MITOCW | watch?v=FU3P-vnGSZ0

PROFESSOR: Good afternoon. We have talked in the last class about magnetic trapping. Today, I want to finish our discussion on magnetic traps with two demonstrations of classical forms of magnetic trapping. But I told you that magnetic trapping is a purely classical phenomenon. The only quantum mechanical aspect is that this angle at which the magnetic dipole is oriented with the magnetic field is quantized.

I also mentioned to you, and this was Wilks' Theorem, that the total magnetic field can have only local minima, not maxima. And therefore, we you can do magnetic trapping only for weak field seeking states. These are states which lower their energy with photomagnetic fields. And as a result, since the spin is up, the magnetic moment is anti-parallel to the magnetic field, those states can always lower their energy by flipping the spin.

I know it's a simple demonstration, but, well, I can't bring a real magnetic trap into the classroom. So what I have here, this is just two little ring magnets, refrigerator magnets. And there is a tube. So I just want to demonstrate the one dimensional form of magnetic trapping to you.

And this is our atom. It's in strong magnet, elongated, and it can be spin up or it can be spin down. So what you're seeing is now you're seeing the two magnets. And you're seeing our dipole, our atom, which is in a stable trapping configuration.

And you can-- am I blocking it? You can see that it's stable by-- I just move it. And you see how it always comes back. So that's one position.

I pull it over. Here, you have a second position where our magnet is trapped. Here is a third one. And here's a fourth one. So there are four positions where we have one dimensional trapping. Well, this is one for one orientation.

Now I take the magnet and flip it over. So we are changing now in the laboratory frame. The sign of the magnetic moment. And there is actually a very nice minimum. You see how clearly it is trapped.

There is one here in the middle. I can also just show you, it really stays put. There's a weak restoring force. And then here's number three.

So it seems very, very rich. I've shown you seven different minima where trapping has happened. And well, just to fill you in, if this is our magnet, these are ring magnets, the magnetic field is about something like that. And if I plot the magnetic field for this configuration, I find indeed that for one orientation, there are four maxima.

So for one orientation-- just give me one second. Yeah, there are four maxima of the magnetic field for this orientation. And therefore, there are four minima of the trapping potential.

So maybe I should have shown it like this. There are one, two, three, four minima. And for the opposite configuration, there are one, two, three trapping minima. But now if you look, there are one, two, three, four, five, six ones marked in red, which are for strong field seekers, because these are maxima at the magnetic field. And there is only one where the magnetic field has a local minimum.

So in the last part, I have shown you the absolute value of the magnetic field. And you realize there are one, two, three, four, five, six maxima. And this is responsible for six trapping configurations. And the seventh one was a minimum of the magnetic field.

So in one dimension, there is no problem. In one dimension, I can get a restoring force for the strong and for the weak field seeker for any orientation of the magnetic dipole. But if we want to have a three dimensional trap, we have to look at the three dimensional stability.

And it would be only this one configuration, which can be stabilised in x and y. What

I've shown you are symmetric fields. So all these minima and maxima only settle points.

They trap in one dimension. But with suitable radial magnetic field, you could turn this one here into a real three dimensional trap. But the other ones you cannot, because this would violate Wilks' Theorem.

I find it amazing that this demonstration has seven different magnetic trapping configurations. That's just how it turns out to be. Questions? Collin.

AUDIENCE: So if you look at the clover leaf coil, the parts of the field that sort of axial [INAUDIBLE] provided by [INAUDIBLE], those are the inner coils?

PROFESSOR: Yeah, so if you look at a clover leaf trap, the pinch coils create such a minimum, which provides confinement in the z direction. But then you add a field, a radially outward quadrupole field, and that overcomes-- in this demonstration, everything is actually symmetric.

So this is a minimum for this direction. But if you would now plot it as a function of x and y and z, it would be a saddle point. So this configuration is trapping along z, but anti-trapping along x and y. But if you add a quadrupole field, a linear field which points outward with these clover leaves or the [INAUDIBLE] bars, then they would provide confinement, which is stronger than the anti-confinement of the saddle point.

If you look at some equations I showed you in the last class, you will actually see that in one of the slides, I had the expression for the magnetic field, where you see how the [INAUDIBLE] bars, the radially linear confinement overcomes the anti-trapping feature of the settle point. Other question?

Then finally, let me give you a demonstration for the following effect. I mentioned to you that magnetic traps become unstable when the magnetic field is very low. And I told you, well, at very low magnetic field, the energy difference between spin up and spin down is very small. At zero magnetic field, it becomes zero. And then, the atom cannot adiabatically stay in its MF state, because we can violate the adiabaticity condition.

So let me now show you an example of purely classical trapping. And I think a number of you now have seen the levitron toy. Actually, a few years ago when my group was one of the first to use magnetic trapping, in every single talk I explained how magnetic trap worked. And I even showed this demonstration live.

But in order to demonstrate it really well, I was building my own levitron. I used machine parts. I used the motor drive. So I hope I can show you one aspect of the levitron demonstration, which is usually not known.

So let me first give you the punchline. Magnetic trapping happens with an orientation of the magnetic dipole which could always lower its energy by flipping over. The reason why it doesn't flip over-- well, quantum mechanically, it adiabatically stays in a quantum state-- but classically, if we have angular momentum, it is gyroscopically stabilized.

The dipole cannot flip over, because it's a gyroscope. It has angular momentum. And actually, it also has angular momentum in quantum physics. So the two explanations sounded different when I said stays in a quantum state, is gyroscopically stabilized. But if you think about it more deeply, they have much more in common than my language suggests.

So I think you know how the levitron works. You spin up the magnet. And I had this nice motor tool to spin it up. And when you have prepared the system. Your atom has now-- atom in quotes-- angular momentum in the magnetic moment. And you bring it to the position where the three dimensional magnetic field fulfills the stability condition. And you can now enjoy magnetic trapping.

This is exactly what your atoms do in your magnetic trap. It is gyroscopically stabilized magnetic levitation. The only difference is that gravity has to be compensated. Gravity is a major player. So the stability point in the three dimensional magnetic field configuration includes the compensation for gravity.

But I have to say, my group was also at some point trapping a Bose-Einstein condensate in a few hertz weak trap, where gravity was the strongest force of all. So what you have seen is an exact demonstration of the principle for magnetic trapping.

But now comes my question. What would you expect, what would happen to our magnetic trap when we spin the levitron faster? Does it help or does it hurt to spin up the levitron, to spin up the gyroscope, to higher angular velocities?

So three possibility. Nothing happens. It doesn't matter as long as it spins. The second one, the trap becomes more stable. The third possibility, the trap becomes unstable. Do you want to offer any opinion? Collin.

AUDIENCE: Well, we have an angular momentum, right, of the really large omega. When you apply torque, it's going to get torque [INAUDIBLE], right. It's going to be the prefactor of omega. So we have a giant omega, we imagine that [INAUDIBLE] small, [INAUDIBLE] torques. No, no, no wait a second. I did this the other way. It was [? healing. ?] Never mind.

PROFESSOR: So do we have to do the experiment? Maybe. So now I put the motor controller to full speed. And I speed it up much faster. You can hear the sound.

I really want to do a careful experiment. So we wait until everything is quiet and has died out. And now, we want to try if we can can do magnetic trapping. And you see it's impossible. The system, when it reaches the point where magnetic trapping would occur, it's unstable.

But then, well, just to prove that it is only the speed of rotation which has caused the instability, I just wait until friction has slowed down the angular velocity. And now again, it works perfectly.

So you see, if you rotate the gyroscope too fast, it's bad. It makes the magnetic trap unstable. Convinced? How would you explain that?

If we have a gyro-- Collin?

AUDIENCE: Why aren't we working in the limit where we assume that the magnetic field generated by our magnet is sort of weak compared to the trap. So the magnet's really modifying [INAUDIBLE].

PROFESSOR: No, we assume here-- and I can give you the analysis-- but yes. No, these are permanent magnets. So the floating magnet is just, you can say, a probe, a test object which is put into the permanent magnetic field of the stronger magnets of the stationary magnet.

AUDIENCE: Get an additional torque to-- because it gets an additional force into the upper state. So it gets an additional torque in towards the center.

PROFESSOR: It's not necessarily the additional torque. Let's put two things together, it's really fascinating, from different principles. The first thing is magnetic trapping requires-- if you have a magnetic trap and you have an inhomogeneous magnetic field. And of course, you need an inhomogeneous magnetic field, the angle cosine theta between the spinning dipole and the magnetic field should stay the same.

And this means quantum mechanically, we stay in the same quantum state. So therefore, a magnetic trap only works because the rapidly precessing spin, when the magnetic field always precesses around the magnetic field. And when the magnetic field tips, the precession keeps the dipole, the magnetic moment, aligned with the magnetic field.

Now, what happens in a gyroscope with a precession frequency when you spin the gyroscope faster? We've seen that in your classical mechanical demonstration, if you had the spinning gyroscope, which was only [INAUDIBLE] suspended with one rope. And then it was precessing in the Earth magnetic field. Does this precession frequency get faster or slower when you spin the gyroscope faster?

Pardon? When the gyroscope spins faster, what happens to the precession frequency?

AUDIENCE: Slower.

PROFESSOR: Slower. Because the torque per unit time adds some angular momentum. This angular momentum adds to the existing angular momentum. But the more angular momentum exists, the smaller is the change. And this lowers the precession.

I can give you another example. If you rotate a coin. When does the coin really wobble very, very quickly? Just when it has slowed down and is about to fall. And this rapid wobbling is the precession frequency.

So the lower the gyroscopic angular momentum is, angular velocity is, the faster is the precession frequency. And fast precession is important for adiabatic following. So in other words, what you saw here in this demonstration was a classical analogy for Majorana Flops.

Now, if we would translate from our classical demonstration to a real atom, what feature, what parameter characterizing the atom, are we changing? So when I spin the magnet faster, what would that correspond to in atomic properties?

AUDIENCE: Higher mu.

PROFESSOR: Higher mu? I'm not changing the permanent magnetic moment of the magnet by spinning it faster-- higher angular momentum. But in this the equation, what corresponds to higher angular momentum?

AUDIENCE: [INAUDIBLE] omega alpha, the precession frequency around the static beam field?

PROFESSOR: Yes. But what I mean-- so the Larmor frequency, this is a precision frequency, becomes lower. So what becomes lower in the atomic property?

AUDIENCE: Sort of the external magnetic--

PROFESSOR: Pardon?

AUDIENCE: h bar?

PROFESSOR: No, let's not mess around with h bar here. h bar is given by nature. We can't change that.

But, I mean, OK. Multiple choice. B, no, no, no, no. It's g. Yeah, it's g. So what you have seen is, you've seen a demonstration where in front of your eyes, I've changed the atomic g factor. And now you sort of see, let me put it together.

What happens is, the mechanical magnet has a given magnetic moment. And if I put much more angular momentum into it, it can sort of-- it has, quantum mechanically, speaking more intermediate states. Because it can change its angular momentum in steps of one.

So if I spin it faster, it has many more intermediate states. And each energy separation has become smaller. And smaller energy separation means I'm getting closer to degeneracy where adiabaticity breaks down.

Anyway, think about it. The analogy is really deep. Questions? OK.

So that's all I wanted to tell you about magnetic trapping. Collin.

AUDIENCE: When you increase the angular momentum, you don't necessarily change the spacing between the levels, though.

PROFESSOR: No. The energy levels is when the magnet is aligned, it has an energy absolute value of mu times b. Here, it has minus absolute value of mu times b. And mu is simply the magnetic moment of the permanent magnet. So I go from here to there. And the number of energy levels in between is the total angular momentum divided by h bar.

So when I give it more angular momentum, in one energy level, in one transition, there is less energy which will be released. And it is actually, you can say, the big note between energy levels, or the difference between energy levels, which is a Larmor frequency.

So therefore, the precession, quantum mechanically, is the energy difference between adjacent energy levels. And if you give it more angular momentum, this energy difference becomes smaller. And this low precession frequency, this small energy difference, is bad for adiabaticity. Yes.

AUDIENCE: Does that have much of an effect on f equals 1 and f equals 2? Do you get more [INAUDIBLE] if f equals 2?

PROFESSOR: Yes. In f equals 2, we have twice the magnetic movement than in f equals 1. But in f equals 2, we have five levels. In f equals 1, we have three levels. So the--

AUDIENCE: Is it that significant of a difference?

PROFESSOR: So I think it just cancels out. The g factor-- I showed you the formula. What matters really is the g factor. And the g factor on f equals 2 is 1/2. In f equals minus 1, it's minus 1/2.

What happens is the magnetic moment in f equals 2 is larger, because everything is stretched. All angular momenta are aligned. But the multiplicity-- you have five levels versus three levels-- and the two effects just cancel. Other questions? Yes.

AUDIENCE: What breaks the system when the magnet spinning gets slower and slower. Now we know why it destabilizes when you spin it too fast. But if you don't spin it at all, it also floats, right?

PROFESSOR: OK, what happens is yes. If I stop spinning it, it will no longer work. So what we have is we have a hierarchy of frequencies. The fastest frequency has to be the spinning frequency. Then we have the trapping frequency. And the precession frequency is one over the spinning frequency.

So you want that the precession frequency is between spinning and one over spinning, because this is the precession. And if you would take the spinning to lower and lower values, you would violate that hierarchy. Yes.

OK, evaporative cooling. Evaporative cooling is a powerful cooling scheme to reach nanokelvin. Actually, I forgot to update this slide. I wanted to say, this is the only technique so far to reach quantum degeneracy for bosons and fermions.

Very recently, people have demonstrated laser cooling of atomic strontium to quantum degeneracy. But if you read the paper, it was laser cooling aided not by evaporative cooling, but by collisional distribution of atoms. It's likely evaporation where you evaporate into-- and you keep-- read the paper. It's--

[LAUGHTER]

The scheme only worked because collisions-- how to say-- you cooled one region and another region was cooled by collisions. And this brings you pretty much back to further [INAUDIBLE] evaporation. Anyway, what I want to say is there is a small footnote. The field is evolving.

You can find now paper's laser cooling to be easy. But if you read the paper carefully, or if you talk to me, I will tell you that there were still collisions necessary, the same kind of collisions which drive the reparative cooling.

But before we go into an expert discussion about variants of evaporate cooling, I should first tell you what evaporative cooling is. But somebody raised his hand.

AUDIENCE: Oh, no I just [INAUDIBLE].

AUDIENCE: You wrote the paper on [INAUDIBLE], for example, that if your quantum degeneracy [INAUDIBLE]. And that would be just purely laser cooling, right?

PROFESSOR: Sub-recoil cooling has not-- any form of laser cooling to sub-recoil temperatures was not compatible with high atomic densities. It only worked at such low densities that they stayed far away from quantum degeneracy. So nanokelvin, yes. Temperature in the nanokelvin range his been reached by laser cooling, but not at sufficiently high densities.

The densities, high density, causes collisions. Those collisions are screwing up laser cooling that it doesn't work anymore. And the only technique which can reach nanokelvin temperatures at sufficiently high density is evaporative cooling. And that applies to strontium. Strontium was laser cooled at low density, and then the low temperature was collisionally transferred to high density region.

So you can say, in a way, it's an oddity of nature that we are now using quantum gases to reveal new features of quantum physics using ultra-cold atoms. But the cooling techniques which gets us into the quantum degenerate regime is pretty much classical.

So this is a cartoon picture how you could think evaporative cooling works. You have a thermal distribution. You remove the high energy tail of the thermal distribution.

And then, you allow the distribution to relax. And it will relax to a Boltzmann distribution, which is a little bit shifted towards the cold, or the low energy, side compared to the original dash distribution. And if you do it again and again and again, you wind up with a distribution of atoms which is colder and colder, because every time you axe away the high energy tail, you remove atoms which have, on average, more than the average energy.

And therefore, the average energy per atom drops and drops and drops. Of course, the atom number also drops rapidly. And I was only able to draw it in this way because I've been on-- I think I-- actually, did I. I forget what the normalization here is in the plot.

AUDIENCE: What is n?

PROFESSOR: n is a number of steps. So after 25 removals of the high energy tail, I'm here. And after 50, I have a very, very narrow distribution. So it tells you already something very powerful, which people were not fully aware before evaporative cooling of atoms was invented, that it doesn't take so long. It doesn't take so many steps.

It takes 50 rethermalization steps. And each rethermalization takes two or three elastic collisions. So what this already demonstrates, if you can keep your atoms and evaporatively cool by removing the high energy atom, after a time which corresponds to a few hundred elastic collision times, just a few hundred collisions, you can go way down in temperature.

So if you have this, I actually used this cartoon picture to write one of the-- to develop a mathematical model for evaporative cooling, which is still the simplest model for evaporative cooling which you can find in the literature. So in any event, but if you think in that way, every time you remove a few percent of the atoms, you realize that you should think in a logarithmic way. Every time you lose a certain percentage of your atoms, you decrease the temperature by a certain percentage. So in the end, if you think either in discrete steps or continuously as a function of time, things should happen exponentially.

So if you want to characterize what happens into this system, we should correlate the percentage of temperature change to the percentage of the change in number. And here we have a coefficient. And this coefficient would give us an exponential-- would give us the power law, how temperature and number are related. Or mathematically alpha, which characterizes how much cooling do you get for which loss in the atom number is the logarithmic derivative of number with temperature.

All other quantities also scale as power loss of the number. Let me just assume we are in a potential. We have T dimensions.

So if you have an harmonic oscillator, it's r square. If you have a linear trap, it's r to the one. And for reasons for simplicity, I took d, the number of dimensions, out. So you choose delta to get the 1d, 2d, or whatever to get a harmonic oscillator or linear potential.

OK, Because of equal partition, the temperature, which is a measure for kinetic energy, is equal to the potential energy. And the potential energy is r to the power d over delta. So therefore, the size of the atoms in the trap scales with temperature with the power law. But if the temperature scales with the power law over the number, then the volume of the atoms, the temperature, everything scales with the number of atoms to some exponent. So everything is sort of exponential according to power laws.

So I mentioned already, the volume and what I've shown here is just the three parameters which determine-- or the two parameters-- which determine all exponential. The volume goes with the product of delta alpha. The density goes with something.

The phase space density is density over temperature to the 3/2. The collision rate is n sigma v. So everything scales with n to the alpha. And all the coefficients are given here.

So the question is, what is alpha? So it seems delta is our trapping potential. Once we know alpha, we have a clear prediction what evaporative cooling can do for us.

Well, alpha was, remember, it told us how much we lower temperature or energy when we lose a number of atoms. So therefore, all we have to figure out is, when we evaporate an atom, how much energy does it take with it? And therefore, alpha will be determined by some people call it the knife edge.

If you truncate the trap at eta kt-- I showed you a cartoon where I used the x to chop off the tail of the Maxwell-Boltzmann distribution at 4kt. So eta would have been 4. So this is what we control experimentally. At what energy do we allow atoms to leak out of the trap.

All right. So if you set a threshold of, let's say, 4kt, 4kt is the minimum energy for atoms to leak out. But some will be a little bit faster. In other words, they are not creeping over the edge. They are jumping, they are zipping over the edge.

But because everything is thermal, this extra energy is on the order of kt, or 1/2 kt. And in the first analysis, we can neglect it. So in other words, we can say that each atom, when it escapes, takes away the energy eta kt. And eta is the famous eta parameter, which we determine experimentally when we evaporate. Any questions so far?

Right now, I've pretty much gone through definitions. And now, I simply look at energy conservation. So now let's look at what is the change of energy during evaporation.

Well, for n atoms, this is the kinetic energy. And the extra potential energy for harmonic oscillator, it would be the same. Equal partition for another power law potential that we have to introduce as delta parameter, which defines the power law of the trapping potential.

So that means the following, that originally, this describes the number of the total energy. After an evaporation step, delta N is negative. I've lost some atoms and delta T is negative. I am now at a lower temperature and a lower atom number. And the difference is simply the energy taken away by the number of atoms TN which have evaporated.

So with that, by just rewriting that, I get a result for the alpha coefficient. The alpha coefficient, which tells us what is the percentage in temperature which-- how many percent is the temperature lowered when I lose a certain percentage of the number of atoms. Actually, 1% drop of the number of atoms gives alpha percent in change in temperature. And this is the alpha coefficient.

So I have an analytic expression for that. And sure, you realize if you put your cut eta not at high energy, if you cut at lower energy, than your alpha coefficient can even turn negative. Because then, the atoms which evaporate do not have more than the average energy. But actually, then the model breaks down.

OK. So alpha characterizes how much more than the average energy is removed by escaping atoms. So [INAUDIBLE] are very simple once we know what is the threshold of the trap, eta kt. At what energy do we leak out atoms?

We have a complete description what is the energy, the phase space density, and such and such after we've lost a number of atoms. But you realize, at least for all of you who do the experiment, something is missing here that sounds almost too ideal. We just evaporate.

And we can freely peek what is the energy, or what is the energy threshold and such. The experiment is more constrained. And we have to work harder to get into this good regime of evaporation.

But let me introduce the experimental constraint which is very important by asking the question how efficient can evaporative cooling be? So based on the idealistic model, which I've presented to you so far, what is the highest efficiency of evaporation you can imagine? Collin.

AUDIENCE: I guess in principle, you could remove just one atom. Then you'd save [INAUDIBLE].

PROFESSOR: Yes. You remove one atom. You wait until you have one atom, which has pretty much all the energy of the system. One atom evaporates and your whole system is as cold as you want to have it. Of course, you are all laughing because this will take much longer than the dwell time of a graduate student at MIT. In other words, what you realize, time is a premium.

And it's not just the dwell time of a graduate student. It's not your patience. What happens is in a real experiment, there are losses. There is some form of technical heating.

Since you don't have a perfect vacuum, residual gas coalitions cause losses. And its clear you have a time budget which is set by losses. And either you evaporate in this time budget, or you've lost your atoms for other reasons.

So that's now what we want to bring in. We can't make a realistic model of evaporative cooling without putting in the constraint of time. And the time constraint is usually determined by losses, by unavoidable atom losses. So now we want to understand what is the speed of evaporation.

So we assume. We truncate. We remove an amount dT dN of atoms above this threshold. And then the question is, how fast can we do it again? How long does it take for collisions to replace the tail of the Maxwell-Boltzmann distribution? But now, you can make an analytic model. I was very pleased when I saw that it is so easy to actually get a precise analytic model of that.

If either in the asymptotic limit that eta is very high, the number of collisions, there is a certain number of collisions which replenish the tail. And you want to know how fast does it happen. But now you can use detail balance. In an equilibrium situation, those atoms will collide live with the bulk of the distribution. And because they are in a highly improbable state, most of the outcomes of the collision will put those energetic atoms back into the bulk of the Maxwell-Boltzmann distribution.

So therefore, the number of particles which arrive in this tail in equilibrium is identical to the number of particles which will leave this tail. And so all we have to do is we have to calculate how many such collisions happen. This is an expression for the fraction of atoms, with the exponential Boltzmann factor, which you can find in this tail.

And those atoms collide with a velocity, which is not the thermal velocity. The velocity is larger by square root eta, because those are fast. So by simply multiplying the fraction of the atoms with the collision rate, we find how many atoms per unit time are removed from this tail.

And in detail balance, it means the same number of atoms is replenished into the tail. So if you now switch to a continuous model of evaporation, where we constantly evaporate the atoms which are produced through elastic collisions with an energy larger than eta kt, then this here is our rate of evaporation.

Since we want to think in terms of time constant, this rate of evaporation is described by a time constant for evaporation. And this time constant is now expressed here by our experimental control parameter eta. Nancy.

AUDIENCE: So when we are saying that the collisions are putting atoms back into the lower velocity states, are we saying that the collisions are more defined than Maxwell-Boltzmann distribution? So when you let the system [INAUDIBLE], it automatically goes into a new Maxwell-Boltzmann distribution, and that's what determines the tail. But then we are saying that the collisions are putting the atoms on the table back into the lower velocities. So the collisions are not [INAUDIBLE] Maxwell-Boltzmann distribution?

PROFESSOR: No. We assume here that the truncation is only weakly perturbing the Maxwell-Boltzmann distribution. And at least the easiest way to figure out how many atoms are produced per unit time, if atoms in this truncated Maxwell-Boltzmann distribution collide, they produce, with a certain time constant, atoms which will populate the tail.

And I can estimate what is this number of atoms which are per unit time fed into the tail by assuming I do not have a truncated Boltzmann distribution. I have a full Boltzmann distribution. And I simply calculate what is the total elastic collision rate of the atoms in this tail.

So in other words, I want to know how many atoms are fed into the tail. I get this number of collisions by saying in detail balance, this number of collisions is the same as the number of collisions in the full Boltzmann distribution which goes backward.

And with that argument, I can immediately write down an expression for what is the collision rate which produces high energy atoms. Think about it. It's subtle, but it's fairly straightforward.

I make the assumption here that eta is sufficiently large, that I can use properties of the equilibrium Maxwell-Boltzmann distribution to estimate those eight constants. And actually, when I found the analytic expression, I could compare to theory, which was much more complicated and used truncated Boltzmann distribution. And in the asymptotic limit of large eta, I was in full agreement with the other results.

So yes, we have an expression now for the time constant of evaporation, how fast evaporation happens because of elastic collisions which populate the high energy tail. But usually, when you have a time constant, you want to express it by another physical time. And the physical time which characterizes a gas is the rate of elastic collisions per atom in a gas.

So therefore, I want to express the rate for evaporation. The rate at which atoms are produced in the high energy tail is a ratio lambda with the time between elastic collisions. And so we realize, of course, that the atoms which have enough energy to evaporate are not produced in every elastic collisions. Actually, there is an exponential factor e to the eta, because it's only a small part of those elastic collisions which happen which produces an high energy atom which can escape.

OK, so with that, we know how many atoms we can lose by evaporation. And this is our expression. But now we have a complete pretty realistic but wonderful toy model to discuss all aspects of evaporative cooling. We have our control parameter eta, which sets the threshold at which atoms can evaporate. And this factor eta determines the two relevant parameters-- alpha, which is the efficiency of evaporation, and lambda, which is the speed of evaporation.

If we set eta very high, following Collin's suggestion, we can put it so high that one lost atom, one evaporated atom, can cool all the other atoms to very low temperature. But we know that this would take too much time.

So in other words, we have a compromise. If you set eta very high, each atom which evaporates provides a lot of cooling power. But high eta means we have exponential slow down in the evaporation rate.

And we have to wait longer and longer, or we never get into evaporation because inelastic collisions and losses has taken its effect. So therefore, it seems clear that this interplay between efficiency and speed is asking for compromise. And this is what we have to realize in the experimental realization.

OK, there is one addition we have to make to the model. And this is the following. We have to introduce losses, losses which do not come from evaporation. It can be losses due to background gas collisions, or losses due to inelastic collisions. So let me just show you how I introduce that.

I mean, I told you that everything is the logarithmic derivative. The logarithmic change of any quantity goes with the logarithmic change in the atom number times the coefficient. And for reasons which become clear in a moment, I'm now not looking at the temperature, or the density, or the phase space density.

I'm really interested in the collision rate, because the collisions rate is what drives evaporative cooling. As long as we have collisions, the cooling can go on. So I want to focus now on how does the collision rate change during evaporation. And during evaporation, what we are changing is the number of atoms, because we evaporated.

So by just putting everything we have said together, I have this expression. I assume that the number of atoms changes as a function of time with the rate at which high energy atoms are produced. The time constant for this was lambda, the ratio between the evaporation time and the elastic collision time times the elastic collision time.

So this is just re-writing what we have discussed so far. But now I introduce that there is another loss rate due to inelastic collision, technical problems and such, which has a time constant of tau loss. And if I now do introduce the famous ratio of good to bad collisions, good collisions are elastic collisions which drive evaporation.

Bad collisions are collisions where atoms are just lost due to technical reasons and inelastic collisions. So if I define this ratio of good to bad collisions, I have now this equation here which tells me how the collision rate changes with time. Just wondering, is a dT missing here or not?

Yes. So there should be a dT. So this tells me how the relative or how the collision rate changes with the function of time.

But now remember, since this is not the derivative of the collision rate, it's the derivative of the collision rate over the collision rate. It's a logarithmic derivative. What we are talking here about it, please add dT to it, we're talking about is the collision rate exponentially growing when this coefficient is positive.

Or is it exponentially decaying when this coefficient is negative? So you realize that we obtain a threshold condition when this coefficient is larger or smaller than zero. And this is called the threshold for runaway evaporation.

So let me just summarize the physics of power loss. The physics of exponential increase and decrease actually means that the experimental situation is often talking about a threshold. If you're above threshold and you get evaporative cooling going, you have a positive exponent.

And it will go faster, and faster, and faster. If this exponent is negative, you have slowed down evaporation, you can evaporate a little bit. But it will pretty much come to a stand still.

So this is why quite often, the experimental realization of evaporative cooling requires to put enough atoms from the laser cooling stage at sufficiently low temperature into a magnetic trap that you fulfill the threshold condition for runaway evaporation. Any questions about that?

So in other words, what we have found out, we have found here an expression for the threshold of runaway evaporation. And it tells us that we need a minimum ratio of good to bad collisions. We may need 100 elastic collision until we have one inelastic collision. And then our ratio is 100. And we will see in a minute if this ratio is 100, if then we can make right hand side in such a way that we run away evaporation.

OK, so the left hand side is the ratio of good to bad collisions is maybe determined how good our vacuum is. What determines the right hand side? Well, we talked about it. Delta is our trapping potential, linear or quadratic. Alpha and lambda depend on the threshold eta. How aggressive are we in setting a threshold in energy for the evaporating atoms?

And what I'm showing you here is the condition. I'm varying the threshold eta. And now I'm figuring out what is the-- if I vary the ratio eta, I call this expression R min. And my ratio of good to bad collision has to be better than that.

So if you just look at the solid curve for parabolic trap, this shows you that you have to pick your parameter eta between five and seven. If you pick eta too fast, evaporation is too slow. And you need a much better ratio of good to bad collisions. In other words, you need a better vacuum, for instance. But if you put eta too low, you're not cooling enough, because you're cutting too deep into the distribution.

You also realize that if you take a linear power law, like a quadrupole trap, you get the dashed line. And the dashed line has a much lower ratio, has a much lower value of R min. In other words, if your vacuum is not good enough and you have losses, in a linear trap, you can still overcome it by picking your eta in this regime, whereas for a parabolic trap, you need at least two or three times better vacuum to get into the runaway regime.

Anyway, this is how you can look at those equations and figure out what is needed to get into evaporative cooling. Collin.

AUDIENCE: Do you need to be in the runway regime to get PEC?

PROFESSOR: Not necessarily. Some early experiments have done, I think, the first experiment of Eric Cornell, I think they never saw this speed up. They had sort of constant evaporation. Maybe the cooling rate was even slightly going down.

So you don't necessarily need the exponential speedup. But you have to be in a regime where at least, if you're not gaining speed, you're not losing too much speed. But, yeah. I mean, you can just take the equations and analyze them and figure out if you're in a favorable regime.

And ultimately, it's fairly easy to integrate those equations as a function of time and have completely realistic models. But what I presented to you here is a simple analytic model. And I used the criterion for runaway evaporation to discuss how do you have to pick your truncation parameter, and what happens if you have a different trapping potential.

So based on those models, you will find out that if you truncated an eta parameter of six, every truncation of the Maxwell-Boltzmann distribution means about 1% loss in the atom number. And after 600 collisions, after 600 elastic collision times, you have lowered the atom number by 100. But your phase space density has increased by six orders of magnitude.

And that means if you have two orders of magnitude in the number and get six orders of magnitude in the phase space density, then your gamma factor, which I haven't really defined it here, but it's the factor which tells you how the phase space density increase. Every order of magnitude in the number boosts the phase space density by three orders of magnitude. And that's regarded as very favorable.

So in other words, we've talked a lot about laser cooling. In laser cooling, the standard laser cooling schemes, you're typically six orders of magnitude away form quantum degeneracy. And this tells you how evaporative cooling can get you there. You should expect to lose approximately a number of 100 in the number of atoms. And that's what it takes you to go from laser cooling to quantum degeneracy.

And you can estimate what your time is by asking what is the elastic collision rate right after laser cooling. If your elastic collision rate is two seconds, and you take 600 collisions to get to PEC, you know that it would take you 20 minutes. And you better work on a vacuum which has 20 minutes lifetime.

Or alternatively, you improve your laser cooling. Or you do some adiabatic compression in your magnetic trap to make sure that your elastic collision time is faster, that you can afford 600 collisions within the lifetime given by other parameters. Any questions?

OK, so evaporative cooling happens in the everyday world. If you have water and you blow at the water, the water evaporates. And by evaporation, the water gets colder. So the water gets colder.

And it has, of course, this process has a lot of common to the evaporative cooling of atoms in an atom trap. But I want to ask you now why is evaporative cooling more dramatic with atoms? We can get the atoms really cold. But I don't think you've ever seen that you can blow at water and the water freezes.

So what is the difference? What is different in evaporation how you encounter it in every day life, and evaporation in the way how I just described it, how we apply it to atoms?

AUDIENCE: Your control of eta. Well, in atoms we can really control eta. But by blowing on it, it's just one level we evaporate.

PROFESSOR: That's very close. We can pick our eta. But even more so, what is constant? Or what is the parameter which describes evaporation in water?

AUDIENCE: Surface tension.

PROFESSOR: Pardon?

AUDIENCE: The surface tension?

PROFESSOR: Some surface tension. But the surface tension always turns into a work function. It tells us if we have water, what is the energy, the work function, for a water molecule to escape?

And this work function, it's an energy, would correspond to eta kt. If the water evaporates, well, the work function stays the same sort. It's sort of electron, whatever fraction of an electron volt or whatever it is. But the temperature gets lower.

And therefore, the number of molecules which, water molecules, which can evaporate, becomes exponentially smaller. I mean, I've sort of lured you into thinking that keeping eta constant is the most natural thing in the world. Yes, it's the most natural world for us who want to-- efficient evaporation of atoms.

But for normal substances, it's the work function which is constant. So therefore, as the system cools, your eta becomes larger and larger. And everything turns into a standstill. So we are actively tuning the work function of our system to sustain a high rate of evaporation.

How do we define eta? How do we select the energy threshold for evaporation? Well, for many, many years, Bose-Einstein condensation was mainly done in magnetic traps. And there were two methods.

One is just lower the magnetic fields. But lowering the magnetic fields is pretty bad, because it weakens the magnetic trap. And if you weaken the trap, you lower the density. And therefore, just because of that, the elastic collision rate slows down.

So what turned out to be by far the superior choice is to remove atoms with our F induced evaporation. So if you have a magnetic trap, you can tune in our F spin flip transition to a certain frequency. But the frequency depends on magnetic field. The magnetic field in a trap depends on position.

So what you're doing is, you're selecting with your RF frequency a certain point in space where the atoms can leave the trap and are transferred by a spin flip RF transition to an anti-trap state, and are then rejected from the trap.

I mean, this is very flexible. You can change the depths of your magnetic trap by just using an RF synthesizer. And you can change the trap depths. You can lower the trap depths without weakening the confinement potential.

I don't have time to go into details, but what I'm showing here is that there is two regimes when you have this magnetic trap. Red is spin up, blue is spin down. And this is your RF transition.

There is a regime where you have strong RF, you should now use dressed energy levels. We've discussed dressed energy levels in the optical domain. These are now dressed energy levels in the RF domain. And there's a wonderful chapter about it in atom photon interaction.

So it really means in the dressed energy levels become something like this. So you have a potential which looks like an inverted W. And you really realize this potential.

But when the RF is weak, you have a certain probability when the atoms go back and forth through the transition region that sometimes they will fall down to the lower state. And this looks more like this. The diabatic picture looks more like that you have a little leak and atoms are trickling down to the non-confining state.

Anyway, there are two regimes. And the experiment is usually somewhere in between. It is not necessary to go to the fully adiabatic. It doesn't pay off to go to the fully adiabatic limit.

So this is how RF evaporation is implemented. But I should say in these days, a lot of evaporation is now done in optical traps. And in optical traps, the method of choice is you just ramp down the optical trapping potential.

Now, when you ramp down the optical trapping potential, addressing Collin's question, you usually do not go into the runaway regime. Because I haven't included that my model assumed we have a constant trap. But if you now add to the model that you are continuously opening up, weakening the trap, you have another effect which makes the exponent for runaway evaporation more and more negative. And I think ultimately, you don't get any runaway evaporation anymore.

On the other hand, optical traps are often more tightly confining than magnetic traps. And you have sufficiently high density to begin with. So therefore, you can tolerate a slow down of evaporation and still reach the destination.

But anyway, I'm now getting more and more into technical aspects. I think I've given you the concepts. But let me just flash you one picture.

This of course, assumes an idealized model where we have only two levels, spin up and spin down. You all know that atoms have hyperfine structure. Sodium or barium has n equals 2.

And if you draw now the stress levels for five hyperfine levels in RF transition, it looks fairly complicated. But the result is fairly beautiful. When an atom in F equals 2 reaches this point, and in the dressed atom picture, the point of evaporation, the point where the atom is in resonance with the RF is the point where this potential bends down.

The atom is adiabatically transferred in a dressed state from MF plus 2 to MF minus 2. So when you evaporate at this point, without maybe you noticing it, you actually do a four photon transition in the dressed atom picture. Again, it's an example where when we teach about the schemes, we can completely neglect about hyperfine structure. And it's just wonderful to see that the actual implementation works as well for complicated atoms then it does for our idealized two level atom.

Final remark. What is the cooling limit for evaporative cooling? When we talked about laser cooling, I derived for you the Doppler limit. And we talked about even improved cooling limits when we discussed sub-Doppler and sub-recoil cooling.

So what is the cooling limit for evaporative cooling? Well, the answer is there is no fundamental limit. There is no h bar or quantized limit there. There is a practical limit, which usually depends on residual heating, on inelastic collisions, and all that. But we have reached an evaporative cooling temperatures as low as 400 picokelvin.

And the limit was not set by anything fundamental. It was more set by our patience. The lower the temperature, the slower the process becomes. And also, by the sensitivity to technical noise and technical perturbations.

Of course, just a final comment which is a segue to what we hopefully do on Friday. For evaporative cooling, I've always assumed that the energy in the trap is continuous. In other words, if you have an harmonic oscillator trapping potential, I've neglected the discrete level structure.

And this is an excellent approximation, because many, many atom traps have dressed frequencies of a few hertz or kilohertz. And at very low temperature, even at nanokelvin temperatures, you populate many levels.

When it comes to the discrete nature of trapping levels, we should use a quantum description of the motion of atoms in the trap. And this is the regime of sideband cooling. Sideband cooling is much more important for ions, charged particles, and for neutral atoms. So therefore, we will discuss sideband cooling on Friday when we discuss ion traps.

Any questions about evaporative cooling?

AUDIENCE: So may be technical, but once you get to the temperatures of picokelvins, how do you maintain-- do you keep cooling to maintain that temperature?

PROFESSOR: Yeah, usually when we reach very low temperature, there is some form of technical heating. And we've often seen when we prepare cloud at nanokelvin temperature, we can only keep it when we allow a little bit of building evaporation.

We've sometimes seen that when we just keep the atoms in an atom trap, they just slowly heat up. But if you keep on evaporating them at a very slow rate, we can maintain low temperatures for much longer. Other questions? OK.

So already pretty late. We just have 10 or 15 minutes left. I was hoping that I could already start earlier today to talk about Bose gasses. Today, I was hoping to talk mainly about Bose gasses, on Wednesday on Fermi gasses, on Friday on ion traps.

We are sort of half an hour behind schedule, and I have to figure out how I make up for it. But let me just start now with Bose gasses. I've taught about this subject several times at summer schools. And what I give you here is a compressed version on it.

I will not talk too much about the ideal Bose gas, because most of you have seen that in statistical mechanics. And I will also omit superfluid hydrodynamics, because well, it's interesting, but it's more special than the other topics.

So what I want to cover here is to give you the main ingredients for the description of weekly interacting homogeneous and inhomogeneous Bose gases. And then as a second part, talk about the superfluid Mott-insulator transition. But well, for those of you who are not working with ultra-cold Bose gasses, maybe some of it sounds like jargon to you.

But there is one sort of overarching concept, which I want to emphasize in the theoretical description. And this is some form of mean field approximation. When we go through the weakly interacting Bose gas, then we go through superfluid Fermi gasses, and we discuss the superfluid Mott-insulator transition.

One theme will be if you have a product of four operators, and that's what you get when you have interacting atoms, you cannot solve anything. And you need a method to go from the product of four operators to the product of two operators. And then you solve a quadratic equation.

And the step to go from four to two is called a mean field approximation. And I want to show you three kinds of mean field approximation, for the one you have often seen for the weakly interacting Bose-Einstein condensate. But then I want to show you mean field approximation for fermions, where it is a pairing field which is a mean field, not your usual mean field energy. And then I want to talk about before that even, I want to talk about the superfluid Mott-insulator transition, where we do a very different mean field approximation.

So maybe you'll realize a little bit by repeating that scheme how theory is done and how you can deal with simple Hamiltonians, but your can't solve them because they contain products of four operators. So there is sort of a [INAUDIBLE] through those three chapters and sort of showing you how you can do interesting many body physics by doing the right approximation.

It also teaches you how many body physics is done. You have an Hamiltonian which you can't solve. And you have to guess the solution and put half of it in an approximation. And once you've done the right approximation, the rest becomes [INAUDIBLE]. So with that spirit, I want to go with you through the Bogoliubov approximation for weakly interacting Bose gas.

I don't think I have to say too much about the ideal Bose-Einstein condensate, because it's dealt in pretty much all undergraduate or graduate text books. There are just two things to remember in terms of a system description. First, whether Bose-Einstein condensation occurs or not depends on the density of states.

And that depends on dimension and confinement. So the fact that you are in a trap, it changes the density of states. And it changes the criterion for Bose-Einstein condensation.

But then in terms of a system description, if you want to describe your Bose gas and its properties, there are aspects of Bose-Einstein condensation which are pretty close to the ideal gas, and others which require many body physics.

What is always close to the ideal gas is the transition temperature and the condensate fraction, because what happens is in almost all experiments, when you reach Bose-Einstein condensation, your gas is to a good approximation non-interacting. Because kt, the transition temperature, is much larger than the interaction energy in the gas.

So before condensation, or at the onset of condensation, your gas is like an ideal gas. And you will only find a few percent corrections to the formula for the ideal gas. So therefore, if you want to know at what temperature do you reach the transition point, or if you're below the transition point, if you're 50% below the transition temperature, what is your condensate fraction, you can simply look up the original formula by Einstein. And it gives you a reasonably accurate answer.

However, for the condensed gas, for the fraction of atoms which are both condensed, those are atoms in one quantum state. For them, there is no other scale than the energy between the atoms. So therefore, for the uncondensed gas, you can get away with an ideal gas approximation. For the condensate itself, we have to put in the many body physics of the interaction.

Well, this slide shows here shadow images of expanding Bose-Einstein condensates. We do evaporative cooling in a magnetic trap. You see the shadow picture of the thermal cloud. And you see the onset of Bose-Einstein condensation as the sudden appearance of, it looks like, a pit in a cherry.

There is a cool down more confined distribution of atoms in the Bose condensed state. I think it should play again. So this gas is pretty much described by an ideal gas where you can put in the g2 function, the onset of quantum degeneracy.

But it looks almost like an ideal gas, whereas in an ideal gas, the Bose-Einstein condensate should be in the lowest energy state of the trap. And I will show you in the next few minutes on Wednesday that we are far away from that. So interactions are negligible for the normal component, for the thermal cloud, but are very important for the Bose condensed gas.

So when you take cross sections through this cloud, or 2-D pictures, you see how the broad thermal distribution turns into a Bose-Einstein condensate. And if you look at such a profile, you can clearly see the normal component.

The normal component can be, with good accuracy, be fitted by a non-interacting model, whereas central peak is a Bose-Einstein condensate, which requires the description I want to show you now. The condensate fraction is shown here. And again, with fairly good accuracy, it follows the description of Einstein, because the condensate fraction is not a property of the condensate. 1 minus the condensate fraction is a property of the thermal gas.

In the spirit of Einstein, Einstein calculated how many atoms can be in the thermal component at a given temperature. So Einstein called it the saturated gas. At a given temperature, you can only keep in thermal equilibrium a certain number of atoms in your gas.

If you have more atoms, they condense into the ground state. This is sort of the statistical description. So therefore, what I'm plotting here is the condensate fraction is actually a property of the normal gas. It shows that the normal gas is saturated, can only hold a certain number of atoms, and the remainder of the number of atoms has to be in the condensate.

AUDIENCE: I thought in the three dimensional gas, it was three halves.

PROFESSOR: Yes, but this is an harmonic trap. And the harmonic trap changes the density of states. Let's talk about the homogeneous Bose-Einstein condensate and weak interactions.

If you write down the Hamiltonian for the interactions, it will appear many, many times. The general way to write down interactions between two particles is you annihilate particles in momentum k and p. And then, they reappear at different momenta.

So one momentum gets upshifted by q. And one momentum gets downshifted by q. This guarantees momentum conservation.

So what I'm showing you here is the elementary process of scattering two particles with momentum k and p scatter, disappear. That's why we have the annihilation operator. And then they reappear to a new momenta.

That's the most general form of a binary interaction. And now we have to make approximations. Nobody can solve this Hamiltonian in the most general way.

So one approximation we make is that since the range of the atomic interactions is much smaller than the distance between atoms or the thermal [INAUDIBLE] wavelengths, we approximate the potential by a short range potential, or delta function.

Well, the Fourier of a delta function is constant. And that would mean in momentum space, this momentum dependent matrix element squared-- yeah, matrix element-- is just constant. So therefore, we can approximate the Hamiltonian by a constant interaction parameter. And then we have the sum over all these creation annihilation operators.

I don't want to go into details of low energy scattering physics, but it is most convenient to describe this parameter by u knot, which is the Fourier transform of the interaction potential. Or very often we parametrize it with the s wave scattering lengths, which is the only relevant parameter for elastic collisions at low temperature.

So with that, we have now a Hamiltonian which has kinetic energy. And here is the potential energy due to the interaction between the atoms. And we have taken this constant Fourier transform u knot out of the summation.

Now, I mentioned already to you that a product of four operators cannot be solved. You need an approximation where you reduce the number of operators from four to two. And then you solve a quadratic equation. And the solution is Bogoliubov solution.

So how do we reduce now this product of four operators to two? Well, when we have a condensate where many, many, atoms are in one quantum state, we can make the Bogoliubov approximation. The Bogoliubov approximation is, well, if the creation annihilation operator in the zero momentum state for the condensate has the following matrix element.

And now you can see, if N knot, the number of atoms in the condensate is large, well, we can neglect the difference between N knot, N knot plus 1, and N knot minus 1. And we simply make the approximation that the operator a knot and a knot dagger is replaced by the square root of N knot.

And then, in this sum of [? our quartic ?] terms, we only keep those terms which have at least two occurrences of the index zero, because the terms which are for instance new occurrence of the index zero don't get this multiplier N knot. So we sort of do an expansion in powers of N knot, and we stop here.

So therefore, now if we make sure that any combination of those two involves the index zero, we factor out N knot. And then we have products of k minus k with dagger, dagger k minus k annihilation operator, or we have mixed term a dagger k, a of k, a minus k dagger, a minus k. But this is what we get.

The next step is purely technical. We want to get rid of N knot and replace it by N. So N is N knot plus the sum of the population or other momentum states. And that's now our Hamiltonian, which still looks complicated. But it can immediately be solved, because all it involves is a quadratic product of operators.

Let me finish this derivation. It takes four to five minutes. I don't think people come in on Mondays. They always come in after us on Wednesday. Is that correct? So let me just continue. I would like to reach the final result with the Bogoliubov transformation.

So what I want to show you is that the moment you have bilinear operators, all you have to do is in essence you have to solve a quadratic equation. And because with all the indices and constant, it looks a little bit complicated. But so let me just say that the structure of this Hamiltonian involves sums which are a of k and a of minus k. k minus k is of course important, because that's important for momentum conservation.

So let me now call a of k a, and a of minus k b. Then this Hamiltonian has the following structure. It has terms a dagger a, b dagger b, b dagger b. But then it has other terms a dagger b dagger plus ba.

Now, let me put it this way. If that wouldn't exist, we would be done. Because an Hamiltonian which has e knot a dagger a is an harmonic oscillator Hamiltonian. It is diagonalized. a and a dagger are just eigenoperators, which create quasi-particles with energy e knot.

So if you could eliminate this term, we would be done. So therefore, let's follow Bogoliubov and say that we introduce new operators, alpha beta. And the alpha beta operators are linear combinations of a and b. Or vice versa, a and b are linear combinations of the new operators, alpha beta.

And since we have a bosonic system, and we think it's a good idea to keep the system bosonic, we require that those new operators fulfill bosonic commutation relations. Those bosonic commutation relations are fulfilled if u square minus v square, u and v are the linear coefficient which express ab in terms of alpha and beta if that is 1.

So in other words, u and v are now two new parameters. One condition for u and v is used up to ensure the bosonic character of alpha and beta. But now, we have a second condition.

We can have two conditions for two parameters. So what we do is we just rewrite this Hamiltonian in terms of alpha beta. And that's what we get. It's a bilinear Hamiltonian where the linear transformation stays bilinear. But now, our second condition which we can impose on u and v is that the prefactor of this cross product is zero.

So then, by using those two conditions for u and v-- and you find the equations in many textbooks. I'm not discussing them in great detail here. We have obtained an Hamiltonian of this kind. And this is diagonalized. We know now that alpha beta create quasi-particles at a certain energy.

So now, by going back, I'm in alpha beta where expressed by a and b. A was a of k. b was a of minus k. I just go back to the original nomenclature. But we have diagonalized the Hamiltonian. So what we have achieved now is our Hamiltonian is written in this following way.

Actually, I've recycled a now. a is no longer a particle in a given momentum state. It's now a quasi-particle. So we have now diagonalized the Hamiltonian for the weakly interacting Bose gas.

And what we have obtained is, well, we have the full solution. Everything you want to know about this system we know. And in particular, we know what are the characteristic excitation energies for quasi-particle.

And this quasi-particle energy gives us the energy as a function of momentum by replacing the parameter u knot by c. c is the speed of sound. I can be rewrite it like this. And you see that it makes sense immediately, because we have now a dispersion relation, which in the limit of high momenta is just the normal kinetic energy.

So the quasi-particle are free particles, whereas at low momentum, when this term dominates, the dispersion relation is linear. And linear means sound and phonons. Bose-Einstein condensation is a low energy phenomena. So you should not expect that you change any characteristics of high energy quasi-particles.

If you wreck a particle in the Bose-Einstein condensate with high energy, it flies out at a high energy particle. But a low energy excitation creates sound waves. So this is what we have found now. And I will show you on Wednesday how it is observed.