PROFESSOR: I have a tiring day today. Just 10 minutes ago, I finished chairing the first part of a meeting of the COA advisory committee. And I gave a talk to them. And right after the lecture, I have to go back. So while they have a lunch break, I have to teach. But that's OK. That's life on the faculty.

Yes. We're talking about applications of the spontaneous light force. We have used optical block recreations and such to learn about the spontaneous force, the stimulated force, in general. And this is our expression for the spontaneous light force. It's pretty much, the momentum of photons is transferred to the atom times to scattering rate.

And what we started last class, and we want to finish today, is we want to discuss applications of this force to three situations which are of experimental importance. One can actually say that slow beams, more or less, in the magneto-optical trap are the key techniques which triggered a revolution in atomic physics. This is really the enabling technology into ultracold atom science.

So I discussed with you, in the last class, molasses. And we discussed what happens when we have two counterpropagating beams. Maybe the key figure to remember is this one. When we have to counterpropagating beams, by just adding up the spontaneous light force of the two beams, we get the red curve, force versus velocity. And the expert eye immediately sees there is a linear slope. And that's a viscous force. This is the viscosity of light interacting with atoms.

Of course, you'd better pick the detuning correctly. If you use the wrong detuning, everything is flipped upside down. And instead of a damping force, you have an anti-damping force.

We then discussed the cooling limit. Any cooling limit is the balance between heating and cooling. And by just setting the heating rate equal to the cooling rate, this describes a steady state. We found the famous Doppler limit.

OK. So that's where I want to continue. But do you have any questions about molasses, the Doppler limit, and what we discussed in the last class?

Well, I know we had discussions about to take two beams which form a standing wave, and take it apart, and calculate the force for each beam. Well, sorry, this is what molasses is about. Molasses is about two beams. And if you want to treat two interfering laser beams in the full generality, there are a lot of other effects. Some of them, we will be uncovering in the next few lectures.

But I'm happy now, when I discuss beam slowing, to give you when there is only one single laser beam decelerating atoms. That is the only situation I know about where you have an exact analytic solution. So I want to present it to you know.

Actually, for a number of years, I presented the exact solution with one laser beam before I presented molasses, where I have to make approximations. But I always felt that beam slowing is more complicated, because you have to go to a decelerating frame and have to add fictitious forces. So some of the intuition is lost.

So that's why I first presented you with molasses, with two beams, for which I need to make approximations. Because the interference pattern between two laser beams has richness and complexity. But now I discuss beam slowing with you.

So this curve here could be the result of the very early beginnings of laser cooling in the late '70s, early '80s. And this is when people had an atomic beam with a Boltzmann distribution, and they had one laser been counterpropagating. But the laser beam had a fixed frequency. And what is this laser beam doing?

Well, it resonantly interacts with the atoms which just have the correct Doppler shift to be in resonance. Those atoms are pushed-- scattered photons-- are pushed to lower frequencies-- to lower velocities. But they are pushed out of resonance. And eventually, the slowing becomes slower and slower. And eventually, things come,

almost, to a standstill. So a laser beam which has a fixed frequency would modify the velocity distribution of an atomic beam in that way.

People called it, already, cooling. And I think some pioneers in Russia who did those experiments were disappointed that they were not honored with a Nobel Prize in laser cooling. This was the first cooling. Because this distribution here is considerably narrower than here. And temperature is nothing else than a measure of the reach of the velocity distribution.

But one can say that those experiments did not trigger a revolution in atomic physics. The revolution came when people eventually figured out how they can narrow down this peak from maybe a hundred or tens of kelvins down to a million microkelvin. So it's clear what is needed. If you want to do more slowing, you have to make sure that you are talking to the other velocities, and that you eventually sweep many velocities into one peak-- clear out the velocity distribution, collect all the atoms, and put them at a monochromatic velocity.

So for that, you can say, well, if you use a comb of frequency, or a broad spectrum of frequencies, you could clean out a larger area. And, indeed, there are techniques like this-- white light slowing or diffuse light slowing. But since those techniques are not anymore of practical importance, I don't want to discuss them.

Let me talk about the simplest technique-- at least, conceptionally-- which is chirped slowing. And this works like follows. You start with a laser beam which is initially in resonance with the velocity group. And as the atoms slow down, you change the detuning of the laser that the laser stays in resonance with the atoms.

So it is a pulsed technique where you have a continuous atomic beam. You start with a laser and say, now I start slowing. And the laser is following a group of atoms as they are decelerated. So you can say it is like the atoms are riding the surf. The laser is always following the atoms as they change their velocity.

And what is, of course, necessary for that to happen is that the chirp and the deceleration of the atoms are synchronized. You will see that in a moment, when I

put this idea into equations.

There is another technique which is used in many labs-- actually, in all my labs and Martin Zwierlein's lab. It's called Zeeman slowing. And this is discussed in your homework. It is a CW version of chirped slowing. And a lot of you know what I mean.

The physics which goes on there is almost identical. But for conceptional reasons, it's easier not to use magnetic fields and such. Let's just keep it simple. Let's have one laser beam talking to a two-level atom. And the only other ingredient is the Doppler shift.

Any questions about-- so I've surveyed, for you, beam slowing. And now I want to give you the exact description for one typical example of a slowing technique. And a lot of the results we find in chirped slowing apply to other slowing techniques. OK.

I know I've taught it several times. And there is the conceptional point that you can only describe chirped slowing well when you go to the decelerating frame of the atoms. You want to describe it in the frame of the atoms which are continuously being decelerated. So the way I can introduce it is the following.

I take the maximum-- I take the spontaneous light force. This is exactly the expression we have discussed several times. And now I'm saying we decide that we want to decelerate atoms with a deceleration, a, which is negative, because we slow them down. Of course, if you pick a too large that it would require more than the maximum spontaneous light force, you will not find a solution. But let's simply assume we want to decelerate atoms with a deceleration a.

Well, what happens is this deceleration, a, requires a certain force. If the intensity of the laser would provide more force on resonance then it's not going to work. Because we want to decelerate the atoms at this deceleration. So therefore, the atoms, if they are decelerated, have to be at a detuning, delta prime, in such a way that the spontaneous light force is exactly providing the acceleration we decided we want to have for our atoms.

4

In other words, this detuning-- in the frame of the atom, the atoms are on resonance when the light force is exactly providing the force to give you the acceleration a. But if you have more laser power than absolutely necessary, in the frame of the atoms, the atoms have to be detuned in such a way that the force, with the detuning, is exactly providing the acceleration we want.

So that's the easiest way to explain it. We have to set the stage. We have to introduce this nominal detuning.

OK. So you can say that this detuning is just a definition. And now we want to describe the slowing process. So we have said we want to start with atoms at an initial velocity-- which we'll actually cancel out in a moment. But let's do it in a systematic way-- and afterwards, we want the velocity to decrease, linearly, in time.

I simply defined this parameter, delta prime, to be a solution of this equation. So this is a well defined quantity. And the laser detuning-- which, I've broken my laser to change its detuning as a function of time-- is now this nominal detuning minus the Doppler shift due to the velocity of the atoms.

I'm really solving the problem backwards for you. I first decide which should be the velocity trajectory of the atoms-- how should the atoms slow down. And then, I define what is the detuning which will result in that. That's the easier way to understand it. Of course, if you pick this detuning, the atoms will exactly do that. OK.

But now what we're doing is the following. This is sort of the acceleration and the velocity we want the atom to have. But there may be some atoms in the beam which do not have this velocity. They deviate from this velocity with v prime. So the goal of cooling would now be to reduce v prime to 0 to make all the atoms follow the trajectory we have designed for them.

OK. Many words, many definitions. Now we can simply substitute what we have defined into the equation. And we get a wonderful result.

So we take equation one, which was the equation for the light force. And we put all those definitions, which I went to great lengths to explain to you, into it. And we get

this result. It's just that the detuning has been expressed by delta prime, and the velocity by this difference velocity v prime. It's pure mathematical substitution.

But now comes the important point that we want to describe everything, now, in the decelerating frame. The atoms are meant to decelerate with a deceleration a. And you know, from classical mechanics, if you describe something in a decelerating frame, you have to add a fictitious force.

This fictitious force, since the acceleration is constant, is a constant force. But I'm expressing it, now, by the other parameters. We've chosen our acceleration. We've chosen the detuning, delta prime, that-- everything is sort of connected. And this is just a way to use the correct-- the most easiest units, or the correct parametrization, for this constant force. So I have done-- and the important thing for you is, this equation is exact. It's exact for an arbitrary velocity v prime of the atoms.

So let me just summarize. I've taken the spontaneous light force. I've given you a bunch of definitions, what delta prime is. That's how you program your laser. But then, in the decelerating frame, we need a fictitious force. And this is nothing else than a mathematically exact rewrite of the spontaneous light force in the decelerating frame.

But now what happens is-- we had a decelerating-- we had a light force where, because we assume the laser is counterpropagating, there's a minus sign. The fictitious force has a plus sign. And that means that we can-- and if v prime is 0, the force is 0. So we can now make a Taylor expansion for small v prime.

So we have, again, a viscous force, a friction force. And it turns out that this friction coefficient for beam deceleration is exactly 1/2 what we got before, for molasses. But it's clear we have one laser beam. Whereas in molasses, we had two laser beams.

If you ask, what is the heating, well, if you have one laser beam and not two laser beams, the heating described by the momentum diffusion coefficient is also just 1/2. One beam heats exactly half of what we got for two beams. And therefore, the

temperature, which characterizes the width of the velocity distribution in the frame of the decelerating atoms is exactly the Doppler temperature we had before. Because the two factors of 1/2 cancel.

So now I have given you-- and I hope you take some pleasure in it-- an example where the laser cooling for a two-level system is exactly described. There's only a single beam. And the Doppler limit is the exact solution. I've not made any approximation other than assuming I have a two-level system for which the spontaneous light force has the expression we've derived.

OK. I know, even with my explanation, you have to read through it once or twice. Because I made a few substitutions. And you really have to digest them at your own pace. So let me give you a summary of what we have done graphically.

One laser beam scatters light. And it is this Lorentzian. And it's a negative Lorentzian because we have a counterpropagating beam. We push the atoms into the negative direction, they fly in the positive direction.

But by going-- by chirping the laser, by going into the decelerating frame, the force doesn't go to 0. I had to add an offset. And this offset is the fictitious force due to the transformation of the decelerating frame. And now you realize that this force-- and it is the correct force in the decelerating frame-- has two 0 crossings. One is a stable point, and one is an unstable point.

And the stable point says, if an atom happens to be too fast, it gets a correcting force. And this is like a lock point. It has a stable lock point. All the atoms which did not have the correct velocity we wanted them to have, they are sort of sucked into this lock point, and eventually, will pile up at a huge peak at this nominal velocity. And the width of this peak is simply described by the Doppler limit.

Before I take your questions, let me show you what happens in the two frames. This is v prime. This is the velocity in the decelerating frame. And let's say our initial velocity is this. So what happens now is all the atoms within the atoms are decelerated towards v prime equals 0. The friction force puts them to v prime

equals 0.

But since we are in a decelerating frame, in the absence of the laser, the whole Maxwell-Boltzmann distribution is just shifted-- not because the velocity of the atoms changed-- because we are in a decelerating frame and we are accelerating away from the atoms. And what I just described for you is the following-- that, at v prime equals 0, we have our friction force. And the atomic distribution in this frame is now pushed through. And everything which goes through this resonant region is sort of collected. And you have a peak which is piling up at v prime equals 0.

Probably, it's easier-- but mathematically, I needed the decelerating frame-- if you use the following picture. And just think of this beam slowing I explained to you. If you just switched on a laser here, you would just burn a hole into the velocity distribution and pile up the atoms.

But now you sweep. And when you sweep, you take the peak, push it further. It gets a bigger peak here. Push it further. So eventually, what you're doing is, by chirping the laser in the correct way with the correct laser power and all that, you're just pushing the atoms all out. And eventually, they are piled up in a narrow distribution which is defined by the time when you just switch off your laser beam.

So one chirp of the laser, from an initial detuning in the lab frame to a final detuning in the lab frame, will just sweep out the whole Maxwell-Boltzmann distribution. And if you suddenly switch off the laser, you have frozen in a velocity distribution which is exactly described by the Doppler limit.

So what is easier to understand-- molasses with two beams, or beam slowing with a single beam? I think molasses is easier, because everything happens in the lab frame. But think about it. It's a nice example for which no approximations have to be made. Questions? Yes, Nancy.

AUDIENCE: It's not really obvious to me how, in this mathematics, we are choosing a time-dependent detuning as opposed to a time-dependent deceleration.

PROFESSOR: The way I presented it logically, we said we want a certain acceleration to happen.

And then, we provide our laser that k/kv is the Doppler shift. ka is how the Doppler shift changes as a function of time. And we chirp our laser exactly in frequency with a chirp which is ka.

So in other words, I started to define a constant a. And that would require that the laser is now chirped in a linear way. I could have started the other around and said, hey, let's assume we chirp the laser in a linear way. And then, we do corresponding substitutions. And the result would be the same.

AUDIENCE: No, actually, like if we did not want to have a detuning which was time-dependent, would we have gotten-- we still have a force. And that force would have been just changing.

PROFESSOR: Well then, it's more complicated, yeah. Then, you have to integrate a differential equation. But what happened is the following. What I've done is, by defining the acceleration and then the chirp, I have the situation that, in the reference frame of the decelerating atoms, the detuning is constant.

So usually, you go to a decelerating frame if something else simplifies. And what simplifies is that, in the frame of the atom, the laser detuning is constant. So you can also say the frame in which the atoms decelerate is exactly the frame in which the chirp has disappeared. Because the decelerating frame compensates the chirp with a linearly varying, time varying Doppler shift. OK.

But the physics, I think, is really-- the moment you have a force which has a 0 crossing, it's a lock point. You determine the slope. And alpha, together with the momentum diffusion, gives you the final temperature. So I hope you enjoy that it's actually the same physics-- exactly the same physics-- we discussed, with some approximation for molasses.

OK. We had a discussion last time about, we have the force, we find exactly what happens to the atoms. But we may have to think about it in different ways to figure out where the kinetic energy goes from. And I think we mentioned it already, in the last class, that what happens is you have a red-detuned laser, which is, the photons

are absorbed. But the atom is scattering light in all directions.

And so if you go at 90 degrees, you don't have a linear Doppler shift. If you go forward-backward with the same probability-- a symmetric pattern-- the Doppler shift averages to 0. So on average, you emit photons which are on resonance. And therefore, every time you scatter a photon, you radiate away the energy which is equal to the Doppler shift.

That means, initially, when the Doppler shift is huge, each photon transports away a lot of energy. And the slower the atoms are, the less a photon transports away. But that makes sense. If you have a huge velocity and you subtract h bar k, the momentum of the photon, the kinetic energy is quadratic in velocity. You know, h bar k, the momentum transfer of a photon, takes out more kinetic energy the faster the velocity of the atoms is. So everything makes sense and is consistent.

OK. Let's go from the-- we started with two beams. We just discussed one beam, chirped slowing. And let's go now to what most of you use in the laboratory-- namely, six beams. You want to cool in x, y, z. So you have one-dimensional molasses with two beams times 3. So in the limit that the intensity is low, you just add up all the forces. You say each beam has a spontaneous light force. And you add the six spontaneous light forces up.

At low intensity-- and you have to take my word for it-- you can ignore all the interference effect between the beams. They can actually lead to very interesting effects. I will give you at least a taste of what will happen-- actually, a pretty good taste-- what happens at higher power. We need the traced atom picture. And we'll discuss how the stimulated force can be used for cooling.

But if you don't have a two-level system, you have sigma plus sigma minus transitions and the beams interfere, you get spatially varying polarizations of your interference patterns, and then you have much [INAUDIBLE] physics. But in the simplest case of a two-level system, there is a regime at low intensity where you can simply take the equations we got for molasses and apply them to XYZ.

I put, on the website, a nice handout, a review paper by Bill Phillips, who is discussing 1D molasses versus 3D molasses, and if you enjoy it, you can read how numerical factors of 1 or 2 or 3 now appear in the friction coefficient, in the heating coefficient, and how they affect the final temperature. But you got the full [INAUDIBLE] concepts and the correct idea for one dimension molasses.

Now, let me give you an outlook for what we will be doing next week. Next week, we want to understand the following-- what happens to the friction coefficient when you increase the power? Well, you remember when we have red-detuned light, then we had one Lorentzian, another Lorentzian. The slope was negative, and negative means positive friction, we cool. And if you put more power into it, the linear slope increases, and for low power, the friction coefficient is proportional to the laser intensity.

But now look what happens. If you approach saturation, then the friction is not only saturating, it suddenly changes the sign, and this is qualitatively new. And you should a little bit wonder about it. Wasn't the reason why we have cooling that you have a laser which is red-detuned? And the photon which is spontaneously emitted is on resonance, and this is the energy difference. But now I'm telling you, at high intensity, it's not the red-detuned laser which cools, it's the blue-detuned laser.

So you should wonder about it. There is really something new to be learned. So I want to motivate you to follow me today and next week through an alternative description of what light does to atoms using the dressed atom picture, and in the dressed atom picture, we will very naturally understand why now blue-detuned light provides cooling. But right now, for you, this should be a mystery, but it should be something where you say, hey, still something qualitative is missing in our understanding of light forces.

OK, I will not be able to explain it [? physically ?] to you. I first have to introduce the dressed atom picture, and I will do that it in 10 minutes or so. But what I want to do here is if I'm telling you something goes terribly wrong and we completely miss even the sign of the force, blue-detuned light is now cooling, not red-detuned light, I at

least want to tell you what went wrong. Because haven't we done an almost exact discussion of light forces?

We used the gradient of d dot e. We used for d the steady state solution of the optical Bloch equations. Haven't we done everything right?

Well, we did approximations. And what happens is, and this is the reason for those qualitatively different things, that we should, in order to get those effects, not use, for the dipole moment, for the off-diagonal matrix elements of the density matrix, the steady state solution for optical Bloch vector. We have to acknowledge that it always takes a finite time to reach steady state.

For instance, if he takes, let's say, a spontaneous emission time, if the atom goes through a standing wave, it goes from high intensity to low intensity, it takes about one spontaneous emission time. When you are highly excited, but now you go through a node of the standing wave, the steady state solution is [INAUDIBLE]. At the node of the standing wave, there is no light. The steady state solution with no light is you better go back to the ground state, but that takes something on the order of the spontaneous emission time.

So therefore, when I am excited because I'm in the middle of the [? pied ?] laser beam and I go fast into the dark region, I'm still excited because I didn't have time to get rid of my excitation. In other words, the internal degree of freedom, the population of grounded excited state, will lag behind the steady state solution by a lag time which is a spontaneous emission time. And that would mean if I multiply with the velocity, it's pretty much that there is a spacial displacement. I'm always a certain distance away from the steady state solution.

And actually, this point, to introduce the lag time, we with do it in a wonderful, physical way when we have the dressed atom picture, but it can be done with optical Bloch equation. It's just a real pain because the optical Bloch equation is sort of like a matrix the equation. It's completely obscure, how it is done, but in order that you convince yourself, I posted the original paper from Gordon and Ashkin where this was done. You read the math. You understand the result, but you don't

understand the derivation, the physical picture behind it. The dressed atom picture will give us a lucid explanation for exactly what a lag time does and why this results in colling.

But anyway, this is the physics we have done wrong. At low power, it's OK. At high power, we really have to take care of effect. And if we would take care of effect-- And as I said, we don't need the dressed atom picture, we can fully do it within the framework of the optical Bloch equation, but it's messy. --we would actually get out, and you see that in the references, the result that the friction coefficient in a standing wave at weak power is exactly two times the friction coefficient for a travelling wave, and that justifies that we describe molasses by the spontaneous light force of two travelling waves, completely ignoring the standing wave.

But the same treatment of the optical Bloch equation will then, eventually, give us the result that alpha changes sign at high intensity. But as I said, I want to describe that after we have introduced the dressed atom picture. Questions?

OK, we have done molasses in one dimension. We have done beam slowing. We just did a short discussion about molasses in arbitrary dimensions. The one technique which is missing now is the magneto-optic trap. Let me just spend five minutes to discuss the background behind the magneto-optic trap because for many of you, I think you were just born when the magneto-optic trap was invented.

Well, I joined the field in 1990, and that was three years after the magneto-optic trap has been invented. And many people say if there is one most singular development in laser cooling and such, it's molasses and the magneto-optical trap. So let me tell you why, the magneto-optical trap, you shouldn't take it for granted.

It was actually in 1983 that Ashkin-- Was it Gordon or Ashkin? I think Ashkin, Art Ashkin. --that he presented an Optical Earnshaw theorem, and it is nicely described in Bill Phillips' Varenna notes which I have posted on the website. And it's pretty much a proof that you cannot build a trap using the spontaneous light force. So in other words, this is a proof which people took as a correct proof that something like the magneto-optic trap could not exist.

So what this is proof like? Well, and that sort of teaches you what was the mental concept or the mental barrier people had to break through before they invented the magneto-optic trap. The proof is the following.

If you think that you have beams-- I mean, a model for the spontaneous light force is you have photons and the photons are sand blasting the atom. Each photon is like a grain of sand, and the atom is in scattering descent in all directions. You can really make this completely classical picture of the spontaneous light force. Almost.

If you do with that you would say the spontaneous light force is proportional to the stream of photons which is described by the Poynting vector. But now, you have in electromagnetism a continuity equation which says the divergence of the Poynting vector plus the change of energy density in a certain volume has to be 0. So if you are in steady state-- If you have an arrangement of laser beams and the energy density of the electromagnetic field is at steady state, that tells you that the divergence of the Poynting vector has to be 0.

Well, if the divergence of the Poynting vector is 0, using this equation that says the divergence of the spontaneous light force has to be 0. But if the divergence of the force has to be 0, that means you can not have a force which is inward from all directions because such a force, a trap needs a force which is inward. A trap requires a negative divergence.

Do you see the wrong assumption in this proof, or did I convince you that the magneto-optic trap is not possible because this proof would say you can not make a trap out of the spontaneous light force?

AUDIENCE: Well, you through out the change in energy density, so that's--

PROFESSOR: But that's OK because if you switch on a laser beam, light travels at the speed of light, and in 10 to the minus 10 seconds, in picoseconds, you reach steady state. And you know when you operate the MOT, you're not switching on and off your lasers at picoseconds. It's not that. Yes, [INAUDIBLE].

AUDIENCE: --the spacial variation of c.

PROFESSOR: Yes, exactly. The force is not just proportional to the Poynting vector. This would be true if I had grains of sand, but photons have to fill a resonance condition.

You can maybe assume I'm an atom, but I have an spatially varying magnetic field around me. At the one magnetic field, I feel the force of the laser beam because I am resonant. At another magnetic field, the laser beam just goes through me because I'm not in resonance. So therefore, the proportionality between the Poynting vector of the laser beam and the light force ignores that the light force is resonant, that the light force depends on polarization and such.

So the solution is the divergence of-- If the c parameter is spatially dependent, then we can have a divergence, and then, we can do trapping. And it was actually Dave Pritchard in '85 or in '86 who wrote a really influential paper. He said the Optical Earnshaw theorem can be circumvented. We can have a spatially dependent factor, and what can be spatially dependent is saturation. If you focus a laser beam, you have a spatially varying saturation parameter. There can be some optical pumping. There can be light--

So he discussed several possibilities how the Optical Earnshaw theorem was circumvented. At this point in 1986, he didn't realize that it can be done with a magnetic field gradient. When he talked about possible ways of circumventing the Optical Earnshaw theorem, Jean Dalibard, who many of who know, who's a famous researcher in our field, had the idea that applying a magnetic field gradient could work. But nobody thought, at this time, that it would work in three dimensions, so the fact that it works in three dimensions is almost a miracle.

So when the idea was born, all those schemes, and also the original idea of applying a magnetic field gradient, were conceived as one dimension scheme. Most of the schemes can not be easily generalized to three dimension, but the MOT scheme works as well in three dimension as it works in one dimension.

Anyway, so this is sort of the background. Since I was a postdoc when there were

only two or three groups who were operating a magneto-optic trap, I put together, with my own hands, the first magneto-optic trap which was ever used at MIT. I mean, it's ancient history, so I wanted to share a little bit this history to you.

But let me now simply describe to you how the magneto-optic trap works. Let me describe it for an atom which we are not using in the lab. We all usually use atoms which have hyperfine structure. But let me just for pedagogical simplicity assume that we have an atom which has an s to p transition, and in the p transition, we have the three different magnetic [? quantum ?] numbers. So it's a simplified model. And what we're doing right now is we are applying molasses, two beams which act as molasses, but there is one difference now. They have different polarizations, sigma plus and sigma minus, and in addition, we apply a linear magnetic field gradient.

So what happens now is the following. The linear magnetic field gradient does Zeeman shifts in the opposite direction to the n plus 1 and n minus 1 level. Now, the transition to m minus 1 is done by sigma plus light. The transition to m plus 1 is done by sigma minus light.

So what happens now is if you have an atom and it moves away from the origin, it will eventually come closer and closer in resonance with the sigma plus light, but the sigma plus light comes from the left side so the atom will be pushed back. If the atom moves out hear, it moves further and further away, in resonance, from the sigma plus light, but it moves now, in resonance, over the sigma minus light, and the sigma minus light is pushing the atoms back. So this geometry has an inward force from both directions and therefore, works as a trap.

So let me now describe it mathematically which is very, very simple. What we have to do is we just have, like in molasses, two laser beams two times the expression for the spontaneous light force. And you'll remember we got the molasses equation by taking the two laser beams and putting in the Doppler shift which is a velocity dependent shift.

Well, all we have to do is, for those two laser beams, we have to put now a spatially dependent Zeeman shift which varies linearly with position. So therefore, all we

have to do is we sum up the total force as the spontaneous force for the right and for the left laser beam, and in addition to the molasses effect, we put now in the spatial dependence.

So we don't have to do the math again. We did already a Taylor expansion in v, but let's now just do a Taylor expansion in the blue expression. And then, we find that there is a linearly restoring force both in velocity space and in configuration space which actually tells you something interesting. The more friction you have, the larger is the restoring force. So the trapping and the cooling are really connected, at least in the simple model of the spontaneous light force.

OK, so if we have this kind of force, a linear force in velocity, a linear force in position, this is the equation of motion of a damped harmonic oscillator. You have a spatial force constant, k, this linear force which is linear in displacement. So this is a damped harmonic oscillator.

Usually, you don't see atoms sloshing in your magneto-optic trap like in an harmonic oscillator because if you look through the [? typical ?] numbers and by just doing the Taylor expansion of the spontaneous light force, you can immediately get an alluded expressions for those numbers. In almost all cases, the harmonic oscillator is overdamped. So that's how the magneto-optic trap works.

However, the magneto-optic trap is not so simple. I would actually say nobody fully understands the magneto-optic trap because we have multi-level structure. We don't have just one count level. A sodium and a rubidium atom has-- f equals one and f equals two-- has eight different hyperfine levels.

So there is not one level and we can excite the atoms. There are eight different levels. And those laser beams to optical pumping among those eight levels, and you have to solve a rate equation for eight different levels. And the amazing thing is it still works. It still works pretty much, at least qualitatively, as well as I just described to you.

So it's a combination of optical pumping and Zeeman shifts. It even works better

than advertised because as I will tell you in the next few weeks, you can reach temperatures-- Because of the multi-level structure, you reach temperatures which are much colder than the Doppler limit. It's called polarization gradient cooling.

It's just if you have laser beams and you have atoms with hyperfine structure, they are colder than the Doppler limit. First, nobody understood it, and then, Bill Phillips figured it out. And that he figured out why atoms get colder that was the main point emphasized by the Nobel Committee when he got the Nobel Prize for laser cooling. Of course, he had already invented Zeeman slowing before, so he had a few other things under his belt.

But anyway, the fact that the MOT works even better than advertised is another miracle, and this is related to the hyperfine structure. And finally, the fact that you can simply extended it to three dimension, is-- I said the first, the second --maybe the third miracle. It also works.

Because when you look at the field of two coils, a so-called Anti-Helmholtz field, there are some directions where the magnetic field is not radially outward. This is a much more complicated magnetic field. It goes outward here, outward here, but in between the quadrupolar field, it sort of curves. There are some directions where the magnetic field is not radial, and you would say, hey, the one dimensional scheme is [INAUDIBLE] [? set. ?]

Actually, before the magneto-optic trap was demonstrated, Dave Pritchard's student, Eric [? Rob ?] did calculations what happens in these diagonal directions which have no resemblance to the one dimension scheme. The solution is in those directions, the trapping form is a little bit weaker, but the big picture is it doesn't matter. So there was an idea conceived in one dimension for a simple atom structure, and this idea worked even better for real atoms and in three dimensions.

OK, in your homework assignment, you will put in some numbers which show what happens if you take a vapor, an alkali vapor at a very low density of 10 to the 8 per cubic centimeter, pressure of, typically, 10 to the minus 8 [? Tor, ?] and you arrange six laser beams around it. Well, those laser beams have a restoring force in velocity.

That means they cool and increase the phase space density over your sample by 8 orders of magnitude.

Those laser beams have also a restoring force in the spatial dimension. Therefore, you get a cloud which has 3 orders of magnitude higher density. So the temperature goes down by 8 orders of magnitude which means the phase space density is enhanced by 12 orders of magnitude. So if you're talking about the goal of reaching Bose-Einstein condensation which requires a phase space density of 1, just having a dilute vapor and these six laser beams with a little bit of magnetic field, increases your phase space density by 15 orders of magnitude.

So this is why molasses and the MOT were a real milestone. Now, with 15 orders of magnitude, you are just 4 or 5 orders of magnitude from Bose-Einstein condensation, and we will discuss, in a week and a half, what techniques were needed to bridge the gap between those temperatures and densities and quantum degeneracy of gasses.

So that's the story of the magneto-optic trap. I'm also telling you that because I'm proud that a lot of the pioneering work in this regard happened at MIT. It was Dave Pritchard who realized at the Optical Earnshaw theorem can be circumvented, and it was him, in a collaboration with Steve Chu, that they together realized the first magneto-optic trap, and this was in 1987. Questions? OK, I think simple applications of the spontaneous light force.

So we have time to talk about the dressed atom. Now, I was waiting until this point to introduce the dressed atom picture for you because the dressed atom picture is the natural picture you want to think about, for instance, for the Mollow triplet, and we had already some discussion about it. But it is also the natural picture you want to think about it when you describe the stimulated light force. And I mentioned already, it is the dressed atom picture which will help me to give you the physical picture behind if you have strong laser beams, why blue-detuned light cools and not red-detuned light. So you should be motivated now to learn about the dressed atom picture.

But first, I have to say it's not anything new, what I'm teaching you. It has exactly the same physics in it as the optical Bloch equations. It's just that the optical Bloch equations are often more complicated, and they don't provide the clear insight. And the reason is the following.

Often when you describe a physical phenomenon, you can choose one basis set, or you can use another basis set. But if you're in the wrong basis, you get the correct results, but in the wrong basis, you just can't describe the physics easily. So in other words, with the dressed atom picture, I want to introduce dressed atom states for you which are no longer the ground state and the excited state. I want to treat the atoms dressed up by the laser beam, which is one mode of the electromagnetic field.

So the atoms, together with this one laser beam, will be partially excited, partially in the ground state. And this partially excited atoms, this is the dressed atom basis. But if you think in terms of those states, you have a much easier way to formulate certain physical processes, exactly processes where you have high laser power. At low laser power, the dress states are identical to the naked states to the ground and excited states. But at high laser power, at infinite laser power, the two dressed states are the symmetric and anti-symmetric equal superposition of grounded and excited state. So that's what we want to describe now.

So the difference to the Optical Bloch Equation is that we are now really looking for the combined state of the atom in the electromagnetic field. So what we want to do is-- this is sort of our picture that we have. This is our Hamiltonian. We have the atom interacting with the vacuum through spontaneous emission with a reservoir of empty modes. But now, it also interacts with the laser beam.

And you remember when we did the master equation, we looked sort of at this. We derived a master equation for this. And then we added in the unitary time evolution due to laser beam. But what we want to do now is we want to first take care of the atom and the laser beam almost exactly in an exact diagonalization.

And once we have re-diagonalized that, we have now eigenstates which we call the

dressed states. Then we allow not the ground and excited state. We then allow the dressed states to emit into the vacuum. And this is the formulation we want to obtain now.

Let me make one comment before I go through a few simple equation, and that's the following. If I want to treat this system almost exactly, I have two choices. I use a strong laser beam.

So I could, as we have discussed several times, simply use the time dependent external electric field. The laser beam is E0 times cosine omega t. And I just introduced that in the Hamiltonian.

And you can get everything I describe now for you out of this Hamiltonian. However, there is an alternative way. And this one is I simply assume that the laser beam is described by the fully quantized field.

And I use the fog state. I just assume N, the number of photons, is very big. And this is, at least for that purpose, a valid description of the laser beam.

So what happens is you really want to look through that. I've really learned a lot by comparing the purely classical treatment of the laser beam and the fully quantized description of the laser beam. The results are the same. They are Absolutely identical.

And what you will see is, by a bunch of approximations, we reduce the total Hamiltonian to-- I've seen it many times-- a two by two matrix. And whether you start with the classical field or whether you start with the fully quantized field, if you make the right approximation, you get the same two by two matrix.

But there's a difference. This Hamiltonian is time independent. a dagger a, time independent Hamiltonian. If you take a classical field, which is cosine omega lt, you don't have a time independent Hamiltonian. You have a driven system, and your wave function has an external time dependence. The laser beam is driving the wave function.

So in other words, you have two choices when you want to describe how an atom interacts with a strong laser beam. You can use the fully quantized field, which is conceptionally a little bit more complicated, because we've quantized the electromagnetic field. But then everything is time independent, and everything I do here is really easy. I can go over it now in five minutes, and you understand everything.

Well, in your homework, I've already asked you at an early homework to look at simply the two level system driven by classical field. And in your next homework, number 10, your last homework, I'm asking you now to interpret the solution with the classical field. And it's actually a little bit a challenge to take the time dependent wave function and read the physics out of this time dependent wave function, which corresponds to the physics we derive in the dressed atom picture.

But you have two ways. One is the classical field. Then the wave function, the Hamiltonian, is time dependent. And the wave function is time dependent. Or here, I start with the quantum field and everything is time independent. But the solutions are identical.

Well, one is a time dependent wave function. One is time independent. But if you interpret the solution in terms of what happens to the physical system, you will find exactly the same.

So in other words, I decided I do the easy part here, which is a fully quantized part. And this is probably what will get stuck in your head. But since I think it's an important pedagogical exercise, you have to go through the homework and do the time dependent part.

OK, so we want to know exactly-- well not-- we want to do the correct approximations. And after some approximations, exactly solve what the atom does with one laser beam. So the laser beam is described by the standard Hamiltonian for single mode light. The atom is described by, well, excitation energy omega 0. And b is the excited state.

So if you don't have any light atom interaction, the eigenstates are just product states, an atom in the excited or the ground state. And N is the photon number. [? Collin? ?]

AUDIENCE: You said that it's usually easier to understand physics in a dressed atom picture versus optical block equations.

PROFESSOR: Yeah.

AUDIENCE: Can you think of-- is that a general rule? Or is this just like it's only applicable in certain areas? [INAUDIBLE]?

PROFESSOR: You're really asking an expert question. And I talked to Jean Dalibard and Claude Cohen-Tannoudji, who invented the dressed atom picture, when I wanted to understand what happens when we have a blue detuned lattice. And I used the dressed atom picture and I got confused.

So what happened is the following. If you pick your right bases, the description is simple. But simple means if all the physics is your atom is in one dressed state or the other dressed state, and it does transitions between the dressed states and that is easy, and this happens in the limit of large detuning.

But if you're in the regime where coherences between the dressed states matter. In other words, the density matrix in the dressed state basis has now off diagonal matrix element. The physics becomes complicated again.

So in other words, what I'm telling you is you are the best description happens when we have a two level system, we always need a two by two matrix. But if you pick the description where the solution is diagonal, if you have to deal with coherence, using coherence is coupled to population, it's complicated. And there's some cases-- one is the blue detuned lattice, which I encounter where you have to consider the coherences to get the correct answer.

And usually, also when the laser beams are weak and on resonance, the dressed atom picture is not providing any advantages. So you want to use the bare picture

for weak excitations. And yes, in the resonance case, the dressed atom picture, if you're very close to resonance, the dressed atom picture gets more complicated.

But to give you the answer in something we have discussed, you know already that when you drive an atom with a laser beam, you get the Mollow triplet. And the Mollow triplet shows you-- I hope we can finish that today-- shows you the energy splitting between the dressed energy levels, between the dressed energy levels.

But if the Mollow triplet is not resolved because the generalized Rabi frequency is small, then you're not resolving your dressed atom levels. And then the dressed atom level is not providing a big advantage. For certain problems, the dressed atom picture will not help you. But I can reassure you, and this is what I learned from Jean Dalibard and Claude Cohen-Tannoudji, if you really look at the full matrix with the off diagonal matrix append and everything, you will always find exactly the same solution.

Anyway, we were just writing down product states of the atoms and the eigenfunctions of the Hamiltonian for the light, which are fog states, number states with N, L N plus 1. But the important thing is now the following.

In the limit of small detuning, the ground state with N plus 1 photon is almost degenerate with the excited state with N photons. And we call these two states a manifold. The manifold E of N. And N is the photon number in the excited state.

So what happens is, if you look at all those states, they are grouped into manifolds. And the energy separation between two manifolds is, of course, I go from N to N plus 1, the energy of a photon. So now I introduce the atom laser coupling in the usual dipole approximation.

And what happens is, this dipole approximation, this dipole light atom coupling will couple, of course, ground and excited state. And it will, of course, have a big effect when the energy splitting is small. Well, you will see that in the next line.

But we have the co and counter rotating terms. This dipole Hamiltonian will also couple the ground state to the excited state in the other manifold. This is the off

resonant term, which will be neglect.

So here are the equations for that. Our electric field is, of course, a plus a dagger. The dipole operator is that.

And if you take what is in front of the electric field and define this as the coupling constant, g, we have what I just told you, that within a manifold, the ground state with N plus 1 photons is coupled to the excited state with N photons. g is simply the coupling constant, which all parameters have been included in that. But then, because of the matrix element of a or a dagger, we have the square root of N plus 1.

OK, so what we are now doing is we are going with this coupling within a manifold, which is strong because the two levels are energetically close. And we will neglect the non-resonant couplings to the other manifolds. And that's just another way of doing, again, the rotating wave approximation.

Now, when we describe an atom in the laser field and the atom scatters light, or we have a laser beam, and I have to tell you, when we have a laser beam, we don't provide exactly N atom, N photons. The laser beam is in a coherent state. And there are Poissonian fluctuations.

But now, we do the next approximation, namely we say when N is large, then delta N can be neglected with regard to N. Then at least for the range of important photon numbers for the problem, this coupling becomes independent of N. And we simply take the average number of photons in our laser beam and define now an electric field, which is independent of N.

And the product with the dipole operator of the atom, that's what we call the Rabi frequency. So at that point, we introduce a classical electric field through this definition, and therefore, the dressed atom picture. Also, we start with the fully quantized Hamiltonian is now identical to the classical field with this electric field amplitude.

But what we have gained is we could discuss everything in a time independent

picture. And now, I promised you, we'll go back to a two by two matrix, which you have seen many times. It's a two level system. We have the detuning. We are in the rotating frame. So that means the ground and excited state with N and N minus 1 photons have the same energy.

But if we have a detuning, there is a difference whether we excite the atom or whether we put a photon in the laser field. But you have seen that many, many times. And the coupling is now described by the Rabi frequency.

So that's kind of our typical two level system where the two levels [INAUDIBLE] the rotating frame, everything is time independent. The two levels are split, and they are coupled. And well, how many times have you seen the solution of two by two matrix? The solution is that the coupling is now increasing the splitting. The splitting between the two levels is, again, the generalized Rabi frequency, which you have seen many, many times.

But now we want to really get the wave function. We want to get something out of the dressed atom picture. And that's the following.

We started out with the bare states-- ground state with N plus 1, excited state with N photons. And the two by two matrix after the diagonalization is now a linear combination of those. And it's convenient to use sine theta, cosine theta, because then the states are automatically normalized.

If I introduce the solution is now that the angle theta is given by this equation, that's just the solution for the two by two matrix. And the physical picture, which we have obtained, is the following. If this here is the energy of the excited state with N photons, I've just sort of drawn this as a reference parallel. But now I want to draw it as a function of the laser frequency.

Well, the ground state with N plus 1 photons has exactly the same energy on resonance. But if the photon is more energetic then this state, because it has N plus 1 photons versus n photons if the photon is more energetic than a resonant photon, this state has higher energy. And here, it has lower energy.

In other words, we have now a level crossing between the two bare states. But a level crossing with interaction turns into an avoided crossing. And the solution which we have just derived are the blue curves.

And if I go far away from resonance, this behavior here, which is perturbative, is nothing else than the AC Stark shift. So this is exactly what we have discussed at this point.

Now, let's get some mileage out of it. And we will talk about light forces. That's really interesting.

And I hope it's another highlight of this course to describe light forces in the dressed atom picture. But what I want to do here is already in preparation for that. I want to not yet introduce forces, simply figure out what happens in the dressed atom picture when the atoms scatter light.

And you know already the solution. The solution is the Mollow triplet. And we have discussed with the Optical Bloch Equations that at low laser power, the Mollow triplet has three curves. It has three broadened curves. And they have different broadening.

And I didn't prove it to you, but I mentioned to you, at low laser power there is-- perturbative laser power, we discussed at great length that there's a delta function. And at low laser power, the delta function remains. So this is sort of the spectrum we get out of it.

And what I want to show you now is how this spectrum is described in the dressed atom picture. Atom photon interaction has a wonderful discussion of all the details. I will share with you a few highlights. But we will go much further in the description of the Mollow triplet then we've gone before. But I will not go into all excruciating details.

So let's discuss first. That's something we hadn't done so far. What are the intensities in this Mollow triplet? So the way how we obtain it is the following.

We have the manifold, the manifold with N minus 1 and the manifold with N photons. And this is the detuning. And if we allow for the atom laser coupling, the splitting here is a generalized Rabi frequency. And well, you immediately see that in this picture, if you allow spontaneous emission from the upper manifold to the lower manifold, you have four different combinations.

One is the long photon is the blue sideband. This is the red sideband. This is the blue sideband. This is the red sideband.

And those two are the carriers. But now, we can do the following. I've just written down for you the solution of the wave function with sine and cosine theta.

If we want to emit the blue sideband, we go from the top to the bottom. But that means, since a photon in the rotating wave approximation can only be emitted when we go from b to a, we find that the matrix element for that involves cosine theta times sine theta.

However, if you want to emit the carrier, you go from here to there. Well, we always have to find the excited state to begin with. And we go-- I must have made a mistake.

Sorry. When I went to emit the blue side, when I go from all the way up to all the way down, and I always have to find an excited state which connects to the ground state. And this gives me a term, which is cosine square. The red detuned sideband gives me sine squared. And if I want to get the carrier, this one, I go from here to here. I get cosine theta sine theta.

So in other words, we can immediately read off the matrix elements. And since the intensity, Fermi's Golden Rule, is the matrix element squared, we find immediately that the upper lower sideband and the carrier involve terms cosine to the 4 theta, sine to the 4 theta, cosine square sine square.

So we got the matrix element. And the only thing we have to know now is that the width or the rate is given by the matrix element. The rate is given by the matrix

element square. But now, what is the intensity of the line?

Well, the upper sideband is emitted by the upper dressed state, because you have to go from the upper dressed state to the lower dressed state to get the blue detuned photon. So therefore, the intensity for this sideband is this rate times the population in the upper dressed state. So what is missing at this point, and we need that urgently for next week, what is the population in the upper and the lower sideband?

Well, we know the transition rates between the upper and the lower sidebands-- well, between the upper and lower dressed state. We know we've just found here the rate. What is the rate? What is the matrix element squared when you go from one is the upper, two is the lower state of the manifold. So we have just gotten those rates.

And rate equation says, the change of population in the upper dressed state is given by the departure from the upper state and what arrives from the lower state. So since we are interested in the steady state solution, we can just set the left side to zero. And then, the rate equations simply give us the population in terms of the rates.

So in other words, if you ask, what is the steady state solution in the upper or lower dressed state, it's given by a ratio of those rate coefficients. But we just got those rate coefficients. And this gives us another bunch of sine and cosine functions.

So in other words, we know now what is the population of the dressed states. I mean, everything is expressed by one non-trivial parameter theta. And all we get is sine square, cosine square, sine to the 4th, we just get those trigonometric functions.

Yes. I think we can finish that. Good. So this is the steady state solution. But now, we can go back to the rate equation and say the time derivative of the rate equation has to relax to the steady state.

This is nothing else than rewriting the equation above. But by rewriting the equation

above in that way, we find now the prefactor in front of-- I mean, here, what you have is the time derivative is given by the difference from steady state. And this is the relaxation coefficient.

And the relaxation coefficient is-- that's the solution. You just have to do the substitutions yourself. The relaxation coefficient, how populations relax to steady state, is now given by the sum of the two rates. One is going from the upper to the lower and from the lower to the upper state. And each of those rates has been expressed by sine square and cosine squares.

I'm not doing it here, because it's a little bit more complicated. But you can also rewrite for the dressed atom picture what happens-- you can rewrite those equations in a way what happens not to the population, but what happens to the coherences, to the off diagram matrix element formulated in the dressed state basis. And what you find is, you find a correlation time, a relaxation time, for the coherences.

Remember, for the naked state, population's relaxed with gamma. Coherence is relaxed with gamma over 2. This is for the Optical Bloch Equations. But now we have reformulated the physics in the dressed atom picture. And what the solution is the following.

This is sort of a weak laser power. But in the most general case, what you will find is that the width of the central peak reflects the relaxation time for the population. And the width for the sidebands is the relaxation time for the coherences in the dressed atom picture.

So we had before with the Optical Bloch Equations, I showed you different cases. We had a three quarter gamma. I need two more minutes and I'm finished with my unit. So we had weird factors of three quarters. And I showed you how they come from the eigenvalues of the matrix of the Optical Bloch Equation.

But now we understand them physically, namely the central feature has a peak, which is the relaxation time for the populations in the dressed atom picture. And the

sidebands have a width, which is the relaxation time for the coherences. They no longer differ by a factor of two, as in the naked bases in which we formulated the Optical Bloch Equations.

By the way, I know a lot of you may be curious about the coherent peak. You can read about it in API. It's 10 pages about correlation function. But in the end, the rate of the delta function can be, again, described by values we have just calculated. We have just calculated steady state populations and transition rates.

So this rate is nothing else than another combination of sine square and cosine squares of this angle theta. So everything is wonderfully and simply described the full physics of that by just this angle theta in the dressed atom picture. OK So with that, we are done. And I think I'll discuss the last page with you on Monday.