Lecture 19: Phases and excitations of the Bose-Hubbard model

In this and the next few lectures, we will study a simple lattice model that will serve as a concrete example within which we will explore the physics of systems of bosons at very low temperatures. In this lecture, we will first develop a very simple and intuitive description of the various $T = 0$ phases of the system and the quantum phase transitions separating them. This will serve as the foundation on which we will build in the next lecture using path integral methods developed in the previous lectures to sketch a derivation of the effective field theory that describes the low-energy, long wavelength properties of the system in the vicinity of these quantum phase transitions. Next, we will introduce external fields into the action of this effective field theory and use our previously developed linear response theory formulae to understand what controls the conductivity of the system by relating it to certain correlators of the effective field theory via a 'Kubo formula'. We will also discuss the connection of the conductivity to the superfluid stiffness.

With this background, consider the Hamiltonian

$$\mathcal{H} \;=\; -w\sum_{\langle ij\rangle}\left(c_i^\dagger c_j + c_j^\dagger c_i\right) - \mu\sum_i n_i + \frac{U}{2}\sum_i n_i(n_i-1)\;, \tag{1}$$

where c_i and $c_i^\dagger$ are the usual boson annihilation and creation operators that obey the commutation relations

$$[c_i\,,\,c_j^\dagger] = \delta_{ij}\;, \tag{2}$$

μ is the chemical potential ($n_i = c_i^\dagger c_i$ is the number operator at site i), $w \geq 0$ denotes the strength of the nearest neighbour hopping term, and $U > 0$ is the magnitude of the on-site repulsion between the bosons. This Hamiltonian defines our simplified Bose-Hubbard model.

We may think of this as the effective Hamiltonian of a regular array of superconducting islands that are only weakly coupled to each other. The bosonic operators then represent the Cooper pairs (pairs of opposite spin electrons bound to each other to form a bosonic composite) of the system, while the hopping term reflects the tunneling of these cooper pairs between neighbouring superconducting islands. Of course, in such a description, we are assuming that it is legitimate to ignore long-range Coulomb interactions as well as processes where a Cooper pair decays into the underlying fermions of the original problem.

Another physical situation that is modeled by this bosonic Hubbard model is a system of ultracold (at nano-Kelvin temperature) neutral atoms trapped in a small region of space by the standing-wave fields produced by counter-propogating laser beams. The effective potential produced by these lasers can be modeled quite well by an overall slowly-varying harmonic confining potential with a periodic lattice potential superposed over this—in typical situations, hundreds of periods of the oscillatory part of the optical potential fit into the overall harmonic envelope, and it is therefore a reasonable first approximation to replace the harmonic confinement by more conventional particle-in-a-box type boundary conditions, and focus on the periodic potential by writing down a tight-binding type description of the motion of particles in this potential (this is similar to the tight-binding model we write down in our usual solid state physics courses for discussing electronic band structure of solids in simple terms).

Let us begin by considering the limiting case $w = 0$. In this limit Eqn. (1) reduces to a collection of decoupled single site problems. Thus, we need to only solve a single site problem to determine the ground state of the entire system. It is extremely straightforward to do this as the Hamiltonian is diagonal in the n representation. Determining the ground state is simply a matter of minimizing the energy as a function of n (allowing only integer valued solutions, of course), and this gives us a many-body ground state that is simply the product state with the same number of particles $n_0(\mu/U)$ on each site. Here, n_0 is an integer valued function equal to zero for $\mu/U < 0$, and equal to m for all μ/U in the interval $m-1 < \mu/U < m$ (where m is any positive integer). When μ/U is precisely equal to any non-negative integer m, the system is free to choose from the states $|m\rangle$ and $|m+1\rangle$ *independently* at each site, leading to a macroscopic degeneracy of 2^M, where $M = (L/a)^d$ is the number of sites in the system (L is the linear dimension of the system and a is the lattice spacing).

Thus, as we increase μ/U at $w = 0$, we sweep through a series of phases in which the number of particles at each site gets pinned at successive positive integers. We may also determine quite easily the lowest lying excited states of the system: Consider the phase in which the ground state density is pinned at some positive integer m. For $m - 1 < \mu/U < m - 1/2$, the lowest lying excitations are hole-like and form a degenerate manifold composed of states in which the occupancy of any one lattice site is decreased from m to $m - 1$.

Similarly, for $m - 1/2 < \mu/U < m - 1$, the lowest lying excitations are particle-like, and consist of states in which the occupancy of a single site is increased by 1 to $m + 1$. Note that both the particle and hole excitations are separated from the ground state by a substantial gap as long as we stay away from integer values of μ/U.

Let us now see what happens to these phases as we turn on a small hopping amplitude w. To begin with, consider the degenerate points $\mu/U = m$. We can handle the effects of small non-zero w by working to lowest order in degenerate perturbation theory. At this level, we merely have to diagonalize the restriction of the hopping term to the degenerate manifold of lowest energy states. Clearly, this problem maps on to a problem of hard-core bosons hopping around at zero chemical potential. The ground state of this problem is certainly a Bose-condensate, familiar from our elementary solid-state physics course. Such a Bose-condensed state is generically also a superfluid, in the sense that it supports the flow of current without any hindrance or dissipation (we will have much more to say on this in the next few lectures, so do not worry if the connection is not immediately obvious). Thus our original system immediately turns superfluid upon turning on w for $\mu/U = m$. The situation is markedly different away from these points of degeneracy. In this case, the unperturbed ground state is separated by a gap E_p (E_h) from the lowest lying particle-like (hole-like) excitations. Since the number operator commutes with the Hamiltonian (and is therefore a conserved quantity), we may continue to label the exact eigenstates of the system in the presence of the hopping term by the corresponding value of the total number of particles in the state. As we turn on w, the energies of the states in the $m(L/a)^d$ particle sector and the sectors with one more or less particle will all evolve smoothly with increasing w. The presence of the energy gap then ensures that the ground state remains strictly within the sector with m particles per site for small w. In other words, *the density remains pinned at* m *as we turn on* w.

Additional insight into the nature of the ground state for small w may be obtained from a perturbative calculation of the amplitude for a boson initially at site i to hop to site j. This may be estimated as follows: When the hopping term acts on the unperturbed ground state, it creates a particle-hole pair in which the particle and the hole are at adjacent sites. As a result, the the amplitude for a boson to have moved by r sites is non-zero only at the r^{th} order of perturbation theory. Each successive order of perturbation theory introduces an additional power of the ratio w/E_{ph}, where $E_{ph} = E_p + E_h$ is

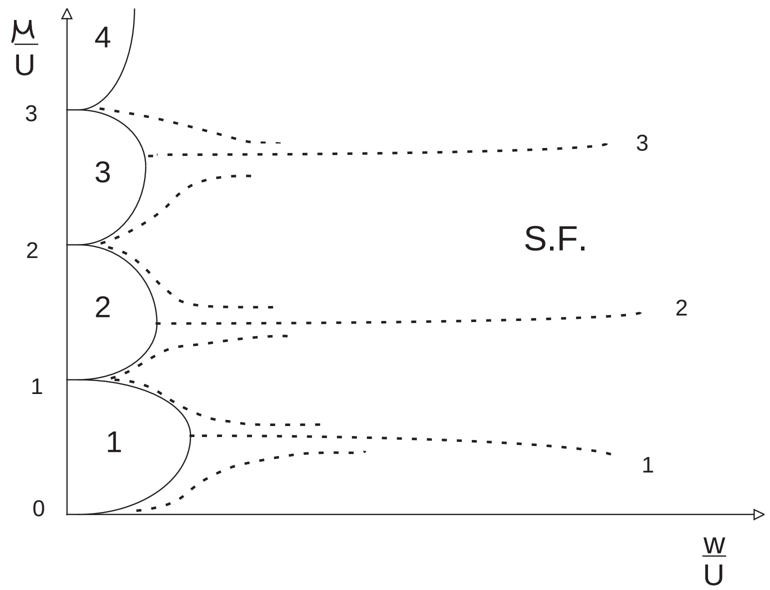

Figure 1: Schematic phase diagram of the model defined by Eqn (1). The lobes labeled by integers are Mott insulating phases with density pinned at the corresponding integer. The system is in the superfluid phase (denoted as S.F in the figure) for large values of w/U. The bold lines represent the position of the phase boundary between the Mott insulating and superfluid phases. Numerical analysis of the mean field theory described in the body of this lecture indicates that the transition is second-order in the sense that the superfluid order parameter goes continuously to zero as we cross over from the superfluid to the Mott insulator. The dotted lines are a schematic rendering of contours of constant density. The contours of constant integer density are expected to hit the corresponding Mott insulating lobe at its tip, and have horizontal tangent at the point of contact (see main text of lecture).

the relevant energy denominator in the perturbative expressions.

The amplitude that a boson has hopped r sites is therefore $\sim \exp(-r/\xi)$, where $\xi \sim [\ln(E_{ph}/w)]^{-1}$. The ground state for small w is now seen to be *insulating*, in addition to having the the density pinned at m. Thus, the phases we found at $w = 0$ extend to Mott-insulating lobes for small w. Of course, as we crank up w, the extent of these Mott-insulating lobes along the μ axis will become smaller and smaller since the gap to creating particle and hole excitations continues to decrease, until finally, at a critical value of w, the lobe pinches off and disappears completely, leaving us with a superfluid phase. A schematic phase diagram that depicts these conclusions is shown in Fig 1.

Now, we will be most interested in a situation in which the density of the system is fixed at some integer in the superfluid phase and the transition to the insulator is driven by tuning some parameter such as the strength of the hopping w (in terms of the physical interpretation we have already discussed, this could represent the effects of reducing the Cooper pair tunneling between adjacent superconducting islands, or adjusting laser beams in the cold-atom context to reduce the overlap between lowest eigenstates localized in successive minima of the periodic optical potential).

Our system will then cross the phase boundary along a contour of constant integer density. It is therefore important for us to know where such an equal density contour intersects the phase boundary. To begin with, let us assume that the contour of constant integer density hits the boundary of the corresponding Mott insulator lobe at some point slightly away from its tip. The constant density contours corresponding to slightly higher or lower densities all have to skirt around the lobe on either side, and finally hit one of the degenerate points marking the boundaries of that particular insulating phase on the $w = 0$ axis.

Now, it is easy to see that this requirement leads to a unphysical negative compressibility in the vicinity of the intersection of the integer density contour and the phase boundary. Thus, our original assumption is untenable, and the constant integer density contour has to intersect the phase boundary at the tip of the corresponding insulating lobe. Moreover, it is clear that the contour has to come in horizontally at this point. We will see in the next lecture that we can use our previously developed path integral description and completely determine the basic form of the low energy theory that describes the system at this transition between a superfluid state and a Mott insulating state at fixed integer density

Here, we conclude by outlining a simple-minded mean-field theory that can be used to back up the strong-coupling expansions and physical arguments employed in our foregoing discussion. The basic idea of this mean-field treatment is to recognize that a Bose-condensed state or a superfluid may be thought of in second-quantized language as a phase in which the creation and annihilation operators a and $a^\dagger$ themselves acquire expectation values.

That this is the case can be argued as follows: Recall that in our previous lecture on bosonic coherent states, we have indicated how the phase variable ϕ which labels a coherent state, is in a certain sense conjugate to the number variable—roughly speaking, having a well-defined phase precludes the possibility of fixing the number variable to a single value, and states with a well-defined phase are formed by superposing many different states with different numbers of particles. Having such large fluctuations in the number and well-defined phase allows current to flow in such a state without dissipation—roughly speaking, the flow of current involves changes in the local particle number, and this is greatly facilitated in states with a well-defined phase. Therefore, superfluids are characterized by having a well-defined phase variable.

With this in hand, we proceed to describe the mean-field theory: The full Bose-Hubbard model Hamiltonian is replaced by a mean-field Hamiltonian:

$$\mathcal{H}_{MF} \;=\; -wz(\psi c^\dagger + \psi^* c) - \mu n + \frac{U}{2} n(n-1) \;, \tag{3}$$

where z is the coordination number of the lattice, and ψ (ψ^*) represents the expectation value of c ($c^\dagger$) averaged over all the neighbours of a given site.

As is usual in mean-field treatments, this Hamiltonian is supplemented with an appropriate self-consistency condition, which, in this case, obviously reads:

$$\psi = \langle b \rangle_{MF} \tag{4}$$

where $\langle\rangle_{MF}$ denotes expectation values in the ground state of the mean-field Hamiltonian. Within this formulation, the system is superfluid for values of μ, U and w for which the mean-field solution has a non-zero value of ψ, while $\psi = 0$ corresponds to the Mott-insulator.

If you have access to some numerical package like Mathematica or can write a small piece of code, you will easily be able to solve this self-consistency equation for yourself and generate the phase diagram of the system, which should look something like the sketch I have drawn in the figure that preceeded this discussion.