

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high-quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

PROFESSOR: OK, good afternoon. So this week we talk about quantum gases, ultra-cold atomic gases. And sure, they're ideal Bose gases, ideal Fermi gases, but I will spend one or two minutes on each, because that's exactly solvable. It's in all the textbooks. That's simple.

Quantum gases become interesting because of interactions. And in my lecture today and my last lecture I want to introduce you to Bose gases and Fermi gases with interactions. And they both turn superfluid, and the superfluid properties are determined by the interactions between the atoms.

So the purpose of those lectures is number one, to acquaint to you with important phenomena in cold gases-- superfluidity, superfluidity in lattices, superfluidity just in free gas, in a normal gas without lattice, and superfluidity of fermions.

But at the same time, I also want to have sort of a theme for you how we deal with interactions. And that kind of theme is how theoretically we reduce unsolvable Hamiltonians to Hamiltonians which can be solved. And you will see that actually appearing in different situations with similarities, but important differences.

So just to remind you, we started out with the interacting Bose gas in a homogeneous system. We have a very general way to describe scattering two particles with initial momentum disappear and two particles with momenta appear. This is a scattering event.

Now this means we have products of four operators, which is very difficult to solve. And the Bogoliubov approximation which we discussed on Monday replaces the operator for the condensate for the zero momentum state with a Z number, saying that $N_0, N_0 + 1$ is the same-- a little bit waving your hands. But it's also you can

say the macroscopic limit when we have a photon field with many, many photons, we can use a Z number in our Hamiltonian to describe the electric field. So that's the same spirit.

So maybe I should emphasize it. This is for atoms. What you're used to do with photons for your whole life. OK.

With that we have transformed the Hamiltonian into a bilinear expression. And so at the moment your bilinear expression-- you do a Bogoliubov transformation. You simply diagonalize it by finding a new set operators where the cross term between a and b or between a of k and a of minus k disappear. And then you've solved it.

And this is what we arrived at at the end of the last lecture. With those approximations, we have diagonalized the Hamiltonian. Our Hamiltonian is now, you can say harmonic oscillator Hamiltonian, or you can see it has become a gas of non-interacting quasi particles. Each of those operators creates a quasi particle, and the quasi particle energy is U of k.

And I explained to you that as expected, the quasi particle energy is simply particles with momentum $\hbar k$ for high momenta, because in high excitation of Bose quantum it is a free particle. But a low-lying excitation is affected in a major way by all of these interactions with the Bose-Einstein condensate. And that turns the quadratic dispersion relation into a linear one, or you can say that turns the free particle into a photon or something.

So that's my review of the last class. Are there any questions before we go further?
Colin.

AUDIENCE: Does this require low density and [INAUDIBLE], or one or the other?

PROFESSOR: We come a little bit later to that, but in the end, there is a small parameter. The small parameter will be $N^{-1/3}$ -- the gas parameter-- N , the density, a cubed, the scattering length. It's usually the dimensionless combination of the two which decides whether we are in the weakly or strongly interacting limit.

There is another assumption which we have made here, but it's related to that-- as you will see a few moments-- that most of the are in the condensate. We assumed N_0 is peak. The condensate depletion-- the number of atoms which are not in the condensate, n equals 0 is small.

But let me just first show you-- I want to mix in theory and experiment how sound can be observed. I should actually say this experiment is sort of dear to my heart, because at some point it clicked to me how the new world of atomic physics connects with condensed matter physics. Soon after we had realized Bose-Einstein condensates, a famous condensed matter theorist said hey, Wolfgang, you should now observe sound. Sound is important.

Yeah, OK, but how do you observe sound? You know, use a piezo and just kick the system and create a sound wave. And I said, oh, gosh, this guy doesn't know with that if you put anything in contact with the quantum gases-- a piezo-- the gases will just stick to it. He has no idea what our system is.

But then I said wait a moment-- we have to translate. He says use a piezo. Well, I have to translate. In atomic physics how do we kick atoms? With a laser beam.

And in that moment, I had the idea that we can take a Bose-Einstein condensate, suddenly switch on a blue de-tuned repulsive laser beam. This will do exactly as a piezo-- create a density perturbation, and then the density perturbation will propagate with the speed of sound.

And that worked. That was just when Bose-Einstein condensations were very fresh. One of the first scientific experiments done-- we switch on the laser and you see color-coded in red the density perturbation which propagates out of it.

And the slope of this line is the speed of sound. And here we determined the speed of sound as a function of density. So that's how phonons-- or at least how the speed of sound and wave packets which propagate with the speed of sound can be prepared. I come back to phonons and collective excitations in a few moments. But let me first say when we have diagonalized the Hamiltonian we know everything we

want to know.

I just focused on the quasi particle energies, but you also know the ground-state energy. And actually here Colin, you see that the corrections to the ground-state energy scale with N a cubed, so this is really the small parameter in this system.

But we can also find out what is the ground state wave function. And let me use it to introduce an important concept to you called the quantum depletion. When you have two atoms in the condensator at zero momentum and you switch on the interactions, the delta function interaction couples zero momentum state to higher momentum states.

So therefore, the effect of interactions is that the condensate is not just at zero momentum. It has some probability, or some admixture, of finite momentum states. This is the ground state. This is how the Hamiltonian is diagonalized.

So from the Bogoliubov approximation, where we introduced this V and U parameter to transfer from one set of Bose operators to another set of Bose operators, those coefficients give us the population of those momentum states in the [INAUDIBLE] of the condensate.

So if I now ask, what is the condensate fraction? What is the number of atoms in the zero momentum states? It's all atoms, but those who have finite momentum. And we find again the small parameter N a cubed.

Now this quantum depletion allows me now to make a distinction between the cold atomic gas condensates and superfluid helium form. In cold atomic gas condensates, this correction is about 1%. So therefore, I can say with 99% probability, or with 99% weight, the condensate, the many-body wave function of the condensate is just the zero momentum state to the power N .

And this very complicated admixture of correlation into the ideal gas wave function is only 1% for alkali gases. But for liquid helium, the condensate fraction-- even at zero temperature-- is only 10%. So when people use neutron scattering-- it would be a long story in itself how this is done-- but when they use neutron scattering to

analyze the liquid helium at low temperature, and figure out what is the fraction of atoms which have zero momentum, they found 10%. The quantum depletion is 90%.

But that's just the difference between a quantum gas and a quantum liquid. And in the quantum liquid-- in liquid helium-- $n \lambda^3$ is on the order of unity.

OK, I mentioned this dispersion relation. These are sort of quasi particles. And I at least showed you how you can measure the slope of the quasi particle dispersion relation by propagating sound waves. But let me now tell you how we can observe quasi particles' elementary excitations directly. And this is actually simply done by light scattering.

If you scatter a photon, and the scattered photon loses energy $\hbar \omega$. And it is scattered at an angle. Therefore, it transfers momentum q . These forces can only happen if there is an elementary excitation with momentum q and energy $\hbar \omega$.

So in other words, on a photon basis, you can see photon by photon, if you scatter a photon, the photon transfers momentum and energy. The process can only happen if you can form a quasiparticle with this momentum and energy.

So since this is sort of the direct way of mapping out whether the system has the possibility to absorb momentum q and energy $\hbar \omega$, this has sort of a name. The scattering probability is called the dynamic structure factor. And the dynamic structure factor is just an integral over all the possibilities that a many-body system can absorb momentum and energy.

Now there is one nice feature which was introduced by us at MIT, and that is if you're going to measure the dynamic structure factor, often you do it with neutron scattering. You scatter neutrons or x-rays and they scatter spontaneously at an angle, and you need a detector to detect the scattered particles.

But in our case, because the gas is so dilute, the distance between atoms is on the order of the optical wavelengths. We don't want neutrons. We don't want x-rays. We want photons, because the wavelengths of photons is perfectly matched to the

properties, to momentum transfer and such we need for our system.

But photons-- we have photon lasers. So instead of in a painstakingly way analyzing the frequency and the momentum of scattered photons, we don't do the spontaneous scattering process. We do a stimulated process. We use two laser beams, and we stimulate a photon to be scattered into the other laser beam.

And by having the two laser beams at frequency difference $\Delta\omega$, we're really asking the system, are you ready to absorb $\Delta\omega$ energy? If yes, then you have a quasiparticle. So this is how we do quasiparticle spectroscopy.

And a few years later, this method was defined by the [INAUDIBLE] Institute. And what you see here is they varied the angles between the two laser beams, realized different momentum transfer, and what you see is the low dispersion, the linear dispersion relation for low momenta, and then the quadratic part at high momentum.

So this is called BEC spectroscopy. It's a variant, you can see, of Raman spectroscopy where you go from a zero quasiparticle state to a one quasiparticle state through a simulated Raman process, and this is how the dispersion relation is determined.

OK so I've so far dealt with aspects of a homogeneous Bose-Einstein condensate where, of course, if you're a homogeneous system, you formulate everything in momentum space. But now we want to deal with the situation that our condensates are finite in size. And in addition, they're usually in a harmonic oscillator potential, and therefore their density is inhomogeneous.

Let me start the discussion with the inhomogeneous Bose-Einstein condensate by showing again a picture which brings in some memories of '90, '96. It was the first time that we could see the Bose-Einstein condensate in the trap. Before people have just seen it by time of flight when it was pretty much already destroyed or it was just flying out.

But here we see the trapped Bose-Einstein condensate and we could even take multiple pictures of the same condensate. So you can see that this was the first time that a condensate was seen alive.

What you see here is actually the bimodal distribution. You see the condensate and then due to speckled pattern, the Fermi cloud looks fragmented, but this is just a speckled pattern due to the signal to noise. But here you really see now-- not in ballistic expansion as I showed you on Monday-- but really insight to the size and shape of the condensate.

OK, how do we describe it? Well, the message is pretty much exactly the same way as we did with the Bogoliubov approximation and such, but instead of in momentum space, we do it now in position space.

Actually, I have to say this week, I go through a lot of material. But what I'm trying is to give you sort of a spirited and animated overview that you really know what is important. Where is the same concept appearing again in a different way.

The details-- I've posted, actually, the original articles, references. Some of are school notes, some of them written by myself on the backside. So yes, it's a little bit different character than other lectures. I want to show you a lot of things, and put a special emphasis on the ideas.

OK, so in second quantization, we are now using field operators which create and annihilate particles at position R . This is the single particle Hamiltonian-- kinetic energy and potential energy. And our interaction term has now, again, four operators-- two interrelation operators, two creation operators, but they're now formulated in position space.

We do exactly the same as we did in the homogeneous gas. We assume the potential is short range. That means delta function. And that means we can get rid of one of the integrations in r and r prime. But we still have product of four operators.

We can formally solve the equation by writing down Heisenberg's situation for the

equation of motion for those field operators, but of course this cannot be solved. Remember products of four operators are hard. We have to reduce it to two operators.

And this is actually done by the Bogoliubov approximation again. But it's done here in the following way. Remember, in the momentum space we replaced a_0 and a_0^\dagger by the square root of a_0 .

What we do here is we say this is sort of a quantum field. And when we have a superfluid, this field operator has an average value. It's actually the macroscopic superfluid order parameter.

So we replace the operator by an average value which we assume is large, because many, many particles are in the superfluid state. And then we have fluctuations which are small. And what we will do is-- it's pretty clear-- when it comes to the fluctuations, we will neglect higher products-- products of four fluctuation operators. And we are down to C numbers and, as you will see, two operators.

We can-- and this is even more dramatic-- it wouldn't make sense in momentum space, but it does make sense in position space-- we can even do a first approximation where we completely neglect those fluctuations. And then what we have is when we simply insert a $\langle Z \rangle$ number into this Hamiltonian. We have then an equation for this number, for this function ψ .

And this is the nonlinear Schrodinger equation, or it's also called Gross-Pitaevskii equation, which is now the analogy to Schrodinger's equation, but now for the macroscopic wave function which is occupied by many, many particles. And in addition to the kinetic energy and the trapping potential in the external potential, it has a term which is proportional to the density. And this is just mean field repulsion which one atom feels exerted by the other atoms.

I should say-- just to connect it to what I've said earlier about the quantum depletion-- that you can regard the function ψ as the best approximation to an ideal Bose-Einstein condensate. So if you want to write your many-body wave

function as a function ψ to the power N , all particles are in the same quantum state, then a variational calculation would say that this function ψ should obey this equation.

If you want to read up what I've said, I've posted a recent paper, which I found very pedagogical, where they derive the Gross-Pitaevskii equation without even using an operator ψ^\dagger , without using any second quantization. They just say we have an Hamiltonian for an interaction system, and we try to write the complicated, many-body wave function as a wave function to the power N . And then you put this ansatz into your Hamiltonian and do a variation optimization which single-particle wave function ψ , if taken to the power N , minimizes the total energy of N particles, including interactions.

And the answer is this equation. So there are many ways this equation is sort of very natural.

Now you all know about single-particle physics-- kinetic energy in an external potential. Let's now learn what this density-dependent term does. Well, it's pretty clear. There is a price to be paid for density.

And while the total integrated density, of course, is given by the total number of particles. That means if you lower the density somewhere, you have to increase the density. But because this is a non-linear term, if you have an average density and you lower it here and increase it here, you have actually more repulsive energy than if the density is flat. So this term has only one goal. It wants to flatten out densities.

So therefore if we have a box potential and you know the ground state in a box is just half a period of a sine function. This is a non-interacting condensate. This would be the macroscopic wave function in a box potential. But if you now put in strong interactions, the strong interaction's going to flatten out the potential.

They flatten out the potential, and only at the very last moment-- all right, now it's time to go down because we have to meet our boundary condition. And the length scale where you eventually go down-- it has a famous name, the healing length-- is

the length scale where the kinetic energy due to the curvature of the wave function is now comparable to the interaction energy.

If the system would curve down earlier, it would cause too much repulsive energy, because the density is not kept flat. If it would curve down too late, the enormous curvature would mean a lot of kinetic energy. And this is just the best compromise between those two criteria.

So that's how you derive the healing length. The healing length is now the length scale over which the system is willing to meet boundary conditions, and not stay flat, as flat as possible. Yes.

AUDIENCE: Just curious-- in that equation, the atom has the mass of each individual--

PROFESSOR: The mass is the mass of a single atom.

AUDIENCE: Single atom, it's not-- sorry, I wasn't here last time. But the quasiparticles are the same mass [INAUDIBLE].

PROFESSOR: In a Bose-Einstein condensate they are. We haven't changed the mass. But we're not talking actually about quasiparticle. We're actually really talking about here qualitative features of the solution of the Gross-Pitaevskii equation.

And the Gross-Pitaevskii equation is sort of a single particle equation for particles of mass, of the original bare mass M . It's a macroscopic wave function, but I sometimes say it describes the wave function of a single particle where all the other particles are included at the mean field level. So therefore, it is really this, the atomic mass, and not any form of collective mass here.

Well, if you would now ask how do interactions transform the first excited state in the box potential? Well, then it's again flat this as much as possible. But then if you want to maintain the parity of the wave function, then it's only close to the zero crossing within a healing length that the system says OK, now I change sign.

So that's sort of what is inside the Gross-Pitaevskii equation. Now once we realize that, we can take it to the next level and say, well, if you neglect-- let's say we're

interested in the ground state and we neglect this boundary region where the kinetic energy becomes important-- maybe we can simplify the Gross-Pitaevskii equation by neglecting the kinetic energy.

If potential energy dominates by far, we can neglect that. And then we should get a good description which will not be valid in the wings of the wave function but in most of space. But now if you look at this equation, it's no longer a differential equation. It has no derivatives. It's just something which applies to the wave function ψ itself.

And we can simply solve that. The solution of that is that ψ^2 -- the density-- is nothing else than there is a constant minus the inverted trapping potential.

So therefore in this Thomas-Fermi approximation where we neglect the kinetic energy completely, you just take your trapping potential, turn it upside down, and then you fill it up with density until you have accommodated the number of atoms you want to accommodate. And this is, of course-- the constraint in the number of atoms is determined by the chemical potential.

Or if you have a more complicated W-shaped potential, the same construction. Flip it over, fill it up, gives you the condensate wave function, the density distribution of the condensate in this potential.

I don't want to belabor it, but coming back to the question of the small parameter, if you look at those equations, you can identify a parameter-- this is now the small parameter, or the parameter in the system which is the important dimensionless parameter. It depends on the number of atoms, and it depends now on the ratio of the scattering length-- which characterizes the interaction-- and the harmonic oscillator length.

You can say the harmonic oscillator length is the ideal wave function harmonic oscillator. So this parametrizes the importance of kinetic energy, whereas the scattering length parametrizes the importance of interactions. And all those solutions can be nicely written as what you would have in an ideal gas, and then this parameter X . So in typical experiments, N is a million. The scattering length is

smaller than the oscillator length, but N is a million, and this parameter X is usually large.

So therefore-- and I'm simply just talking about this solution-- we have the situation that the chemical potential is larger by a power of X than the ideal gas solution, which would just be the ground state with its zero point energy in the harmonic oscillator potential, or the width-- the size-- of the wave function is larger than the ground state of the harmonic oscillator, but only with an exponent, which is one-fifth.

Well, we can see that. These are now, again, somewhat improved pictures of condensates inside the trap. For the expert, the previous picture was dark-ground imaging. This is now phase-contrast imaging.

And if you take a profile, we clearly see the condensate wave function and we see the thermal wings. When we look at the size of the condensate wave function, you realize what I just meant-- that the ground-state wave function, the harmonic oscillator length is 7 micron in the axial direction, but here you have 300 micron. So this condensate is completely dominated by interactions. And it fulfils very nice, and pretty much the whole shape, except maybe some details in the wings, are quantitatively described by the simple approximation I have explained to you.

So just a little bit show and tell. We have this Gross-Pitaevskii equation. The Gross-Pitaevskii equation has, as Schrodinger's equation, a time independent form to get the ground state. It has also a time dependent form.

You simply replace the energy by the derivative of the wave function. This is the time-dependent form. Everything is very simple, and you can do many-body physics, but on your computer, you pretty much look for single-particle wave function and take them to the power N .

So some areas where the Gross-Pitaevskii equation has quantitatively explained experiments-- one is the expansion of a Bose-Einstein condensate. It's this famous situation when you have an elongated condensate and let it expand. It expands faster in the radial direction.

One simple argument is the pressure of the mean field is larger and leads to faster acceleration. So therefore, if you have a cigar-shaped condensate and release it, it turns into a disk. It inverts the aspect ratio, going from a cigar shape to a disk shape.

And that has been beautifully and quantitatively described by the Gross-Pitaevskii equation. Here we have measured the interaction energy as a number of condensed atoms. And I mentioned to you that this X parameter comes often with power $1/5$ and $2/5$. And this here is a fit to the power $2/5$, beautifully confirming the theory.

So the Gross-Pitaevskii equation was invented in 1962, about 50 years ago by Gross and Pitaevskii to describe vortices. Actually, Lev Pitaevskii is still alive, going strong, and publishing papers. One of his latest predictions was solitons in Fermi gases, and I know in Professor Zwierlein's group, one of his lab just looked at the same physics, and compared to the theory of Lev Pitaevskii. And this year we are celebrating-- I forgot, the 85th or 90th birthday of him.

AUDIENCE: [INAUDIBLE]

PROFESSOR: 90?

AUDIENCE: There's a poster.

PROFESSOR: I know there's a poster next to my door, but I forgot which anniversary. So I mean, he's an legend. But he's still walking. He's still doing science. So if you meet him, you go back to 50 years in history. Anyway, it is this Lev Pitaevskii, and he invented his equations to describe vortices.

So anyway so we have this nonlinear Schrodinger equation. Let me just show you what vortices are and how they are formed. Vortices come-- if you solve the Gross-Pitaevskii equation-- if you ever any quantum fluid or quantum gas, and you add angular momentum, the angular momentum cannot lead to rigid body rotation. This would violate the fact that the velocity field has to be irrotational to make sure that the phase of the wave function is well defined.

So if you rotate the system, it can absorb angular momentum only by forming vortices. And vortices are singular points of the wave function-- singular points where the density is 0. In other words-- I don't have time to be exhaustive here-- but when something rotates, there is a dynamic phase you can say, because there are matter waves going in circles. And the integral of the phase has to be an integer number of 2π , otherwise you would not have a well-defined phase of the wave function.

And if you now say you make a circle, which is 2π . You make the circle smaller, smaller, smaller. In the middle of the circle, you go around. The wave function changes by 2π , but what should the wave function do on one point? Which number between 0 and 2π should the wave function peak?

Well, the wave function says I can't peak. I just go to 0. And a 0 value has no phase and I'm fine. That's exactly what the wave function tells you. So therefore if you want to describe vortices, you want to now describe the Gross-Pitaevskii equation, but with the boundary condition that the wave function goes to 0. And if you want to describe one vortex, you want the wave function to go 0 in the center of the cloud.

And so you make a corresponding Ansatz. You allow the phase to wrap around by 2π , when the angle ϕ is varied. And then when you solve it, you have, of course, put in that there should be a 0 of the density at the center. And it's now nice for me to-- there's an important review paper which describes all that and more-- but it's nice to show you now the two extreme cases of the ideal condensate without interaction, and the strong interacting condensate.

The ideal condensate-- well, we are in the first excited state of the harmonic oscillator potential. And this is the dashed line. And of course, the size of the whole, the 0, is pretty much on the order of the oscillator length, because there is no other length scale in the ideal harmonic oscillator.

But if you have the interacting system, remember what interactions are doing. They want to keep the density as constant as possible. Well, we are not in a box potential.

We're in an inverted parabola potential.

And remember, our Thomas-Fermi solution, which neglects kinetic energy, is the inverted parabola up to here. And then the tail is when we can no longer neglect the kinetic energy. But if we now say, OK, fine, but now in addition, we want a vortex in the center, the wave function out there says no.

We do what minimizes kinetic energy, and what minimizes repulsive energy. We follow the inverted parabola, and only at the very last moment-- on the scale of the healing length, the system meets the required boundary condition that the density goes to 0.

Anyway, with this qualitative understanding, you can get a lot out of those. You can immediately understand the salient feature of the solution.

Question?

AUDIENCE: I know we're not talking details, but is there a simple, maybe clear reason why when you spin a Bose-Einstein condensate you get many vortices, but when you spin your coffee in your cup, you just get one big one?

PROFESSOR: Yes. So the question is, if I have a wrap-around of 4π in phase, whether the system should have a doubly charged vortex, or two singly charged vortices? What is the difference in energy?

Well, what happens is doubly charged vortices are unstable. My group, at some point, were the first to create doubly charged vortices, but we saw that they immediately decayed.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Pardon?

AUDIENCE: [INAUDIBLE]?

PROFESSOR: Initially, we couldn't observe. We just saw that it was unstable. But it's a decay into

two vortices.

But it's easy to understand. If you put two vortices on top of each other, what is the energy of the system compared to one vortex? Well, the energy of a vortex is the rotational field around it, and if you superimpose two vortices, the velocity around the doubly charged vortices is twice the velocity of one vortex. And therefore, the kinetic energy is four times.

If you have two vortices which are far separated that it has its own velocity field. It has its own velocity field here. And when the two velocity fields come together, the velocity is already so low that you don't have to consider that, then those two vortices have an energy which is two times the energy of a single vortex.

So therefore, when the two vortices, they start out with four times the energy of a single vortex, and when they dissociate and repel each other, they have shed half of their energy. So this argument tells you immediately that vortices are in effect repulsive. And therefore, any multiple-charged vortex will spontaneously decay.

It's also this net repulsion between vortices which makes the vortices arrange in a regular lattice. The regular lattice which you saw before is Nature's answer to how can we minimize the energy of all those vortices? And the idea is let's keep the average distance between them as large as possible, and the answer is a hexagonal lattice. Yeah.

AUDIENCE:

So a few slides back, when you write the G-P equation, when you plotted sort of for the box potential, the ground state and the excited state comparing the G-P equation to the typical single-particle [INAUDIBLE]. And so I guess I'm a little bit confused about how excitations manifest themselves in the system.

Because in some sense, if you were to solve the G-P equation, would you arrive at a spectrum of solutions? Then I guess I'd be confused at whether excitations manifest themselves into that spectrum of solutions, or rather than they become excitations, sort of deviations from the mean field as including the [INAUDIBLE].

PROFESSOR:

You are now asking about something more complicated. You're asking what are the

excited states of the many-body system? And actually, we have already found one answer. And this is if you want these small excitations, it's one particle becomes a phonon, or one particle becomes a quasiparticle.

What we are talking about here is what happens if the whole macroscopic wave function is in an excited state. So we are asking here in the box potential what happens if you force all N atoms to have one node in the wave function. And this is what I'm talking about here.

So the excitation energy of this state here is much, much higher than of single quasiparticle excitations. I'm not sure if I'm addressing your question here, but--

AUDIENCE: So with single quasiparticle excitations, are those essentially sort of deviations from the mean field?

PROFESSOR: No, they are-- you really look for many-body physics-- main field deviations-- probably. I mean, what you do is in that sense, yes. You're allowing in this Ansatz that ψ operator is ψ average plus fluctuations. You now look for fluctuations, and you're looking for the energy eigenspectrum of those fluctuations. And the answer are quasiparticles.

So it is, actually, the Bogoliubov solution for the spectrum of the fluctuations. And the answer is the dispersion relation I presented to you. Whereas here we are asking what are excited states of the macroscopic wave function? It's a very, very different question which we've addressed here. Yes.

AUDIENCE: So when you stir the condensate, do you view it as-- do you stir, sort of, at the trap frequency, so I do sort of N single excitations? Or do you stir at N times the trap frequency, so you get one sort of N particle excitation?

PROFESSOR: Well, the experimental answer is you want to stir at the quadrupole frequency, which is square root 2 times the trap frequency. Then you create quadrupolar excitations. The quadrupolar excitation can be regarded as a standing sound wave, but you create a macroscopic number of those excitations.

So you make, actually, the whole condensate wave function oscillate in a quadrupolar pattern, but it's a rotating quadrupole, and that eventually then rearranges itself and leads to vortices. So I think the correct answer is the most efficient way to create vortices would be to excite quasiparticles, but then create so many-- and this may relate to Matt's question-- create so many quasiparticles that you have really a time-dependent and oscillating macroscopic wave function.

So eventually, you have a coherent excitation of quasiparticles. And that eventually means the condensate moves in a quadrupolar pattern, but it's a quadrupolar pattern with rotation, and that eventually turns into many vortices. But that's really a very rich question which you're asking, which has been studied-- a lot of different aspects of that have been studied.

OK, so we've talked about Bose-Einstein condensates in a homogeneous system, just sort of to lay the groundwork. We've talked about Bose-Einstein condensates in traps in inhomogeneous system. Now we want to talk about Bose-Einstein condensates in optical lattices.

Well, there are two reasons why we want to do that. One is we want to use the Bose-Einstein condensate to obtain deeper insight into the properties of matter. And a lot of forms of matter appear in periodic lattices.

So if you put a Bose-Einstein condensate into a periodic potential, we can at least understand some of the properties of crystalline matter, or electrons which are block waves in a periodic potential. So this is one reason why.

The other reason why we want to go to optical lattices is the following-- ideal Bose-Einstein condensates are trivial. Weakly interacting Bose-Einstein condensates are entertaining, and you can write a lot of papers, have a lot of fun with it, develop your methods-- also mildly intellectually interesting because how those weak interactions manifest itself in vortices and all. It's really rich and interesting.

But the conceptual problems appear when you go to strong interactions. Strongly correlated matter is where mean field descriptions no longer work. This is really the

frontier of our understanding of many-body systems.

And when you want to be there, you want to create a strongly correlated system. And strongly correlated systems means that the interaction energy dominates over kinetic energy.

In Bose-Einstein condensates you can only go so far with increasing interactions-- why are Feshbach resonances? Because you get into some inelastic collisions. So that's one knob to turn.

You go to larger and larger of scattering lengths, and crank up the interactions. But if atoms strongly interact, they start to do bad chemistry. They start to spin flip. They do other things.

And in some cases you keep it under control. In others, you just can't keep it under control.

But another way to get to strong interactions means you reduce the kinetic energy. It's the ratio of the two which matters. And when you put particles in a lattice, well, the lattice actually reduces the kinetic energy. You may know from condensed matter physics that in a lattice there is an effective mass, which is higher than the [INAUDIBLE] mass. Therefore you've reduced the kinetic energy.

Or if you want another hand-waving approximation, the kinetic energy is given by the bandwidths. And if the tunneling becomes slower and slower, the width of your band becomes narrower, and your kinetic energy is less. So anyway, I can give you many hand-waving approximations why in a lattice, kinetic energy is reduced and repulsive energy is probably enhanced, because instead of having your atoms spread out, they're bunched up at each lattice site into a higher density cloud.

So anyway, maybe you're interested in parity potentials, or you're interested in quantum systems with strong correlations. For whatever reasons, you want optical lattices.

Now I have here a few slides which introduce periodic potentials, but it is really just

the single particle physics in a periodic potential. It has nothing do with quantum gases. This could be the first lecture of how to describe electrons in a metal.

So let me just quickly go through assuming that almost all of you are familiar. But I use those slides also just to give you a few definitions and introduce a few symbols. So this is simple, boring, exactly understood physics. We have a Hamiltonian which has kinetic energy, and a periodic potential which is our lattice potential.

It's rather trivial to solve, but exactly for your wave function, you use Bloch's theorem and divide it by an exponential factor with quasi momentum times a periodic function. And if you now solve Schrodinger's equation, you want to use Fourier space.

You Fourier analyze the wave function. You Fourier analyze the potential. And since the potential is periodic, sine square potential has only three Fourier components at 0 plus/minus 1 times the periodicity of the lattice.

You do a Fourier expansion for your periodic wave function. And if you insert that into Schrodinger's equation, well, the Fourier transform has turned the differential equation into an algebraic equation, because the second derivative simply becomes now k squared or q squared.

So in other words, you have a set of linear equations. There is an index which is the band index-- how high do you want to go? And usually you truncate it. But it's the same. It is that same trivial story which is told in all textbooks of condensed matter physics.

If you have no band structure, well, you have the parabola, but to prepare for band structure, I've opened it down here in the first Brillouin zone. And if you now introduce a lattice, you introduce band gaps, and you go from the left to the middle to the right for stronger and stronger lattices.

So the case which I will focus on, because it is the most extreme case away from free space, is this case which is called the tight binding limit, where the potential energy of the lattice is large. And large means compared with the kinetic energy at

the Brillouin Zone, which is k^2 -- k is the lattice momentum-- k^2 over $2m$. And that's the recoil energy. So that our dimensionless parameter here is the depth of the lattice measured in recoil energies of the photon, because that is the kinetic energy of the free gas at the Brillouin zone.

OK there are few things which immediately simplify that. Once we're in the tight binding limit, our lattice is really deep, and each site forms a harmonic oscillator. And the harmonic oscillator frequency is analytically given by the depth of the lattice.

And the solution here for the lowest band is that the lowest band-- the energy in the lowest band, the dispersion relation is-- well, we have a harmonic oscillator in each site at x, y, z . So the average energy in the lowest band is the 0 point energy in $x, y,$ and z -- $3/2 \hbar \omega_0$. And then we have a cosinusoidal band structure where q is the quasimomentum. And what appears here as the only interesting parameter is J . And J is I think this $4J$ should be $4J$ is the bandwidth.

So what appears here now for the first time is J . It appears here as the bandwidth. But let me immediately give it another interpretation as a tunneling matrix element in the following way-- right now, we have formulated the physics in Bloch wave functions which are infinitely extended. The Bloch wave functions are for the lattice what plane waves are for free space.

But if you have tight binding limit, there is another limit which is important. Namely, a particle is localized and hops around in the lattice. The localized particle, of course, is in free space. It would be a wave packet-- a superposition of plane waves.

So let's do the same in the lattice. Let's construct superpositions of Bloch waves. And these are our wave packets. And the wave packet is now called the Wannier function.

And there is a mathematical procedure how you should pick the phases here to get the maximal normalization. But the simple picture is those Wannier functions are very, very close to the Gaussian ground state solution of the harmonic oscillator on each site. The wings are different, but I don't want to go into that.

You transform from an orthonormal basis of Bloch wave function to another orthonormal basis Wannier function. And the Wannier functions are as well localized as possible. That's the procedure.

And now we can simply rewrite our total Hamiltonian or everything we're interested in, not in Bloch wave functions but in Wannier wave functions. And what comes out now is, is that the bandwidth j is nothing else than our Hamiltonian with kinetic energy in the periodic potential. But j becomes now a matrix element between two Wannier functions.

But the Wannier functions have now-- I'll put in some indices in a moment-- connect now two different sides. So it is-- you have Wannier function. You have the Hamiltonian and connect it to another Wannier function. So it is the amplitude that, with the Hamiltonian, the particle can hop from one side to another side.

So therefore j , which was the bandwidth, is now the tunneling energy divided by \hbar -- the tunneling rate from on one side to the other side. For very deep lattice, everything is analytic and it can easily be solved. And I mentioned already in the tight binding approximation, you should think about your Wannier function as just localized Gaussian eigensolutions of the harmonic oscillator. Yes?

AUDIENCE: So qualitatively, this question sounds sort of silly, but normally when we write down the solution for j , we're only considering nearest neighbor. But from the math, I don't immediately see why we wouldn't include i equal to j .

PROFESSOR: Give me one more slide. So I should've actually-- what I should have done is that-- j has an index here. I just didn't want to overload you with indices. I mean, this is sort of just telling you what j is in its simplest form. It is loaded with indices and I will show you in a moment where those indices come in.

So I just wanted to give you the idea if we want to have Wannier function. We hop from one Wannier function to the other one, and the operator is the Hamiltonian. And therefore, there should be an index j and l . And maybe-- let me just do the next step first and then come back to it, but I wanted to tell you here what I'm aiming at--

namely interpretation of j , of tunneling from one side to the next.

But right now I haven't really told you which j I really mean. There should be-- based on the right-hand side, there is a j which has two indices. And I will make the indices disappear in a moment.

But before I make those indices disappear, let me introduce the other relevant parameter, which will also have indices, and that is we have to bring in the interactions. We want to describe an interacting system. We describe an interacting system using the short-range approximation by assuming that two particles interact with a delta function.

And if you have two particles on site, each of them has a density which is the Wannier function squared. And the product of the two densities integrated gives us the expectation value for the repulsive energy. And this is given here.

So the moment we introduce interactions, we are now interested in the interaction energy between two particles which in this case occupy the same site. OK so I've tried to introduce what sort of j is. j is a matrix element between two Wannier functions with a Hamiltonian in between. And u is the matrix element of two Wannier functions with the interaction operator in between.

And now I want to use that concept to take my full Hamiltonian and transform from field operators localized at x to Wannier functions. So they B operators are now creation operators. They create an atom in a Wannier function. In other words, B^\dagger means you have a particle in a Wannier function at site i . And I just use that as a basis transformation and the exact transformation of this Hamiltonian is now into this form. So

What I have right now is, I have the tunneling matrix element between particles at site i and j , and I have to sum over all of them. And in terms of interaction, I can calculate this matrix element by using Wannier function with four different indices and ask what happens.

And this is simply an exact way of rewriting it. And here I've given you the definition.

So if you want to forget everything I told you about j and u , I've done an exact transformation from field operators to Wannier function creation operators. And this introduces tunneling terms like this from site i to site j .

And these here includes products of four Wannier functions which are responsible for the interaction term. Actually, if this is not complicated enough, I've suppressed band indices here. I should also now sum over all possible bands.

But OK, I want to come now to the leading approximation in a tight binding model. And that is where-- I mean those Wannier functions are overlapped. Two neighboring Wannier functions barely overlap. If I go further away, the overlap becomes even smaller. So the most dominant terms are nearest neighbor interactions. And the nearest neighbor interaction is where i and j differ by 1. And this is what I call j without indices.

And similarly, when it comes to the interaction term where we have products of four Wannier functions and we want to get the overlap of all four, and then multiply it with g , the prefactor of our delta function to get an interaction term, well the best overlap is if all indices are the same. And this is what I call U .

So in that limit, in that tight binding limit, my Hamiltonian is now very simple. It consists of a tunneling term parametrized with j , and an on site interaction term parametrized with U . Yes.

AUDIENCE: So you're saying i , j and k are all the same, so the interaction is with itself?

PROFESSOR: No, two particles per site. When the four particles are the same, it turns into-- this Hamiltonian here has all indices are the same. And if you calculate that, it turns into the product of $2 b_i b_i$ dagger is the occupation number at each site. But if you're careful with commutators, it becomes occupation number times occupation number minus 1. If you have only one particle per site, this term is 0.

So technically U is, if you put in 2×1 , if you have two particles per site, U is the interaction energy between two particles. Just use this expression to figure out what it is for three or four. But for one particle, you get 0.

The self interaction of a particle is absolutely 0. One particle does not interact with itself.

AUDIENCE: That gets back to my question with the [INAUDIBLE]. So when you have i equals j , the [INAUDIBLE] term, do you just--

PROFESSOR: No, I assume, i j is nearest neighbor. I assume that the index i and j differ by 1.

AUDIENCE: Yeah, I mean your argument was that the best overlap is between nearest neighbors.

PROFESSOR: Yes.

AUDIENCE: So an overlap with the Wannier functions at the same site is sort of the kinetic energy term. Is it approximately 0, or you just absorb that into the chemical potential?

PROFESSOR: This is a 0.

AUDIENCE: [INAUDIBLE] So we're only looking at nearest neighbor, so are we essentially saying the case where i equals j is approximately 0 [INAUDIBLE].

AUDIENCE 2: Because it's the matrix element of a Hamiltonian. So p , the candidate, plus the lattice operator. So if I take my Wannier function, which is built out of eigenvalues for that equation, [INAUDIBLE] bunch of energy terms. On every site [INAUDIBLE] they're the same. Because you have a common [INAUDIBLE] energy that you can [INAUDIBLE].

AUDIENCE: Well then just multiple it by [INAUDIBLE].

PROFESSOR: My gut feeling is-- and this is why nobody considers it-- it's just a constant energy, which is probably something like the 0 point energy times the number of particles, which is not affecting any dynamics. It's pretty much a constant which can be simply dropped.

Let me just go back to that slide. I hope-- let me just look up the reference. I haven't

looked at it recently. When we do this exact transformation, there should be a reference whether i and j , what happens when i equals j . I think it's just a constant term.

AUDIENCE: Because if you're summing over every single lattice site, so you get the number of particles times 0 [INAUDIBLE].

PROFESSOR: Yeah, but even mathematically it gives just a constant term, here, which is [INAUDIBLE]. Thanks. It was good to clarify it.

OK, I think within the next 20 minutes, I can step you through the superfluid to Mott insulator transition. First, references for what I've just said are given here.

But let me now come take this Bose-Hubbard model and discuss its two limiting cases. One case is where U is much larger than j . The other case is where j is much larger than U . These are the two limiting cases, and it will turn out that one is an insulator and one is a superfluid.

And that makes perfect sense, of course. If U is much larger than j , I can set j equals to 0 if I can neglect it. No tunneling means no transport, and that means an insulator. It's also clear that when there is no tunneling, that the system is really described by a product of so and so many particles per site.

So it could be one particle per site. And I have a product over all sites, or two particles per site. So this will be the ground state. And it's called the ground state of the Mott insulator.

It's also trivial to discuss what happens if j is much, much larger than U , because then I simply neglect U , and I have a free gas. Well, a free gas of Bloch wave function, but that's the same as a free gas of planar wave function. It's just that quasimomentum replaces momentum. And if you have just an ideal Bose gas in a periodic lattice instead of Bose-Einstein condensation in the lowest momentum state, you have Bose-Einstein condensation in the lowest quasimomentum state.

The lowest quasimomentum state is a superposition of all Wannier function-- I

mean, the 0 momentum state in free space is a superposition of all position delta functions. The plane wave is delocalized. And the lowest quasimomentum state is just completely delocalized over all Wannier functions. But this is nothing else than the $q = 0$ quasimomentum Bloch wave.

The interesting question is-- and this has led to hundreds if not more papers in the literature-- how do we go from one limit to the other limit? And now you would say, well, maybe we should use what has worked so well. We should use the Bogoliubov approximation. Just assume we have a condensate, replace all those kind of a operators for the lowest Bloch waves.

However, this doesn't work, because the interesting thing here is to find the transition from the superfluid from the Bose-Einstein condensate to an insulating state. The nature of this approximation is that you need N_0 to be large. But we're interested when the condensate wave function turns to 0, and in the insulator it becomes 0.

So we are actually interested in the opposite limit. And indeed, if you would ignore everything I've just said that doesn't make sense to make this approximation and make it nevertheless, you will find that you never get the insulating state because you've pretty much eliminated the possibility to describe an insulating state by doing this approximation.

So I want to show you now that we have to do another mean field approximation, which is actually nice. It's very different from this Bogoliubov approximation, but it's also a mean field approximation which will describe our system. So the goal is now that I want to find an effective Hamiltonian which describes the transition from here to there.

And the important approximation I will use is the following-- again, I have to get rid of operators. Products of two many operators cannot be solved. And so what I will do is I will use products of operators, write them as average value plus fluctuations.

And then when I multiply that out, I take the product of the average values, and I include the fluctuations in leading order. So I take δA times B and δB times A average, but I neglect the product of those fluctuations. You can say I neglect the correlation of fluctuations here.

So this is spelled out here. But the sign is important. Just look at this equation-- $A \delta B + \delta B$ with B plus A average times B average.

If I absorb A times B by upgrading δB to B , but I do the same here, I have to subtract 1 product of the two average values. It's actually this minus sign which will play a role later, but here you see already that I will make this decoupling approximation-- that I decouple the fluctuations from each other and I write it in this way, there is an important minus sign.

OK, so we want to start in the insulating state. And we're going to figure out how does the system develop superfluidity out of the insulation state? And the perturbation operator which takes me out of the insulating state is tunneling.

So therefore, the operator which is responsible for breaking out of the insulating state is the operator which induces tunneling between neighboring sites. And the tunneling operator in this Hamiltonian-- remember, the Bose-Hubbard model had j times $B^\dagger B$ -- this was our tunneling term-- this involves now products of operators.

And I told you we want to get rid of products of operators to get something we can easily solve. So we use now this product of operators on two neighboring sites. And we use exactly this decoupling approximation.

So therefore, we replace each operator by an average value, and we neglect the product of the fluctuations. And then we obtain this. And I explained to you where the minus sign came from. So now I call this average value of the operator B ψ , I call the superfluid order parameter ψ .

I think I could have chosen ψ to be complex, and then $B^\dagger B$ would have complex conjugate. But it's sufficient here to restrict the discussion on real numbers,

and it makes the notation simpler. So anyway, I introduce that.

And what happens now is the following-- that I have my Hamiltonian. The Hamiltonian had the interaction energy. But the tunneling term is now very simplified, because instead of tunneling from one site to a neighboring site, the other side is sort of absorbed by the mean field, by the superfluid order parameter ψ .

And therefore, and this here gives me a ψ squared term which appears here. Trust me, this is just an identical rewrite of the previous Hamiltonian by using this decoupling approximation. So what we have gained now is something really dramatic.

We had many sites and tunneling from site to site. But if you look at it now, we simply sum the effective Hamiltonian over site index l . So our effective Hamiltonian is now the sum of an identical Hamiltonian per site.

The sites no longer interact with each other. Each site interacts with all the other sites described by the mean field by the superfluid order parameter ψ .

So therefore, our many-body problem, which is still a many-body problem, but has turned into an effective Hamiltonian for each site, because each site has the same Hamiltonian. Colin.

AUDIENCE: What happened to the j ψ squared term?

PROFESSOR: This here?

AUDIENCE: Oh, OK.

PROFESSOR: OK so I can say instead of solving for the sum, I can just solve for each site individually. And this is now my effective Hamiltonian for each site. Now I want to catch the onset of superfluidity. So I want to get the system when ψ is small.

And therefore, I can just ask-- I don't know what ψ is. It's part of my solution. But I'm interested in the moment when ψ begins to take off from 0, when superfluidity

emerges.

So what I can therefore do is, I can regard ψ as an epsilon parameter, as a small parameter. And the ψ parameter comes with an operator V . And this operator V is nothing else than $B^\dagger + B$.

So in other words, what I'm doing is, I'm separating my Hamiltonian into a Hamiltonian which is diagonal in the quantum numbers of the isolating state-- just one, two, three particles per site. And ψ^2 is the Z number. ψ^2 is also diagonal in that.

And now the possibility of tunneling, the possibility of superfluidity is now perturbative in this term ψV .

OK, I don't want to explain, actually, this expression. It just formalizes [INAUDIBLE] intermediate step. When we have the chemical potential, and we raise the chemical potential-- we go from zero to one particle per site to two particles per site. And whatever the chemical potential is determines whether we have one or two particles per site.

This is just telling me as a function of chemical potential, what is the ground state of the insulator? So now we take this ground state-- it's actually much easier described in words than by this formula-- we take this ground state and do perturbation theory in our term ψV . Remember, the operator is $B^\dagger + B$. And the epsilon is ψ .

So in second-order perturbation theory, we get ψ^2 . And then V , because $B^\dagger + B$ is very simple. It only couples one occupation number N to $N + 1$ and $N - 1$, because B and B^\dagger destroy or create a particle per site.

So therefore, I can immediately write down what this matrix element is in second-order perturbation theory. I mean, these are all-- sorry, it's all defined here. I know I'm losing you now. Nobody will tell me what is the difference between \bar{U} and U , but it's trivially defined.

So the idea is we have the isolating system. We do perturbation theory in tunneling. And the perturbative operator is ψ times B plus B^\dagger . The B plus B^\dagger matrix elements are trivial, because they admix to N particles per site-- $N + 1$ and $N - 1$. And this is what we've done here.

AUDIENCE: And the j is just occupation number?

PROFESSOR: The j is the occupation number of our site.

AUDIENCE: I have two questions. So in the bottom equation, what happened to the ψ squared?

PROFESSOR: Sorry, this is the sum. The ψ square is missing.

AUDIENCE: I guess I had the same question for the equation next to the green thing. Is there supposed to be a ψ squared in there? Because originally, there's a ψ squared and an h_0 .

PROFESSOR: No, this is the ground state where ψ is 0. And now we do perturbation theory in ψ V . These are the unperturbed energies which appear in the energy denominator.

The wave functions we are using are Fock states-- number states-- per site. And here we couple occupation number j to all possible N 's. But because of B and B^\dagger , N has only two values-- $j + 1$ and $j - 1$. And this is what's given here.

It's mathematically trivial. The notation is more complicated than the physics behind it. But now comes again-- I think I need five minutes and I'm done. But now comes the interesting physics.

We have to ask what are we doing here? It's all the mathematical, and the math is really simple here. What we have done is we have looked at the isolating state, and we have done perturbation theory in ψ times B plus B^\dagger . And now we get an energy correction which is ψ squared. And what are we really doing here?

Well, you can say the following-- we started with a hypothesis that we have a superfluid state characterized by ψ and ψ , at the onset of superfluidity-- is small.

But now we have done the calculation assuming that there is a ψ . But now we are turning around and said have we really done the system a favor by introducing superfluidity? In other words, has our perturbation theory in ψ lowered the energy of the state or raised the energy of the state.

So in other words, we've done a hypothetical calculation. Hey, what would the system feel like if there were a little bit of superfluidity? If the system says great, I've lowered my energy, then we know we are in the superfluid state.

When the system says no, I raised my energy because of the ψ , then the system has rejected our idea to introduce superfluidity. So therefore, the question we are raising now is after we have done the calculation, for what values of U and j is it favorable to introduce a ψ or not?

Now I was expecting the question of some of you that in second-order perturbation theory, second-order perturbation theory always lowers the energy. But remember, this is why I emphasized the minus sign-- we had a ψ squared term which came from that, which came from the last term of the decoupling approximation, which had a minus sign.

And therefore, we have in ψ squared one term which came from this special ψ -- and I emphasized in the decoupling approximation. And we have a contribution of ψ squared which comes from perturbation theory. And the two together can actually change their sign.

So what we have right now is if you describe the ground state as a function of ψ , we have our unperturbed energy of the Mott insulating states, and then we have a term in ψ squared. And we should-- and we could, but we don't-- calculate the next order in ψ to the four. And it turns out in fourth-order perturbation theory, this term is always positive.

So what happens now to the total energy when this term A_2 , which we have exactly analytically calculated-- if this term A_2 is larger or smaller than 0. Well, if you have a parabolic term and a quartic term, in this case, both the quadratic and the quartic

term are opening up like in a U shape, but here the total energy turns into W shape.

And the interpretation is fairly simple. Under those conditions, the system prefers to have ψ equal 0, whereas here, it's like the system wants to have a ψ parameter which is finite, either positive or negative. And whether it's positive or negative is sometimes called spontaneous symmetry breaking between two degenerate solutions.

What I'm showing here is, of course, very similar to the Ginzburg-Landau theory of phase transitions, where you have an effective potential and the phase transition takes place when the second-order coefficient changes sign, and the effective potential turns from U-shaped into W-shaped. So that's what we have done.

So anyway, I think this is a nice problem where the interpretation is, I think, much more subtle than the calculation itself. But we've calculated the phase transition. Everything is analytic. Here is the analytic result, which I know with very indigestible notation.

I followed exactly the paper by Vanderstraeten, [Stouff ?] and collaborators, which is posted on the website. You can plot this phase diagram in this way. But if you use what is more common, normalize the chemical potential by U and normalize j by U , you get these wonderful lobes of the Mott insulator where you see that, if you increase the chemical potential, you have N equals 1, N equals 2, N equals 3, Mott insulator, and in between you go through superfluid regions.

However if your tunneling is larger, if your tunneling is too large, you're only superfluid. So this is sort of the way how you derive this rather rich phase diagram of bosonic atoms in an optical lattice. Green is an insulating state. And white is the superfluid state.

Let me just conclude by showing a few slides how this can be observed. So in one case, we have an insulating state with a definite number of particles per site. And here we have superfluid state, which has the normal fluctuations in number.

Here's another cartoon picture of an condensate in the lowest Bloch wave function.

And here you have localized number states. Often, you observe the transition by taking the system and releasing it in time of flight into ballistic expansion.

If the superfluidity is coherent, the wave function on the different sites is coherent. And what you get is a multi-slit interference pattern, which is a diffraction pattern. And you see these characteristic diffraction patterns. It's an in-slit diffraction pattern which characterizes the coherence in the initial state.

However if you have a completely isolating state, you have a Gaussian wave packet on each side, which has nothing to do with the next neighbor. And then the Gaussian wave function simply expands in a structureless way. And so the transition from here to there happens exactly when the ratio of U over j crosses the value we have derived.

There's another technique of observing that which is more recent. This is a quantum gas microscopes where you do the same physics in two dimensions, so you have only one plane where this physics takes place. And you can observe atoms now with a microscope.

So each green dot here is an atom. If you sort of zoom in. And if you now study the region where you go through the superfluid to the Mott insulator, you see that you have exactly one atom per site.

If you increase the atom number, you have an insulator with one atom per site. And in the middle, you have state with two atoms per site, which-- for reasons I don't want to discuss-- are color-coded here in black. And that sort of goes on and goes on. So you can really resolve site per site the occupation of the number of atoms per site.

So I think I'll stop here with the Bose gases. When you think the Mott insulator is the end of the story, all motion has been frozen out and you have one particle per site, well, it could be the beginning of a new story. Because if you use two different hyperfine states, spin up and spin down, we can talk about spin ordering.

But this would be a whole different lecture.

