propto \alpha^2$.

4 Imaging, Control, and Characterization of a Quasi-2D Bose-Einstein Condensate

Producing a BEC and confining it to the quasi-2D potential is only the first step. Studying the dynamics of a quasi-2D BEC relies on precise control, a high-fidelity imaging system, and a well-characterized initial state defined by the trap parameters and the condensate atom number.

In this chapter, I outline the basics of imaging theory that is necessary to understand how aberrations affect image quality, followed by a detailed discussion of characterizing our imaging system. An important control knob for the binary condensate experiments in later chapters is the magnetic field gradient $\partial_x B_z, \partial_y B_z$. I describe our experimental protocols for their calibration. At the end, I present the measurement of condensate atom number in the quasi-2D trap —a number that significantly impacts the condensate dynamics.

4.1 Imaging a BEC

Usually, the only way to detect atoms in the vacuum system is through their interaction with near-resonant laser light. While Section 3.3.1 introduced the basic principle of absorption imaging—capturing the shadow cast by the condensate on the probe beam—this section focuses on how to image these shadows *clearly*. The first problem is that, as stated in Section 3.3.1, a large condensate absorbs on-resonant light too well, often leaving no light penetrating through it. The camera images a complete dark shadow—capturing no information of the features inside the condensate at all. At the same time, a condensate is tiny. The typical size of a condensate in our experiment is $\sim 100\,\mu\text{m}$—roughly the width of a human hair. To resolve the fine spatial features of condensates inside the vacuum system, we use a microscopic imaging system, centering around a homemade microscope

objective lens with a designed numerical aperture (NA) of 0.28 [1]. This number is modest among modern quantum gas experiments; nonetheless, constructing such an imaging system is far from trivial. First, the atoms are located deep within the vacuum chamber, requiring the microscope objective to have a long working distance. Second, a thick vacuum window separates the atoms and the objective, causing spherical aberrations difficult to compensate for. Beyond these two hurdles lie many more technical challenges that make imaging quantum gases a rich field of study in its own right.

The central topic of this thesis is the experimental study of interface phenomena in a binary condensate. A quantitative understanding of the interface dynamics hinges on resolving the fine structural details at the interface, and properly measuring and modeling the attenuation induced by our finite imaging resolution.

4.1.1 Partial transfer absorption imaging (PTAI) technique

We first address the problem of a condensate absorbing too much probe light. As shown in Eq. (3.13), the light intensity decreases exponentially as the column density n_{2D} increases. Increasing the probe light intensity is not always a viable option: first, the camera can easily saturate in regions of lower atomic density; second, Eq. (3.13) breaks down if the probe intensity is too strong regardless of the cloud density, requiring extra correction terms from Eq. (3.12).

We circumvent these problems by imaging only a small fraction of the condensate atoms. In our experiments, condensates form in the atomic ground state $|F = 1, m_F = -1\rangle$. To image them, we first transfer atoms from the $|F = 1\rangle$ ground state to the $|F = 2\rangle$ imaging state, then pulse a probe beam resonant with the $|F = 2\rangle \longrightarrow |F' = 3\rangle$ transition. Standard absorption imaging typically employs an optical pumping ("repump") pulse that indiscriminately transfers nearly all atoms into the imaging state. To exercise better control over this transfer process, we instead excite a microwave transition (see Section 3.2.2) between $|F = 1, m_F = -1\rangle$ and $|F = 2, m_F = -2\rangle$ states. We typically transfer 5 % to 10 % of the population when imaging the ground state condensate component, which casts a bright-enough shadow to be imaged by the imaging system.

[1] The design of the microscope objective lens is documented in Monica's thesis [63]. The NA measured from bench testing ≈ 0.21.

4.1.2 Imaging theory: a quick tour

Now, we are well-positioned to discuss how to image the shadow of atoms with high fidelity. To this end, this section presents the fundamentals of imaging theory that govern how the imaging system impacts image quality.

The cornerstone of modern imaging theory is Fourier optics [80]. Under this framework, the electric field at the object plane $\vec{E}_o$ is decomposed into plane waves at various spatial frequencies $\vec{k}$; the imaging system is then modeled by an amplitude transfer function $H(k_x, k_y)$ that describes how the amplitude and phase of a plane wave at each k are modified. Consequently, the electric field at the image plane $\vec{E}_i$ is the convolution of $\vec{E}_o$ with the transfer function in real space $h(x, y)$—the impulse response function—whose form is directly determined by the shape of the imaging system's *pupil*.

Amplitude transfer function for free space propagation

As a concrete example, we consider the free-space propagation of a plane-wave electric field with wave vector $\vec{k}$, governed by the Helmholtz equation

$$(\nabla^2 + k^2)\vec{E}(\mathbf{r}) = 0. \tag{4.1}$$

Here, k is linked to the angular frequency $\omega = 2\pi f$, the refractive index n, and the speed of light c by $k = n\omega/c$. It is straightforward to determine how a plane wave propagates in the z-direction from Eq. (4.1): one calculates the longitudinal wave vector k_z and simply multiplies the input field by the phase factor $e^{ik_z z}$. A more general electric field emanating from an arbitrary object can be decomposed into plane waves of spatial frequencies $\vec{k}$ by Fourier transformation $\tilde{E}(k_x, k_y) = \mathcal{F}[E(x, y)]$, where E denotes one of the scalar components (E_x or E_y) of the electric field.

Let the object be at $z = 0$. For each of these decomposed plane waves, we write down

$$\tilde{E}(k_x, k_y; z) = \int_{-\infty}^{+\infty} \mathrm{d}x\mathrm{d}y\, E(x, y, z) e^{-i(k_x x + k_y y)} \tag{4.2}$$

$$= \tilde{E}(k_x, k_y; z = 0) e^{ik_z z}. \tag{4.3}$$

The amplitude transfer function in this case is simply

$$H(k_x, k_y; z) = \exp(ik_z z) \tag{4.4}$$

where $k_z = (k^2 - k_x^2 - k_y^2)^{1/2}$. To obtain the real-space electric field at z, we perform an inverse Fourier transformation and find

$$E(x,y,z) = \frac{1}{4\pi^2}\int_{-\infty}^{+\infty} \mathrm{d}k_x \mathrm{d}k_y\, H(k_x,k_y;z)\tilde{E}(k_x,k_y;z=0)e^{i(k_x x+k_y y)} \tag{4.5}$$

$$= \frac{1}{4\pi^2}\int_{-\infty}^{+\infty} \mathrm{d}x'\mathrm{d}y'\, E(x',y',z=0)\int_{-\infty}^{+\infty} \mathrm{d}k_x\mathrm{d}k_y H(k_x,k_y;z)e^{i[k_x(x-x')+k_y(y-y')]} \tag{4.6}$$

which can be viewed as a convolution between the initial object field $E(x,y,z=0)$ and the "propagator" $h(x,y;z) = (1/4\pi^2)\int_{-\infty}^{+\infty}\mathrm{d}k_x\mathrm{d}k_y\ \exp\big[i(k_x x + k_y y)\big]H(k_x,k_y;z)$

$$E(x,y,z) = E(x,y,z=0) * h(x,y;z). \tag{4.7}$$

It should be clear from this expression that $h(x,y)$ acts as the response of the $E(x,y,z)$ to a δ function-like $E(x,y,z=0)$ (the "impulse response").

In most practical cases, we consider only the electric field near the optical axis, or the "paraxial" case, in which the k_z component dominates the magnitude of $\vec{k}$. By a Taylor expansion, we write down the paraxial approximation $k_z \approx k - (k_x^2 + k_y^2)/2k$, leading to the amplitude transfer function

$$H(k_x,k_y;z) = \exp\left[i\left(k - (k_x^2+k_y^2)/2k\right)z\right], \tag{4.8}$$

and its real-space counterpart—the "paraxial propagator"

$$h(x,y;z) = \frac{1}{i\lambda z}\exp\left(\frac{ik(x^2+y^2)}{2z}\right)\exp(ikz), \tag{4.9}$$

which is very useful for numerical purposes[2].

Amplitude transfer function of an imaging system

By the same token, we aim to extract the amplitude transfer function of an imaging system at the image plane—the plane of the camera sensor. As shown in Ref. [80, Chapter 6], the image of an ideal point source is simply the Fraunhofer diffraction pattern of the exit "pupil" of the imaging system[3]. Here, the term "pupil" refers not merely to a physical

[2] This formalism is exactly equivalent to solving the paraxial wave equation for the electric field $(\nabla_\perp^2 + 2ik\partial_z + 2k^2)\vec{E} = 0$.

[3] One can think of the imaging system as an equivalent ideal 4-f setup but with a "complex" aperture placed immediately before the final lens, which focuses the Fraunhofer diffraction pattern of this "aperture" onto the imaging plane.

aperture, but to the overall attenuation and phase shift imparted by the imaging system on waves at different $\vec{k}$. If the coordinates in the object plane and the image plane are (u, v) and (ξ', η'), respectively, the image of a δ function-like point at the object plane, the "impulse response", of a point source at (u, v) is

$$h(u, v; \xi', \eta') = \frac{A}{\lambda z_i} \int_{-\infty}^{+\infty} \mathrm{d}x' \mathrm{d}y' \; \mathcal{P}(x', y') \exp\left(-i \frac{2\pi}{\lambda z_i} ((u - M\xi')x' + (v - M\eta')y')\right) \tag{4.10}$$

where A is an overall amplitude scaling factor, z_i is the distance between the exit pupil and the image plane, M is the magnification, and $\mathcal{P}(x', y')$ is the complex pupil function. Defining $\xi \equiv M\xi', \eta \equiv M\eta'$ and setting $x \equiv u - \xi, y \equiv v - \eta$, we arrive at

$$h(x, y) = \frac{A}{\lambda z_i} \int_{-\infty}^{+\infty} \mathrm{d}x' \mathrm{d}y' \; \mathcal{P}(x', y') \exp\left(-i \frac{2\pi}{\lambda z_i} (xx' + yy')\right), \tag{4.11}$$

which itself mimics the form of a Fourier transform[4]. On the other hand, its Fourier transform—the amplitude transfer function $H(k_x, k_y)$—reads

$$H(k_x, k_y) = \frac{A}{\lambda z_i} \int \mathrm{d}x \mathrm{d}y \mathrm{d}x' \mathrm{d}y' \; \mathcal{P}(x', y') \exp\left(-i \frac{2\pi}{\lambda z_i} (xx' + yy')\right) e^{-i(k_x x + i k_y y)}. \tag{4.12}$$

Collecting the exponential factors into delta functions of x, y and integrating them out yields

$$H(k_x, k_y) = A\lambda z_i \mathcal{P}(-\lambda z_i k_x / 2\pi, -\lambda z_i k_y / 2\pi). \tag{4.13}$$

This is a central result of imaging theory: the amplitude transfer function is directly proportional to a scaled version of the pupil function.

[4] Fraunhofer diffraction is mathematically equivalent to the Fourier transform of the pupil function. This is, again, not a coincidence: after being clipped by the pupil the plane-wave components of the remaining electric field at different $\vec{k}$—which constitute its Fourier transform—propagate along their respective k direction; the final lens focuses every plane wave down to the focal plane. Thus, the resulting image is the Fourier transform of the pupil.

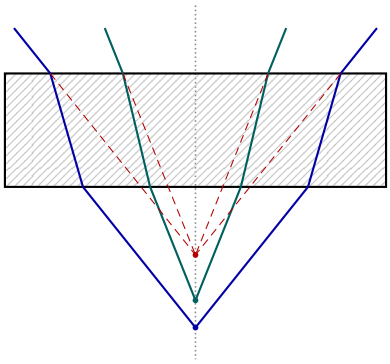

Figure 4.1: Spherical aberration introduced by a thick glass slab. Without the slab, both incoming rays at different incident angles converge at the same point (red). The slab causing them to focus at different axial positions.

4.1.3 Aberrations in the imaging system

In the presence of aberrations—the departure of a point-source wavefront from the ideal spherical form after traversing the imaging system—the pupil function becomes complex,

$$H(k_x, k_y) = A\lambda z_i \mathcal{P}(-\lambda z_i k_x/2\pi, -\lambda z_i k_y/2\pi) \tag{4.14}$$

$$= A\lambda z_i P(-\lambda z_i k_x/2\pi, -\lambda z_i k_y/2\pi) e^{ikW(-\lambda z_i k_x/2\pi, -\lambda z_i k_y/2\pi)}. \tag{4.15}$$

Here, $kW(x, y)$ represents the wavefront distortion caused by aberrations, where $W(x, y)$ is interpreted as an optical path length error across the pupil plane.

A common type of aberration is *defocus*, which occurs when the imaging plane is away from the focus by a distance d. The wavefront error can be readily found by using Eq. (4.8) to propagate the electric field away from the actual image plane by d, yielding

$$W(\rho, \theta) = \frac{d}{2z_i^2}\rho^2, \tag{4.16}$$

which is quadratic in radial pupil coordinates ρ. When x- and y-directions experience different amounts of defocusing,

$$W(x, y) = \frac{d_x}{2z_i^2}x^2 + \frac{d_y}{2z_i^2}y^2, \tag{4.17}$$

which represents *astigmatism*.

Another common type of aberration is *spherical aberration*, which occurs when a converging beam passes through a thick glass slab (for example, a vacuum window) in the beam path. The refraction through the slab not only shifts the focus, but also causes rays from different angles to focus at different axial positions, as illustrated in Fig. 4.1. Let the thickness of the slab be t, and let the refractive indices for the free space and the slab be

n_0 and n, respectively, with their ratio defined as $p = n_0/n$. The wavefront error caused by the slab is[5]

$$W(\rho,\theta) = \frac{t}{p}\left[1 - \frac{1}{2}\left(p\frac{\rho}{z_i}\right)^2 - \frac{1}{8}\left(p\frac{\rho}{z_i}\right)^4\right]. \tag{4.18}$$

The first term represents a constant phase shift that does not affect image formation; the second term reflects the overall focal shift for all rays, which can be eliminated by shifting the image plane by an amount of $\delta = t(1-p)$; the third term directly contributes to the spherical aberrations, leading to[6]

$$W_{\text{sph}}(\rho,\theta) = \frac{p(1-p^2)}{8}\left(\frac{\rho}{z_i}\right)^4 t, \tag{4.19}$$

which is the wavefront error induced by the spherical aberration of the slab. Setting $n_0 = 1$ reduces Eq. (4.19) to the well-known Seidel coefficient $W_{040} = (n^2-1)t\,\text{NA}^4/8n^3$ of a plane-parallel plate [81, Eq. (73)].

4.1.4 How optical density (OD) is modified by the imaging system

After examining how the imaging system modifies the electric field at the image plane, we now look at how the optical density (OD)—what absorption imaging actually measures—is affected.

As discussed previously, absorption imaging relies on the Beer-Lambert law (Eq. (3.10)). The physical process underlying the Beer-Lambert law is the absorption of the probe light as it propagates inside the atomic ensemble. The electric field of the probe light inside the atomic ensemble is governed by the same Eq. (4.1), modified to account for a slowly varying atomic susceptibility χ

$$\left(\nabla^2 + [1+\chi(\mathbf{r})]k_0^2\right)\vec{E}(\mathbf{r}) = 0 \tag{4.20}$$

[5] The transfer function through a thick slab of thickness t and refractive index n is $H(k_\perp) = \exp(ik'\cos\theta' t)$, where $k' = k/p$ and θ' is the refraction angle in the slab. Using Snell's law, $\cos\theta' = p\cos\theta[1+(1-p^2)/(p^2\cos^2\theta)]^{1/2}$. Simplifying and expanding around $k_\perp/k$ yields $H(k_\perp) = \exp\left[ik(t/p)(1-(pk_\perp/k)^2/2-(pk_\perp/k)^4/8)\right]$.

[6] Because spherical aberration scales as ρ^4, when shifting the image plane, the transfer function $\exp(ik_z z)$ must be expanded to ρ^4, giving a change $\Delta W = -(t-\delta)[1-(\rho/z_i)^2/2-(\rho/z_i)^4/8]$. Combining this with Eq. (4.18) yields Eq. (4.19).

where $\vec{k}_0$ is the wave vector in vacuum. The atomic susceptibility χ is given by [82, Ex. 19] [83]

$$\chi(\mathbf{r}) = \frac{\sigma_0}{k_0} \frac{i - 2\delta/\gamma}{1 + I/I_s + (2\delta/\gamma)^2} \rho(\mathbf{r}) \tag{4.21}$$

The imaginary part accounts for absorption, yielding Eq. (3.10), while the real part imprints a phase shift onto the probe beam when detuned from the atomic transition ($\delta \neq 0$).

A reasonable premise is that the object is thin along the probe beam axis, allowing us to neglect the transverse propagation of the light ($\nabla^2 = \nabla_\perp^2 + \partial_z^2 \approx \partial_z^2$). Let the object thickness be Δz, spanning from $z = -\Delta z$ to 0. The electric field of the probe light after passing through the object becomes

$$\vec{E}(x, y, 0) = \vec{E}_0(x, y, -\Delta z) \exp\left(i k_0 \int_{-\Delta z}^{0} \mathrm{d}z \, \sqrt{1 + \chi(x, y, z)} \right) \tag{4.22}$$

$$\approx \vec{E}_0(x, y, 0) \exp\left(i \frac{k_0}{2} \int_{-\Delta z}^{0} \mathrm{d}z \, \chi(x, y, z) \right). \tag{4.23}$$

We restrict the present discussion to on-resonant absorption only and let

$$-\alpha(x, y) \equiv i \frac{k_0}{2} \int_{-\Delta z}^{0} \mathrm{d}z \, \chi(x, y, z). \tag{4.24}$$

In the low intensity limit $I/I_s \ll 1$, $\alpha = \sigma_0 \rho_{2\mathrm{D}}/2$ where ρ_{2D} is the column density. The electric field immediately upon exiting the object (at $z = 0$) satisfies [7]

$$E(x, y) = \exp[-\alpha(x, y)] E_0(x, y) \tag{4.25}$$

$$= [1 - \alpha(x, y)] E_0(x, y) \tag{4.26}$$

$$= [1 - \sigma_0 \rho_{2\mathrm{D}}(x, y)/2] E_0(x, y), \tag{4.27}$$

where the expansion is valid for an optically thin ensemble (OD $\ll 1$). Similar to Eq. (4.7), after traversing the imaging system with an impulse response function $h(x, y)$, the electric field at the image plane becomes

$$E_i(x, y) = E(x, y) * h(x, y) = [1 - \sigma_0 \rho_{2\mathrm{D}}(x, y)/2] E_0(x, y) * h(x, y) \tag{4.28}$$

$$\equiv E_{i0}(x, y) + \Delta E_i(x, y), \tag{4.29}$$

[7] For simplicity, we let E_0 be real and analyze a single scalar component. Generalization is straightforward.

which yields the measured OD at the imaging plane

$$\mathrm{OD}(x,y) = -\ln \frac{|E_{i0}(x,y) + \Delta E_i(x,y)|^2}{|E_{i0}(x,y)|^2} \tag{4.30}$$

$$\approx -2\,\mathrm{Re}\,\frac{\Delta E_i(x,y)}{E_{i0}(x,y)} = \mathrm{Re}\,\frac{\sigma_0\rho_{2\mathrm{D}}(x,y)E_0(x,y) * h(x,y)}{E_0(x,y) * h(x,y)} \tag{4.31}$$

$$\approx \sigma_0\rho_{2\mathrm{D}}(x,y) * \mathrm{Re}[h(x,y)] \tag{4.32}$$

where the approximation in the last line assumes that E_0 varies slowly enough to remain constant over the spatial extent where $h(x,y)$ is non-zero. Rewriting Eq. (4.32) in Fourier space yields

$$\mathcal{F}\{\mathrm{OD}(x,y)\}(k_x,k_y) \tag{4.33}$$

$$= \mathcal{F}\{\sigma_0\rho_{2\mathrm{D}}(x,y)\}\,\mathcal{F}\{\mathrm{Re}\,h(x,y)\} \tag{4.34}$$

$$= \mathcal{F}\{\sigma_0\rho_{2\mathrm{D}}(x,y)\}\,\mathcal{F}\left\{\mathrm{Re}\,\mathcal{F}^{-1}\left\{H(k_x',k_y')\right\}\right\} \tag{4.35}$$

$$\propto \mathcal{F}\{\sigma_0\rho_{2\mathrm{D}}(x,y)\}\,\mathcal{F}\left\{\mathrm{Re}\,\mathcal{F}^{-1}\left\{\mathcal{P}(-\lambda z_i k_x'/2\pi, -\lambda z_i k_y'/2\pi)\right\}\right\} \tag{4.36}$$

which matches the form reported in Ref. [84]. The second term is known as the modulation transfer function (MTF)[8]

$$\mathcal{M}(k_x,k_y) = \mathcal{F}\left\{\mathrm{Re}\,\mathcal{F}^{-1}\left\{\mathcal{P}(-\lambda z_i k_x'/2\pi, -\lambda z_i k_y'/2\pi)\right\}\right\} \tag{4.37}$$

$$= \frac{1}{2}\left[\mathcal{P}(-\lambda z_i k_x'/2\pi, -\lambda z_i k_y'/2\pi) + \mathcal{P}^*(\lambda z_i k_x'/2\pi, \lambda z_i k_y'/2\pi)\right]. \tag{4.38}$$

Eq. (4.36) is significant—it describes how the imaging system's pupil function $\mathcal{P}$ attenuates the actual atomic absorption signal. It can be used both to characterize imaging system aberrations and to partially correct the attenuation of high-k features once the pupil function is known, which is the subject of the next section.

[8] This daunting combination of $\mathcal{F}$ and Re can be simplified using $\mathrm{Re}\,\mathcal{F}^{-1}\mathcal{P} = [\mathcal{F}^{-1}\mathcal{P}(\mathbf{r}) + \mathcal{F}^{-1}\mathcal{P}^*(-\mathbf{r})]/2$ $\implies \mathcal{M} = [\mathcal{P}(\mathbf{r}) + \mathcal{P}^*(-\mathbf{r})]/2$. When $\mathcal{P}$ has only even-order aberrations $[\mathcal{P}(\mathbf{r}) = \mathcal{P}(-\mathbf{r})]$, $\mathcal{M} = [\mathcal{P}(\mathbf{r}) + \mathcal{P}^*(\mathbf{r})]/2 = \mathrm{Re}\,\mathcal{P}(\mathbf{r})$.

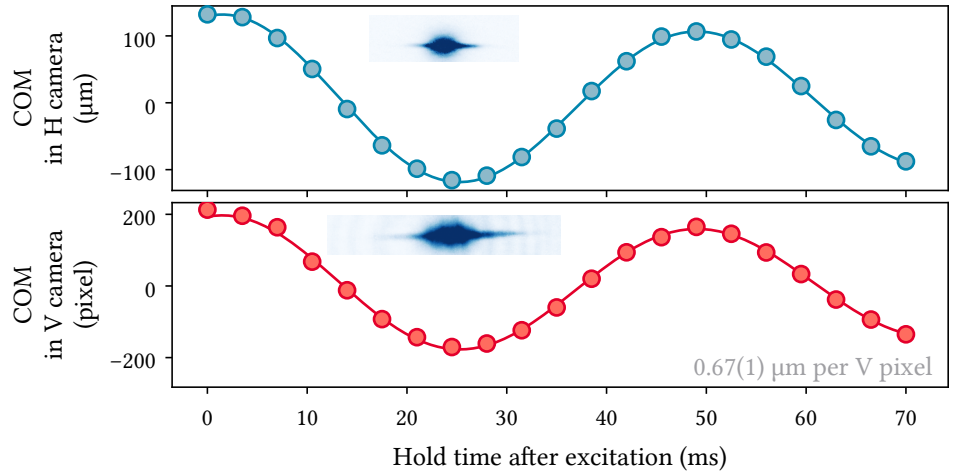

Figure 4.2: Calibration of the vertical imaging system magnification. The dipole mode of the condensate is excited inside the hybrid trap and recorded on both the horizontal and vertical imaging systems. The magnification $M_v = 5.60(9)$ is extracted from the ratio of the oscillation amplitudes measured by the two cameras.

4.2 Characterizing the imaging system

Having introduced basic imaging theory and its impact on measured optical density (OD), we are well positioned to discuss the characterization of aberrations in our imaging system.

In this section, I discuss the magnification calibration of the vertical imaging system before presenting the characterization of the aberrations.

4.2.1 Calibrating the magnification M

The magnification factor M converts the measured length on the camera sensor to the physical length at the plane of atoms. Its accurate determination is extremely crucial for the quantitative study of any condensate dynamics.

For the horizontal imaging system, the camera observes the xz-plane. One common method for determining the magnification M_h is to record the condensate center of mass at different times-of-flight t: as the condensate accelerates under gravity, its trajectory follows $z = gt^2/2$, where $g = 9.8\,\mathrm{m\,s^{-2}}$ is the local gravitational acceleration. Our measured $M_h = 0.98$ agrees closely with the design value of 1.00.

Because the vertical imaging system cannot observe the freefall of a condensate, the calibration of its magnification M_v must rely on an alternative reference. In our case, we excite the dipole mode of the condensate inside the hybrid trap and record the oscillating center-of-mass position on both the horizontal and vertical cameras with standard absorption

imaging. Partial transfer absorption imaging (PTAI) cannot be performed due to the inhomogeneous Zeeman shifts caused by the quadrupole magnetic gradient field in the hybrid trap, which leads to degraded image quality on the vertical camera. To reliably determine M_v, we extract the ratio of the oscillation amplitudes from a joint fit to the center-of-mass trajectories on both cameras, as this method is more robust against noise when determining the center-of-mass positions from highly saturated vertical images. This measurement is shown in Fig. 4.2.

4.2.2 Characterizing aberrations of the imaging system

Quantitative characterization of imaging system aberrations is important both for optical alignment and data analysis: it aids alignment by precisely pinning down errors caused by defocus and astigmatism— which can be corrected by fine-tuning the focus and centering optics along the principal axis—and it enables reconstruction of signals attenuated by aberrations.

Calibrating and aligning an imaging system in ultracold atomic experiments is usually challenging, as the optics need to focus on the atoms inside the vacuum chamber, where placing a standard resolution test chart is impractical. This task is further complicated by the transient nature of the ultracold gas samples to image—they do not last long in the vacuum, at least compared to the time scale of the manual alignment process. Furthermore, measurements are usually destructive, meaning the feedback rate of each alignment move is limited by the experiment's cycle time (~20 s for our typical experiments).

To align and characterize the aberrations of our imaging system, we adopt the method pioneered in Ref. [84]. Instead of imaging a resolution test chart, we image the noise inherent to ultracold gas: thermal density fluctuations and quantum projection noise arising from the PTAI process. The former was utilized in Ref. [84]. At relatively high temperature, the random density fluctuations of a thermal Bose gas exhibit a flat power spectrum up to the thermal de Broglie wavelength λ_{th} in Fourier space (or a δ-function like density-density correlation in real space). According to Eq. (4.36), the average power spectrum of OD images resulting from thermal noise is directly proportional to the squared modulation transfer function $\mathcal{M}$:

$$|\mathcal{F}\{\mathrm{OD}(x,y)\}(k_x,k_y)|^2 \propto |\mathcal{M}(k_x,k_y)|^2 |\mathcal{F}\{\rho_{2\mathrm{D}}(x,y)\}|^2 \tag{4.39}$$

$$\propto |\mathcal{M}(k_x,k_y)|^2. \tag{4.40}$$

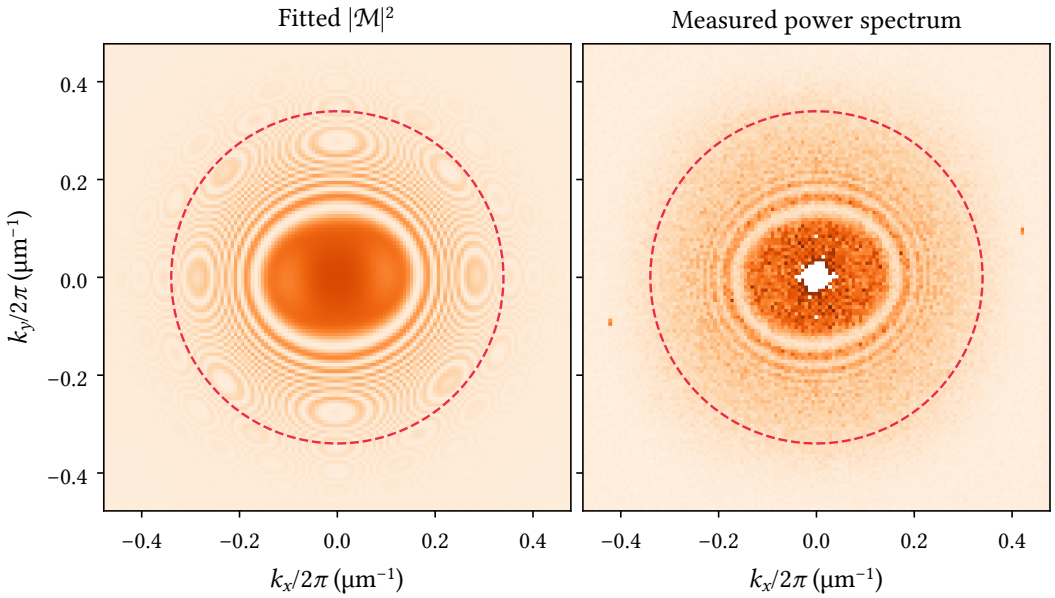

Figure 4.3: Characterizing the aberrations of the imaging system from the noise power spectrum. Left: the best fit of $|\mathcal{M}|^2$ that includes defocus, astigmatism, spherical aberration, and a Gaussian that models the effect of the finite NA. Right: average power spectrum of ~30 OD images acquired using PTAI, after removing leading principal components. Both panels share the same color scale. The saturated low-k region is excluded both from the plot and the fit. The dashed circle marks the diffraction limit $k/2\pi = \mathrm{NA}/\lambda = 0.34\,\mu\mathrm{m}^{-1}$, for NA = 0.2 and $\lambda = 589\,\mathrm{nm}$. The fine rings in $|\mathcal{M}|^2$ result from the spherical aberration, which dominates the total wavefront error in our imaging system. From the fit, we can infer the camera has a mean defocus of 1.16 mm with an astigmatic separation of 1.03 mm. Due to the strong spherical aberration in our system, the best performance is achieved when slightly defocused.

Similarly, when performing PTAI, we drive atoms into a superposition of the "spin-↑" and "spin-↓" states:

$$|\Psi\rangle_{\mathrm{atom}} = c_\downarrow|\downarrow\rangle_{\mathrm{atom}} + c_\uparrow|\uparrow\rangle_{\mathrm{atom}}. \tag{4.41}$$

The measurement collapses each atom into a definite state $|\uparrow\rangle$ or $|\downarrow\rangle$ with probability $|c_{\uparrow/\downarrow}|^2$, independently of the other atoms. This spatially random noise also has the same flat power spectrum whose amplitude is maximized at $|c_\uparrow|^2 = 0.5$.

In practice, we repeatedly prepare ultracold gas samples and let it expand in the sheet trap (similar to Fig. 3.9) before performing PTAI, accumulating ~30 OD images. Before calculating the power spectrum, we first perform principal component analysis (PCA) and remove the first few principal components from each image to separate the structured noise (typically caused by the probe beam fringes and trapping potential imperfections) from the desired random noise. The computed average power spectrum of the images is

then fitted to a model of $\mathcal{M}$ [Eq. (4.38)] derived from a pupil function that accounts for defocus, astigmatism, and spherical aberration. The result of this fit is shown in Fig. 4.3.

4.3 Control of magnetic field gradients in optical traps

In the following chapters, I present studies of a binary condensate consisting of atoms in states $|\downarrow\rangle = |F = 1, m_F = -1\rangle$ and $|\uparrow\rangle = |F = 2, m_F = -2\rangle$. Atoms in these two spin states couple differently to the magnetic field: having distinct magnetic dipole moments $\mu_{\uparrow/\downarrow} = g_F \mu_B m_{F,\uparrow/\downarrow}$, they experience different forces $\mathbf{F}_{\uparrow/\downarrow} = \mu_{\uparrow/\downarrow} \nabla B_z$ where g_F is the Landé g-factor, μ_B is the Bohr magneton and ∇B_z is the gradient of the magnetic field along the magnetization direction z. In our experiments, we trap the condensate in the $x - y$ plane and control $\partial_x B_z$ and $\partial_y B_z$ using two sets of four coils (the "gradient compensation coils") as described in Section 3.2.2. The calibration of these coils has two objectives: first, to measure the current-to-magnetic field gradient ratios for $\partial_x B_z$ and $\partial_y B_z$; second, to determine the currents that cancel any residual magnetic field gradients at the atomic position (likely induced by the long leads of the Z-bias coils and nearby 2D MOT permanent magnets).

This task is complicated by the minuscule variation of B_z across the trapping region and concurrent shifts in the DC magnetic field at different currents, making the most straightforward approach—spectroscopic measurement of spatial variations in magnetic hyperfine transition frequencies—impractical.

First, to measure the current-to-gradient ratios, our approach is to create an isotropic harmonic potential in the $x-y$ plane with varying trap frequencies and measure the shift in the condensate's equilibrium position, as shown in Fig. 4.4. The harmonic trap frequencies are calibrated by abruptly shifting the trap center and observing the center-of-mass oscillations as a function of hold time. Then the current-to-gradient ratio along the x-direction $\partial(\partial_x B_z)/\partial I_i$ (for coil sets $i = 0, 1$) is read out via

$$\frac{\partial}{\partial I_i} \partial_j B_z = m\omega^2 \frac{\partial_{I_i} x_0}{\mu}, \tag{4.42}$$

where I_i is the current in coil set i, x_0 is the equilibrium x-position of the condensate, μ is the magnetic moment of the state, and ω is the harmonic trap frequency; an identical relation applies along the y-direction. The unit vector pointing along the gradient direction produced by coil set i is thus

$$\hat{n}_i \propto \left(\frac{\partial}{\partial I_i} \partial_x B_z, \frac{\partial}{\partial I_i} \partial_y B_z \right). \tag{4.43}$$

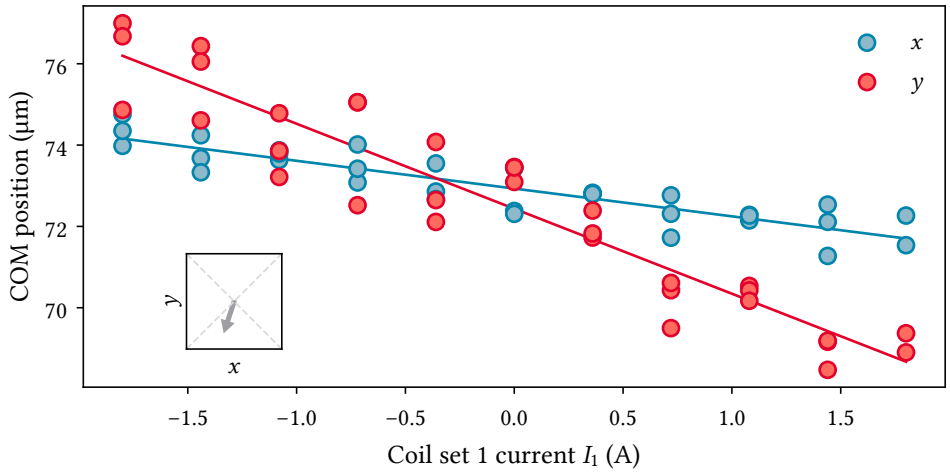

Figure 4.4: Shift of the condensate's equilibrium position in an isotropic harmonic trap as the current in one set of gradient compensation coils is swept. Markers are measured center-of-mass positions at different currents; the two fitted slopes give $\partial(\partial_x B_z)/\partial I_1$ and $\partial(\partial_y B_z)/\partial I_1$. The measured gradient-to-current ratio is $0.044(2)\ \mathrm{G\,cm^{-1}\,A^{-1}}$, pointing at $-117°$ relative to the x-axis. Inset: the unit vector $\hat{n}_1$ points along the gradient direction produced by this coil set.

Second, to find the currents in the two coil sets that cancel the residual gradient at the atomic position, we transfer half of the condensate atoms from the $|\downarrow\rangle$ ground state to the $|\uparrow\rangle$ state and wait for the binary condensate to separate under strong repulsive interactions (see Chapter 5). The two spin states experience magnetic gradient forces of opposite signs, creating a configuration highly sensitive to residual magnetic gradients: even a subtle symmetry-breaking gradient "drags" the $|\uparrow\rangle$ state along its direction, manifesting as a less random spatial separation that aligns with the gradient direction. Experimentally, we sweep the currents inside the two coil sets independently and measure the x- and y-positions of the $|\uparrow\rangle$'s center of mass (Fig. 4.5). We then fit x- and y-positions across different currents to a tanh function to identify the sign-flipping current $I_{i,\mathrm{flip}}$, which corresponds to the current that minimizes the residual magnetic field gradient. In principle, $I_{i,\mathrm{flip}}$ reflects the magnitude of the residual gradient projected onto the direction of the gradient produced by that coil set. Thus $I_{0,\mathrm{flip}}$ and $I_{1,\mathrm{flip}}$ allow us to reconstruct both the magnitude and the direction of the residual gradient.

For our setup, we measured the current-to-gradient ratio for the two coil sets 0 and 1 to be $0.039(2)\ \mathrm{G\,cm^{-1}\,A^{-1}}$ and $0.044(2)\ \mathrm{G\,cm^{-1}\,A^{-1}}$, pointing at $142°$ and $-117°$ in the x-y plane relative to the x-axis. We aligned coil set 0 to the direction of the residual gradient at the atomic position, which has a magnitude of $0.044(2)\ \mathrm{G\,cm^{-1}}$.

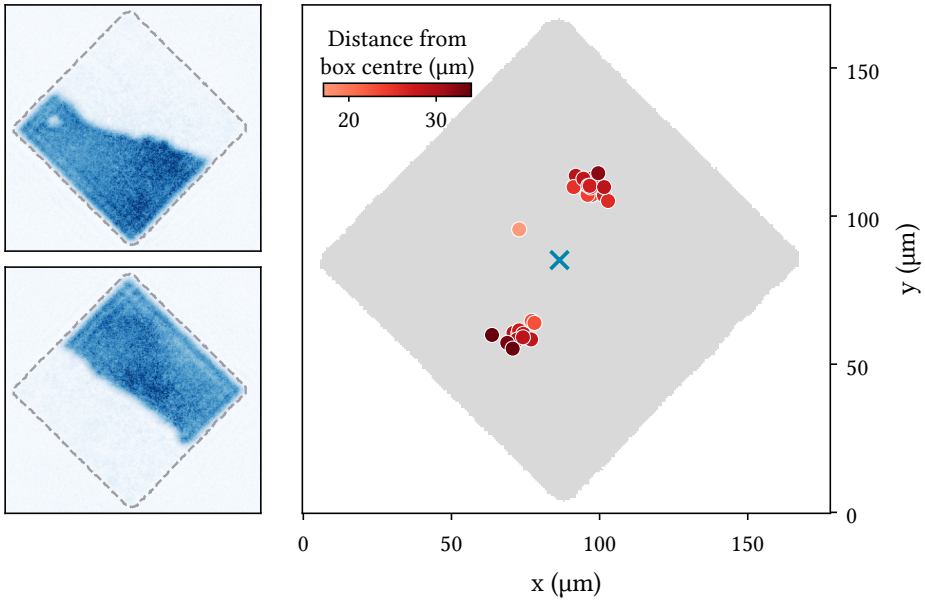

Figure 4.5: Identifying the current in compensation coil set 1 (I_1) that cancels the residual magnetic gradient. Left: Images of the $|\uparrow\rangle$ component at the two ends of the current sweep, which lie on either side of the sign-flipping current $I_{1,\mathrm{flip}}$. The dashed outline marks the box trap region. Right: center of mass of the $|\uparrow\rangle$ component across the whole sweep, colored by its distance from the box center (indicated by ×). The marker closest to the center identifies $I_{1,\mathrm{flip}}$.

4.4 Calibrating the atom number by phonon excitations

The atom number sets the strength of the interaction term in the Gross–Pitaevskii equation, fundamentally dictating condensate dynamics and demanding precise calibration. The most direct route—converting measured optical density (OD) to atomic density—rarely performs well in the real world. Systematic errors quickly corrupt this method: imperfect probe polarization induces optical pumping, the photon scattering cross section σ shifts from its ideal value, and the atomic transition saturates at $I \sim I_{\mathrm{sat}}$. Consequently, raw OD measurements yield notoriously poor results unless σ is independently recalibrated under specific imaging conditions [43]. In Section 3.3.2, we circumvented these issues by inferring the atom number from the expansion dynamics of the condensate. However, this method is incompatible with the quasi-2D box potential, where the specific trap geometry alters the expansion dynamics. An alternative approach is warranted.

Decades of experimental practice have taught us, as Jun Ye put it (and William D. Phillips would surely agree): "The easiest thing to measure as a physical quantity is a frequency." For a condensate trapped in a quasi-2D potential, the most accessible frequency that ties to the atom number is the frequency of the sound modes.

Theory

As laid out in Section 2.6, the sound excitations are described by Bogoliubov quasi-particles with a linear dispersion relation $\epsilon_p \approx c_s p$ at low p. Here, the speed of sound c_s in a homogeneous system is given by

$$c_s = \sqrt{gn/m}. \tag{4.44}$$

In our experiments, the finite-size box potential imposes hydrodynamic boundary conditions on the sound modes (requiring the normal velocity components $\mathbf{v}(\mathbf{r})$ to vanish at the walls). For excitations along a box axis of length L, this discretizes the allowed wave vectors to $k_j = j\pi/L$, where $j = 1, 2, \cdots$. Experimentally, these sound modes manifest as density standing waves in the condensate, oscillating in time at the frequency ω_j. By measuring this oscillation frequency $f_j = \omega_j/2\pi$ (and thus the mode energy $\epsilon_j = \hbar\omega_j$), we extract the speed of sound via

$$c_s = \frac{\omega_j}{k_j} = \frac{2Lf_j}{j}. \tag{4.45}$$

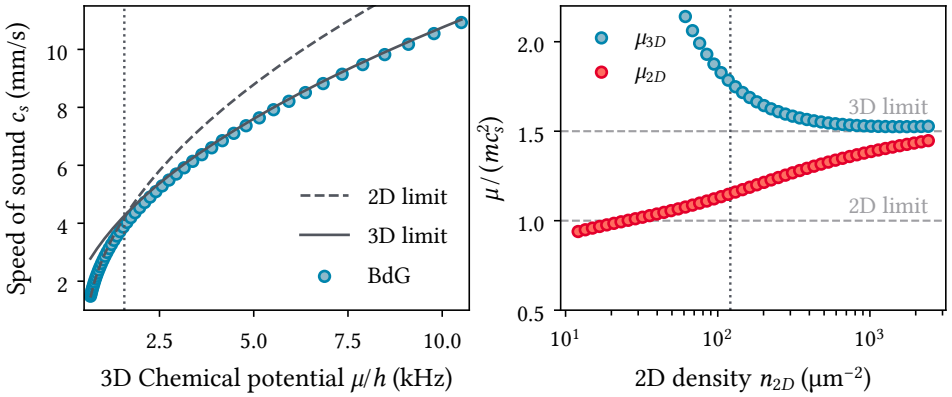

Figure 4.6: Speed of sound in a quasi-2D condensate across the 2D–3D crossover, obtained from the sound mode of the 3D BdG equation reduced to the z-dimension only (see Appendix). The vertical trap frequency is $\omega_z/2\pi = 1100$ Hz. Left: the speed of sound c_s versus chemical potential μ/h; dashed and solid grey curves show the 2D [Eq. (4.46)] and 3D [Eq. (4.51)] limits. Right: the ratio $\mu/(mc_s^2)$ versus 2D density n_{2D}, evaluated using both the 3D chemical potential μ and the 2D-effective chemical potential $\mu - \hbar\omega_z/2$; dashed lines mark the 2D (1) and 3D (3/2) limits. The dotted line marks the typical experimental condition ($n_{2D} \approx 120\,\mu\text{m}^{-2}$, corresponding to 1.6×10^6 atoms in a 115 μm-by-115 μm square box potential).

Now, it is tempting to directly derive the atomic density from Eq. (4.44). However, the condensates are inhomogeneous along the vertical direction, tightly confined by a harmonic trap along z.

Consider first the extreme 2D limit, where vertical harmonic confinement is sufficiently tight so that all motion in the z-direction is frozen in the harmonic oscillator ground state. In this regime, the system reduces strictly to two dimensions, governed by the effective 2D GPE [Eq. (2.57)]. The speed of sound is then simply expressed in terms of the effective 2D interaction g_{2D}, the 2D density n_{2D}, and the 2D chemical potential μ_{2D} [43, 55]:

$$c_s = \sqrt{g_{2D} n_{2D}/m} = \sqrt{\mu_{2D}/m}. \tag{4.46}$$

In the opposite limit, where the vertical confinement $U(z) = m\omega_z^2 z^2/2$ is weak compared to the condensate chemical potential, the system enters the 3D Thomas-Fermi regime along the z-direction, with density

$$n(z) = \begin{cases} \frac{\mu - U(z)}{g} = \frac{\mu - m\omega_z^2 z^2/2}{g} & \text{if } |z| \le z_0 \\ 0 & \text{if } |z| > z_0 \end{cases}, \tag{4.47}$$

where $z_0 = \sqrt{2\mu/m\omega_z^2}$ is the vertical Thomas-Fermi radius. Here, the hydrodynamic approach from Section 2.7 applies: substituting $n(z)$ into the hydrodynamic secular equation Eq. (2.129) yields:

$$\omega^2 \delta n = \omega_z^2 \left[z\, \partial_z \delta n - \frac{1}{2}(z_0^2 - z^2)\nabla^2 \delta n \right]. \tag{4.48}$$

Using the separable ansatz $\delta n = \delta n_0 R_\perp(x, y) Z(z)$ where the planar part $R_\perp \propto e^{i\mathbf{k}_\perp \cdot \mathbf{r}_\perp}$ represents a traveling wave in the xy-plane, we obtain:

$$\gamma^2 Z = uZ' - \frac{1}{2}(1 - u^2)(Z'' - z_0^2 k_\perp^2 Z), \tag{4.49}$$

where dimensionless variables are defined as $\gamma = \omega/\omega_z, u = z/z_0$, with primes denoting derivatives with respect to u.

In the low energy limit $k_\perp \to 0$, Eq. (4.49) reduces to Legendre's equation, with solutions given by Legendre polynomials $Z(z) \propto P_l(z/z_0)$, where l indexes vertical excitations. We set $l = 0$ for a pure planar density wave[9]. The key is to treat the $k_\perp^2$ term perturbatively: substituting Z back to Eq. (4.49) and taking the inner product with $P_l(z/z_0)$ determines the mode energy γ. To leading order in $k_\perp z_0 \ll 1$, straightforward algebra yields:

$$\omega = k_\perp \sqrt{\frac{g\bar{n}}{m}}. \tag{4.50}$$

Here, $\bar{n} = 2n_0/3$ is the vertically averaged condensate density and $n_0 = \mu/g$ is the peak density at $z = 0$. Expressed directly in terms of the chemical potential, the speed of sound becomes:

$$c_s = \sqrt{\frac{g\bar{n}}{m}} = \sqrt{\frac{2}{3}\frac{\mu}{m}}. \tag{4.51}$$

Notably, this result matches the conclusion in Ref. [85] that $c_s = \sqrt{g\bar{n}/m}$ even though the geometry under consideration is entirely different.

However, our actual experiments sit in neither of the two limits. As a result, the factor μ/mc_s^2 lies between the 2D value of 1 and the 3D Thomas-Fermi value of 3/2. Determining the exact relation between μ and c_s requires numerical calculation. In Fig. 4.6, we plot the 3D chemical potential and the speed of sound alongside the ratio μ/mc_s^2 obtained from a numerical BdG calculation (see Appendix). The vertical dotted line marks the typical con-

[9] For $l = 1$, the same treatment yields a mode dispersion relation $\omega^2 = (2\mu/5m)k_\perp^2 + \omega_0^2$, indicating that the energy of any $l = 1$ excitation will be at least one quanta of the sheet trap energy $\Delta E/\hbar = \omega_0 \sim 1\,\mathrm{kHz}$.

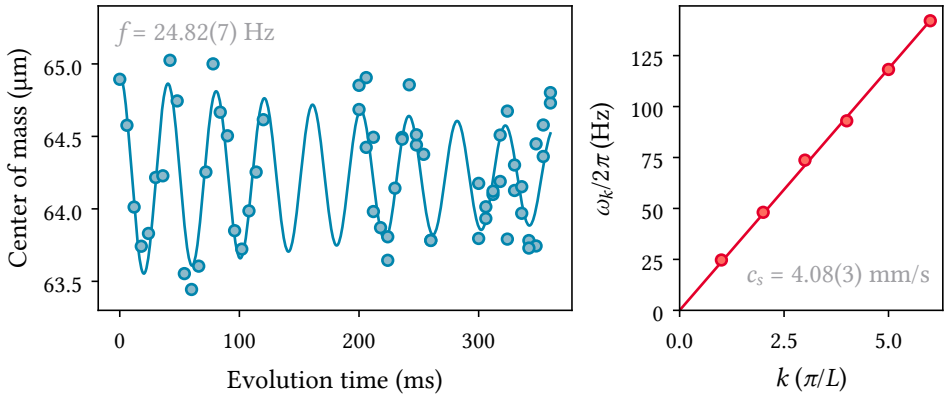

Figure 4.7: Fundamental mode evolution and dispersion in the quasi-2D box potential. Left: Condensate center of mass after imprinting the slosh pattern, oscillating at the frequency ω_1; the solid line is a damped-sinusoid fit, with the fitted frequency annotated on the plot. Right: Phonon frequency measured at various k (in units of π/L). The solid line is a linear fit through the origin, yielding $c_s = 4.08(3)\,\mathrm{mm\,s^{-1}}$.

densate 2D density in our quasi-2D box potential, which clearly falls within the crossover regime. For dynamic simulations of the condensate where only the correct μ is needed, one can circumvent knowing N entirely by constructing a lookup table that directly maps the measured speed of sound to the chemical potential μ. However, the thermometry protocol developed in the following chapter demands an accurate absolute density calibration. Establishing N via sound mode measurements is therefore indispensable.

Experimental realization

To excite the sound modes experimentally, we imprint a subtle sinusoidal density modulation inside the box potential by displaying a half-toned sinusoidal pattern (the "slosh pattern") via the DMD during the evaporation and transfer process. Following a brief hold time, we abruptly switch the DMD pattern to a uniform box potential and track the ensuing time evolution of the condensate density. In the quasi-particle picture, this protocol initializes a coherent state of both $-k$ and $+k$ phonons, whose amplitude is set by the amount of the initial density imbalance. This is, by far, the simplest method for exciting phonons in the condensate. Alternatively, one can apply a magnetic gradient impulse [56], or oscillate either the phase of the slosh pattern or the amplitude of the magnetic gradient to map out the frequency response [86].

When imprinting the slosh pattern, the spatial periodicity of the sinusoidal potential projected by the DMD selects the wavevector $k_j = j\pi/L$ of the excited mode. The mode amplitude as a function of evolution time is extracted by performing either a discrete co-

sine transform (DCT) or principal component analysis (PCA) on the OD images—or, for the fundamental mode ($j = 1$), by simply tracking the center-of-mass motion. Figure 4.7 shows a representative measured center-of-mass oscillation for the fundamental mode alongside the full phonon dispersion relation $\omega(k)$ across multiple wavevectors k. Extracting the slope of a linear fit to $\omega(k)$ directly yields the speed of sound c_s, from which the atom number is determined by comparing with the numerical results shown in Fig. 4.6. The damping visible in Fig. 4.7 is a common feature in similar experiments (usually has a lifetime $\tau \sim 10^2$ ms) and is commonly attributed to the finite temperature of the system [55]. Empirically, the observed damping behavior varies across experiments, sensitive to the quality of the evaporation and transfer protocols. A complete model of the damping remains an open challenge [56, 86]. Fortunately, because our condensates operate deep in the degenerate regime $T < 0.1T_c$, temperature-induced frequency shifts are mostly negligible.

One benefit of this "slosh" measurement is that the extracted atom number achieves decent accuracy compared to other methods (for example, by integrating OD images). The atom number scales as $N \propto c_s^2$ in the 2D limit or $N \propto c_s^3$ in the 3D Thomas-Fermi regime; therefore a 0.7 % error in c_s corresponds to a fairly minimal 1.5 % to 2.2 % in N. In practice, the total accuracy is usually not limited by the uncertainty in the speed of sound, but by the error in the magnification calibration of Section 4.2, which dictates the box size L.

5 Study of Ripplons and the Rayleigh-Taylor Instability in a Binary Condensate

This chapter presents the first of two original scientific contributions of this thesis: the study of ripplon dynamics and the Rayleigh-Taylor instability (RTI) in an immiscible binary condensate of ^{23}Na atoms, published in Ref. [33]. Ripplon excitations and the Rayleigh-Taylor instability represent two sides of the same coin—the stable and unstable branches of the interfacial dispersion relation. Theoretical predictions for both phenomena in binary condensates date back to the early 2000s [32, 87]; our experiments provide their first direct observation. These results demonstrate the emergence of fundamental fluid instabilities in quantum settings, advancing our understanding of quantum hydrodynamics in immiscible condensates.

We established the theoretical foundation for immiscible BECs in Section 2.5. In this work, we select the two states from the ground-state hyperfine manifold, $|F = 1, m_F = -1\rangle \equiv |\downarrow\rangle$ and $|F = 2, m_F = -2\rangle \equiv |\uparrow\rangle$, to form our binary quantum fluid. This choice offers two distinct advantages. First, as "stretched states", their inter-spin inelastic collisions are inhibited, granting the system a long coexistence lifetime. Second, the intra- and inter-spin scattering lengths—$a_{\uparrow\uparrow} = 64.30(40)\ a_B$, $a_{\downarrow\downarrow} = 54.54(20)\ a_B$, $a_{\uparrow\downarrow} = 64.3\ a_B$ [88–90] yield an interaction parameter ratio of $g_{\uparrow\downarrow}/\sqrt{g_{\uparrow\uparrow}g_{\downarrow\downarrow}} \approx 1.1$ [1]. Consequently, this combination readily satisfies the immiscibility condition, Eq. (2.75), while forming a "sharp" interface with minimal interpenetration between two condensates[2].

In this chapter, we focus on the dynamical properties of this interface. The distinct energy cost the interface carries—the interfacial tension—governs the system's characteristic

[1] The scattering length connects to interaction strength via $g = 4\pi\hbar^2 a/m$ [Eq. (3.5)] where m is the ^{23}Na atomic mass. This value of $a_{\uparrow\downarrow}$ was first reported in Ref. [88], derived from existing Feshbach spectroscopy data [90]. We assume an uncertainty of $0.4a_B$, consistent with other scattering lengths reported in Ref. [90].

[2] By comparison, the combination of $|F = 1, m_F \pm 1\rangle$ and $|F = 1, m_F = 0\rangle$ is also immiscible, but their $g_{\uparrow\downarrow}/\sqrt{g_{\uparrow\uparrow}g_{\downarrow\downarrow}} \approx 1.02$. Consequently, the interpenetration depth is large, resulting in a spatially broad, "fuzzy" interface. Also, the low interfacial energy scale makes the initial preparation of a clean, flat interface experimentally challenging.

dynamics. The underlying physical picture mirrors the familiar oil-water interface in classical fluids. Indeed, the dispersion relation of the quantum fluid interface closely tracks its classical counterpart across both stable and unstable regimes. This correspondence motivates our investigation of the ripplons—the quantum analogue of capillary waves—and the quantum realization of the ubiquitous Rayleigh-Taylor instability.

The material for this chapter is adapted from our work published in Ref. [33]. The chapter is organized as follows: first, I motivate the ripplon dispersion relation by analyzing the interfacial motion under classical approximations. Next, I present our experimental measurement of the ripplon dispersion relation, followed by the observation and characterization of the Rayleigh-Taylor instability. Finally, I discuss the matter-wave interferometry technique used to visualize the local counterflow velocity across the quantum fluid interface.

5.1 Dispersion relation of a quantum fluid interface

In Section 2.7, we discussed the similarities between Bose-Einstein condensates and classical fluid dynamics. Here, we exploit this analogy to motivate the interfacial dispersion relation from the dynamics of a classical fluid-fluid interface.

We begin by recalling that the Gross-Pitaevskii equation maps directly onto Euler's equation in the classical fluid dynamics, Eq. (2.127):

$$m\frac{\partial \mathbf{v}}{\partial t} = -\nabla\left(\frac{1}{2}m\mathbf{v}^2 + U + gn - \frac{\hbar^2}{2m\sqrt{n}}\nabla^2\sqrt{n}\right), \tag{5.1}$$

where m is the atomic mass and n is the local number density, and the last term represents the quantum pressure.

For low-amplitude perturbations along a 1D interface aligned in the x-direction, bulk density variations are negligible; the condensate is thus mostly incompressible and the quantum pressure can be neglected away from the interface, which is treated as a hard, sharp boundary. Under these assumptions, Eq. (5.1) reduces to the classical Euler equation for a potential flow, allowing us to adapt the standard hydrodynamic derivation for the capillary wave dispersion relation. In classical fluid dynamics, a potential flow refers to the motion of an inviscid fluid in the bulk. Its flow is irrotational ($\nabla \times \mathbf{v} = 0$), enabling the definition of a scalar velocity potential ϕ (hence the "potential" in its name) such that $\mathbf{v} = \nabla\phi$. This relation closely maps to Eq. (2.126), which connects the local condensate velocity and phase by $\mathbf{v} = \hbar\nabla\varphi/m$.

If the fluid experiences a body force F per particle opposite to the y-direction (e.g., gravity or a magnetic gradient force), the potential takes the form $U = F\eta$, where η measures the displacement along the y-direction. Linearizing Euler's equation for small flow velocities $\mathbf{v}$ and making use of the incompressibility yields [91, §12]:

$$\frac{\partial \mathbf{v}}{\partial t} + \nabla\left(\frac{p}{mn} + \frac{F}{m}\eta\right) = 0, \tag{5.2}$$

where $p = gn^2/2$ is the pressure.

Mechanical equilibrium requires equal pressure across the interface, except when the interface stores energy—in which case the interfacial tension enters the pressure balance equation. Near equilibrium, balancing the work required to stretch or shrink the interface and the work needed to displace the fluids leads to the Young-Laplace equation, relating pressure difference $p_2 - p_1$ (with fluid 1 below fluid 2) to the local interface curvature [91, §61]:

$$p_2 - p_1 = \sigma \frac{\partial^2 \eta}{\partial x^2}. \tag{5.3}$$

Substituting this boundary condition into Eq. (5.2) for the two fluids of equal particle mass and using $\partial_t \eta = v_y = \partial_y \phi$ at the interface yields the wave equation

$$-n_2 F_2 \frac{\partial \phi_2}{\partial y} - m n_2 \frac{\partial^2 \phi_2}{\partial t^2} + n_1 F_1 \frac{\partial \phi_1}{\partial y} + m n_1 \frac{\partial^2 \phi_1}{\partial t^2} - \sigma \frac{\partial^2}{\partial x^2}\frac{\partial}{\partial t}\eta = 0, \tag{5.4}$$

which is solved together with the matching condition for normal velocity across the interface, $\partial_y \phi_1 = \partial_y \phi_2$. Under the oscillatory ansatz $\phi \propto e^{-k|y|}\cos(kx - \omega t)$, we obtain the dispersion relation:

$$\omega^2 = \frac{1}{2m}\left(\mathcal{F}k + \frac{\sigma}{\bar{n}}k^3\right), \tag{5.5}$$

with density-weighted force difference defined as:

$$\mathcal{F} \equiv F_1 \frac{n_1}{\bar{n}} - F_2 \frac{n_2}{\bar{n}}, \tag{5.6}$$

where $\bar{n} = (n_1 + n_2)/2$ is the average number density.

Now, let us take a step back and ask: how much of this classical derivation is applicable to the binary quantum fluid? First, Eq. (5.2) assumes the fluid to be incompressible. A condensate, however, is compressible, and the spatial interpenetration of the two condensates across the interface changes slightly under ripplon excitations. This occurs not

only at high ripplon energies (high k) but also in the long-wavelength limit ($k \to 0$): the incompressibility approximation requires sub-sonic ripplon phase velocities $\omega/k \ll c$, a condition that fails at low k when $\mathcal{F} > 0$, which we examine further in the next section. Furthermore, this hydrodynamic approach treats the interfacial tension σ as a purely phenomenological parameter. In a quantum fluid, however, the quantum pressure accounts for a major portion of interfacial tension (in fact, half of it when the bulk is homogeneous)—an effect omitted in Eq. (5.2) and absorbed implicitly into σ.

Nonetheless, extensive theoretical work [32, 60, 92, 93] demonstrates that a dispersion relation mostly identical to Eq. (5.5) can be derived *ab initio* from the Gross-Pitaevskii equation and the Bogoliubov-de Gennes framework. Detailed numerical studies [87, 92, 93] confirm the validity of Eq. (5.5) in a binary quantum fluid at low energies.

It is clear from Eq. (5.5) that the right-hand side can be either positive or negative, determined by the sign of $\mathcal{F}$. In the former case, the excitation frequency is real, describing stable ripplon modes. In the latter case, the forces push the two components together, creating an unstable configuration with a purely imaginary ripplon mode frequency. The perturbation amplitude grows exponentially in time, leading to the Rayleigh-Taylor instability. We present the experimental results for both cases in the following sections.

5.2 Experimental measurement of the ripplon dispersion in a binary condensate

Our experimental investigation begins with a homogeneous quasi-2D condensate in the $|{\downarrow}\rangle$ state, prepared in a square box potential of side length $L = 114(3)$ μm in the xy-plane. We apply a $\pi/2$ microwave pulse resonant with the $|{\downarrow}\rangle \to |{\uparrow}\rangle$ transition to create an equal superposition of both $|{\downarrow}\rangle$ and $|{\uparrow}\rangle$ states. As the two spins begin to phase-separate, we apply a small magnetic field gradient $B' \equiv \partial_y B_z$ along the y-direction through our gradient compensation coils (Sec. 3.2.2). This magnetic gradient exerts state-dependent Stern-Gerlach forces on the two components, establishing the interface normal along y. In this case, the differential force $\mathcal{F}$ takes the form:

$$\mathcal{F} = \mu_{\uparrow} B' \frac{n_{\uparrow}}{\bar{n}} - \mu_{\downarrow} B' \frac{n_{\downarrow}}{\bar{n}} \tag{5.7}$$

$$\approx \mu_{\uparrow} B' - \mu_{\downarrow} B', \tag{5.8}$$

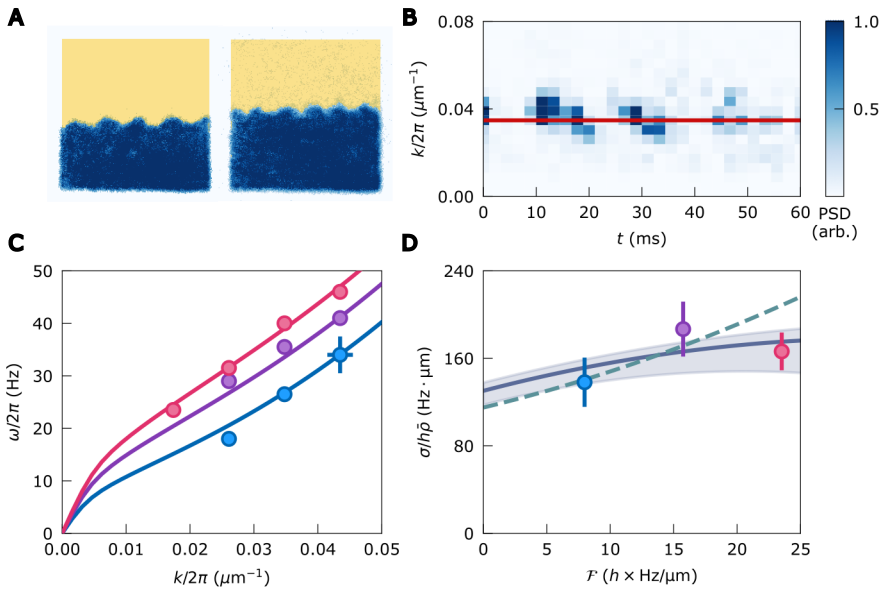

Figure 5.1: Excitation and detection of ripplons (reproduced from Ref. [33]). (A) Parametrically excited ripplons with $\mathcal{F}/h = 15.4(8)\,\mathrm{Hz\,\mu m^{-1}}$. Excitation frequencies $\omega_\mathrm{d}/2\pi = 74\,\mathrm{Hz}$ (left) and 87 Hz (right) yield $k/2\pi = 0.035\,\mathrm{\mu m^{-1}}$ and $0.043\,\mathrm{\mu m^{-1}}$, respectively. (B) PSD of the height field of a freely evolving interface after excitation. The solid line marks the k of the excited mode. (C) Ripplon spectrum measured at $\mathcal{F}/h = 7.7(4)\,\mathrm{Hz\,\mu m^{-1}}$ (blue), $15.4(8)\,\mathrm{Hz\,\mu m^{-1}}$ (purple) and $23(1)\,\mathrm{Hz\,\mu m^{-1}}$ (red). Dots mark experimental measurements, and curves denote numerical simulations. (D) Normalized interface tension. The figure labels the average number density $\bar{n}$ as $\bar{\rho}$, following the notation of Ref. [33]. The blue solid curve results from fitting Eq. (5.5) to the Bogoliubov spectrum. The green dashed line is the interfacial tension computed from the free-energy difference, Eq. (5.10). Error bars and bands mark 1σ uncertainties.

where $\mu_\uparrow$ and $\mu_\downarrow$ denote the magnetic moments of the two spin states, and the second line is justified by the negligible density difference between the two components[3].

We deploy a parametric driving scheme to excite ripplons with a well-defined wavevector k and frequency ω. Starting from an almost flat interface, we modulate the magnetic field gradient B' at frequency ω_d, introducing a force modulation $\delta\mathcal{F} \propto \cos(\omega_\mathrm{d} t)$. At resonant drive frequencies ω_d, standing-wave ripplon modes emerge, manifesting as sinusoidal deformations of the interface height $\eta(x)$. Unlike standard textbook parametric resonances, here $\omega_\mathrm{d} \neq 2\omega$ due to the complex dynamics of the drive process. To measure ω, we halt the drive once the desired mode is well-established, allowing the ripplons to evolve freely. We then capture absorption images of the $|\uparrow\rangle$ condensate at various evolution times [Fig. 5.1(A)]. Since the parametric drive amplifies initial interfacial fluctuations with arbi-

[3] For the experiments reported here, assuming $n_\uparrow = n_\downarrow$ introduces an error in $\mathcal{F}$ of $\approx 2\%$, well within uncertainties from other sources.

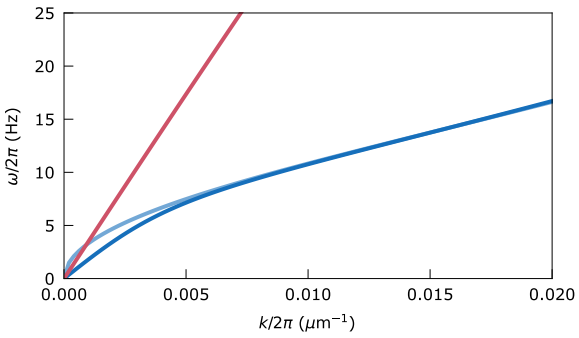

Figure 5.2: Dispersion of the two lowest gapless modes (reproduced from the Supplementary Material of Ref. [33]). The dark blue and red curves show the dispersion relations of the ripplon and phonon modes, respectively. The light blue curve is a fit of the ripplon dispersion to Eq. (5.5), yielding $\sigma/\bar{n}$.

trary spatial phases, single-shot measurements vary from run to run. We therefore analyze the power spectral density $\mathrm{PSD}(k) = |\tilde{\eta}(k)|^2$ (with the tilde denoting the Fourier transform), which is inherently phase-independent. As shown in Fig. 5.1(B), tracking the PSD over time reveals both the wavevector k and the frequency ω of the freely evolving ripplon mode. Figure 5.1(C) displays the ripplon dispersion relation at three values of B'. These results show good agreement with the Bogoliubov spectra (solid curves) computed for the experimental conditions with no adjustable parameters. The discrete wavevectors $k = j\pi/L$ ($j = 1, 2, \ldots$) reflect the finite geometry imposed by the box potential.

It is worth noting that the measured dispersion relation deviates from the prediction of Eq. (5.5) at high and low k. At short wavelengths (high k), deformation of the finite-width interfacial layer alters the mode energy, while at long wavelengths (low k), coupling to the phonon branch flattens the dispersion relation. To illustrate the latter, Fig. 5.2 plots a zoomed-in view of the Bogoliubov spectrum for both the phonons and ripplons near $k = 0$, alongside a fit to Eq. (5.5). At small k, the phonon mode intersects the classical curve given by Eq. (5.5), signaling the breakdown of the incompressibility assumption invoked in Section 5.1. The resulting avoided crossing forces a departure from the characteristic $\omega \propto \sqrt{k}$ scaling at very small k.

5.3 Experimental observation of the Rayleigh-Taylor instability in a binary condensate

As established in Sec. 5.1, reversing the sign of the differential force ($\mathcal{F} < 0$) causes the mode energies to become purely imaginary across a band of k, signaling a dynamical in-

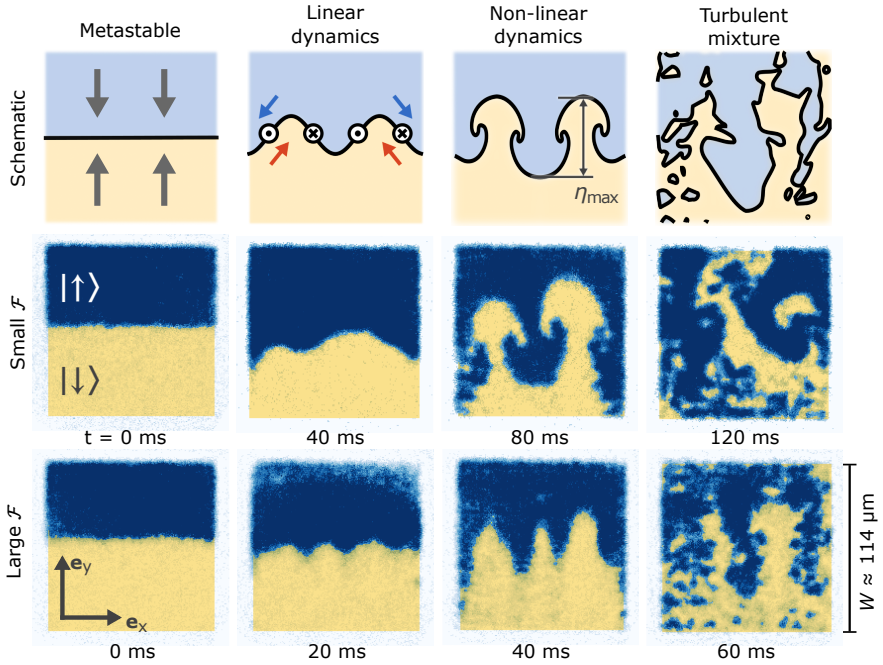

Figure 5.3: Progression of the Rayleigh-Taylor instability (reproduced from Ref. [33]), showing the atomic density in $|\uparrow\rangle$ (blue) and $|\downarrow\rangle$ (yellow). From top to bottom: schematic illustrations, experiments with $\mathcal{F}/h = -3.1(2)\,\mathrm{Hz\,\mu m^{-1}}$ (small $\mathcal{F}$) and experiments with $\mathcal{F}/h = -7.7(4)\,\mathrm{Hz\,\mu m^{-1}}$ (large $\mathcal{F}$). Initial forces (grey arrows) destabilize a fluid interface. The interface then develops nominally sinusoidal modulations where counterflowing currents in the two fluids (colored arrows) induce vorticity at the interface (black symbols). The modulations subsequently acquire characteristic mushroom- or spike-like shapes before dissolving into a turbulent mixture.

stability in the system—the quantum counterpart of the well-known Rayleigh-Taylor instability (RTI) in classical fluid dynamics [94–96].

Instabilities are commonplace across scales: from droplet formation via the Plateau-Rayleigh instability [97, 98] and galactic structure formation via fluid-gravitational instabilities [99], to the destabilization of fusion plasmas via magnetohydrodynamic instabilities [100, 101]. In particular, the RTI manifests in settings ranging from laboratory experiments to astronomical phenomena [96, 102, 103]. At its root, the RTI is triggered when two immiscible fluids are "pressed" together. In this configuration, the two fluids tend to interpenetrate; any small perturbations of the interface grow exponentially. Shear generated along the deformed interface feeds the secondary Kelvin-Helmholtz instability, turning the interface into characteristic bubble-, spike-, and mushroom-like structures before it dissolves into a turbulent mixture [104]. We illustrate this progression in Fig. 5.3.

In our experiment, the binary condensate is initially held under $B' = 4.4\,\mu\mathrm{T\,cm^{-1}}$. Because of the opposite signs of the magnetic moments $\mu_\uparrow, \mu_\downarrow$, the magnetic field gra-

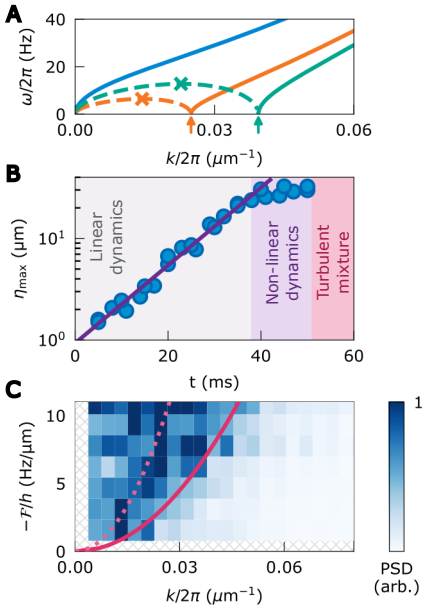

Figure 5.4: Growth of unstable interface modes (reproduced from Ref. [33]). (A) Dispersion of interface excitations computed for $\mathcal{F}/h = 16\,\mathrm{Hz\,\mu m^{-1}}$, $-3.1\,\mathrm{Hz\,\mu m^{-1}}$ and $-7.7\,\mathrm{Hz\,\mu m^{-1}}$ (blue, red and green, respectively) using Eq. (5.5). Solid and dashed curves plot the real and imaginary components, respectively, with k_c marked by arrows and k_{max} by crosses. The ratio $\sigma/\bar{n}$ used for evaluating k_c and k_{max} is the average value computed at a few small $|\mathcal{F}|$. (B) Interface height η_{max} as a function of evolution time t for $\mathcal{F}/h = -7.7(4)\,\mathrm{Hz\,\mu m^{-1}}$. The solid curve is an exponential fit to data in the "linear dynamics" regime. (C) Power spectral density of the interface height measured during the exponential growth regime. The red curve plots the threshold wavevector k_c and the dashed curve plots the wavevector with maximum gain k_{max}.

dient exerts opposing Stern-Gerlach forces that stabilize the interface, producing $\mathcal{F}/h = 9.3(5)\,\mathrm{Hz\,\mu m^{-1}}$. At $t = 0$, we trigger the RTI process by rapidly inverting the magnetic field gradient B'. Figure 5.3 shows the resulting evolution of the RTI at $\mathcal{F}/h = -3.1(2)\,\mathrm{Hz\,\mu m^{-1}}$ (small $\mathcal{F}$) and $\mathcal{F}/h = -7.7(4)\,\mathrm{Hz\,\mu m^{-1}}$ (large $\mathcal{F}$). Upon entering the unstable configuration, sinusoidal modulations emerge on the interface and grow rapidly in amplitude.

For large $\mathcal{F}$, the dominant perturbation shifts toward higher k with steeper growth rates. This behavior follows directly from Eq. (5.5), which predicts unstable mode amplitudes grow exponentially with rate $\Gamma = \mathrm{Im}(\omega)$. The band of unstable modes terminates at an upper threshold $k_c \equiv \sqrt{-\mathcal{F}\bar{n}/\sigma}$, beyond which the positive interfacial tension in Eq. (5.5) dominates, restoring real frequencies ω [Fig. 5.4(A)]. The maximal growth rate occurs at $k_{max} = k_c/\sqrt{3}$, marked by crosses in Fig. 5.4(A). To quantify this growth, we extract the peak-to-peak interface height η_{max}, which exhibits clear exponential growth during early times, as shown in Fig. 5.4(B) for $\mathcal{F}/h = -7.7(4)\,\mathrm{Hz\,\mu m^{-1}}$. An exponential fit in the early linear regime yields $\Gamma = 88(3)\,\mathrm{s^{-1}}$, agreeing well with the theoretical value of $80(5)\,\mathrm{s^{-1}}$

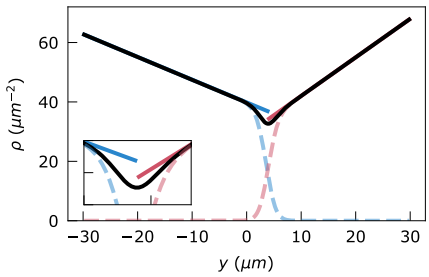

Figure 5.5: Spin-resolved density profile under $\mathcal{F}/h = 7.3(4)\,\mathrm{Hz\,\mu m^{-1}}$ (reproduced from the Supplementary Material of Ref. [33]), with the $|\uparrow\rangle$ and $|\downarrow\rangle$ densities and their sum shown in blue, red and black, respectively. The numerically computed profiles (dashed blue and red) are overlaid with the corresponding LDA estimates (solid blue and red), which ignore the inter-spin interaction at the interface. The inset shows an expanded view of the region near the interface.

computed from Eq. (5.5). At longer times, the linearized model breaks down as evidenced by the appearance of mushroom- and spike-like structures in Fig. 5.3. In this non-linear regime, the interface growth rate saturates from exponential to terminal linear growth before the interface completely disintegrates, as annotated in Fig. 5.4(B).

We analyze the amplitude of all unstable modes triggered across different $\mathcal{F}$ by evaluating their PSD, which quantifies the mode amplitude distribution across spatial frequencies. As shown in Fig. 5.4(C), the band of unstable modes expands with increasing $|\mathcal{F}|$, with its upper threshold closely tracking the calculated k_c.

5.4 Obtaining the interfacial tension

Extracting interfacial tension σ from Eq. (5.5) is subtle: fitting Eq. (5.5) yields the ratio $\sigma/\bar{\rho}$ [Fig. 5.1(D)]; an accurate determination of the mean density $\bar{\rho}$ is warranted to isolate σ.

Defined as the free energy cost of creating an interface, the interfacial tension for a homogeneous binary condensate is given by [31, 32, 92] (see Appendix A)

$$\sigma = 2\int_{-\infty}^{+\infty} \mathrm{d}y \frac{\hbar^2}{2m}\left(|\partial_y\psi_\uparrow|^2 + |\partial_y\psi_\downarrow|^2\right) \tag{5.9}$$

where $\psi_\uparrow$ and $\psi_\downarrow$ are the condensate wavefunctions for the two spin components. However, atomic BECs are compressible; the magnetic field gradient induces significant bulk density gradients. As shown in Fig. 5.5, $\partial_y\psi_{\uparrow,\downarrow}$ remains non-zero even deep in the bulk, causing the integral in Eq. (5.9) to accumulate spurious bulk contributions. To isolate the interfacial tension, these bulk contributions need to be excluded.

Prior theoretical treatments [31, 105] circumvent this difficulty by invoking a form of local density approximation (LDA)—identifying an effective atomic density at the interface and applying Eq. (5.9) to an equivalent homogeneous profile. However, this method underestimates σ for our system due to the significant density inhomogeneity created by the magnetic gradient.

To avoid this difficulty, we directly compute σ as the per-unit-length difference in free energy ($\Omega = E - \sum_j \epsilon_j N_j$) between the full condensate wavefunctions and the Thomas-Fermi baseline given by the LDA under experimental conditions[4]:

$$\sigma = \left[\Omega(\psi_\uparrow, \psi_\downarrow) - \Omega(\psi_{\mathrm{LDA},\uparrow}, \psi_{\mathrm{LDA},\downarrow})\right]/L, \tag{5.10}$$

where $\psi_{\mathrm{LDA},j} = \sqrt{(\epsilon_j - U_j)/g_{jj}}$ inside the domain of spin-j and zero elsewhere, and ϵ_j is the chemical potential for spin-j.

The underlying free energy functional $\Omega(\phi_\uparrow, \phi_\downarrow)$ takes the standard form:

$$\Omega(\phi_\uparrow, \phi_\downarrow) = \int \mathrm{d}x\,\mathrm{d}y \Big\{ \frac{\hbar^2}{2m}|\nabla\phi_\uparrow|^2 + \frac{\hbar^2}{2m}|\nabla\phi_\downarrow|^2 \tag{5.11}$$

$$+ \left[V_\uparrow(y) - \epsilon_\uparrow\right]|\phi_\uparrow(y)|^2 + \left[V_\downarrow(y) - \epsilon_\downarrow\right]|\phi_\downarrow(y)|^2 \tag{5.12}$$

$$+ \frac{g_{\uparrow\uparrow}}{2}|\phi_\uparrow(y)|^4 + \frac{g_{\downarrow\downarrow}}{2}|\phi_\downarrow(y)|^4 + g_{\uparrow\downarrow}|\phi_\uparrow(y)|^2|\phi_\downarrow(y)|^2 \Big\}. \tag{5.13}$$

Taking $\bar{\rho}$ as the LDA density at the interface, the resulting ratio $\sigma/\bar{\rho}$ is plotted as the green dashed curve in Fig. 5.1(D).

5.5 Visualizing counterflow velocity at the interface

Unlike the abstract velocity potential ϕ of a classical fluid, the relative phase φ of condensates is directly accessible through interference experiments. To demonstrate this, we express the wavefunctions of a binary BEC as a two-component vector $\Psi(\mathbf{r}) = [\psi_\uparrow(\mathbf{r}), \psi_\downarrow(\mathbf{r})]^T$, where $\psi_j(\mathbf{r}) = \sqrt{\rho_j(\mathbf{r})}\exp\left[i\varphi_j(\mathbf{r})\right]$ represents the wavefunction of an individual component $j = \uparrow, \downarrow$. This vector is conveniently represented on the Bloch sphere [Fig. 5.6(B)]. In

[4] Whether the interfacial tension should be computed via the energy E or the free energy Ω was historically debated [30, 31]. In both schemes, the calculation finds the excess E or Ω of an actual system with an interface against a reference system "without" an interface—where the kinetic energy term is dropped and the density is solely given by LDA (Thomas-Fermi approximation). While both yield expressions structurally similar to Eq. (5.9), they differ by a factor of 4. However, the depletion of density at the interface results in a reduced total atom number in the realistic system relative to the LDA baseline. An injudicious subtraction in energy alone inadvertently incorporates the cost of adding atoms to the trap (the chemical potential ϵ). Defining σ via $\Omega = E - \epsilon N$ removes this particle-number mismatch naturally. See Appendix A for a detailed explanation.

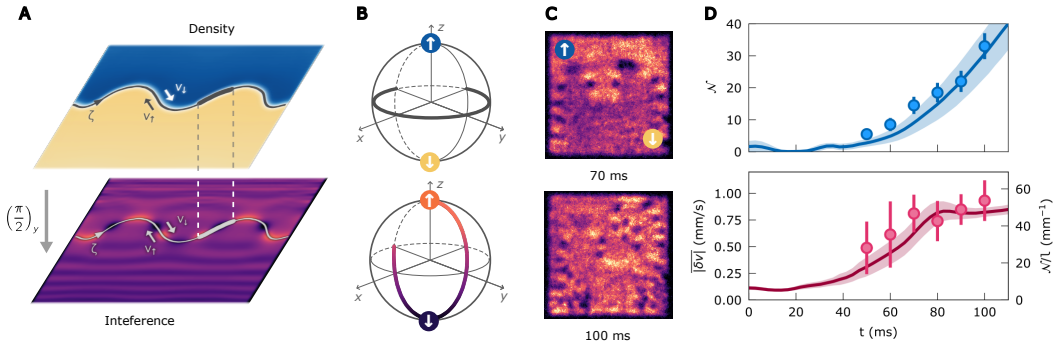

Figure 5.6: Interface velocimetry (reproduced from Ref. [33]). (A) Interferometry scheme showing the original density distributions $\rho_{\uparrow,\downarrow}$ (top) and a simulated interference pattern in $\rho_\uparrow$ after the $\pi/2$-pulse (bottom). The solid curve marks the interface; ζ is the spatial coordinate defined on the interface, and the arrows indicate the fluid velocity on each side of the interface. (B) Spin vector Ψ on the Bloch sphere before (top) and after (bottom) the $\pi/2$-pulse. The arcs colored gray (top) and orange/purple (bottom) identify the spin orientation on the bold segments in (A). (C) Experimental images of $\rho_\uparrow$ after the microwave $\pi/2$-pulse. (D) Time evolution of the vortex number (top) and the average magnitude of the interfacial counterflow velocity $\overline{|\delta v|}$ (bottom). The markers represent measurements and the solid curves are numerical results; error bars and bands mark 1σ uncertainties.

the bulk $|\uparrow\rangle$ domain, this Bloch vector points to the north pole, whereas in the bulk $|\downarrow\rangle$ domain, it points to the south pole. At the interface, where the two densities coincide ($\rho_\uparrow = \rho_\downarrow$), the Bloch vector lies in the equatorial plane with its azimuthal angle equaling the phase difference $\delta\varphi = \varphi_\uparrow - \varphi_\downarrow$.

On the other hand, the velocity difference—or counterflow—across the interface is given by:

$$\Delta\mathbf{v} \equiv \mathbf{v}_\uparrow - \mathbf{v}_\downarrow = \hbar\nabla(\varphi_\uparrow - \varphi_\downarrow)/m \tag{5.14}$$

$$= \hbar\nabla\delta\varphi/m. \tag{5.15}$$

As a result, variations in $\delta\varphi$ along the interface directly encode the local counterflow velocity $\delta v = \hbar\partial_\zeta\delta\varphi/m$, where ζ is the interface coordinate, as indicated in Fig. 5.6(A).

Experimentally, we interfere the spin-$\uparrow$ and -$\downarrow$ wavefunctions by driving a microwave $\pi/2$ pulse, which rotates the Bloch vector about the y-axis. This rotation maps the flow pattern in the vicinity of the interface into a chain of vortices pinned along the interface[5]. Specifically, phase variations $\delta\varphi$ near the interface transform into observable spatial density modulations, $\rho_\uparrow = \bar{\rho}\cos^2(\delta\varphi/2)$, which are subsequently recorded by absorption

imaging. Figure 5.6(C) displays typical interference patterns captured at two hold times during the RTI evolution [for $|\mathcal{F}|/h = 4.6(3)\,\mathrm{Hz\,\mu m^{-1}}$].

The total vortex count $\mathcal{N}$ measures the net 2π accumulation of $\delta\varphi$ along the interface (of length ℓ). This value directly yields the average magnitude of the counterflow velocity:

$$\overline{|\delta v|} \approx h\mathcal{N}/(m\ell) \tag{5.16}$$

where the primary uncertainties stem from the limited vortex density and the ambiguous regions where δv switches its direction. Figure 5.6(D) presents the result: at short times (linear dynamics), $\mathcal{N}$ increases as counterflowing currents develop along the interface. At late times (non-linear and turbulent dynamics), $\overline{|\delta v|}$ saturates, while increasingly intricate velocity fields precipitate a vortex proliferation that reveals the flow patterns associated with interface disintegration.

5.6 Conclusion

In this chapter, we have systematically investigated the stable and unstable aspects of the interfacial dispersion relation in a binary quantum fluid. The stable interfacial modes correspond to ripplons, whose spectrum we mapped by tracking their free evolution after parametric excitation. Conversely, the unstable modes give rise to the Rayleigh-Taylor instability, which we triggered by a magnetic field gradient inversion. We observed its characteristic mushroom- and spike-like features, and quantified its early-stage progression via its growth rate. Taken together, these results mark the first experimental observations of both ripplons and the RTI in ultracold atomic systems.

Beyond the excellent agreement with theory, our observation of the RTI expands the growing family of fluid-dynamical instabilities realized in ultracold gases. Apart from RTI, related instabilities such as the Kelvin-Helmholtz instability have also been attributed to the development of conventional vortex chains seen in single-component [106, 107] as well as more exotic spin textures in multi-component systems [34].

Because the RTI is often detrimental in applied contexts, controlling its growth remains a key question. The same parametric scheme for exciting ripplon modes could potentially be inverted to stabilize the RTI [108]—a critical goal in inertial confinement fusion [35], reminiscent of recent realizations at classical fluid-gas interfaces [109].

[5] Slightly away from the interface, counterflowing currents carry opposite bulk velocities in the two components. The $\pi/2$ pulse puts the condensates into an equal superposition of spin-↑ and spin-↓, with each spin component inheriting the opposing local velocity on either side of the interface. This rapid velocity flip manifests as a phase gradient inversion across the interface, producing an array of vortices centered at the destructive interference density minima along the interface.

Finally, the long, well-aligned interfaces demonstrated in this work also point beyond the instability itself: they offer a promising platform for the analogue simulation of relativistic field theories [110, 111]. Furthermore, the thermal population of ripplons, quantified by the initial fluctuations on the interface, offers a route to study the non-equilibrium thermodynamics of BECs—a foundation we build upon in the next chapter to explore prethermalization of ripplons in a binary condensate.

6 Study of Prethermal Ripplons on a One-Dimensional Quantum Fluid Interface

In this chapter, I present the study of the post-quench prethermal state of ripplons on a 1D quantum fluid interface [112], marking the second major scientific contribution of this thesis.

How many-body systems relax toward thermal equilibrium remains an open question spanning both classical and quantum physics. Prethermalization, a fast decay into a long-lived quasi-steady state followed by slow relaxation to actual thermal equilibrium, often occurs in a wide range of systems [113]: from astrophysical processes [114, 115] to classical Fermi-Pasta-Ulam-Tsingou (FPUT) chains [116–118]. Ultracold quantum gases with limited degrees of freedom serve as an ideal platform to probe this physics; experiments on 1D Bose gases [119–122], trapped-ion spin chains [123], and Floquet lattices [124] have observed prethermal behavior. However, prethermalization in these settings is tied to their near-integrability and accompanying approximate conservation laws.

At its root, prethermalization stems from a separation of timescales: rapid initial system reconfiguration establishes a quasi-steady state followed by a slow relaxation process. This raises a natural question: can prethermalization exist in complex systems without clear integrability? Here, we demonstrate that prethermalization also emerges at a 1D quantum interface—ripplon modes on the spin domain wall embedded in a binary condensate [33]. Rather than being isolated, this interface is in thermal contact with the bulk phonon modes, which act as a well-defined thermal bath.

Following a quench, the system settles into a long-lived state with persistent non-equipartition of energy in long-wavelength ripplon modes, surviving far beyond the intrinsic thermalization times of the bath. Remarkably, while each ripplon mode independently arrives at a thermal state, global thermal equilibrium across modes is not established. We trace this prethermalization to kinematic isolation: an energy mismatch between the long-wavelength ripplons and the phonon bath hindering energy exchange [86, 118, 125, 126]. Thus, prethermalization in our interface system arises from kinematic constraints rather

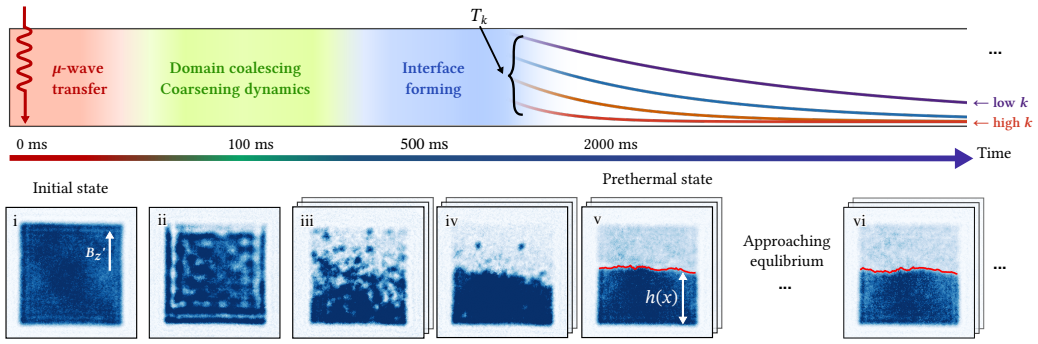

Figure 6.1: The long road to equilibrium (reproduced from Ref. [112]). A schematic illustration of the relaxation process (top) is shown along with representative spin-selective images of the $|\uparrow\rangle$ density (bottom). (i) A microwave pulse at $t = 0$ transfers an initially spin-polarized BEC into an equal superposition of $|\uparrow\rangle$ (dark blue) and $|\downarrow\rangle$. The orientation of a magnetic field gradient is indicated. (ii) Initial spin structure forms nearly deterministically; (iii) small initial fluctuations then seed random droplet patterns that coalesce and coarsen. (iv) After about 500 ms an interface forms and (v) becomes single-valued by 2000 ms (red curves). Ripplons with different wavenumber k combine to give the observed interface profiles; their occupation numbers N_k (solid curves in the top panel) relax with drastically different timescales to their respective equilibrium values (dashed lines).

than integrability—a mechanism analogous to slow thermalization in classical FPUT systems [118].

Our study is enabled by a novel thermometry technique we developed for binary condensates. Unlike immiscible spin condensate components that separate in space, the thermal components do not—they possess high enough energy to overcome the inter-spin energy barrier, permeating both spin domains. Combining this natural spatial separation and our spin-selective imaging allows precise measurement of the thermal density, paving the way for the extraction of bulk temperature. We exploit this capability to compare the temperature of short-wavelength ripplons and that of the bulk phonons. Thermometry in ultracold gases is notoriously challenging [55, 127]; by overcoming this long-standing obstacle, our approach offers a powerful, robust diagnostic tool broadly applicable to multicomponent quantum gases.

The material for this chapter is adapted from our work published in Ref. [112].

6.1 Fluctuations on a quenched quantum interface

Our experiment begins with a ^{23}Na BEC of $\sim 1.6\times10^6$ atoms in the $|F = 1, m_F = -1\rangle$ ground state, prepared in a quasi-2D optical box potential (Sec. 3.4). At $t = 0$, a microwave $\pi/2$ pulse quenches the system into an equal superposition of $|\downarrow\rangle = |F = 1, m_F = -1\rangle$ and

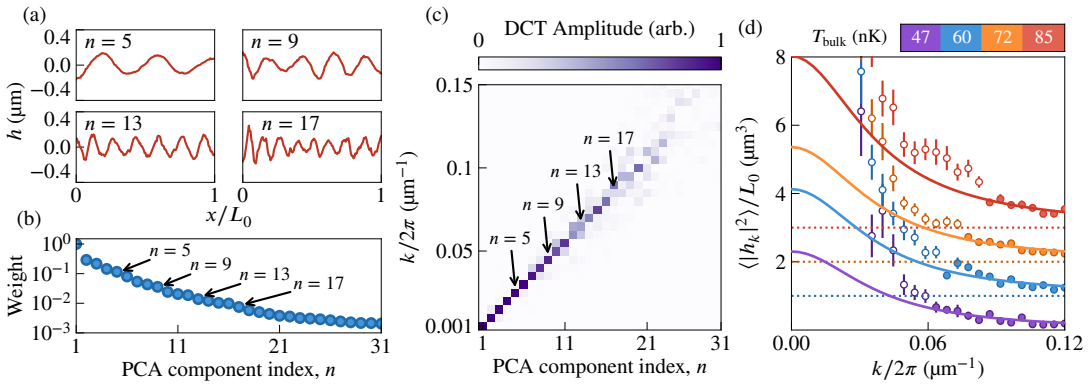

Figure 6.2: Fluctuations of a prethermal interface, with $t_{\mathrm{hold}} = 2.0\,\mathrm{s}$ (reproduced from Ref. [112]). (a) PCA components of $h(x)$ for $n = 5$, 9, 13, and 17. (b) Normalized PCA weights. (c) DCT amplitudes of each PCA component. (d) Power spectrum $\langle|h_k|^2\rangle$, measured at four bulk temperatures (vertically offset for clarity) using ≈ 70 experimental realizations per temperature, plotted along with the corresponding equilibrium prediction (no free parameters).

$|\uparrow\rangle = |F = 2, m_F = -2\rangle$ states. Driven by strong inter-spin repulsive interactions, the immiscible mixture undergoes rapid domain formation and coarsening [128], as shown in Fig. 6.1. By $t \approx 500\,\mathrm{ms}$, two spin domains form on opposite sides of the box potential, creating a one-dimensional interface along the x-direction [33] stabilized by a small magnetic field gradient $B'_z = 2.9(1) \times 10^{-4}\,\mathrm{T\,m^{-1}}$ along y. As the interface forms, we continue evaporative cooling by slowly ramping down the box trap depth to U_D, which controls the final temperature of the condensate bulk. Following a variable hold time t_{hold}, we selectively image the $|\uparrow\rangle$ density *in situ* by absorption imaging and extract the interface height profile $h(x)$.

To uncover the nature of interfacial fluctuations without *a priori* physical assumptions, we perform principal component analysis (PCA) [129–131] on ensembles of measured $h(x)$. As a model-agnostic approach, PCA yields a set of statistically uncorrelated, orthogonal fluctuating components, together with their corresponding explained variance. We order and index these extracted PCA components by an integer $n \geq 1$ in descending order of their explained variance (weights) across the dataset.

Representative PCA components and their corresponding weights are displayed in Fig. 6.2(a) and (b). These PCA components closely mirror the cosinusoidal standing-wave ripplon modes dictated by the imposed boundary conditions (see Sec. 5.2). The discrete cosine transform (DCT) of each component [Fig. 6.2(c)] reveals a nearly one-to-one mapping between the PCA index n and the dominant wavenumber k. Those observations strongly confirm the ripplon nature of those fluctuating interface modes, laying the foundation for our power spectral density (PSD) analysis in Section 6.3. Beyond $k/2\pi \approx 0.10\,\mu\mathrm{m}^{-1}$,

imaging aberrations attenuate fluctuation amplitudes, degrading the signal-to-noise ratio and inducing mode mixing across high-k PCA components. To restore true mode amplitudes, we characterize and compensate for these imaging aberrations with the approaches outlined in Sec. 4.2.

6.2 Theory of thermal ripplons

Here, we construct a phenomenological model for the fluctuating ripplons to quantify the fluctuations across wavenumbers k, starting with an evaluation of the potential energy of the interface system.

The total interfacial potential energy H_{potn} comprises contributions from the interfacial tension H_{intf} and the magnetic gradient force H_{mag}. Expanding interface length to the second order in the fluctuating interface height $h(x)$, the surface tension energy reads:

$$H_{\text{intf}} = \sigma \int_0^{L_0} \mathrm{d}x \sqrt{1 + h'(x)^2}, \tag{6.1}$$

$$\approx \sigma L_0 + \frac{\sigma}{2} \int_0^{L_0} \mathrm{d}x\, h'(x)^2. \tag{6.2}$$

Here, σ is the interfacial tension [31, 33], assumed to be well-defined for the measured k range, $h'(x) \equiv \partial_x h(x)$, and L_0 is the box dimension. The magnetic gradient energy H_{mag} represents the work done against the gradient when displacing the interface by $h(x)$. Integrating along the interface yields:

$$H_{\text{mag}} = \int_0^{L_0} \mathrm{d}x \int_0^{h(x)} \mathrm{d}y\, \bar{n}\mathcal{F} y \tag{6.3}$$

$$= \frac{1}{2}\bar{n}\mathcal{F} \int_0^{L_0} \mathrm{d}x\, h(x)^2 \tag{6.4}$$

where $\mathcal{F}$ is the differential force per atom [Eq. (5.8)], and $\bar{n}\mathcal{F}$ is the force per unit area.

Using the Fourier transform $h(x) = (1/L_0)\sum_k h_k \exp(ikx)$ and $h_{-k} = h_k^*$, the total potential energy (omitting the constant offset σL_0) is:

$$H_{\text{potn}} = \frac{1}{2}\int_0^{L_0} \mathrm{d}x\left[\sigma h'(x)^2 + \bar{n}\mathcal{F}h(x)^2\right] \tag{6.5}$$

$$= \frac{1}{2L_0^2}\sum_{k,k'}\int_0^{L_0} \mathrm{d}x[-\sigma k k' + \bar{n}\mathcal{F}]h_k h_{k'}\, e^{i(k+k')x} \tag{6.6}$$

$$= \frac{1}{2L_0}\sum_k (\sigma k^2 + \bar{n}\mathcal{F})|h_k|^2. \tag{6.7}$$

where we used mode orthogonality $(1/L_0)\int_0^{L_0}\mathrm{d}x\, \exp[i(k+k')x] = \delta_{k,-k'}$ and $h_{-k} = h_k^*$. We see that H_{potn} completely decomposes into independent quadratic modes h_k. In the semiclassical limit ($\hbar\omega_k/k_\mathrm{B}T \to 0$, where $\hbar\omega_k$ is the quantized ripplon energy), the classical equipartition theorem [1]dictates the distribution of h_k and its second moment $\langle|h_k|^2\rangle$ [132, 133]:

$$P(h_k) \propto \exp\left[-\frac{\sigma(k^2+\xi^{-2})|h_k|^2}{2L_0 k_\mathrm{B}T}\right], \tag{6.8}$$

$$\frac{\langle|h_k|^2\rangle}{L_0} = \frac{k_\mathrm{B}T}{\sigma(k^2+\xi^{-2})}, \tag{6.9}$$

where T is the interface temperature and $\xi^2 \equiv \sigma/(\bar{n}\mathcal{F})$ is the capillary length [133]. Crucially, $\langle|h_k|^2\rangle$ corresponds directly to the ensemble-averaged power spectral density (PSD) extracted from experimentally measured $h(x)$ across wavenumbers k. Eq. (6.9) provides a solid theoretical basis for interpreting measured $\langle|h_k|^2\rangle$ and extracting the mode temperature T_k in Sec. 6.3.

While Eq. (6.9) relies on the phenomenological potential energy Eq. (6.2) and the semiclassical limit ($\hbar\omega_k/k_\mathrm{B}T \to 0$), it holds well across our experimental regime, as validated by the full quantum treatment in Sec. 6.6.

6.3 Quantifying prethermal ripplon modes

The theoretical PSD of interfacial fluctuations $\langle|h_k|^2\rangle$ [Eq. (6.9)] decreases monotonically with k, agreeing qualitatively with the mapping between the dominant k and the decreasing PCA weights in Fig. 6.2(b) and (c). Figure 6.2(d) compares the measured $\langle|h_k|^2\rangle$ aver-

[1]The equipartition theorem requires only that the kinetic energy be quadratic in the mode momenta (the kinetic terms) and independent of h_k (satisfied near equilibrium). The kinetic contribution does not enter $P(h_k)$ or $\langle|h_k|^2\rangle$.

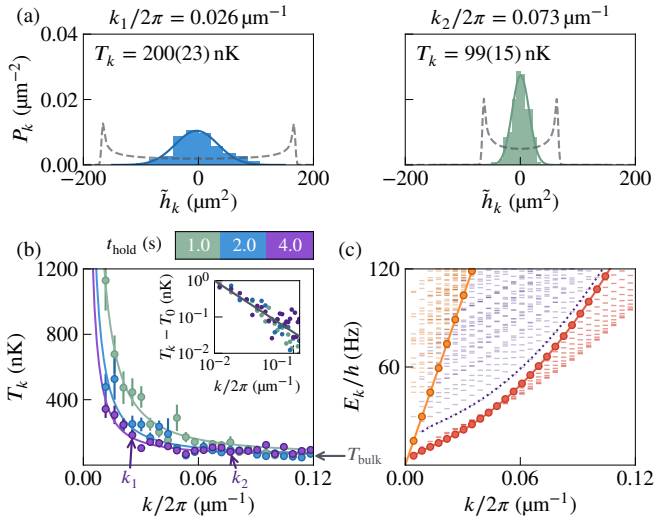

Figure 6.3: Distribution functions and equilibration at $T_{\text{bulk}} = 56\,\text{nK}$ (reproduced from Ref. [112]). (a) PDFs P_k for $k_1/(2\pi) = 0.026\,\mu\text{m}^{-1}$ and $k_2/(2\pi) = 0.073\,\mu\text{m}^{-1}$ at $t_{\text{hold}} = 4\,\text{s}$. The histograms plot experimental data with corresponding normal distributions (solid curves). Dashed curves show the PDF corresponding to a phase-averaged ripplon coherent state with the same ripplon occupation. (b) Mode temperature as a function of k for three hold times. Solid curves show power-law fits with a common exponent and offset. Inset: log-log plot in terms of dimensionless quantities $\bar{T}_k = (T_k - T_0)/T_0$ and $\bar{k} = k/k_0$ showing data collapse after dividing the fitted amplitude for each hold time. (c) Scattering processes in our $L_0 \times L_0$ box. Circles connected by curves plot the energy of ripplons (red) and phonons (yellow) as a function of k. Ticks indicate the total energy of momentum-conserving 3-field processes starting from a momentum-k ripplon and ending in a pair of phonons (yellow), a phonon and a ripplon (purple), or two ripplons (red). The dotted curve traces out the lowest-energy phonon-producing channel for each k.

aged over $\sim$ 70 experimental runs (markers) against the parameter-free thermal predictions [Eq. (6.9)] evaluated at the independently measured bulk temperatures T_{bulk} (solid curve). The PSD measured across all temperatures shows a consistent trend: modes at high-k coincide with the thermal equilibrium prediction, whereas low-k modes exhibit a pronounced excess population far exceeding thermal expectation.

To probe the statistical nature of these fluctuations, we extract the full probability density function (PDF) $P_k \equiv P(h_k)$[2] from an ensemble of $\sim$ 70 experimental realizations. Figure 6.3(a) displays P_k for two representative modes with low- and high-k at $T_{\text{bulk}} = 56\,\text{nK}$, respectively. The increased width of P_k at low-k agrees with the observed elevated $\langle |h_k|^2 \rangle$,

[2] As observed in Sec. 5.2, the standing wave ripplon modes in our box potential follow cosinusoidal profiles $\sim \cos(kx)$ with $k = j\pi/L_0$ for $j = 1, 2, \ldots$. This justifies using the discrete cosine transform (DCT), equivalent to Fourier transforming a mirror-extended profile on $[-L_0, L_0]$ with normalization $h_k = \sqrt{2}\int_0^{L_0} \mathrm{d}x\, h(x)\cos(kx)$. Under this convention, each standing wave retains the same energy form of Eq. (6.7), and Eq. (6.9) holds with $\langle |h_k|^2 \rangle = \langle h_k^2 \rangle$.

which corresponds directly to the second moment of P_k. Across all wavenumbers, P_k is statistically indistinguishable from a normal distribution [solid curves in Fig. 6.3(a)]; higher-order moments $\langle h_k^m \rangle$ match the Gaussian expectation within uncertainties up to $m = 6$, before the expectation values are overwhelmed by the experimental error. Because Gaussian statistics are expected for thermal states with a quadratic Hamiltonian [Eq. (6.7), see Sec. 6.6], we can unambiguously assign an effective mode temperature T_k to each individual ripplon mode [Fig. 6.3(b)].

Figure 6.3(b) tracks T_k across hold times t_{hold} from 1 s to 4 s, shortly before technical atom loss becomes significant. The extracted $T_k \equiv \sigma(k^2 + \xi^{-2})\langle |h_k|^2 \rangle / (L_0 k_{\mathrm{B}})$ further demonstrates that only high-k modes equilibrate with the bulk bath, while low-k modes remain strongly out-of-equilibrium. At low-k, the excessive temperature follows a power-law $(T_k - T_{\mathrm{bulk}})/T_0 \propto (k/k_0)^{-\alpha}$, with the exponent $\alpha = 1.42(5)$, and $T_0 \equiv T_{\mathrm{bulk}}$ and $k_0/2\pi = 0.01\,\mu\mathrm{m}^{-1}$ define the reference scales. The long decay lifetime 5(1) s of long-wavelength ripplons—much slower than the 1×10^{-1} s equilibration time of the phonon bath—strongly highlights the prethermal nature of the interface system.

The large excess ripplon population raises the intriguing possibility of a transient ripplon condensate [38, 127]: although ripplon numbers are not strictly conserved, slow relaxation could enforce quasi-conservation on intermediate timescales, analogous to photon condensates (Sec. 2.1). Such a condensate would manifest as either a number state or a coherent ripplon state with a spontaneously broken temporal oscillation phase differing across experiments. In the regime of large ripplon occupation, both condensed states yield an identical PDF [Fig. 6.3(a)]. Our experimentally observed Gaussian PDFs explicitly rule out this possibility.

6.4 Proposed prethermal mechanism: kinematic isolation

We attribute the prethermalization in our system to the kinematic isolation between the long-wavelength ripplons and the bulk phonon bath. Since the fluctuation amplitudes remain small at the interface, the system is described by Bogoliubov–de Gennes (BdG) excitations (Sec. 2.6): the interfacial ripplons and the bulk phonons. Figure 6.3(c) plots the discretized mode energies for both ripplons (red) and phonons (yellow) in the box potential as solid circles. The substantial energy mismatch between these two branches strongly suppresses energy exchange beyond intrinsically weak non-linear interactions [134–136]: at lowest order, the ripplon-ripplon coupling and ripplon-to-phonon decay are three-quasiparticle scattering processes—the scattering between a high-energy quasiparticle and a pair of low-energy quasiparticles or its inverse. This leads to Beliaev damping [86, 125, 126].

The tick markers in Fig. 6.3(c) represent the total energy of a quasiparticle pair produced by ripplon scattering channels that conserve longitudinal momentum k. Marker colors distinguish the quasiparticle species in the pair: phonon-phonon (yellow), phonon-ripplon (purple), and ripplon-ripplon (red). On-shell scatterings—processes conserving both energy and momentum at tree level—are only possible where tick markers and circles coincide. Crucially, at low-k, where prethermal ripplon populations are observed, all phonon-producing channels are separated from the ripplon branch by a gap of $\approx 20\,\mathrm{Hz}$. This gap precludes the tree-level decay of low-k ripplons into phonons, which we propose as the primary origin of their kinematic isolation; at high-k, this energy gap closes, naturally accounting for the rapid thermalization observed in short-wavelength ripplon modes. Similar mechanisms have also been proposed for FPUT systems [118]. While thermally assisted Landau processes, in which a ripplon is absorbed by a thermal phonon, remain kinematically possible, the measured long lifetime demonstrates that their effect is strongly suppressed; a full higher-order calculation remains a subject for future theoretical work.

6.5 Measuring the bulk temperature of the condensate

Quantifying prethermalization at the interface requires an independent measurement of the bulk temperature to establish whether the interface is in thermal equilibrium with the bulk phonon bath. In typical condensate systems, temperature measurement is notoriously challenging because the minuscule thermal excitations are largely eclipsed by the dense condensate background.

We circumvent this issue by probing the thermal excitations of minority spin-↑ atoms "floating" atop the majority spin-↓ condensate. Together with our spin-selective imaging, these thermal populations appear as a faint box-shaped "shadow" visible in Fig. 6.1(v) and (vi) that can be reliably quantified.

The density of these thermal spin-↑ atoms is given by

$$n_{T,\uparrow} = g_1\left[\exp\left[\mu_{\mathrm{loc},\uparrow}(\mathbf{r})/k_{\mathrm{B}}T\right]\right]/\lambda_T^2, \tag{6.10}$$

where $\mu_{\mathrm{loc},\uparrow}(\mathbf{r}) = \mu_\uparrow - V_{\mathrm{eff},\uparrow}(\mathbf{r})$ is the local chemical potential, $\mu_\uparrow$ is the spin-↑ chemical potential, $\lambda_T \equiv (2\pi\hbar^2/mk_{\mathrm{B}}T)^{1/2}$ is the thermal de Broglie wavelength, and g_n are Bose functions.

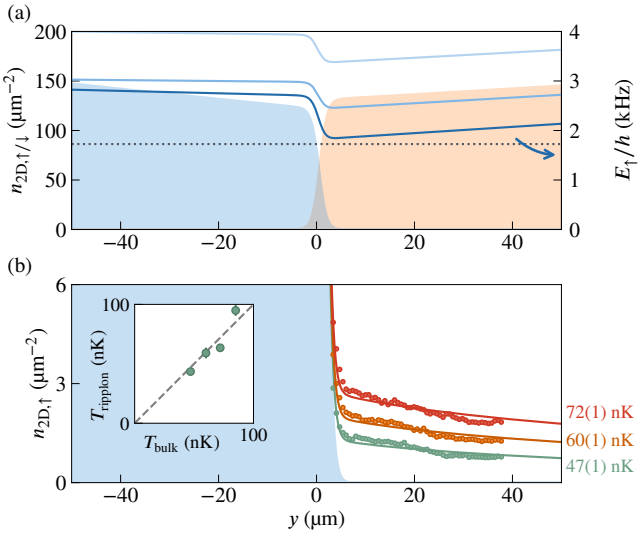

Figure 6.4: Minority thermal atom thermometry (reproduced from Ref. [112]). (a) Effective potentials for thermal spin-↑ atoms in the lowest three local transverse modes (solid curves), compared with the chemical potential $\mu_\uparrow$ (dotted line). The shaded curves plot the condensate densities $n_{\uparrow/\downarrow}$ in blue/yellow respectively. (b) Spin-↑ density profiles. Markers plot the measured total density (thermal plus condensed) along with fits (solid curve) to our theoretical model, yielding the annotated temperatures. The shaded region shows the computed condensate density and highlights its vast difference in scale. Inset: T_{bulk} versus T_{ripplon} obtained by fitting the high-k ripplon tail in Fig. 6.2(d) using Eq. (6.9) with T_{ripplon} as the only free parameter; the dashed line indicates equality.

Neglecting the interactions within the low-density thermal fraction, the effective potential experienced by thermal spin-↑ atoms is:

$$V_{\text{eff},\uparrow}(\mathbf{r}) = V_{\text{ext},\uparrow}(\mathbf{r}) + 2g_{\uparrow\uparrow}n_\uparrow(\mathbf{r}) + g_{\uparrow\downarrow}n_\downarrow(\mathbf{r}), \tag{6.11}$$

where the external potential $V_{\text{ext},\uparrow}$ includes the magnetic gradient, the second and the third terms represent the interaction energies between same- and opposite-spins with densities $n_\uparrow$ and $n_\downarrow$, respectively. The factor of two in the same-spin term arises from the Hartree-Fock exchange contribution.

Inside the spin-↓ domain, where $n_\uparrow = 0$, the inter-spin interaction sets up a barrier that thermal spin-↑ atoms must overcome. For a homogeneous bulk, $-\mu_{\text{loc},\uparrow} = (g_{\uparrow\downarrow}/\sqrt{g_{\uparrow\uparrow}g_{\downarrow\downarrow}} - 1)\,\mu_\uparrow \approx 0.09\,\mu_\uparrow$ (∼ a few nK). Because the mixture is weakly immiscible [Eq. (2.75)], thermal atoms (with $k_{\text{B}}T > 40$ nK in our experiment) readily permeate the opposite-spin domain while the condensates remain separated.

In practice, the magnetic gradient and the $n_\downarrow$ condensate slope induce an effective energy gradient of $\sim 6\,\text{Hz}\,\mu\text{m}^{-1}$, causing significant spatial deformation of $n_{\uparrow,\downarrow}$ across the box, as

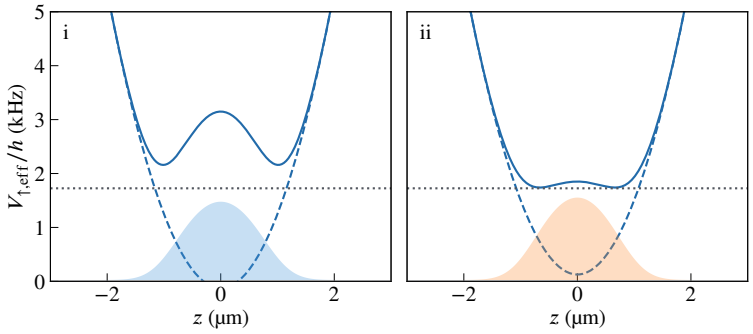

Figure 6.5: Effective-potential slices used to calculate n_T (reproduced from Ref. [112]). Panels (i) and (ii) show $V_{\mathrm{eff},\uparrow}$ (solid curves) along z at $y = -31\,\mu\mathrm{m}$ (same-spin condensate region) and $y = 31\,\mu\mathrm{m}$ (opposite-spin condensate region), respectively. Dashed curves show the bare external potential, the dotted line marks the chemical potential $\mu_\uparrow$, and the shaded curves plot the condensate densities $n_\uparrow$ (blue) and $n_\downarrow$ (orange) in arbitrary units.

shown in Fig. 6.4(a). Moreover, the bulk temperatures in our experiments range from 50 nK to 90 nK, comparable to the energy spacing of the vertical confining potential $\hbar\omega_z/k_B \approx 50\,\mathrm{nK}$ (see Sec. 3.4), necessitating a full quasi-2D treatment. Assuming $V_{\mathrm{eff},\uparrow}$ varies slowly in the y-direction, we solve a 1D single-particle Schödinger equation (with the ∂_y^2 term dropped) on a z slice at each y location (Fig. 6.5):

$$\left[-\frac{\hbar^2}{2m}\partial_z^2 + V_{\mathrm{eff},\uparrow}(y,z)\right]\phi_n(z;y) = \epsilon_n(y)\,\phi_n(z;y). \tag{6.12}$$

Numerical diagonalization yields the energy levels $\epsilon_n(y)$ (light curves in Fig. 6.4) and transverse wavefunctions $\phi_n(z;y)$. Integrating over the z-coordinate and summing over the transverse states yields the 2D projected thermal density:

$$n_{T,\uparrow}^{\mathrm{2D}}(y) = \sum_n \int \mathrm{d}z|\phi_n(z;y)|^2 \frac{1}{\lambda_T^2}\, g_1\Big(e^{[\mu_\uparrow - \epsilon_n(y)]/k_\mathrm{B}T}\Big), \tag{6.13}$$

$$= \frac{1}{\lambda_T^2}\sum_n g_1\Big(e^{[\mu_\uparrow - \epsilon_n(y)]/k_\mathrm{B}T}\Big). \tag{6.14}$$

Fitting $n_{T,\uparrow}^{\mathrm{2D}}(y)$ to experimental images except near the box edges (where condensate densities vary too quickly) yields the bulk temperature T_{bulk}, as shown in Fig. 6.4(b). As an independent validation, the inset of Fig. 6.4(b) compares this T_{bulk} against T_{ripplon} obtained by fitting Eq. (6.9) to the high-k tail of $\langle|h_k|^2\rangle$ [Fig. 6.2(d)]; the two independent temperature measures agree within experimental error.

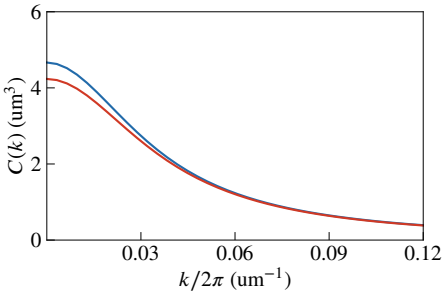

Figure 6.6: Classical and quantum predictions for $C(k) \equiv \langle |h_k|^2 \rangle / L_0$ (reproduced from Ref. [112]). Red and blue curves show the classical prediction [Eq. (6.9)] and quantum prediction [Eq. (6.16)] at 80 nK, a typical experimental temperature. Their small residual difference stems from the BdG-computed versus phenomenological interface parameters.

6.6 Quantum statistics of thermal ripplons

Having derived Eq. (6.9) from energy equipartition in Section 6.2, we now provide a microscopic quantum treatment, which confirms that this semiclassical result holds to high accuracy across our experimental parameter regime.

We numerically solve for the ripplon modes under the BdG framework (see Sec. 2.6). As the Goldstone modes arising from the translational symmetry broken by the interface (Sec. 2.6), ripplons have a spatial profile closely tracking the derivative of the condensate wavefunctions $\partial_y \psi_0$. In our box, the ripplons are standing waves (Sec. 6.3). Let $\hat{b}_{0,k}$ and $\hat{b}^\dagger_{0,k}$ be annihilation and creation operators for a ripplon at wavenumber k with energy $\hbar\omega_k$. The second-quantized ripplon Hamiltonian is:

$$\hat{H} = \hbar \sum_k \omega_k \hat{b}^\dagger_{0,k} \hat{b}_{0,k}. \tag{6.15}$$

The interface height operator is constructed as $\hat{h}(x) = \sqrt{2} \sum_k \eta_k (\hat{b}_{0,k} + \hat{b}^\dagger_{0,k}) \cos(kx)$, with its cosine transform $\hat{h}_k = L_0 \eta_k (\hat{b}_{0,k} + \hat{b}^\dagger_{0,k})$. Here, η_k denotes the interface displacement under a single ripplon quantum, which is evaluated from the spatial overlap between the ripplon mode function $u_{0,k} - v_{0,k}$ and $\partial_y \psi_0$, leading to the power spectral density $\langle |h_k|^2 \rangle = \eta_k^2 L_0^2 (2n_{0,k} + 1)$, where $n_{0,k}$ is the ripplon occupation number at k.

At thermal equilibrium, the ripplon occupation follows the Planck distribution $n_{0,k} = 1/[\exp(\hbar\omega_k / k_B T) - 1]$, yielding:

$$\langle |\hat{h}_k|^2 \rangle = \eta_k^2 L_0^2 \coth\left(\frac{\hbar\omega_k}{2k_B T}\right). \tag{6.16}$$

Figure 6.6 compares $\langle|h_k|^2\rangle$ computed from Eq. (6.9) and Eq. (6.16).

Furthermore, for the ripplon mode at a fixed k, the thermal density matrix is:

$$\hat{\rho}_k = (1/Z)\sum_m \exp(-m\hbar\omega_k/k_\mathrm{B}T)|m\rangle\langle m|. \tag{6.17}$$

Projecting $\hat{\rho}_k$ onto the h_k basis (analogous to the displacement basis of a harmonic oscillator) yields the probability distribution $P(h_k) = \langle h_k|\hat{\rho}_k|h_k\rangle$, where $|m\rangle$ is the number basis of the corresponding ripplon mode. $P(h_k)$ is evaluated by:

$$P(h_k) = \frac{1}{Z}\sum_m e^{-m\hbar\omega_k/k_\mathrm{B}T}|\langle h_k|m\rangle|^2 \tag{6.18}$$

$$= \frac{1}{Z}\sum_m e^{-m\hbar\omega_k/k_\mathrm{B}T}\frac{1}{2^m m!}\frac{1}{\sqrt{2\pi}\eta_k L_0}\exp\left(-\frac{h_k^2}{2\eta_k^2 L_0^2}\right)H_m^2\left(\frac{h_k}{\sqrt{2}\eta_k L_0}\right), \tag{6.19}$$

Using Mehler's formula [137, (18.18.28)], this becomes:

$$P(h_k) \propto \exp\left[-\tanh\left(\frac{\hbar\omega_k}{2k_\mathrm{B}T}\right)\frac{1}{2\eta_k^2 L_0^2}h_k^2\right] \tag{6.20}$$

$$\to \exp\left(-\frac{\hbar\omega_k}{4k_\mathrm{B}T\eta_k^2 L_0^2}h_k^2\right) \quad \text{as } \frac{\hbar\omega_k}{k_\mathrm{B}T} \to 0.$$

Therefore, the quantum $P(h_k)$ distribution is strictly Gaussian at all temperatures. In the high-temperature limit $\hbar\omega_k/k_\mathrm{B}T \to 0$, its second moment $\langle|h_k|^2\rangle$ maps to the classical form if the relation $\eta_k^2 = \hbar\omega_k/[2L_0\sigma(k^2+\xi^{-2})]$ holds. We compare the classical distribution with the quantum result Eq. (6.20) in Fig. 6.7. Across our experimental parameters, quantum corrections to $\langle|h_k|^2\rangle$ are below 0.2 %, rigorously justifying the classical equipartition framework employed throughout this chapter.

6.7 Conclusion

In this chapter, we have investigated the post-quench prethermalization of ripplons at a 1D quantum fluid interface embedded in a binary condensate. Spectral analysis of the interfacial fluctuations reveals a striking bifurcation of thermalization behavior: high-k ripplons rapidly equilibrate with the bulk phonon bath, whereas low-k ripplons remain pinned at elevated temperatures for seconds—a long-lived prethermal state. Across both regimes, the mode amplitudes follow a thermal distribution, explicitly ruling out a ripplon condensate. We attribute the prethermal state to kinematic isolation: because no energy- and

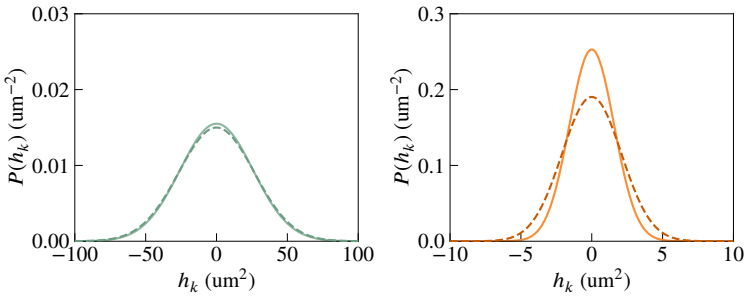

Figure 6.7: Classical and quantum probability distributions for h_k at $k/(2\pi) = 0.026\,\mu\text{m}^{-1}$ [k_1 in Fig. 6.3(a)] (reproduced from Ref. [112]). Solid and dashed curves show the classical prediction and Eq. (6.20), respectively, at 80 nK (green, typical experimental temperature) and 0.3 nK (orange). Only at the lower temperature do quantum (zero-point) fluctuations visibly broaden the distribution.

momentum-conserving three-field process allows the low-k ripplons to decay directly into the phonon bath, energy remains trapped within those long-wavelength modes, shielding them from rapid equilibration.

Whereas prethermalization in 1D Bose gases [120, 121] is protected by near-integrability, the prethermal state observed here arises from kinematic constraints in a system coupled to a well-defined thermal reservoir—closer in spirit to slow thermalization in FPUT chains [118]. Nevertheless, the observed state shares a key feature with the generalized Gibbs ensemble [113, 121]: each mode exhibits thermal statistics with its own mode temperature T_k, as if the low-k ripplon occupations were quasi-conserved—a direct consequence enforced by the proposed kinematic isolation mechanism. Several open questions naturally emerge: what microscopic pathways govern the fast initial scrambling into this prethermal state, and what processes drive its ultimate relaxation to global equilibrium? Within the kinematic picture, a natural candidate is thermally assisted three-field scattering: a low-k ripplon absorbed by a thermal phonon. Its contribution is more nuanced and remains a subject for future theoretical work. Furthermore, by bringing the system to even lower temperatures where quantum fluctuations dominate, the same mode-resolved analysis could reveal the quantum statistics of ripplons discussed in Sec. 6.6. Finally, the utility of our robust minority-component thermometry extends far beyond quantum interfaces. By opening a new path toward ultracold temperature measurement, this methodology promises to become a standard diagnostic across the broader family of multi-component quantum fluids.

7 Conclusion and Outlook

In this thesis, starting from the historical context, I outlined the basic condensate theory, followed by the experimental aspects including the design, optimization, and characterization of the apparatus, and culminating in the two scientific projects focused on the non-equilibrium properties of immiscible binary condensates. First, we made the initial observations of ripplons and the Rayleigh-Taylor instability (RTI) in an ultracold gas: we measured the ripplon dispersion relation, extracted the interfacial tension, and quantified the RTI growth (Chapter 5). Second, we studied the prethermalization of the 1D quantum fluid interface following a quench, attributed to kinematic isolation from the phonon bath (Chapter 6).

The immiscible binary condensate offers a versatile platform for rich physics, ranging from the domain-wall coarsening [128] to false vacuum decay analogies [138]. In retrospect, our hydrodynamics-themed research particularly benefited from the large condensate atom number and strong inter-spin interactions. Our microscopic imaging system also enabled us to capture the density pattern of the condensate, despite its limited imaging resolution.

The findings of this thesis open several avenues for future research. First, the parametric driving scheme used to excite ripplons is predicted to stabilize the RTI [108], directly analogous to Kapitza stabilization at classical fluid interfaces [109]. Furthermore, the thermal fluctuations on the interface could be amplified by RTI—conceptually similar to classical squeezing—positioning the RTI as a potential thermometer for the interface systems. These future directions critically rely on a well-prepared initial state whose fluctuations are well understood. Ideally, preparing the interface with an engineered vector light-shift potential that acts like a spatially resolved synthetic magnetic field could bypass the violent quench, which has been previously demonstrated in ^{87}Rb systems [139, 140].

Moreover, transitioning to atomic species (for example, ^{39}K or ^{7}Li) with accessible Feshbach resonances [141] would provide a tunable control knob over inter- and intra-spin interaction strengths. This capability would not only benefit the studies of spinor dynamics, but also unlock opportunities for probing novel quantum effects in strongly interacting regimes.

Appendix

A Proper Definition of Interfacial Tension

As explained in Sec. 5.4, there was a moment of confusion about the definition of interfacial tension σ in binary condensates around the turn of the century.

To recount the details: in 1998, Ao and Chui produced the first calculation of the interfacial tension [30]. They derived σ from the energy excess due to the interface—the difference of the total system energy ("with interface") and the Thomas-Fermi system energy ("without interface"), which omits the kinetic term contribution. Ten years later, Van Schaeybroeck published a more in-depth analysis [31], in which he explicitly defined σ using the excess in "grand potential" instead of in energy, resulting in a value that is four times larger than what Ao and Chui initially obtained.

The interfacial tension, though not directly measurable by itself, connects the pressure difference and interface curvature through the Young-Laplace equation; it also enters the ripplon dynamics through the dispersion relation [Eq. (5.5)], rendering this question of definition not merely a semantic one. On ripplon dynamics, both Mazets [32] and Barankov [142] derived the ripplon dispersion in binary condensates. The former gave the correct dispersion relation but had no link to the interfacial tension, whereas the latter used the same expression from Ref. [30]. Later, both numerical studies [87, 92] and our experimental measurements (Sec. 5.2) confirmed the results from Van Schaeybroeck, who defined the interfacial tension using the grand potential. His original argument invokes the language of thermodynamic ensembles under fixed chemical potentials. However, all the numerical studies above are based on the GPE framework that is derived for an isolated system at zero temperature; our experiment is also performed deep in the degenerate regime and is well described by GPE. The choice of thermodynamic ensemble hardly makes sense for explaining this factor of 4 difference in those cases, where thermodynamics is mostly irrelevant to the system dynamics.

As a matter of fact, there is no need to invoke the grand canonical ensemble to explain this factor of 4. Ao and Chui's argument fails because it overlooks the atom number difference when comparing energy between the system with and without the interface—at the

interface, the reference Thomas-Fermi system has a hard density step where the density for each spin abruptly drops to zero from its bulk value. In reality, the inter-spin repulsion forces a reduction in density on the scale of the spin healing length, with a density tail extending slightly beyond the interface. At identical chemical potential, this repulsion always reduces the total atom number around the interface [Eq. (A.19) below]. Computing the excess "bare" energy would incorporate the chemical potential of those removed atoms (Sec. 2.2) into the interfacial tension, whereas those atoms are *displaced*, instead of removed, in the interface-forming process. The grand potential *free energy* $\Omega = E - \mu N$ properly tracks the energy of removing atoms in this process, and thus is able to isolate the interface contribution.

To illustrate the differences in these two definitions of σ, we derive them both below. Consider a binary condensate in a homogeneous trap separated by a flat interface in the x-direction. The condensate wavefunctions for the two spins can be expressed as $\psi_{\uparrow,\downarrow} = \psi_{\uparrow,\downarrow}(y)$. Deep inside the bulk, the wavefunctions are uniform and $\psi'_{\uparrow,\downarrow} \equiv \partial_y \psi_{\uparrow,\downarrow} = 0$. The energy of the binary condensate per unit length in the x-direction is

$$E = \int \mathrm{d}y \Big\{ \frac{\hbar^2}{2m}|\psi'_\uparrow|^2 + \frac{\hbar^2}{2m}|\psi'_\downarrow|^2 \tag{A.1}$$

$$+ \frac{g_{\uparrow\uparrow}}{2}|\psi_\uparrow|^4 + \frac{g_{\downarrow\downarrow}}{2}|\psi_\downarrow|^4 + g_{\uparrow\downarrow}|\psi_\uparrow|^2|\psi_\downarrow|^2 \Big\} \tag{A.2}$$

$$\equiv A_\uparrow + A_\downarrow + \frac{1}{2}B_\uparrow + \frac{1}{2}B_\downarrow + C, \tag{A.3}$$

where $A_{\uparrow,\downarrow} \equiv \int \mathrm{d}y\, (\hbar^2/2m)|\Psi_{\uparrow,\downarrow}(\mathbf{r})|^2$, $B_{\uparrow,\downarrow} \equiv \int \mathrm{d}y\, g_{\uparrow\uparrow,\downarrow\downarrow}|\Psi_{\uparrow,\downarrow}(\mathbf{r})|^4$, and $C \equiv \int \mathrm{d}y\, g_{\uparrow\downarrow}|\Psi_\uparrow(\mathbf{r})|^2|\Psi_\downarrow(\mathbf{r})|^2$.

On the other hand, the energy for the reference Thomas-Fermi system is

$$E_{\mathrm{TF}} = \int \mathrm{d}y \Big\{ \frac{g_{\uparrow\uparrow}}{2}|\psi_{\uparrow,0}|^4 + \frac{g_{\downarrow\downarrow}}{2}|\psi_{\downarrow,0}|^4 \Big\} \tag{A.4}$$

$$= \frac{1}{2}\big(\mu_\uparrow N_{\uparrow,0} + \mu_\downarrow N_{\downarrow,0}\big), \tag{A.5}$$

where $\psi_{\uparrow/\downarrow,0} = \sqrt{\mu_{\uparrow/\downarrow}/g_{\uparrow\uparrow/\downarrow\downarrow}}$ is the Thomas-Fermi wavefunction inside the domain of its own spin and zero elsewhere, and $N_{\uparrow/\downarrow,0}$ is the atom number in the reference Thomas-Fermi system.

Meanwhile, the time-independent GPE for spin-↑ reads

$$\mu_\uparrow \psi_\uparrow = -\frac{\hbar^2}{2m}\psi''_\uparrow + g_{\uparrow\uparrow}|\psi_\uparrow|^2\psi_\uparrow + g_{\uparrow\downarrow}|\psi_\downarrow|^2\psi_\uparrow, \tag{A.6}$$

and likewise for spin-↓. Multiplying Eq. (A.6) by $\psi_\uparrow^*$ and integrating over y yields:

$$\mu_\uparrow N_\uparrow = \int \mathrm{d}y \left\{ \frac{\hbar^2}{2m}|\psi_\uparrow'|^2 + g_{\uparrow\uparrow}|\psi_\uparrow|^4 + g_{\uparrow\downarrow}|\psi_\uparrow|^2|\psi_\downarrow|^2 \right\} \tag{A.7}$$

$$= A_\uparrow + B_\uparrow + C \tag{A.8}$$

and vice versa for spin-↓. Here, $N_{\uparrow/\downarrow}$ explicitly refers to the atom number in the actual interface system. This allows us to express the energy excess of the interface system as

$$\Delta E = E - E_{\mathrm{TF}} \tag{A.9}$$

$$= A_\uparrow + A_\downarrow + \frac{1}{2}(-A_\uparrow - A_\downarrow + \mu_\uparrow N_\uparrow + \mu_\downarrow N_\downarrow) - \frac{1}{2}\left(\mu_\uparrow N_{\uparrow,0} + \mu_\downarrow N_{\downarrow,0}\right) \tag{A.10}$$

$$= \frac{1}{2}(A_\uparrow + A_\downarrow) + \frac{1}{2}(\mu_\uparrow \Delta N_\uparrow + \mu_\downarrow \Delta N_\downarrow). \tag{A.11}$$

If the difference $\Delta N_{\uparrow/\downarrow} = N_{\uparrow/\downarrow} - N_{\uparrow/\downarrow,0}$ is overlooked, we obtain $\Delta E = (A_\uparrow + A_\downarrow)/2$—the same result Ao and Chui derived in Ref. [30].

To correctly account for the atom number difference, we should instead compute

$$\Delta E - \mu_\uparrow \Delta N_\uparrow - \mu_\downarrow \Delta N_\downarrow = \frac{1}{2}(A_\uparrow + A_\downarrow) - \frac{1}{2}(\mu_\uparrow \Delta N_\uparrow + \mu_\downarrow \Delta N_\downarrow). \tag{A.12}$$

To evaluate the second term, we take the first integral of the time-independent GPEs. Multiplying both sides of Eq. (A.6) by $\psi_\uparrow'^*$ and summing it up with its complex conjugate yields:

$$\mu_\uparrow \left(|\psi_\uparrow|^2\right)' = -\frac{\hbar^2}{2m}\left(|\psi_\uparrow'|^2\right)' + \frac{g_{\uparrow\uparrow}}{2}\left(|\psi_\uparrow|^4\right)' + g_{\uparrow\downarrow}|\psi_\downarrow|^2\left(|\psi_\uparrow|^2\right)', \tag{A.13}$$

and likewise for spin-↓. Summing up the above equations for both spin-↑ and -↓, we obtain the first integral:

$$\mu_\uparrow|\psi_\uparrow|^2 + \mu_\downarrow|\psi_\downarrow|^2 \tag{A.14}$$

$$= -\frac{\hbar^2}{2m}|\psi_\uparrow'|^2 - \frac{\hbar^2}{2m}|\psi_\downarrow'|^2 + \frac{1}{2}g_{\uparrow\uparrow}|\psi_\uparrow|^4 + \frac{1}{2}g_{\downarrow\downarrow}|\psi_\downarrow|^4 + g_{\uparrow\downarrow}|\psi_\uparrow|^2|\psi_\downarrow|^2 + d. \tag{A.15}$$

Here, d is an integration constant that can be evaluated anywhere in the bulk region (ψ' vanishes, the other spin has zero density), which yields $d = \mu_\uparrow |\psi_{\uparrow,0}|^2/2$ or $d = \mu_\downarrow |\psi_{\downarrow,0}|^2/2$[1]. Its spatial integral is thus:

$$\mu_\uparrow N_\uparrow + \mu_\downarrow N_\downarrow \tag{A.16}$$

$$= -A_\uparrow - A_\downarrow + \frac{1}{2}(B_\uparrow + B_\downarrow) + C + \frac{1}{2}(\mu_\uparrow N_{\uparrow,0} + \mu_\downarrow N_{\downarrow,0}) \tag{A.17}$$

$$= -A_\uparrow - A_\downarrow + \frac{1}{2}(-A_\uparrow - A_\downarrow + \mu_\uparrow N_\uparrow + \mu_\downarrow N_\downarrow) + \frac{1}{2}(\mu_\uparrow N_{\uparrow,0} + \mu_\downarrow N_{\downarrow,0}), \tag{A.18}$$

$$\implies \frac{3}{2}(A_\uparrow + A_\downarrow) = -\frac{1}{2}(\mu_\uparrow \Delta N_\uparrow + \mu_\downarrow \Delta N_\downarrow), \tag{A.19}$$

where the third line is obtained by summing Eq. (A.8) over both spins. Inserting Eq. (A.19) into Eq. (A.12) yields:

$$\sigma = \Delta E - \mu_\uparrow \Delta N_\uparrow - \mu_\downarrow \Delta N_\downarrow = 2(A_\uparrow + A_\downarrow), \tag{A.20}$$

which is exactly Eq. (5.9), and four times larger than the value $(A_\uparrow + A_\downarrow)/2$ obtained by overlooking $\Delta N_{\uparrow/\downarrow}$ in Eq. (A.11).

[1] This also leads to the pressure balance condition $\mu_\uparrow |\psi_{\uparrow,0}|^2 = \mu_\downarrow |\psi_{\downarrow,0}|^2$ (see Sec. 2.5).

B Numerical Computation of BEC Dynamics and Quasi-Particle Spectrum

The experimental results in this thesis are heavily complemented by simulations of condensate time evolution (such as the Rayleigh-Taylor instability in Chapter 5) and numerically calculated quasi-particle spectra (Chapter 5 and 6). These numerical approaches rest on the mean-field framework and the Bogoliubov approach established in Sec. 2.4 and 2.6: the condensate dynamics is governed by a single complex order parameter (or "wave-function") whose evolution follows the Gross-Pitaevskii equation (GPE); on a mean-field ground state, the quasi-particle spectrum is acquired by diagonalizing the Hamiltonian quadratic in field perturbation.

This appendix details the numerical methods used to tackle both problems. We begin with the dimension reduction strategies and dimensionless units, followed by a survey of numerical methods for solving the GPE. Finally, we discuss the numerical diagonalization schemes for computing quasi-particle spectra.

B.1 Further reduction of dimension by symmetry

The full 3-dimensional simulation of our experimental system is generally prohibitively expensive at a workstation level. The quasi-2D geometry of our system strongly motivates us to find an equivalent numerical treatment in lower dimensions. In Sec. 2.4, we have derived an effective GPE in 2D [Eq. (2.57)], where the reduction of dimension comes from assuming the dynamics is essentially frozen in the tightly confined direction. The same spirit applies even in the quasi-2D regime: we have shown in Sec. 4.4 that the excitations along the tight direction (with a trap frequency of ω_z) have an energy cost of at least $\hbar\omega_z$ under the full 3D hydrodynamic treatment, justifying the separate treatment of low-energy dynamics.

In our actual numerical workflow, the quasi-2D condensate dynamics is always simulated with the 2D GPE [Eq. (2.57)]. However, the 3D nature of our system becomes important when the effective 2D coupling (or, in numerical terms, the conversion between the effective chemical potential and the actual chemical potential) is in question. Fortunately, we do not need to perform a full 3D ground state calculation for this purpose. In our homogeneous quasi-2D system, $\phi(x, y, z) = \phi(z)$. Both ∂_x and ∂_y terms in the 3D GPE [Eq. (2.43)] vanish, rendering the 3D GPE essentially 1D. Similarly, for the sound mode calculation, one can apply the traveling wave ansatz $\delta\phi_i(x, y) = u_i(y)e^{i(kx-\omega t)} - v_i^*(y)e^{i(-kx+\omega t)}$ and turn the BdG equation [Eq. (2.105)] into a 1D equation. Figure 4.6 is acquired through this method; the same approach applies equally well in computing the ripplon spectrum, which we will detail later.

Sometimes, even when the system is not homogeneous in the quasi-2D plane, it is still possible to perform calculations in lower dimensions if the in-plane trap geometry varies slowly enough. In this case, we can similarly drop the ∂_x and ∂_y terms in the 3D GPE and perform the calculation in 1D locally to a coordinate (x, y)—commonly known as the *local density approximation (LDA)*. To piece together the ground state wavefunctions computed across the trap region, a common chemical potential μ has to be used at all (x, y) coordinates. Usually not known *a priori*, μ has to be determined by matching the computed ground state profile to experimental measurements.

B.2 Defining the numerical units

It is convenient to define a set of units so that the physical quantities in the numerical calculations are of order one. Working with quantities in reduced units avoids underflow and overflow errors in the floating-point representation; it also helps the choice of the grid and timestep when integrating differential equations without the need for unit bookkeeping.

To examine how the unit conversion works, take the velocity as an example: it has a dimension of LT^{-1}, which corresponds to $\mathrm{m\,s}^{-1}$ in SI units. If we define l_0 and t_0 to be the units for length and time in the numerical routine, the unit of velocity becomes $l_0 t_0^{-1}$. Converting between SI and our reduced units is straightforward: we express l_0 in terms of m for all numerical quantities or vice versa. Usually, we fix t_0, l_0, m_0 to the characteristic values of the experimental system. They can also be chosen to cancel a certain physical constant. For example, by choosing appropriate t_0, l_0, and m_0, we can set $\hbar = 1 \times l_0^2 m_0 / t_0$, which is commonly referred to as "setting $\hbar$ to one". The reduced unit choices of our numerical routines and some common units derived from them are listed in Table. B.1.

Table B.1: Mapping between physical (SI) quantities and the reduced units used in our numerical routines. The last column gives each reduced unit in SI.

Quantity	SI unit $[Q]$	reduced unit $[\tilde{Q}]$	$\tilde{Q}$ expressed in SI unit
Base units			
Mass	kg	$m_0 = m_{^{23}\mathrm{Na}}$	$3.82 \times 10^{-26}\,\mathrm{kg}$
Time	s	t_0	1 ms
Action	J s	$\hbar$	$1.05 \times 10^{-34}\,\mathrm{J\,s}$
Length	m	$l_0 = \sqrt{\hbar t_0/m_0}$	1.66 µm
Derived units			
Energy	J	$\hbar/t_0$	$1.05 \times 10^{-31}\,\mathrm{J} = h \times 159\,\mathrm{Hz}$
Temperature	K	$\hbar/(k_B t_0)$	7.64 nK
Wavevector / 1D density	m^{-1}	$1/l_0$	$0.602\,\mu\mathrm{m}^{-1}$
2D density	m^{-2}	$1/l_0^2$	$0.362\,\mu\mathrm{m}^{-2}$
3D density	m^{-3}	$1/l_0^3$	$0.218\,\mu\mathrm{m}^{-3}$

Defining $\tilde{x} \equiv x/l_0$, $\tilde{y} \equiv y/l_0$, $\tilde{t} \equiv t/t_0$, and the reduced 2D wavefunction $\Phi \equiv l_0\phi$, we acquire the unitless 2D GPE:

$$\partial_{\tilde{t}}\Phi = i\left(\frac{1}{2}(\partial_{\tilde{x}}^2 + \partial_{\tilde{y}}^2) - \frac{t_0}{\hbar}V(\tilde{x},\tilde{y}) - \frac{\sqrt{8\pi}a}{z_0}N|\Phi|^2\right)\Phi, \tag{B.1}$$

$$\implies \partial_{\tilde{t}}\Phi = i\left(\frac{1}{2}(\partial_{\tilde{x}}^2 + \partial_{\tilde{y}}^2) - \tilde{V}(\tilde{x},\tilde{y}) - \tilde{g}N|\Phi|^2\right)\Phi, \tag{B.2}$$

where $z_0 = \sqrt{\hbar/m\omega_z}$ is the oscillator length of the tightly confined third dimension and a is the s-wave scattering length of ^{23}Na atoms; the unitless potential and interaction strength are defined as $\tilde{V}(\tilde{x},\tilde{y}) = (t_0/\hbar)V(\tilde{x},\tilde{y})$ and $\tilde{g} = \sqrt{8\pi}a/z_0$.

B.3 Solving for the condensate ground state

In Sec. 2.3, we have shown that minimizing the mean-field energy of the system yields the condensate ground state wavefunction. This variational problem yields the time-independent GPE:

$$\tilde{\mu}\Phi = -\left(\frac{1}{2}(\partial_{\tilde{x}}^2 + \partial_{\tilde{y}}^2) - \tilde{V}(\tilde{x},\tilde{y}) - \tilde{g}N|\Phi|^2\right)\Phi, \tag{B.3}$$

where all quantities are written in unitless form. Here, the chemical potential $\tilde{\mu} = \mu t_0/\hbar$ serves as the Lagrange multiplier that fixes the normalization.

$$\int |\Phi(\tilde{x}, \tilde{y})|^2 \, \mathrm{d}\tilde{x} \, \mathrm{d}\tilde{y} = 1. \tag{B.4}$$

Newton's method offers a direct route to root-finding for Eq. (B.3) and Eq. (B.4). After spatial discretization on a grid (Sec. B.4), we evaluate the system's Jacobian matrix and iteratively solve the linearized problem to reach the actual solution.

In practice, the standard workhorse for finding the condensate ground state is the *imaginary time propagation*. Rather than tackling Eq. (B.3) directly, the imaginary time propagation solves the pseudo-time-dependent problem:

$$\partial_{\tilde{\tau}}\Phi = \left(\frac{1}{2}(\partial_{\tilde{x}}^2 + \partial_{\tilde{y}}^2) - \tilde{V}(\tilde{x}, \tilde{y}) - \tilde{g}N|\Phi|^2\right)\Phi. \tag{B.5}$$

The wavefunction is explicitly normalized after every timestep $\Delta\tau$. High-energy excited modes decay exponentially faster than the lowest-energy ground state. After enough iterations, the solution is dominated by the true ground state wavefunction, provided the initial guess has non-zero spatial overlap with the ground state.

The imaginary time propagation method owes its popularity to its robustness and algorithmic simplicity: it bypasses the need to construct or invert a Jacobian matrix. However, it struggles to converge when the lowest energy levels are near degenerate, and demands careful timestep control when targeting metastable states (for example, the metastable magnetic gradient configuration in Sec. 5.3) to prevent overshooting to the true global ground state. By contrast, Newton's method converges rapidly near the solution and reliably settles into metastable states, at the cost of assembling the Jacobian and solving a linear system at every iteration.

B.4 Survey of numerical methods for solving GPE

Simulating condensate time evolution requires numerical integration of the GPE. Mathematically, the GPE is a nonlinear partial differential equation (PDE), first-order in time and second-order in space. To solve it numerically, we approximate continuous time evolution by discrete timesteps. The simplest choice is the explicit Euler method, which steps forward as:

$$\Phi(\tilde{\mathbf{r}}, \tilde{t} + \Delta t) = \Phi(\tilde{\mathbf{r}}, \tilde{t}) + \Delta t \dot{\Phi}(\tilde{\mathbf{r}}, \tilde{t}), \tag{B.6}$$

where the dot denotes the time derivative, and Δt is the timestep in units of t_0.

The spatial degrees of freedom are similarly handled by discretizing the continuous field Φ onto a grid. The Laplacian term $\nabla^2\Phi$ can be evaluated in several ways. The most direct approach approximates the spatial derivatives with local finite differences. For instance, on a Cartesian grid with dimensionless spacing Δx, the second-order spatial derivative takes the form:

$$\partial_x^2\Phi(x,y) \approx \frac{\Phi(x+\Delta x,y)+\Phi(x-\Delta x,y)-2\Phi(x,y)}{\Delta x^2}. \tag{B.7}$$

Spatial discretization converts the continuous PDE into a system of coupled ordinary differential equations (ODEs), which are readily solved by standard ODE integrators. This framework is commonly known as the finite-difference time-domain (FDTD) scheme.

Alternatively, the Laplacian can be evaluated in momentum space using the Fourier transform. Expanding $\Phi(x,y)$ into a Fourier series $\Phi(x,y) = \sum \tilde{\Phi}(k_x,k_y)\exp\big[i(k_x x + k_y y)\big]$ transforms spatial derivatives into simple multiplications:

$$\partial_x^2\Phi(x,y) = \sum(-k_x^2)\tilde{\Phi}(k_x,k_y)\exp\big[i(k_x x + k_y y)\big]. \tag{B.8}$$

Two methods follow. The first is the pseudo-spectral method, which uses the standard time-stepping method but replaces the finite difference with spectral derivative evaluation. The second is Strang splitting—a Trotterization scheme that sequentially applies the Laplacian step (diagonal in Fourier space) and the potential/interaction step (diagonal in real space) within a single timestep. Each of the three methods has certain strengths and drawbacks. In our case, Strang splitting offers good efficiency at moderate precision, whereas the pseudo-spectral method excels at high precision, as discussed below.

Before detailing each method, it is instructive to highlight the central problem shared by all explicit numerical schemes. While mapping a continuous field onto a discrete grid is intuitive, the stability criteria of the resulting numerical schemes are surprisingly subtle. In explicit schemes that rely on spatial discretization, the timestep is tied to grid spacing: the finer the grid, the smaller the maximum stable timestep for which the scheme does not diverge. This is the essence of the Courant-Friedrichs-Lewy (CFL) condition: the numerical domain of dependence must encompass the PDE's domain of dependence. Or, in physical terms: the timestep of a numerical scheme cannot outpace the information propagation across grid sites[1]. One might naively expect that stepping faster than the grid's fastest dy-

[1] The CFL condition was originally formulated for hyperbolic PDEs with linear dispersion like the wave equations, where the signal propagation speed is constant. However, the dispersion relations of the Schrödinger equation and GPE are not linear. The fastest grid mode travels at the highest group velocity $k_{\max} \sim 1/\Delta x$;

namics simply introduces benign inaccuracies in those fast modes, about which we rarely care. But the problem extends beyond inaccuracy: if the timestep fails to resolve these fast-oscillating modes[2], local errors compound exponentially across timesteps, which seeds the eventual catastrophic divergence of the solution. PDEs of this kind are classic examples of *stiff* problems—a numerical bottleneck that identically plagues the GPE.

As spatial resolution increases ($\Delta x \to 0$), the timestep required merely to prevent numerical divergence becomes prohibitively small, motivating the development of *implicit methods*. Unlike the explicit schemes, which approximate the next step by current derivatives (for example, Eq. (B.6)), implicit methods evaluate $\dot{\Phi}$ at future time $t + \Delta t$:

$$\Phi(\tilde{\mathbf{r}}, \tilde{t} + \Delta t) = \Phi(\tilde{\mathbf{r}}, \tilde{t}) + \Delta t \dot{\Phi}(\tilde{\mathbf{r}}, \tilde{t} + \Delta t). \tag{B.9}$$

Practically, $\dot{\Phi}(\tilde{\mathbf{r}}, \tilde{t} + \Delta t)$ is unknown at the time $\tilde{t}$; Eq. (B.9) thus defines a nonlinear system that must be solved at every single timestep. The reward for this computational burden is *unconditional stability*: the timestep is freed from the CFL constraint. For the typical grid dimensions used in our work, the overhead of implicit solvers outweighs the stability penalty; GPU-accelerated explicit schemes outperform implicit solvers. In the following sections, we focus exclusively on explicit schemes.

Finite-difference time-domain (FDTD) method

FDTD represents the battle-tested workhorse for partial differential equations, enjoying sustained popularity across computational fluid dynamics and electromagnetic field simulations. Because FDTD acts locally, the parallelized evaluation of the right-hand side (RHS) of Eq. (B.2) maps naturally onto GPU architectures. In particular, a massive grid can be decomposed into subgrids across multiple GPUs or compute nodes, requiring only boundary-value exchanges between adjacent subgrids. Listing B.1 illustrates a representative GPU kernel implementation of the RHS in Julia.

Listing B.1: Pseudo Julia code of a GPU kernel that calculates the right- hand side of GPE in a FDTD scheme.

```
@inline function lapl_2d_periodic_elem(u, iy, ix, Nx, Ny, delta_x, delta_y)
   # Helper function that helps compute the Laplacian on grid site (ix, iy)
   # under periodic boundary condition.
   # It can be run from both CPU and GPU.

```

the minimal propagation time between grid sites is thus $\propto \Delta x^2$, mirroring the stability bound derived via von Neumann stability analysis shown below.

[2] The high-k modes of the GPE are purely oscillatory. For diffusion problems (including the imaginary-time GPE problem), the fastest dynamics are exponential damping. However, the same stability bound exists: the timestepping error turns exponential damping into growth, which also leads to divergence.

```
    ix_p1 = mod1(ix + 1, Nx)
    ix_m1 = mod1(ix - 1, Nx)
    iy_p1 = mod1(iy + 1, Ny)
    iy_m1 = mod1(iy - 1, Ny)

    return (-2*u[iy, ix] + u[iy_m1, ix] + u[iy_p1, ix]) / delta_y^2 +
           (-2*u[iy, ix] + u[iy, ix_m1] + u[iy, ix_p1]) / delta_x^2
end

@kernel function gpe_full_rhs_kernelfunc!(dR, dI, R, I, p)
    # GPU kernel that computes the right-hand side (RHS) of the GPE.
    # Input: R, I - real and imaginary of the wavefunction
    # Output: dR, dI - real and imaginary part of the RHS

    iy, ix = @index(Global, NTuple) # retrieve the grid indices this core is at

    u2 = R[iy, ix]^2 + I[iy, ix]^2

    # compute the laplacian for real and imaginary part
    dR[iy, ix] = gpe_dR(lapl_2d_periodic_elem(I, iy, ix, Nx, Ny, delta_x, delta_y), u2, ...)
    dI[iy, ix] = gpe_dI(lapl_2d_periodic_elem(R, iy, ix, Nx, Ny, delta_x, delta_y), u2, ...)
end
```

Once spatially discretized, time-stepping is handled by standard ODE integrators, such as those provided by the DifferentialEquations.jl / OrdinaryDiffEq.jl ecosystem [143]. In practice, we adopt Verner's "Most Efficient" 7/6-order Runge-Kutta method (Vern7 in OrdinaryDiffEq.jl) with adaptive time-stepping [144], which provides an optimal balance between speed and accuracy.

In FDTD, the solution accuracy is fundamentally limited by the truncation error of the discretized Laplacian. The maximum stable timestep is dictated by the numerical growth of errors in the highest-k mode on the grid[3]. To quantify the compounding of time-stepping error across steps, we evaluate the amplification factor R of a small perturbation in the highest-k mode; stability requires $|R| \leq 1$ [145, Chapter 9]. For real-time GPE evolution, von Neumann stability analysis reveals that the forward Euler method is unstable at any Δt; the 4th-order Runge-Kutta (RK4) method yields a stability bound of $\Delta t \leq (\sqrt{2}/d)\Delta x^2 \approx (1.41/d)\Delta x^2$, where d is the spatial dimension[4].

Pseudo-spectral method

Unlike FDTD, the pseudo-spectral method evaluates the Laplacian via Fourier transforms [Eq.(B.8)]. This spectral approach (Listing B.2) works best for smoothly varying, periodic fields, but develops ringing artifacts when the spatial features approach grid resolution.

[3] Also known as the zig-zag or the checkerboard mode, $\Phi_j \propto (-1)^j$ (j is the grid index), corresponding to the Nyquist limit $k_{\max}/2\pi = 1/(2\Delta x)$.

[4] This stability bound Δt (and those below) assumes a flat potential and that the kinetic energy dominates the interaction at the scale Δx. In practice, Δt may be further restricted by interaction energy and inhomogeneous potentials.

Listing B.2: Pseudo Julia code that evaluates the Laplacian with FFT.

```
buf = ... # The array that holds the intermediate value of the RHS

k = π2 .* collect(FFTW.fftfreq(N, 1 / dx))
k2 = [kx^2 + ky^2 for ky in k, kx in k] # |k|^2 on the 2D grid, [iy, ix]
pf = plan_fft!(buf) # Make FFT plans
pb = plan_ifft!(buf)

@inline function lapl_2d_periodic_fft!(buf, u, k2, pf, pb)
   # Laplacian of u on the whole grid (periodic boundary), written to buf.

   copyto!(buf, u)
   pf * buf
   @. buf *= -k2
   pb * buf
end
```

Compared to the FDTD method, eliminating truncation errors yields far better accuracy at the same timestep. Counterintuitively, however, the maximum stable timestep for the pseudo-spectral method is actually *smaller*: stability for RK4 requires $\Delta t \leq (4\sqrt{2}/d\pi^2)\Delta x^2 \approx (0.57/d)\Delta x^2$. It turns out that the truncation error in FDTD distorts dynamics of the mode at $k_{\max}$, lowering its frequency $\omega_{\mathrm{FDTD}}(k_{\max}) < \omega(k_{\max})$ and inadvertently relaxes the stability constraint. By contrast, the spectral approach computes the kinetic energy of that mode exactly; the true mode frequency $\omega(k_{\max})$ now limits the stability, forcing the integrator to take finer timesteps.

Strang splitting method

The "surprise" from the pseudo-spectral analysis highlights a key distinction: spatial discretization error dictates accuracy, whereas compounding timestepping error governs numerical stability. We seek a better way of timestepping: instead of the linear approximation [Eq. (B.6)], we consider the formal wavefunction evolution of a single timestep:

$$\Phi(\tilde{\mathbf{r}}, \tilde{t} + \Delta t) \approx \exp\left[i\Delta t\left(\frac{1}{2}(\partial_{\tilde{x}}^2 + \partial_{\tilde{y}}^2) - \tilde{V}(\tilde{x}, \tilde{y}) - \tilde{g}N|\Phi|^2\right)\right]\Phi(\tilde{\mathbf{r}}, \tilde{t}). \tag{B.10}$$

To actually evaluate the matrix exponential, we split it into two terms:

$$\exp\left[i\Delta t\left(-\tilde{V}(\tilde{x}, \tilde{y}) - \tilde{g}N|\Phi|^2\right)\right]\exp\left[i\frac{\Delta t}{2}\left((\partial_{\tilde{x}}^2 + \partial_{\tilde{y}}^2)\right)\right]. \tag{B.11}$$

Here, the kinetic term is diagonal in Fourier space and the potential/interaction term is diagonal in real space, both of which can be evaluated exactly. This factorization, known as Lie splitting [146], only incurs a commutation error given by the Baker–Campbell–Hausdorff formula. While Lie splitting is unconditionally stable for the linear Schrödinger

equation, the nonlinear interaction term in the GPE reintroduces a stability bound $\Delta t \leq (2/d\pi)\Delta x^2 \approx (0.64/d)\Delta x^2$.

There is no particular reason to order the two exponents the way Lie splitting does. We can symmetrize it by breaking the kinetic step into two halves and wrapping them around the potential and interaction step, yielding:

$$\exp\left[i\frac{\Delta t}{4}\left((\partial_{\tilde{x}}^2+\partial_{\tilde{y}}^2)\right)\right]\exp\left[i\Delta t\left(-\tilde{V}(\tilde{x},\tilde{y})-\tilde{g}N|\Phi|^2\right)\right]\exp\left[i\frac{\Delta t}{4}\left((\partial_{\tilde{x}}^2+\partial_{\tilde{y}}^2)\right)\right]. \tag{B.12}$$

This is known as Strang splitting [146]. While seemingly more computationally demanding, adjacent kinetic steps from the two consecutive iterations collapse into a single full step. Consequently, Strang splitting requires virtually no extra computation over Lie splitting, apart from the two additional kinetic half steps at the beginning and the end of the entire evolution. The payoff is substantial: Strang splitting removes the offset error caused by the term order in Eq. (B.11), resulting in a local splitting error of $O(\Delta t^3)$ per step (globally $O(\Delta t^2)$), compared to the $O(\Delta t^2)$ (globally $O(\Delta t)$) of Lie splitting. Strang splitting is not included in OrdinaryDiffEq.jl at the time of writing. We implement it manually. Listing B.3 provides a custom Julia implementation of the Strang splitting scheme.

Listing B.3: Pseudo Julia code that integrates the RHS using Strang splitting.

```
buf = ... # The array that holds the intermediate value of the RHS

k = π2 .* collect(FFTW.fftfreq(N, 1 / dx))
k2 = [kx^2 + ky^2 for ky in k, kx in k] # |k|^2 on the 2D grid, [iy, ix]
pf = plan_fft!(buf) # Make FFT plans
pb = plan_ifft!(buf)

function strang!(u, k2, pf, pb, V, g, mu, dt, nsteps)
   # Integrate the RHS of GPE using Strang splitting

   kf = cis.(-k2 .* (dt * 0.5)) # full Laplacian step
   kh = cis.(-k2 .* (dt * 0.25)) # half Laplacian step

   pf * u
   u .*= kh # Prepend with a half Laplacian step
   pb * u

   for s in 1:nsteps
      # Lie splitting for each dt step
      @. u *= cis(-(V + g * abs2(u) - mu) * dt)
      pf * u
      u .*= (s < nsteps ? kf : kh) # Do a half Laplacian step at the end
      pb * u
   end

   return u
end
```

Although the maximum stable timestep of Strang splitting is tighter than FDTD paired with RK4, each step requires far fewer operations while delivering better accuracy.

Performance comparison across numerical schemes

We benchmark the efficiency and numerical precision of each scheme across two common experimental setups: dipole oscillations in a 2D harmonic trap and the fundamental sound mode in a 2D box trap (see Sec. 4.4). We discretize the spatial domain on a 512×512 grid with $\Delta x = 0.153 l_0 = 0.254\,\mu\text{m}$ and perform time integration up to $t = 100 t_0 = 100\,\text{ms}$ across various timesteps within the stability constraint[5]. Figure B.1 compiles the resulting errors in center-of-mass position, oscillation frequency, and wavefunction normalization against total runtime.

An optimal numerical scheme takes less time and achieves better precision, occupying the lower-left corner of each plot. Generally, refining the time step improves accuracy at the expense of runtime. However, a few cases in the plot depart from this trend. At single precision (`Float32`), Strang splitting shows worse results at finer timesteps: rounding errors from repeated FFTs scale with the total step count, eventually outweighing the refinement from a smaller timestep. Several `Vern7` runs show clustered points across different precisions at similar runtimes, indicating that the integrator timestep is stability-limited rather than tolerance-limited.

Because consumer NVIDIA GPUs deliberately throttle the `Float64` throughput, evaluating the trade-offs between `Float32` and `Float64` is warranted. The FDTD scheme is memory bandwidth-limited due to its simple, local nature, hence showing the least sensitivity to the choice of floating-point precision. However, the FFT-based pseudo-spectral and Strang splitting schemes become computation-bound at 64-bit precision. Switching to `Float64` increases runtime by a factor of 3–6 on our 512×512 grid, and up to a factor of 11 on a 2048×2048 grid.

In practice, Strang splitting operating at a timestep adequately below its stability bound (for example, at 2/3 of $\Delta t_{\text{max}} = (2/d\pi)\Delta x^2$) achieves the highest speed for low-energy dynamics at moderate precision, except in hard-wall box potentials where FDTD shows a modest advantage. When high precision is required, both Strang splitting and pseudo-spectral methods dominate, though the optimal choice depends on problem-specific details and requires case-by-case benchmarking. Over-refining the timestep in 32-bit precision Strang splitting must be avoided, as it degrades the solution accuracy rather than improving it. Regardless of floating-point precision, FDTD schemes show similarly high error

[5] The error tolerance is used as a control knob for the timestep of `Vern7`. However, its actual timestep is mostly stability-limited (see below).

due to the Laplacian truncation error and are the least favorable choice among the three for smooth potentials.

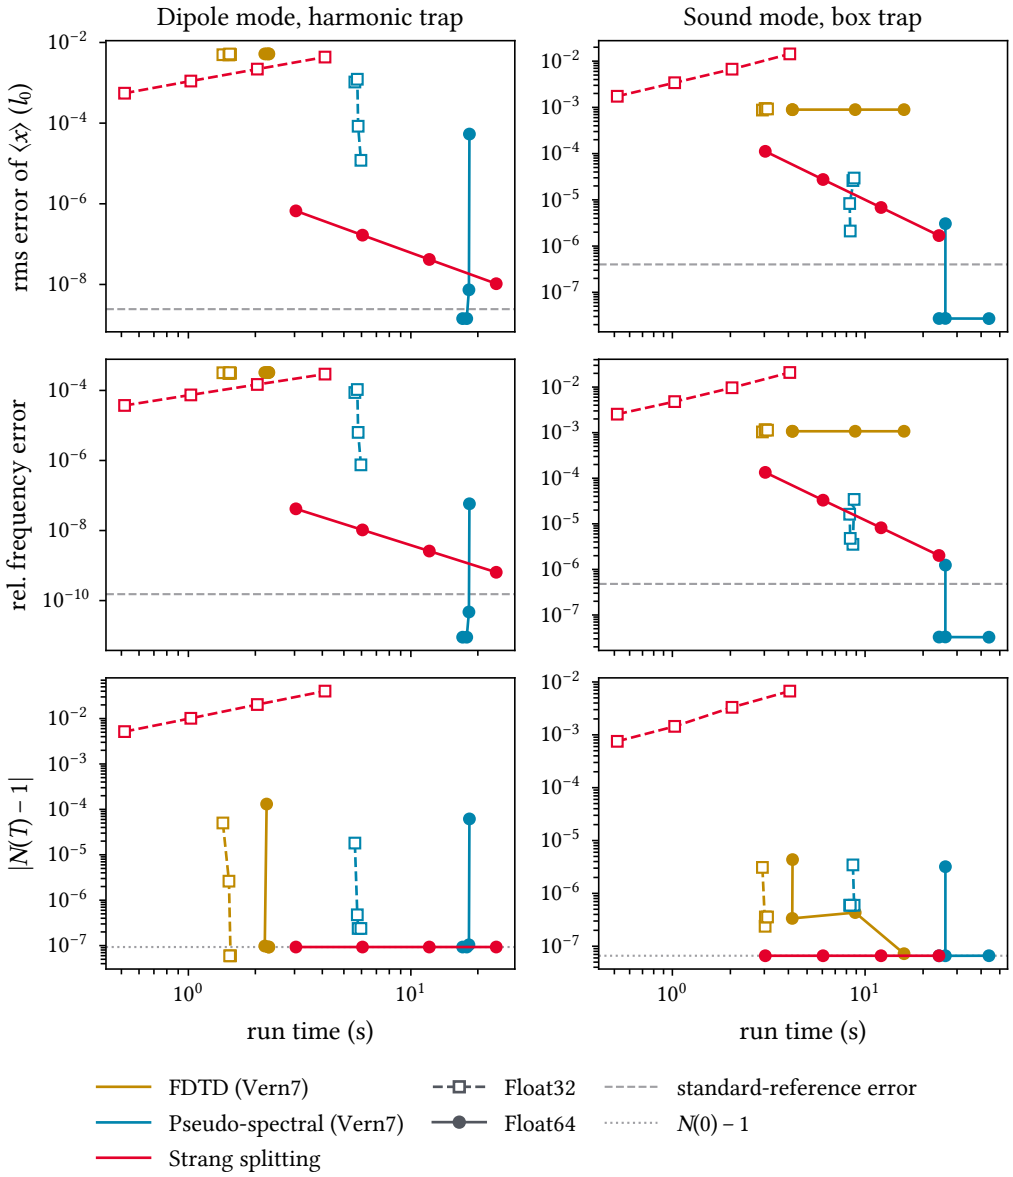

Figure B.1: Errors of the 2D GPE solvers against wall time for the dipole mode in a harmonic trap (left) and the lowest sound mode in a box trap (right), integrated for $t = 100t_0$ on a 512×512 grid on a single NVIDIA RTX 4080 GPU. Rows: rms error of the center of mass $\langle x \rangle(t)$, relative error of the fitted oscillation frequency, and the drift of the normalization $|N(T) - 1|$. Colors mark the schemes; open squares with dashed lines are Float32, filled circles with solid lines Float64. FDTD and pseudo-spectral runs use adaptive Vern7 with error tolerances from 10^{-3} to 10^{-6} (Float32) and 10^{-3} to 10^{-9} (Float64); Strang runs use fixed timesteps $\Delta t = (0.625\text{–}5) \times 10^{-3} t_0$. The errors are measured against Float64 Strang-splitting references: Float32 runs against one at $\Delta t = 3.125 \times 10^{-4} t_0$, Float64 runs against a finer one at $\Delta t = 7.8 \times 10^{-5} t_0$. Dashed gray line: the error of the coarser reference itself, below which the Float32 errors are not resolved. Dotted gray line: $N(0) - 1$, the normalization offset of the initial state.

B.5 Computing the quasi-particle spectrum

Determining the quasi-particle spectrum amounts to finding the eigenvalues and eigenvectors of the Bogoliubov–de Gennes (BdG) matrix linearized about the condensate ground or metastable state [Eq. (2.105)]. In this section, we focus on the single-component condensate; the generalization to the two-component case follows analogously. In dimensionless units, the BdG matrix for a single-component condensate reads

$$\sigma_z \tilde{M}(\tilde{\mathbf{r}}) \begin{pmatrix} u_i(\tilde{\mathbf{r}}) \\ v_i(\tilde{\mathbf{r}}) \end{pmatrix} = \epsilon_i \begin{pmatrix} u_i(\tilde{\mathbf{r}}) \\ v_i(\tilde{\mathbf{r}}) \end{pmatrix}, \tag{B.13}$$

where $\sigma_z = \mathrm{diag}(1, -1)$ and

$$\tilde{M}(\tilde{\mathbf{r}}) = \begin{pmatrix} -\frac{1}{2}\tilde{\nabla}^2 - \tilde{\mu} + \tilde{V}(\tilde{x}, \tilde{y}) + 2\tilde{g}N|\Phi_0(\tilde{x}, \tilde{y})|^2 & -\tilde{g}N\Phi_0(\tilde{x}, \tilde{y})^2 \\ -\tilde{g}N\Phi_0^*(\tilde{x}, \tilde{y})^2 & -\frac{1}{2}\tilde{\nabla}^2 - \tilde{\mu} + \tilde{V}(\tilde{x}, \tilde{y}) + 2\tilde{g}N|\Phi_0(\tilde{x}, \tilde{y})|^2 \end{pmatrix}. \tag{B.14}$$

Here, $\tilde{\nabla}^2$ denotes the dimensionless Laplacian, and Φ_0 is the condensate ground state or metastable state wavefunction.

To compute the spectrum of $\sigma_z \tilde{M}(\tilde{\mathbf{r}})$ numerically, we discretize it onto a 2D grid, approximate the Laplacian via a discrete form, and flatten the grid points into a single column vector. This yields a large matrix:

$$\mathcal{M} = \sigma_z \begin{pmatrix} \mathcal{P} & \mathcal{Q} \\ \mathcal{Q}^* & \mathcal{P} \end{pmatrix}. \tag{B.15}$$

Here, the block matrices $\mathcal{P}$ and $\mathcal{Q}$ have elements

$$\mathcal{P}_{ij} = -\frac{1}{2}\tilde{\nabla}^2_{ij} + \delta_{ij}(\tilde{V}_i + 2\tilde{g}N|\Phi_{0;i}|^2 - \tilde{\mu}), \tag{B.16}$$

$$\mathcal{Q}_{ij} = -\delta_{ij}\tilde{g}N\Phi^2_{0;i}, \tag{B.17}$$

where i, j are the row and column indices of the sub-blocks. A standard choice for the discretized Laplacian is the finite-difference scheme [Eq. (B.7)], which produces a sparse, banded matrix (tridiagonal in 1D) except for the boundary terms.

Numerical method for solving BdG eigenproblem

Solving the full eigenproblem of a 2D BdG matrix can be challenging. In practice, only the few lowest-lying quasi-particle modes are physically of interest. This is a common theme

in many computational problems: the matrix is large and probably sparse; the goal is not to find the entire spectrum but a handful of extremal eigenpairs. The classic archetype for solving such problems is power iteration: to find the single eigenvector $v_{\max}$ corresponding to the eigenvalue of largest magnitude $\lambda_{\max}$ of a matrix A, we start from an initial guess b and compute $A^n b$. For sufficiently large n, only $v_{\max}$ will survive: $A^n b \propto \lambda_{\max}^n v_{\max}$. Modern eigensolvers generalize power iteration via the Arnoldi method. Instead of finding the single dominant eigenvector, the Arnoldi method constructs an orthonormal basis $q_1, \ldots, q_m$ spanning the m-dimensional Krylov subspace $\mathcal{K}_m = \mathrm{span}\{b, Ab, A^2 b, \ldots, A^{m-1} b\}$. We expect that vectors $A^n b$ at large n will align with the subspace spanned by the eigenvectors with dominant eigenvalues. We can therefore project A into $\mathcal{K}_m$ and solve an eigenproblem of much smaller dimension. Collecting the basis vectors as columns of the matrix $Q_m = (q_1, \ldots, q_m)$, we find the projection of A onto the Krylov subspace, $H_m = Q_m^\dagger A Q_m$. The extremal eigenvalues of H_m rapidly converge to the largest-magnitude eigenvalues of A, while transforming the eigenvectors of H_m back to A's original basis yields accurate approximations to the corresponding eigenvectors of A. The cost of all the above is merely a number of matrix–vector products (MVPs) Ax, making the method especially well suited to large sparse matrices.

For our purposes, however, we need the eigenvalues of the BdG matrix closest to zero. This is accomplished by applying the Arnoldi method to $\mathcal{M}^{-1}$, which maps the smallest-magnitude eigenvalues to the largest of the spectrum. In actual implementation, we never construct the matrix inverse explicitly since the algorithm only requires the matrix–vector product $y = \mathcal{M}^{-1} x$. Evaluating it amounts to solving the linear system $\mathcal{M} y = x$, which can be efficiently handled by computing and caching the sparse LU decomposition of $\mathcal{M}$ once and reusing it at every iteration. In practice, the BdG matrix constructed above possesses zero modes, rendering it singular and causing numerical difficulties for the LU decomposition. To circumvent this singularity, we employ a shift-and-invert technique: we shift the matrix by a small diagonal term σI and calculate the eigenvalues λ_i of $(\mathcal{M} - \sigma I)^{-1}$ instead. The actual BdG excitation energies are then recovered via $\epsilon_i = \sigma + 1/\lambda_i$. Listing B.4 shows our implementation in Julia, using the Krylov–Schur restarted Arnoldi method of ArnoldiMethod.jl [147].

Listing B.4: Julia code that finds the BdG eigenmodes closest to σ with the shift-and-invert Arnoldi method.

```
using LinearAlgebra, SparseArrays
using ArnoldiMethod, LinearMaps

function si_eigs(M, sigma; nev = 10, tol = 1e-12)
   # Find the nev eigenvalues of M closest to sigma by Arnoldi + shift-and-invert.

```

```
    F = lu(M - sigma * I) # cached sparse LU factorization, computed once

    # Linear map x -> (M - sigma I)^{-1} x that computes the MVP,
    # using the cached LU factors
    T = complex(eltype(M))
    op = LinearMap{T}((y, x) -> (y .= F \ x), size(M, 1); ismutating = true)

    # Calling partialschur from ArnoldiMethod.jl
    decomp, history = partialschur(op; nev = nev, tol = tol,
                            restarts = 300, which = :LM)
    theta, X = partialeigen(decomp)

    return sigma .+ 1 ./ theta, X # BdG energies and eigenvectors
end

# Small complex shift to avoid the singularity;
# add a small imaginary part to avoid running into an actual real eigenvalue
energies, modes = si_eigs(M_bdg, 1e-4 + 1e-4im; nev = 12)
```

Traveling-wave ansatz

Running the Arnoldi method on a 512×512 2D BdG problem is a non-issue even on a laptop nowadays. We can even reduce the 2D problem to 1D if the system is translationally invariant along one direction, as in both the phonon and ripplon cases. Assuming homogeneity along the x-direction, the longitudinal wavevector k_x serves as a good quantum number. Imposing the traveling-wave ansatz $\delta\phi(x, y, t) = u(y)e^{i(k_x x - \omega t)} - v^*(y)e^{i(-k_x x + \omega t)}$ separates out the longitudinal plane-wave motion, reducing the eigenproblem to one dimension. For a given k_x, the sub-blocks of the reduced dimensionless BdG matrix take the form

$$\mathcal{P}_{k_x;i} = \frac{1}{2}(-\partial_{\tilde{y}}^2 + \tilde{k}_x^2) - \tilde{\mu} + \tilde{V}_i + 2\tilde{g}N|\Phi_{0;i}|^2. \tag{B.18}$$

Tracking the eigenvalue of the lowest-lying mode as a function of k_x directly yields the dispersion relation of the system's Nambu-Goldstone mode. It is also worth noting that k_x here can be varied continuously. This is a clear advantage over solving the 2D problem directly—the finite-size momentum quantization is bypassed without the need to make the system infinitely long.

Acronyms

AOM	Acousto-optic modulator
BdG	Bogoliubov–de Gennes
BEC	Bose-Einstein condensate
BKT	Berezinskii-Kosterlitz-Thouless
CCD	Charge-coupled device
CFL	Courant-Friedrichs-Lewy
DAC	Digital-to-analog converter
DAMOP	Division of Atomic, Molecular and Optical Physics (American Physical Society)
DC	Direct current
DCT	Discrete cosine transform
DMD	Digital micromirror device
FDTD	Finite-difference time-domain
FFT	Fast Fourier transform
FM	Ferromagnetic
FPUT	Fermi-Pasta-Ulam-Tsingou
GP	Gaussian process
GPE	Gross-Pitaevskii equation
GPM	Gallons per minute
GPU	Graphics processing unit
IGBT	Insulated-gate bipolar transistor
IR	Infrared
JILA	Formerly the Joint Institute for Laboratory Astrophysics
JQI	Joint Quantum Institute
KAIST	Korea Advanced Institute of Science and Technology
LDA	Local density approximation
LED	Light-emitting diode
LIAD	Light-induced atomic desorption

LM	Lagrange multiplier
LU	Lower-upper (matrix factorization)
MIT	Massachusetts Institute of Technology
MOT	Magneto-optical trap
MTF	Modulation transfer function
NA	Numerical aperture
NIST	National Institute of Standards and Technology
OD	Optical density
ODE	Ordinary differential equation
ODT	Optical dipole trap
PCA	Principal component analysis
PDE	Partial differential equation
PDF	Probability density function
PGC	Polarization gradient cooling
PM	Paramagnetic
PSD	Power spectral density
PTAI	Partial transfer absorption imaging
RF	Radio frequency
RHS	Right-hand side
RTI	Rayleigh-Taylor instability
SatAbs	Saturated absorption spectroscopy
SI	International System of Units
SLM	Spatial light modulator
SP	Spin polarization
SPOT	Spontaneous-force optical trap
TF	Thomas-Fermi
UHV	Ultra-high vacuum
UMD	University of Maryland
UV	Ultraviolet

Bibliography

[1] D. M. Stamper-Kurn, *Peeking and poking at a new quantum fluid: studies of gaseous Bose-Einstein condensates in magnetic and optical traps*, Ph.D. thesis, Massachusetts Institute of Technology, Dept. of Physics (2000).

[2] A. Einstein, Quantentheorie des einatomigen idealen Gases. Zweite Abhandlung, in *Sitzungsberichte der Preußischen Akademie der Wissenschaften (Berlin)* (Physikalisch-mathematische Klasse, 1925) pp. 3–14.

[3] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, Science **269**, 198–201 (1995).

[4] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Stamper-Kurn, and W. Ketterle, Physical Review Letters **75**, 3969–3973 (1995).

[5] W. Ketterle, Reviews of Modern Physics **74**, 1131–1151 (2002).

[6] E. A. Cornell and C. E. Wieman, Nobel lecture: Bose-Einstein condensation in a dilute gas: the first 70 years and some recent experiments, in *Collected Papers of Carl Wieman* (World Scientific, 2008) p. 626–644.

[7] P. Kapitza, Nature **141**, 74–74 (1938).

[8] J. F. Allen and A. D. Misner, Nature **142**, 643–644 (1938).

[9] F. London, Nature **141**, 643–644 (1938).

[10] O. Penrose and L. Onsager, Physical Review **104**, 576–584 (1956).

[11] T. R. Sosnick, W. M. Snow, P. E. Sokol, and R. N. Silver, Europhysics Letters (EPL) **9**, 707–712 (1989).

[12] R. T. Azuah, W. G. Stirling, H. R. Glyde, M. Boninsegni, P. E. Sokol, and S. M. Bennington, Physical Review B **56**, 14620–14630 (1997).

[13] I. F. Silvera and J. T. M. Walraven, Physical Review Letters **44**, 164–168 (1980).

[14] R. W. Cline, D. A. Smith, T. J. Greytak, and D. Kleppner, Physical Review Letters **45**, 2117–2120 (1980).

[15] N. Masuhara, J. M. Doyle, J. C. Sandberg, D. Kleppner, T. J. Greytak, H. F. Hess, and G. P. Kochanski, Physical Review Letters **61**, 935–938 (1988).

[16] D. G. Fried, T. C. Killian, L. Willmann, D. Landhuis, S. C. Moss, D. Kleppner, and T. J. Greytak, Physical Review Letters **81**, 3811 (1998).

[17] T.-L. Ho, Physical Review Letters **81**, 742 (1998).

[18] J. Stenger, S. Inouye, D. M. Stamper-Kurn, H.-J. Miesner, A. P. Chikkatur, and W. Ketterle, Nature **396**, 345–348 (1998).

[19] D. M. Stamper-Kurn, M. R. Andrews, A. P. Chikkatur, S. Inouye, H.-J. Miesner, J. Stenger, and W. Ketterle, Physical Review Letters **80**, 2027–2030 (1998).

[20] L. D. Kapitza, The Journal of Physics, USSR **5**, 71 (1941).

[21] L. Pitaevskii and S. Stringari, *Bose-Einstein Condensation and Superfluidity*, 1st ed. (Oxford University Press, 2016).

[22] Lord Rayleigh, Proceedings of the London Mathematical Society **s1-15**, 69–78 (1883).

[23] C. C. Grimes and G. Adams, Physical Review Letters **42**, 795–798 (1979).

[24] D. S. Fisher, B. I. Halperin, and P. M. Platzman, Physical Review Letters **42**, 798–801 (1979).

[25] Y. P. Monarkha and V. B. Shikin, Soviet Journal of Low Temperature Physics **8**, 279–298 (1982).

[26] V. B. Shikin, Soviet Physics Uspekhi **32**, 452–455 (1989).

[27] M. Dykman, K. Kono, D. Konstantinov, and M. Lea, Physical Review Letters **119**, 256802 (2017).

[28] C. J. Myatt, E. A. Burt, R. W. Ghrist, E. A. Cornell, and C. E. Wieman, Physical Review Letters **78**, 586–589 (1997).

[29] D. S. Hall, M. R. Matthews, J. R. Ensher, C. E. Wieman, and E. A. Cornell, Physical Review Letters **81**, 1539–1542 (1998).

[30] P. Ao and S. T. Chui, Physical Review A **58**, 4836 (1998).

[31] B. Van Schaeybroeck, Physical Review A **78**, 023624 (2008).

[32] I. E. Mazets, Physical Review A **65**, 033618 (2002).

[33] Y. Geng, J. Tao, M. Zhao, S. Mukherjee, S. P. Eckel, G. K. Campbell, and I. B. Spielman, Science Advances **11**, 10.1126/sciadv.adw9752 (2025).

[34] S. Huh, W. Yun, G. Yun, S. Hwang, K. Kwon, J. Hur, S. Lee, H. Takeuchi, S. K. Kim, and J.-y. Choi, Nature Physics **21**, 1398–1403 (2025).

[35] J. P. Boris, Comments Plasma Phys. Controlled Fusion **3**, 1 (1977).

[36] C. J. Pethick and H. Smith, *Bose–Einstein Condensation in Dilute Gases*, 2nd ed. (Cambridge University Press, 2008).

[37] J. J. Sakurai and J. Napolitano, *Modern Quantum Mechanics*, 2nd ed. (Cambridge University Press, 2018).

[38] J. Klaers, J. Schmitt, F. Vewinger, and M. Weitz, Nature **468**, 545–548 (2010).

[39] K. Huang, *Statistical mechanics*, 2nd ed. (Wiley, New York, 1987).

[40] C. Weiss and M. Wilkens, Optics Express **1**, 272 (1997).

[41] M. Christensen, T. Vibel, A. Hilliard, M. Kruk, K. Pawłowski, D. Hryniuk, K. Rzążewski, M. Kristensen, and J. Arlt, Physical Review Letters **126**, 153601 (2021).

[42] R. J. Fletcher, M. Robert-de Saint-Vincent, J. Man, N. Navon, R. P. Smith, K. G. Viebahn, and Z. Hadzibabic, Physical Review Letters **114**, 255302 (2015).

[43] J.-L. Ville, *Quantum gases in box potentials : sound and light in bosonic Flatland*, Ph.D. thesis, Agence Bibliographique de l'Enseignement Supérieur.

[44] S. Giorgini, L. P. Pitaevskii, and S. Stringari, Journal of Low Temperature Physics **109**, 309–355 (1997).

[45] J. W. Negele and H. Orland, *Quantum Many-Particle Systems*, Advanced Books Classics (CRC Press, 1998).

[46] F. Dalfovo, S. Giorgini, L. P. Pitaevskii, and S. Stringari, Reviews of Modern Physics **71**, 463–512 (1999).

[47] A. J. Leggett, Reviews of Modern Physics **73**, 307–356 (2001).

[48] M. Lewenstein and L. You, Physical Review Letters **77**, 3489–3493 (1996).

[49] Y. Castin and R. Dum, Physical Review A **57**, 3008–3021 (1998).

[50] M. R. Andrews, C. G. Townsend, H.-J. Miesner, D. S. Durfee, D. M. Kurn, and W. Ketterle, Science **275**, 637–641 (1997).

[51] Y. Castin and J. Dalibard, Physical Review A **55**, 4330–4337 (1997).

[52] B. V. Svistunov, E. S. Babaev, and N. V. Prokof'ev, *Superfluid States of Matter* (CRC Press, 2015).

[53] C. R. Cabrera, L. Tanzi, J. Sanz, B. Naylor, P. Thomas, P. Cheiney, and L. Tarruell, Science **359**, 301–304 (2018).

[54] M. R. Andrews, D. M. Kurn, H.-J. Miesner, D. S. Durfee, C. G. Townsend, S. Inouye, and W. Ketterle, Physical Review Letters **79**, 553–556 (1997).

[55] J. Ville, R. Saint-Jalm, E. Le Cerf, M. Aidelsburger, S. Nascimbène, J. Dalibard, and J. Beugnon, Physical Review Letters **121**, 145301 (2018).

[56] S. J. Garratt, C. Eigen, J. Zhang, P. Turzák, R. Lopes, R. P. Smith, Z. Hadzibabic, and N. Navon, Physical Review A **99**, 021601 (2019).

[57] R. Lopes, C. Eigen, N. Navon, D. Clément, R. P. Smith, and Z. Hadzibabic, Physical Review Letters **119**, 190404 (2017).

[58] T. D. Lee, K. Huang, and C. N. Yang, Physical Review **106**, 1135–1145 (1957).

[59] J.-P. Blaizot and G. Ripka, *Quantum Theory of Finite Systems* (MIT Press, Cambridge, Massachusetts, 1986).

[60] H. Takeuchi and K. Kasamatsu, Physical Review A **88**, 043612 (2013).

[61] D. C. Aveline, J. R. Williams, E. R. Elliott, C. Dutenhoffer, J. R. Kellogg, J. M. Kohel, N. E. Lay, K. Oudrhiri, R. F. Shotwell, N. Yu, and R. J. Thompson, Nature **582**, 193–197 (2020).

[62] S. Banik, *Analogue Cosmology Experiments with Sodium Bose-Einstein Condensates*, Ph.D. thesis, University of Maryland (2021).

[63] M. Gutierrez Galan, *Bose Einstein Condensates for Analogue Cosmology Experiments*, Ph.D. thesis, University of Maryland (2021).

[64] Y. Geng, S. Mukherjee, S. Banik, M. G. Galan, M. Anderson, H. Sosa-Martinez, S. Eckel, I. B. Spielman, and G. K. Campbell, Review of Scientific Instruments **96**, 093202 (2025).

[65] D. A. Steck, *Sodium D Line Data*, Alkali D Line Data (2025).

[66] H. J. Metcalf and P. Van Der Straten, *Laser Cooling and Trapping*, Graduate Texts in Contemporary Physics (Springer, New York, NY, 1999).

[67] M. Zhao, A. Restelli, J. Tao, Q. Liang, and I. B. Spielman, AIP Advances **13**, 10.1063/5.0138145 (2023).

[68] G. Lamporesi, S. Donadello, S. Serafini, and G. Ferrari, Review of Scientific Instruments **84**, 063102 (2013).

[69] W. Ketterle, K. B. Davis, M. A. Joffe, A. Martin, and D. E. Pritchard, Physical Review Letters **70**, 2253 (1993).

[70] P. D. Lett, R. N. Watts, C. I. Westbrook, W. D. Phillips, P. L. Gould, and H. J. Metcalf, Physical Review Letters **61**, 169–172 (1988).

[71] Y.-J. Lin, A. R. Perry, R. L. Compton, I. B. Spielman, and J. V. Porto, Physical Review A **79**, 063631 (2009).

[72] L. Chomaz, L. Corman, T. Yefsah, R. Desbuquois, and J. Dalibard, New Journal of Physics **14**, 055001 (2012).

[73] L. Corman, *The two-dimensional Bose Gas in box potentials*, Ph.D. thesis, Agence Bibliographique de l'Enseignement Supérieur (2017).

[74] Y. Castin and R. Dum, Physical Review Letters **77**, 5315–5319 (1996).

[75] 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, Scientific Reports **6**, 10.1038/srep25890 (2016).

[76] C. E. Rasmussen and C. K. I. Williams, *Gaussian processes for machine learning*, 3rd ed., Adaptive computation and machine learning (MIT Press, 2008).

[77] P. Zupancic, P. M. Preiss, R. Ma, A. Lukin, M. Eric Tai, M. Rispoli, R. Islam, and M. Greiner, Optics Express **24**, 13881 (2016).

[78] L. Feng, *Coherent Nonequilibrium Many-body Dynamics in Driven Bose Condensates*, Ph.D. thesis, University of Chicago (2019).

[79] A. Mazurenko, C. S. Chiu, G. Ji, M. F. Parsons, M. Kanász-Nagy, R. Schmidt, F. Grusdt, E. Demler, D. Greif, and M. Greiner, Nature **545**, 462–466 (2017).

[80] J. Goodman, *Introduction to Fourier Optics*, Electrical Engineering Series (McGraw-Hill, 1996).

[81] J. C. Wyant and K. Creath, in *Applied Optics and Optical Engineering*, Vol. XI (Academic Press, 1992) Chap. 1.

[82] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Atom-Photon Interactions: Basic Processes and Applications* (Wiley, 1998).

[83] A. Putra, D. L. Campbell, R. M. Price, S. De, and I. B. Spielman, Review of Scientific Instruments **85**, 10.1063/1.4862046 (2014).

[84] C.-L. Hung, X. Zhang, L.-C. Ha, S.-K. Tung, N. Gemelke, and C. Chin, New Journal of Physics **13**, 075019 (2011).

[85] E. Zaremba, Physical Review A **57**, 518–521 (1998).

[86] J. Zhang, C. Eigen, W. Zheng, J. A. Glidden, T. A. Hilker, S. J. Garratt, R. Lopes, N. R. Cooper, Z. Hadzibabic, and N. Navon, Physical Review Letters **126**, 10.1103/physrevlett.126.060402 (2021).

[87] K. Sasaki, N. Suzuki, D. Akamatsu, and H. Saito, Physical Review A **80**, 063611 (2009).

[88] G. Lamporesi, Two-component spin mixtures, in *Proceedings of the International School of Physics "Enrico Fermi", Course 211* (Società Italiana di Fisica, 2022).

[89] C. Samuelis, E. Tiesinga, T. Laue, M. Elbs, H. Knöckel, and E. Tiemann, Physical Review A **63**, 012710 (2000).

[90] S. Knoop, T. Schuster, R. Scelle, A. Trautmann, J. Appmeier, M. K. Oberthaler, E. Tiesinga, and E. Tiemann, Physical Review A **83**, 042704 (2011).

[91] L. D. Landau and E. M. Lifshitz, *Fluid Mechanics, Second Edition: Volume 6*, Course of Theoretical Physics, Vol. 6 (Pergamon Press, Oxford, 1987).

[92] D. Kobyakov, V. Bychkov, E. Lundh, A. Bezett, V. Akkerman, and M. Marklund, Physical Review A **83**, 043623 (2011).

[93] D. A. Takahashi, M. Kobayashi, and M. Nitta, Physical Review B **91**, 184501 (2015).

[94] Lord Rayleigh, Proceedings of the London Mathematical Society **s1-14**, 170 (1883).

[95] G. I. Taylor, Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences **201**, 192 (1950).

[96] A. Banerjee, Journal of Fluids Engineering **142**, 120801 (2020).

[97] J. A. F. Plateau, *Statique expérimentale et théorique des liquides soumis aux seules forces moléculaires*, Vol. 2 (Gauthier-Villars, 1873).

[98] Lord Rayleigh, The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science **34**, 177 (1892).

[99] C. J. Jog and P. M. Solomon, The Astrophysical Journal **276**, 114 (1984).

[100] J. P. Graves, I. T. Chapman, S. Coda, M. Lennholm, M. Albergante, and M. Jucker, Nature Communications **3**, 1 (2012).

[101] J. D. Kilkenny, S. G. Glendinning, S. W. Haan, B. A. Hammel, J. D. Lindl, D. Munro, B. A. Remington, S. V. Weber, J. P. Knauer, and C. P. Verdon, Physics of Plasmas **1**, 1379 (1994).

[102] A. J. Allen and P. A. Hughes, Monthly Notices of the Royal Astronomical Society **208**, 609 (1984).

[103] S. I. Abarzhi, A. K. Bhowmick, A. Naveh, A. Pandian, N. C. Swisher, R. F. Stellingwerf, and W. D. Arnett, Proceedings of the National Academy of Sciences **116**, 18184 (2018).

[104] V. N. Goncharov, Physical Review Letters **88**, 134502 (2002).

[105] H. Takeuchi, N. Suzuki, K. Kasamatsu, H. Saito, and M. Tsubota, Physical Review B **81**, 094517 (2010).

[106] D. Hernández-Rajkov, N. Grani, F. Scazza, G. Del Pace, W. J. Kwon, M. Inguscio, K. Xhani, C. Fort, M. Modugno, F. Marino, and G. Roati, Nature Physics **20**, 939 (2024).

[107] B. Mukherjee, A. Shaffer, P. B. Patel, Z. Yan, C. C. Wilson, V. Crépel, R. J. Fletcher, and M. Zwierlein, Nature **601**, 58 (2022).

[108] D. Kobyakov, V. Bychkov, E. Lundh, A. Bezett, and M. Marklund, Physical Review A **86**, 023614 (2012).

[109] B. Apffel, F. Novkoski, A. Eddi, and E. Fort, Nature **585**, 48 (2020).

[110] Z. Fifer, T. Torres, S. Erne, A. Avgoustidis, R. J. A. Hill, and S. Weinfurtner, Physical Review E **99**, 031101 (2019).

[111] V. S. Barroso, C. R. D. Bunney, and S. Weinfurtner, Journal of Physics: Conference Series **2531**, 012003 (2023).

[112] Y. Geng, S. Eckel, G. K. Campbell, and I. B. Spielman, Prethermal ripplons in quenched binary Bose-Einstein condensates (2026), arXiv:2609.09464 [cond-mat.quant-gas] .

[113] T. Mori, T. N. Ikeda, E. Kaminishi, and M. Ueda, Journal of Physics B: Atomic, Molecular and Optical Physics **51**, 112001 (2018).

[114] J. Berges, S. Borsányi, and C. Wetterich, Physical Review Letters **93**, 142002 (2004).

[115] L. Kofman, A. Linde, and A. A. Starobinsky, Physical Review Letters **73**, 3195–3198 (1994).

[116] E. Fermi, J. Pasta, S. Ulam, and M. Tsingou, *Studies of nonlinear problems. I, Technical Report LA-1940* (Los Alamos Scientific Laboratory, Los Alamos, 1955).

[117] G. Benettin, H. Christodoulidi, and A. Ponno, Journal of Statistical Physics **152**, 195 (2013).

[118] M. Onorato, L. Vozella, D. Proment, and Y. V. Lvov, Proceedings of the National Academy of Sciences **112**, 4208–4213 (2015).

[119] T. Kinoshita, T. Wenger, and D. S. Weiss, Nature **440**, 900–903 (2006).

[120] M. Gring, M. Kuhnert, T. Langen, T. Kitagawa, B. Rauer, M. Schreitl, I. Mazets, D. A. Smith, E. Demler, and J. Schmiedmayer, Science **337**, 1318–1322 (2012).

[121] T. Langen, S. Erne, R. Geiger, B. Rauer, T. Schweigler, M. Kuhnert, W. Rohringer, I. E. Mazets, T. Gasenzer, and J. Schmiedmayer, Science **348**, 207–211 (2015).

[122] Y. Tang, W. Kao, K.-Y. Li, S. Seo, K. Mallayya, M. Rigol, S. Gopalakrishnan, and B. L. Lev, Physical Review X **8**, 021030 (2018).

[123] B. Neyenhuis, J. Zhang, P. W. Hess, J. Smith, A. C. Lee, P. Richerme, Z.-X. Gong, A. V. Gorshkov, and C. Monroe, Science Advances **3**, 10.1126/sciadv.1700672 (2017).

[124] A. Rubio-Abadal, M. Ippoliti, S. Hollerith, D. Wei, J. Rui, S. L. Sondhi, V. Khemani, C. Gross, and I. Bloch, Physical Review X **10**, 021044 (2020).

[125] S. T. Beliaev, Soviet Physics JETP **7**, 299 (1958).

[126] J. Dereziński, B. Li, and M. Napiórkowski, Journal of Statistical Physics **191**, 110 (2024).

[127] F. Fang, R. Olf, S. Wu, H. Kadau, and D. M. Stamper-Kurn, Physical Review Letters **116**, 095301 (2016).

[128] S. Huh, K. Mukherjee, K. Kwon, J. Seo, J. Hur, S. I. Mistakidis, H. R. Sadeghpour, and J.-y. Choi, Nature Physics **20**, 402–408 (2024).

[129] S. R. Segal, Q. Diot, E. A. Cornell, A. A. Zozulya, and D. Z. Anderson, Physical Review A **81**, 053601 (2010).

[130] R. Dubessy, C. De Rossi, T. Badr, L. Longchambon, and H. Perrin, New Journal of Physics **16**, 122001 (2014).

[131] M. Gałka, P. Christodoulou, M. Gazo, A. Karailiev, N. Dogra, J. Schmitt, and Z. Hadzibabic, Physical Review Letters **129**, 190402 (2022).

[132] E. G. Flekkøy and D. H. Rothman, Physical Review Letters **75**, 260 (1995).

[133] D. G. A. L. Aarts, M. Schmidt, and H. N. W. Lekkerkerker, Science **304**, 847 (2004).

[134] R. Barnett, A. Polkovnikov, and M. Vengalattore, Physical Review A **84**, 023606 (2011).

[135] Z.-X. Gong and L.-M. Duan, New Journal of Physics **15**, 113051 (2013).

[136] C. Yin and A. Lucas, Physical Review Letters **131**, 050402 (2023).

[137] F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, and B. V. Saunders, *NIST Digital Library of Mathematical Functions* (2026), release 1.2.6 of 2026-03-15.

[138] A. Zenesini, A. Berti, R. Cominotti, C. Rogora, I. G. Moss, T. P. Billam, I. Carusotto, G. Lamporesi, A. Recati, and G. Ferrari, Nature Physics **20**, 558 (2024).

[139] M. Zhao, *Turbulence and Superfluidity in the Atomic Bose-Einstein Condensate*, Ph.D. thesis, University of Maryland (2024).

[140] J. Tao, *How Are Non-Hermitian Superfluids Special? Theory and Experiment*, Ph.D. thesis, University of Maryland (2024).

[141] C. Chin, R. Grimm, P. Julienne, and E. Tiesinga, Reviews of Modern Physics **82**, 1225 (2010).

[142] R. A. Barankov, Physical Review A **66**, 013612 (2002).

[143] C. Rackauckas and Q. Nie, Journal of Open Research Software **5**, 15 (2017).

[144] J. H. Verner, Numerical Algorithms **53**, 383–396 (2009).

[145] R. J. LeVeque, *Finite Difference Methods for Ordinary and Partial Differential Equations: Steady-State and Time-Dependent Problems* (Society for Industrial and Applied Mathematics, 2007).

[146] R. I. McLachlan and G. R. W. Quispel, Acta Numerica **11**, 341–434 (2002).

[147] H. T. Stoppels and L. Nyman, ArnoldiMethod.jl: Arnoldi method with Krylov-Schur restart, natively in Julia.

Artificial Intelligence Usage Disclosure

The author has used generative artificial intelligence tools (Claude Opus 5.5, Gemini 4, DeepSeek V4.1 Flash) for checking the correctness of physical arguments and equations, identifying grammatical errors, and providing language suggestions during the writing process.

The author has not used generative artificial intelligence tools for directly generating content for this dissertation or for analyzing data within this work.