

Phase transitions in quasiperiodic and driven optical lattices

Edward Alexander Carter

This thesis is submitted for the degree of
Doctor of Philosophy.

Downing College, University of Cambridge

November 2020



Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the preface and specified in the text.

It is not substantially the same as any work that has already been submitted before for any degree or other qualification except as declared in the preface and specified in the text.

It does not exceed the 60,000 word limit prescribed by the Faculty of Physics and Chemistry, including abstract, tables and footnotes and appendices.



This thesis is dedicated to those who read it.



Abstract

Phase transitions in quasiperiodic and driven optical lattices

Edward Alexander Carter

Ultracold atoms in optical lattices are a versatile tool for precisely-controlled quantum simulation of a range of condensed matter phenomena. This thesis describes work to apply this tool in two areas: the simulation of quasicrystalline materials with a novel lattice geometry, and the use of Floquet driving to engineer behaviours qualitatively different from static systems. In particular I will present our work using periodic driving to make a continuous quantum phase transition become discontinuous, which to our knowledge has never been done before.

Quasicrystals are a fascinating but still relatively under-explored class of materials existing as an intermediate between periodic and disordered systems. This makes them an ideal context for studying non-ergodic behaviour, especially with the tunable interactions allowed by ultracold atom physics. At the same time quasicrystals have an intriguing link to higher dimensions, so that a two-dimensional quasicrystal can be used to simulate systems with more than three spatial dimensions. I will describe two experiments performed by our group that use a quasicrystalline optical lattice to explore higher dimensions and non-ergodic states respectively, which we hope are a stepping stone towards future work with many-body localisation and higher-dimensional topological effects.

The other part of the story is Floquet physics: the study of time-periodic Hamiltonians. Time dependence allows us to break many of the usual rules of quantum systems, adding a whole new set of control parameters and allowing qualitatively different behaviours. In this thesis I will discuss the application of Floquet driving to a periodic optical lattice and present our experimental observation of a discontinuous form of the well-known Mott insulator to superfluid quantum phase transition. While interesting in itself, we view this work too as a stepping stone towards Floquet engineering with the full optical quasicrystal.

Contents

Introduction	12
1 Phase transitions	15
1.1 What is a phase?	15
1.2 What is a phase transition?	18
1.3 Classification of phase transitions	20
1.3.1 Discontinuous phase transitions	20
1.3.2 Continuous phase transitions	22
1.4 Quantum phase transitions	26
1.4.1 A quantum picture of phase transitions	26
1.4.2 Dynamics of quantum phase transitions	27
2 Quasicrystals and localisation	29
2.1 Physics of periodic systems	29
2.1.1 Quasimomentum	29
2.1.2 Bloch waves	31
2.1.3 Tight binding	32
2.1.4 Semiclassical dynamics of particles	33
2.1.5 The Bose-Hubbard model	35
2.1.6 Adding random disorder	37
2.2 Quasicrystals	39
2.2.1 What is a quasicrystal?	39

2.2.2	Incommensurable length scales in quasiperiodic systems	40
2.2.3	Links to higher dimensions	41
2.2.4	Between order and disorder	43
2.2.5	Localisation in quasicrystals	45
3	An optical quasicrystal for ultracold atoms	46
3.1	Cooling down atoms	46
3.1.1	The light-atom interaction	47
3.1.2	Early cooling stages: spontaneity	47
3.1.3	Late cooling stages: coherence	48
3.2	The BEC transition	49
3.2.1	Statistical mechanics	49
3.2.2	Imaging the transition	53
3.3	Interactions	54
3.3.1	The scattering length	54
3.3.2	Feshbach resonances	54
3.3.3	Quantum depletion	56
3.4	Summary of experimental sequence	56
3.5	Quantum simulation	58
3.5.1	Optical lattices	58
3.5.2	Time of flight imaging in an optical lattice	61
3.5.3	Kapitza-Dirac diffraction	62
3.6	The optical quasicrystal	65
3.6.1	Four dimensional quantum walks	66
3.6.2	Localisation experiments	69
4	The Mott insulator-superfluid transition	74
4.1	The superfluid phase	74

4.1.1	Superfluidity in the weakly interacting Bose gas	74
4.1.2	Properties of superfluids	78
4.1.3	Superfluids in lattices	79
4.1.4	Imaging superfluids	80
4.2	The Mott insulator	80
4.2.1	Qualitative description and general features	80
4.2.2	Incoherence	82
4.2.3	Incompressibility	82
4.2.4	Imaging the Mott insulator	83
4.3	The Mott insulator-superfluid transition	83
4.3.1	The mean-field approach	84
4.3.2	Mean-field treatment of the Bose-Hubbard Hamiltonian	85
4.3.3	Imaging the transition	87
5	Floquet physics	88
5.1	Driven systems	88
5.2	Floquet theory	89
5.2.1	Hopping renormalisation in a single band	90
5.2.2	Floquet states and quasienergies	92
5.3	Driving an optical lattice	95
5.3.1	The Hamiltonian	95
5.3.2	Space-time Bloch waves	96
5.4	Using Floquet theory	97
5.4.1	Floquet heating	97
5.4.2	Single-band Floquet physics	99
5.4.3	Band engineering	100
6	The shaken Mott insulator experiment	101

6.1	Theory of the shaken Mott insulator	101
6.1.1	Motivation	101
6.1.2	Calculated phase diagrams	104
6.2	Technical implementation	104
6.2.1	Adiabatic loading of an optical lattice	105
6.2.2	Methods of lattice shaking	106
6.3	The experiment	107
6.3.1	Preliminary work and characterisation	107
6.3.2	One dimension	110
6.3.3	Two dimensions	113
7	Future directions	115
7.1	Technical additions	115
7.1.1	Anticonfinement compensation	115
7.1.2	DMDs and designer potentials	117
7.2	Experimental directions	119
7.2.1	The disordered Bose-Hubbard phase diagram	119
7.2.2	Floquet driving of an optical quasicrystal	122
7.2.3	Topology	122
	Summary	125
	Acknowledgements	126
	A Dimension scaling in quantum walks	128
	B Theorem regarding the monodromy operator	129
	C The shaken lattice Hamiltonian	131
	References	133

Introduction

Ultracold atoms

This thesis describes quantum simulation experiments performed using quantum degenerate or "ultracold" clouds of ^{87}Rb and ^{39}K atoms loaded into an optical lattice. As such it is part of a rapidly evolving field with an extensive history, some idea of which I will try to sketch here.

Quantum simulation refers to the practice of studying interesting but computationally intractable quantum systems by building simpler quantum systems that share their essential features and studying those instead [1–3]. This requires two things: an interesting system to simulate, and an experimental platform with which to do so. Several such platforms exist besides ultracold atoms, including trapped ions [4], photons [5], photonic crystals [6], and increasingly ultracold molecules [7, 8], between them creating a diverse range of precisely controlled quantum simulators. However, in this thesis we will focus exclusively on ultracold atom experiments.

The field of experimental ultracold atom physics began with the first observations of Bose-Einstein condensation in 1995 [9, 10] following earlier breakthroughs in laser cooling [11], with the first ultracold atomic Fermi gas coming soon after in 1999 [12]. Since then ultracold atoms have found numerous applications. For example, atomic clouds can be used as instruments for precision measurements – often involving interferometry [13] – in fields as diverse as time-keeping [14] and gravitational wave detection [15, 16].

There is also a great deal of rich and interesting physics to be studied by investigating ultracold atomic gases for their own sake, rather than as simulators. For example, they have been used to study quantum turbulence [17], construct "atomtronic" analogues of electric circuits [18], and realise novel supersolid phases [19, 20]. They are also an ideal platform for observing the Kibble-Zurek mechanism [21], the BEC-BCS crossover in fermions [22], quantum droplet formation [23, 24], and the BKT crossover [25].

However, we are interested specifically for their use in quantum simulation. As noted above, once we have chosen an experimental platform, we then need an interesting system to simulate. There are in fact many possible choices, some quite unexpected: ultracold atomic systems have been used to simulate Hawking radiation from black holes [26], the dynamical Casimir effect [27], and the anisotropy of the cosmic microwave background [28].

We will restrict ourselves, however, to a more conventional choice: quantum simulation of condensed matter systems, or alternatively many-body quantum dynamics (and statics). These experiments commonly involve *optical lattices*, which are laser interference patterns that serve as potential surfaces for ultracold atoms. This allows us to simulate the behaviour of electrons moving through the potential surface of a material, though on a much larger scale and thousands of times more slowly. In the case of quantum gas microscopes, such experiments can even be conducted with single-site resolved imaging [29, 30]. An enormous variety of many-body quantum phenomena have been observed with optical lattices already, with highlights including magnetism [31], transport phenomena [32], and most significantly for us the quantum phase transition between the Mott insulator and superfluid states [33]. Disordered many-body systems are also a ripe area of experimentation, with effects of interest including Anderson [34–36] and many-body localisation [37–39], as well as the Bose glass phase [40–42]. I will also mention the potential of optical lattices to study topologically non-trivial condensed matter systems, often via synthetic gauge fields [43, 44], with both the Hofstadter [45] and Haldane [46] models having been realised.

Quasicrystals

One area of condensed matter physics where quantum simulation shows great promise is the study of quasicrystals. Quasicrystals are ordered aperiodic solids [47, 48] proposed as a mathematical curiosity by Penrose [49] and discovered in real materials by Shechtman [50]. Their links to topological states [51] and higher-dimensional physics [52] are strong incentives to understand them better, and by various means. Quasicrystalline systems have been studied experimentally in fields as diverse as photonics [6] and colloid physics [53, 54], and while the first work with atoms in quasicrystalline potentials dates back to 1997 [55, 56], it was our own group that built the first quasicrystalline optical lattice for ultracold atoms [57, 58].

Floquet physics

The tools of ultracold atom physics, especially relating to quantum simulation, are greatly expanded by the use of Floquet physics. This refers to the study of periodically-driven quantum systems, which has extensive analogies to spatially periodic systems and offers a similar degree of richness and complexity [59–62]. A wide range of work has already been done using Floquet physics to simulate condensed matter effects [63–65] and also more surprising systems such as black holes [66]. A particularly powerful tool, and of central importance to the second half of this thesis, is Floquet band engineering [67], which opens up new frontiers in systems that are periodic in both space and time. I particularly wish to emphasise the proposal [68] for using band engineering to tune a continuous quantum phase transition into becoming discontinuous, the first experimental realisation of which will be described in this thesis.

Layout of this thesis

This thesis has two focuses. The first is our experimental work on ultracold atoms in an optical quasicrystal, which has previously been published in [57, 58] and is described extensively in the theses of my colleagues Matteo Sbroscia [69] and Konrad Viebahn [70], but which I will describe again here, more briefly and within the context of ongoing work. The other is the use of Floquet band engineering to realise a discontinuous form of the Mott insulator-superfluid quantum phase transition, as proposed in [68]. With these goals in mind, chapter 1 lays out the physics of classical and quantum phase transitions, while chapter 2 describes quasicrystals and contrasts them with the periodic systems more familiar in solid-state physics. Chapter 3 then explains how we use ultracold atoms to investigate quasicrystals, culminating in the experimental results from [57, 58]. Chapters 4 and 5 describe the Mott insulator to superfluid transition and the physics of Floquet band engineering respectively, laying the basis for chapter 6 which presents our realisation of a discontinuous form of the transition. Finally, chapter 7 presents our plans to build on these results, especially regarding the potential for probing the phases of interacting particles in the optical quasicrystal and ultimately adding Floquet shaking to this system as well.

1

Phase transitions

1.1 What is a phase?

A thermodynamic phase¹ may be defined as a system throughout which all physical properties are uniform, ignoring variations on the smallest scales [71]. Mathematically this corresponds to moving from atomic degrees of freedom (such as the spin of a single particle) to slowly-varying hydrodynamic fields that represent the collective behaviour of many interacting particles (such as the magnetisation, which averages over many spins). For example, a single crystal of silicon is a phase, even though we would see variations in the density if we sampled on scales smaller than the interatomic spacing. A mixture of water and ethanol is also a phase, because the liquids are completely miscible down to the molecular level. A glass of water containing ice cubes, however, is a mixture of two phases, because the variations in density and structure are across "macroscopic" scales far bigger than the size of a water molecule.

This definition is simple and consistent with everyday intuition. However, there are many systems where intuition is a poor guide and a more rigorous approach is needed. To see this, consider the question: how can we define whether two phases are different from each other (say water and ice), or two instances of the same phase (say two ice cubes)?

¹Not to be confused with a phase in the sense of an oscillator's phase, represented by a factor $e^{i\phi}$. In a thesis dealing with both thermodynamics and quantum physics there is unfortunately no avoiding this ambiguity.

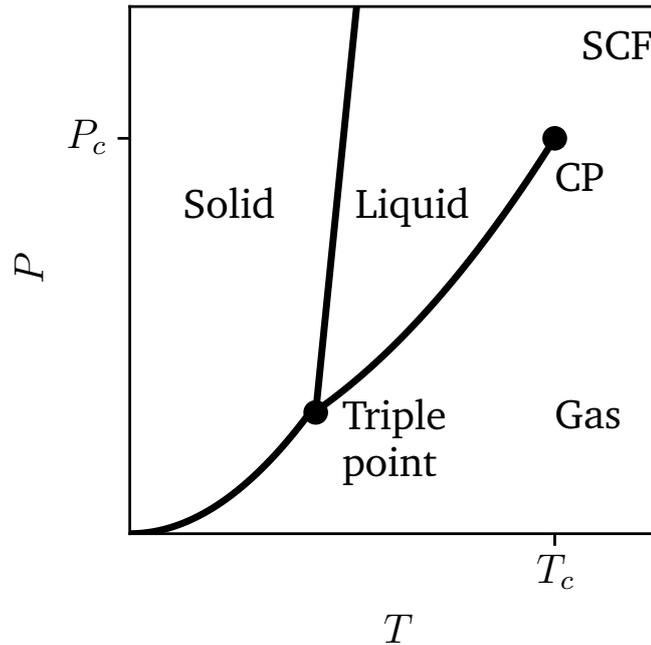


Figure 1.1: The solid-liquid-gas phase diagram, showing the qualitative regions of pressure P and temperature T that result in each phase. The solid-liquid and solid-gas transitions are always discontinuous (section 1.3.1), but at the critical point (T_c, P_c) , labelled CP, the liquid-gas transition is continuous (section 1.3.2). In the supercritical region above this ($T > T_c$ and $P > P_c$) the liquid and gas phases are no longer distinguishable and there is no transition, which is possible because the liquid and gas phases have no difference in fundamental symmetry (section 1.1). The phase observed in this region is called a supercritical fluid (SCF). The triple point is the only point in phase space where all three regions meet, resulting in simultaneous coexistence of liquid, solid and gas.

Distinguishing phases by physical properties

In some cases phases are "obviously" different on the basis of their physical properties. This is true for the three traditional states of matter: solids are distinguished from liquids by their structural rigidity, and gases are distinguished from both by their compressibility and low density. The phase diagram showing the qualitative regimes of pressure and temperature that result in each of these phases is given in figure 1.1.

However, this method of distinguishing phases is too naïve to work in all cases, because physical properties can vary widely for many reasons. A piece of metal at room temperature is "obviously" different from a red-hot piece of metal in many respects (strength, conductivity, brightness etc.) but they are the same phase. A quantum example is the BEC-BCS crossover in ultracold fermions [72], wherein the physical model used to describe the system changes drastically but gradually, so that we cannot delineate two distinct phases.

Distinguishing phases by symmetry

Phases may often be distinguished by some difference in fundamental symmetry. For example, the liquid and gas phases both have continuous translational symmetry, which is broken in the solid phase to become discrete. (The implications of this for the physics of solids are discussed in section 2.1.) Bose-Einstein condensation, discussed in section 3.2, breaks the gauge symmetry of the wavefunction's quantum phase in a similar way; this symmetry is also broken in the Mott insulator-superfluid transition of chapter 4. A process where a system moves from a phase with some symmetry (the *disordered phase*) to a phase without (the *ordered phase*) is called *spontaneous symmetry breaking*², and is discussed below in section 1.3.2.

This criterion is dependable when it applies: if two states have different symmetry, they are sure to be distinct phases. However, we have already met a case, the liquid-gas transition, where the two phases have no difference in underlying symmetry. This demonstrates that symmetry breaking, while a sufficient condition for distinguishing phases, is not a necessary one.

Distinguishing phases by the presence of a phase transition

A third possible definition is that the distinctness of two phases depends on what happens as we move between them. If tuning between two states is a smooth process with no sudden changes at any value of our control parameter, they are not distinct phases. If the two states have very different physical properties, we call this a *crossover*. On the other hand, a *phase transition* (PT) is a non-analytic change in the properties of a system resulting from a gradual change in some parameter (often temperature). For example, the properties of liquid water at atmospheric pressure change slowly and smoothly at all temperatures between 0 °C and 100 °C, but as soon as we reach either of those limits there is a large, abrupt jump in the density as we obtain ice or water vapour. At both of these temperatures the system undergoes a phase transition. We could say that two states are distinct phases only if moving between them involves a PT.

However, this definition is once again not comprehensive. Consider again the liquid and gas phases of figure 1.1. At temperatures below T_c and pressures below P_c the system undergoes a PT when tuning between those phases, but by passing through the supercritical regime we can move from a liquid to a gas without any non-analytic behaviour at any point. We can do the same thing with an Ising ferromagnet [73] by applying a strong external magnetic field aligned with the spins, heating above T_c , reversing the applied field, and then cooling down again. This is generically possible whenever the initial and final states have no difference in underlying symmetry.

I believe the unsatisfactory nature of all three of these criteria demonstrates that the

²Strictly, the symmetry breaking is "spontaneous" only if nothing in the environment leads to a particular choice of direction. For example, the ferromagnetic phase transition shows spontaneous symmetry breaking only if there is no external magnetic field. However, in this thesis we will largely consider transitions where there is no practical way of imposing a direction.

concept of a phase has no rigorous definition. Instead there are a number of overlapping intuitive ideas of what constitutes a phase, each of which breaks down in different circumstances. While this may be disappointing, to dwell on it is to miss a far more productive area of study: the nature of PTs themselves. Despite the vague nature of phases, the transitions between them can be described in very precise terms.

1.2 What is a phase transition?

Much like phases themselves, PTs have historically been considered in several different ways across different fields. Fortunately these definitions end up being fairly compatible in most cases, with the result that PTs are well-defined under ideal circumstances. For now this discussion will be entirely classical, with quantum descriptions of PTs left to section 1.4.

The order parameter

During our discussion of symmetry I mentioned that we often define the two phases on either side of a PT as the "ordered" and "disordered" phases. We can make this idea quantitative by introducing a (generally complex) quantity called the *order parameter* ψ , which vanishes in the disordered phase and has some finite value in the ordered phase [71]. For example, at the paramagnetic-ferromagnetic PT the order parameter is the magnetisation in the absence of an external field.

It may sometimes be necessary to choose a rather contrived order parameter in order to meet these requirements. For example, at the liquid-gas phase transition the physical quantity that most noticeably changes is the density, ρ , which leaps from the liquid value ρ_l to the gaseous value ρ_g , but is not zero in either phase. We can therefore define the order parameter as $\psi := \rho - \rho_g$, so that it is zero in the gas and positive in the liquid³.

Classic thermodynamic treatment

The starting point for PTs in classical thermodynamics is the partition function [74], defined for a system with Hamiltonian \hat{H} and eigenenergies E_i ⁴ as

$$Z := \text{Tr} \left[e^{-\frac{\hat{H}}{k_B T}} \right] = \sum_i e^{-\frac{E_i}{k_B T}} \quad (1.1)$$

³This designates the liquid as the ordered phase. Our previous definition of the ordered phase was based on symmetry and so doesn't apply here, but since liquids are generally more structured than gases due to the increased interactions between particles, it feels intuitive to define the order parameter this way around.

⁴One can also stick to purely classical terms with an integral over phase space, but I have sacrificed historical accuracy for the sake of clarity.

where T is the temperature and k_B is Boltzmann's constant. At thermodynamic equilibrium Z is approximately equal to the number of eigenstates that are significantly occupied. However, its more important use is the derivation of various *thermodynamic potentials*, for example the *Helmholtz free energy* given by

$$F = -k_B T \ln(Z). \quad (1.2)$$

For a system with fixed temperature and volume, it can be shown that the minimal value of F , F_{\min} , corresponds to the maximum value of the Boltzmann entropy S^5 . By the Second Law of thermodynamics S takes its maximum value, and so F its minimum, at thermal equilibrium. A PT therefore occurs when the value F_{\min} changes non-analytically, such as at a singularity or a cusp, and in classical thermodynamics we *define* these non-analytic points to be PTs. This behaviour is related back to the order parameter ψ by the Ginzburg-Landau theory of sections 1.3.1 and 1.3.2.

Spontaneous symmetry breaking

It was noted in section 1.1 that often different phases will exhibit different symmetries. If this is the case, the symmetry of the system must change at the PT (and can't change anywhere else). This is intrinsically a sudden, discontinuous process because there is no middle ground: either the system has a symmetry or it doesn't, and moving between the two requires spontaneous symmetry breaking.

For example, consider the liquid-to-solid PT, which as noted above involves a change from continuous to discrete translational symmetry. This requires the solid phase to "choose" a privileged point to become the origin for a periodic lattice. Placing one lattice site at this point determines the positions of all others, but the initial choice is completely arbitrary: indeed, if the system is large and the phase transition is crossed quickly, different choices may be made in different regions, which is called the Kibble-Zurek mechanism [75, 76]. The boundaries between these regions are called *domain walls* and generally persist as excitations, while the size of the domains is determined by the speed of crossing the transition⁶. Another example of spontaneous symmetry breaking is the Bose-Einstein condensation PT, which requires the condensate to pick a value for its overall phase from the interval $[0, 2\pi)$.

In terms of the order parameter, the phase with higher symmetry is always the disordered phase with $\psi = 0$. (To see this, consider that total randomness has every possible symmetry.)

⁵ $S = k_B \ln W$ for a system with W macroscopically indistinguishable configurations. S and F are then related by $F = U - TS$, where U is internal energy and T temperature. This relation and other similar ones allow us to calculate all the classical thermodynamic properties of a system starting from equation 1.2.

⁶This is ultimately a relativistic argument: the maximum possible domain size is the distance information can travel at the speed of light, or equivalent maximum speed in a material, during the time of crossing.

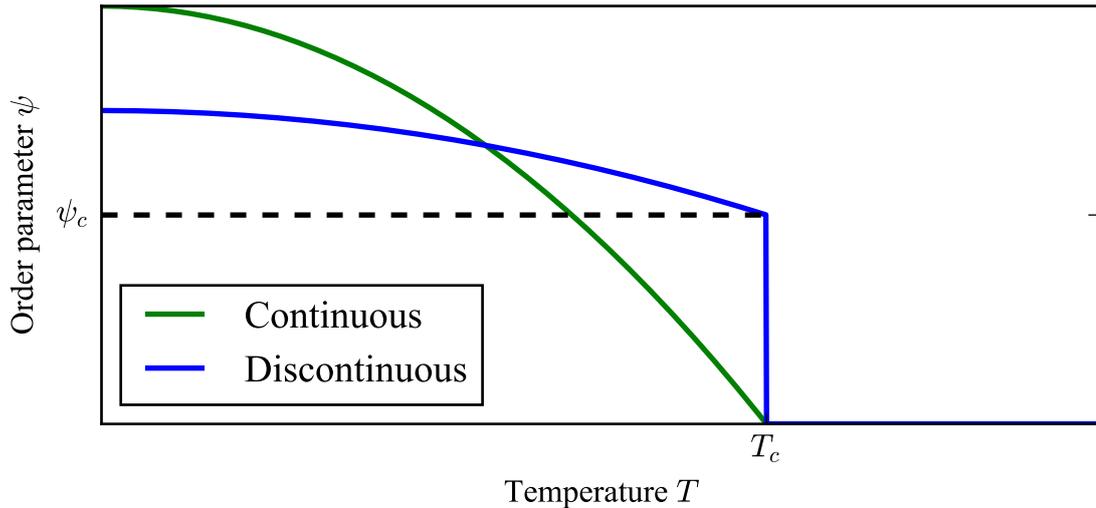


Figure 1.2: A graph sketching the generic behaviour of the order parameter ψ at discontinuous and continuous phase transitions (PTs). Both transitions are assumed to occur at a temperature T_c , above which the order parameter is 0. The discontinuous PT shows a jump in ψ at T_c up to some value ψ_c , while the continuous transition has a discontinuity only in the gradient.

1.3 Classification of phase transitions

PTs are usually classified into two categories: in *discontinuous* PTs there is a discontinuous change in the value of ψ , while in *continuous* PTs the discontinuity is in one of the derivatives of ψ ⁷. Examples of these behaviours are sketched in figure 1.2. We will examine each case in turn; the reader is also directed to [73] for an alternative account.

1.3.1 Discontinuous phase transitions

Characteristic features

Discontinuous PTs involve a sudden change in the value of the order parameter ψ (figure 1.2). For clarity, let us call the ordered phase A (where $\psi \neq 0$) and the disordered phase B (where $\psi = 0$). A and B will in general have different free energies because of their different values of ψ , which means transforming between them will either absorb or release energy. This energy is called the *latent heat* of the transition, and the need for it to be supplied or released via the system's edges typically means the

⁷These are sometimes called first- and second-order PTs respectively. This is because ψ is often related to the derivative of the free energy F , so that $\partial_r \psi = \partial_r^2 F$. This terminology is suggestive of the existence of PTs of third order or greater, with the discontinuity in $\partial_r^3 F$ or a higher derivative, and such phenomena do indeed exist [77, 78]. However, their properties are typically very similar to those of second-order PTs (see e.g. [79]). For the purposes of this thesis we will consider PTs of second and all higher orders to be included within the "continuous" category.

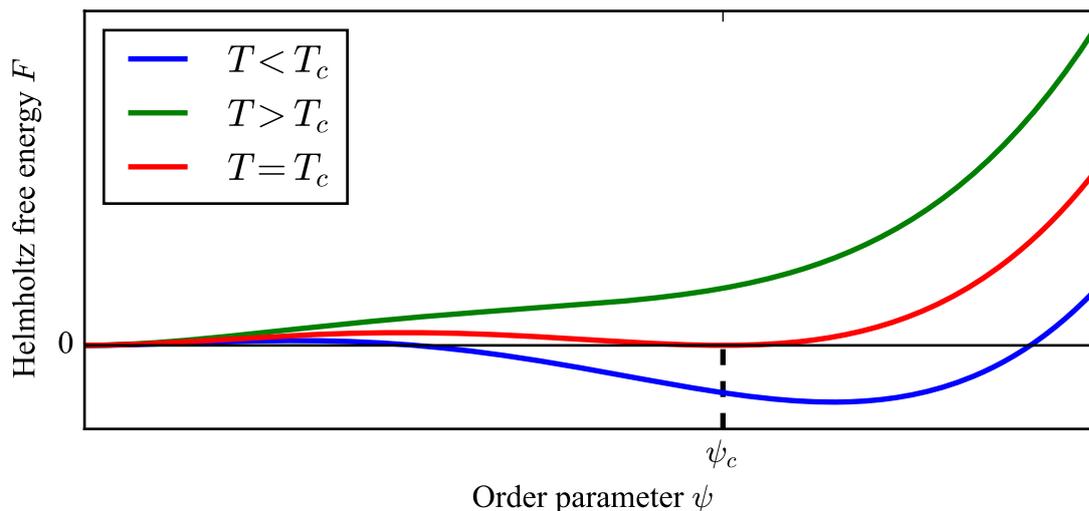


Figure 1.3: The Ginzburg-Landau polynomial from equation 1.4 plotted for three different values of $c_2(T)$. $F(\psi)$ always displays a minimum at $\psi = 0$, but at some particular $T = T_c$ the minimum at finite ψ becomes just as favourable (phase coexistence). Below this value $\psi = 0$ is no longer the global minimum and the phase transition has occurred. The label ψ_c refers to figure 1.2, and shows the value that ψ will jump to at the transition.

PT takes some amount of time, during which the system sits at the phase boundary. This may be observed as *phase coexistence*, meaning that A and B are each present in clearly distinct regions simply because not all the A has had time to change into B (or vice-versa). Note that this implies a very inhomogeneous system at the phase boundary.

As an example, consider the liquid-to-gas phase transition observed when boiling water at atmospheric pressure. The order parameter is the density, which is much higher in the liquid than in the gas. Once the water reaches 100°C the temperature will stop climbing, as the energy of the heating element is instead going towards the latent heat needed to break hydrogen bonds in the liquid. Bubbles will form within the water, demonstrating phase coexistence. No part of the system ever has density intermediate between the liquid and solid phases: when a bubble forms the local density jumps discontinuously to the equilibrium density of the gas at ambient pressure, resulting in a turbulent swirl of bubbles and liquid water.

Ginzburg-Landau phenomenology

These observations may be made quantitative with a phenomenological model due to Ginzburg and Landau. Suppose we wish to write the free energy of equation 1.2 as a function of the order parameter. In general this may be very complicated, but we can reasonably suppose that there is a Taylor expansion

$$F(T, \psi) = \sum_j c_j(T) \psi^j. \quad (1.3)$$

If ψ is small we can truncate this expansion. The minimal model of a discontinuous phase transition in this model occurs for the truncation

$$F(T, \psi) \approx c_2(T)\psi^2 - c_3\psi^3 + c_4\psi^4, \quad (1.4)$$

where I have taken the coefficients c_j to be real and positive. Let us also assume for simplicity that $c_2(T)$ increases monotonically with temperature. An F of this form may have two minima, as shown in figure 1.3, and their relative heights will change as the value of c_2 changes. (Recall that the global minimum of F is the equilibrium state of the system.) At high T the global minimum of F is at $\psi = 0$. As we cool below some temperature, the other minimum comes down to $F = 0$ and then below, so that the equilibrium value of ψ jumps discontinuously from 0 to a finite value. Phase coexistence occurs when the two minima are both exactly at $F = 0$, and the latent heat is the energy to jump from one minimum to another. Note that this discontinuous jump occurs without any non-analytic behaviour in the function $F(\psi)$.

Despite ignoring all the microscopic features of the system, this minimal phenomenological model gives a rationalisation for how discontinuous PTs can arise from a very generic form of F .

1.3.2 Continuous phase transitions

The critical point

Unlike discontinuous PTs, for which the two phases A ($\psi \geq 0$) and B ($\psi = 0$) can always be clearly distinguished, for a continuous PT A and B must become indistinguishable at the point where the transition occurs. This is because the value of ψ cannot jump, and so $\psi = 0$ for both A and B when $T = T_c$. This is illustrated in figure 1.2. It is only when we move away, and ψ grows for A but stays at 0 for B, that we can easily distinguish the phases. This has a number of implications. First of all there is no latent heat, because there is no sudden change in structure. Similarly it makes no sense to speak of phase coexistence because A and B are indistinguishable.

To put this in precise terms, we call the point in phase space where the PT occurs a *critical point*, and close to this point we observe a *critical phase* that cannot be clearly identified as either A or B. The characteristic feature of the critical phase is that thermal fluctuations in the system become very important, and in particular the correlation length of thermal fluctuations, ξ , diverges according to some power law

$$\xi \propto |t|^{-\nu} \quad (1.5)$$

where t is a dimensionless measure of temperature⁸, $t := |T - T_c|/T_c$, and ν is referred to as a *critical exponent*. This divergence indicates that thermal fluctuations of the system become global, and it is this that destroys the order present in A and results in the disordered phase B, with ψ decreasing to 0 as ξ diverges. Many other physical

⁸More generally, of distance from the critical point. If the transition is driven not by temperature, but by some other parameter r , we simply define t as $|r - r_c|/r_c$ and the rest of the discussion is unchanged.

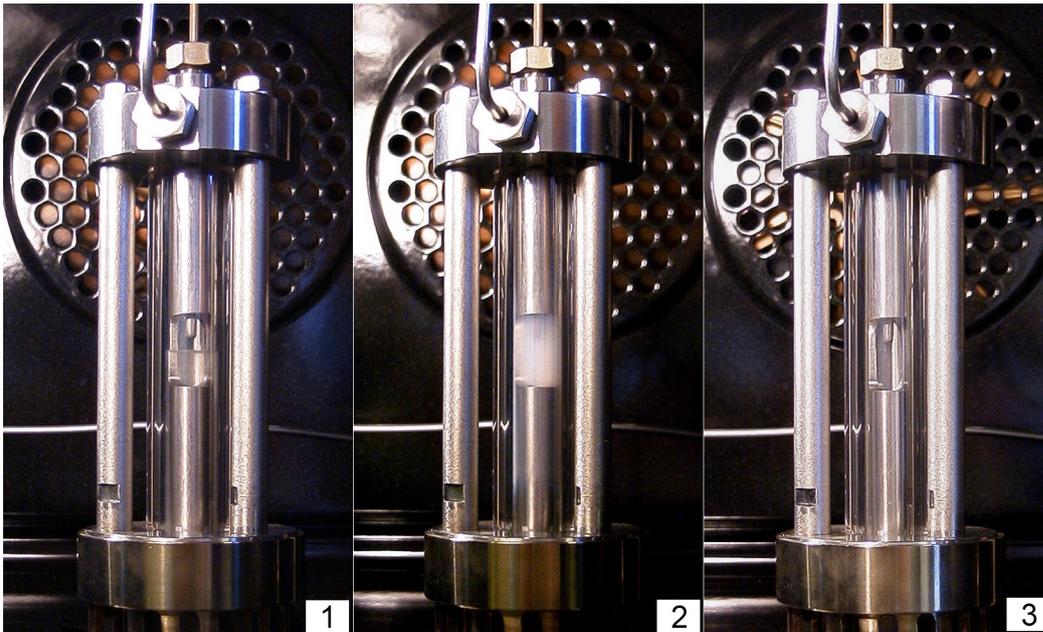


Figure 1.4: A demonstration of critical opalescence in ethane, which has a phase diagram much like figure 1.1. Panel 1 shows the liquid and gas phases coexisting with a visible boundary. Panel 2 is taken at the critical point of 32.17°C , 4.87 MPa , where the liquid-gas mixture becomes a single opaque fluid due to macroscopic density fluctuations. Finally, panel 3 is taken in the supercritical region above the critical point where the liquid and gas phases are indistinguishable, so there is no opalescence but also no visible boundary. This figure is taken from [81].

quantities also diverge at the critical point (e.g. susceptibilities), and each has its own critical exponent.

An important consequence of the diverging correlation length is *universality*. The critical phase is insensitive to microscopic details because global fluctuations result in averaging of physical properties over the entire system. If we know only the symmetry of the order parameter and the number of spatial dimensions we can calculate the critical exponents, and these in turn determine everything about the critical phase [80]. Systems with the same critical exponents are said to belong to the same *universality class*, allowing their properties near criticality to be found by studying a minimal model.

An example is the liquid-to-gas transition at the critical point (figure 1.1), the only point on the solid-liquid-gas phase diagram with a continuous transition. Near this point we can observe critical opalescence (figure 1.4), where the fluid becomes opaque even if both the gas and the liquid are transparent. This is because of density fluctuations on scales comparable to the wavelength of visible light: a visible manifestation of the diverging correlation length.

Another continuous PT is the BEC transition (figure 1.5). The order parameter in this case is the occupation of the single-particle ground state⁹, N_0 , which grows smoothly

⁹Strictly this is true for a noninteracting BEC; if there are repulsive interactions we instead use the occupation of a modified ground state derived from mean-field theory. This is discussed a little more

from 0 above a critical value of the phase-space density. In this case T_c depends on the number density n , so that we can choose to define a critical temperature at fixed density or a critical density n_c at fixed temperature. The BEC is the ordered phase because it breaks global phase symmetry by picking a macroscopic phase for the wavefunction.

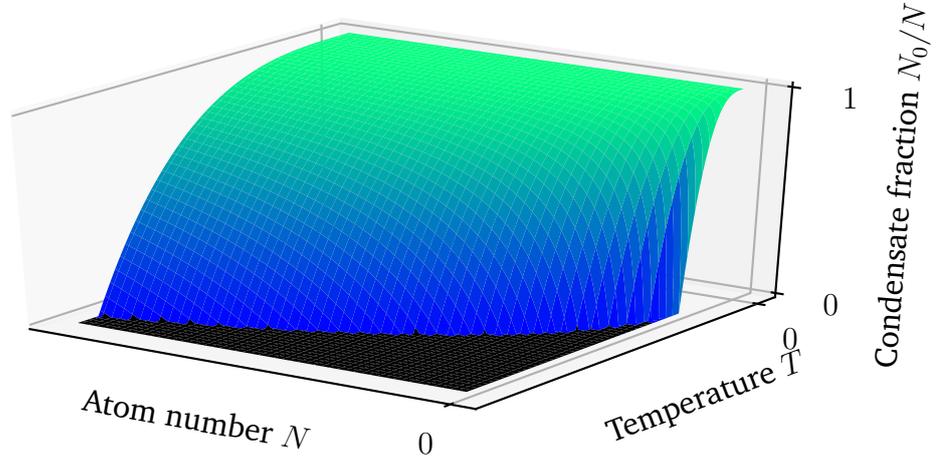


Figure 1.5: A phase diagram for the BEC phase transition (PT) in a three-dimensional harmonic trap. This is an example of a continuous (section 1.3.2) quantum (section 1.4) PT, and the order parameter for the transition (see section 1.2) is the condensate fraction N_0/N , which grows from 0 (black) below a temperature $T_c(N) \propto N^{\frac{1}{3}}$. Comparing with figure 1.1, the phase boundary here marks a much less abrupt change, with no discontinuity in the order parameter. In addition, unlike the continuous liquid-to-gas transition at point C in figure 1.1, there is no way to go from the classical gas to the BEC without a PT, because the two phases have different fundamental symmetry (the BEC breaks a gauge symmetry by picking a macroscopic phase). This PT is discussed in more detail in section 3.2.

Ginzburg-Landau phenomenology

We can once again apply the empirical method of Ginzburg and Landau to visualise a continuous phase transition. This time the minimal model requires that we omit the third-order term in the Taylor expansion, obtaining

$$F(T, \psi) \approx c_2(T)\psi^2 + c_4\psi^4. \quad (1.6)$$

We then suppose that c_2 is negative below some T_c , which we call the critical temperature, and positive above. As shown in 1.6, when c_2 is positive the minimum of F is always at $\psi = 0$. However, when c_2 becomes negative there are instead two minima at $\psi = \pm\sqrt{-c_2/2c_4}$ which move smoothly away from 0 as c_2 becomes more negative. This captures both the continuous change in ψ and the phenomenon of spontaneous symmetry breaking, as the system must choose which minimum to occupy (assuming ψ is a quantity that is able to be negative).

in section 4.1.1, but for the sake of the present discussion the transition is the same.

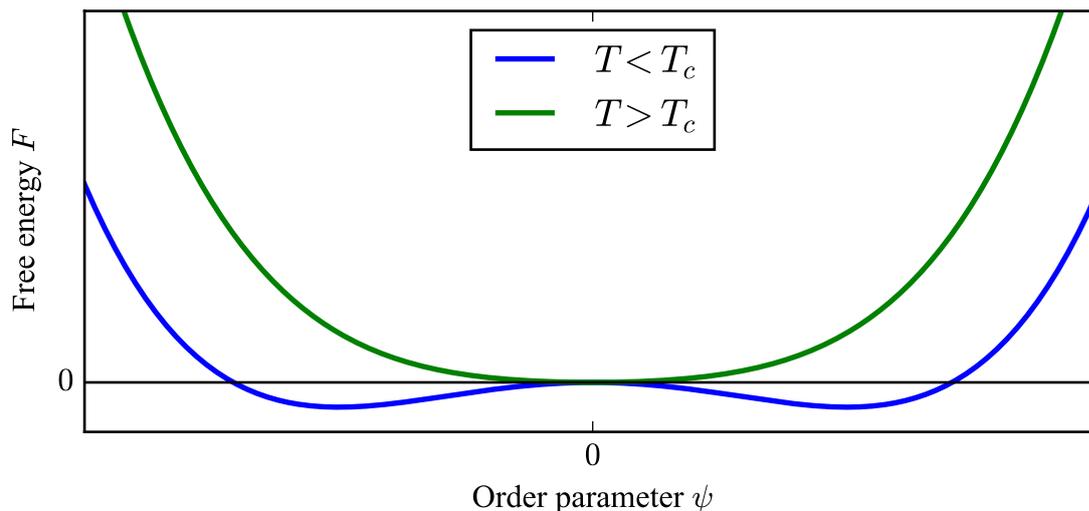


Figure 1.6: The Ginzburg-Landau polynomial for continuous PTs, given in equation 1.6, plotted above and below the transition by changing the sign of c_2 . Below the transition there are two symmetric global minima at equal $|\psi|$, causing spontaneous breaking of the symmetry between $\pm\psi$ as the system seeks to minimise F .

The thermodynamic limit

Comparing figures 1.6 and 1.2 raises an important point. While the polynomials in figure 1.6 will certainly result in ψ increasing from 0, it is not obvious that they will produce a kink in $\psi(T)$ of the kind shown in figure 1.2. Might the order parameter not increase smoothly from zero, in which case the system undergoes not a PT but a crossover?

The answer to this lies in the partition function of equation 1.1. The partition function is a sum of exponentials, which are analytic everywhere. The only way to construct discontinuous behaviour from a sum of smooth functions is if the sum is infinite. For us to obtain the kink in figure 1.2, therefore, Z must be a sum over an infinite number of eigenstates, requiring an infinitely large physical system. Such a system is known as the *thermodynamic limit*. This leads to the conclusion that all continuous PTs become crossovers in finite-sized systems, and since all real systems are finite (except, perhaps, the entire universe), this means that continuous PTs are impossible in the strictest sense. In practice the distinction between crossovers and continuous PTs is essentially a matter of resolution: if the smoothing of the kink in 1.2 is so small as to be invisible, as is the case in any system with a number of particles on the order of Avogadro's number, this detail can be ignored. A typical particle number for ultracold atom experiments is on the order of 10^5 , which is generally enough to reach a confident conclusion, but for numerical calculations that can only simulate a few tens of particles the distinction is much more ambiguous.

Note the contrast with discontinuous PTs, which can occur perfectly well in finite systems. This is ultimately because the ordered and disordered phases remain distinguishable.

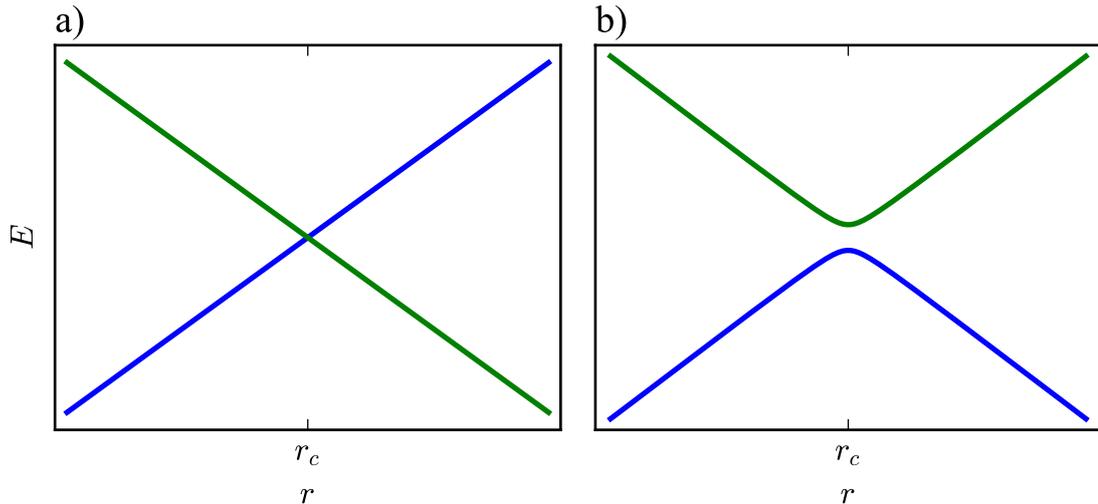


Figure 1.7: Sketched level diagrams for a) discontinuous and b) continuous quantum phase transitions. In each case the energies of two states are plotted as a function of some control parameter r (which is not temperature) in the vicinity of a level crossing at $r = r_c$. In a) the two levels are uncoupled and simply cross without interacting, while in b) a coupling term results in an avoided crossing and a change in our identification of the states above and below r_c .

1.4 Quantum phase transitions

1.4.1 A quantum picture of phase transitions

In this section I will offer an intuitive picture for extending our understanding of classical PTs to quantum systems via an analogy with Landau-Zener transitions [82]. This description is sufficient for understanding the experiments involving QPTs described in chapters 3 and 6, though there is also precedent for quantitative use [83]. However, it is only the start of a full discussion of QPTs. For a much fuller account I suggest [80] and especially [84].

Suppose we treat the two lowest eigenstates of some Hamiltonian \hat{H} , $|\phi_A\rangle$ and $|\phi_B\rangle$, as distinct phases, with the ground state being whichever of the two is lower in energy. A phase transition is then a level crossing where $E_A = E_B$, as illustrated in figure 1.7. For now we will ignore any issues of dynamics and assume that $T = 0$, so our only consideration is the ground state.

If a symmetry of \hat{H} prevents any transitions between $|\phi_A\rangle$ and $|\phi_B\rangle$ the two levels will cross and the ground state will jump suddenly from A to B. The ground state value of an observable described by some operator \hat{O} will jump suddenly from $\langle\phi_A|\hat{O}|\phi_A\rangle$ to $\langle\phi_B|\hat{O}|\phi_B\rangle$. If that observable is the order parameter, this is the signature of a discontinuous QPT. This behaviour is extremely unusual, as even the tiniest coupling between $|\phi_A\rangle$ and $|\phi_B\rangle$ will destroy this behaviour if we have sufficient experimental resolution. However, discontinuous QPTs have been observed or predicted in a handful

of systems [85–88].

If $|\phi_A\rangle$ and $|\phi_B\rangle$ are coupled by some term in the Hamiltonian they will show an avoided crossing that lifts the degeneracy at r_c . Near r_c the system will be in a superposition of $|\phi_A\rangle$ and $|\phi_B\rangle$ with intermediate properties, and only as we move away will it distinctly occupy either one of the two eigenstates. This process doesn't sound very abrupt, and indeed the situation is the same as in section 1.3.2: in finite systems we will see a crossover between the states $|\phi_A\rangle$ and $|\phi_B\rangle$, and this will become a quantum PT only in the thermodynamic limit. This is overwhelmingly the most common type of PT observed in quantum systems, with the most relevant example for this thesis being the transition between the Mott insulator and superfluid states for bosons in a lattice [89], and many more examples can be found in e.g. [80].

1.4.2 Dynamics of quantum phase transitions

Up until now we have considered PTs only in an entirely static sense, where we assume the system is always at equilibrium and all changes happen infinitely slowly. The model of QPTs presented in figure 1.7, however, has important implications when crossing the transition experimentally by tuning r . Let us for now consider the continuous case, which we have approximated as an avoided crossing between two energy levels.

The dynamics of avoided crossings is given by the *Landau-Zener* formula [82]. Let the energy gap between the two levels be $\Delta(t)$, controlled by varying r . Suppose we prepare the system at some $r_i \ll r_c$, sweep through the transition at a constant rate $\dot{\Delta}$, and end at $r_f \gg r_c$. If the minimum value of $\Delta(t)$ is Δ_0 , the probability of tunnelling across the gap to the excited state is

$$P \approx e^{-\frac{\Delta_0^2}{\dot{\Delta}}}. \quad (1.7)$$

From this we see that a sufficiently slow sweep, $\dot{\Delta} \rightarrow 0$, results in $P \approx 0$, meaning that we stay in the ground state and observe the PT we expect. (This is called *adiabatic passage*.) A sweep that is not sufficiently fast, however, allows some amount of tunnelling across the gap, resulting in a superposition of the two levels, and in the opposite limit of $\dot{\Delta} \gg \Delta_0^2$ we will jump entirely to the excited state and see no sign that there is a QPT at r_c .

Extension to infinite systems

As previously noted, continuous PTs only truly occur in the thermodynamic limit of infinite system size. However, in an infinite system there is not necessarily a gap between the ground state and the first excited state, as an infinite set of eigenstates is able to form a continuum: see for example the discussion of the gapless superfluid state in chapter 4. As the value of Δ_0 for the lowest excited state tends to zero, so does the rate $\dot{\Delta}$ required for adiabatic passage. This implies that in an infinite system we will be unable to cross a continuous QPT adiabatically, which is the basis of the experiment

described in section 3.6.2. Similarly, in a finite but large system the required $\dot{\Delta}$, while non-zero, may be prohibitively slow for experimental purposes.

However, note that the lowest excited state in a continuum most likely has properties extremely similar to those of the ground state. Rather than the upper state in figure 1.7, which represents a different *phase* and is therefore presumably very different in its properties, the state with $\Delta_0 \rightarrow 0$ is more likely to be the same phase as the ground state with some low-energy excitation. To avoid an admixture of *this* state we require $\dot{\Delta} = 0$, but the other phase is likely to be at significantly higher energy and consequently to have a more relaxed time scale for adiabatic passage.

This insight makes physical sense when we recall our mention of spontaneous symmetry breaking and the Kibble-Zurek mechanism in section 1.2. Crossing a continuous PT from a disordered phase to an ordered phase at a finite rate may result in different choices of the order parameter being made in different regions, with the boundaries between these regions then persisting as lasting excitations called domain walls. This is represented in our level crossing analogy by imperfect adiabatic passage. The infinitely long crossing time required to avoid any excitations reflects the time needed to establish correlations across an infinite system: see e.g. [90] for an investigation of this process.

Discontinuous quantum phase transitions

Let us now consider the discontinuous case from figure 1.7, in which $\Delta_0 = 0$ and there is no coupling between the two phases. In this case $P = 1$ regardless of $\dot{\Delta}$, matching our expectation that we cannot transition between two states that are not coupled. We cannot observe a discontinuous QPT experimentally by sweeping across r_c , because we will always stay on the same branch with no population in the new equilibrium phase. This phenomenon is called *metastability*. An external reservoir may allow a path back down to the ground state, but if the system is isolated we will stay in the metastable state and have no idea that we crossed a QPT at all. This is the basis of the experiment described in section 6.3.

I will emphasise the contrast between the continuous and discontinuous cases, as this will be a crucial point in later chapters. If we cross a continuous QPT between two phases A and B, going sufficiently slowly will allow us to reach B, though always with some excitations: the slower we go the fewer excitations we expect, but in an infinite system we can never go slowly enough to eliminate them. If the transition is discontinuous, we will simply stay in phase A and be unable to cross to phase B, even though A is no longer the ground state. Accessing phase B requires an additional degree of freedom to go around the discontinuous QPT (see e.g. section 6.3).

Note also, as with the classical case, that the discontinuous QPT can occur in finite systems and does not require the thermodynamic limit.

2

Quasicrystals and localisation

This chapter aims to describe quasicrystals and their relation to localisation. To do so, it is helpful to first review both periodic and randomly disordered systems, since quasicrystals are an intermediate between these two. The physics of periodic systems will also be a useful basis for discussing Floquet physics in chapter 5.

2.1 Physics of periodic systems

In systems governed by spatially periodic Hamiltonians, linear momentum p (or equivalently the wavenumber $k = p/\hbar$) is not a conserved quantity. Instead we introduce a new conserved quantity, quasimomentum, which is defined only within a specific range of values. To see why, it is useful to consider how momentum conservation appears in quantum mechanics in the first place.

2.1.1 Quasimomentum

Noether's theorem

Noether's theorem [91] states that any continuous symmetry of a physical system must give rise to a conservation law. In the specific context of quantum mechanics this is seen most easily using Ehrenfest's theorem [92], which states the following:

$$\frac{d}{dt}\langle\hat{O}\rangle = -\frac{i}{\hbar}\langle[\hat{O},\hat{H}]\rangle + \left\langle\frac{\partial\hat{O}}{\partial t}\right\rangle \quad (2.1)$$

Here, \hat{O} is a unitary operator corresponding to some observable, and the angle brackets $\langle \rangle$ refer to the expected value in any arbitrary wavefunction. If \hat{O} describes a conserved quantity the LHS must vanish, and this is related to the symmetries of the Hamiltonian by the terms on the RHS.

As an example, consider energy, in which case we let $\hat{O} = \hat{H}$. Since $[\hat{H}, \hat{H}]$ must vanish, we see that energy is conserved ($dE/dt = 0$) if and only if the Hamiltonian is constant ($\partial_t \hat{H} = 0$). Equivalently, we could say that symmetry under translation in time gives rise to conservation of energy.

Alternatively we could consider the linear momentum operator, \hat{p} . In this case $\partial_t \hat{p}$ vanishes, so momentum is conserved if and only if $[\hat{p}, \hat{H}] = 0$. Since $\hat{p} = -i\hbar\partial_x$, this commutator will only vanish if \hat{H} is independent of x . Therefore we conclude that momentum conservation requires translational invariance of the Hamiltonian.

Discrete translational symmetry

In periodic systems, the Hamiltonian obeys the condition

$$\hat{H}(x) \equiv \hat{H}(x + a) \quad (2.2)$$

for some distance a , often called the *lattice constant*. However, we do not require that the Hamiltonian is constant over distances less than a . The result is that the *continuous* translational invariance required for momentum conservation is replaced by *discrete* translational invariance. The result of this is that the linear momentum¹ $p = \hbar k$ is conserved only up to a change

$$k \rightarrow k \pm \frac{2n\pi}{a}, \quad (2.3)$$

where n is any integer. We can essentially think of momentum in quantities of $2n\pi\hbar/a$ as being freely available².

Given equation 2.3 we can always map k onto a value q in the interval $[-\frac{\pi}{a}, \frac{\pi}{a})$, which is called the *first Brillouin zone* (first BZ), by choosing an appropriate value of n . We then call the quantity $\hbar q$ the *quasimomentum*, which is conserved no matter how many times we apply the transformation 2.3.

Note that a particular q corresponds to an infinite set of ks , spaced by $2\pi/a$. This is shown experimentally in the Kapitza-Dirac images of section 3.5.3, in which every peak is a k value corresponding to $q = 0$: this particular set of k values is known as the *reciprocal lattice*.

¹For simplicity we consider only one spatial dimension. The generalisation to any number of dimensions is straightforward [93].

²These packets of momentum are zero-energy acoustic phonon modes.

2.1.2 Bloch waves

The eigenstates of periodic Hamiltonians are described by *Bloch's theorem*. I will not present a rigorous proof here – see [93] for an authoritative account – but I will give a plausibility argument that may offer some physical insight. For simplicity we work in one dimension, but it is trivial to extend to higher dimensions provided they are *separable*, i.e. the Hamiltonian can be written as a sum of independent terms $\hat{H}(x, y, z) = \hat{H}_x(x) + \hat{H}_y(y) + \hat{H}_z(z)$. The following discussion then applies to each term independently with no additional complications.

Consider a Hamiltonian \hat{H} of the form

$$\hat{H}(p, x) = \hat{T}(p) + \hat{V}(x), \quad (2.4)$$

where $\hat{T}(p)$ is the kinetic energy operator, a function of momentum, and $\hat{V}(x)$ denotes potential energy as a function of position. A periodic system has a Hamiltonian of this form with the additional property that

$$\hat{V}(x + a) \equiv \hat{V}(x) \quad (2.5)$$

for some *lattice constant* a . In our later discussion of optical lattices we will be interested in the case where $V(x)$ is a cosine,

$$\hat{V}(x) = V_0 \cos(k_{\text{lat}} x), \quad (2.6)$$

which we will refer to as a one-dimensional lattice potential. We might reasonably assume that all physical properties of the eigenstates $\phi_j(x)$ of this Hamiltonian will share this periodicity, so that

$$|\phi_j(x + a)|^2 \equiv |\phi_j(x)|^2. \quad (2.7)$$

The phase of the eigenstate is free to change between sites, leading us to suspect the most general form will be

$$\phi_{n,q}(x) = e^{iqx} u_{n,q}(x) \quad (2.8)$$

where $u_{n,q}(x)$ is a periodic function with period a , and we have split the label j to represent the wavenumber q and any other quantum numbers n that might be important³. Wavefunctions with this form are called *Bloch waves*. It is also useful to define the *Wannier states* given by

$$w_{n,j}(x - x_j) = \frac{1}{\sqrt{N}} \sum_q e^{iqx_j} \phi_{n,q}(x), \quad (2.9)$$

where x_j is the position of the j th lattice site (the j th minimum of the potential in equation 2.6) and N is the total number of sites.

There are two principle strategies for finding the eigenstates $\phi_{n,q}$ of a Hamiltonian of the form in equation 2.4, depending on whether kinetic energy dominates (the *nearly free* regime) or potential energy (*tight binding*). In subsequent chapters we will consider experiments in the tight binding regime, so I will present that approach here.

³Note that the switch from p to q here makes no difference to the value of e^{iqx} , since $e^{2\pi i} = 1$.

2.1.3 Tight binding

The tight binding regime refers to the case where V_0 from equation 2.6 is "large". (It is frequently the case that this condition holds only for $n = 0$, so I will at this point drop the index n and assume that $n = 0$ from here on.) In this limit the ground band Wannier states (equation 2.9) exist almost entirely on a single site, so that $w(a) \ll w(0)$ and $w(2a) \approx 0$. We can therefore imagine that particles in the lattice will mostly sit on one site in a Wannier state, before occasionally tunnelling (*hopping*) to the adjacent site, with no possibility of hopping two sites. This assumption is called *nearest-neighbour hopping*. In this approximation we can reverse equation 2.9 to write the Bloch wave eigenstates as

$$\begin{aligned}\phi_{0,q}(x_m) &= \sum_j e^{iqx_j} w_j(x_m - x_j) \\ &= w_m(0) + e^{iqa} w_{m+1}(-a) + e^{-iqa} w_{m-1}(a) \\ &= w(0) + 2\cos(qa)w(a).\end{aligned}\tag{2.10}$$

In the last line we have also assumed that the Wannier functions are symmetric, $w_j(x - x_j) \equiv w_j(-x + x_j)$, and that the Wannier states on each site are equivalent by translational symmetry. To find the energy of this state we must then find the expectation value of \hat{H} :

$$\begin{aligned}E_0(q) &= \int \phi_{0,q}^*(x) \hat{H} \phi_{0,q}(x) dx \\ &= \int \left[w^*(x) \hat{H} (w(x) + 2\cos(qa)w(x+a)) \right] dx \\ &= \int w^*(x) \hat{H} w(x) dx + 2\cos(qa) \int w^*(x) \hat{H} w(x+a) dx\end{aligned}\tag{2.11}$$

We then define some additional terms:

$$E_0(q) =: E_c - 2J\cos(qa)\tag{2.12}$$

where E_c is the energy of the Wannier state $w(x)$ and J is known as the *overlap integral*⁴. $E_0(q)$ is known as the *ground band*, and the sinusoidal shape we have derived here is characteristic of the tight-binding regime.

Higher bands

Repeating this calculation for $n \neq 0$ defines a spectrum of bands $E_n(q)$ at higher and higher energies. More and more bands will adopt the form 2.12 as the lattice depth V_0 grows deeper, but for this thesis we are concerned only with the case that E_0 is in the tight-binding regime and no other bands are. A numerical plot of $E_n(q)$ for a particular set of parameters is shown in figure 2.1.

⁴Sometimes also called the hopping integral or transfer integral. This quantity is often given the variable t in conventional condensed matter physics.

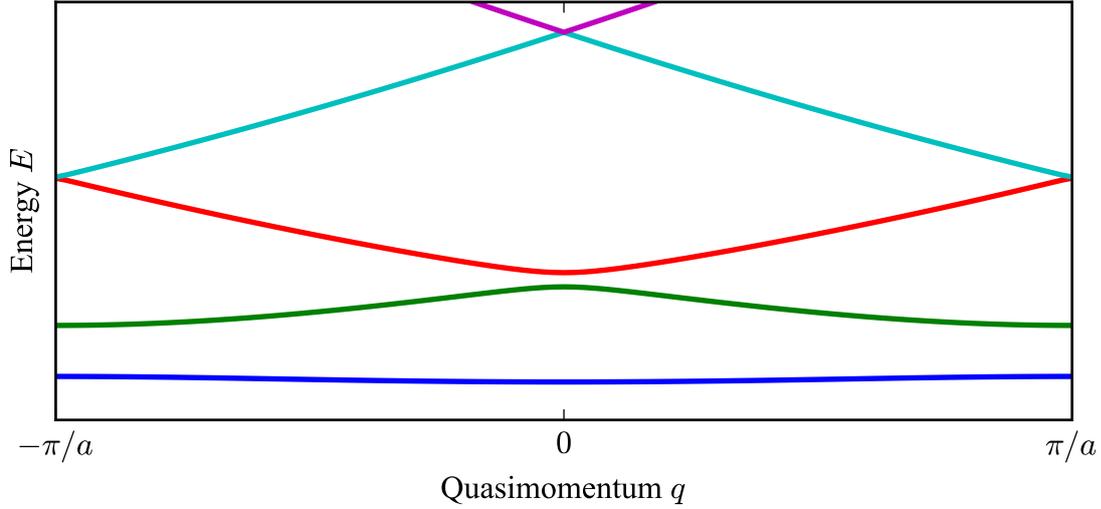


Figure 2.1: Band structure $E_n(k)$ of a 1D cosine lattice with depth $V_0 = 5E_r$, the approximate lower value for the tight-binding regime (see text). The lowest band is approximately a sinusoid (equation 2.12), though it appears quite flat on this scale. The method followed to calculate this figure is given in section 2 of [67].

Note that the validity of the tight binding regime is band-dependent. We usually assume that the lowest band is in the tight binding regime for $V_0 \gtrsim 5E_r$, where

$$E_r := \frac{\hbar^2 k_{\text{lat}}^2}{2m} \quad (2.13)$$

is known as the recoil energy for a particle of mass m in the lattice. Deeper lattices will be required for higher bands to show tight binding behaviour: this is why the higher bands in figure 2.1 do not have a sinusoidal shape.

2.1.4 Semiclassical dynamics of particles

Let us consider how to represent a particle using the Bloch basis. The Bloch waves of equation 2.8 are extended over all space, so to construct a localised wavefunction $\Psi_n(x)$ we consider a wavepacket

$$\Psi_n(x) = \int_{-\frac{\pi}{a}}^{\frac{\pi}{a}} g(q) \phi_{n,q}(x) dq \quad (2.14)$$

where $g(q)$ is a distribution peaked at some q_0 . We can now treat this wavepacket as corresponding to a classical particle, with velocity given by the *group velocity*

$$v_g = \frac{1}{\hbar} \partial_q E_n(q) |_{q_0}. \quad (2.15)$$

This allows us to derive the effect of an external force F on such a particle. First we equate two different expressions for the rate of work done by such a force:

$$\begin{aligned}
F v_g &= \partial_t E_n(q)|_{q_0} \\
&= \partial_q E_n(q) \partial_t q|_{q_0} \\
&= \hbar v_g \partial_t q|_{q_0} \\
\implies F &= \hbar \partial_t q|_{q_0}.
\end{aligned} \tag{2.16}$$

Recalling that the quasimomentum is given by $\hbar q$, what we have derived is precisely Newton's Second Law with quasimomentum replacing momentum. To gain further physical insight we can consider the effect on the group velocity:

$$\begin{aligned}
\partial_t v_g &= \frac{1}{\hbar} \partial_t \partial_q E_n(q) \\
&= \frac{1}{\hbar} \partial_q^2 E_n(q) \partial_t q && \text{Use equation 2.16} \\
&= \frac{1}{\hbar^2} \partial_q^2 E_n(q) F \\
\implies F &= \frac{\hbar^2}{\partial_q^2 E_n(q)} \partial_t v_g \\
&=: m^* \partial_t v_g
\end{aligned} \tag{2.17}$$

where in the last step we have defined the *effective mass* of the particle, m^* , by comparison with the other form of Newton's Second Law, $F = m \partial_t v$. This result, together with equation 2.16, is essential for our physical intuition into what band structure actually means. We find that the curvature of the band directly translates to how "heavy" a particle occupying it appears to be.

Bloch oscillations

This semiclassical model has interesting implications when we consider the finite range of possible values for quasimomentum. If we take a particle initially at $q = 0$ and apply a constant force F_0 , we see from equation 2.16 that the value of q will change linearly,

$$q(t) = \frac{F_0}{\hbar} t. \tag{2.18}$$

Once q grows to reach π/a , it can grow no larger. Instead we redefine it by subtracting $2\pi/a$, in accordance with our earlier discussion of conservation laws, and q instantaneously jumps to $-\pi/a$ to begin increasing again. Meanwhile the velocity of the particle is given by

$$v_g(t) = \frac{1}{\hbar} \partial_q E_n(q(t)). \tag{2.19}$$

If the particle is in the lowest band of a tight-binding lattice, equation 2.12, we obtain

$$\begin{aligned}
v_g(t) &= -\frac{2J}{\hbar} \partial_q \cos(q(t)a) \\
&= \frac{2Ja}{\hbar} \sin(q(t)a) \\
&= \frac{2Ja}{\hbar} \sin\left(\frac{F_0 a}{\hbar} t\right).
\end{aligned} \tag{2.20}$$

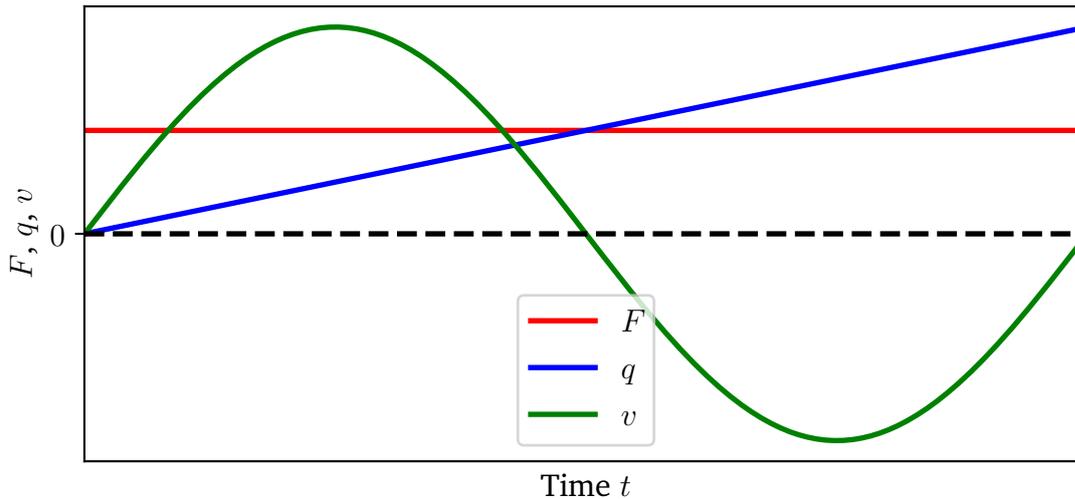


Figure 2.2: An illustration of Bloch oscillations (equations 2.18 and 2.20) in a periodic potential. Under a constant external force F the quasimomentum q of a particle will increase linearly, but the velocity v shows oscillations that result in no net motion.

This result, together with $q(t)$, is plotted in figure 2.2. We see applying a constant force to a particle in a periodic potential causes the particle to oscillate back and forth, spending half the time moving against the force, so that on average it goes nowhere! This surprising phenomenon is called a *Bloch oscillation*, and has been observed in both optical lattices and conventional condensed matter systems [94, 95].

2.1.5 The Bose-Hubbard model

We now add an additional complexity: interactions between particles. This results in all kinds of fantastically rich physics, encompassing everything from magnetism to superconductivity to Mott insulators, so we will only consider a very specific model called the *Bose-Hubbard model* (see e.g. [2]). We make the following assumptions:

- The particles are identical bosons.
- They occupy a 1D lattice in the tight-binding regime with nearest-neighbour hopping.
- They are in the ground band (equation 2.12) and the other bands are far away in energy.
- The interparticle interaction has a very short range, so that it only applies between bosons occupying the same site.

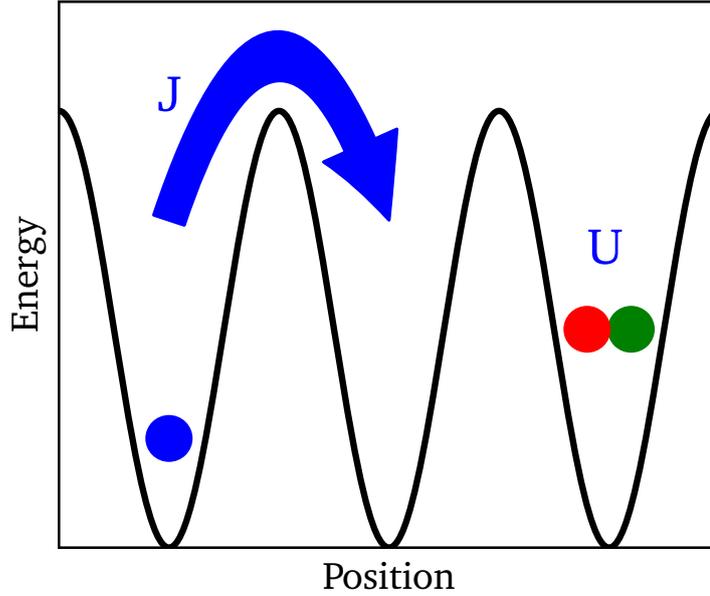


Figure 2.3: An illustration of the parameters in the Bose-Hubbard Hamiltonian of equation 2.21. J is the energy scale of particles hopping between sites, while U is the interaction energy incurred when two particles occupy the same site. Note that the third parameter, μ , is not illustrated (see main text).

Under these conditions we can use second-quantised notation to write the Bose-Hubbard Hamiltonian in the form

$$\hat{H}_{\text{BH}} = \underbrace{-J \sum_{\langle i,j \rangle} (\hat{a}_i^\dagger \hat{a}_j + \hat{a}_j^\dagger \hat{a}_i)}_{\text{Hopping}} + \underbrace{\frac{1}{2} U \sum_i \hat{n}_i (\hat{n}_i - 1)}_{\text{Interactions}} - \underbrace{\mu \sum_i \hat{n}_i}_{\text{Chemical potential}}. \quad (2.21)$$

There are many terms to be defined here. \hat{a}_i is the operator that removes a boson from the i th lattice site, known as an *annihilation operator*, while its Hermitian conjugate \hat{a}_i^\dagger , the *creation operator*, adds a boson to the same site. The number operator $\hat{n}_i := \hat{a}_i^\dagger \hat{a}_i$ gives the number of bosons on site i , which can range from 0 to ∞ .

J is the energy scale of the hopping process as before. The angle brackets under the sum in the hopping term indicates that the sum is only over nearest neighbours (i.e. only over $i = j \pm 1$).

U quantifies the on-site interaction between particles and is given by

$$U = g \int_{-\infty}^{\infty} |w_0(x)|^4 dx, \quad (2.22)$$

where g captures the strength of the interaction itself and the integral over the Wannier state quantifies how tightly the particles are confined on top of each other. We usually consider repulsive interactions, in which case $U > 0$.

Finally, μ is the chemical potential. This term is often omitted, but it will be required in our discussion of the Mott insulator in chapter 4, and also for understanding the

BEC transition (see discussion of equation 3.14). It captures how "full" the lattice is of particles, indicating for example how many bosons there are per lattice site on average. This is typically important when atoms are able to enter and leave the system (usually because we are considering only a part of the lattice), in which case whether the system is more or less "full" will determine whether particles are inclined to leave or enter. This situation is known as the *grand canonical ensemble*.

Phases of the Bose-Hubbard Hamiltonian ground state

The Bose-Hubbard Hamiltonian of equation 2.21 has two ground-state regimes, determined by the ratio J/U and separated by a quantum phase transition (section 1.4). If $J \gg U$ then hopping dominates and all parts of the lattice are able to remain coherent, resulting in a *superfluid* phase. If on the other hand $J \ll U$, the Hamiltonian is dominated by interactions and the system may become a *Mott insulator* depending on the value of μ . This behaviour will be discussed at length in chapter 4.

2.1.6 Adding random disorder

As a last note on periodic systems I will briefly and qualitatively describe the effects of adding random disorder to the system, meaning the appearance in the Hamiltonian of a random variable drawn from some distribution. In conventional condensed matter physics disorder is commonly introduced by doping, though myriad other sources of disorder exist including grains and dislocations. In optical lattice experiments random disorder can be introduced with a DMD (section 7.1.2 and see e.g. [38]) projecting a pseudo-randomly generated pattern of pixels, or by passing light through a diffuser to introduce speckle.

The effect of random disorder depends strongly on the presence or absence of interactions, so I will consider each case separately. First, however, it is instructive to consider what "localisation" means, since it is a common feature of disordered systems.

Localisation and ergodicity

A system is said to be *ergodic* if, in the thermodynamic limit and given infinite time, it will lose all memory of its initial conditions and fully explore all regions of its phase space that are not forbidden by conservation laws. For example, in a bottle of water at room temperature the molecules will eventually move through every possible arrangement of positions that doesn't defy gravity, and do so in much the same way no matter how they were arranged to begin with.

By contrast, a non-ergodic system retains a memory of its initial conditions for infinite time. There is no good example in common experience, but the simplest case may be a collisionless gas: in the absence of interactions particles cannot exchange energy and the initial distribution of velocities persists forever. Another simple case is a system

where every particle is constrained to stay within a certain distance of its starting position, so that long-range transport is impossible and an initially uneven density distribution cannot fully even out; this is called a *localised* system, and is the most commonly studied case of non-ergodicity in condensed matter physics.

Anderson localisation

The effect of adding random disorder to a non-interacting periodic system is called *Anderson localisation* [96]. As noted previously the eigenstates in such a system are Bloch waves, which have a constant amplitude out to infinity. Adding random disorder, however, with correlation length identically 0, will cause *all* eigenstates⁵ to acquire a term $e^{-\frac{r}{\xi}}$ that decays exponentially over the *localisation length* ξ . The magnitude of ξ is determined by the range of the random variable, but even the tiniest random perturbation will still cause ξ to become non-zero and the system to become localised. This prevents long-range transport, as no particle in any state can move further than a few ξ from its initial location [97, 98].

There is an experimental complication if ξ is larger than the size of the system – unlikely in a real crystal but very possible in an optical lattice. This introduces a practical limitation on the smallest amount of disorder that will make localisation apparent in transport measurements. Note also that truly random disorder is essentially impossible to achieve – for example, any optical system must have finite resolution that sets a minimum disorder correlation length.

Many-body localisation

Many-body localisation (MBL) refers to a different scenario. A system with both random disorder and a repulsive interaction is said to be many-body localised if there are high-energy excited states that are non-ergodic [99]. This will *not* generally happen for an infinitesimally small disorder term; rather, disorder and interactions typically act in opposition to one another, with a repulsive interaction encouraging the system to spread out as far as possible through phase space and become ergodic. This tends to mean there is a transition from an extended state to an MBL state when the disorder is increased or the interaction reduced (though if the interaction vanishes we recover Anderson localisation instead). See e.g. [100].

Note that the Bose glass of section 7.2.1 is closely related to MBL states, but is distinguished by being a ground state phase while MBL is a feature of excited states.

Glass

The extreme case of random disorder is a *glass* or amorphous solid, in which there is no lattice and all particles are randomly distributed. These systems are typically

⁵In 1D and 2D.

insulators with all eigenstates localised, and the comparison between glasses and MBL states is a field of current research [101].

2.2 Quasicrystals

2.2.1 What is a quasicrystal?

Within this thesis, I will use the term *quasicrystal* (QC) to mean a system that has three properties:

1. Long-range order, such that knowing the local arrangement of a small region determines the structure out to infinity.
2. Aperiodicity (a complete lack of translational symmetry).
3. Rotational symmetry of an order other than 1, 2, 3, 4 or 6.

The related term *quasiperiodicity* requires only the first two and can therefore apply to one-dimensional systems. There are broadly speaking two reasons why QCs are interesting, which I will summarise briefly here before covering them in more detail below – see also e.g. [47, 48].

First, QCs do not fit neatly into most discussions of either ordered or disordered systems. Ordered systems in condensed matter physics are usually also periodic, and considered using the treatment of section 2.1 that begins with Bloch’s theorem. Since this theorem does not apply in aperiodic systems, none of this treatment is applicable to QCs. On the other hand, the long-range order of QCs means they cannot be considered randomly disordered. This has profound implications for localisation in quasicrystalline systems, discussed below.

Secondly, QCs are intrinsically linked to higher-dimensional systems. More precisely, any QC may be viewed as a slice through a fictional *parent* or *ancestor* crystal that is higher-dimensional and periodic [51, 102, 103]. The reason this is interesting is that QCs can be constructed for which the parent has more than three spatial dimensions: as we shall see in section 3.6, the optical quasicrystal in our experiment has a four-dimensional parent. This enables experimental access to systems that cannot otherwise be built owing to the unfortunate shortage of spatial dimensions in our universe⁶, and is of particular interest because it may allow experimental observation of the four-dimensional quantum Hall effect [52] (see section 7.2.3).

Quasicrystals in nature

Although the concept of QCs can be explored in the abstract, it is perhaps of benefit to briefly note that they can exist in practice. The first experimental observation of a

⁶Extended ones, anyway.

This figure has been redacted for copyright reasons. The copyright owner is the American Physical Society. See print version, or top inset from figure 2 in [50].

Figure 2.4: Reproduced from Shechtman’s paper [50], this X-ray diffraction pattern from an aluminium-manganese alloy was the first proof of a real quasicrystalline material. Note the ten-fold symmetry of diffraction peaks, precluding periodicity even while the sharp peaks confirm long-range order.

QC was made in 1982 by Shechtman when recording the X-ray diffraction pattern of an aluminium-manganese alloy: see figure 2.4.

Quasicrystals in mathematics

Interest in quasiperiodic and quasicrystalline patterns in pure mathematics pre-dates their discovery in nature, especially in the context of tessellating the 2D plane. The most famous example is the fivefold-symmetric Penrose tiling [49], but I will also mention the Ammann-Beenker tiling [104] as its eightfold symmetry matches the optical QC described in chapter 3.

2.2.2 Incommensurable length scales in quasiperiodic systems

A reliable method for creating a quasiperiodic pattern is to superpose two periodic patterns with periods p_1 and p_2 that are *incommensurable*⁷. This means that the equation

$$ap_1 = bp_2 \tag{2.23}$$

has no solutions for integer a and b , or equivalently that the ratio p_1/p_2 is an irrational number such as $\sqrt{2}$. The significance of this is that if we move by some distance $np_1 + mp_2$, for integer n and m , we can always find the unique values of n and m

⁷The word *incommensurate* is often used instead; they are synonyms.

afterwards. Equivalently we could observe that we can never return to our starting position unless both n and m are 0.

Let's see how this produces a quasiperiodic system. Consider the function

$$f(x) := \sin\left(\frac{x}{2\pi}\right) + \sin\left(\frac{x}{2\sqrt{2}\pi}\right). \quad (2.24)$$

By inspection this function is long-range ordered, since it's just a superposition of two repeating patterns. But is it periodic? Suppose it is, with some period x_0 . In that case we know that $f(x_0) = f(0)$, but also that $f'(x_0) = f'(0)$ and so on for all higher derivatives. Since this gives infinite simultaneous equations to satisfy and our only variable is x_0 , the only solution is that x_0 is a whole number of periods for *both* terms in $f(x)$ individually. This gives the conditions

$$\begin{aligned} x_0 &= n \\ x_0 &= \sqrt{2}m \\ \implies \frac{n}{m} &= \sqrt{2} \end{aligned} \quad (2.25)$$

for integer n, m . Since $\sqrt{2}$ is irrational, there is no choice of m and n that satisfies this last line, and we conclude that $f(x)$ can't be periodic: instead, it is quasiperiodic.

The relation between incommensurable length scales and quasiperiodic systems is very general, and is central both to the relationship between QCs and higher dimensions discussed in the next section, and to experimental realisations of quasiperiodic optical lattices.

2.2.3 Links to higher dimensions

A very productive way to consider QCs is as the projection of a periodic lattice – the parent crystal – onto a lower-dimensional space, an approach known as the *cut-and-project method* [105]. The parent crystal must have at least d_p dimensions, where

$$d_p = n_i d \quad (2.26)$$

for a d -dimensional quasicrystal with n_i incommensurable length scales [102]. Most quasiperiodic optical lattices have $d = 1$ or 2 and $n_i = 2$, resulting in $d_p = 2$ or 4 . The OL can then be modelled as a d -dimensional section through a space with d_p dimensions, which allows the simulation of d_p -dimensional systems (see e.g. [103]).

There is another, more intuitive way of understanding this mapping to higher dimensions. Consider again the function $f(x)$ from equation 2.24. This potential has two length scales, 1 and $\sqrt{2}$, meaning that its reciprocal lattice⁸ will also have two length scales, 1 and $1/\sqrt{2}$. Any position in the reciprocal lattice will therefore have the form

$$\mathbf{r} = n\mathbf{g} + m\mathbf{h} \quad (2.27)$$

⁸Which technically is not a lattice, since it is not periodic, but there is no agreed-upon alternative term. Reciprocal quasilattice?

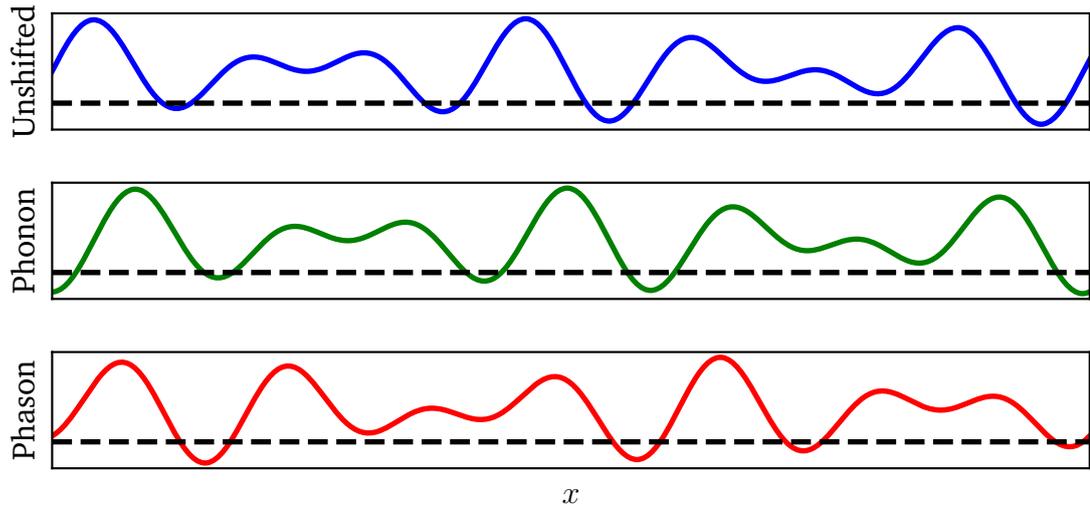


Figure 2.5: Illustration of the effects of phonons and phasons on a quasiperiodic lattice (see main text). These three graphs show the function $f(x, x_1, x_2) = \sin(k_1(x + x_1)) + \sin(k_2(x + x_2))$, where $k_2/k_1 = \sqrt{2}$: a 1D quasiperiodic lattice. The black dotted line is at the same height in each case, and represents an arbitrary chemical potential μ such that each minimum dipping below μ represents an occupied lattice site. The first graph is $f(x, 0, 0)$, the unshifted lattice. The second is $f(x, -x_0, -x_0)$ for a constant x_0 . Note that the green curve is identical to the blue curve, but shifted to the right: this is an example phononic strain. The third curve is $f(x, 0, -x_0)$, with one sublattice shifted but not the other. Comparing the blue and red curves note that the second minimum from the left has risen above μ , while a new minimum has dipped below on the right-hand side. In other words, a lattice site has disappeared and a new one has appeared to replace it at a different position. This abrupt shift is totally different from the smooth phononic case, and is characteristic of quasiperiodic functions.

where $|\mathbf{h}| = \sqrt{2}|\mathbf{g}|$. Although the two vectors are collinear, the incommensurability of their moduli means that n and m always have unique values, as discussed above. However, equation 2.27 is identical to the reciprocal of a 2D periodic lattice, which also has two linearly independent vectors. Therefore, it is natural to say that the potential expressed by equation 2.24 is a 1D quasiperiodic lattice with a 2D periodic parent: if we just "forget" where equation 2.27 comes from, we have no way of knowing we aren't dealing with the 2D system.

Phasons

The link between quasiperiodic systems and higher dimensions has one particularly striking implication: a collective excitation, or quasiparticle⁹, unique to quasiperiodic systems. These excitations are called *phasons*, and they can be intuitively understood in several ways.

⁹This use of the prefix quasi- is totally unrelated to quasicrystals, and is used generally for many-body systems – an unfortunate overlap of terminology.

First, consider again the 2D square lattice suggested by equation 2.27. This lattice has no phasons, but it does have phonons (normal mode collective oscillations of the lattice) along two orthogonal directions. These two phonon modes must be reflected in two independent quasiparticle degrees of freedom in the 1D quasiperiodic system as well, since they are described by the same reciprocal lattice, but a 1D system will have only one phonon mode. The "extra" degree of freedom becomes a phason mode. While this argument gives no clues as to what phasons look like, it gives the general result that the number of phason modes is $d_p - d = d(n_i - 1)$. Therefore, for the systems with $n_i = 2$ that are our chief concern, there are d phason modes to match the d phonon modes.

Alternatively, let us see how this extra mode arises from equation 2.24. To see lattice excitations we must generalise the equation to

$$f_\phi(x) := \sin\left(\frac{x}{2\pi} + \phi_1\right) + \sin\left(\frac{x}{2\sqrt{2}\pi} + \phi_2\right), \quad (2.28)$$

where ϕ_1 and ϕ_2 range between 0 and 2π . Collective oscillations of the lattice correspond to time-dependent perturbations $\delta\phi_1$ and $\delta\phi_2$ of these phases, since that will cause the positions of lattice sites to oscillate. Since we have two phases, it's immediately clear that there are two independent excitation modes. Suppose we parametrise them as follows:

- A mode with $\delta\phi_1 = \delta\phi_2$. This corresponds to a translation of the lattice without its form changing: a phonon.
- A mode with $\delta\phi_1 = 0$. In this case only one of the sublattices oscillates, while the other stays fixed. This is the phason mode.

The reader is left to satisfy themselves that a linear combination of the above two modes allows any combination of ϕ_1 and ϕ_2 .

Phasons have been observed both in real materials [106] and in quasiperiodic optical lattices [107]. See [108] for an excellent numerical visualisation of phasons.

2.2.4 Between order and disorder

The other broad reason for interest in quasicrystals lies in their intermediate status between periodic and disordered systems, as a result of which they are sometimes referred to as "quasi-disordered". To see this in precise terms, we introduce a quantity known as the *structure factor*,

$$F(\mathbf{k}) := \int d^3\mathbf{r} \rho(\mathbf{r}) e^{-i\mathbf{k}\cdot\mathbf{r}} \quad (2.29)$$

where $\rho(\mathbf{r})$ is the density distribution of a lattice (or random system) in real space. A more detailed calculation for the case of a lattice will be given in section 3.5.2, but

for now let us note qualitatively that the structure factor looks like the interference pattern produced by a plane wave with a specific wavevector \mathbf{k} originating from every part of the system, weighted according to density. There are three interesting cases.

- For a periodic lattice $\rho(\mathbf{r}) = \sum_n \delta(\mathbf{r} - \mathbf{r}_n)$ for a periodic array of sites \mathbf{r}_n with spacing a . $F(\mathbf{k})$ will then be another periodic lattice with spacing $\propto 1/a$. The lattice sites of $F(\mathbf{k})$ occur where every plane wave interferes constructively, and the spaces in between have perfect destructive interference. (This is the reciprocal lattice mentioned in section 2.1.1.)
- If $\rho(\mathbf{r})$ is randomly distributed $F(\mathbf{k})$ has no sharp features. There is no chance for overall constructive or destructive interference because there are always waves with every possible phase.
- If $\rho(\mathbf{r})$ is a quasiperiodic array of Dirac deltas, $F(\mathbf{k})$ will be a dense¹⁰ set of Dirac deltas with varying strength.

The third case is simply a restatement of equation 2.27, and is illustrated by the experimental pictures of section 3.5.3. The n and m of equation 2.27 must always be uniquely identifiable, meaning that the infinity of peaks indexed as (n, m) has to fit on the 1D line of wavenumbers k . Fitting infinite values for both n and m onto one line is only possible if the spacing between values is arbitrarily small.

The sharp peaks of the QC structure factor resemble the periodic case, but the lack of finite gaps recalls the case of random disorder. This leads us to suspect that QCs may have other behaviours intermediate between these two systems, for example in the context of localisation.

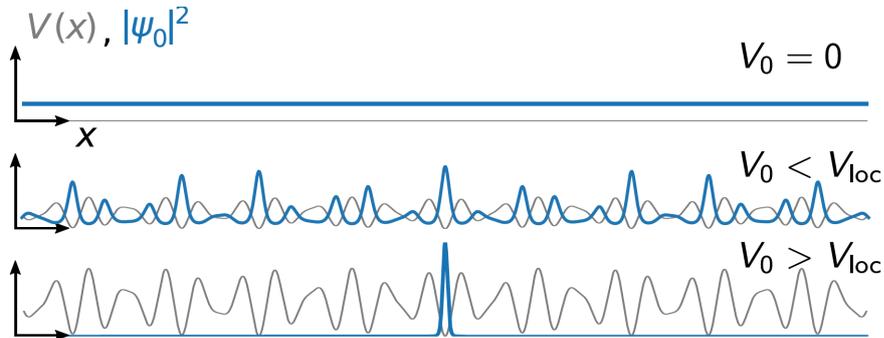


Figure 2.6: Reproduced from [57] copyright APS. A sketch of the eigenstates of a quasiperiodic lattice of depth V_0 , for $V_0 = 0$ and two finite values well below and well above the critical value V_{loc} for the localisation transition. The $V_0 = 0$ case is uniform, corresponding to a Bloch wave. For $0 < V_0 < V_{loc}$ the wavefunction is substantially perturbed by the quasiperiodic potential but nonetheless extends to infinity. Well above the transition the wavefunction localises to a single lattice site. (Note that if $V_0 \gtrsim V_{loc}$ the particle may still extend over more than one lattice site, but certainly does not reach infinity.)

¹⁰Dense here means that for every element e of the set you can find another element as close to e as you like, so there is no minimum spacing between elements. The rational numbers are the archetypal example of a dense set.

2.2.5 Localisation in quasicrystals

Much like random disorder (section 2.1.6), quasi-disorder can cause both Anderson localisation in non-interacting systems and MBL in interacting systems. This has been studied both numerically, for example with the Aubry-André model [109, 110], and experimentally [34, 37–39].

In the case of Anderson localisation there is an important distinction between the effects of random disorder and quasi-disorder. As previously noted, even an infinitesimally small random term in the Hamiltonian will cause Anderson localisation, albeit with a long localisation length. However, an infinitesimal quasiperiodic term will only perturb the extended eigenstates, not localise them (figure 2.6). Instead there is a *localisation transition* when the quasiperiodic term reaches a certain critical size, with the ground state undergoing a continuous QPT from extended to localised. This is another example of behaviour intermediate between ordered systems (always extended) and randomly disordered (always localised). Adding repulsive interactions will have the usual effect of opposing localisation, which here means increasing the critical strength [69].

3

An optical quasicrystal for ultracold atoms

In this chapter I will summarise the basic operation of our ultracold atom machine and the experiments we have performed with it so far. I joined the group one year into the building of this machine, when the first ^{87}Rb MOT had just been achieved by my former colleagues Matteo Sbroscia and Konrad Viebahn. Their theses [69, 70] contain a much more detailed description of the apparatus, including the work done before my involvement, but in later chapters we will require only the brief account presented here.

3.1 Cooling down atoms

In this section I will briefly describe how we take ^{87}Rb and ^{39}K atoms from room temperature to the "ultracold" quantum degenerate regime. The relevant physical quantity for quantum degeneracy is the *phase space density* $\langle n \rangle$, defined as the expected occupation of any particular single-particle state. In a classical gas this is very close to 0, so that no two particles ever occupy the same state and their fermionic or bosonic nature is never relevant. If $\langle n \rangle$ for any state approaches 1 the gas begins to display non-classical behaviour, which for bosons means Bose-Einstein statistics and ultimately condensation. This condition is also referred to as *quantum degeneracy*.

In achieving a value of $\langle n \rangle$ close to 1 our principal tool is light, in the form of laser beams. I will first discuss how light interacts with atoms and the ways in which we use this to achieve cooling. An authoritative account of this material can be found in [111].

3.1.1 The light-atom interaction

A laser beam incident on an atom will exert a force

$$\mathbf{F} = \underbrace{\hbar \mathbf{k} \frac{\Gamma}{2} \frac{I}{I_{\text{sat}}}}_{F_{\text{scatter}}} \frac{1}{1 + \frac{I}{I_{\text{sat}}} + \left(\frac{2\delta}{\Gamma}\right)^2} - \underbrace{\frac{\hbar \delta}{2} \frac{\nabla\left(\frac{I}{I_{\text{sat}}}\right)}{1 + \frac{I}{I_{\text{sat}}} + \left(\frac{2\delta}{\Gamma}\right)^2}}_{F_{\text{dipole}}} \quad (3.1)$$

where \mathbf{k} is the wavevector of the light; I is its intensity; δ is the *detuning*, or frequency difference between the light and the closest atomic transition; Γ is the width of the transition; and I_{sat} is the saturation intensity, also a property of the transition. The detuning is referred to as *blue* if $\nu_{\text{laser}} > \nu_{\text{atom}}$, and *red* if the reverse is true. F_{scatter} is parallel to \mathbf{k} , and represents a push along the direction of the light caused by scattered photons that are spontaneously re-emitted. F_{dipole} arises from the second-order process where the atom coherently absorbs and re-emits a photon through stimulated emission, and which causes a shift in the ground state energy depending on the sign of the detuning. In the limit of $\delta \gg \Gamma$, these equations become

$$\begin{aligned} \mathbf{F}_{\text{scatter}} &= \hbar \mathbf{k} R_{\text{scatter}} \\ \mathbf{F}_{\text{dipole}} &= -\nabla U_{\text{dipole}} \end{aligned}$$

where

$$\begin{aligned} R_{\text{scatter}} &\approx \frac{\Gamma}{8} \left(\frac{\Gamma}{\delta}\right)^2 \frac{I}{I_{\text{sat}}} \\ U_{\text{dipole}} &\approx \frac{\hbar \Gamma}{8} \frac{\Gamma}{\delta} \frac{I}{I_{\text{sat}}}. \end{aligned}$$

The most important behaviour to note here is that $R_{\text{scatter}} \propto \frac{1}{\delta^2}$, while $U_{\text{dipole}} \propto \frac{1}{\delta}$. This means that we can tune the relative importance of these two effects by changing the detuning of the laser, with spontaneous effects more important near resonance but vanishing at large detunings to be replaced by coherent ones.

3.1.2 Early cooling stages: spontaneity

The scattering force can only be used to cool atoms above a temperature known as the *recoil limit*. This is because spontaneous emission results in a momentum kick from the departing photon. The accompanying *recoil energy* E_r sets the temperature limit T_r :

$$\begin{aligned} k_B T_r &\approx E_r \\ &= \frac{\hbar^2 k_r^2}{2m}, \end{aligned} \quad (3.2)$$

where m is the atomic mass and k_r is the wavenumber of the atomic transition. For ^{87}Rb this has the value 180 nK, while for ^{39}K it is 415 nK. In practice we expect to achieve temperatures on the order of 100 times this limit, perhaps $\sim 50 \mu\text{K}$ for both species.

We use two techniques, described more fully in [111], to cool atoms above T_r with the scattering force:

- In a magneto-optical trap (MOT), the Doppler and Zeeman effects are used to reduce the detuning for faster atoms and atoms further from the trap centre, respectively. This results in atoms being trapped and slowed.
- The optical molasses technique also uses the Doppler effect, but achieves lower temperatures using Sisyphus cooling, which involves optical pumping in a light field with spatially varying polarisation.

3.1.3 Late cooling stages: coherence

Once we have reached the limits of spontaneous cooling we switch to using far-detuned lasers and magnetic fields to create conservative potentials for the atoms. We can then employ *evaporative cooling*, in which the barrier to escape from a trapping potential is gradually lowered. The most energetic atoms escape first, and if the evaporation is slow enough the cloud continuously re-thermalises at progressively lower temperatures. We use two principal experimental methods to achieve evaporative cooling:

- Atoms in certain angular momentum states may be trapped by a quadrupole magnetic field. Evaporation is achieved by shining in radio- or microwave-frequency radiation to drive transitions between trapped and untrapped states. Only the most energetic atoms reach the regions with the highest magnetic field, so the difference in detuning can be used to target them first.
- Sufficiently cold atoms can be trapped in a powerful red-detuned laser beam, known as an optical dipole trap (ODT), by the restoring force arising from the dipole potential. Evaporation is accomplished by decreasing the intensity of the beam.

There is no fundamental limit to the temperature achievable with evaporative cooling (though with fewer and fewer atoms), and we continue to use this method until the gas becomes degenerate.

Sympathetic cooling

While evaporative cooling can be effective for ^{39}K , we avoid using it for as long as possible. This is because the earlier cooling stages are much less efficient for ^{39}K than for ^{87}Rb ¹, so that we generally have a much smaller population of potassium and do not wish to lose nearly all of it by evaporating. Instead we use a technique called *sympathetic cooling*, which involves trapping a large population of ^{87}Rb along with our

¹This is ultimately because the hyperfine splitting in ^{39}K is much smaller, making it impossible to address individual transitions.

^{39}K and evaporating only the rubidium. The ^{39}K atoms are then cooled by collisions with ^{87}Rb . We resort to direct evaporation of ^{39}K , aided by a Feshbach resonance (section 3.3.2), only when all the rubidium has fallen out of the ODT under gravity.

This method has another advantage: we can also choose to keep some ^{87}Rb around to the end of the cooling process and have both species present. We may eventually want to work with mixtures for some unforeseen future experiment, so sympathetic cooling keeps our options open.

3.2 The BEC transition

The Bose-Einstein condensation (BEC) transition is a continuous QPT (quantum phase transition) of the kind discussed in sections 1.3.2 and 1.4. The ordered phase is the BEC itself, while the disordered phase is a dilute classical gas. We identify the control parameter r from figure 1.7 with $1/N$ (N is the number of particles). The order parameter (section 1.2) can be chosen as the condensate fraction N_0/N , which is the proportion of particles in the lowest single-particle state. The usual treatment of the transition, which I will present here, focusses on the regime $N > N_c(T = 0)$ and $T > 0$, and the phase diagram in this region is depicted in figure 1.5. This material is drawn in large part from [112] and [113].

3.2.1 Statistical mechanics

The starting point for the statistical mechanics of the BEC transition is the Bose distribution,

$$f(E) = \frac{1}{e^{\frac{E-\mu}{k_B T}} - 1}, \quad (3.3)$$

which gives the occupation of a single-particle state with energy² E at chemical potential $\mu \leq 0$ and temperature T for a system of bosonic particles. Figure 3.1 shows this distribution for a few values of the parameters. From this definition we can immediately write down the relation

$$\begin{aligned} N &= \int_0^\infty g(E)f(E)dE \\ &= \int_0^\infty \frac{g(E)}{e^{\frac{E-\mu}{k_B T}} - 1} dE, \end{aligned} \quad (3.4)$$

where $g(E)$ is the density of states. The precise expression for $g(E)$ depends on the number of spatial dimensions and the trapping potential that confines the gas, but it is typically of the form

$$g(E) = C_\alpha E^{\alpha-1}, \quad (3.5)$$

²If we wish the ground state to have some energy $E_0 \neq 0$ we replace E with $E - E_0$.

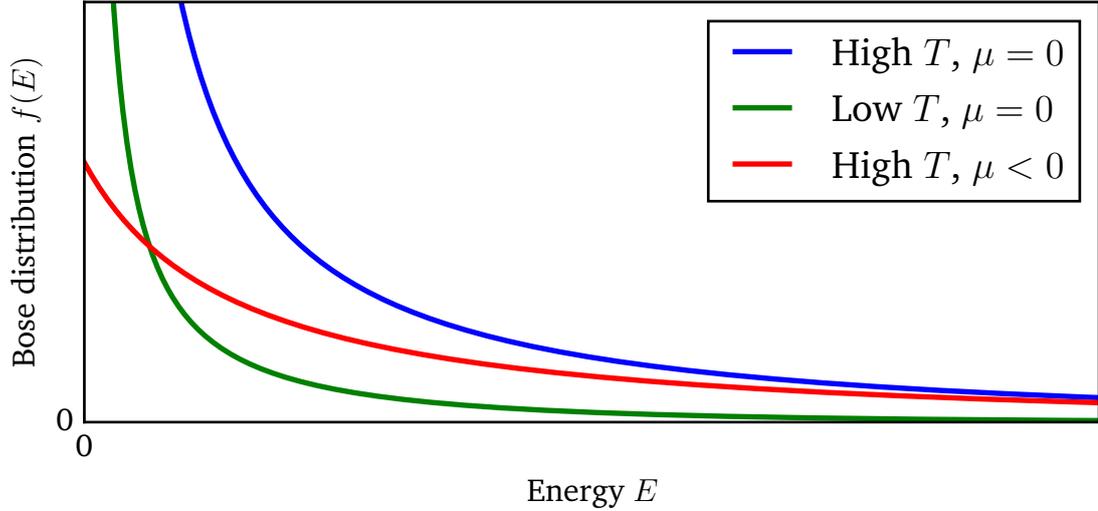


Figure 3.1: The Bose distribution of equation 3.3 for three different values of (T, μ) . The blue curve is for $\mu = 0$ and an arbitrary "high" T , showing an infinity at $E = 0$. It is interesting to contrast this with the other curves, which show the effect of lowering T (green) or allowing μ to become negative (red). Note that in both cases the area under the curve has visibly decreased, and in the case of $\mu < 0$ the divergence as $E \rightarrow 0$ has also disappeared. The area is relevant because it will determine the value of N_c (equation 3.8), and the infinity at $E = 0$ is the key to Bose-Einstein condensation (see text).

where α and C_α must be calculated for the particular system under consideration. In our experiment the BEC is in a three-dimensional harmonic trap, for which case³ $\alpha = 3$ and

$$C_3 = \frac{1}{2\hbar^3 \omega_x \omega_y \omega_z} \quad (3.6)$$

for trapping frequencies $\omega_{x,y,z}$ along each direction. Knowing the energy dependence of $g(E)$ allows us to make progress with equation 3.4, which now takes the form

$$N = C_\alpha \int_0^\infty \frac{E^{\alpha-1}}{e^{\frac{E-\mu}{k_B T}} - 1} dE. \quad (3.7)$$

This expression for N reaches its largest value for a given T when $\mu = 0$ (this is illustrated in figure 3.1; recall that $\mu \leq 0$ for a Bose gas). Let us call this maximum particle number N_c , given by

$$N_c = C_\alpha \int_0^\infty \frac{E^{\alpha-1}}{e^{\frac{E}{k_B T}} - 1} dE. \quad (3.8)$$

However, in an experiment we can always increase N further by simply adding more particles to the system. The Bose distribution will be violated if the atoms enter any state with $E \neq 0$, but $f(0) \rightarrow \infty$ as $\mu \rightarrow 0$, meaning that there is always space for

³Another important case is the 3D box trap, in which particles are confined by rigid walls to a region of uniform potential. In that case $\alpha = \frac{3}{2}$ and $C_{3/2} = \frac{Vm^{3/2}}{\sqrt{2\pi^2\hbar^3}}$ for box volume V and particle mass m .

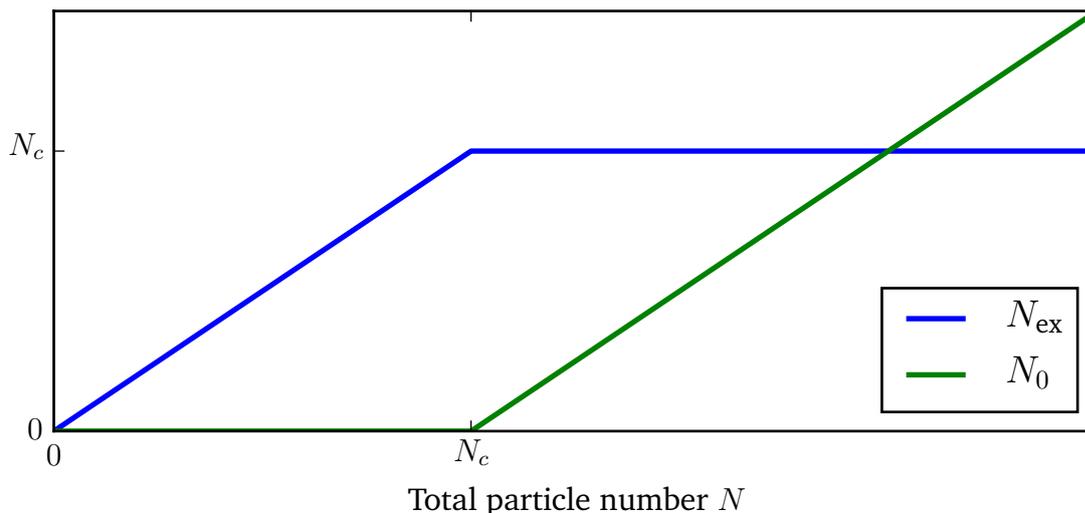


Figure 3.2: A plot of the particle numbers in the ground state (N_0) and all other states (N_{ex}) as a function of the total particle number N for a Bose gas. At some critical number N_c (equation 3.8) the Bose distribution forbids any further particles from entering the excited states and additional particles must go to the ground state. The sudden change in the gradient of N_0 constitutes a continuous phase transition (compare with figure 1.2).

more atoms in the ground state $E = 0$. (Equation 3.8, being a continuum approximation, neglects this contribution: it is a single point in the integral.) Therefore, once $\mu = 0$, any further atoms will have to enter the single-particle ground state, and the expression for $N(> N_c)$ becomes

$$N = N_0 + N_c \quad (3.9)$$

with N_0 atoms in the ground state and N_c atoms across all excited states. Therefore the BEC order parameter N_0/N increases from 0 as soon as $N > N_c$, which we call the *critical number*. This behaviour is shown in figure 3.2.

In order to evaluate N_c , start from equation 3.8 and make the substitution $x := E/k_B T$, obtaining

$$\begin{aligned} N_c &= C_\alpha (k_B T)^\alpha \int_0^\infty \frac{x^{\alpha-1}}{e^x - 1} dx \\ &= C_\alpha \Gamma(\alpha) \zeta(\alpha) (k_B T)^\alpha \\ &\propto T^\alpha, \end{aligned} \quad (3.10)$$

where $\Gamma(\alpha)$ is the gamma function and $\zeta(\alpha)$ is the Riemann zeta function. Despite the intimidating coefficients, this expression has a simple temperature dependence: for example, in a 3D harmonic trap we have $N_c \propto T^3$. It's important to note that $\Gamma(\alpha)\zeta(\alpha)$ does not converge for all values of α : in particular, $\zeta(1) = \Gamma(0) = \infty$. In these cases there's no BEC transition.

There is another very common way of describing the same physics. Suppose that instead of fixing T and adding particles to the system, we instead fix N and lower T ,

which in turn lowers N_c (see figure 3.1). If we reduce the temperature far enough we can make N_c equal to N , at which point further cooling will force atoms to condense into the ground state. In this picture we identify a temperature T_c at which $N = N_c$, and all the above equations can be rewritten by moving the c subscript from N to T . Equation 3.10 therefore becomes

$$k_B T_c = \frac{N^{1/\alpha}}{[C_\alpha \Gamma(\alpha) \zeta(\alpha)]^{1/\alpha}}, \quad (3.11)$$

and we can combine both equations to write

$$N_c = N \left(\frac{T}{T_c} \right)^\alpha. \quad (3.12)$$

Recalling equation 3.9 we at last obtain an expression for the order parameter,

$$\frac{N_0}{N} = 1 - \left(\frac{T}{T_c} \right)^\alpha. \quad (3.13)$$

This is the result that was used to plot figure 1.5, with $\alpha = 3$ as for a 3D harmonic trap.

A note on the chemical potential

In the above derivation we have made use of the fact that μ is never positive for a Bose gas. It is worth considering why, in terms of the physical meaning of μ . In classical thermodynamics μ is defined by

$$\mu = \left. \frac{\partial E_{\text{tot}}}{\partial N} \right|_{S,V} \quad (3.14)$$

where E_{tot} is the total energy of the system, S is entropy and V is volume. Therefore we can think of μ as the energy at which we can add a particle without the entropy of the system changing⁴. At high temperatures $\mu < 0$, which we interpret to mean that a classical gas is so chaotic, with so many possible configurations of particles, that we have to slow all the particles down to balance the entropy of adding one more⁵.

When we raise N to access the quantum regime, the system has to increase the area under the Bose distribution shown in figure 3.1. If we hold T constant it does this by increasing μ , and eventually reaches the point where $\mu = 0$ and the Bose distribution becomes singular at $E = 0$. This singularity means that the distribution can now accommodate an infinite number of particles (as long as they go into a BEC), and this is why the chemical potential never increases further: if I add particles to try and push μ up, I'll find they all fit in the $\mu = 0$ distribution already. Rather than saying that μ can't be positive for a Bose gas, it's more intuitive to say that we have no way to make it go positive, because the distribution becomes singular for $\mu = 0$ ⁶.

⁴The requirement that the volume remain constant is much less illuminating.

⁵This also applies to Fermi gases, since quantum statistics are only important at low temperatures.

⁶This can be contrasted with the case of fermions. The Fermi-Dirac distribution always remains finite, so by adding particles we can make the chemical potential as positive as we like.

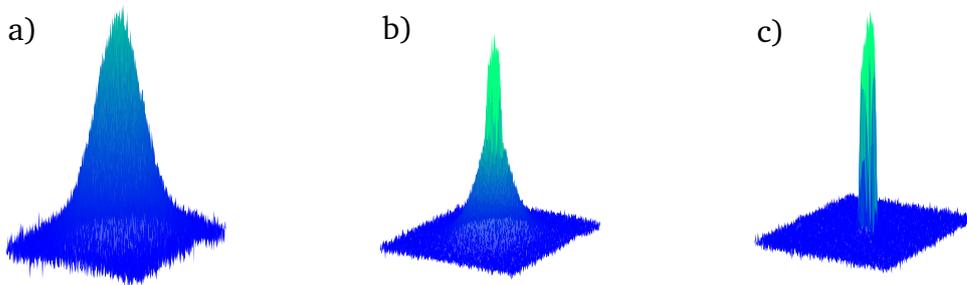


Figure 3.3: Three time-of-flight images of atomic clouds that have a) $N_0/N = 0$, a purely classical cloud with a Gaussian Maxwell-Boltzmann momentum distribution; b) $0 < N_0/N < 1$, with the sharp peak of the condensate at $k = 0$ visible above a Gaussian background; and c) $N_0/N = 1$, with no Gaussian background left. The colour scale is the same in all three images and indicates the optical density of the cloud.

3.2.2 Imaging the transition

The BEC transition can be detected experimentally using time of flight (TOF) imaging, which involves abruptly releasing the cloud from its harmonic trap and allowing it to fall for some time t_{TOF} before taking a picture. A particle at initial position \mathbf{r}_0 with momentum $\hbar\mathbf{k}$ will arrive at a position

$$\mathbf{r}_{\text{image}} = \mathbf{r}_0 + \frac{\hbar\mathbf{k}t_{\text{TOF}}}{m} + \frac{1}{2}\mathbf{g}t_{\text{TOF}}^2 \quad (3.15)$$

under gravity \mathbf{g} . Our image will record the projection of \mathbf{r} onto the imaging plane⁷. If t_{TOF} is chosen to be sufficiently long the first term can be neglected and we can directly image the momentum distribution of a cloud. Because in the BEC a significant proportion of (or all) particles are in the ground state, which has $k = 0$, the condensate fraction will appear as a sharp peak at the centre of the cloud in the TOF image, appearing above the Gaussian background of atoms that aren't in the condensate⁸. Figure 3.3 presents data from our experiment showing this peak appearing as the cloud crosses the transition.

⁷All TOF images in this thesis are taken from directly below the atoms, so that the contribution of gravity is invisible.

⁸Strictly speaking this isn't quite right, as while $k = 0$ is the ground state in free space, the BEC forms in a trap where plane waves are not stationary states. This results in a slight broadening, consistent with Heisenberg's uncertainty principle and the finite size of the trap ground state, so we do not truly obtain a Dirac delta at $k = 0$. In practice, however, other effects like weak interactions or the initial position \mathbf{r}_0 will be more significant.

3.3 Interactions

3.3.1 The scattering length

So far we have only considered the non-interacting⁹ BEC, but the behaviour of the gas is much more interesting if we introduce a weak repulsive interaction. We can do this most simply with the Hamiltonian

$$\hat{H} = \sum_i^N \left(\frac{\mathbf{p}_i^2}{2m} + V(\mathbf{r}_i) \right) + g \sum_{i=1}^N \sum_{j=1}^{i-1} \delta(\mathbf{r}_i - \mathbf{r}_j), \quad (3.16)$$

where we have added a point-contact interaction with strength g . This is somewhat physically justified, for the following reasons:

- If the gas is dilute, we can assume that the interparticle spacing d is much greater than the range r_{int} of any interparticle interaction ($d \gg r_{\text{int}}$). This means that even a finite-range interaction might as well be point-contact, since the particles have to be so much closer than average already.
- The size λ of the wavepacket describing each particle becomes large at low temperatures, and so we can assume $\lambda \gg r_{\text{int}}$. It follows that any interaction between particles must be very close to a head-on collision characterised by just one number, g . This assumption is called *s-wave scattering*.

The implications of this will be developed more fully in section 4.1.1, but now I will simply note that we often express g as a *scattering length* given by

$$a := \frac{m}{4\pi\hbar^2} g, \quad (3.17)$$

usually measured in units of the Bohr radius a_0 . Note that $a > 0$ for a repulsive interaction.

3.3.2 Feshbach resonances

Many atoms display one or more features called *Feshbach resonances*, wherein a depends on the ambient magnetic field \mathbf{B} with the functional form

$$a(B) = a_{\text{bg}} \left(1 - \frac{\Delta}{B - B_0} \right) \quad (3.18)$$

where $B = |\mathbf{B}|$, Δ is the width of the resonance, B_0 is the field at which the interaction is strongest¹⁰, and a_{bg} is the scattering length far from any resonances. This is plotted in figure 3.4.

⁹Strictly speaking, we have assumed that the interparticle interactions are strong enough for the cloud to thermalise but weak enough that the eigenstates of the system are the same as the single-particle eigenstates.

¹⁰The vicinity of this point is known as the unitary regime.

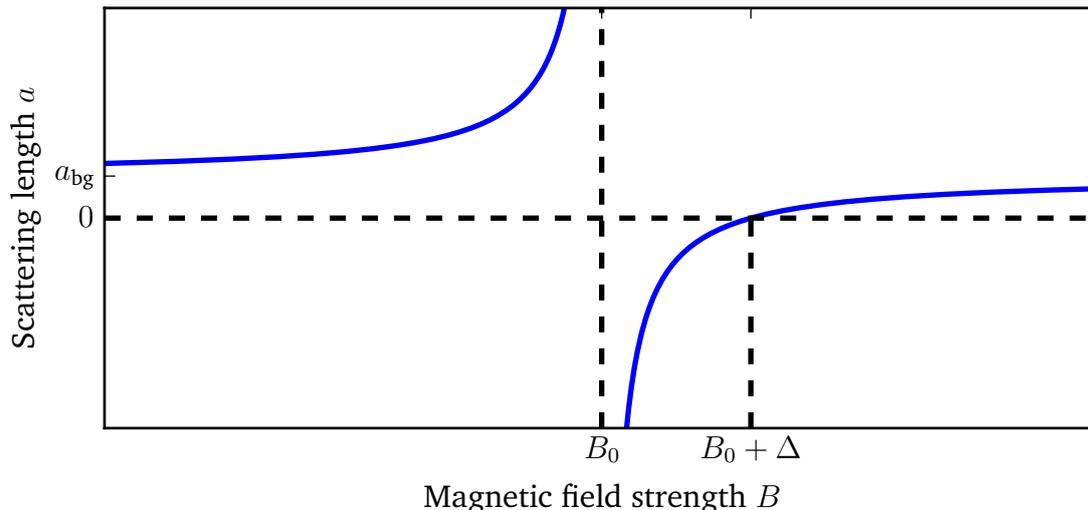


Figure 3.4: Behaviour of the interatomic scattering length a at a Feshbach resonance (equation 3.18). a diverges at $B = B_0$ and vanishes at $B = B_0 + \Delta$, giving access to every possible strength of interaction provided a magnetic field $\approx B_0$ controlled to a precision $\ll \Delta$ is experimentally achievable.

The origin of these resonances is beyond the scope of this thesis and definitive accounts already exist [114], so I will simply note that they are one of the most important advantages of ultracold atomic physics: if we choose an atomic species with such a resonance, we can freely tune the interparticle interaction to whatever value we want provided the values of Δ and B_0 are experimentally reasonable.

In our experiment we have so far implemented two different atomic species, with one of the most significant differences between them being their behaviour regarding Feshbach resonances¹¹:

- ^{87}Rb has several Feshbach resonances. However, they all have extremely small values of Δ ($\ll 1$ G), and in several cases very large values of B_0 . This renders their use impractical, and experimentally we consider ^{87}Rb to have a fixed scattering length $a = 100a_0$.
- ^{39}K has a Feshbach resonance at $B_0 = 402.4$ G, with $\Delta = -52$ G and $a_{\text{bg}} = -29a_0$. These parameters make its use in experiments straightforward.
- Additionally, there is an interspecies Feshbach resonance between ^{87}Rb and ^{39}K at $B_0 = 317.9$ G, with $\Delta = 7.6$ G and $a = 34a_0$, allowing us to tune the interaction between the two atoms without affecting either intraspecies scattering length.

The upshot of this is that we can tune interactions for ^{39}K but not ^{87}Rb , allowing for many more experiments with ^{39}K .

¹¹The values quoted here assume the atoms are in the state $|F = 1, m_F = 1\rangle$.

3.3.3 Quantum depletion

In the interacting BEC it is no longer true that all particles occupy the $|\mathbf{k} = \mathbf{0}\rangle$ state at zero temperature. Instead there will be some population n_{ex} in the "excited" states, not due to any thermal excitation. This effect is called *quantum depletion*, or *condensate depletion*, and if sufficiently pronounced is visible on TOF images such as those in 3.3 as a Gaussian background. The effect of significant quantum depletion, in fact, is indistinguishable in TOF from the effect of heating above the BEC transition, as the same states are occupied in each case. Provided that n_{ex} is small it can be derived [112] that

$$\frac{n_{\text{ex}}}{n} = \frac{8}{3\sqrt{\pi}} \sqrt{na^3}. \quad (3.19)$$

In BECs of ultracold atoms the depletion is typically $\lesssim 1\%$ ¹², but note that $n_{\text{ex}} \propto \sqrt{n}$: increasing the density, for example by strong compression in an ODT, can lead to huge quantum depletion and an apparently thermal cloud.

3.4 Summary of experimental sequence

A combination of all the above techniques allows us to produce a BEC of either ⁸⁷Rb or ³⁹K, perform an experiment and take an image in 25 s. We then repeat the process over and over. I will now summarise the experimental steps required for a ³⁹K BEC and the time afforded to each (as previously mentioned, see [69, 70] for a much more detailed discussion). Note that making an ⁸⁷Rb BEC does not require ³⁹K, whereas to make a ³⁹K BEC we must cool both species and remove the ⁸⁷Rb at the end.

Note that in this summary I will refer a few times to the internal hyperfine state of the atoms using the quantum numbers F and m_F . For our purposes it is sufficient to understand that these quantities refer to angular momentum, with m_F determining the magnitude of the Zeeman shift in a magnetic field. The curious reader may consult [111] for the full story.

- We start by loading MOTs of both species for 2.5 s. The atoms come from solid samples inside the vacuum chamber and are transferred to the MOT by a separate "2D MOT"¹³ for each species.
- In the last 100 ms or so of the MOT we ramp down the power of the repump¹⁴ light for ³⁹K, allowing the atoms to fall into a dark state to reduce scattering. This "temporal dark spot MOT" is necessary because ³⁹K's closely-spaced hyperfine states make optical molasses less effective.

¹²This is not true of all degenerate Bose systems. In superfluid ⁴He, for example, strong interactions result in a depletion of $\sim 90\%$ even at $T = 0$ [113].

¹³A 2D MOT is a MOT with no magnetic field gradient along one direction. This produces a marginally focussed beam of atoms along that direction at a cooler temperature than the initial vapour, and the atoms in this beam are then trapped by the MOT.

¹⁴Repump is a key ingredient for making a MOT in practice, preventing the atoms from falling into a dark state during Doppler cooling.

- Switching off the quadrupole magnetic field in the MOT chamber (and switching back on the ^{39}K repump), we cool using optical molasses for 5 ms. During this time the atoms are not trapped, and so the clouds expand, but sub-Doppler cooling allows us to achieve a temperature of around $50\ \mu\text{K}$ for both species. The repump light is switched off slightly before the cooling light for both species, so that all atoms end up in the $|F = 1\rangle$ hyperfine state.
- The quadrupole field is switched back on, trapping those atoms that are in the $|F = 1, m_F = -1\rangle$ state with a magnetic field gradient. (Consequently the $\frac{2}{3}$ of atoms with $m_F = 0, 1$ are lost.) The atoms are then transported from the MOT chamber to the science chamber by ramping a series of coil pairs to move the magnetic field minimum, a process which lasts 5 s.
- Evaporative cooling of ^{87}Rb is achieved using microwave radiation to couple the $|F = 1, m_F = -1\rangle$ state to the untrapped $|F = 2, m_F = -2\rangle$ state¹⁵, gradually decreasing the microwave frequency over 8 s. The hyperfine splitting of the two species is different by an order of magnitude so we do not lose any ^{39}K atoms, which are instead cooled sympathetically by ^{87}Rb .
- We ramp down the magnetic trap and switch on the ODT, applying a weaker uniform guide field to keep the atoms polarised in the $|F = 1, m_F = -1\rangle$ state.
- A ~ 100 ms Landau-Zener sweep of RF radiation takes the atoms from the $|F = 1, m_F = -1\rangle$ state to $|F = 1, m_F = 1\rangle$ (via $|F = 1, m_F = 0\rangle$). Both species have the same Zeeman shift of 0.7 MHz G^{-1} , so this works for both at once. We use $|F = 1, m_F = 1\rangle$ because it has the lowest energy of the m_F states and consequently the longest lifetime in the dipole trap and lattice.
- We evaporatively cool once again by reducing the intensity of the ODT beams, employing the interspecies Feshbach resonance to enhance the rate of collisions between ^{87}Rb and ^{39}K . Because $m_{\text{Rb}} \approx 2m_{\text{K}}$, ^{87}Rb falls out of the dipole trap significantly before ^{39}K , and we use this to achieve a pure potassium cloud. We then switch to the intraspecies Feshbach resonance to allow ^{39}K to continue evaporating by itself. After 3 s in the ODT we achieve a BEC with $N_0/N = 1$.
- We perform an experiment by adjusting the Feshbach field and switching on the optical lattice beams, either before or after we switch off the ODT fully. This can be as brief as a few μs for Kapitza-Dirac diffraction or as long as a few hundred ms for adiabatic loading of the lattice.
- We switch off the optical lattice and the ODT, allowing the atoms to fall for a time on the order of 30 ms before taking a TOF absorption image. We typically shift the Feshbach field to the zero crossing at $B_0 + \Delta$ at the start of TOF to allow momentum peaks to separate, before increasing the scattering length again to "puff up" the clouds and avoid optical densities too high to measure.

The output of the sequence is always one or more photographs of the atoms in TOF. We can also image the atoms without letting them fall first (*in-situ imaging*), but so far we haven't used this technique for any experiments.

¹⁵We do not reach cold enough temperatures for Majorana losses [69, 70] to become important, so we do nothing to prevent them.

3.5 Quantum simulation

There are several reasons to create BECs. For example, the physics of the BEC itself is a rich topic of research, and cold atoms have growing applications in metrology and sensing. However, our interest is in using them for *quantum simulation*, which means attempting to understand a complex quantum system by creating a more tractable one governed by the same idealised Hamiltonian [1–3]. This is necessary because quantum systems with many interacting degrees of freedom are too difficult to simulate on existing (or foreseeable) computers.

In this section I will present methods particularly focused on simulating the solid state physics discussed in chapter 2, which revolve around *optical lattices* (OLs), and I will then sketch how these methods have been used by our group in two previous experiments.

3.5.1 Optical lattices

An OL is a laser interference pattern $I(\mathbf{r})$ that creates a Hamiltonian for ultracold atoms via the dipole force 3.1.1, which is proportional to $-\nabla I(\mathbf{r})/\delta$. Interest has typically been in cases where

$$I(\mathbf{r} + \mathbf{a}) = I(\mathbf{r}) \quad (3.20)$$

with

$$\mathbf{a} = (na_x, ma_y, la_z) \quad (3.21)$$

for integers n, m, l and lattice constants a_x, a_y, a_z in three dimensions. In these cases the periodic physics discussed in section 2.1 applies and the spectrum of the system will have band gaps.

A rich variety of Hamiltonians can be studied in this way, depending on several key features of the lattice design.

Lattice detuning

Nowadays it is usual that OLs are made using far-detuned lasers to avoid heating from spontaneous emission. However, experiments conducted with non-degenerate clouds have frequently used near-resonant lasers and included active cooling such as optical molasses while the atoms were in the lattice: see e.g. [55, 56]. This allowed the investigation of diffusive and dispersive dynamics that I will not discuss in this thesis.

An important consideration for far-detuned lasers is the the *sign* of the detuning, i.e. blue or red. If the detuning is blue ($\delta > 0$) the minima of the potential will coincide with the minima of intensity, while for a red-detuned lattice ($\delta < 0$) the lattice sites will be maxima of intensity. This has two main consequences.

- If the detuning is not so large as to make scattering completely negligible, the

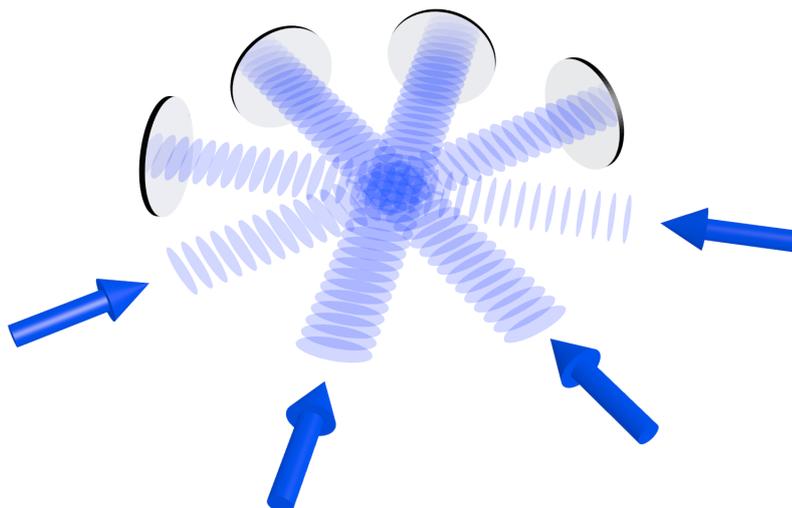


Figure 3.5: Reproduced from [57], copyright APS. An illustration of the lattice geometry used for our optical quasicrystal. Four laser beams are retro-reflected from perpendicular mirrors, each forming a 1D standing wave. In the region where these four standing waves superpose a 2D eightfold-symmetric optical lattice is formed, though the standing waves are mutually detuned so that they do not interfere with each other. Similar geometries can simulate a large variety of lattices, by changing the number and angle of laser beams and possibly by removing retro-reflection or allowing the various beams to interfere. A single standing wave may be added perpendicular to the 2D plane to provide confinement to a plane.

red-detuned case may have a marginally shorter lattice lifetime as the atoms are exposed to more light.

- The lattice potential always has a Gaussian envelope due to the intensity profile of the interfering beams. In the red-detuned case this envelope will tend to pull atoms towards the centre of the lattice (*confining*), while in the blue-detuned case they will be pushed to the edges (*anticonfining*). See section 7.1.1 for more detail, as well as strategies for compensating this effect.

In our experiment we use a blue-detuned lattice at 725 nm, far detuned from the ^{87}Rb and ^{39}K resonances at 780 nm and 767 nm respectively.

Lattice geometry

The "geometry" of the lattice refers to the number and orientation of laser beams used to create it, which in turn determines the reciprocal lattice vectors. An example of a fairly complex lattice is given in figure 3.5, but the simplest way to create an OL is to reflect a single laser beam back on itself with a perpendicular mirror. This is called *retro-reflection*, and produces a 1D standing wave of nodes and antinodes. While the electric field $\mathbf{E}(t)$ averages to zero everywhere, the average intensity $\langle I(t) \rangle := \langle \mathbf{E}(t) \cdot \mathbf{E}(t) \rangle$ vanishes only at the nodes, with maxima at the antinodes. This results in

a 1D periodic pattern $I(x) = I_0 \cos^2(k_L x)$ that will attract atoms towards its maxima or minima for red- or blue-detuned light respectively.

More complex lattices are often achieved by adding similar standing waves along additional directions, possibly with different wavelengths. A 2D square lattice requires two perpendicular retro-reflected pairs, and a third pair along the final direction gives a 3D simple cubic lattice. Different symmetries within the 2D plane are achieved by matching that symmetry with the arrangement of beams: a triangular lattice, for example, requires three lattice beams in the same plane at 60° angles. This enables simulation of many different condensed matter systems.

Note that in practice a 1D or 2D OL experiment will often use a 3D arrangement of lasers to prevent dynamics in the unwanted directions. For example, a spherical BEC loaded into a 1D OL will be split into discs, but to observe strongly-correlated physics we usually want only a few atoms on each site. Adding 1D OLs along the other two directions that are deep enough to prevent hopping will allow truly 1D dynamics along the direction of interest. Similarly, 2D OLs of any geometry will often have a 1D OL along the perpendicular direction to confine atoms to the plane.

Coherence of lattice beams

In the 1D OL just described, interference between the forward- and backward-propagating beams is guaranteed, because they are the same beam separated only by a single mirror. For more complex arrangements, however, we may not want all pairs of lattice beams to be mutually coherent. A 2D square lattice, for example, may be constructed by superposing two 1D lattices of the kind described above, but there is no need for the two standing waves to interfere: they'll create the necessary pattern of maxima and minima regardless. On the other hand, three lattice beams can create a triangular or hexagonal lattice without any retro-reflection, in which case each beam needs to interfere stably with both of the other two and their relative phases will determine the resulting geometry. If interference is not desired, it can be avoided by using perpendicular polarisations, detuning the lattice beams by a few tens of MHz so that the interference terms vary too fast to affect the atoms, or both.

The interfering and non-interfering cases for a particular geometry will generally have different physics. For example, interference would add a huge number of additional peaks to the Kapitza-Dirac diffraction experiments shown below, as two-photon processes are possible between any pair of beams. The real-space geometry of the lattice can also be strongly affected, though the fundamental symmetry of the arrangement of beams will still apply.

From an experimental perspective it is simplest to use schemes requiring as few pairs of beams to be coherent as necessary, because every phase is a degree of freedom that may cause heating if not carefully controlled. In periodic lattices slow drifts can be absorbed as phononic strain and followed adiabatically by the atoms, but in quasiperiodic lattices this will also cause phasonic strains (section 2.2.3) and so not even slow drifts can be tolerated. This can be a major factor in experiment design: we originally intended to build a five-fold symmetric optical quasicrystal rather than eight-fold, but

eight-fold can be constructed from four mutually incoherent 1D standing waves, while five-fold requires five beams that are all mutually coherent. The former case requires only the retro mirrors to be stabilised, while the five beams in the latter case must all pass through separate optical fibres and then be stabilised against the noise this introduces. This was ultimately the main reason why we selected eight-fold instead, even though it involves more beams and seems superficially more complex to build.

3.5.2 Time of flight imaging in an optical lattice

In section 3.2.2 above we introduced the idea of time of flight (TOF) imaging, observing that a BEC will remain small while falling but a thermal cloud will expand. We revisit this idea now because TOF imaging can produce much more varied results in OLs and is a key tool for probing the atomic wavefunction, which requires a more quantitative understanding. The following derivation is based on [115].

Suppose the density distribution of atoms in a lattice is $n(\mathbf{r})$. We can express this as

$$n(\mathbf{r}) = \langle \hat{\Psi}_r^\dagger \hat{\Psi}_r \rangle, \quad (3.22)$$

where $\hat{\Psi}_r^\dagger$ is the field operator that creates the infinitesimal "slice" of wavefunction $\Psi(\mathbf{r})$ at position \mathbf{r} .

For simplicity, let us restrict ourselves to the case where the average density of atoms is uniform across the lattice (i.e. the system is homogeneous and at equilibrium). We can then express $\hat{\Psi}_r^\dagger$ in terms of more familiar operators as

$$\hat{\Psi}_r^\dagger = \sum_j w^*(\mathbf{r} - \mathbf{r}_j) \hat{a}_j^\dagger \quad (3.23)$$

where $w(\mathbf{r} - \mathbf{r}_j)$ is the Wannier function centred on the j th lattice site, and the sum is over all lattice sites. We can understand intuitively why this is the right form: if an atom occupies each lattice site equally, the wavefunction at position \mathbf{r} is the sum of contributions from every Wannier state.

This defines the situation in real space. To calculate the result of TOF imaging, we must then calculate the quantity

$$\tilde{n}(\mathbf{k}) = \langle \hat{\Psi}_k^\dagger \hat{\Psi}_k \rangle \quad (3.24)$$

which is the equivalent of $n(\mathbf{r})$ in momentum space. Since we have an expression for $\hat{\Psi}_r^\dagger$, we can find $\hat{\Psi}_k^\dagger$ by means of a Fourier transform:

$$\begin{aligned}
\hat{\Psi}_{\mathbf{k}}^\dagger &= \frac{1}{\sqrt{2\pi\hbar^3}} \int d^3\mathbf{r} e^{i\mathbf{k}\cdot\mathbf{r}} \hat{\Psi}_{\mathbf{r}}^\dagger \\
&= \frac{1}{\sqrt{2\pi\hbar^3}} \int d^3\mathbf{r} e^{i\mathbf{k}\cdot\mathbf{r}} \sum_j w^*(\mathbf{r} - \mathbf{r}_j) \hat{a}_j^\dagger \\
&= \frac{1}{\sqrt{2\pi\hbar^3}} \sum_j \int d^3\mathbf{r}' e^{i\mathbf{k}\cdot(\mathbf{r}'+\mathbf{r}_j)} w^*(\mathbf{r}') \hat{a}_j^\dagger \\
&= \left[\frac{1}{\sqrt{2\pi\hbar^3}} \int d^3\mathbf{r}' e^{i\mathbf{k}\cdot\mathbf{r}'} w^*(\mathbf{r}') \right] \times \sum_j e^{i\mathbf{k}\cdot\mathbf{r}_j} \hat{a}_j^\dagger \\
&= \tilde{w}^*(\mathbf{k}) \times \sum_j e^{i\mathbf{k}\cdot\mathbf{r}_j} \hat{a}_j^\dagger
\end{aligned} \tag{3.25}$$

where we have introduced $\tilde{w}(\mathbf{k})$ as the Fourier transform of the Wannier function. We can substitute this expression into equation 3.24 to obtain

$$\tilde{n}(\mathbf{k}) = \underbrace{|\tilde{w}(\mathbf{k})|^2}_{\text{Wannier envelope}} \times \underbrace{\sum_{j,m} e^{i\mathbf{k}\cdot(\mathbf{r}_j - \mathbf{r}_m)} \langle \hat{a}_j^\dagger \hat{a}_m \rangle}_{\text{Structure factor}}. \tag{3.26}$$

The second term harks back to the structure factor introduced in section 2.2.4, though here it is the structure factor for the *wavefunction*, not the lattice. We can apply the same physical interpretation of a plane wave originating at each lattice site with an amplitude and phase given by $\langle \hat{a}_j^\dagger \hat{a}_k \rangle$ and consequently depending on the coherence of adjacent lattice sites. This will be important for detecting the phase transition discussed in section 4.

The prefactor, known as the Wannier envelope, is independent of $\Psi(\mathbf{r})$ and captures the structure of the lattice itself. For an incoherent wavefunction and therefore a constant structure factor, a TOF image will simply show the Wannier envelope with no sharp features. (Recall that the structure factor of a random system, discussed in section 2.2.4, is also smooth and continuous.)

3.5.3 Kapitza-Dirac diffraction

Kapitza-Dirac diffraction (KDD) [105] is a technique used to image the reciprocal of an OL (section 2.1.1), otherwise known as the structure factor (section 2.2.4). KDD involves switching on the OL abruptly, allowing a few μs of evolution, then switching it off again to image in TOF. The resulting image shows the reciprocal lattice: see figure 3.6 for examples. Let us initially assume for simplicity that the lattice is periodic, in which case the reciprocal lattice is as well.

To understand this, let us take a moment to consider the effect of suddenly switching on the lattice. Let \hat{H} and \hat{H}' represent the Hamiltonians of the dipole trap and lattice respectively, with eigenbases $|\phi_j\rangle$ and $|\phi'_j\rangle$. Before we switch on the lattice the atoms are in the state $|\Psi\rangle = |\phi_0\rangle$, since the BEC is the ground state of the dipole trap. When

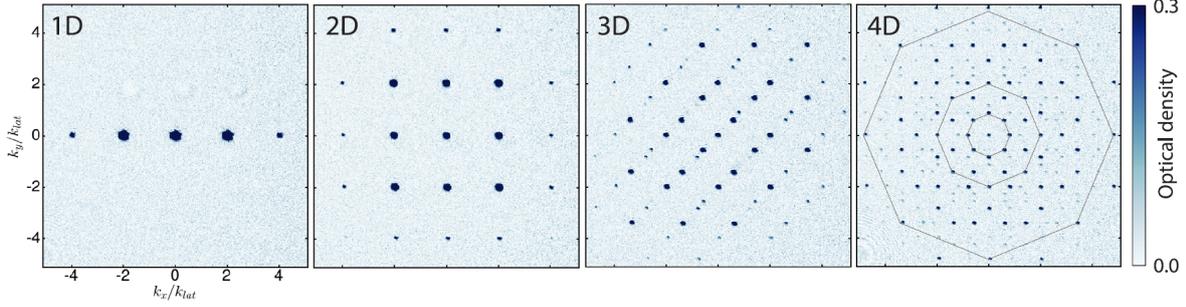


Figure 3.6: Reproduced from [58], copyright APS. Examples of time-of-flight images taken after Kapitza-Dirac diffraction in the optical quasicrystal (see main text). From left to right the lattices used were a single 1D standing wave, a 2D square lattice, a periodic array of 1D quasiperiodic lattices, and the full eightfold quasicrystal. The central peak in each image is the BEC, while each other peak is a momentum state populated by scattering and corresponds to a point in the reciprocal lattice. Note the striking similarity of these figures to the X-ray diffraction pattern in figure 2.4 – the two techniques are direct equivalents across two fields.

\hat{H} suddenly becomes \hat{H}' the value of $|\Psi\rangle$ does not instantaneously change, but since $|\phi_0\rangle$ is no longer an eigenstate we must express $|\Psi\rangle$ in the eigenbasis of the new Hamiltonian as

$$\begin{aligned} |\Psi\rangle &= \sum_j c_j |\phi'_j\rangle \\ \implies |\Psi(t)\rangle &= \sum_j c_j |\phi'_j(t=0)\rangle e^{i\omega_j t} \end{aligned} \quad (3.27)$$

where $\omega_j := E_j/\hbar$ is the angular frequency of the state $|\phi'_j\rangle$. Note, however, that this is not a totally general superposition: the new \hat{H}' may not conserve momentum $\hbar\mathbf{k}$, but it does conserve *quasimomentum* $\hbar\mathbf{q}$. The BEC has $\mathbf{k} = \mathbf{0}$ and hence $\mathbf{q} = \mathbf{0}$, so only the infinite set of states with $\mathbf{q} = \mathbf{0}$ can contribute to the sum in equation 3.27. States with equal \mathbf{q} can have values of \mathbf{k} separated by any combination of reciprocal lattice vectors $\mathbf{g} = 2\hbar\mathbf{k}_{\text{lat}}$, where \mathbf{k}_{lat} is the wavevector of any lattice beam, and so the states that make up $|\Psi\rangle$ could in principle include the entire reciprocal lattice. Over the duration of the pulse these states will all evolve with different frequencies, and the plane wave states with corresponding \mathbf{k} will be visible in TOF.

There is a complication, however. Even if the real-space lattice has no trapping potential, in momentum space the atoms will always experience a term $\propto k^2$ arising from kinetic energy. The momentum lattice is therefore analogous to a real-space lattice with a harmonic trap added over it, and this system does *not* have perfect translational symmetry. This enters our picture as a cutoff, k_{max} , limiting which states can be coupled to $|\mathbf{k} = \mathbf{0}\rangle$. In TOF we see this as a circular "window" through which reciprocal lattice peaks with $k < k_{\text{max}}$ are visible. The value of k_{max} is determined by the duration of the KDD pulse via the energy-time uncertainty relation, with shorter pulses giving a proportionally wider window and more accessible states, but also allowing less time for them to dephase and become visible. This can be countered by using a more intense pulse to make the evolution faster.

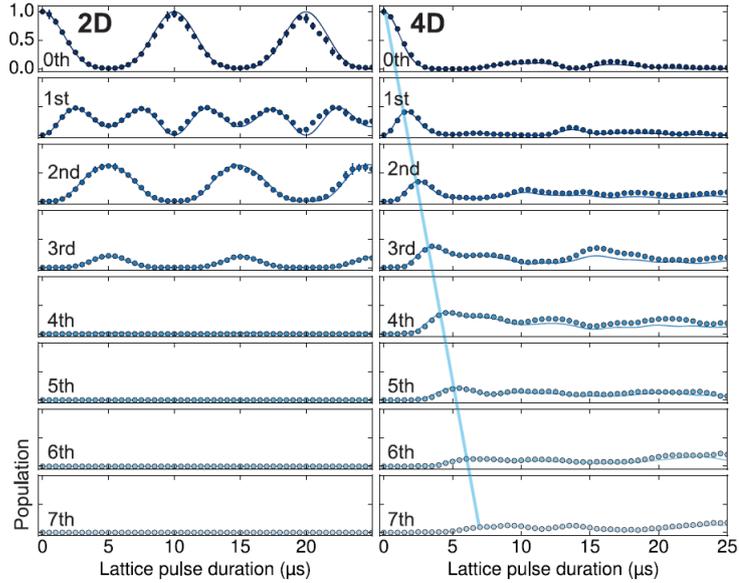


Figure 3.7: Reproduced from [58], copyright APS. This plot shows the populations of the various "generations" of peaks visible in TOF images as the duration of the the Kapitza-Dirac pulse changes. (The generation refers to the number of scattering vectors required to connect each peak to the BEC, so that generation 0 is the BEC itself.) Points represent data while the solid lines are numerical calculations from the time-dependent Schrödinger equation. The left-hand side is a 2D square lattice and the right-hand side is the optical quasicrystal, representing a 4D lattice. Observe that the 2D structure shows time-periodic dynamics, with generations above the third never becoming significantly populated: both effects are due to the kinetic energy cutoff described in the main text and the resulting finite number of accessible states. The 4D case, however, shows no revival of the BEC and no generational cutoff, because the quasicrystalline nature of the reciprocal lattice results in an infinite number of accessible states within the kinetic energy window. The blue line joining the peaks in the 4D curves is a guide to the eye, intended to make the reader agree that the dynamics resembles ballistic expansion.

The net result of this is that the *duration* of the KDD pulse determines which states are *accessible*, while the *area* of the pulse (the product of duration and intensity) determines which states are *reached*, i.e. which have sufficient time to become populated. The simplest method for observing dynamics experimentally is to keep the pulse intensity constant and scan the duration, but the reader should bear in mind that this conflates the two effects if the pulses become long enough: the window becomes narrower even as we have time to explore more of it.

There is another note to be made here. The kinetic energy window allows only a finite number of the periodically-spaced reciprocal lattice vectors to be populated, with the result that the sum in equation 3.27 has a finite number of frequencies ω_j . This results in *collapse and revival*, meaning that the initial BEC is recovered when the pulse duration is a multiple of all such frequencies and the time evolution carries on periodically until the narrowing of the energy window forces it to break down. (See the left-hand side of figure 3.7.)

Kapitza-Dirac diffraction in quasicrystalline lattices

We realised previously, in section 2.2.4, that the structure factor of a QC is a dense set of sharp peaks. The same applies to the reciprocal "lattice" observed by KDD and TOF imaging of a quasicrystalline lattice. This is because the incommensurate length scales discussed in section 2.2.2 also arise in momentum space via the reciprocal lattice vectors, with the result that a QC in real space also gives a QC in momentum space. This has a major effect on the dynamics just described, as although kinetic energy still limits the maximum achievable k for a given pulse length, it does *not* follow that the total number of accessible states is finite. Instead the kinetic energy window is filled with a dense set of states extending to arbitrarily small kinetic energies, with higher and higher numbers of reciprocal lattice vectors allowing states closer and closer to the BEC. As a result, KDD in quasicrystalline lattices does not show collapse and revival (right-hand side of figure 3.7).

This can be understood quite intuitively using the higher-dimensional picture of QCs introduced in section 2.2.3. The finite number of states in periodic lattices arises because new states can only be reached by moving to higher values of k , but doing so also increases the kinetic energy, eventually reaching the cutoff k_{max} . In the QC, however, it is always possible to reach new states by moving mostly in the additional higher dimensions, with only a small component projecting onto the 2D plane of the "real" lattice. These directions will not experience the harmonic confinement imposed by kinetic energy on the 2D plane, but will see some more complicated energy dependence that never increases significantly above zero as the projected position in the 2D plane stays close to the BEC. This allows access to momentum states far from the BEC, as long as that distance is chiefly in the extra dimensions of the parent crystal and not in real space.

3.6 The optical quasicrystal

While up until now I have discussed OLs only in the abstract, here I will describe the OL implemented in our experiment, shown in figure 3.5, which consists of four retro-reflected lattice beams at 45° angles¹⁶. Retro-reflection makes this an eight-fold symmetric arrangement. Our lattice wavelength is 725 nm, far to the blue of the atomic resonances of ^{87}Rb at 780 nm and ^{39}K at 767 nm. Each 1D standing wave is made incoherent from the others by detuning using AOMs, with frequency differences on the scale of tens of MHz.

Referring back to section 2.2.1, long-range order and a rotational symmetry other than 1, 2, 3, 4 or 6 are sufficient for a pattern to be considered a quasicrystal. Since the arrangement of lasers is eight-fold symmetric, and the temporal coherence of a laser

¹⁶For the experiments in chapter 6 we also added a deep 1D lattice in the vertical direction, providing confinement to a single 2D plane. This is necessary to observe strongly correlated states like the Mott insulator. However, the 1D lattice was not yet present, nor is it required, for the experimental results shown later in this chapter, which do not involve any such states.

ensures long-range order in its interference pattern, we refer to the resulting OL¹⁷ as our optical QC. We also have the option of turning on only a subset of the four 1D OLs, which we usually use to create either a 1D OL or a 2D square OL using two of the perpendicular pairs.

The parent crystal of the optical QC is four-dimensional, as may be seen from equation 2.26: the OL is two-dimensional ($d = 2$), and the $\pi/4$ angle between beams adds an incommensurate length scale of $\sin(\pi/4) = 1/\sqrt{2}$ ($n_i = 2$). The equivalent expression to equation 2.27 therefore has four terms [58],

$$\mathbf{r} = n_1\mathbf{g}_1 + n_2\mathbf{g}_2 + n_3\mathbf{g}_3 + n_4\mathbf{g}_4, \quad (3.28)$$

which each \mathbf{g}_j running parallel to one of the four standing waves. Since the incommensurate length scales make these four vectors linearly independent, equation 3.28 describes a 4D generalisation of a simple cubic lattice, which we are fond of referring to as a simple "hypercubic" lattice. Switching off one or more standing waves allows us to reduce the parent dimension to 3, 2 or 1, allowing us to measure scalings with d from 1-4.

Our group has so far used this optical QC in two main experimental investigations, resulting in the publications [58] and [57]. I will describe them here, as I was involved in both, though I will not attempt to reproduce the comprehensive treatment of my colleagues Konrad Viebahn and Matteo Sbroscia in their respective theses [69, 70]. The first focusses on the link between the QC and its higher dimensional parent, while the second focusses on the QC's intermediate nature between ordered and disordered systems. Both experiments were performed with ³⁹K, and we used that atom's Feshbach resonance to remove all interactions except where noted otherwise.

3.6.1 Four dimensional quantum walks

The reciprocal lattice of the optical quasicrystal shows all the features described in section 3.5.3: a dense set of sharp peaks with eight-fold symmetry, more and more of which become visible within a circular cutoff defined by pulse duration when imaged using KDD and TOF. This momentum-space QC is more convenient to work with than the real-space QC as TOF imaging requires less precision than high-resolution in-situ imaging, so our first QC experiments treated the reciprocal lattice as our primary lattice. This allows quantum simulation of one, two, three or four dimensions based on equation 3.28. KDD in such a lattice, starting from the central BEC, can then be considered as a continuous quantum random walk in up to four dimensions.

Random walks

A random walk is the motion of a particle that keeps randomly selecting its direction, altering course either at discrete intervals (a discrete random walk) or constantly and

¹⁷In the strictest sense it isn't a lattice, since that term usually implies periodicity, but OL is such a familiar term that we use it anyway.

smoothly (a continuous random walk). A classical particle undergoing such motion tends to move rather inefficiently, as it will keep doubling back on itself. The average displacement of such a particle after time t , moving at constant speed, is therefore only $\propto \sqrt{t}$. (This may be referred to as *diffusive* motion.) A quantum particle, however, makes no definite decisions unless it is measured, and so each time it selects a new direction its wavefunction will split along all available courses. Because there will always be a part of the wavefunction that has not yet doubled back, the distance spanned by the wavefunction scales much more favourably with time, growing $\propto t$. A quantum random walk is therefore quite different from the classical case, and in fact is analogous instead to classical ballistic expansion. That is the case we observe using KDD in the momentum QC, with each atom spreading out coherently until forced to pick a final destination by the imaging beam.

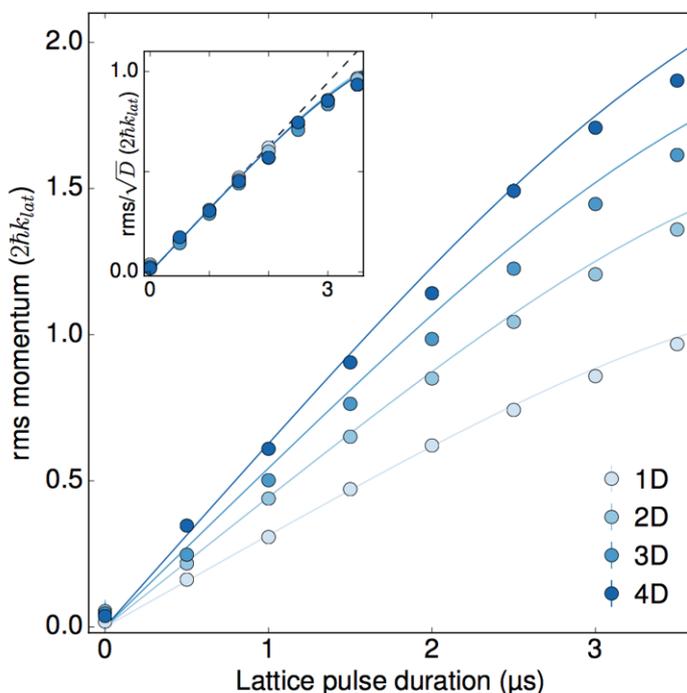


Figure 3.8: Reproduced from [58], copyright APS. This figure shows a continuous quantum walk being carried out in 1-4 dimensions using Kapitza-Dirac diffraction: points are data and solid lines are numerical calculations from the time-dependent Schrödinger equation. The number of effective dimensions is controlled by changing the number of linearly independent scattering vectors (equation 3.28), which in practice means the number of standing waves from figure 3.5 that are switched on. The vertical axis then plots the RMS momentum of the visible peaks, but note that this is calculated in the coordinates of the *parent* by indexing each peak according to equation 3.28 and treating each direction equivalently. The horizontal axis plots the duration of the KDD pulse. For the first 2 μs or so the plot is linear, indicating ballistic expansion in momentum space, and the gradient is proportional to \sqrt{d} for d dimensions (see inset). The curving of the lines after $\sim 2 \mu\text{s}$ is due to kinetic energy effects (see main text).

Results

The results of momentum-space random walks in 1-4D are shown in figure 3.8. We measure the expansion using the root-mean-square momentum of all atoms observed during TOF. In the quasicrystalline 3D and 4D cases this quantity grows linearly with time until kinetic energy becomes significant after $\sim 2\mu\text{s}$, at which point the rate of growth slows ("Single Pulse" line in figure 3.9). In the 1D and 2D cases we observe periodic dynamics much as in figure 3.7. The initial slope scales $\propto \sqrt{d}$ for d dimensions, which is simple to verify by considering that the atoms have d directions in which to spread out (see appendix A for a proof). This correct scaling with d demonstrates the point that our optical QC is able to simulate additional spatial dimensions. While kinetic energy makes the picture more complicated, even this can be ignored until the wavefunction spreads far enough to reach the kinetic energy window.

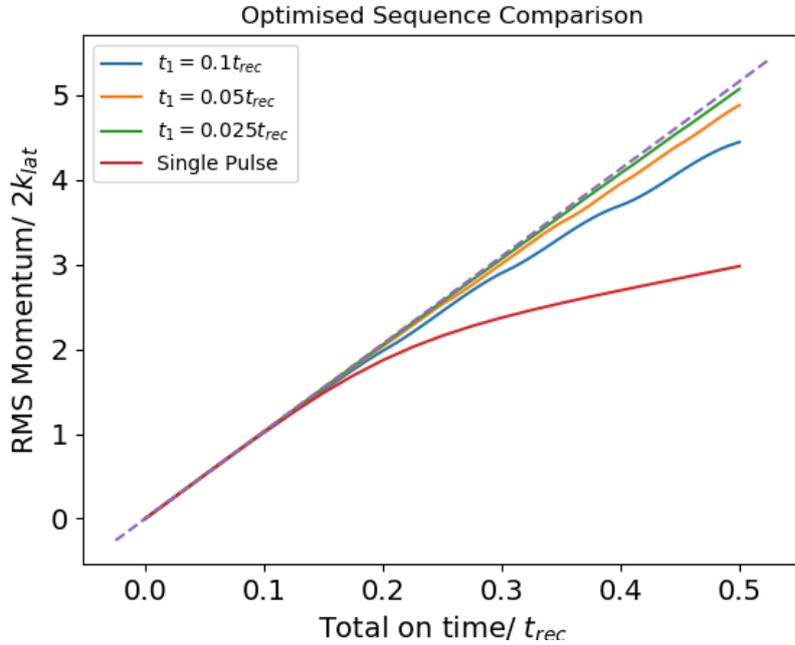


Figure 3.9: This figure, produced by Jonathan Mortlock [116], illustrates the benefit of using a sequence of pulses rather than a single long pulse in the quantum walk of figure 3.8. The lines show the evolution of the RMS momentum for a series of pulses of length t_1 (legend), with the gaps between pulses having been optimised for the longest possible linear region and the sum of durations of all pulses shown on the horizontal axis. Note that all times are given in units of the recoil time $t_{rec} \approx 103\mu\text{s}$ for ^{39}K in a 725 nm lattice. As t_1 becomes shorter the linear region extends to longer durations, pushing back the effects of kinetic energy.

Mitigating kinetic energy effects

It would be interesting if we could increase the size of the kinetic energy window and extend the straight-line region of figure 3.8. As noted, the window is larger for shorter

pulses, so in principle we could move to shorter and shorter pulses with higher intensity to compensate. This is impossible in practice, however: the above experiments are already close to the limit of the shortest, most intense pulses we can create. To overcome this limitation, we can move away from a single pulse of increasing duration and instead use sequences of short pulses, so that the sum of their durations increases but the individual pulses still give a large window. A simple brute-force numerical optimisation using the TDSE can be used to find the optimum combination of pulse durations and gaps for the longest possible linear region. This simulation was carried out under my supervision by Jonathan Mortlock with promising initial results [116] (figure 3.9), but has not yet been implemented on the experiment.

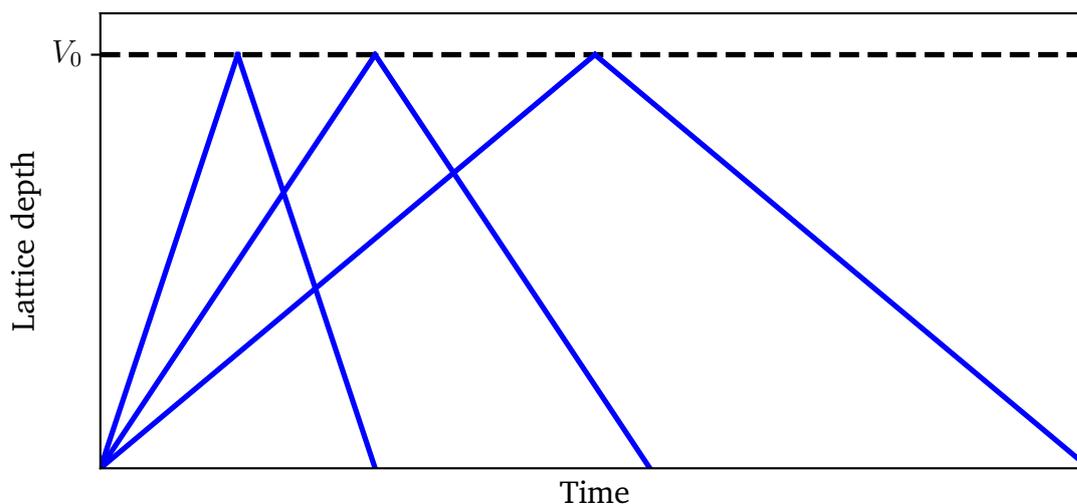


Figure 3.10: An illustration of the triangular lattice pulses used to measure the localisation transition in the optical quasicrystal. The lattice is ramped linearly up to a final depth V_0 in a time τ , then back down to 0. If the process is sufficiently slow for the BEC to adiabatically follow the ground state, a pulse of this shape causes no excitations. However, the time scale of adiabaticity diverges if the ground state at V_0 is localised.

3.6.2 Localisation experiments

Method: triangular pulses

In this section I will describe a method, published in [57], for imaging the QC localisation transition (section 2.2.5) using only TOF imaging. Rather than the abrupt KDD pulses used for the previous experiment, this investigation uses triangular pulses of the kind shown in figure 3.10. These pulses are parametrised by a maximum lattice depth, V_0 , and the time τ taken to reach that depth (with the same time then taken to ramp back down). If τ is sufficiently large, we would naively expect the wavefunction Ψ to adiabatically follow the ground state $|\phi_0\rangle$ of the Hamiltonian as the lattice ramps up and down, returning to a BEC in the ODT when the lattice is fully switched off again[117]. There are two major cases where this does not happen.

- If τ is too small, terms related to the gradient of the Hamiltonian will become important (this is explored in section 5.1). $|\Psi\rangle$ will not be able to follow quickly enough to keep up with $|\phi_0\rangle$ and will not make it back to the BEC by the time the lattice is switched off, resulting in excitations.
- If there is a continuous QPT at some lattice depth $V < V_0$, the considerations of section 1.4.2 will apply and we do not expect to be able to cross the transition adiabatically. This results in excitations no matter the value of τ .

We can therefore detect phase transitions as a function of lattice depth using a two-dimensional scan of both τ and V_0 followed by TOF imaging. A sufficiently small τ is expected to result in excitations for any value of V_0 , but a phase transition will be indicated by a value of V_0 above which no τ is long enough to achieve adiabaticity. (Note, however, that since we cannot send τ to ∞ in the experiment we have only finite resolution, and ultimately cannot distinguish a QPT from a sufficiently rapid crossover.)

Results

The results of scanning V_0 and τ are shown in figure 3.11, which plots the fraction of atoms that successfully return to the BEC as a function of τ for several values of V_0 . (This figure was created by my colleague Matteo Sbroscia and is reproduced from [57].) The essential result is that for $V_0 \lesssim 1.8E_r$ there are excitations for sufficiently short τ that disappear for longer ramps, while above that critical value the final BEC population only decreases as we make τ larger. This indicates a QPT at a lattice depth of $\sim 1.8E_r$, consistent with our expectation of a localisation transition: physically, the wavefunction of a single particle would have to change from spanning the entire lattice to occupying a far smaller localised region in order to remain adiabatic. The precision of the critical V_0 value increases as τ is made longer, as the divergence between lines returning to 1 and those continuing downwards is made clearer. While this data stops at $\tau = 250\mu\text{s}$, this is not due to any technical limitation and I expect we will probe longer times with future experiments¹⁸.

Note that figure 3.11 does not show *exactly* the behaviour predicted in section 1.4.2. While we indeed observe that we cannot make τ long enough to cross the QPT adiabatically, we observe *more* excitations as we make τ longer, where in section 1.4.2 we expected fewer. While we do not know for certain why this is the case, there are two likely explanations, both related to the dense set of momentum states available in the quasicrystal (see e.g. figure 3.6).

- It may be that we would see normal Kibble-Zurek behaviour of increasing adiabaticity if we used much longer sweeps. (In the finite experimental system this must certainly be the case.) Even the curves below the transition in figure

¹⁸We stopped at 250 μs when taking this data because we were rudely interrupted by a faulty viewport causing a vacuum breach and have not yet returned to the project, otherwise we would certainly have continued to longer ramp times.

3.11 show complicated dynamics at short enough τ , including regions where the sweep becomes less adiabatic as τ increases. It could be that this simply continues to much longer values of τ for the localised state. Numerical data in [57] shows the negative gradient continuing to $\tau > 10$ ms at least, but the asymptotic behaviour as $\tau \rightarrow \infty$ is not known.

- It may be a consequence of the unique momentum structure of the QC. Not only are there momentum states at arbitrarily small kinetic energies, these states take arbitrarily large sums of reciprocal lattice vectors to access, requiring more and more time in the lattice before they can be populated. (Compare with the ballistic spread to higher generations of peaks in figure 3.7.) If this is the case, it would mean that in the thermodynamic limit the Kibble-Zurek mechanism does not apply to this system and the decreasing adiabaticity will continue forever.

Whichever of these is the case, we believe that this departure from the expected behaviour at a continuous QPT is an exciting sign that more investigation is needed, and an example of the remarkable nature of QC physics.

We also investigated the role of interactions, using the Feshbach resonance of ^{39}K to add a repulsive interaction. The only data taken for this investigation¹⁹ is shown in figure 3.12, again reproduced from Matteo Sbroscia's paper [57]. For this measurement we started deep in the localised regime – the data point at $V_0 = 2.4$ and $\tau = 250\mu\text{s}$ in figure 3.11 – and increased the scattering length a from 0 up to $200a_0$. This had a dramatic effect, increasing the BEC fraction after the pulse almost all the way back up to unity. This is consistent with GPE results (inset to figure 3.12) showing that the localisation transition is pushed to higher lattice depths by the addition of repulsive interactions. This is fairly generic behaviour, as repulsive interactions tend to favour spreading out to reduce interaction energy and therefore oppose localisation.

This measurement raises the possibility of using the optical QC to investigate MBL, since the necessary ingredients of (quasi-)disorder, interactions, and excitations are all present. The exact boundary between MBL and the Bose glass state (see chapter 7) is still poorly defined, but this is likely to be a productive future project regardless of which name is most appropriate.

¹⁹Further measurements were prevented by the same equipment failure.

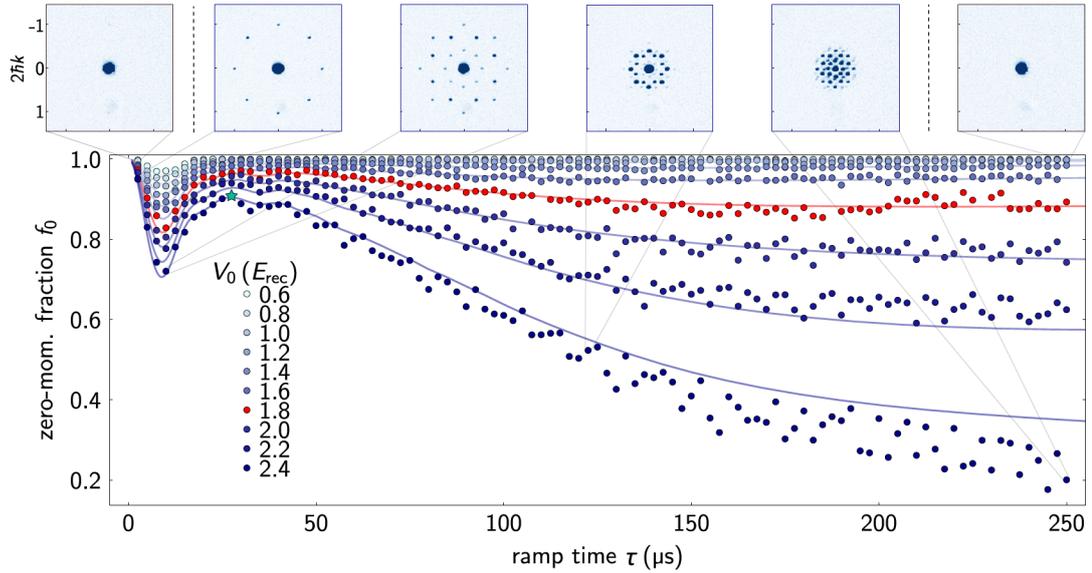


Figure 3.11: Reproduced from [57], copyright APS. This figure shows the results of adiabatically loading and unloading the optical quasicrystal at various lattice depths. (See main text.) Points are data and lines are calculated from the TDSE. The horizontal axis is the duration of the loading process, equal to half the total duration of the triangular lattice pulse of figure 3.10. After loading and unloading the atoms are imaged in time of flight (pictures along the top). The vertical axis records the proportion of atoms found in the central peak (the BEC), and each curve is a different lattice depth (height of the triangle peak). If the loading and unloading processes are adiabatic, all atoms return to the BEC and the result is 1, but if we ramp too quickly there will be occupation of the excited states (satellite peaks). For lattice depths lower than $\sim 1.8E_r$, there is some initial dynamics for ramps shorter than about $50\mu\text{s}$, then the ramp becomes adiabatic and all atoms return to the BEC for all subsequent times. Above about $1.8E_r$, the proportion returning to the BEC only decreases as the ramp becomes slower, and we never achieve adiabaticity. This behaviour indicates that the system has undergone a phase transition to the localised state, which causes too drastic and abrupt a change to the ground state wavefunction to be followed adiabatically.

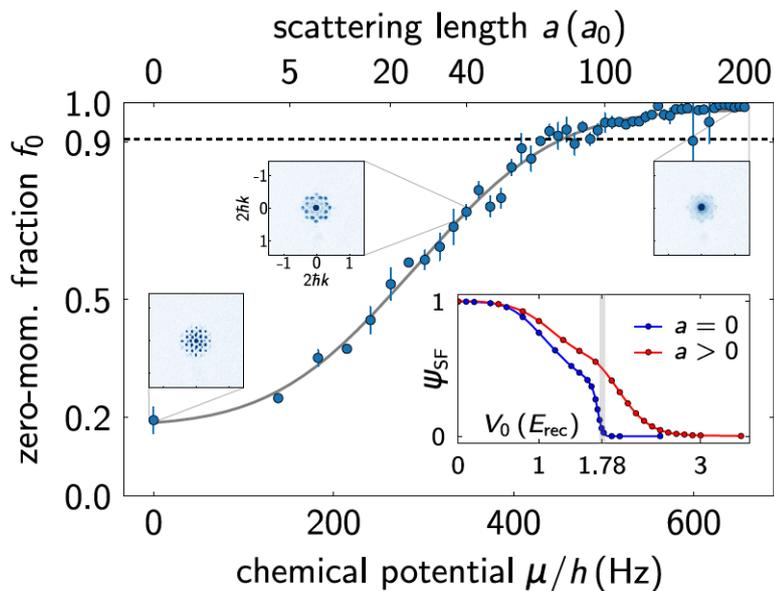


Figure 3.12: Reproduced from [57], copyright APS. This figure shows the effect of repulsive interactions on the localisation transition of figure 3.11. The horizontal axis shows the scattering length a , or equivalently the chemical potential μ ($\mu \propto a^{2/5}$ in a harmonic trap from Thomas-Fermi theory [112]). This figure tracks the lower right-most data point of figure 3.11 ($\tau = 250 \mu\text{s}$, $V = 2.4E_r$) as a repulsive interaction is switched on, showing almost a full recovery up to 1 by the time a reaches $200a_0$. This indicates that the localisation transition has moved to higher values of V_0 , supported by the inset showing a numerical calculation of superfluid stiffness (characteristic of the superfluid phase; see section 4.1.2) done by our collaborator Alex Gaunt using the Gross-Pitaevskii equation. The superfluid stiffness falls quite sharply to zero in the noninteracting system (blue) at $V_0 = 1.78E_r$, suggestive of a kink in the gradient rounded by finite system size effects, but stiffness persists for much longer in the interacting case (red).

4

The Mott insulator-superfluid transition

In this chapter I will describe the physics of the quantum phase transition between the Mott insulator and superfluid states, since that will play a key role in chapter 6. I will start by outlining the physics of the two states individually, then give an account of the transition between them.

4.1 The superfluid phase

A thorough account of the material covered in this section can be found in [112].

4.1.1 Superfluidity in the weakly interacting Bose gas

Recall the Hamiltonian 3.16 for a gas of ultracold bosons with a weak contact interaction. We will now see that the interaction results in superfluidity. By calculating the ground-state energy of this Hamiltonian with a variational method [112] we can arrive at the time-dependent Gross-Pitaevskii equation,

$$\left[-\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{r}) + g|\Psi(\mathbf{r}, t)|^2 \right]\Psi(\mathbf{r}, t) = i\hbar\partial_t\Psi(\mathbf{r}, t), \quad (4.1)$$

which resembles the time-dependent Schrödinger equation with the addition of a non-linear term. Here we have assumed that all particles occupy the same state $\Psi(\mathbf{r}, t)$, with no significant correlations caused by interactions: this is called the *mean-field* approximation. However, we do *not* assume that $\Psi(\mathbf{r}, t)$ is the same as the non-interacting wavefunction.

The ground-state solution to this equation has the form

$$\Psi_0(\mathbf{r}, t) = \sqrt{n(\mathbf{r})} e^{-\frac{i\mu t}{\hbar}} \quad (4.2)$$

where $n(\mathbf{r}) := |\Psi(\mathbf{r}, t)|^2$. Note that this solution has a single phase factor, $e^{-\frac{i\mu t}{\hbar}}$. Because μ is constant across the system at thermal equilibrium, the entire wavefunction may be thought of as having a single, "macroscopic" phase. It is often useful to define a phase correlation length ξ , which I mention only to note that in this case $\xi = \infty$ [90].

To find the excitation spectrum we consider a perturbation,

$$\Psi(\mathbf{r}, t) = \Psi_0(\mathbf{r}, t) + \delta\Psi(\mathbf{r}, t), \quad (4.3)$$

which we take to have the form

$$\delta\Psi(\mathbf{r}, t) = e^{-\frac{i\mu t}{\hbar}} \left(u(\mathbf{r}) e^{-i\omega t} - v^*(\mathbf{r}) e^{i\omega t} \right). \quad (4.4)$$

Substituting this form into 4.1 and collecting terms proportional to $e^{\pm i\omega t}$ we recover a pair of coupled equations

$$\begin{aligned} \left[-\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) + 2gn(\mathbf{r}) - \mu - \hbar\omega \right] u(\mathbf{r}) &= gn(\mathbf{r})v(\mathbf{r}) \\ \left[-\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) + 2gn(\mathbf{r}) - \mu + \hbar\omega \right] v(\mathbf{r}) &= gn(\mathbf{r})u(\mathbf{r}). \end{aligned} \quad (4.5)$$

For a general potential $V(\mathbf{r})$ these equations have no analytic solution. The usual approach is to take $V(\mathbf{r}) = 0$ and assume that the qualitative features of this result will carry over to other cases (see note about non-uniform systems below). In a uniform system it's reasonable to guess that $u(\mathbf{r})$ and $v(\mathbf{r})$ have plane-wave solutions,

$$\begin{aligned} u(\mathbf{r}) &= u_k e^{ik \cdot \mathbf{r}} \\ v(\mathbf{r}) &= v_k e^{ik \cdot \mathbf{r}} \end{aligned} \quad (4.6)$$

for some wavenumber k . The equations 4.5 become

$$\begin{aligned} \left[E_k + gn - \hbar\omega \right] u_k &= gn v_k \\ \left[E_k + gn + \hbar\omega \right] v_k &= gn u_k \end{aligned} \quad (4.7)$$

where

$$E_k := \frac{\hbar^2 k^2}{2m}. \quad (4.8)$$

Now that we have reduced the problem to a pair of linear algebraic equations we can find the condition for non-trivial solutions with the matrix equation

$$\begin{vmatrix} E_k + gn - \hbar\omega & -gn \\ -gn & E_k + gn + \hbar\omega \end{vmatrix} = 0 \quad (4.9)$$

from which we find that

$$\hbar\omega = \sqrt{E_k(E_k + 2gn)}. \quad (4.10)$$

This result is known as the Bogoliubov excitation spectrum of the weakly-interacting Bose gas, and is plotted in figure 4.1. To understand this dispersion relation, consider the limits of long- and short-wavelength excitations (small and large k respectively).

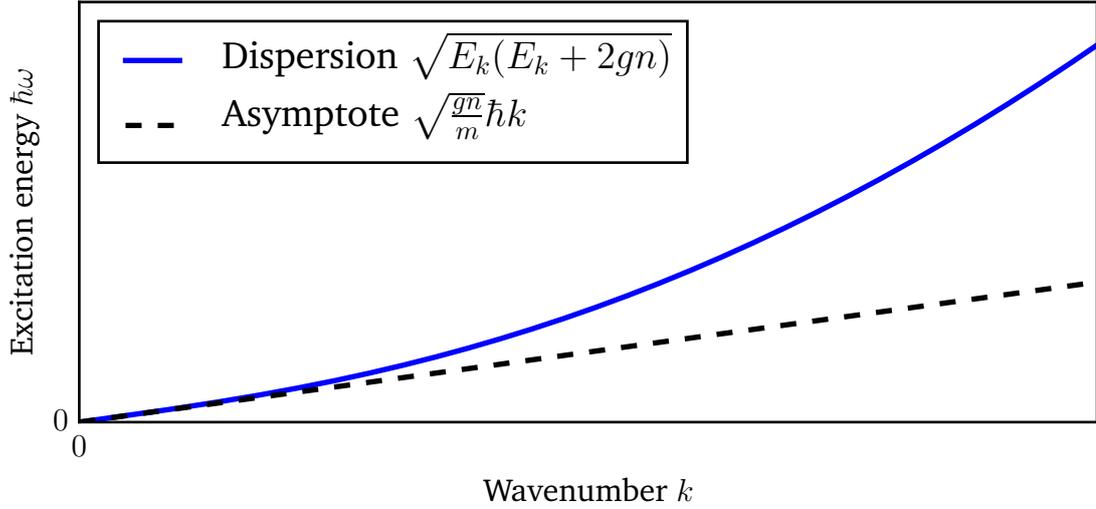


Figure 4.1: Excitation spectrum of the weakly-interacting Bose gas in a uniform system (equation 4.10). In the short-wavelength limit $k \rightarrow \infty$ the curve becomes a parabola, which is the behaviour of a free particle. On the other hand, in the long-wavelength limit $k \rightarrow 0$ the dispersion is linear, a characteristic behaviour of acoustic phonons (collective excitations in crystalline materials).

- For large k , $E_k \gg 2gn$ and $\hbar\omega = E_k$, resulting in the parabolic dispersion of a free particle. This makes sense: for wavelengths on the scale of individual particles we won't notice the influence of interactions.
- For small k , $E_k \ll 2gn$ and $\hbar\omega = \sqrt{\frac{gn}{m}}\hbar k$, a linear dispersion. This suggests that long-wavelength excitations of the system, which presumably involve collective behaviour of many particles, have a characteristic speed $v_c := \sqrt{\frac{ng}{m}}$. We identify these as *acoustic phonons*, otherwise known as sound waves.

Note that the gradient $\partial_k \omega \geq v_c$ everywhere. Recall from section 2.1.4 that the group velocity of a wavepacket with wavenumber k_0 is

$$\begin{aligned} v_g &= \frac{1}{\hbar} \partial_k E(k)|_{k_0} \\ &= \partial_k \omega(k)|_{k_0}. \end{aligned} \tag{4.11}$$

We deduce from this that the gradient of the dispersion relation gives the speed of an excitation, and therefore that the system can't support excitations moving at speeds $< v_c$. This in turn means that an impurity particle moving through the gas slower than v_c can't cause excitations; in the lab frame, then, there can be no drag if the gas flows around an obstruction provided that the speed of flow is $< v_c$ everywhere. Flowing without viscosity below a certain speed is a characteristic property of *superfluids* (SFs), and v_c is called the *Landau critical velocity*. Note that $v_c \propto \sqrt{g}$, so without weak interactions the dispersion becomes purely parabolic and the system isn't an SF.

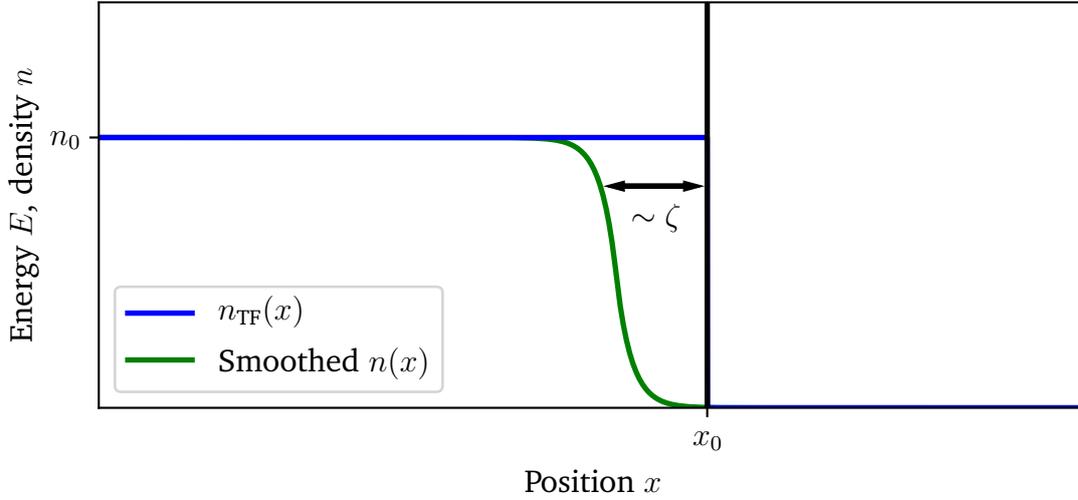


Figure 4.2: A BEC in a uniform potential encountering a hard wall (infinite jump) in potential. The Thomas-Fermi approximation (blue curve) predicts the distribution $n_{\text{TF}}(x)$, where the particle density jumps discontinuously from some n_0 to 0 at the wall. In practice the recovery will happen over a distance comparable to the healing length ζ (green curve).

Extension to non-uniform systems

In the above derivation of the Bogoliubov dispersion we made the assumption that $V(\mathbf{r}) = 0$. The simplest way to extend the same result to non-uniform systems is to replace n with $n(\mathbf{r})$ without any other changes, which is called the *local density approximation* (LDA). This works well provided that the density never varies too quickly, so that any terms depending on $\nabla n(\mathbf{r})$ can be neglected. This is equivalent to saying that the kinetic energy $E_{\text{kin}} \propto \nabla^2 \Psi(\mathbf{r}, t)$ always remains negligible, which is also referred to as the *Thomas-Fermi approximation*.

Substituting equation 4.2 into equation 4.1 and neglecting the spatial derivative we obtain

$$V(\mathbf{r}) + gn(\mathbf{r}) = \mu, \quad (4.12)$$

where the chemical potential is the same everywhere at equilibrium. (Note that this means our assumption of infinite phase correlation length ξ still holds.) The LDA then effectively defines a *local* chemical potential

$$\mu(\mathbf{r}) := \mu - V(\mathbf{r}) = gn(\mathbf{r}), \quad (4.13)$$

giving a density distribution that is simply the trapping potential turned upside down. A smooth potential will then give a smooth density that satisfies the LDA, except perhaps at the points where $V(\mathbf{r}) = \mu$. This method can result in quite unlikely kinks in the gradient of $n(\mathbf{r})$ because kinetic energy isn't being penalised.

If the potential $V(\mathbf{r})$ does change abruptly, the density $n(\mathbf{r})$ will respond more slowly because a rapid change would incur a large kinetic energy penalty. The simplest case

to consider is a hard wall where the potential jumps from 0 to ∞ , and $n(\mathbf{r})$ must decrease from some bulk value n_0 to 0 at the boundary (figure 4.2). We define the *healing length* ζ as the distance over which $n(\mathbf{r})$ recovers from 0. To evaluate this distance, we equate

$$\text{Kinetic energy at boundary} = \text{Interaction energy in bulk.} \quad (4.14)$$

The kinetic energy operator is

$$\hat{T} = \frac{\hbar^2}{2m} \nabla^2, \quad (4.15)$$

and so for a variation over a length ζ it is reasonable to assume that the kinetic energy is on the order of

$$E_{\text{kin}} \approx \frac{\hbar^2}{2m\zeta^2}. \quad (4.16)$$

We can then write

$$\begin{aligned} \frac{\hbar^2}{2m\zeta^2} &= gn_0 \\ \implies \zeta &= \frac{\hbar}{\sqrt{2mgn_0}}. \end{aligned} \quad (4.17)$$

We can combine the LDA with this value of the healing length to create a general strategy for finding $n(\mathbf{r})$ and hence the critical velocity: first apply the naïve result 4.13 to find an initial guess for $n(\mathbf{r})$, then look for any discontinuities in either $n(\mathbf{r})$ or $\nabla n(\mathbf{r})$ and smooth them out over a length ζ . We can then apply the result

$$v_c(\mathbf{r}) = \sqrt{\frac{gn(\mathbf{r})}{m}} \quad (4.18)$$

to a wide variety of commonly used trapping potentials. The case of a harmonic trap is shown in 4.3 with and without this smoothing.

4.1.2 Properties of superfluids

Consider again the SF ground state 4.2. Due to its importance, I will reiterate that the most remarkable feature of this wavefunction is that its phase $\phi = \mu t / \hbar$, referred to as the macroscopic phase, is independent of \mathbf{r} , meaning that the same phase is shared by the entire wavefunction. This is also captured by the statement that the phase correlation length, ξ , is infinite in the SF. When a classical gas undergoes the BEC transition described in section 3.2 it must pick an initial value for ϕ , which is an example of spontaneous symmetry breaking (section 1.2).

The macroscopic phase ϕ also exhibits a property called *superfluid stiffness*, which is an example of the general correspondence between broken symmetries and rigidities [118]. If ϕ is made to rotate by some amount θ over some distance, we can define the superfluid stiffness (up to a numerical factor) as [113]

$$\rho_s := \partial_\theta^2 E_0 |_{\theta=0}. \quad (4.19)$$

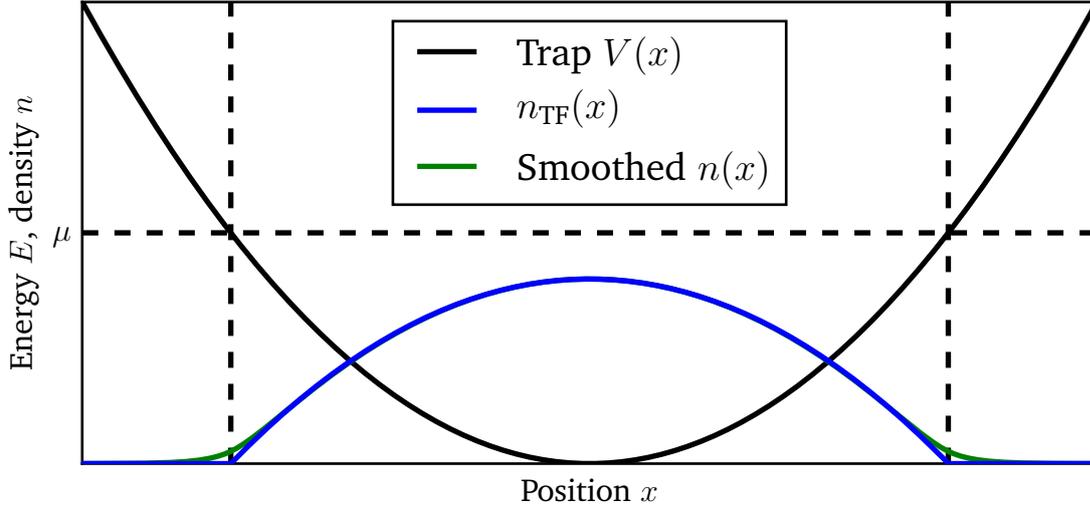


Figure 4.3: The density distribution of a weakly-interacting BEC in a harmonic trap $V(x)$. The blue curve shows the naïve Thomas-Fermi calculation $n_{\text{TF}}(x)$, with sharp kinks where $V = \mu \implies n = 0$. The green curve (which overlaps with the blue everywhere except at the kinks) shows the same density distribution convolved with a point spread function $e^{-|x|/\zeta}$, where ζ is the healing length. This smooths out the kinks and results in a much more physically reasonable density distribution.

where E_0 is the ground-state energy of the SF. While not explored any further in this thesis, ρ_s is also relevant to the study of superconductors, where it can be measured experimentally via the London penetration depth [119, 120], and in quantum magnetism [121].

4.1.3 Superfluids in lattices

As described in section 2.1.5, superfluidity in the Bose-Hubbard model corresponds to the case where $J \gg U$ in equation 2.21. If we take the limiting case $U = 0^1$, the Bose-Hubbard Hamiltonian \hat{H}_{BH} becomes

$$\hat{H}_{\text{SF}} = -J \sum_{\langle i,j \rangle} (\hat{a}_i^\dagger \hat{a}_j + \hat{a}_j^\dagger \hat{a}_i) - \mu \sum_i \hat{n}_i. \quad (4.20)$$

The ground state of this non-interacting Hamiltonian is simply a product state of all atoms in the lowest-energy Bloch wave, $|q = 0\rangle$, which may be written for N atoms on M sites as

$$\begin{aligned} |\Psi_{\text{SF}}\rangle &= |q = 0\rangle^N \\ &\propto \left(\sum_{i=1}^M a_i^\dagger \right)^N |0\rangle, \end{aligned} \quad (4.21)$$

¹As mentioned previously, this case should not strictly be called an SF as its critical velocity vanishes. However, all the results that are usually of interest to us for SFs in lattices continue to apply in this regime, with the result that the non-interacting gas is often referred to as an SF anyway. If the reader prefers, they may assume that $U > 0$ but is small enough to be neglected for mathematical convenience.

with $|0\rangle$ referring here to an unoccupied lattice.

4.1.4 Imaging superfluids

Time of flight (TOF) imaging of the SF state produces a density distribution given by equation 3.26. Recall that this distribution consists of a broad "Wannier" envelope arising from the on-site wavefunction, and a structure factor that allows sharp features. The width of such features is limited by the phase correlation length ξ , and cannot be smaller than $\sim 1/\xi$. Since ξ is infinite in the SF, this provides no limit at all, and the SF state will show sharp peaks in TOF with a symmetry characteristic of the lattice². The Wannier envelope is still present, becoming a "window" within which peaks are visible, but under ideal conditions there is no sign of an incoherent background.

4.2 The Mott insulator

The Mott insulator (MI) is the other possible ground state of the Bose-Hubbard Hamiltonian given in equation 2.21. It becomes the ground state for a sufficiently high ratio U/J , with the critical value depending strongly on μ . The MI is in striking contrast with the SF, being its opposite in almost every way; I will illustrate this first qualitatively, then quantitatively, and finally explain how we can detect the MI state experimentally. Accounts of the physics of MIs can be found in e.g. [89, 115, 122].

4.2.1 Qualitative description and general features

As its name suggests, the MI is an insulating state, which means that no long-distance transport is possible. Insulating states are common in conventional condensed matter physics, where the carriers are fermions, allowing insulating states known as band insulators to occur even in clean, noninteracting systems. The MI, however, requires a strong interparticle interaction and is therefore a distinctive feature of strongly correlated many-body systems. Unlike band insulators, MIs occur for both bosons [89] and fermions [122], though in this thesis we are interested only in the bosonic case.

The physical intuition for the MI is to imagine a tight-binding lattice with the same number n of particles on each lattice site (figure 4.6). From the Bose-Hubbard Hamiltonian this results in an interaction energy of $\frac{1}{2}Un(n-1)$ per site. If a single particle hops to an adjacent site, those two sites will now have $n+1$ and $n-1$ particles re-

²The width of the peaks does not go to zero when ξ is infinite, but is instead limited by other effects. These may include the initial radius of the BEC, the strength of repulsive interactions while the atoms are falling, or quantum depletion (section 3.3.3).

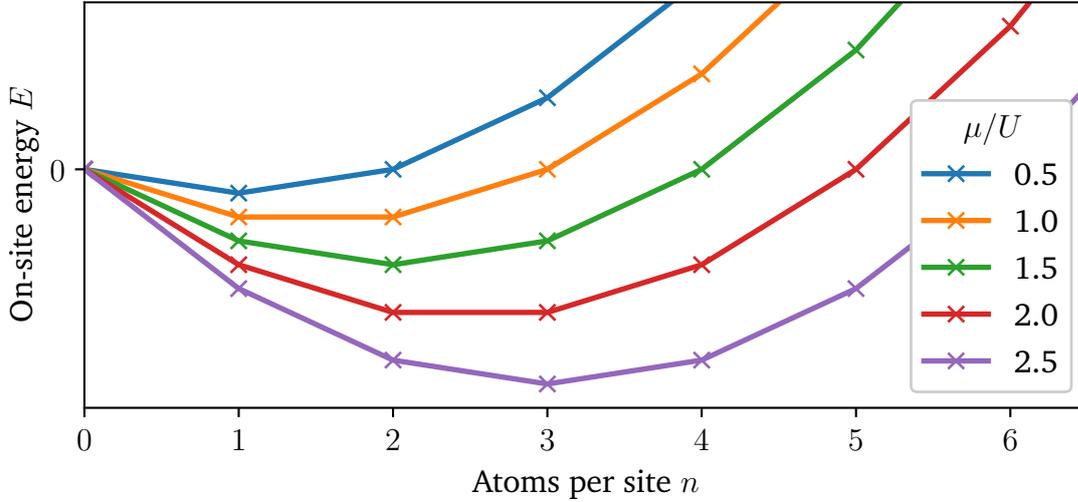


Figure 4.4: On-site energy of the Mott insulator as a function of occupation number, using equation 4.23, plotted for several values of the chemical potential μ (given in units of the interaction energy U). For the half-integer values of μ , E shows a minimum exactly at integer n , giving a clear equilibrium filling. However, integer μ results in two ns having the same E and there is no Mott insulating state.

spectively. The change in interaction energy is therefore

$$\begin{aligned}
 \Delta E &= \frac{1}{2}U(n+1)n + \frac{1}{2}U(n-1)(n-2) - 2 \times \frac{1}{2}Un(n-1) \\
 &= \frac{1}{2}U[n^2 + n + n^2 - 3n + 2 - 2n^2 + 2n] \\
 &= \frac{1}{2}U[2] \\
 &= U.
 \end{aligned} \tag{4.22}$$

If U is sufficiently large compared to J , the energy scale of hopping, this energy cost will become prohibitive. The state is very reminiscent of a traffic jam: the carriers could hop freely if there were a clear path, but in every direction they are blocked by others. Note that placing even a single additional particle into such a system, so that a single site has $n+1$ carriers while the others still have n , technically destroys the MI, as that additional carrier can hop without any energy penalty and carry a current. The removal of a single carrier has an identical effect, allowing particles to shuffle the resulting vacancy across the system. This strongly indicates that the existence of an MI state is dependent on the chemical potential μ .

The number of carriers on each site can be derived from a similar calculation. From the Bose-Hubbard Hamiltonian, the energy per site with n carriers is given by

$$E = -\mu n + \frac{1}{2}Un(n-1). \tag{4.23}$$

At equilibrium n will take the (integer) value that minimises this expression: see the plot in figure 4.4. Note that if μ is an integer multiple of U , two values of n will give equal results for E : in that case the MI state does not exist unless J is precisely zero, as atoms can hop freely between sites with those occupations.

4.2.2 Incoherence

The MI ground state may be written most simply in the limit of $J = 0$ (but $U > 0$). For N particles on M sites (N/M must be an integer), the state may be written as

$$|\Psi_{\text{MI}}\rangle \propto \prod_{j=1}^M (\hat{a}_j^\dagger)^{\frac{N}{M}} |0\rangle \quad (4.24)$$

where $|0\rangle$ is the "vacuum" state that describes an empty lattice site. Contrast this with the SF ground state of equation 4.21. For the SF we placed particles into the Bloch wave $|q = 0\rangle$, delocalised over the entire lattice, whereas for the MI we place particles into Wannier states localised to a single site. The complete absence of hopping means that there is no admixture of Wannier states between sites; to put it another way, the particles do not talk to their neighbours. This means that particles on adjacent sites have no phase coherence, and the correlation length ξ that was infinite in the SF is 0 in this state.

If we allow a non-zero J , but still $\ll U$, it will no longer be true that $\xi = 0$: this is not a general characteristic of the MI in the same way that infinite ξ is for the SF. Instead ξ will take a non-zero value that diverges to infinity at the transition to the SF. It is finite, rather than vanishing, ξ that indicates the MI phase [90]. This can be observed as faint peaks in TOF even deep in the insulating regime [123].

4.2.3 Incompressibility

Great insight into the MI state can be gained using the idea of *compressibility*, which we may define as [124]

$$\kappa := \frac{\partial n}{\partial \mu} \quad (4.25)$$

for total atomic density n and chemical potential μ . This corresponds to the *electronic compressibility* in conventional condensed matter physics, and its value depends on the energy required to add a single additional particle to the system. Since the MI will always require at least an energy U to add a particle, the MI is *incompressible*, with $\kappa = 0$. This breaks down only when μ is an integer multiple of U : as figure 4.4 shows, in that case there are two different occupations with the same interaction energy, allowing us to freely add particles to the less occupied sites. This results in a series of Dirac deltas for κ positioned at integer μ/U .

This form of κ can also be verified using equation 4.23, which allows us to plot the on-site occupation n as a function of μ and observe the gradient to read off the compressibility (see figure 4.5). Similar behaviour has been predicted [125] and observed [126] for ultracold atom MIs in inhomogeneous OLs, resulting in a characteristic "wedding cake" structure of plateaus in density separated by sharp jumps. (Note that non-zero J will smooth out the sharp jumps between plateaus: see [126] and section 7.1.2.)

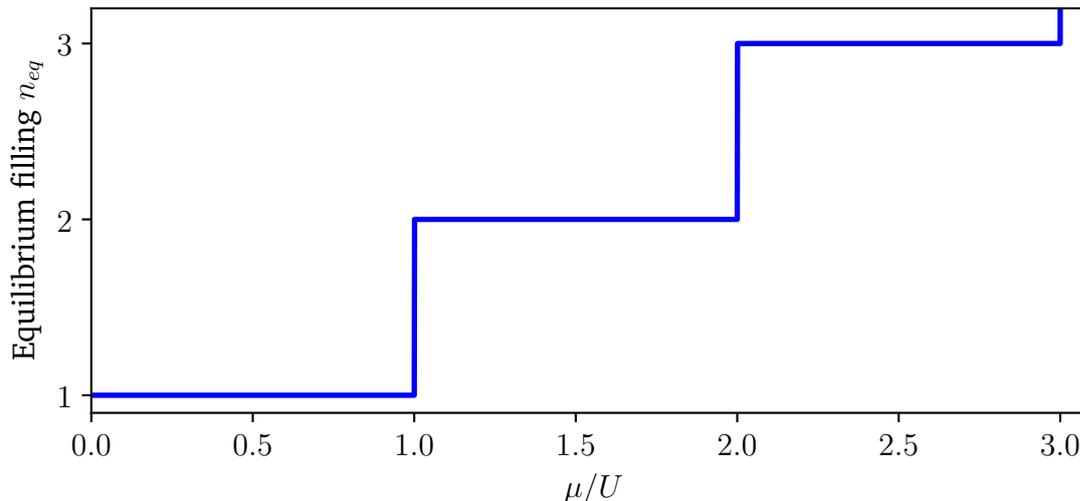


Figure 4.5: Occupation of each lattice site, n_{eq} , as a function of chemical potential μ (in units of the interaction U) in the limit $J = 0$. This was calculated by finding the integer value of n that minimises equation 4.23 at each value of μ . Note the discontinuous jumps whenever μ reaches an integer multiple of U , with plateaus in between. The gradient of this plot gives the compressibility κ , which is zero on the plateaus and spikes to infinity at each jump.

4.2.4 Imaging the Mott insulator

As noted when discussing the SF, the distribution of atoms visible in TOF is greatly affected by the correlation length ξ . Since ξ can take any non-infinite value in the MI, there is more variety in MI TOF images than in the SF case. Close to the transition the distinction is only that the sharp peaks of the structure factor begin to broaden, while the broad Wannier background becomes faintly visible. As U/J increases the peaks broaden until they are no longer visible, leaving a single broad smudge that indicates near-total incoherence: see figures in e.g. [89, 90].

This requires caution: the lack of any sharp features can also indicate that experimental problems have heated the cloud back above the BEC transition. To be sure the state is genuinely an MI it is necessary to lower U/J again and confirm that sharp peaks reappear, as was done in [89].

4.3 The Mott insulator-superfluid transition

The previous sections have discussed the two ground states of the Bose-Hubbard Hamiltonian. Which of these has lower energy depends on the ratio J/U , with the MI occurring below some critical value and the SF above it. We will now discuss the transition between these states, which is most easily treated in the limit of mean-field theory. I will largely follow the approach of [84] and [115], but see also e.g. [89, 127].

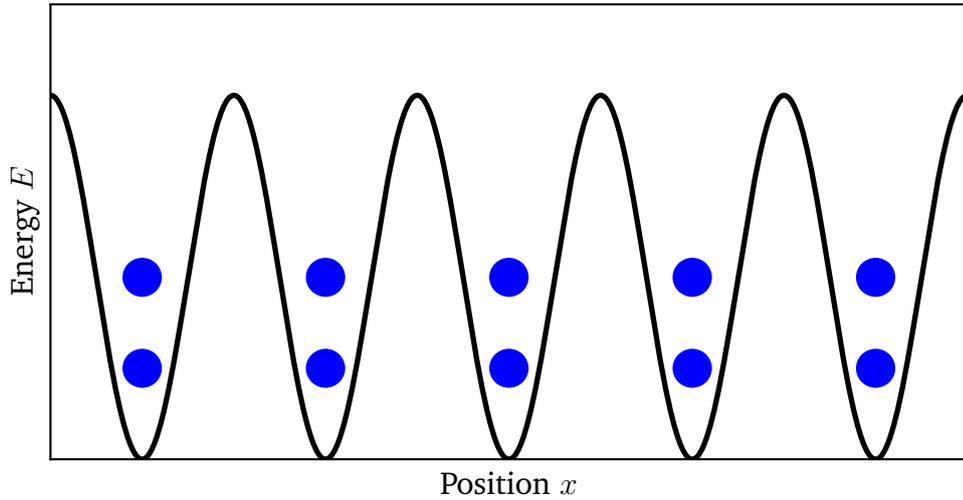


Figure 4.6: Illustration of the Mott insulating phase with two particles per site. In the absence of interactions particles would be able to hop from site to site, but the energy penalty of U for having three particles on a site is too large to be overcome. Note that this only works perfectly if the number of particles is exactly an integer multiple of the number of lattice sites: one extra particle would be able to ignore the fixed background and hop freely.

4.3.1 The mean-field approach

Since the Bose-Hubbard Hamiltonian is in general quite intractable, it is easiest to approach the MI-SF transition using *mean-field theory*. In the most general sense, this involves assuming that the wavefunction $\Psi(\{\mathbf{r}_i\})$ of a many-body interacting system, for some set of parameter vectors \mathbf{r}_i , can be written as a product state

$$\Psi(\{\mathbf{r}_i\}) = \prod_i \psi(\mathbf{r}_i) \quad (4.26)$$

with all particles in the same state $|\psi\rangle$. This greatly simplifies the role of interactions by allowing each particle to interact only with the *average* of all other particles. There are several different ways of saying this:

- For non-interacting bosons, the ground state will have all particles in the single-particle ground state. The mean-field approach keeps this picture fundamentally the same, but allows the state they all occupy to differ from the single-particle ground state.
- Although the mean-field picture allows fluctuations, it does not allow them to couple to each other, only to the mean wavefunction. So while it allows interactions between particles, it forbids interactions between *quasiparticles*.
- In the Bose-Hubbard model specifically, nearest-neighbour hopping means that particles can only be influenced by adjacent lattice sites. The mean-field picture is therefore equivalent to averaging over the nearest neighbours of every lattice site, and assuming that this is the same as averaging over all lattice sites.

The last of these points has implications for the validity of the mean-field approach depending on the coordination number³ z of the lattice: the larger the value of z , the better the mean-field approximation becomes, because the effect of fluctuations on the nearest-neighbour sites will become less significant. Since we usually have $z \propto d$ for d dimensions, this means that mean-field methods are more accurate as the dimensionality increases, becoming exact in the limit of infinite dimensions.

4.3.2 Mean-field treatment of the Bose-Hubbard Hamiltonian

As the Bose-Hubbard Hamiltonian is a tight-binding Hamiltonian, it's natural to apply the general mean-field approach of equation 4.26 by introducing a local Hamiltonian \hat{h}_i for the i th lattice site,

$$\hat{H}_{MF} = \sum_i \hat{h}_i. \quad (4.27)$$

Referring to the Bose-Hubbard Hamiltonian of equation 2.21, we might guess that \hat{h}_i will be something like

$$\hat{h}_i = -Jz\hat{a}_i^\dagger\hat{a}_j + \frac{1}{2}U\hat{n}_i(\hat{n}_i - 1) - \mu\hat{n}_i \quad (4.28)$$

where \hat{a}_j refers to one of site i 's nearest neighbours, and a factor of the coordination number z has been added to replace the sum over $\langle i, j \rangle$. We can simplify this further by separating \hat{a}_i into a mean-field term and a fluctuation:

$$\hat{a}_i =: \langle \hat{a}_i \rangle + \hat{\delta}_i. \quad (4.29)$$

Now let us assume that $\langle \hat{a}_i \rangle$ is the same for all sites i ⁴, taking some constant value that we call ψ (anticipating that this will become the order parameter of the transition). We can then insert this into the hopping term of \hat{h}_i :

$$\begin{aligned} \hat{a}_i &= \psi + \hat{\delta}_i \\ \implies \hat{a}_i^\dagger\hat{a}_j &= (\psi^* + \hat{\delta}_i^\dagger)(\psi + \hat{\delta}_j) \\ &= |\psi|^2 + \psi^*(\hat{a}_j - \psi) + (\hat{a}_i^\dagger - \psi^*)\psi + \underbrace{\hat{\delta}_i^\dagger\hat{\delta}_j}_{\text{Negligible}} \\ &= \psi^*\hat{a}_j + \hat{a}_i^\dagger\psi - |\psi|^2 \end{aligned} \quad (4.30)$$

where we have assumed sufficiently small fluctuations that $\hat{\delta}_i^\dagger\hat{\delta}_j$ is negligible. We now arrive at the form for \hat{h}_i

$$\hat{h}_i = -Jz(\psi^*\hat{a}_i + \psi\hat{a}_i^\dagger - |\psi|^2) + \frac{1}{2}U\hat{n}_i(\hat{n}_i - 1) - \mu\hat{n}_i. \quad (4.31)$$

At this point we must turn to numerical calculations, using ψ as a variational parameter to minimise $\langle \hat{H}_{MF} \rangle$.

³Number of nearest neighbours of each site. In a 1D lattice $z = 2$, while in a 2D square lattice $z = 4$.

⁴This doesn't hold for an SF with $q \neq 0$, or for a non-uniform system. This approach can be extended to the former case with an extra factor of e^{iqr} , and to the latter with the local density approximation (section 4.1.1).

There is, however, one more analytical point to be made. Note that in the limit where $J = 0$ this precisely recovers the expression 4.23 we wrote down from inspection of the Bose-Hubbard Hamiltonian. However, we now gain an additional insight: equation 4.31 also becomes equation 4.23 if $J \neq 0$ but $\psi = 0$. This confirms our expectation that ψ is the order parameter for the MI-SF transition: $\psi = 0$ implies the MI state, while non-zero ψ results in the SF as long as $J \neq 0$ (and if $J = 0$ we expect the MI to be the only ground state anyway). We therefore deduce that the MI is the disordered state in this transition, with the long-range phase coherence of the SF representing increased order.

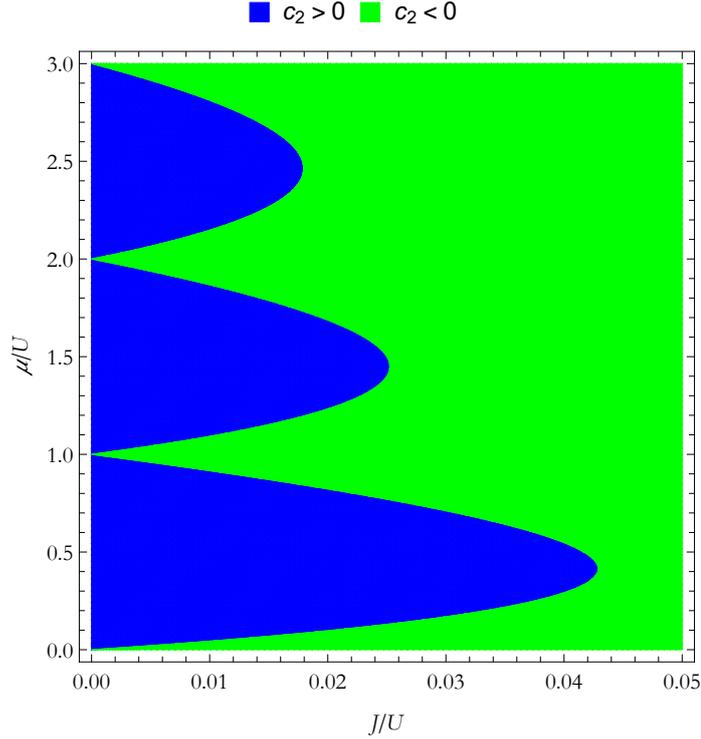


Figure 4.7: A plot showing the sign of the Ginzburg-Landau coefficient c_2 (equation 4.33) with $z = 4$, the coordination number of a 2D square lattice. A change in sign of c_2 indicates a continuous quantum phase transition between the Mott insulator (blue) and superfluid (green). The Mott insulator region is split into several lobes which touch only at $J = 0$, and which can be indexed by the integer value of $[\mu/U]$.

Ginzburg-Landau theory

Now let J have some finite value and consider a small ψ close to the transition. We can use the Ginzburg-Landau theory of continuous PTs introduced in section 1.3.2 to guess that the free energy has the form

$$\langle \hat{H} \rangle = E_0 + c_2 |\psi|^2 + c_4 |\psi|^4 + \dots \quad (4.32)$$

Recall that the sign of c_2 determines the position of the global minimum, with $c_2 = 0$ marking the phase boundary. It can be calculated [84] that

$$c_2 = f\left(\frac{\mu}{U}\right) \left[1 - zJf\left(\frac{\mu}{U}\right) \right] \quad (4.33)$$

where

$$f\left(\frac{\mu}{U}\right) := \frac{\lceil \mu/U \rceil + 1}{U\lceil \mu/U \rceil - \mu} + \frac{\lceil \mu/U \rceil}{\mu - U(\lceil \mu/U \rceil - 1)}. \quad (4.34)$$

Figure 4.7 shows the regions of phase space where $c_2 < 0$ and $c_2 > 0$. The lobes where $c_2 > 0$ can be indexed by the integer value of $\lceil \mu/U \rceil$, which is exactly the MI occupation of figure 4.5. The mean-field approach therefore recovers the exact phenomenology of our earlier intuitive picture.

4.3.3 Imaging the transition

The phase transition between the SF and MI states is readily observed in time of flight by combining the observations of sections 4.1.4 and 4.2.4. As U/J is increased across the phase boundary, the sharp peaks of the SF begin to broaden and an incoherent background appears. The peaks ultimately disappear as coherence is completely lost. This method does not allow precise identification of the phase boundary, but certainly verifies that the transition is occurring⁵.

More quantitatively, such images can be used to measure ξ as a function of U/J . There is no simple way of extracting ξ from TOF images (though assuming the peak widths to be $1/\xi$ is a reasonable approximation), but simulated TOF images can be calculated from equation 3.26 by assuming a value of ξ and then compared to the experimental results. The location of the asymptote, where ξ diverges to infinity, is the phase boundary.

⁵As noted in section 4.2.4, verification requires crossing back from the MI to the SF to be sure that the loss of coherence is reversible and not due to heating.

5

Floquet physics

5.1 Driven systems

In this section I will briefly review the physics of periodically driven quantum systems. By periodic driving I mean that the Hamiltonian is time-dependent with the property $\hat{H}(t + T) \equiv \hat{H}(t)$ for some period T . This closely resembles the starting point for Bloch's theorem mentioned in section 2.1.2, and results in essentially the same physics. We saw that a system with discrete rather than continuous translational symmetry conserves momentum only modulo $\frac{2\pi}{a}$ for some repeat distance a . In direct analogy we will discover that periodically driven systems conserve energy only modulo $\frac{2\pi\hbar}{T}$.

We can naturally divide the physics of periodically driven systems into two limits: the *adiabatic* limit of slow driving, which I will address briefly, and the *Floquet* regime of fast driving, which will be the subject of the rest of this chapter. These two limits correspond to periods T respectively much slower and much faster than some *adiabaticity time scale* T_a , which I will now derive.

The time scale of adiabaticity

Consider a system with static Hamiltonian \hat{H} that has eigenbasis $|\phi_i\rangle$. A general wavefunction of this system may be expressed as

$$|\Psi\rangle = \sum_i c_i |\phi_i\rangle \tag{5.1}$$

for some coefficients c_i . The simplest way to add time dependence to $|\Psi\rangle$ would be to modify this equation so that it reads

$$|\Psi(t)\rangle = \sum_i c_i(t) |\phi_i(t)\rangle, \quad (5.2)$$

where $|\phi_i(t)\rangle$ is the instantaneous eigenbasis of $\hat{H}(t)$ and the coefficients $c_i(t)$ are found by integrating the phase from the time-dependent eigenenergies E_i ,

$$c_i(t) = \exp\left(-\frac{i}{\hbar} \int_0^t E_i(t') dt'\right) c_i(0). \quad (5.3)$$

However, this is in general too naïve, as there may be terms depending on $\partial_t \hat{H}$. These will affect the time evolution if $\partial_t \hat{H}$ causes significant couplings between eigenstates. From perturbation theory we expect the coupling term between two states $|\phi_i\rangle$ and $|\phi_j\rangle$ to be proportional to

$$\frac{\langle \phi_i | \partial_t \hat{H} | \phi_j \rangle}{E_i - E_j}, \quad (5.4)$$

where all variables are time-dependent. The natural scale for these terms to be significant is the energy gap between states, $E_i - E_j$. Therefore, we expect them to be negligible if we meet the adiabaticity condition [128]

$$\hbar |\langle \phi_i | \partial_t \hat{H} | \phi_j \rangle| \ll |E_i - E_j|^2 \quad (5.5)$$

for all choices of i and j . This will typically be limited by whichever two states have the smallest energy gap. (Recall the Landau-Zener formula of section 1.4.2, which is closely related to this result.)

Adiabatic driving

In the case of periodic driving that meets the adiabaticity condition, no transitions between eigenstates occur as the Hamiltonian completes its cyclic excursion from $\hat{H}(t)$ back to $\hat{H}(t+T) = \hat{H}(t)$. Therefore we would expect that there is no observable change in the system once we return to the initial Hamiltonian, so that for any initial $|\Psi(0)\rangle$ and integer n ,

$$||\Psi(nT)\rangle|^2 = ||\Psi(0)\rangle|^2. \quad (5.6)$$

This allows only for the wavefunction to accumulate some phase ϕ_B , known as the *Berry phase* or *geometric phase*, with each periodic cycle [129, 130]. Calculating this phase is beyond the scope of our interest here – we simply observe that it is the only novel physical effect arising from periodic driving if the driving is adiabatic.

5.2 Floquet theory

We will now consider the physics of systems subject to periodic driving that is much faster than the adiabaticity condition in equation 5.5. There are essentially two possible effects:

- The system will behave as though it is governed by the time-averaged Hamiltonian, because the wavefunction can't respond within one period of the driving.
- The system will transition between different energy states.

If the frequency of driving is slow enough and the amplitude weak enough, it may be possible to ignore the second of these effects, and in this case the effects of Floquet driving can be derived quite straightforwardly. To illustrate this I will begin in section 5.2.1 by considering a single-band model of a driven 1D lattice, which results in renormalisation of the hopping J but nothing more intricate. The most interesting applications of Floquet theory, however, require taking both effects into account, and I will develop this more involved theory in section 5.2.2.

5.2.1 Hopping renormalisation in a single band

Equation 2.12 gives the form of the ground band of a 1D lattice. Suppose the system is driven by a homogeneous force $F(t)$, such that $F(t + T) \equiv F(t)$ for some period T . This will not affect E_c , a , or t in 2.12, but equation 2.16 tells us that

$$\partial_t q(t) = \frac{1}{\hbar} F(t). \quad (5.7)$$

To make things simple let us assume that the force is monochromatic, $F(t) = F_0 \cos(\omega t)$. We can then integrate to find

$$q(t) = q_0 + \frac{F_0}{\hbar\omega} \sin(\omega t) \quad (5.8)$$

for some initial q_0 . We can substitute this back into equation 2.12 to find the shape of the ground band as a function of time.

$$\begin{aligned} E_0(q_0, t) &= -2J \cos \left[\left(q_0 + \frac{F_0}{\hbar\omega} \sin(\omega t) \right) a \right] \\ &\equiv -2J \cos(q_0 a) \cos(K_0 \sin(\omega t)) \\ &\quad + 2J \sin(q_0 a) \sin(K_0 \sin(\omega t)), \end{aligned} \quad (5.9)$$

where $K_0 := F_0 a / \hbar\omega$ is a dimensionless measure of the driving and we have used a trigonometric identity for $\cos(\theta + \phi)$. Our argument is that provided ω is large enough (while still being slow enough to avoid interband transitions), we must be able to use the time-averaged form

$$\begin{aligned} \bar{E}_0(q_0) &\equiv \frac{1}{T} \int_0^T E_0(q_0, t) dt \\ &= -2J \cos(q_0 a) \underbrace{\frac{1}{T} \int_0^T \cos(K_0 \sin(\omega t)) dt}_{\equiv \mathcal{J}_0(K_0)} \\ &\quad + 2J \sin(q_0 a) \underbrace{\frac{1}{T} \int_0^T \sin(K_0 \sin(\omega t)) dt}_{\equiv 0} \\ &\equiv -2J \mathcal{J}_0(K_0) \cos(q_0 a), \end{aligned} \quad (5.10)$$

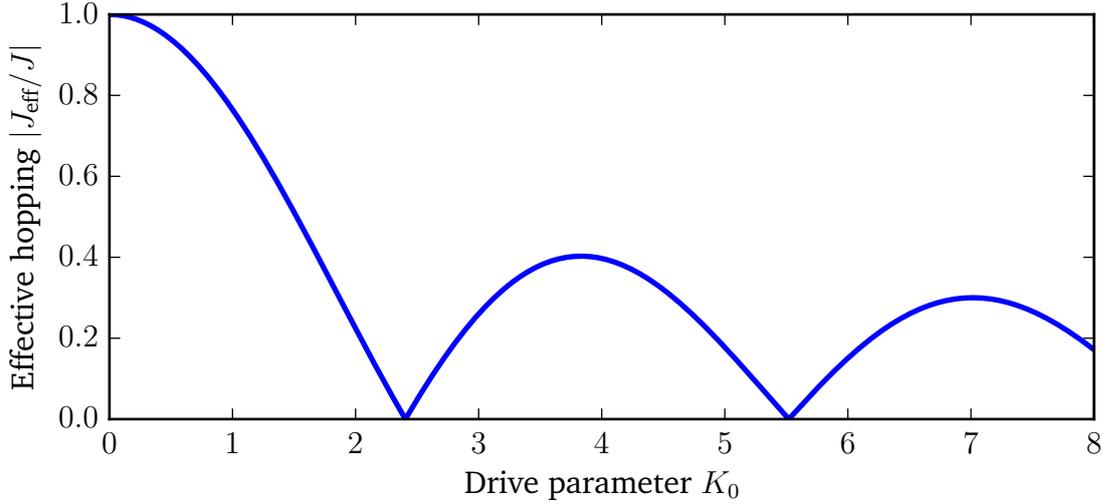


Figure 5.1: A plot showing the absolute value of the effective hopping J_{eff} in the ground band of a driven 1D lattice as a function of the dimensionless drive parameter K_0 (equation 5.11). At various values of K_0 the hopping vanishes, an effect called dynamic localisation. Note that what is plotted here is the *absolute value* of the hopping, as the Bessel function changes sign at each zero crossing. This was done to make this figure resemble the experimental result from [64].

where $\mathcal{J}_0(x)$ is a zeroth-order Bessel function of the first kind. Comparing this result to equation 2.12 we see that it has an identical form, but with J replaced by an effective hopping term

$$J_{\text{eff}} := \mathcal{J}_0(K_0)J. \quad (5.11)$$

J_{eff} is plotted in figure 5.1. The most striking feature to note is that the Bessel function vanishes at various values of K_0 , meaning that for these drive parameters there is no hopping in the lattice, an effect called *dynamic localisation* [64]. The corresponding shape of the ground band is

$$E_0(q) = E_c - 2J_{\text{eff}}\cos(qa), \quad (5.12)$$

which is plotted in figure 5.2 for a few representative values of K_0 .

Both dynamic localisation and the inverted band shown in figure 5.2 have interesting experimental consequences. Dynamic localisation can be used to drive the Mott insulator-superfluid transition (chapter 4) by changing the ratio J/U in the Bose-Hubbard model (section 2.1.5) [63, 131], an approach equivalent to increasing the lattice depth or enhancing U with a Feshbach resonance. An inverted ground band, meanwhile, results in the lowest-energy Bloch wave being $q = \pm\pi/a$ rather than $q = 0$. Bose-Einstein condensation (section 3.2) in this system will therefore result in atoms gathering at the edges of the first Brillouin zone rather than the centre, a state known as a π -superfluid.

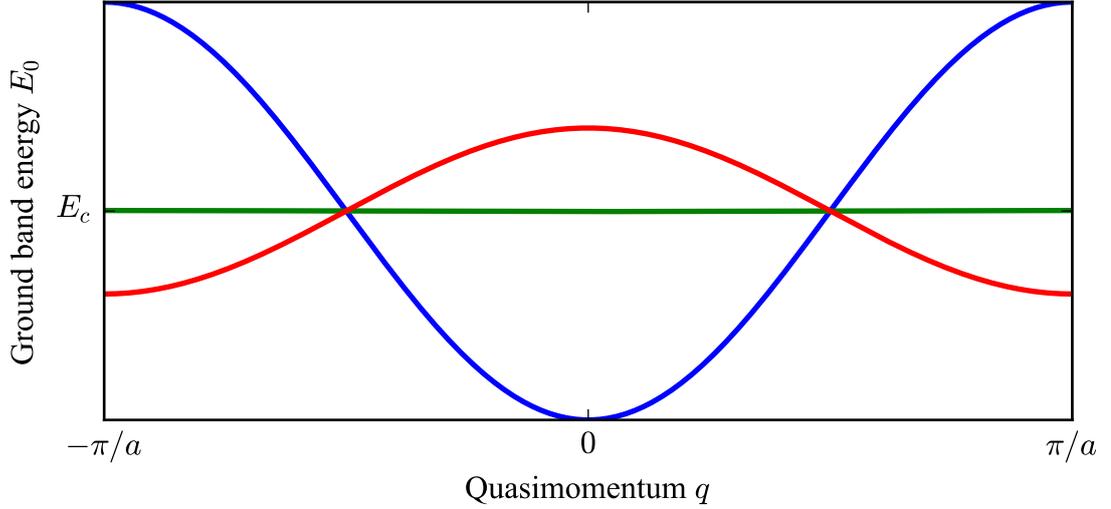


Figure 5.2: The ground band (equation 5.12) of a driven 1D lattice for three different values of the drive parameter: $K_0 = 0$ (blue) is identical to the undriven lattice of equation 2.12, $K_0 = 2.4$ (green) is close to the first zero of the Bessel function in figure 5.1 and results in a flat band with almost no hopping, and $K_0 = 4$ (red) shows band inversion as the effective hopping becomes negative.

5.2.2 Floquet states and quasienergies

I will now present a more complete form of Floquet theory that takes into account multiple bands. I will largely follow the treatment of Holthaus in [67], but see also [62, 132]. We begin from the time-dependent Schrödinger equation,

$$i\hbar\partial_t |\Psi(t)\rangle = \hat{H}(t) |\Psi(t)\rangle. \quad (5.13)$$

To find a general solution to this equation we must introduce the *propagator* $\hat{U}(t_1, t_2)$, which has the property

$$|\Psi(t)\rangle = \hat{U}(t, 0) |\Psi(0)\rangle, \quad (5.14)$$

which by substitution into 5.13 must itself satisfy

$$i\hbar\partial_t \hat{U}(t, 0) = \hat{H}(t) \hat{U}(t, 0). \quad (5.15)$$

It must also have the property that $\hat{U}(0, 0) = \hat{I}$, where \hat{I} is the identity operator, and furthermore we require $\hat{U}(t_1 + t_2, 0) = \hat{U}(t_1 + t_2, t_1) \hat{U}(t_1, 0)$ based on physical intuition. From these properties and the periodicity of \hat{H} it can be proven (appendix B) that

$$\hat{U}(t + T, 0) = \hat{U}(t, 0) \hat{U}(T, 0), \quad (5.16)$$

implying that we can fully understand the time evolution of the system if we know $\hat{U}(t, 0)$ only over the range $0 \leq t \leq T$. This suggests that $\hat{U}(T, 0)$, known as the *monodromy operator*, must be important. To explore its structure, we can start by expressing $\hat{U}(t, 0)$ in the form

$$\hat{U}(t, 0) =: \hat{P}(t) e^{-i\hat{G}t/\hbar} \quad (5.17)$$

where \hat{P} is time-periodic and \hat{G} is a Hermitian operator satisfying

$$e^{-i\hat{G}T/\hbar} := \hat{U}(T, 0). \quad (5.18)$$

If ϵ_n and $|n\rangle$ are the eigenvalues and eigenstates of \hat{G} , we can then write the monodromy operator as

$$U(\hat{T}, 0) = \sum_n |n\rangle e^{-i\epsilon_n T/\hbar} \langle n|, \quad (5.19)$$

where the factors $e^{-i\epsilon_n T/\hbar}$ are known as *Floquet multipliers*. To solve the general dynamics we can expand $|\Psi(0)\rangle$ and $|\Psi(t)\rangle$ in this basis, since the eigenstates of a Hermitian operator are complete and can be chosen to be orthonormal. First, let

$$|\Psi(0)\rangle = \sum_n c_n |n\rangle, \quad (5.20)$$

defining the probability amplitudes c_n . We can then substitute 5.17 and 5.20 into 5.14 to obtain

$$\begin{aligned} |\Psi(t)\rangle &= \hat{U}(t, 0) |\Psi(0)\rangle \\ &= \sum_n c_n \hat{P}(t) e^{-i\hat{G}t/\hbar} |n\rangle \\ &= \sum_n c_n \hat{P}(t) e^{-i\epsilon_n t/\hbar} |n\rangle \\ &=: \sum_n c_n e^{-i\epsilon_n t/\hbar} |u_n(t)\rangle \\ &=: \sum_n c_n |\psi_n(t)\rangle, \end{aligned} \quad (5.21)$$

where we have defined the *Floquet functions* $|u_n(t)\rangle := \hat{P}(t) |n\rangle$ and the *Floquet states* $|\psi_n(t)\rangle := e^{-i\epsilon_n t/\hbar} |u_n(t)\rangle$, which are the solutions to 5.13. Note that the probability amplitudes c_n are *constant*: all the time-dependence is carried by the Floquet states. This provides the central strategy for understanding the dynamics of driven systems: if we can find $|n\rangle$ and ϵ_n from the Hamiltonian, we need only make the expansion 5.20 to find the coefficients c_n .

There is an important physical insight here. The quantities ϵ_n are time-independent and play the same role as the energies of stationary states for a static Hamiltonian. Therefore we give them the name *quasienergies*, a direct analogue to quasimomentum in spatially periodic systems.

Quasienergy Brillouin zones

Consider again the Floquet multipliers $e^{-i\epsilon_n T/\hbar}$, which are the only way the quasienergies ϵ_n enter the dynamics of $|\Psi(t)\rangle$. While these multipliers are uniquely defined, their logarithm, used to extract the quasienergies, is not. In particular, we can shift any quasienergy by any multiple of $2\pi\hbar/T$, which we will define in terms of an angular frequency as $\hbar\omega$, without changing the value of the multiplier. We will therefore introduce the notation

$$\epsilon_{(n,m)} := \epsilon_n + m\hbar\omega, \quad (5.22)$$

where m is any positive or negative integer, or zero. In this way we see that each quasienergy ϵ_n is a representative of an *equivalence class* rather than a unique choice. This allows us to choose a representative from each class that lies in the range $-\hbar\omega/2 < \epsilon_n \leq \hbar\omega/2$, which we define as the first quasienergy Brillouin zone in analogy to Bloch waves.

Extending Hilbert space: the Kamiltonian

Further insight into this energy structure can be obtained by extending our formalism to remove the artificial distinction between quasienergy representatives. The centre-piece of this extension is the quasienergy operator, or *Kamiltonian*,

$$\hat{K} := \hat{H}(t) - i\hbar\partial_t. \quad (5.23)$$

See e.g. [133]. There are several important comments to make here.

- The Floquet functions $|u_n(t)\rangle$ satisfy

$$\hat{K} |u_n(t)\rangle = \epsilon_n |u_n(t)\rangle, \quad (5.24)$$

which can be seen by substituting the Floquet states $|\Psi_n(t)\rangle$ into the time-dependent Schrödinger equation. This justifies regarding \hat{K} as the quasienergy operator.

- In defining \hat{K} like this we aim to change the way time is represented in our problem. Typically in non-relativistic quantum mechanics time is promoted above space as a parameter of evolution rather than a coordinate. By introducing ∂_t as an operator we undo this, treating time as an additional dimension. (This is unrelated to the concept of space-time in relativity.)
- \hat{K} operates on a different Hilbert space from \hat{H} , which has been extended to include all T -periodic functions [134]. If \hat{H} acts on \mathcal{H} , the space that \hat{K} acts on is designated $L_2[0, T] \otimes \mathcal{H}$.
- The new derivative ∂_t appears in the Hamiltonian only to linear order. Contrast this with ∂_x^2 , which appears in the one-dimensional kinetic energy operator. A linear derivative in the Hamiltonian allows a spectrum that is bounded neither from above nor from below, which we have already seen in equation 5.22. This is a central reason why the physics of Floquet systems is different from anything that can be achieved in static condensed matter systems.

To illustrate the correspondence between \hat{K} and \hat{H} , consider that

$$\hat{K} |u_n(t)\rangle e^{im\omega t} = (\epsilon_n + m\hbar\omega) |u_n(t)\rangle e^{im\omega t}, \quad (5.25)$$

indicating that each representative of a quasienergy equivalence class 5.22 is distinct in $L_2[0, T] \otimes \mathcal{H}$. However, this distinctness drops out when we use the state $|u_n(t)\rangle e^{im\omega t}$ to find a Floquet state in \mathcal{H} :

$$\begin{aligned} |\Psi_n(t)\rangle &= e^{-i(\epsilon_n + m\hbar\omega)t/\hbar} (|u_n(t)\rangle e^{im\omega t}) \\ &= e^{-i\epsilon_n t/\hbar} |u_n(t)\rangle, \end{aligned} \quad (5.26)$$

which is our original definition of the Floquet state corresponding to $|u_n(t)\rangle$.

Fully exploring the application of this operator is beyond the scope of this thesis. However, let us pursue it a little further. Consider a state in $L_2[0, T] \otimes \mathcal{H}$. We can't consider this state as evolving with time t , because t is a coordinate in this space. Therefore we introduce a new variable τ , which will fill the role of evolution parameter that t has vacated. The state is then written as $|\Psi(\tau, t)\rangle\rangle$, with the double bracket indicating it does not belong in \mathcal{H} . We expect this state to satisfy

$$i\hbar\partial_\tau |\Psi(\tau, t)\rangle\rangle = \hat{K} |\Psi(\tau, t)\rangle\rangle, \quad (5.27)$$

and we recover the corresponding physical state by setting τ and t equal:

$$|\Psi(t)\rangle = |\Psi(\tau, t)\rangle\rangle|_{\tau=t}. \quad (5.28)$$

These conditions can be used to recover the conventional Schrödinger equation for $|\Psi(t)\rangle$. We therefore reach a picture where the Kamiltonian governs the true dynamics, while we observe only the projections of the full wavefunctions $|\Psi(\tau, t)\rangle\rangle$ onto a lower-dimensional space.

We can also obtain a tangible mathematical insight from this excursion into a higher realm. Note that the Kamiltonian is not invariant under the conventional parity transformation $x \rightarrow -x$. It instead requires an accompanying transformation in time,

$$\begin{aligned} x &\rightarrow -x \\ t &\rightarrow t + \frac{T}{2}. \end{aligned} \quad (5.29)$$

This means that a state with quasienergy $\epsilon_{(n,m)}$, which has parity $(-1)^{n+1}$ in \mathcal{H} , will instead have parity $(-1)^{n+m+1}$ in $L_2[0, T] \otimes \mathcal{H}$. Therefore, provided the driving is strong enough to be comparable to other terms in the Hamiltonian, the selection rules for transitions will change to satisfy this extended parity and potentially violate the conventional definition.

5.3 Driving an optical lattice

5.3.1 The Hamiltonian

I will now apply the theoretical framework developed in section 5.2.2 to a concrete example, namely a longitudinally driven one-dimensional optical lattice. I will continue to follow the example of [67]. Begin by considering the Hamiltonian

$$\hat{H}_{lab}(x, t) = \frac{p^2}{2m} + \frac{1}{2}V_0\cos(2k_{\text{lat}}[x - \Delta L\cos(\omega t)]) \quad (5.30)$$

for an atom in a shaken optical lattice. Here p and m are the momentum and mass of the atom, V_0 and k_{lat} are the lattice depth and wavenumber, and ΔL and ω are the amplitude and frequency of shaking. There are several ways to accomplish such

shaking experimentally – a few will be discussed in section 6.2.2. A series of unitary transformations (appendix C) may be used to put this Hamiltonian in a more tractable form:

$$\hat{H}(x, t) = \frac{1}{2m}(p + m\Delta L\omega\sin(\omega t))^2 + \frac{1}{2}V_0\cos(2k_{\text{lat}}x) - \frac{1}{4}m(\Delta L\omega)^2. \quad (5.31)$$

The first term now resembles the typical form for a particle in an electromagnetic field, $(p - eA)^2$ for some charge e and vector potential A , or more generally a gauge-invariant coupling between a particle and a field, which is the driving force in this case. The second term is the unmodified spatial potential of the undriven lattice. The third term captures the motion of the particle during each cycle of the driving (*micromotion*), and is called the *ponderomotive energy*: it is equal to the average kinetic energy of a particle undergoing classical simple harmonic motion.

5.3.2 Space-time Bloch waves

The eigenstates of this Hamiltonian can be easily guessed, and have the form

$$\Psi_{n,k}(x, t) = e^{ikx - i\epsilon_n(k)t/\hbar} u_{n,k}(x, t) \quad (5.32)$$

where the functions $u_{n,k}(x, t)$ are periodic in both space and time and are known as the *Bloch-Floquet functions*. The wavefunctions themselves may be described as space-time Bloch waves, extending the familiar Bloch waves of undriven periodic systems. This extension corresponds also to semiclassics (see equation 2.14): we may describe a particle in the n th quasienergy band by the wavefunction

$$\Psi_n(x, t) = \sqrt{\frac{a}{2\pi}} \int dk g_n(k) e^{ikx - i\epsilon_n(k)t/\hbar} u_{n,k}(x, t), \quad (5.33)$$

where $g_n(k)$ is a continuous form of the probability amplitudes c_n from 5.21. Results for group velocity and the effects of forces may similarly be directly transferred to the driven system.

To focus on the temporal component of the space-time Bloch waves, let us define

$$\phi_{n,k} = e^{-ikx} \Psi_{n,k}(x, t). \quad (5.34)$$

We can then define new coordinates $z := k_{\text{lat}}x$ and $\tau := \omega t$, and make use of the recoil energy E_r , to write the rescaled Schrödinger equation in the form

$$iE\partial_\tau\phi_{n,k}(z, \tau) = \left[-\partial_z^2 + \kappa^2 + \frac{1}{2}\beta^2 + \frac{1}{2}V\cos(2z) - 2i(\kappa + \beta\sin(\tau))\partial_z \right] \phi_{n,k}(z, \tau), \quad (5.35)$$

where we have defined the dimensionless constants $E := \frac{\hbar\omega}{E_r}$, $\kappa := \frac{k}{k_{\text{lat}}}$, $V := \frac{V_0}{E_r}$, and $\beta := \frac{m\Delta L\omega}{\hbar k_{\text{lat}}}$. The solutions to this equation are the quasienergy bands of the lattice. To solve it we take the undriven Bloch waves as the initial condition for $\phi_{n,k}(x, t = 0)$. We then propagate each state by one period T . The resultant states form the columns of the monodromy operator $U(T, 0)$, which has as its eigenvalues the quasienergies.

This figure has been redacted for copyright reasons. The copyright owner is IOP Publishing Ltd. See print version, or figure 6 in [67].

Figure 5.3: Reproduced from [67]. The lowest three quasienergy bands of a 1D OL with $V = 4$ and $E = 0.5$, considering only the point $\kappa = 0$ of each (see equation 5.35). The horizontal axis is β from that equation. The lowest band is marked with an arrow at $\beta = 0$. See how, as β increases, that band initially continues unperturbed, then passes rapidly through six avoided crossings and eventually becomes difficult to follow. The number of avoided crossings only increases as more quasienergy bands are considered.

5.4 Using Floquet theory

With equation 5.35 our derivation is finished, and we can now attempt to understand what this result means intuitively. We may also wonder whether this result can be connected back to the single-band calculation of section 5.2.1 in appropriate limits, and finally we may ask how control of the shaking parameters can be used to design band structures of our choosing. I will address these questions in this section. In doing so I will reproduce several figures from [67] to illustrate the discussion.

5.4.1 Floquet heating

Equation 5.22 describes an infinite set of quasienergy bands corresponding to every Floquet multiplier, allowing us to fold the entire spectrum of the shaken system into the first quasienergy Brillouin zone $-\hbar\omega/2 < \epsilon_n \leq \hbar\omega/2$. Since this implies folding a very large or infinite number of curves into a finite region, it is no surprise that the first quasienergy Brillouin zone tends to be rather crowded, with a huge number of avoided level crossings. Figure 5.3 is a general example of this with just three quasienergy bands included.

This suggests that the quasienergy picture will become useless if we try to consider all bands, with avoided crossings at every point of the Brillouin zone. The saving grace is that the crossings vary wildly in strength depending on how many units of $\hbar\omega$ separate the bands in the unfolded picture, or equivalently how many phonons must be absorbed from the lattice to couple the two states. A single-phonon process will

likely always result in a prominent avoided crossing, but a twenty-phonon process would likely be invisible on the scale of figure 5.3. This allows an approach called *coarse graining* [67], where we ignore all avoided crossings weaker than a certain energy scale and recover a manageable picture.

At sufficiently high β coarse-graining typically breaks down, however, and a large number of crossings will appear that are too large to ignore. If we drive in this regime the system will rapidly absorb energy from the shaking process and the atomic cloud will be heated severely, typically limiting or preventing experimental use. This is called *Floquet heating*, and can be prevented for a non-interacting system by staying well within the coarse-grained regime.

This figure has been redacted for copyright reasons. The copyright owner is IOP Publishing Ltd. See print version, or figures 7 and 9 in [67] for a) and b) respectively.

Figure 5.4: Reproduced from [67]. Illustration of the effect of increasing lattice depth on the structure of the quasienergy bands (without interactions). Both figures are solutions to equation 5.35 under the same conditions except for differing values of V as noted, and each of them shows the energy of eleven points in the lowest quasienergy band, $\kappa = 0, 0.1, 0.2, \dots, 1.0$, with many other quasienergy bands fainter in the background. First consider a). The eleven darker lines begin spread out across a wide range of E , representing the curvature of the ground band. They collapse down to a single point at $\beta \approx 0.7$, which is dynamic localisation (section 5.2.1). However, before the band can become flat a second time a profusion of avoided crossings destroys the structure of the lowest band and it becomes invisible. This is the end of the coarse graining regime and the beginning of uncontrolled Floquet heating. Contrast this with b). In this case the lowest band remains intact right up to $\beta = 2$ and beyond, with two collapses visible, and Floquet heating does not become important within the range of the calculation. This illustrates the central role of lattice depth in determining the importance of Floquet heating in a non-interacting system.

Interacting systems

Adding interactions generically makes Floquet heating much more important. This is because even if β is far from a value that readily couples two bands, an interacting system is able to support collective excitations that divide the energy of a single phonon across many particles [135]. Given sufficient time, the possibility of both multiphonon

processes and collective excitations will allow any drive to heat any interacting system to infinite temperature. Since interesting experimental systems tend to be interacting, the best we can do is often simply to choose the β that gives as long as possible for the experiment to take place before this heating becomes too severe. One possible solution to this is MBL; see section 7.2.2.

5.4.2 Single-band Floquet physics

To recover the single-band results of section 5.2.1 in this more general picture we must increase the lattice depth V_0 and keep ω away from any resonances of the undriven system (i.e. band gaps). This allows the coarse-grained regime to continue to much higher values of β , as illustrated in figure 5.4.

This allows us to separate the effects of Floquet driving in non-interacting systems into two broad categories. If we stay in the coarse-grained regime, the time-averaging analysis of section 5.2.1 is sufficient and we can ignore the more in-depth analysis in the rest of this chapter. This in itself allows for interesting experimental possibilities [63, 64]. If, however, we are driving near interband transitions, those averaging effects will be superposed with countless transitions at the avoided crossings and constant, uncontrolled heating. Adding interactions results in both effects always being present, even if we try to stay within the single-band regime.

This figure has been redacted for copyright reasons. The copyright owner is IOP Publishing Ltd. See print version, or figure 12 in [67].

Figure 5.5: Reproduced from [67]. Calculation from equation 5.35 of the ground quasienergy band of a 1D OL for four values of β , increasing from bottom to top. The parameters are $V = 7$, $E = 5.51$, and $\beta = 0.17, 0.35, 0.52, 0.69$. Note the progression from the classic ground band single well (though perhaps already with a small central maximum) to a double-welled structure as β increases. There are multiphonon avoided crossings near the centre and edges of the quasimomentum Brillouin zone that will cause heating if particles move near them, but the wells themselves are intact. This type of structure can be used to investigate ferromagnetism, but is also closely related to the scheme described in section 6.1.

5.4.3 Band engineering

For the purposes of this thesis, the most important consequence of the physics described in this chapter is that quasienergy bands can be made to have shapes quite different from the band structure of the static system by careful use of avoided crossings. This technique is known as *Floquet band engineering*, and figure 5.5 shows a calculation of great relevance to later chapters: a quasienergy band made to change from a single central well to a double well by increasing β . (See also [65].) This will be discussed again in section 6.1 using a more intuitive picture of hybridising the static bands, but the underlying maths is always equation 5.35 and the rest of this chapter.

6

The shaken Mott insulator experiment

In this chapter I will present the shaken Mott insulator experiment, starting with design and motivation, then covering some technical details added to the experimental procedure of section 3.4 to realise Floquet shaking, and finally presenting the current state of our results. All the experiments discussed in this chapter were performed with ^{87}Rb .

6.1 Theory of the shaken Mott insulator

6.1.1 Motivation

The Mott insulator to superfluid transition described at length in chapter 4 is a continuous QPT in the ground state of the Bose Hubbard Hamiltonian 3.16, occurring at a particular value of J/U . Given a Mott insulating state of ultracold atoms in an optical lattice, there are two experimentally straightforward ways to "melt" the Mott insulator and obtain a superfluid:

- Use a Feshbach resonance to decrease the value of U .
- Reduce the intensity of the lattice laser, therefore reducing the lattice depth. This will increase J and reduce U (since the Wannier states become less tightly localised, see equation 2.22), increasing J/U .

Both these methods can reduce J/U to the critical point required for the continuous QPT. The first method is usually better as it offers independent control over J and

U , while the second method is necessary if the species in question has no Feshbach resonance. Otherwise they are interchangeable.

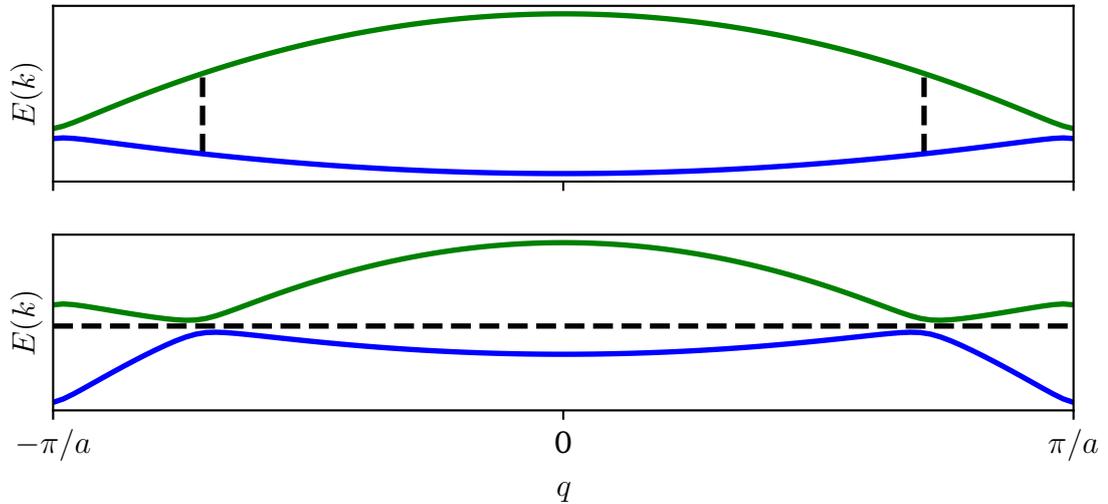


Figure 6.1: Toy illustration of our band engineering scheme, as a summary of the Floquet physics covered in chapter 5. This figure shows the two lowest energy bands of an optical lattice (compare with figure 2.1). Shaking the lattice at a frequency ν effectively raises the ground band by an energy $h\nu$. This allows us to bring the bands close together (top) and then cause them to cross (bottom), with the points joined by dotted lines in the top plot brought into coincidence. This results in avoided crossings, with a new ground band forming from a mixture of both static bands. Note that the minimum of the new ground band is at $q = \pm\pi/a$ rather than 0.

Here I will present an alternative method, proposed in [68], of melting a Mott insulator: coupling the ground band to the first excited band by periodic driving, a procedure called Floquet band engineering. (The theory underlying this was developed at length in chapter 5.) While we often ignore the first excited band in the context of OLs, it can be described in much the same terms as the ground band, except that it is inverted and has a stronger curvature (upper half of figure 6.1). This results in the value of J/U being generally higher in the excited band¹.

The effect of coupling these bands, as sketched in the lower half of figure 6.1, is that a ground band can be produced with global minima at $q = \pm\pi/a$ rather than $q = 0$. The value of J/U in this new minimum is larger than at $q = 0$, since it has the character of the first excited band. This opens up the possibility of choosing parameters such that the ground state of the static system is an MI, but the ground state of the driven system is a superfluid at $q = \pm\pi/a$: the π -SF mentioned in section 5.2.1. We simply need a ground band with J/U in the MI regime and an excited band with J/U sufficient for an SF.

This intuitive picture already suggests that the transition may be made discontinuous using this method: figure 6.1 resembles the discontinuous Ginzburg-Landau poly-

¹Although it describes the ground band, the result from equation 2.12 that J is proportional to the total bandwidth holds for the excited bands as well.

mials of figure 1.3, with two minima at different values of some parameter passing each other in energy. As we shall see, this prediction is born out both by more sophisticated simulations and by experimental evidence.

Contrast with dynamic localisation

There is a risk of confusing this scheme with the hopping renormalisation described in section 5.2.1. In that section we showed that Floquet driving can renormalise the hopping energy J by a Bessel function, allowing it to be tuned to arbitrarily small values (dynamic localisation). This will cause interacting bosons in a superfluid state to make the transition to a Mott insulator, as predicted in [131] and observed in [63].

The key difference is that this hopping renormalisation is a single-band effect, i.e. the driving frequency is much smaller than the band gap. This allows rescaling of the ground band but cannot fundamentally change its shape, and is no different in effect from the Feshbach and lattice depth methods described above: in particular, it does not allow the transition to be made discontinuous, the chief reason for our interest in the band engineering scheme.

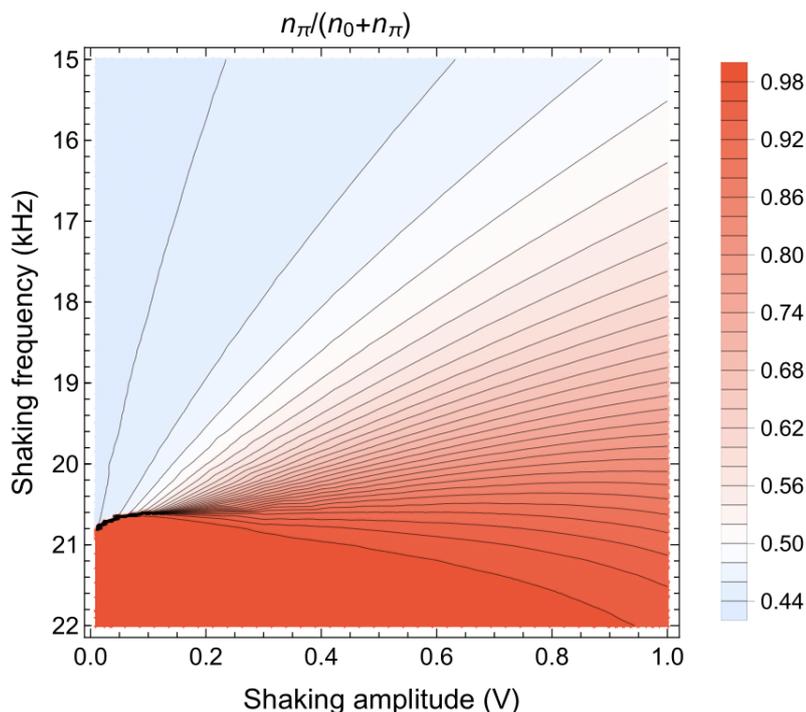


Figure 6.2: Shovan Dutta's calculated phase diagram for the effectively 1D lattice, with lattice depths 30, 10 and 30 along X, Y and Z. The colour bar indicates the population of the quasimomentum peaks at $q = \pm\pi/a$ as a proportion of the total atom number, which we expect to be ~ 1 in the π -SF and much lower in the MI which does not have well-defined peaks. The abrupt change at the black phase boundary for $V < 0.1$ V suggests a discontinuous phase transition, while the more gradual change at higher amplitudes is presumably a continuous phase transition.

6.1.2 Calculated phase diagrams

Figure 6.2 shows Dr Shovan Dutta’s numerical simulation of the phase diagram that is the focus of this chapter. It corresponds to figure 6.1 in the case where the ground band J/U is small enough to create an MI, but the first excited state has a J/U in the SF regime. By controlling the shaking parameters we can move between two phases: the MI and π -SF. By its very nature we expect the π -SF to be associated with most atoms having $q \approx \pi/a$, while the MI has no such structure and remains centred around $q = 0$. We define the ratio

$$R := \frac{n_\pi}{n_0 + n_\pi}, \quad (6.1)$$

where n_π is the population of atoms with $q \approx \pi/a$ and n_0 the population close to $q = 0$, as a metric for distinguishing these phases²: $R \approx 1$ indicates a π -SF, while the MI has a much lower value that we find both experimentally and numerically to be ~ 0.4 to 0.5 .

Note in particular that at low (but non-zero) shaking amplitudes the change from blue to red in figure 6.2 is extremely abrupt, suggesting a very rapid jump in R as we scan the frequency. This is suggestive of a discontinuous QPT as hinted above. At higher amplitudes, however, the change becomes smooth, suggesting behaviour similar to the usual continuous MI-SF transition. Dr Dutta has carried out a more involved analysis of correlation functions around the transition, not presented here, that supports these conclusions.

Section 1.4.1 noted that discontinuous QPTs are far less commonly observed than the continuous variety. Given that rarity, we believe that this work is the first instance of a continuous QPT that can be tuned to become discontinuous by Floquet shaking. In section 6.3 I will endeavour to convince the reader that we have realised this behaviour, but before that I will describe some of the methods used in doing so.

6.2 Technical implementation

In section 3.4 I gave a brief summary of the steps involved in creating a BEC of ^{39}K or ^{87}Rb atoms. The shaken MI experiment requires a couple of extra steps, which I will explain here.

²Note that R is not the order parameter ψ , which was identified in section 4.3.2. R is merely a convenient quantity for distinguishing the phases experimentally (figure 6.5).

6.2.1 Adiabatic loading of an optical lattice

Adiabatic and diabatic changes

Recall the discussion of section 5.1. There we argued that the time dependence of a state

$$|\Psi\rangle = \sum_i c_i |\phi_i\rangle, \quad (6.2)$$

subject to a time-dependent Hamiltonian $\hat{H}(t)$, is given by

$$|\Psi(t)\rangle = \sum_i c_i(t) |\phi_i(t)\rangle, \quad (6.3)$$

provided that $\hat{H}(t)$ changes slowly enough to satisfy the adiabaticity condition 5.5. If we consider the special case where $|\Psi\rangle = |\phi_0\rangle$, the ground state of $\hat{H}(0)$, we see that $|\Psi\rangle$ will always be the instantaneous ground state of $\hat{H}(t)$. This is called *adiabatic following* of the ground state³.

Now recall instead the discussion of section 3.5.3. There we discussed the effect of suddenly changing the Hamiltonian between two values, \hat{H} and \hat{H}' , with corresponding eigenbases $|\phi_j\rangle$ and $|\phi'_j\rangle$. If initially $|\Psi\rangle = |\phi_0\rangle$, the ground state of \hat{H} , an instantaneous change will not change $|\Psi\rangle$ itself, but will project it onto the eigenbasis of \hat{H}' as some general $|\Psi\rangle = \sum_j c_j |\phi'_j\rangle$. We would call this change *diabatic*, indeed the most diabatic change possible, by contrast with the adiabatic case.

Adiabatic loading

Kapitza-Dirac diffraction requires the latter, and is achieved experimentally by switching the lattice on and off suddenly to create a highly non-equilibrium state. However, the superfluid and Mott insulator are both ground states (depending on U/J), so we need to employ an adiabatic change to move from a BEC in the ground state of the dipole trap to a superfluid or Mott insulator in the ground state of the lattice.

This requires the intensity $I(t)$ of the lattice laser to switch on slowly. While various time dependences are possible, the simplest is $I \propto t$, with a slope chosen to be slow enough for adiabaticity. The required speed can be found by loading and then unloading the lattice using a triangular pulse, which will cause excitations in the BEC if done too fast (see [57, 117] and section 3.6.2). We achieve this using an intensity regulating control loop, measured by a pickoff plate and photodiode and with the set point given by a function generator. This allows arbitrary control of $I(t)$, with the experiment beginning once the desired lattice depth has been reached adiabatically.

³Strictly speaking, since the experiment is not at $T = 0$ we are not fully in the ground state, so adiabaticity means more that no significant excitations are created above the existing thermal ones. One could say we adiabatically follow almost the ground state.

6.2.2 Methods of lattice shaking

A 1D optical lattice consists of a laser beam retro-reflected by a mirror to create a standing wave. There are three experimentally straightforward ways of introducing a time-dependent modulation to such a lattice.

- Modulate the intensity of the light.
- Modulate the frequency of the light to change the number of wavelengths that fit along the beam path.
- Modulate the position of the retro mirror.

In our experiment we modulate the frequency, at least for now. I will discuss these methods one by one to explain why.

Amplitude modulation

This is most easily achieved using an acousto-optic modulator (AOM). Modulating the RF power delivered to the AOM changes the intensity of the diffraction order that reaches the experiment, therefore allowing arbitrary control over the laser intensity. This doesn't result in any change in alignment so it's very easy to implement.

However, this method is unsuitable because of parity. Rather than a term of the type Fx considered in chapter 5, which has odd parity, amplitude modulation adds a time-dependent prefactor to the term $V(x)$ in the Hamiltonian. This has no directionality, and therefore the time-dependent term has even parity. Since our intention is to couple the even-parity ground band to the odd-parity first excited band, we require an odd-parity perturbation.

Frequency modulation

Frequency modulation changes the relative phase of the incident and retro-reflected lattice beams, which in turn changes the positions of the standing wave maxima and minima. This is an odd-parity perturbation capable of coupling the ground and first excited bands, and can be achieved straightforwardly by modulating the frequency of the RF signal to the AOM described above.

This method has a minor technical complication. Changing the frequency of an AOM also changes the deflection angle of the diffraction orders, and consequently the alignment of the laser after the AOM. However, this can be overcome by retro-reflecting light to pass through the AOM twice, a technique known as a double-pass AOM, so that the frequency shifts add but the deflections cancel. This is the method used to obtain all the results shown in this chapter.

A more significant limitation of this method is the amplitude of modulation. A double-pass AOM can achieve a frequency modulation on the scale of 10 MHz without significant alignment issues. This proved more than sufficient for the experiments described here, since our shaking is resonant with the band gap and only small amplitudes are required. However, it may well be that future work, especially relating to phasons (section 2.2.3), will benefit from much larger modulations, perhaps of several wavelengths, corresponding to amplitudes orders of magnitude higher than can be achieved this way.

Modulation of the retro mirror position

This method requires a significant time investment: the retro mirror must be mounted on a piezoelectric actuator to allow rapid and controlled modulation of its position. This actuator acts as a capacitor, requiring high currents in order to be driven at high frequencies. Precise control requires an interferometer and a control loop, which must be positioned between the atoms and the retro mirror – a region where space is at a premium. Due to this barrier to entry we have not yet implemented any piezo-actuated retro mirrors in our experiment, and we continue to use frequency modulation for now.

However, as noted above, I believe this method has considerable promise in the long term for enabling larger modulations. Achieving large modulations at high frequencies requires very large currents, and this has proven challenging in initial tests, but at some point these challenges will likely be worth overcoming.

6.3 The experiment

I will now present the results of our investigation of the phase diagram in figure 6.2. While I was involved in planning this work, and helped prepare most of the required apparatus – for example, I designed the layout of optics for the vertical lattice – delays caused by the pandemic resulted in this work being carried out by my colleagues Bo Song, Jr-Chiun Yu, and Shaurya Bhave while I wrote this thesis, and so I did not make the upcoming measurements myself. My colleagues investigated the shaken lattice initially in 1D, and at the time of writing have just begun to extend their work to the 2D system.

6.3.1 Preliminary work and characterisation

Lattice loading

We begin by loading the BEC adiabatically into a 3D square lattice – call the three lattice directions X, Y and Z. A lattice depth of $30E_r$ is used in the X and Z directions, much deeper than the $5E_r$ required for the tight-binding regime, so that hopping in these directions is negligible and does not contribute to the dynamics. This creates

an effective 1D system along the Y direction, with a lattice depth that can be varied independently. An effective 2D system is similarly achieved by using a lower lattice depth along both X and Y while keeping the Z lattice deep.

Using such a strong blue-detuned lattice along all three directions will tend to push the BEC out of the lattice, as I will describe in section 7.1.1. Since this experiment does not involve transport there is no need for careful balancing of this effect, and we roughly compensate simply by keeping a strong harmonic trap that confines the BEC within the lattice. Because the BEC is interacting, tight confinement causes quantum depletion (section 3.3.3) which is visible as a classical-appearing cloud in the images below rather than the sharp peaks of previous chapters.

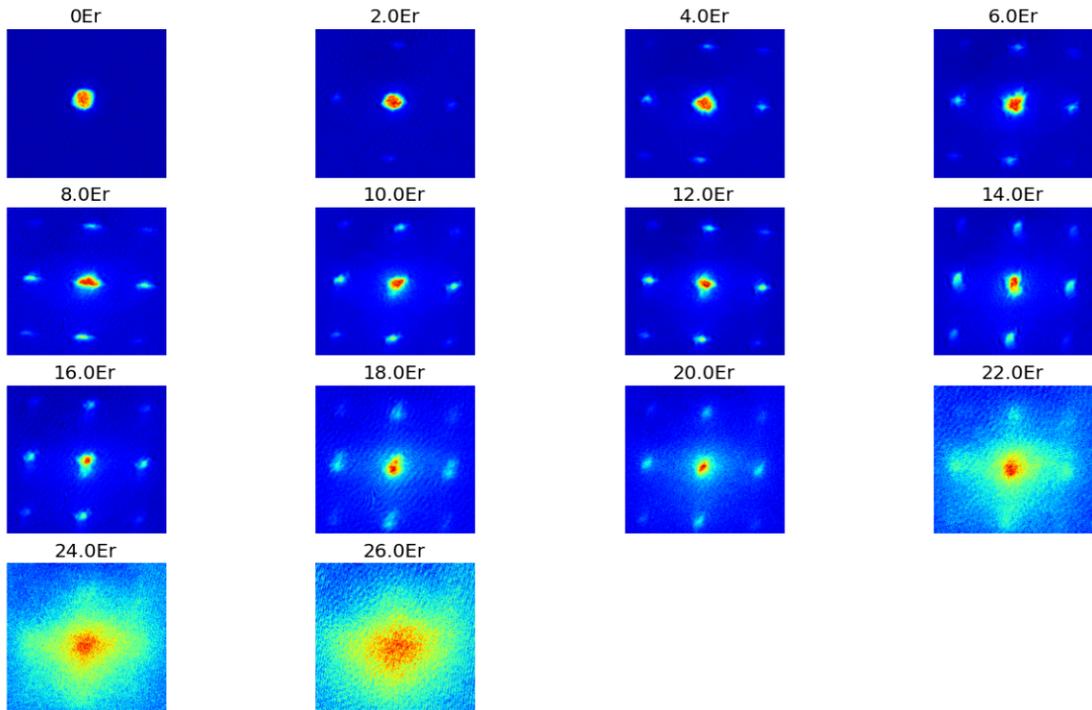


Figure 6.3: A series of TOF images taken after ramping up the 3D cubic lattice to the depth shown over 100 ms and holding for 25 ms. The images up until about $12E_r$ are characteristic of the superfluid, with sharp peaks and fairly constant background. From perhaps $14E_r$, but certainly by $18E_r$, the sharpness of the peaks has markedly decreased and the background is much more prominent. This indicates an MI close to the phase boundary with finite correlation length. At $24E_r$ there is no longer any sign of peaks and the system is deep into the MI.

Characterisation without shaking

To confirm the correctness of the static lattice, we first observe the conventional MI-SF transition, as laid out in chapter 4. To check the transition both ways, first we ramp the 3D lattice linearly up from 0 to a variable final depth and hold there before imaging in time of flight (figure 6.3), then we pick a maximum depth in the MI regime and image after ramping back down to various depths (figure 6.4). This confirms that the phase

transition is real and reversible and that the loss of coherence is not simply due to heating. These measurements were chosen to be directly comparable to the original observation of this transition in an optical lattice in [89].

Note that figures 6.3 and 6.4 appear not quite to give a consistent critical lattice depth for the MI-SF transition: it seems to occur around $\sim 13E_r$ in figure 6.3, but a little lower, around $\sim 8.5E_r$, on the way down in figure 6.4. This probably indicates simply that the cloud has heated slightly, pushing the ordered SF phase to higher J/U and lower lattice depths.

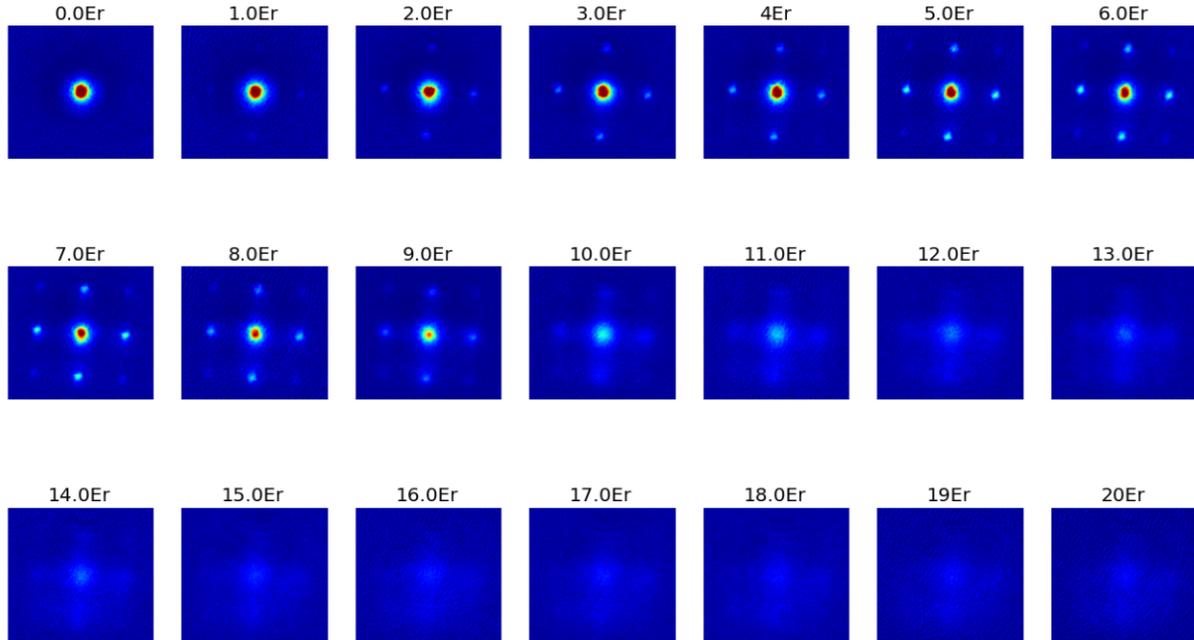


Figure 6.4: These images were taken by ramping the 3D lattice up to $20E_r$ as for figure 6.3, then ramping back down to the indicated depth over 20 ms. The reappearance of sharp peaks around $8E_r$ indicates that the system has re-established long-range phase coherence and re-entered the SF phase, and the eventual return to a BEC with no excitations when the lattice is switched off indicates that the entire process was adiabatic.

Implementation of shaking

Now we load the cloud directly into the MI regime, using a lattice depth along Y of $10E_r$. Then we begin to shake. This can be done in two ways:

- Suddenly. This entails beginning to shake abruptly at the desired frequency and amplitude.
- Gradually (or adiabatically). Both the frequency and the amplitude of the modulation are (say linearly) ramped up to the desired value.

Both these methods were tried, and found to produce broadly consistent experimental results. However, the lifetime of atoms in the lattice is shorter in the case of sudden shaking, presumably because the abrupt change introduces an admixture of excited states in the driven basis. For this reason a gradual start to shaking was used instead.

Once we reach the desired frequency and amplitude of modulation with a linear ramp, we abruptly switch off all potentials and image in time of flight.

Distinguishing phases

Section 6.1.2 identified the ratio of populations, $R := n_\pi/(n_0 + n_\pi)$, as a useful parameter for distinguishing the MI and π -SF phases, and so we require a method for measuring this experimentally. Figure 6.5 illustrates how R can be extracted from a TOF image, such that this is all we require to distinguish between the MI and π -SF states.

6.3.2 One dimension

Our 1D measurements centred around two different paths in phase space, illustrated in figure 6.6, which I will call the "direct" and "indirect" paths. The purpose of this relates back to the discussion of section 1.4.2 about the dynamics of crossing QPTs in Landau-Zener sweeps. Since we expect the PT between the MI and π -SF to be discontinuous at low amplitudes, crossing the boundary should result in metastable persistence of the MI phase and no sign of a PT. Therefore, we expect that the results of using the direct path will differ markedly from the calculated phase diagram. The results of this method, both numerical and experimental, are in figure 6.7. Note that the MI continues right down to 23 kHz at low amplitudes, contradicting the phase diagram of figure 6.2.

By contrast, the indirect path avoids crossing the discontinuous QPT, instead moving through the region where we expect the transition to be continuous. The results of numerics and experiment for this path are in figure 6.8, once again showing good agreement, but showing a different result from figure 6.7 at high frequencies and low

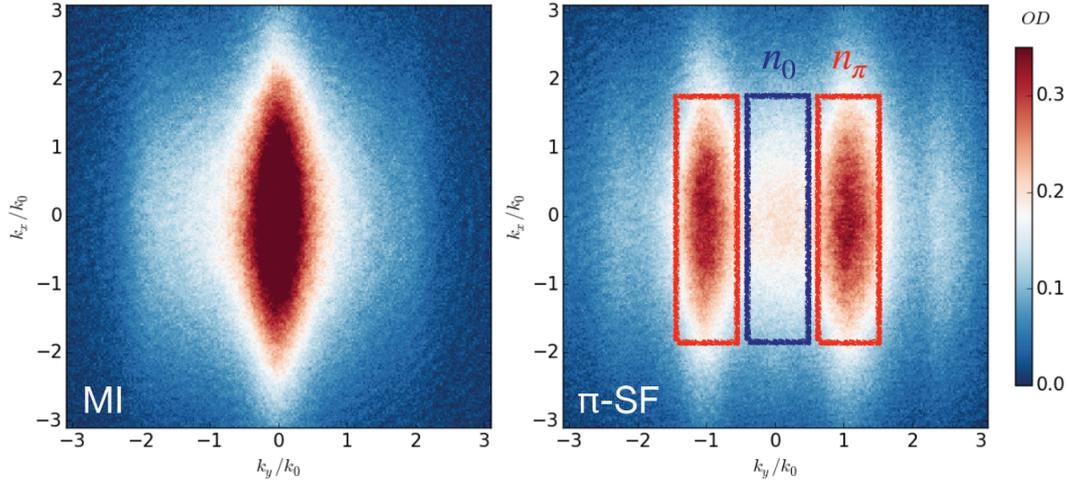


Figure 6.5: Illustrative TOF images of the MI and π -SF. The MI shows a single broad peak, elongated in the vertical direction due to compression by the dipole trap and vertical lattice, centred at $\mathbf{k} = \mathbf{0}$. The π -SF, on the other hand, has a large majority of atoms at $k_y = \pm \frac{\pi}{a}$, with a much smaller population in the centre. (Note that even in this image the momentum peaks are rather broad: this is either a result of quantum depletion following compression in the dipole trap, or a consequence of crossing the QPT too quickly to establish long-range coherence.) To distinguish between these states we use the ratio R (equation 6.1), calculated by using fits to find the number of atoms in the three boxes marked on the right-hand image.

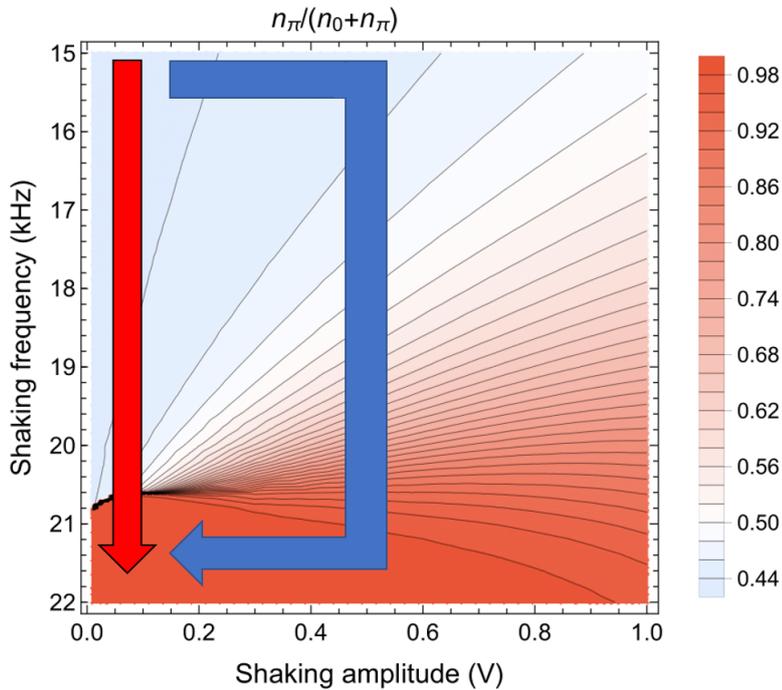


Figure 6.6: In order to locate the discontinuous phase transition we used two different paths from the MI to the π -SF: the "direct" path (red arrow) and "indirect" path (blue arrow). The direct path crosses the discontinuous phase boundary, while the indirect path avoids this and moves through the region where we expect a continuous transition.

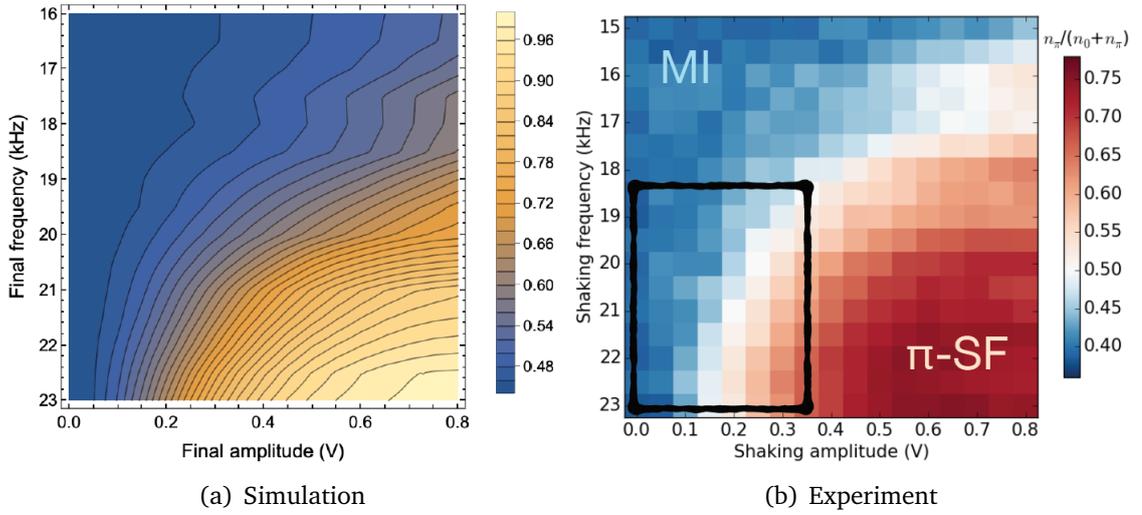


Figure 6.7: This figure shows the results of using the direct path from figure 6.6, with Shovan Dutta’s numerical prediction on the left and our measurements on the right. In both cases we ramp the amplitude of shaking from 0 to the indicated value over $125 \mu\text{s}$ with a fixed frequency of 15 kHz, then ramp the frequency from 15 kHz to the indicated value in $600 \mu\text{s}$ before taking a time of flight image and measuring R as indicated in figure 6.5. Numerics and theory show very good qualitative agreement. The black box is included for comparison with the same region in figure 6.8.

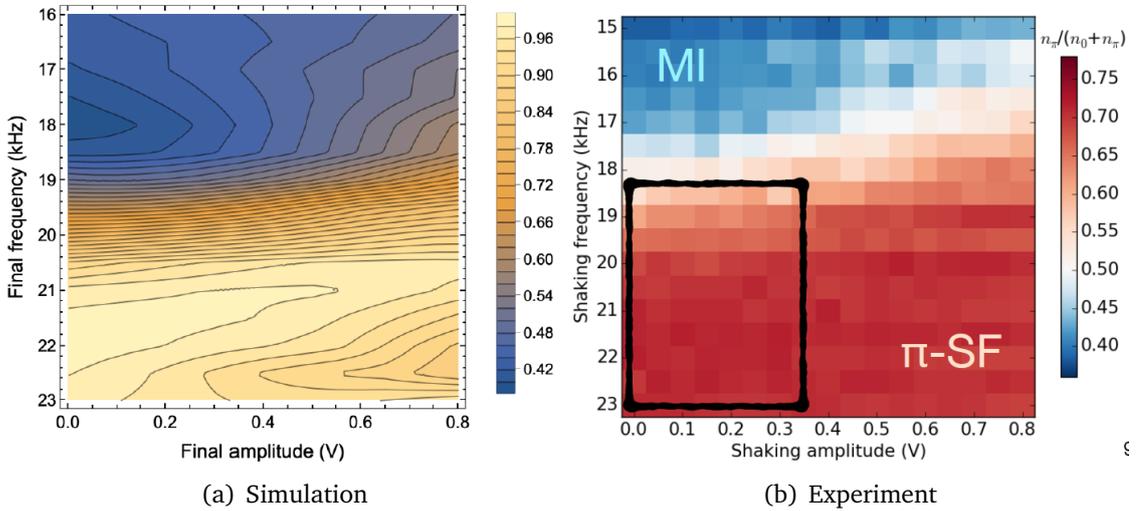


Figure 6.8: This figure is similar to figure 6.7, but using the indirect path from figure 6.6 for both numerics and experiment. To avoid the expected location of the discontinuous phase transition, we first ramp the shaking amplitude to 500 mV in $125 \mu\text{s}$ at 15 kHz, then ramp the frequency to its final value in $600 \mu\text{s}$ at constant amplitude, then finally sweep the amplitude of shaking to the indicated value in $300 \mu\text{s}$. Once again there is good agreement between numerics and experiment, and the results are very similar to figure 6.7 except for the region within the black box. This is consistent with the presence of a discontinuous phase transition: see main text.

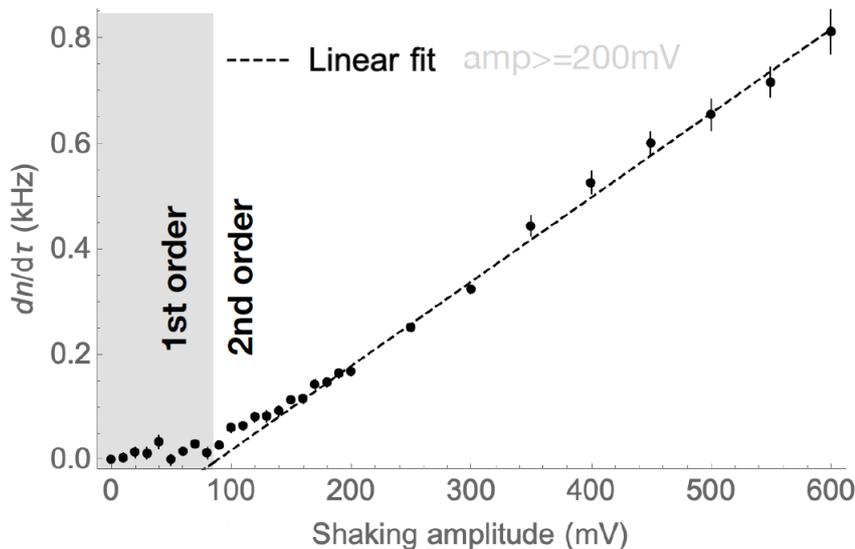


Figure 6.9: This figure offers an attempt to locate where the discontinuous phase transition suggested by figures 6.7 and 6.8 becomes continuous. For this measurement the frequency was swept from 15 kHz to 21 kHz at a series of different amplitudes (horizontal axis). The vertical axis shows the gradient when R is plotted against the duration $0 < \tau < 1$ ms of the frequency sweep, expressed here as $dn/d\tau$. We observe an initial flat region where the sweep duration has no effect, followed by a linear increase as we move to higher amplitudes. This indicates that above about 100 mV there is some finite adiabaticity time scale, so that while all of these sweeps are too fast they at least become *more* adiabatic for longer sweeps. Below this value, however, R does not increase no matter how slowly we go, showing that adiabaticity is impossible.

amplitudes, looking instead much more like figure 6.2. Now that we cross only a continuous PT, it becomes possible to adiabatically follow the true ground state.

Figure 6.9 investigates the time scales of adiabaticity more closely, searching for the amplitude at which the transition switches from discontinuous to continuous. It shows that if we cross the transition at an amplitude below ~ 100 mV we cannot get any population into the π -SF ground state, at least for the rates of crossing used here. Above this point, on the other hand, we begin to see a linear time dependence, with better and better adiabatic following as we go slower. This qualitative difference in behaviour is an indication of the expected change from a continuous QPT to discontinuous, achieved entirely by Floquet driving.

6.3.3 Two dimensions

Method

The lattice configuration described above can be made effectively two-dimensional simply by lowering the lattice depth of X to match Y , leaving Z as the only lattice deep enough to prevent hopping. However, the shaking can now be performed in two

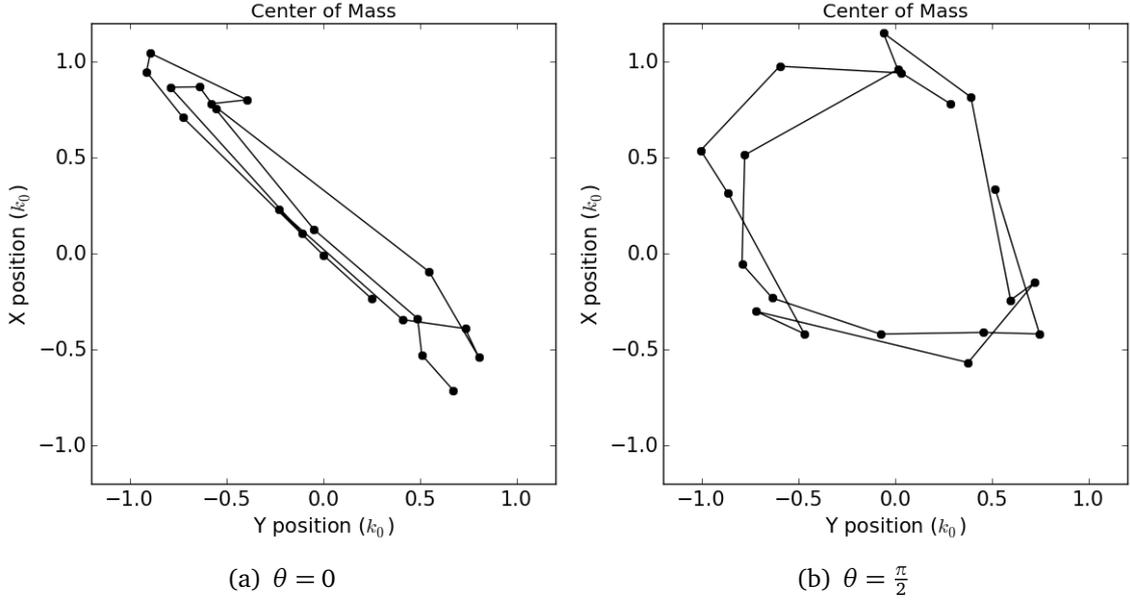


Figure 6.10: Micromotion during shaking along both the X and Y axes for two different values of the relative phase θ . The data points show the "centre of mass", or weighted average, of TOF images taken every $0.1T$ (where T is the period of shaking) for durations between 5 and $7T$. In the case that the two drives are in phase the net effect is linear shaking along the diagonal, which we see in figure a): this is equivalent to shaking along one direction. However, b) shows the case where the two drives are out of phase by $\frac{\pi}{2}$, and in this case the micromotion is circular as we might expect. Though we don't yet know what implications this could have for the long-term dynamics, it's at least interesting to see that we can bring about qualitatively different micromotion from the 1D case.

different ways: we can shake along only one direction as before, or shake along both X and Y. The former is a more direct generalisation of the 1D case, and so we have decided to explore the latter, which has more potential for qualitatively new effects as the phase θ between the two modulations becomes a new control parameter.

From a technical viewpoint this requires only the construction of a second double-pass AOM to allow frequency modulation along a second direction, or a second piezo-actuated mirror should we one day move to that method.

Results

Results of the 2D shaking are extremely preliminary at the time of writing, but we have already observed behaviour in the micromotion of the shaken system consistent with our expectation. Figure 6.10 shows how the distribution of atoms in momentum space varies over a few cycles of shaking for $\theta = 0$ and $\theta = \frac{\pi}{2}$. While this provides no great insight yet into the 2D phase diagram, it at least suggests that the experimental system is working and hints that θ could be an interesting variable to play with in future measurements.

7

Future directions

In this chapter I will outline our future plans for the quasicrystal machine as they currently stand. This includes both technical additions and experimental proposals, which will build on the work presented in this thesis to open up a wide range of possibilities.

7.1 Technical additions

There are two primary technical additions that still remain to be fully implemented: a vertical red-detuned beam to compensate for the anticonfining effect of the blue-detuned lattice, and a digital micromirror device (DMD) for the projection of custom optical potentials.

7.1.1 Anticonfinement compensation

I will only sketch the principles of anticonfinement and how we plan to address it, as a far more authoritative discussion may be found in the thesis of my colleague Matteo Sbroscia [69].

What is anticonfinement?

Future experiments involving real-space transport will run into the issue of *anticonfinement*. To understand this effect, consider a 1D OL made by retro-reflecting a single blue-detuned laser (figure 7.1). Along the direction of propagation of the beam, the

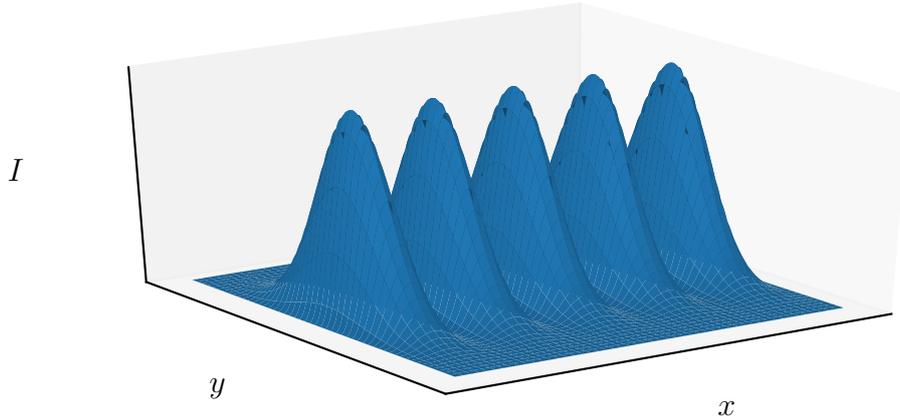


Figure 7.1: The intensity distribution $I(x, y)$ of a retro-reflected laser beam propagating along x , taking into account the beam's Gaussian profile in the transverse direction y (chosen arbitrarily; the picture is similar along any transverse direction). If the laser is blue-detuned atoms will tend to be pushed out of the beam, away from the highest intensity region, while the opposite will happen for red detuning.

potential is a sinusoid with constant amplitude¹. However, in both transverse directions this sinusoid is modulated by the beam's Gaussian profile.

Now, since blue-detuned light has a repulsive effect via the dipole force (section 3.1.1), atoms will occupy the *minima* of this potential and will not see the shape of $I(x, y)$ around the maximum. However, the curvature at the minima decreases away from the axis of the beam, and that in turn reduces the on-site ground state energy². The energy gradient gives rise to a force pushing the atoms to the edges of the beam: this is anticonfinement. (In a red-detuned lattice this is offset by the dipole force pulling atoms into the maxima of intensity, resulting in an overall *confining* effect instead.)

In our experiment the situation is made more complex by the addition of more lattice beams, but the result is similar. In the 2D plane of the experiment, a number of blue-detuned retro-reflected beams interfere to produce our OL. While on the atomic scale this produces some complicated arrangement of minima and maxima, the intensity envelope is simply a maximum at the centre (where all the beams have their individual intensity maxima) that tapers off towards the edges. The overall effect is to push atoms towards the edges of the 2D lattice.

¹Or, more precisely, a very slowly changing amplitude with its maximum at the beam's focus. We neglect this because we choose lenses such that our lattice beams have a Rayleigh range far longer than any experimentally relevant distance.

²Analogous to a simple harmonic oscillator with decreasing ω .

Method for compensation

Anticonfinement prevents accurate measurements of atomic transport. For example, an initially confined cloud allowed to expand outwards (in the manner of [32]) will expand faster than it should. We can eliminate this effect by adding an additional red-detuned beam (leftover light from our ODT laser at 1064 nm) perpendicular to the plane of the experiment. Such a beam will have a confining effect, tending to pull atoms towards its axis, and its waist and intensity can be optimised to compensate for the leading terms of the anticonfining potential. With only two degrees of freedom this method can never provide a perfectly flat lattice, but it can expand the region around the centre where anticonfinement is negligible, and is recommended by its simplicity.

Experimental relevance

While section 2.2.5 laid out a method of observing a localisation transition using only time-of-flight imaging, we expect to use in-situ observations for related measurements in the future. For example, we hope to try preparing atoms in an out-of-equilibrium state in an OL, such as a density wave or a tightly confined peak, to then observe whether they equilibrate, and how quickly, in the spirit of e.g. [38]. These measurements are impossible if the effects of the lattice are conflated with the anticonfining potential, so they require it to be compensated as described above.

7.1.2 DMDs and designer potentials

Here I will outline the operation and potential uses of a *digital micromirror device* (DMD) for our experiment. I expect the forthcoming thesis of my colleague Jr-Chiun Yu to cover this in much more detail, so I will not try to do more than give the general idea.

The operation of a DMD

A DMD is an array of tiny³ mirrors which can be individually positioned at one of two angles (see figure 7.2). Arbitrary patterns can be made using the mirrors as pixels, and an imaging system can be used to magnify or demagnify this pattern and project it onto another plane. While the most common use for DMDs is in projectors, which use a large magnification to make an image on a wall or screen, we are interested in using a large demagnification to project arbitrary optical potentials onto the 2D plane of our experiment. We will use a wavelength of 532 nm, creating a repulsive potential.

It is worth asking why, if DMDs make it possible to construct arbitrary potentials, we bothered building an OL by interfering lasers in the first place rather than projecting a square or quasicrystalline lattice with a DMD. The answer is that the fidelity of the

³On our DMD, a Vialux V-9001 VIS, the micromirrors are 7.56 μm square.

DMD's projection is never perfect, as it will be limited either by the size and number of pixels on the DMD or by the resolution and aberrations of the imaging system used for projection, and this would prevent taking images with the clarity and precision visible in section 3.5.3.

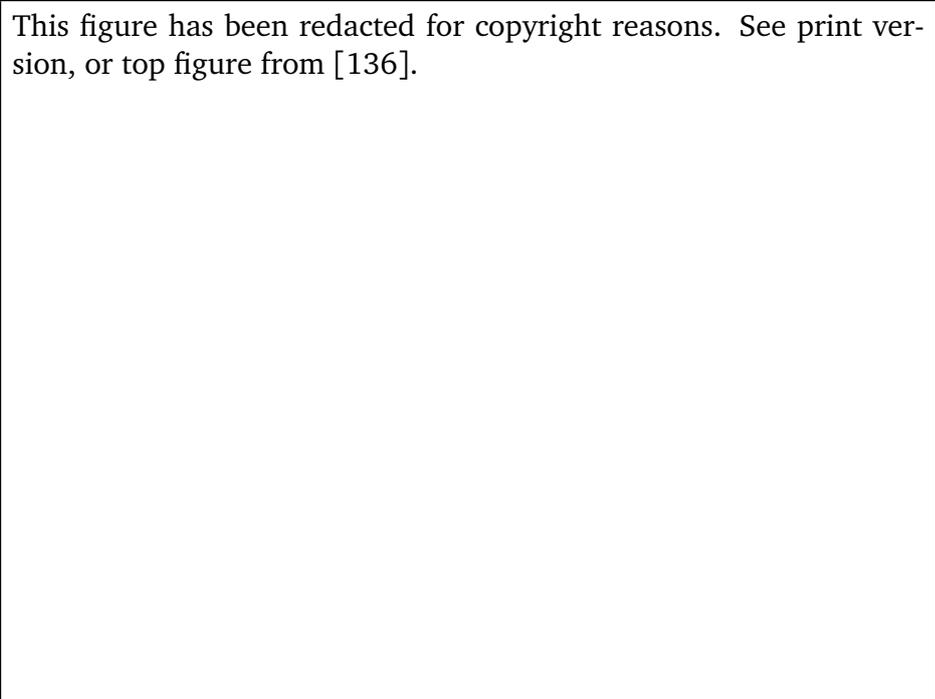


Figure 7.2: A microscope image of the surface of a DMD showing the micromirrors, taken from [136]. Note the nearer mirror tilted to the right while the two behind are tilted to the left. During operation all mirrors are tilted, and only one of the two directions sends light to the image plane. This allows each mirror to act as a pixel that's either on or off, drawing a 2D image. The resolution is limited either by the number of micromirrors or by the fidelity of the imaging system that projects them.

Patterns of interest

There are three patterns we are initially interested in using a DMD to project onto our existing OL: a sinusoidal modulation, a linear gradient, and a ring.

A sinusoidal modulation with a period of many lattice sites can be used to create a density wave in the OL. This is experimentally interesting as a way of investigating many-body localisation (MBL) by suddenly switching off the modulation (see e.g. [37, 39]). The appeal of this method is that if the cloud is able to homogenise it tends to do so rather quickly, since the maximum transport distance required is one period of the density wave, allowing a striking contrast with the MBL state that retains the modulation indefinitely.

A linear potential gradient can be used to measure compressibility (section 4.2.3). A compressible state like the SF or BG will show a linear density gradient that mimics the potential, but the incompressible MI will instead show a "wedding cake" structure

of plateaus joined by smooth SF gradients [126]. (Deep in the MI regime the SF regions disappear, resulting in the abrupt jumps of figure 4.5.) This is one way of distinguishing between the BG and MI states in particular (see below).

A ring of blue-detuned light is interesting because it provides hard-wall boundary conditions, which are impossible to achieve in ultracold atom physics without some kind of beam shaping. This is primarily of interest for the topological effects outlined in section 7.2.3 below, as hard walls are required for the observation of chiral edge states.

7.2 Experimental directions

In this section I will describe three major directions for our planned future investigations, which expand on work discussed in earlier chapters by combining different effects and introducing new elements to the Bose-Hubbard Hamiltonian. Our main ingredients are quasiperiodicity, Floquet driving, and custom potentials from a DMD.

7.2.1 The disordered Bose-Hubbard phase diagram

Our largest and most immediate experimental interest is adding another dimension to the SF-MI transition: quasi-disorder. Consider the following Hamiltonian, which has been altered from the Bose-Hubbard Hamiltonian of equation 2.21:

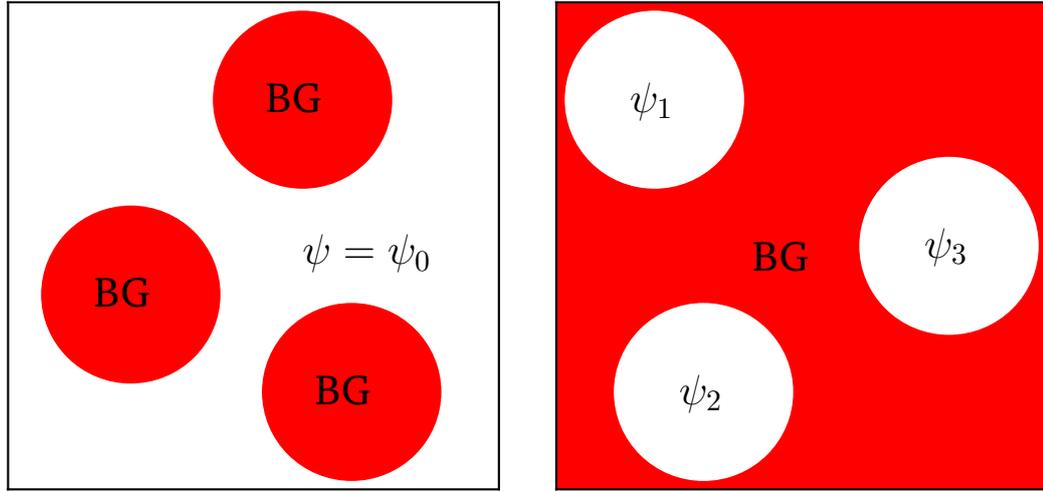
$$\hat{H} = -J \sum_{\langle i,j \rangle} (\hat{a}_i^\dagger \hat{a}_j + \hat{a}_j^\dagger \hat{a}_i) + \frac{1}{2} U \sum_j \hat{n}_j (\hat{n}_j - 1) - \mu \sum_j \hat{n}_j + \sum_j \delta_j \hat{n}_j. \quad (7.1)$$

The new term δ_j is a (pseudo-)random variable with mean 0 and half-width on the order of some value Δ^4 . This adds a third axis to the traditional MI/SF phase diagram of figure 4.7. There has been much exploration of this phase diagram for a 1D quasiperiodic lattice, both numerically [137–139] and experimentally [41]. However, it is only during the writing of this thesis that calculations have been made for a 2D quasicrystal, in fact with our specific lattice geometry in mind [140]. We would very much like to explore the 2D phase diagrams experimentally and verify these predictions.

The Bose glass

The most striking feature of the disordered phase diagram is that it contains not two, but *three* phases: the SF and MI are joined by the so-called *Bose glass* (BG), which begins to appear when $\Delta > 0$. The BG is a disordered, localised, compressible, insulating state with a gapless excitation spectrum [40]. While it occurs only in the presence of interactions and disorder, it remains unclear whether we should think of the BG as

⁴True randomness would have no correlation between the value δ_j on adjacent lattice sites, while in the quasicrystal the values are perfectly correlated but never repeat. Here we are interested only in the quasiperiodic case.



(a) Superfluid with insulating patches

(b) Bose glass with coherent patches

Figure 7.3: Artist's impression of the emergence of the Bose glass phase within a superfluid. In a) local insulating regions have begun to form at sites of unusually high disorder, but the superfluid maintains phase coherence everywhere else. In b) the insulating regions have joined up and the few remaining patches of phase coherence have lost touch with each other, in principle resulting in different phases. Note, however, that because the coherent regions are at locations of similar chemical potential, their wavefunctions will evolve with the same frequency and the phase difference can only be at most a constant offset rather than incoherence.

an example of MBL: MBL refers to a highly excited state that is unable to thermalise, while the BG is the ground state for appropriate values of Δ , U and J .

Physical intuition about the BG can be gained by considering the transition from each of the other two states. At low U , the SF becomes the BG when Δ increases above some critical value, which grows higher as U increases [140]. Close to the transition, local insulating regions are thought to appear where there is the most local disorder. (This is a special case of a general result for disordered systems called the *theorem of inclusions* [141]. See also [142] and references within.) These insulating regions grow larger as Δ increases⁵. Once they grow large enough to join up – more precisely, once it is no longer possible for a signal to percolate through the system and maintain phase coherence – the system transitions to the BG. The remaining isolated regions of phase coherence lose touch with each other, shrink, and ultimately disappear as Δ continues to grow. This behaviour is sketched in figure 7.3.

The MI behaves similarly, becoming a BG when Δ increases above some value that itself increases with U . There is some suggestion [138] that in this case the boundary is linear, which suggests that the condition is simply $\Delta \gtrsim U$. Again, this has an intuitive

⁵This sounds superficially reminiscent of the phase coexistence observed in discontinuous PTs. However, this transition is continuous, and there are several key differences. The insulating regions do not nucleate randomly or independently, appearing at locations dictated by the distribution of δ_j , and this further means that their wavefunctions evolve coherently and with a common frequency. Any difference in initial phase will be an excitation in the eventual BG that can be lost to a reservoir.

picture. Transport is impossible in the MI because there is an energy penalty U associated with a single hopping event. If $\Delta \sim U$, there will be adjacent sites with an energy gap comparable to U . This allows an atom to hop from the higher site to the lower site with no penalty, as the energy offset balances the increased interaction energy. Once this process becomes possible, the state is no longer an MI. This also helps us understand why the BG is compressible, unlike the incompressible MI: a sufficiently large Δ can compensate for the energy penalty of adding an additional atom.

Finally, we can try to understand why the BG energy spectrum is gapless like the SF, rather than gapped like the MI: an excitation at some infinitesimal energy ϵ is always possible because there will always be two adjacent sites with energy offset equal to ϵ , with the excitation consisting of one atom hopping across. The closing of the characteristic MI energy gap at the transition is another key distinction between the BG and MI states.

Methods of distinguishing phases

Experimental investigation of the quasi-disordered Bose-Hubbard phase diagram requires experimental methods for identifying the three phases: SF, MI and BG. Since each phase boundary involves only two phases, we propose three different techniques for investigating each transition.

The SF-MI transition may be imaged as before using the time-of-flight techniques discussed in section 4.3.3: the presence of sharp peaks is characteristic of the SF, and in the MI these peaks slowly broaden into a completely incoherent envelope as the correlation length shrinks to zero.

The MI-BG transition can be probed by periodic driving, not to perform any Floquet physics, but simply to measure the excitation spectrum [41]. A perfect MI will show sharp excitation peaks at U and $2U$. We expect the BG to show a broad, gapless feature from 0 to $\sim 2\Delta$, centred at Δ . The experimental method is to load the lattice adiabatically at some value of (U, Δ) , drive at a frequency ν for a fixed duration, then adiabatically ramp the lattice back down. Any excitations caused this way will result in a higher temperature of the final BEC, so that an excited state at energy $h\nu$ will be visible as a dip in the condensate fraction.

We may also distinguish these phases by compressibility (section 4.2.3). For example, in section 7.1.2 we considered how a linear potential projected by a DMD results in quite different behaviours for each phase. A similar measurement could be made without using a DMD by loading the lattice from the harmonic trap with a series of different trap frequencies (similar to e.g. [126, 143]).

The SF-BG transition is both an extended-localised transition and a conducting-insulating transition. Due to the former, the time-of-flight method described in section 2.2.5 should be directly applicable, with the time-scale of adiabatic lattice loading diverging in the BG. The latter can be observed with transport measurements. For example, the dipole trap can be compressed during lattice loading to give a tightly confined initial distribution in the lattice, after which the dipole trap may be switched off (see e.g.

[32]). This will result in ballistic expansion in the SF regime while the insulating BG phase should remain stationary, with the difference detectable by in-situ imaging. Alternatively, once the DMD is available, an initial density wave distribution (described in section 7.1.2) should also show quite different behaviour in each phase. Ideally we will be able to perform a mixture of loading and transport measurements to confirm the location of the phase boundary.

7.2.2 Floquet driving of an optical quasicrystal

In noninteracting systems, Floquet driving can mostly be prevented from causing destructive heating by a careful choice of shaking parameters, avoiding any resonances of the undriven system. Interacting systems, however, will typically be heated by driving at any frequency given enough time (see section 5.4.1). Intriguingly, though, it has been observed that this does *not* apply to MBL systems, which may show "transparency" to Floquet driving at sufficiently high frequency [144]. As is often the case, the effect of interactions (enabling heating) is counteracted by the effects of disorder (encouraging localisation).

We would like to see how this extends to localised states in the optical QC. While the effect is interesting enough to replicate for its own sake, especially since we can do so in 2D, it is also expected to have considerable experimental importance as a technique to study stable driven systems. There has accordingly been a large amount of theoretical work on the novel phases that may be accessed by Floquet driving of MBL systems [145–148], and exploring these phases is likely to be another promising avenue of future research.

7.2.3 Topology

Another area of future interest for our experiment is likely to be the simulation of topological phases. Recall our definition of the order parameter ψ from section 1.2. While in this thesis we have discussed numerous PTs, each with a different ψ , so far they have all had in common the feature that they are *local*, that is, ψ can be evaluated independently at each point in space. A topological phase, by contrast, has a *global* order parameter [149], also called a topological invariant. A phase where such an invariant is non-zero is then said to be *topologically nontrivial*.

The classification of topological invariants is a vast subject that I will not discuss here, but it can be done, with the result that the possible nontrivial invariants for a system are determined by its dimensionality and symmetries [150–152]. It is therefore natural for us to consider topological applications given our capability to simulate four spatial dimensions (section 3.6.1). I will discuss here some of our ideas for such experiments, though they are still at a very early stage of planning.

The quantum Hall effect

A common focal point of interest in topological phases is the (integer) quantum Hall effect (QHE), which gives rise to quantised conductance in 2D systems in the presence of a magnetic field [153]. This is associated with a topological invariant called the Chern number [154, 155], which is given by the integral of Berry curvature over the Brillouin zone.

A Chern number can be constructed in any even number of dimensions [156], with the $2n$ -dimensional version being called the n th Chern number. Just as the first Chern number is associated with the 2D QHE, it has been shown that a 4D generalisation of the QHE relates to the second Chern number [157, 158]. There has already been interest in investigating this with ultracold atoms, using for example topological charge pumping [52, 149] or synthetic dimensions based on internal atomic degrees of freedom [159, 160]. We feel that our experimental system may be able to make productive contributions in this area. For instance, our higher dimensions are extended, whereas synthetic dimensions are typically compact, with only a handful of states to move between in the added direction.

A characteristic experimental signature of the QHE is the presence of chiral edge states, which are lossless transport channels arising at the boundary between regions with differing topological properties [161, 162]. (The broader relation between a topologically nontrivial region and interesting behaviour at its edges is called the *bulk-boundary correspondence*, and also gives rise, for instance, to surface modes in 3D topological insulators [163, 164].) If we are able to create a QHE state in our optical QC, we hope to use hard wall boundary conditions projected with a DMD (section 7.1.2) to observe such edge states.

Topology in the quasicrystal

Recall from section 3.6 that the parent crystal of our optical QC is a 4D simple "hyper-cubic" lattice. Such a lattice does not have any interesting topological properties, and so we expect that our 2D slice through it must also be topologically trivial. This seems somewhat validated by the finding that a QC with even rotational symmetry, such as our eightfold case, cannot have non-zero Berry curvature [165], and therefore cannot have a Chern number, as long as time-reversal symmetry holds.

However, this limitation may also point the way forwards: the use of Floquet shaking along two directions to break time-reversal symmetry is well established in ultracold atom experiments [46]. The principle is essentially illustrated by the circular micromotion we have already observed (figure 6.10), with the chirality of the circle changing if the direction of time is reversed. While it's important to note that time-reversal symmetry is not the only important symmetry – as [150–152] cited above note, particle-hole symmetry also plays an important role – at least shaking may allow us to achieve non-zero Berry curvatures, and we hope that interesting topological effects will become accessible in the driven QC that are impossible in the static system.

This essentially reiterates the conclusion of section 7.2.2: Floquet shaking of interacting systems makes a wide variety of novel phases accessible, including ones with interesting topological properties. It is very possible that MBL and the accompanying transparency to heating will come in useful when pursuing these experiments.

Summary

This thesis has presented results from three experiments investigating three quite different problems: a 4D quantum walk, the localisation transition in a 2D quasicrystal, and the use of Floquet shaking to realise a discontinuous form of the Mott insulator to superfluid quantum phase transition. In each case we have explored the underlying physics and technical implementation, and we have seen that each experiment shows abundant promise for future investigation:

- The 4D quantum walk may be extended to longer times using a sequence of pulses, adding a new axis of experimental control.
- The localisation transition showed quite unexpected dynamics, appearing to depart completely from the Kibble-Zurek mechanism.
- Additionally, the effect of interactions on the localisation transition has only just begun to be explored.
- The work of extending the shaken Mott insulator from a quasi-1D system to quasi-2D is currently under way, and the relative phase of shaking along two directions adds a qualitatively new degree of freedom.

However, I believe there is still more promise in the fact that all three experiments were realised on the same machine, raising the possibility of combining all the techniques used to achieve them and simulating an increasingly rich array of systems. I expect the group to struggle in choosing what to pursue first: the quasi-disordered Bose-Hubbard phase diagram, potential MBL states in the quasicrystal and their transparency to Floquet heating, and the still-forming plans to simulate a 4D quantum Hall state will all soon be within reach. The addition of the DMD and possible future work with mixtures of atomic species open up a dizzying array of techniques and directions, all available on this one quantum simulator. Our goal from the beginning was to keep the machine as flexible as possible, and I believe this future scope, not to mention the range of different physics present in this thesis, is optimistic of success.

Returning briefly to Earth, the work in this thesis, while interesting in itself, is a sequence of stepping stones and preliminary sorties. We have shown that the optical quasicrystal can be used to simulate four spatial dimensions and that it offers an excellent system for studying disordered physics, and we have made use of Floquet driving to realise a quantum phase transition that can be tuned from continuous to discontinuous. The work of adding new features to the experiment still continues, though with the end now perhaps in sight, and I believe the most exciting physics is yet to come.

Acknowledgements

I'd like to take a moment to thank the people without whom this thesis couldn't have been written. First of all I am grateful to my supervisor, Ulrich Schneider, for giving me this opportunity in the first place, and for being a patient, considerate guide over the last four years. His calm, open-minded approach and strong belief in working as a team have been very valuable in tough times.

Next I would like to thank those lab colleagues I worked with the most. Matteo Sbroscia and Konrad Viebahn taught me a vast amount during the first half of my PhD, and those lessons will always stay with me. Later on, Jr-Chiun Yu and Shaurya Bhawe were stalwart companions during a disastrous period of equipment failure and the long road to recovery, and I will be eternally grateful for their determination and patience. I am also grateful to Bo Song for his invaluable experience and insight, and to Georgia Nixon for many fascinating discussions.

I would also like to thank those who have made invaluable contributions to the technical side of our experiments. Stephen Topliss has offered essential technical assistance and expertise, and displayed endless tolerance for replacing burned (or variously exploded) MOSFETs. I am also grateful to Shovan Dutta for his theoretical work, which underlies part of this thesis, and for being generally delightful and always happy to help. Gary Large provided a great deal of technical support and advice, and a lot of entertainment and humour while doing so, and Alan Turner saved us during many crises and provided endless bizarre plumbing attachments. I am also grateful to Chris Haniff, who has been a source of advice and guidance throughout my entire time at Cambridge, for his dedication to his students.

I would also like to acknowledge our friendly acquaintances⁶ in the kagome team. Dan Reed and Max Melchner offered endless and unhesitating assistance through the quasicrystal team's many ups and downs, and I have greatly appreciated their companionship and conversation, as well as that of Vanessa Koh and Luca Donini. I would like to thank Tiffany Harte in particular for endless conversations and a lot of extremely good vegetarian food, quite apart from her thoughtful insights and ready assistance in the lab.

⁶And occasional adversaries in the arena of shared equipment.

Supervising project students has been a consistently rewarding and intriguing part of my work as well, and I would like to briefly mention those I was most involved with: Callum Stevens, Marius Weber, Jonathan Mortlock, Robin Croft, Elvinas Ribinskas, Attila Szabó, Jared Jeyaretnam, Krzysztof Zamarski and Lee Reeve all brought their own personalities and approaches, and generally made our lab a more pleasant place to be. I also owe a debt of gratitude to Oliver Brix and Hendrik von Raven, without whom I am certain our experiment would never have functioned in the first place.

Finally I would like to offer thanks to my parents for their constant and unwavering support, and to a certain Dr J.F. Hughes for making all this worth it.



Dimension scaling in quantum walks

Here I will prove that the group velocity of a particle on a continuous quantum random walk in d dimensions is proportional to \sqrt{d} , as claimed in section 3.6.1. The formula for the speed of a wavepacket in a 1D system is given in equation 2.15, and we can generalise this to a separable d -dimensional lattice by applying it to each component v_j of the group velocity \mathbf{v} individually:

$$v_j = \frac{1}{\hbar} \partial_{q_j} E_n(\mathbf{q})|_{q_0} \quad (\text{A.1})$$

where the quasimomentum \mathbf{q} is now also a vector, with components q_j . The dispersion relation is likewise a higher-dimensional generalisation of equation 2.12:

$$E(\mathbf{q}) = E_c - 2J \sum_{j=1}^d \cos(q_j a). \quad (\text{A.2})$$

Substituting into equation A.1 we obtain

$$v_j = \frac{2Ja}{\hbar} \sin(q_j a) \quad (\text{A.3})$$

from which we can calculate the root-mean-square as

$$\begin{aligned} \sqrt{\overline{\mathbf{v}^2}} &= \sqrt{\sum_{j=1}^d \overline{v_j^2}} \\ &= \frac{2Ja}{\hbar} \sqrt{\sum_{j=1}^d \overline{\sin^2(q_j a)}} \quad \overline{\sin^2(x)} \equiv \frac{1}{2} \\ &= \frac{2Ja}{\hbar} \sqrt{\frac{d}{2}} \\ &\propto \sqrt{d}. \end{aligned} \quad (\text{A.4})$$

B

Theorem regarding the monodromy operator

I will prove the assertion 5.16

$$\hat{U}(t + T, 0) \equiv \hat{U}(t, 0)\hat{U}(T, 0) \quad (\text{B.1})$$

that underlies the central importance of the monodromy operator $\hat{U}(T, 0)$ in Floquet theory. Let us define an operator

$$\hat{O}(t) := \hat{U}(t + T, 0)\hat{U}^{-1}(T, 0) \quad (\text{B.2})$$

which then has the property

$$\hat{O}(0) = \hat{I} = \hat{U}(0, 0) \quad (\text{B.3})$$

where \hat{I} is the identity operator in the appropriate Hilbert space, and the second equality follows because the propagator \hat{U} can change nothing if no time has passed. This establishes that at time $t = 0$, \hat{O} and \hat{U} coincide.

Now, the time evolution of $\hat{U}(t, 0)$ is given by equation 5.15. Consider the following:

$$\begin{aligned} i\hbar\partial_t\hat{O}(t) &= i\hbar\partial_t\hat{U}(t + T, 0)\hat{U}^{-1}(T, 0) \\ &= \hat{H}(t + T)\hat{U}(t + T, 0)\hat{U}^{-1}(T, 0) && \text{By equation 5.15} \\ &= \hat{H}(t)\hat{U}(t + T, 0)\hat{U}^{-1}(T, 0) && \text{By periodicity of } \hat{H}(t) \\ &= \hat{H}(t)\hat{O}(t). \end{aligned} \quad (\text{B.4})$$

So $\hat{O}(t)$ obeys the same differential equation as $\hat{U}(t, 0)$. Since they also start from the same initial condition, we must then have that

$$\hat{O}(t) \equiv \hat{U}(t, 0). \quad (\text{B.5})$$

We can use this to prove the assertion 5.16:

$$\begin{aligned}\hat{U}(t,0)\hat{U}(T,0) &= \hat{O}(t)\hat{U}(T,0) \\ &= \hat{U}(t+T,0)\hat{U}^{-1}(T,0)\hat{U}(T,0) \\ &= \hat{U}(t+T,0)\end{aligned}\tag{B.6}$$

which is the assertion we were aiming to prove.

C

The shaken lattice Hamiltonian

In this appendix I will show how the Hamiltonian 5.30 can be changed into the much more useful form 5.31. The first step is to transform to the co-moving frame by defining

$$\Psi_{cm}(x, t) := e^{\frac{i}{\hbar}\Delta L \cos(\omega t)p} \Psi_{lab}(x, t). \quad (\text{C.1})$$

To see that this puts us in the co-moving frame, note that the expected value of position x using this wavefunction is

$$\begin{aligned} \langle x \rangle &= \int \Psi_{cm}^*(x', t) x' \Psi_{cm}(x', t) dx' \\ &= e^{-\frac{i}{\hbar}\Delta L \cos(\omega t)p} \langle x \rangle_{lab} e^{\frac{i}{\hbar}\Delta L \cos(\omega t)p} \\ &\equiv \langle x \rangle_{lab} - \frac{i}{\hbar} \Delta L \cos(\omega t) [p, x] \\ &= \langle x \rangle_{lab} - \Delta L \cos(\omega t). \end{aligned} \quad (\text{C.2})$$

We can now substitute Ψ_{cm} into the time-dependent Schrödinger equation, $i\hbar\partial_t\Psi_{lab}(x, t) = \hat{H}(x, t)\Psi_{lab}(x, t)$. First let's evaluate

$$i\hbar\partial_t\Psi_{lab}(x, t) = e^{-\frac{i}{\hbar}\Delta L \cos(\omega t)p} (i\hbar\partial_t\Psi_{cm}(x, t) - \Delta L \omega \sin(\omega t)p \Psi_{cm}(x, t)). \quad (\text{C.3})$$

The second term in this equation becomes part of a new co-moving Hamiltonian

$$\hat{H}_{cm}(x, t) = \frac{p^2}{2m} + \frac{1}{2}V_0 \cos(2k_{lat}x) + \Delta L \omega \sin(\omega t)p, \quad (\text{C.4})$$

where we have also redefined x according to equation C.2. The next step is to complete the square, obtaining

$$\hat{H}_{cm}(x, t) = \frac{1}{2m}(p + m\Delta L \omega \sin(\omega t))^2 + \frac{1}{2}V_0 \cos(2k_{lat}x) - \frac{1}{2}m(\Delta L \omega)^2 \sin^2(\omega t). \quad (\text{C.5})$$

We're now very close to the desired form 5.31, but the final term is still time-dependent. To remove this we have to absorb $\sin^2(\omega t)$ into Ψ_{cm} , but we have to do this carefully because $\langle \sin^2(\omega t) \rangle = \frac{1}{2}$ which can result in linearly growing, non-periodic terms. To fix this we use the fact that $\sin^2(x) \equiv \frac{1}{2}(1 - \cos(2x))$ and ignore the constant term. We can then define

$$\begin{aligned}\Psi(x, t) &:= e^{\frac{im}{4\hbar}(\Delta L\omega)^2 \int_0^t \cos(2\omega\tau) d\tau} \Psi_{cm}(x, t) \\ &= e^{\frac{i}{8\hbar} M \Delta L^2 \omega \sin(2\omega t)} \Psi_{cm}(x, t),\end{aligned}\tag{C.6}$$

after which the Hamiltonian takes the form 5.31.

References

- [1] J. Ignacio Cirac and Peter Zoller. “Goals and opportunities in quantum simulation”. en. In: *Nature Physics* 8.4 (Apr. 2012) (cit. on pp. 12, 58).
- [2] Immanuel Bloch, Jean Dalibard, and Sylvain Nascimbène. “Quantum simulations with ultracold quantum gases”. en. In: *Nature Physics* 8.4 (Apr. 2012) (cit. on pp. 12, 35, 58).
- [3] Christian Gross and Immanuel Bloch. “Quantum simulations with ultracold atoms in optical lattices”. en. In: *Science* 357.6355 (Sept. 2017) (cit. on pp. 12, 58).
- [4] Justin G. Bohnet et al. “Quantum spin dynamics and entanglement generation with hundreds of trapped ions”. en. In: *Science* 352.6291 (June 2016) (cit. on p. 12).
- [5] Jan Klaers et al. “Bose–Einstein condensation of photons in an optical microcavity”. en. In: *Nature* 468.7323 (Nov. 2010) (cit. on p. 12).
- [6] Walter Steurer and Daniel Sutter-Widmer. “Photonic and phononic quasicrystals”. en. In: *Journal of Physics D: Applied Physics* 40.13 (June 2007) (cit. on pp. 12, 13).
- [7] Cindy A. Regal et al. “Creation of ultracold molecules from a Fermi gas of atoms”. en. In: *Nature* 424.6944 (July 2003) (cit. on p. 12).
- [8] Roahn Wynar et al. “Molecules in a Bose-Einstein Condensate”. en. In: *Science* 287.5455 (Feb. 2000) (cit. on p. 12).
- [9] M. H. Anderson et al. “Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor”. en. In: *Science* 269.5221 (July 1995) (cit. on p. 12).
- [10] K. B. Davis et al. “Bose-Einstein Condensation in a Gas of Sodium Atoms”. en. In: *Physical Review Letters* 75.22 (Nov. 1995) (cit. on p. 12).
- [11] Steven Chu et al. “Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure”. In: *Physical Review Letters* 55.1 (July 1985) (cit. on p. 12).
- [12] B. DeMarco and D. S. Jin. “Onset of Fermi Degeneracy in a Trapped Atomic Gas”. en. In: *Science* 285.5434 (Sept. 1999) (cit. on p. 12).
- [13] Alexander D. Cronin, Jörg Schmiedmayer, and David E. Pritchard. “Optics and interferometry with atoms and molecules”. In: *Reviews of Modern Physics* 81.3 (July 2009) (cit. on p. 12).
- [14] Andrei Derevianko and Hidetoshi Katori. “Colloquium: Physics of optical lattice clocks”. In: *Reviews of Modern Physics* 83.2 (May 2011) (cit. on p. 12).

-
- [15] Peter W. Graham et al. “New Method for Gravitational Wave Detection with Atomic Sensors”. In: *Physical Review Letters* 110.17 (Apr. 2013) (cit. on p. 12).
- [16] L. Badurina et al. “AION: an atom interferometer observatory and network”. en. In: *Journal of Cosmology and Astroparticle Physics* 2020.05 (May 2020) (cit. on p. 12).
- [17] Nir Navon et al. “Emergence of a turbulent cascade in a quantum gas”. en. In: *Nature* 539.7627 (Nov. 2016) (cit. on p. 12).
- [18] B. T. Seaman et al. “Atomtronics: Ultracold-atom analogs of electronic devices”. en. In: *Physical Review A* 75.2 (Feb. 2007) (cit. on p. 12).
- [19] Jun-Ru Li et al. “A stripe phase with supersolid properties in spin–orbit-coupled Bose–Einstein condensates”. en. In: *Nature* 543.7643 (Mar. 2017) (cit. on p. 12).
- [20] Julian Léonard et al. “Supersolid formation in a quantum gas breaking a continuous translational symmetry”. en. In: *Nature* 543.7643 (Mar. 2017) (cit. on p. 12).
- [21] Jérôme Beugnon and Nir Navon. “Exploring the Kibble–Zurek mechanism with homogeneous Bose gases”. en. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 50.2 (Jan. 2017) (cit. on p. 12).
- [22] C. A. Regal and D. S. Jin. “Experimental Realization of the BCS-BEC Crossover with a Fermi Gas of Atoms”. en. In: *Advances In Atomic, Molecular, and Optical Physics*. Ed. by P. R. Berman, C. C. Lin, and E. Arimondo. Vol. 54. Academic Press, Jan. 2007 (cit. on p. 12).
- [23] C. R. Cabrera et al. “Quantum liquid droplets in a mixture of Bose-Einstein condensates”. en. In: *Science* 359.6373 (Jan. 2018) (cit. on p. 12).
- [24] Igor Ferrier-Barbut et al. “Observation of Quantum Droplets in a Strongly Dipolar Bose Gas”. In: *Physical Review Letters* 116.21 (May 2016) (cit. on p. 12).
- [25] Zoran Hadzibabic et al. “Berezinskii–Kosterlitz–Thouless crossover in a trapped atomic gas”. en. In: *Nature* 441.7097 (June 2006) (cit. on p. 12).
- [26] Jeff Steinhauer. “Observation of quantum Hawking radiation and its entanglement in an analogue black hole”. en. In: *Nature Physics* 12.10 (Oct. 2016) (cit. on p. 12).
- [27] J.-C. Jaskula et al. “Acoustic Analog to the Dynamical Casimir Effect in a Bose-Einstein Condensate”. In: *Physical Review Letters* 109.22 (Nov. 2012) (cit. on p. 12).
- [28] Chen-Lung Hung, Victor Gurarie, and Cheng Chin. “From Cosmology to Cold Atoms: Observation of Sakharov Oscillations in a Quenched Atomic Superfluid”. en. In: *Science* 341.6151 (Sept. 2013) (cit. on p. 12).
- [29] Jacob F. Sherson et al. “Single-atom-resolved fluorescence imaging of an atomic Mott insulator”. en. In: *Nature* 467.7311 (Sept. 2010) (cit. on p. 13).
- [30] Waseem S. Bakr et al. “A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice”. en. In: *Nature* 462.7269 (Nov. 2009) (cit. on p. 13).

-
- [31] Frederik Görg et al. “Enhancement and sign change of magnetic correlations in a driven quantum many-body system”. en. In: *Nature* 553.7689 (Jan. 2018) (cit. on p. 13).
- [32] Ulrich Schneider et al. “Fermionic transport and out-of-equilibrium dynamics in a homogeneous Hubbard model with ultracold atoms”. en. In: *Nature Physics* 8.3 (Mar. 2012) (cit. on pp. 13, 117, 122).
- [33] M. Greiner et al. “Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms”. In: *Nature* 415 (2002) (cit. on p. 13).
- [34] Giacomo Roati et al. “Anderson localization of a non-interacting Bose–Einstein condensate”. en. In: *Nature* 453.7197 (June 2008) (cit. on pp. 13, 45).
- [35] Juliette Billy et al. “Direct observation of Anderson localization of matter waves in a controlled disorder”. en. In: *Nature* 453.7197 (June 2008) (cit. on p. 13).
- [36] Donald H. White et al. “Observation of two-dimensional Anderson localisation of ultracold atoms”. en. In: *Nature Communications* 11.1 (Oct. 2020) (cit. on p. 13).
- [37] Michael Schreiber et al. “Observation of many-body localization of interacting fermions in a quasirandom optical lattice”. en. In: *Science* 349.6250 (Aug. 2015) (cit. on pp. 13, 45, 118).
- [38] Jae-yoon Choi et al. “Exploring the many-body localization transition in two dimensions”. en. In: *Science* 352.6293 (June 2016) (cit. on pp. 13, 37, 45, 117).
- [39] Pranjal Bordia et al. “Probing Slow Relaxation and Many-Body Localization in Two-Dimensional Quasiperiodic Systems”. In: *Physical Review X* 7.4 (Nov. 2017) (cit. on pp. 13, 45, 118).
- [40] L. Fallani et al. “Ultracold Atoms in a Disordered Crystal of Light: Towards a Bose Glass”. en. In: *Physical Review Letters* 98.13 (Mar. 2007) (cit. on pp. 13, 119).
- [41] Chiara D’Errico et al. “Observation of a Disordered Bosonic Insulator from Weak to Strong Interactions”. In: *Physical Review Letters* 113.9 (Aug. 2014) (cit. on pp. 13, 119, 121).
- [42] M. White et al. “Strongly Interacting Bosons in a Disordered Optical Lattice”. In: *Physical Review Letters* 102.5 (Feb. 2009) (cit. on p. 13).
- [43] A. Celi et al. “Synthetic Gauge Fields in Synthetic Dimensions”. In: *Physical Review Letters* 112.4 (Jan. 2014) (cit. on p. 13).
- [44] M. Mancini et al. “Observation of chiral edge states with neutral fermions in synthetic Hall ribbons”. en. In: *Science* 349.6255 (Sept. 2015) (cit. on p. 13).
- [45] M. Aidelsburger et al. “Realization of the Hofstadter Hamiltonian with Ultracold Atoms in Optical Lattices”. In: *Physical Review Letters* 111.18 (Oct. 2013) (cit. on p. 13).
- [46] Gregor Jotzu et al. “Experimental realization of the topological Haldane model with ultracold fermions”. en. In: *Nature* 515.7526 (Nov. 2014) (cit. on pp. 13, 123).

-
- [47] Dov Levine and Paul Joseph Steinhardt. “Quasicrystals: A New Class of Ordered Structures”. In: *Physical Review Letters* 53.26 (Dec. 1984) (cit. on pp. 13, 39).
- [48] W. Steurer. “Quasicrystals: What do we know? What do we want to know? What can we know?” en. In: *Acta Crystallographica Section A: Foundations and Advances* 74.1 (Jan. 2018) (cit. on pp. 13, 39).
- [49] Roger Penrose. “The role of aesthetics in pure and applied mathematical research”. en. In: (Jan. 1974) (cit. on pp. 13, 40).
- [50] D. Shechtman et al. “Metallic Phase with Long-Range Orientational Order and No Translational Symmetry”. In: *Physical Review Letters* 53.20 (1984) (cit. on pp. 13, 40).
- [51] Yaacov E. Kraus et al. “Topological States and Adiabatic Pumping in Quasicrystals”. In: *Physical Review Letters* 109.10 (Sept. 2012) (cit. on pp. 13, 39).
- [52] Yaacov E. Kraus, Zohar Ringel, and Oded Zilberberg. “Four-Dimensional Quantum Hall Effect in a Two-Dimensional Quasicrystal”. In: *Physical Review Letters* 111.22 (Nov. 2013) (cit. on pp. 13, 39, 123).
- [53] Jules Mikhael et al. “Archimedean-like tiling on decagonal quasicrystalline surfaces”. en. In: *Nature* 454.7203 (July 2008) (cit. on p. 13).
- [54] Jules Mikhael et al. “Proliferation of anomalous symmetries in colloidal monolayers subjected to quasiperiodic light fields”. en. In: *Proceedings of the National Academy of Sciences* 107.16 (Apr. 2010) (cit. on p. 13).
- [55] L. Guidoni et al. “Atomic diffusion in an optical quasicrystal with five-fold symmetry”. In: *Physical Review A* 60.6 (Dec. 1999) (cit. on pp. 13, 58).
- [56] L. Guidoni et al. “Quasiperiodic Optical Lattices”. In: *Physical Review Letters* 79.18 (Nov. 1997) (cit. on pp. 13, 58).
- [57] Matteo Sbroscia et al. “Observing localisation in a 2D quasicrystalline optical lattice”. en. In: *arXiv:2001.10912 [cond-mat, physics:quant-ph]* (Jan. 2020) (cit. on pp. 13, 14, 44, 59, 66, 69–73, 105).
- [58] Konrad Viebahn et al. “Matter-Wave Diffraction from a Quasicrystalline Optical Lattice”. In: *Physical Review Letters* 122.11 (Mar. 2019) (cit. on pp. 13, 14, 63, 64, 66, 67).
- [59] Ennio Arimondo et al. “Chapter 10 - Kilohertz-Driven Bose–Einstein Condensates in Optical Lattices”. en. In: *Advances In Atomic, Molecular, and Optical Physics*. Ed. by Paul Berman, Ennio Arimondo, and Chun Lin. Vol. 61. *Advances in Atomic, Molecular, and Optical Physics*. Academic Press, July 2012 (cit. on p. 13).
- [60] André Eckardt and Egidijus Anisimovas. “High-frequency approximation for periodically driven quantum systems from a Floquet-space perspective”. en. In: *New Journal of Physics* 17.9 (Sept. 2015) (cit. on p. 13).
- [61] Michael Messer et al. “Floquet Dynamics in Driven Fermi-Hubbard Systems”. In: *Physical Review Letters* 121.23 (Dec. 2018) (cit. on p. 13).
- [62] N. Goldman and J. Dalibard. “Periodically Driven Quantum Systems: Effective Hamiltonians and Engineered Gauge Fields”. en. In: *Physical Review X* 4.3 (Aug. 2014) (cit. on pp. 13, 92).

- [63] Alessandro Zenesini et al. “Coherent Control of Dressed Matter Waves”. en. In: *Physical Review Letters* 102.10 (Mar. 2009) (cit. on pp. 13, 91, 99, 103).
- [64] H. Lignier et al. “Dynamical Control of Matter-Wave Tunneling in Periodic Potentials”. en. In: *Physical Review Letters* 99.22 (Nov. 2007) (cit. on pp. 13, 91, 99).
- [65] Colin V. Parker, Li-Chung Ha, and Cheng Chin. “Direct observation of effective ferromagnetic domains of cold atoms in a shaken optical lattice”. en. In: *Nature Physics* 9.12 (Dec. 2013) (cit. on pp. 13, 100).
- [66] Bastien Lapiere et al. “Emergent Black Hole Dynamics in Critical Floquet Systems”. en. In: *arXiv:1909.08618 [cond-mat, physics:hep-th]* (Sept. 2019) (cit. on p. 13).
- [67] Martin Holthaus. “Floquet engineering with quasienergy bands of periodically driven optical lattices”. en. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 49.1 (Jan. 2016) (cit. on pp. 13, 33, 92, 95, 97–99).
- [68] Christoph Sträter and André Eckardt. “Orbital-driven melting of a bosonic Mott insulator in a shaken optical lattice”. In: *Physical Review A* 91.5 (May 2015) (cit. on pp. 13, 14, 102).
- [69] Matteo Sbroscia. *Between order and disorder: ultracold atoms in a quasicrystalline optical lattice (University of Cambridge PhD thesis)*. Tech. rep. 2019 (cit. on pp. 14, 45, 46, 56, 57, 66, 115).
- [70] Konrad Viebahn. *Quasicrystalline optical lattices for ultracold atoms (University of Cambridge PhD thesis)*. Tech. rep. 2018 (cit. on pp. 14, 46, 56, 57, 66).
- [71] Mário J. de Oliveira. *Equilibrium Thermodynamics*. Graduate Texts in Physics. Springer, 2013 (cit. on pp. 15, 18).
- [72] Wilhelm Zwerger. *The BCS-BEC Crossover and the Unitary Fermi Gas*. Springer, 2012 (cit. on p. 16).
- [73] Claudio Castelnovo. *University of Cambridge Part III Physics lecture course on Phase Transitions and Collective Phenomena*. 2020 (cit. on pp. 17, 20).
- [74] Katherine Blundell and Stephen Blundell. *Concepts in Thermal Physics*. Oxford University Press (cit. on p. 18).
- [75] T. W. B. Kibble. “Topology of Cosmic Domains and Strings”. en. In: *J.Phys.A* 9 (1976) (cit. on p. 19).
- [76] W.H. Zurek. “Cosmic strings in laboratory superfluids and the topological remnants of other phase transitions”. In: *Acta Physica Polonica B* 24.7 (1993) (cit. on p. 19).
- [77] Han Wang, Luigi Delle Site, and Pingwen Zhang. “On the existence of a third-order phase transition beyond the Andrews critical point: A molecular dynamics study”. en. In: *The Journal of Chemical Physics* 135.22 (Dec. 2011) (cit. on p. 20).
- [78] Tian Ma and Shouhong Wang. “Third-order gas-liquid phase transition and the nature of Andrews critical point”. en. In: *AIP Advances* 1.4 (Dec. 2011) (cit. on p. 20).

-
- [79] E. C. Ekuma, G. C. Asomba, and C. M. I. Okoye. “Thermodynamics of third order phase transition: A solution to the Euler–Lagrange equations”. en. In: *Physica B: Condensed Matter* 405.9 (May 2010) (cit. on p. 20).
- [80] Matthias Vojta. “Quantum phase transitions”. en. In: () (cit. on pp. 23, 26, 27).
- [81] Dr. Sven Horstmann / CC BY-SA (<https://creativecommons.org/licenses/by-sa/3.0>). “Theoretische und experimentelle Untersuchungen zum Hochdruckphasengleichgewichtsverhalten fluider Stoffgemische für die Erweiterung der PSRK-Gruppenbeitragszustandsgleichung”. PhD thesis. Universität Oldenburg, 2000 (cit. on p. 23).
- [82] Amar C. Vutha. “A simple approach to the Landau–Zener formula”. en. In: *European Journal of Physics* 31.2 (Feb. 2010) (cit. on pp. 26, 27).
- [83] Wojciech H. Zurek, Uwe Dorner, and Peter Zoller. “Dynamics of a Quantum Phase Transition”. In: *Physical Review Letters* 95.10 (Sept. 2005) (cit. on p. 26).
- [84] Subir Sachdev. *Quantum Phase Transitions*. Cambridge University Press, 1999 (cit. on pp. 26, 83, 86).
- [85] Lih-King Lim, C. Morais Smith, and Andreas Hemmerich. “Staggered-Vortex Superfluid of Ultracold Bosons in an Optical Lattice”. In: *Physical Review Letters* 100.13 (Apr. 2008) (cit. on p. 27).
- [86] J. Struck et al. “Engineering Ising-XY spin-models in a triangular lattice using tunable artificial gauge fields”. en. In: *Nature Physics* 9.11 (Nov. 2013) (cit. on p. 27).
- [87] L.-Y. Qiu et al. “Observation of generalized Kibble-Zurek mechanism across a first-order quantum phase transition in a spinor condensate”. en. In: *Science Advances* 6.21 (May 2020) (cit. on p. 27).
- [88] Kenji Kamide et al. “First-order superfluid-Mott-insulator transition for quantum optical switching in cavity QED arrays with two cavity modes”. en. In: *Physical Review A* 87.5 (May 2013) (cit. on p. 27).
- [89] Markus Greiner et al. “Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms”. en. In: *Nature* 415.6867 (Jan. 2002) (cit. on pp. 27, 80, 83, 109).
- [90] Simon Braun et al. “Emergence of coherence and the dynamics of quantum phase transitions”. en. In: *Proceedings of the National Academy of Sciences* (Mar. 2015) (cit. on pp. 28, 75, 82, 83).
- [91] Emmy Noether and M. A. Tavel. “Invariant Variation Problems”. In: *arXiv:physics/0503066* (May 2018) (cit. on p. 29).
- [92] Stephen Gasiorowicz. *Quantum Physics, 3rd Ed.* Wiley, 2003 (cit. on p. 29).
- [93] John Singleton. *Band Theory and Electronic Properties of Solids*. Oxford Master Series in Physics. Oxford, New York: Oxford University Press, Aug. 2001 (cit. on pp. 30, 31).
- [94] J. Feldmann et al. “Optical investigation of Bloch oscillations in a semiconductor superlattice”. en. In: *Physical Review B* 46.11 (Sept. 1992) (cit. on p. 35).
- [95] Maxime Ben Dahan et al. “Bloch Oscillations of Atoms in an Optical Potential”. en. In: *PHYSICAL REVIEW LETTERS* 76.24 (1996) (cit. on p. 35).

-
- [96] P. W. Anderson. “Absence of Diffusion in Certain Random Lattices”. In: *Physical Review* 109.5 (Mar. 1958) (cit. on p. 38).
- [97] E. Abrahams et al. “Scaling Theory of Localization: Absence of Quantum Diffusion in Two Dimensions”. In: *Physical Review Letters* 42.10 (Mar. 1979) (cit. on p. 38).
- [98] R. C. Kuhn et al. “Localization of Matter Waves in Two-Dimensional Disordered Optical Potentials”. In: *Physical Review Letters* 95.25 (Dec. 2005) (cit. on p. 38).
- [99] Rahul Nandkishore and David A. Huse. “Many-Body Localization and Thermalization in Quantum Statistical Mechanics”. en. In: *Annual Review of Condensed Matter Physics* 6.1 (Mar. 2015) (cit. on p. 38).
- [100] Vadim Oganesyan and David A. Huse. “Localization of interacting fermions at high temperature”. In: *Physical Review B* 75.15 (Apr. 2007) (cit. on p. 38).
- [101] James M. Hickey, Sam Genway, and Juan P. Garrahan. “Signatures of many-body localisation in a system without disorder and the relation to a glass transition”. en. In: *Journal of Statistical Mechanics: Theory and Experiment* 2016.5 (May 2016) (cit. on p. 39).
- [102] Dov Levine et al. “Elasticity and Dislocations in Pentagonal and Icosahedral Quasicrystals”. en. In: *Physical Review Letters* 54.14 (Apr. 1985) (cit. on pp. 39, 41).
- [103] Nicolas Macé, Anuradha Jagannathan, and Michel Duneau. “Quantum Simulation of a 2D Quasicrystal with Cold Atoms”. en. In: *Crystals* 6.10 (Oct. 2016) (cit. on pp. 39, 41).
- [104] F. P. M. Beenker. “Algebraic theory of non-periodic tilings of the plane by two simple building blocks : a square and a rhombus”. English. In: (1982) (cit. on p. 40).
- [105] Michel Duneau and André Katz. “Quasiperiodic Patterns”. In: *Physical Review Letters* 54.25 (June 1985) (cit. on pp. 41, 62).
- [106] Keiichi Edagawa, Kunio Suzuki, and Shin Takeuchi. “High Resolution Transmission Electron Microscopy Observation of Thermally Fluctuating Phasons in Decagonal Al-Cu-Co”. In: *Physical Review Letters* 85.8 (Aug. 2000) (cit. on p. 43).
- [107] Shankari V. Rajagopal et al. “Phasonic Spectroscopy of a Quantum Gas in a Quasicrystalline Lattice”. In: *Physical Review Letters* 123.22 (Nov. 2019) (cit. on p. 43).
- [108] Justus A. Kromer et al. “What Phasons Look Like: Particle Trajectories in a Quasicrystalline Potential”. In: *Physical Review Letters* 108.21 (May 2012) (cit. on p. 43).
- [109] Attila Szabó and Ulrich Schneider. “Mixed spectra and partially extended states in a two-dimensional quasiperiodic model”. en. In: *arXiv:1909.02048 [cond-mat, physics:physics, physics:quant-ph]* (Sept. 2019) (cit. on p. 45).
- [110] Attila Szabó and Ulrich Schneider. “Non-power-law universality in one-dimensional quasicrystals”. In: *Physical Review B* 98.13 (Oct. 2018) (cit. on p. 45).

-
- [111] C. J. Foot. *Atomic Physics*. 2005 (cit. on pp. 46, 48, 56).
- [112] C. J. Pethick and H. Smith. *Bose-Einstein Condensation in Dilute Gases*. 2008 (cit. on pp. 49, 56, 73, 74).
- [113] James F. Annett. *Superconductivity, Superfluids and Condensates*. Oxford University Press, 2004 (cit. on pp. 49, 56, 78).
- [114] C. Chin et al. “Feshbach Resonances in Ultracold Gases”. In: *Rev Mod Phys* 82.2 (2010) (cit. on p. 55).
- [115] Ulrich Schneider. *University of Cambridge Part III Physics lecture course on Quantum Simulation*. 2020 (cit. on pp. 61, 80, 83).
- [116] Jonathan Mortlock. *Pulsed Optical Lattices (University of Cambridge Part III report)*. Tech. rep. 2019 (cit. on pp. 68, 69).
- [117] J Hecker Denschlag et al. “A Bose-Einstein condensate in an optical lattice”. en. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 35.14 (July 2002) (cit. on pp. 69, 105).
- [118] PW Anderson. *Basic Notions of Condensed Matter Physics*. 1984 (cit. on p. 78).
- [119] Y. J. Uemura et al. “Basic similarities among cuprate, bismuthate, organic, Chevrel-phase, and heavy-fermion superconductors shown by penetration-depth measurements”. In: *Physical Review Letters* 66.20 (May 1991) (cit. on p. 79).
- [120] O. Simard et al. “Superfluid stiffness in cuprates: Effect of Mott transition and phase competition”. en. In: *Physical Review B* 100.9 (Sept. 2019) (cit. on p. 79).
- [121] S. E. Krüger et al. “Direct calculation of the spin stiffness of the spin-1/2 Heisenberg antiferromagnet on square, triangular, and cubic lattices using the coupled-cluster method”. In: *Physical Review B* 73.9 (Mar. 2006) (cit. on p. 79).
- [122] Robert Jördens et al. “A Mott insulator of fermionic atoms in an optical lattice”. en. In: *Nature* 455.7210 (Sept. 2008) (cit. on p. 80).
- [123] Fabrice Gerbier et al. “Phase Coherence of an Atomic Mott Insulator”. In: *Physical Review Letters* 95.5 (July 2005) (cit. on p. 82).
- [124] Immanuel Bloch, Jean Dalibard, and Wilhelm Zwerger. “Many-body physics with ultracold gases”. In: *Reviews of Modern Physics* 80.3 (July 2008) (cit. on p. 82).
- [125] B. DeMarco et al. “Structure and stability of Mott-insulator shells of bosons trapped in an optical lattice”. In: *Physical Review A* 71.6 (June 2005) (cit. on p. 82).
- [126] Nathan Gemelke et al. “In situ observation of incompressible Mott-insulating domains in ultracold atomic gases”. en. In: *Nature* 460.7258 (Aug. 2009) (cit. on pp. 82, 119, 121).
- [127] D. Jaksch et al. “Cold Bosonic Atoms in Optical Lattices”. In: *Physical Review Letters* 81.15 (Oct. 1998) (cit. on p. 83).
- [128] Austen Lamacraft. *University of Cambridge Part III Physics lecture course on Theoretical Physics II: Topics in Quantum Theory*. 2019 (cit. on p. 89).
- [129] Michael Victor Berry. “Quantal phase factors accompanying adiabatic changes”. In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 392.1802 (Mar. 1984) (cit. on p. 89).

-
- [130] Eliahu Cohen et al. “Geometric phase from Aharonov–Bohm to Pancharatnam–Berry and beyond”. en. In: *Nature Reviews Physics* 1.7 (July 2019) (cit. on p. 89).
- [131] André Eckardt, Christoph Weiss, and Martin Holthaus. “Superfluid-Insulator Transition in a Periodically Driven Optical Lattice”. en. In: *Physical Review Letters* 95.26 (Dec. 2005) (cit. on pp. 91, 103).
- [132] André Eckardt. “Colloquium: Atomic quantum gases in periodically driven optical lattices”. en. In: *Reviews of Modern Physics* 89.1 (Mar. 2017) (cit. on p. 92).
- [133] J. Howland. “Scattering theory for Hamiltonians periodic in time”. In: *arXiv:1212.2931 [math]* (Dec. 2012) (cit. on p. 94).
- [134] Hideo Sambe. “Steady States and Quasienergies of a Quantum-Mechanical System in an Oscillating Field”. In: *Physical Review A* 7.6 (June 1973) (cit. on p. 94).
- [135] Martin Reitter et al. “Interaction Dependent Heating and Atom Loss in a Periodically Driven Optical Lattice”. en. In: *Physical Review Letters* 119.20 (Nov. 2017) (cit. on p. 98).
- [136] *The pioneering work that led to the DMD - DLP[®] technology - Technical articles - TI E2E support forums* (cit. on p. 118).
- [137] Robert Roth and Keith Burnett. “Phase diagram of bosonic atoms in two-color superlattices”. en. In: *Physical Review A* 68.2 (Aug. 2003) (cit. on p. 119).
- [138] G. Roux et al. “Quasiperiodic Bose-Hubbard model and localization in one-dimensional cold atomic gases”. In: *Physical Review A* 78.2 (Aug. 2008) (cit. on pp. 119, 120).
- [139] Hepeng Yao, Thierry Giamarchi, and Laurent Sanchez-Palencia. “Lieb-Liniger Bosons in a Shallow Quasiperiodic Potential: Bose Glass Phase and Fractal Mott Lobes”. In: *Physical Review Letters* 125.6 (Aug. 2020) (cit. on p. 119).
- [140] Ronan Gautier, Hepeng Yao, and Laurent Sanchez-Palencia. “Strongly-Interacting Bosons in a Two-Dimensional Quasicrystal Lattice”. In: *arXiv:2010.07590 [cond-mat, physics:quant-ph]* (Oct. 2020) (cit. on pp. 119, 120).
- [141] L. Pollet et al. “Absence of a Direct Superfluid to Mott Insulator Transition in Disordered Bose Systems”. In: *Physical Review Letters* 103.14 (Sept. 2009) (cit. on p. 120).
- [142] Peter B. Weichman. “Large rare patches of order in disordered boson systems”. en. In: *Physics* 2 (Sept. 2009) (cit. on p. 120).
- [143] U. Schneider et al. “Metallic and Insulating Phases of Repulsively Interacting Fermions in a 3D Optical Lattice”. en. In: *Science* 322.5907 (Dec. 2008) (cit. on p. 121).
- [144] Pranjal Bordia et al. “Periodically driving a many-body localized quantum system”. en. In: *Nature Physics* 13.5 (May 2017) (cit. on p. 122).
- [145] Andrew C. Potter, Takahiro Morimoto, and Ashvin Vishwanath. “Classification of Interacting Topological Floquet Phases in One Dimension”. In: *Physical Review X* 6.4 (Oct. 2016) (cit. on p. 122).

-
- [146] C. W. von Keyserlingk and S. L. Sondhi. “Phase structure of one-dimensional interacting Floquet systems. I. Abelian symmetry-protected topological phases”. In: *Physical Review B* 93.24 (June 2016) (cit. on p. 122).
- [147] Dominic V. Else and Chetan Nayak. “Classification of topological phases in periodically driven interacting systems”. In: *Physical Review B* 93.20 (May 2016) (cit. on p. 122).
- [148] Vedika Khemani et al. “Phase Structure of Driven Quantum Systems”. en. In: *Physical Review Letters* 116.25 (June 2016) (cit. on p. 122).
- [149] Michael Lohse et al. “Exploring 4D quantum Hall physics with a 2D topological charge pump”. en. In: *Nature* 553.7686 (Jan. 2018) (cit. on pp. 122, 123).
- [150] Alexei Kitaev. “Periodic table for topological insulators and superconductors”. In: *AIP Conference Proceedings* 1134.1 (May 2009) (cit. on pp. 122, 123).
- [151] Shinsei Ryu et al. “Topological insulators and superconductors: tenfold way and dimensional hierarchy”. en. In: *New Journal of Physics* 12.6 (June 2010) (cit. on pp. 122, 123).
- [152] Max McGinley and Nigel R. Cooper. “Classification of topological insulators and superconductors out of equilibrium”. In: *Physical Review B* 99.7 (Feb. 2019) (cit. on pp. 122, 123).
- [153] K. v. Klitzing, G. Dorda, and M. Pepper. “New Method for High-Accuracy Determination of the Fine-Structure Constant Based on Quantized Hall Resistance”. In: *Physical Review Letters* 45.6 (Aug. 1980) (cit. on p. 123).
- [154] D. J. Thouless et al. “Quantized Hall Conductance in a Two-Dimensional Periodic Potential”. In: *Physical Review Letters* 49.6 (Aug. 1982) (cit. on p. 123).
- [155] Yasuhiro Hatsugai. “Chern number and edge states in the integer quantum Hall effect”. In: *Physical Review Letters* 71.22 (Nov. 1993) (cit. on p. 123).
- [156] Dan-Wei Zhang et al. “Topological quantum matter with cold atoms”. In: *Advances in Physics* 67.4 (Oct. 2018) (cit. on p. 123).
- [157] Shou-Cheng Zhang and Jiangping Hu. “A Four-Dimensional Generalization of the Quantum Hall Effect”. en. In: *Science* 294.5543 (Oct. 2001) (cit. on p. 123).
- [158] Xiao-Liang Qi, Taylor L. Hughes, and Shou-Cheng Zhang. “Topological field theory of time-reversal invariant insulators”. In: *Physical Review B* 78.19 (Nov. 2008) (cit. on p. 123).
- [159] H. M. Price et al. “Four-Dimensional Quantum Hall Effect with Ultracold Atoms”. In: *Physical Review Letters* 115.19 (Nov. 2015) (cit. on p. 123).
- [160] Fangzhao Alex An, Eric J. Meier, and Bryce Gadway. “Direct observation of chiral currents and magnetic reflection in atomic flux lattices”. en. In: *Science Advances* 3.4 (Apr. 2017) (cit. on p. 123).
- [161] B. I. Halperin. “Quantized Hall conductance, current-carrying edge states, and the existence of extended states in a two-dimensional disordered potential”. In: *Physical Review B* 25.4 (Feb. 1982) (cit. on p. 123).
- [162] T. Patlatiuk et al. “Evolution of the quantum Hall bulk spectrum into chiral edge states”. en. In: *Nature Communications* 9.1 (Sept. 2018) (cit. on p. 123).

- [163] M. Z. Hasan and C. L. Kane. “Colloquium: Topological insulators”. In: *Reviews of Modern Physics* 82.4 (Nov. 2010) (cit. on p. 123).
- [164] Xiao-Liang Qi and Shou-Cheng Zhang. “Topological insulators and superconductors”. In: *Reviews of Modern Physics* 83.4 (Oct. 2011) (cit. on p. 123).
- [165] Stephen Spurrier and Nigel R. Cooper. “Semiclassical dynamics, Berry curvature, and spiral holonomy in optical quasicrystals”. en. In: *Physical Review A* 97.4 (Apr. 2018) (cit. on p. 123).