

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

PROFESSOR:

OK. Here is the menu for today. We're discussing line shifts and line broadening. And I want to finish up today this chapter by describing collisional narrowing, also called Dicke narrowing. Then I want to have two more shorter topics on two other aspects which lead to important line shapes and line broadening.

So I want to quickly discuss the spectrum of emitted light by an atom. And I want to discuss collisional broadening. None of that will be done in-depth. The spectrum of emitted light is really open-ended, and we will have a more advanced treatment in 8.422. But I do feel if I show you all are kinds of line shifts and line broadening, I should at least mention here the basic things.

And collisional broadening, I'm not sure how many atomic physics courses you will find which teach about collision broadening, because this is the physics from gas discharge lamps, the old-fashioned physics. However, I've realized that a lot of people know now about clock shifts, and mean field broadening, and mean field, and the ultra-cold gases.

And they have no idea that similar physics actually happens in ordinary gas. So at least in terms of broadening your understanding, I want to talk just 10 minutes about collisional broadening. I've completely eliminated from this course the quantitative description about collisional broadening, but I want to show you a few cartoons and put some pictures into your mind.

And probably, we have time to start the next chapter, which is actually a pretty short one, on two-photon transitions. For me, actually, Dicke narrowing is really the highlight of the course, because it provides conceptual insight into what really line broadening is, and to realize that collisions can narrow lines and not just broaden

them. This is sort of subtle and insightful.

And similar to photon transitions, it's short, but I hope it's also a highlight, because there's so many people-- you and other people-- who are sometimes struggling-- a photon is absorbed and is emitted. Usually, the photon is not absorbed, the photon is scattered. And whenever you think about photon in, photon out, you really should think about two-photon transitions. So the framework of two-photon transitions allows me now to give you the tools how you should really think about whenever you have atoms interacting with light. The light is not absorbed, the light is just scattered. And so you need that.

On the other hand, based on all the descriptions I've given you about light-atom interaction, two-photon transitions would just mean we need one more order of perturbation theory, and then it's the same thing you have already learned. So it's a highlight, and it's to some extent also a review. You will actually recognize, in some situations, the difference between two-photon processes and one-photon processes. It's not so big, it's just you have to-- you'll see. You have to use a different Rabi frequency and some different concepts. OK.

So that's, I think, really an agenda of highlights. Let's go back to the physical picture we drew up on Wednesday about Dicke narrowing. I just explained to you that when we have an atom which is trapped and tightly confined, that the spectrum consists of a sharp light and sidebands.

And now I was addressing the situation-- what happens if you have an atom which is surrounded by buffer gas? I would say, well, that's a cheap trap created by nature, because when the particle wants to fly away, it collides with a buffer gas atom. So it stays put. So it's a cheap trap, but it's also a lousy trap, because there is some randomness in the number of collisions.

But I'm waving my arm, so that maybe you can go along with the picture that it is an ensemble of traps which have very different trap frequencies. And then we would expect, based on our understanding of the spectrum of confined particles, that we have a carrier, which is always the same, at the electronic excitation frequency. But

then we have sidebands at the trap frequency. But those sidebands are now smeared out, because we do not have a defined harmonic oscillator potential here.

So therefore, I sort of tried to lead you in the way that that may be one way how you can think about the situation. And well, now I want to give you a different way to look at it. And let's see how those things come together.

So if I would ask you, give me a quantitative estimate how wide this line is, do you have any idea how we can do that? Or let me even put it this way. It's again one of those kind of things-- you have the knowledge, but to put it together is hard. But I would guarantee everybody in this room has the knowledge to write down in one line what is the width of this line.

What is the width of this line? It's an inverse time. What time? Coherence time. Yeah. But now, if you have an atom which is now starting and it hits buffer gas, how would you estimate the coherence time?

AUDIENCE: Mean free pass.

PROFESSOR: Mean free pass will be important. But the coherence time-- how long will the atom talk to the laser beam in a phase-coherent way? So what do we have to compare to the wavelengths?

AUDIENCE: Mean free pass.

PROFESSOR: The mean free pass. This will be the important parameter. But if the mean free pass is much shorter than the wavelengths? Well, the atom has just moved a tiny bit, one mean. Is it still coherent at this time?

AUDIENCE: It's shorter than the wave.

PROFESSOR: If the mean free path is shorter than the wave. So it will do many collisions. But since it stays localized, it will still coherently interact. When will it stop coherently interacting?

AUDIENCE: Mean free path.

PROFESSOR: No, the mean free path is always smaller. Let's just assume that. I hear "diffusion."

AUDIENCE: [INAUDIBLE].

PROFESSOR: Yes. The atom has a short mean free path, but it will do a random walk. And the movement it diffuses by more than a wavelength, it has randomly changed its position by a wavelength. And that would mean it experiences the phase of the drive field in a random way. End of coherence. 1 over this time is its line width. OK.

So let me write that down. So based on the concepts we have learned by looking at Doppler broadening and all that, we realize the important aspect is, when do atoms in an ensemble randomly move by a wavelength?

So therefore, our estimate is now-- estimate the widths of the sharp peak. We use a model which is diffusion. We know that in diffusion, the random walk, the RMS position in atom has moved away. Ballistic motion is linear in time. Diffusive motion is the square root of time, or z squared is linear in time. And we have the diffusion constant.

Diffusion by, well, λ or λ bar after time, which is wavelength squared. So therefore, we would expect that the full width at $1/2$ maximum of our peak is k squared-- k , the wave vector of the light, and D is the diffusion constant.

Well, since you mentioned the mean free path, let me, already at this point-- I wanted to do it earlier, but it fits in very well here. In an ideal gas, the diffusion constant is given by some thermal average speed times the mean free path. And if you look up some textbooks, there's a factor of 3 .

So our $\Delta\omega$, which is k squared D . Let me write it as k times \bar{v} . And the other k I write as 1 over λ bar. So this is k squared, and I need l . So this is nothing else than k squared-- I take this expression, k squared l . This is one k , this is the other k . And \bar{v} times l is the ideal gas expression for the diffusion constant.

But now you realize that $k \cdot \bar{v}$ -- since in a gas, the most probably velocity is also the momentum spread, this is nothing else than the Doppler width. So therefore, we

find that the line widths in Dicke narrowing-- if we have buffer gas and we have diffusive motion-- is much smaller-- and this is why it's called Dicke narrowing-- than the Doppler broadening if the mean free path is much smaller than the wavelength. So this is where the mean free path comes in.

What happens if the mean free path is much longer than the wavelength? What line widths do we then get? Do we then have a line width which is larger than the Doppler broadening, or do we get the Doppler broadening?

Let's have a clicker question. So if l is much smaller than λ , is $\Delta\omega$ equal to? Or option B, larger? OK. So what do you think? It's always a question when we derive something, how seriously you should take what we derived.

So this expression for $\Delta\omega$ Dicke shows that the line width would get larger and larger the longer the mean free path is. And the question is, is that correct or not? OK. Any more takers? Whoops. Oops. Sorry, press again. I erased it by clicking the wrong button. You've already made up your mind. You know the answer. So stop, display. Yes. The majority answer is definitely correct.

What happens until the first collision happens, you just have normal Doppler-- you don't have diffusive motion, you have ballistic motion. And if the line width is already determined once the atoms have spread out by a wavelength, that's it already. And then if then the particles collide, it doesn't matter. So what we have assumed is we've assumed that the relevant model for the spread out to a wavelength is a diffusive model. And if we are past that-- because the mean free path is larger than the wavelength-- we have to go back to normal Doppler broadening. OK.

But now let me calculate Dicke narrowing by using the formalism we have developed. So let's use the correlation function for that. And we know that the line width was nothing else than the Fourier transform of the correlation function, how the atoms experience the drive field at two different times.

So we had here the matrix element, the Rabi frequency squared. We have between

time t and $t + \tau$, the phase of the drive field accumulates $e^{i\omega\tau}$. But now we have the e^{-ikr} factor. And r , or the position-- in diffusion, you often call it s -- changes now, because the particle undergoes a random walk.

So what we have to do is to describe collisional narrowing, we have to take this factor and average it over our ensemble of diffusing particles. And well, diffusion means that if particles start out at $t = 0$ at the origin, the probability that we find the particle at time t , a distance s away from the origin-- so I'm depicting this one here.

The particle does sort of a random walk. And after time t , it is out at a position s . And the probability for that is $e^{-s^2/4Dt}$. And you see it's s^2 is quadratic over t . This is a random walk. s increases s 's square root of time. And the probability is normalized by this expression.

So therefore, this red average in the correlation function is calculated by convoluting with this probability distribution for the random walk. So it is e^{iks} , $e^{-s^2/4Dt}$. The time is now called τ . We used to integrate over all possibilities. And $e^{i\omega\tau}$ is a common factor.

My notes show that the integral is done from minus infinity to plus infinity. Either this is right, or there's a factor of two to be accounted, because well, it depends. Now, is s a coordinate, or is s in radians? In one case, it has a sign, in the other case, it has not. I'm not able to reconstruct it now, but it's just a numerical factor, which would be affected by that. So the result of this integral is that we obtain an exponential function.

So what we find out is that there is now an exponentially decaying function.

Remember, we had situations where the correlation function for how the atom experiences a coherent field was decaying exponentially because of spontaneous lifetime? And now we have an exponential decay, because the particle is diffusing around.

But we also know, of course, that exponential decay, when we Fourier analyze it,

gives us a Lorentzian. So when we ask what is the rate, Fermi's golden rule's rate expression, the rate for excitation. We want to do the Fourier transform of this correlation function. And what we obtain is, well, a Lorentzian which looks like the Lorentzian for spontaneous emission, except for that we have a different width now. It's a Lorentzian with a width $2k^2 D$.

So pretty straightforward. But I hope you've seen and you've enjoyed there's a very intuitive picture. But the correlation function-- if you just look at it from the perspective of the atoms, how do I experience a coherent field? You put in simply the diffusive motion. You can exactly calculate what the line shape is. And that's definitely something you would not have known how to do it without this formalism. OK.

I have a question for you. And this is, how does the spectrum for diffusive motion look like? What we have just calculated was that it is a Lorentzian with the line widths we've just determined. But I want to give you another choice. And this is what we discussed at the end of the last class. This picture where we took the confined particles to the ridiculous limit-- where there was no trap anymore, just collisions-- suggested actually that we have some sharp line. But then this envelope of those sidebands give some unresolved pedestal.

So the question is, what is correct? We have just done a quantitative calculation using the model, a diffusion propagator. But we also had this intuitive picture, which had more this kind of bimodal distribution. A broad pedestal and a sharp peak. And I want you to think about it for a few seconds. What would you expect to be the correct answer?

Is the line shape just a Lorentzian, or does the line shape have two different parts to it? OK. Does somebody want to speak out in favor of his or her choice?

Well, one argument is, when you derive something, it must be more correct than when you wave your hands. Therefore, we derived a Lorentzian, and the other picture was just waving our hands and using some analogy. OK. That would be one argument to vote for A. Somebody wants to defend B? Pardon?

AUDIENCE: More intuitive.

PROFESSOR: More intuitive. Yes. You know, something must be right about it. I mean, in some limit, this must be like a trap, and there should be some sidebands. Yes. OK. But now I would ask you-- Nancy, if you say it's more intuitive, why don't we get those broadening? Why don't we get this extra pedestal in our quantitative derivation?

AUDIENCE: [INAUDIBLE].

PROFESSOR: But what? Didn't we lose the exact probability distribution for diffusive process?

AUDIENCE: There's no harmonic oscillator. The sidebands [INAUDIBLE] harmonic oscillator.

PROFESSOR: So that's an argument also for A. We don't have a harmonic oscillator. And when I said there is sort of this trap-type feature, I'm really over-extending the analogy.

AUDIENCE: The mean free path is not--

PROFESSOR: Pardon?

AUDIENCE: The mean free path is not infinitesimally small.

PROFESSOR: The mean free path is not infinitesimally small. Now we're getting close. We assumed diffusive motion. We put in the exact expression for diffusive motion. And of course, you also know, when you describe a line width, what happens in the middle at detuning 0 is more what happens in the limit of long times. What happens further and further out happens at shorter times.

What is the motion of atoms at short times? Before the first collision happens, it moves straight. So the diffusive propagator which we put in is only valid after the first collision. Until the first collision happens, we have free motion. And free motion should give rise to simple Doppler broadening. So until the first collision happens, we should get a little bit of Doppler broadening. But once the collisions happened, we should be very well within the description of the diffusion operator.

And to address your concern, actually, in the limit of many sideband, often, the

envelope of all those sidebands is actually determined by the Doppler profile. There is a limit, if you have a large modulation index, that you have many sidebands. And in some semi-classical limit, the sidebands have an envelope which is the Doppler profile. So now I think the two pictures agree.

If we had used a propagator which would interpolate between the first moment where the particle moves straight, and then diffusion, we would have gotten little pedestals here. And so that sort of tells us that the smeared outside bands, yes, strictly speaking, they are linked to harmonic motion. But they are sort of the leftover of the motional effect, which is the Doppler effect. OK.

Good. So we have the correct answer here. And here, it is the model neglects ballistic motion until the first collision. Questions? OK.

Let's now spend 10 minutes discussing the fluorescent spectrum of an atom. As I pointed out in my overview, we teach much more about it in 8.422. We use a dressed atom picture. But something would be missing if I wouldn't do it also in this course, because we have discussed to quite some extent what happens when we excite atoms with light.

And what we have discussed so far is that we excite atom and scan the laser, or the excitation frequency. And when we are plotting the intensity of the fluorescence, we are looking at the number of scattered photon. And what we are scanning is the detuning of the laser. We would expect, in the case of a motionless atom, simply a Lorentzian, or in the general limit, a power-broadened Lorentzian.

Just to be clear, what I'm discussing in those 10 minutes is I'm completely ignoring that the atom can move. So you should think that it's either an atom with infinite mass, or it's an atom tightly localized in the Lamb-Dicke limit, and all we are looking at is the structure of the central peak. So no motional effects. We just look at the pure kind of intrinsic line widths of the electronic transition. OK.

So we have discussed that. But now I want to look at another aspect of spectroscopy. And this is another scan we can do. We want to have a fixed

detuning. And we look at the spectrum of the emitted light.

So let's assume that we have a laser which is at a detuning δ . The laser is fixed. The light is emitting. But we are now dispersing with a spectrograph. We are analyzing what is the frequency of the emitted light. And we determine the spectrum.

And yes, the question is, how does it look like? And I want to give you four options. So this is 0 detuning at resonance. Oops, I need one more. Let me relabel it to make it clear.

So our option one is our spectrum is a delta function at the resonance of the atom. Option two is it is a Lorentzian centered around the resonance. Option three is it is a delta function at the laser frequency. And option four is it is a Lorentzian with line width γ .

So we have two options. Is it a delta function at ω_0 or ω_l ? Or is it a Lorentzian-broadened function at either ω_0 or ω_l . And since we want to keep things simple, we want to first discuss the case of the perturbative case that the laser which excites the atom has very low power.

So what would you expect? An atom is excited, it's a little light bulb. You analyze the spectrum of the light bulb. Which of the four spectrum will you measure?

OK. Let's try again.

[LAUGHTER]

And maybe I should, like our online learning system, try to give you one hint. I want you to really think hard what energy conservation means in this problem. OK.

OK, we're getting closer. Now, the way how I want you to think about it is if you take the limit of low power, you should really think about it that there is only one photon in your whole laboratory. This photon is scattered of an atom, and then you measure the frequency of the photon. And energy conservation clearly tells us that the frequency of the scattered photon cannot be at the resonance frequency. It has

to be at the laser frequency. Otherwise, we would violate energy.

But it seems-- still, the more subtle thing is, is it a delta function? Of course, a delta function will be broadened. If you do the experiment for one second, you will have a Fourier line width which is 1 hertz. But that's, for practical reasons, almost like a delta function, because the spectral line widths, or the natural line widths of your favorite atom is 10 megahertz. So there will be some temporal broadening, which is in case in any realistic experiment. But we are talking about, are you limited by the experiment and by technical noise, or are you limited by the natural line widths?

Now, I would argue, just with energy conservation-- but I'll give you another argument in a second-- that if I have a monochromatic laser, the photon which has to come out has to be exactly the same frequency. Because we talked about energy conservation. And if I would start with a photon at the laser frequency and I would measure a photon which has 10 megahertz away, I would violate energy conservation by 10 megahertz. So the argument is actually correct. It's a delta function at the drive frequency.

But I want to give you another argument. It's, again, something you have heard in the course. But you should really take it seriously, and apply to that situation. And that's the following. Remember when we talked about the AC stark effect. I told you that you should really think about the atom as a harmonic oscillator. We even introduced for you the oscillator strengths that we could even quantitatively describe the response of the atom to the scattering of light by pretty much a model where we have a mechanical model where an electron is attached with a spring to an origin. OK.

Let me now paraphrase the question. Well, I've given you the answer. Before giving you the answer, I should have asked you the following-- if you now make the assumption that you have a harmonic oscillator, the harmonic oscillator has a resonance frequency of ω_0 . And it has a damping rate of γ . But we are driving the harmonic oscillator not at ω_0 ; we were driving it at a frequency ω_l . At what frequency does the harmonic oscillator respond? At the drive

frequency or at its own frequency?

AUDIENCE: Drive.

PROFESSOR: Of course. It's a drive frequency. And if you have a CW experiment, I said let's wait a second, let's really do a long experiment, if you have a driven harmonic oscillator, and you drive it for a long time, and you analyze the spectrum of the motion, is the frequency spectrum of the driven harmonic oscillator absolutely sharp at the drive frequency? Or is it broadened by γ , the damping constant of the harmonic oscillator?

AUDIENCE: Sharp.

PROFESSOR: Who thinks it's sharp? Who thinks it's broad? Hands up for sharp. Hands up for broad. OK.

AUDIENCE: [INAUDIBLE].

PROFESSOR: Look, what you are overlooking is the following. There is the difference between a transient, which is damped out at a rate γ . But then there is the CW response. If you have an ultra-weak drive, you leave it on for an hour, the harmonic oscillator will reach steady state.

And it will just oscillate driven by your drive. And it's actually monochromatic. And if you analyze the motion, it's a harmonic oscillator in steady state with a drive, and it just moves with a fixed amplitude and fixed frequency. I mean, the way how I shake my hands, this is a delta function at the drive frequency.

So in this simple harmonic oscillator model, γ and damping comes in when you do what I discussed earlier, when you change the drive frequency. But this is a completely different experiment. We are not changing the drive frequency and looking for the response. We are at a fixed drive frequency, and we are analyzing what the motion is which comes out.

And then, as you know from differential equation, you may have a grade transient. You may have a transient at the resonance frequency which dies out with a rate

gamma. This is when you suddenly switch on your drive, and you're not adiabatically switching on your drive. But this is just a transient.

But what we are asking here when you do a CW experiment, you drive it for a long time, you look what happens. And the motion of the driven harmonic oscillator is a delta function at the drive. And what is valid for a harmonic oscillator is also valid for the atom, pretty much for the same reasons. Any questions? Yes.

AUDIENCE: Regarding the energy conservation in D, wouldn't the photon processes explain energy conservation?

PROFESSOR: OK. You are now really asking-- Nancy is asking about two-photon processes. Well, we don't want to stop here. We want to see something more interesting. This is just sort of the trivial, simple harmonic oscillator. The question is now, if you want to see something richer, if we want to see a little bit of broadening, or we want to see something which is not as boring as just a classical harmonic oscillator, what do we have to do?

AUDIENCE: Increase the strength.

PROFESSOR: Increase the strengths of the field. And I heard somebody saying, two photons. And that may be--

AUDIENCE: Next-order perturbation theory.

PROFESSOR: Pardon?

AUDIENCE: It's like next-order perturbation theory.

PROFESSOR: Next-order perturbation theory. But just in the sort of intuitive picture, if you have two photons which come quickly enough-- and how quickly the two photons come will be debated. Of course, it's parametrized by the Rabi frequency, the strengths of the drive.

But if two photons come quickly, each photon has to be scattered with a delta function because of energy conservation. But if we have two photons, suddenly, it's

possible that one is scattered here and one is scattered here. And we still conserve energy. So if you want to see some form of line broadening, if you want to see all the things you have mentioned, we have to go away from the low power limit.

And this is what we do next. So we assume that we have higher power. We are on resonance. And let me assume that the Rabi frequency is larger than gamma. So what is the physics now of the atom if the Rabi frequency is larger than gamma? It can do Rabi oscillation before it's damped. So now we have a system which has an internal dynamics at the Rabi frequency.

And you know if you have an emitter, if you have an antenna which emits light, and you move the antenna around at a frequency-- we had the example of a trapped particle which has a trapped frequency of omega trap-- the spectrum, if you Fourier transform, it leads to the basal function, it leads to the sideband.

So therefore, we know now, based on all the analogies, if you have a modulated emitter, we obtain sidebands. That's what we would expect classically. So let me now ask you the three following possibilities. Maybe we still get a delta function. Maybe we get three peaks where the splitting is the Rabi frequency. So the two outer peaks have a-- and the third option is that yes, we observe the Rabi frequency, but the Rabi frequency is the splitting between the outer peaks.

So the first and second answer differ by a factor of 2 in the sideband spacing. So we have three choices. Do we expect sidebands? That's A and B. Or would you expect that we still have a delta function, because, well, maybe energy has to be conserved at the single photon level? And then is the splitting Rabi frequency on each side, or Rabi frequency between the two sidebands? All right.

Yes. It's indeed the situation. For those who picked A, it's pretty much the definition of what the Rabi frequency is. After one cycle of the Rabi frequency, the atom has gone from the ground state to the excited state and back to the ground state. So the model you should have is that you have an object which is emitting light, but the object has some internal modulation at the Rabi frequency.

Whatever factor of 2 you had in the amplitude versus probability or something, all this has been factored into the definition of the Rabi frequency in such a way the atom is really blinking between ground and excited state with a frequency which is the Rabi frequency. And this frequency leads-- and the spectrum is now the sum and difference of the relevant frequencies. And the relevant frequency is the resonance frequency and the Rabi frequency. So the answer is B.

But this was sort of more in form of a stick diagram. Now I want to bring in in the next question the line broadening. So the question is-- and there are again four choices. The question is, we have now our three sticks.

One is a carrier, and these are the two sidebands split off by a Rabi frequency. Are all three sticks now sharp delta functions? Are all three sticks broadened by the natural line widths? Is the central part sharp, and only the sidebands are broadened? Or do we have a sharp stick with a pedestal, and then two sidebands?

I'm not actually expecting all of you to know the answer, because this is really now getting into more subtle things. But just in terms of show-and-tell, and attract your curiosity to the second part of the course, what would you expect? OK.

Yes. Good. A is definitely eliminated. I mean, if we scatter two photons simultaneously, there will be sort of broadening of gamma, because the atom has a natural broadening. So we can't expect that the two photon processes are sharp. There's no reason to expect that.

The answer between B, C, and D depends now. If we have very, very little power, what we observe is, of course, the low-power delta function. And then we observe very, very small sidebands which are broadened. So answer C is correct. If you just think about it-- what is the structure of infinitesimal peaks?

But when you crank up the power, and what you have is you have an elastic scattering peak, which has a delta function, energy conservation at the single photon level. And then you have those-- they are called inelastic peaks. But when you crank up the power, then the central feature has also an inelastic component.

You can sort of argue that in Rabi oscillation, you are in the ground state, excited state, ground state, excited state. But if you do now light scattering in the excited state, you automatically broaden it by the lifetime of the excited state. And if you crank up the power higher and higher, the elastic peak will be more and more suppressed. And you find actually a spectrum which has only three broadened peaks, and the delta function has disappeared.

It will actually be something where we need more knowledge. The broadening of those peaks is not gamma anymore. One of them will be $3/2$ gamma, the other one will be $1/2$ gamma. But this now really requires a deeper understanding in terms of the atom picture. But the scale of all the broadening is a factor on the order of unity times gamma. Any questions?

So the general answer, sorry, would be-- the generic picture which you should keep in your mind is that. You have inelastic scattering on all three peaks. But the limit of D at very low power is that. And at very high power, is that. The elastic component of the central peak is either 100% or 0%. OK.

So our last topic for line shapes and line broadening is pressure broadening. And as I mentioned in the introduction, pressure broadening has made modern appearances in the form of clock shift, and mean field shifts, and mean field broadenings in Bose-Einstein condensates.

But let me sort of tell you how you should understand pressure broadening. I'm using here the semi-classical picture that an excited atom acts as an oscillator.

And what happened is the atom-- you can say the oscillation is superposition of ground and excited state. It oscillates. But after some time tau, there may be a quenching collision, a de-excitation collision. Then the atom is sort of in the ground state. And then it waits until it's excited again by the laser, and the atomic dipole is oscillating again.

In a situation like this, you would expect that the total line widths is actually the sum of two rates. One is the spontaneous emission rate, and the other one would be the

collision rate, which is this. And in general, this is sort of how people looked at it.

You vary the pressure in your buffer gas cell, and you find that the line width has a component which increases linearly with pressure. So this is one model how you can imagine what happens in a collision. These are like knock-out collisions. When an atom collides, the excitation, the energy disappears. The atom is quenched to the ground state.

Well, we can draw up another model, where we have an oscillation of our atomic oscillator. But then after the same time τ , there is now a hiccup in the phase. It collides with another atom. The other atom is not de-exciting. It's not removing the energy. But after the collision, the atom continues to oscillate, but with a very, very different phase.

And sort of what I assumed here is this time where the collision happens is very short. It's a very short collision time. That I can approximate the collisions as simply imparting random phase jumps to the atom. Well, if you ask, what is the widths of the spectroscopic line, it's exactly the same result. It just means we have a different rate. It's now a de-phasing rate, which is $1/\tau$. And this is added to the spectroscopic widths. OK.

This is sort of just creating some phenomenological picture. Let's now ask a little bit more microscopically. How can it happen? How can it happen that there is a phase change, that there's a change in the phase of the atomic oscillator?

And this leads us now to consider the interaction potential as a function of r between two atoms. We have one atom which is sort of our active atom, which has an excited and ground state. And this atom is now getting closer to, let's say, an argon atom, which acts as buffer gas. And what I'm plotting in this graph now is the interaction potential between our light-emitting, or light-absorbing, atoms and the buffer gas atom.

And let's just genetically assume there is sort of something like a molecular potential. But in general, the molecular potentials will be different in the ground and

excited state. And of course, if you do hyperfine transitions between atoms in a Bose-Einstein condensate, you may have a scattering length, which is different in the two states. And I think you see the connection.

So in general, if the interaction environment is different between ground and excited state, we expect that the interaction potential causes certain shifts. So we could use the picture that we have a phase evolution. And the phase evolution is a frequency difference which is simply given by the difference of the two potentials.

So the picture is a little bit-- which is actually very valid for cold collisions, which-- well, really a new chapter in atomic physics opened up by laser cooling and magneto-optic traps. An atom, when it emits a photon, would emit here at the resonance frequency. But if the collision is very, very slow, it will actually emit, with a certain probability, a photon which is shifted by exactly this expression.

And in that limit, that could be very interesting, because by analyzing the spectrum of the immediate light, you learn something about the interaction potential of two atoms. But of course, this argument has a little bit of a flaw, because you can observe a frequency not instantaneously. You can observe-- if a collision happens very fast, you would actually go through those frequency changes so fast that you cannot resolve them. The question is, which one is larger?

So if the frequency shift is larger than the inverse time for the duration of the collision, then you can observe it. Otherwise, you can't. So the picture we should draw is now the following. And this is sort of a microscopic picture on those phase jumps, which I mentioned earlier. That you have your atomic oscillator. The atom is oscillating.

But now it comes close to another atom. And let's assume there is an energy increase between ground and excited state due to the presence of the buffer gas atom. Then you would say you get a quick oscillation. Depending on the impact parameter, it can last various amount of time. And then later, the oscillation starts. But essentially, this causes a random phase jump. So this time here is the collision time τ_c .

What we expect now is we have to now interpolate between two models. I'm not taking it any further than that. We have an interesting line shape. There may be a central part which is simply collisionally broadened by the interruptions of the phase, by the number of collisions with the collision time τ . But if when the atoms collide, there is a huge frequency shift momentarily, that this will actually affect the wings.

So for frequency shifts which come from the interaction potential which are much larger than the collision time, than the time between collisions, we will observe something which goes by the name "far wing broadening."

And also, the central part of the potential is simply, you can say, it's a Lorentzian which is simply broadened by the collision rate. There is no microscopic information - what is the nature, what is going on during the phase jumps. The wings will have actually information about the molecular potential.

And I just wanted to present it to you in this way to sort of show how you actually have, in a line shape, often two effects. One is pretty much just the interruption of the coherence which gives the central line. But you still have, at least in some limit, information about what causes those phase jumps. And this appears in the wings.

20 years ago in the atomic physics course, we taught you a theory how to describe that. It could be probably taught in one hour. There are some links on the Wiki, but I don't want to carry that further. Any questions about that?

AUDIENCE: So why has the importance decreased?

PROFESSOR: The importance has decreased of that, because I mean, who of you is studying atoms and gas cells? Nobody. The frontier has moved on. We are much more in a regime where we are not observing atoms in certain environment. We are creating an interesting system out of atoms by putting the atoms in a well-defined environment where those things are absent.

Or in the ultra-cold domain, and de Broglie wavelengths is so long that this kind of model for collisions is no longer applicable. We are in the extreme case of a single

partial wave, where a single parameter, the scattering lengths, describes all of it.
OK. Cory?

AUDIENCE: Yeah. Could you talk for a little bit about how this model of collisions breaks down and allows for Dicke narrowing?

PROFESSOR: Yes. Thank you. That was one comment I forgot. What happens is you have Dicke-- thanks for this question. You have Dicke-- I drew up the two potential curves. If the two potential curves are absolutely identical, the atoms can approach each other and can collide.

But there is never any perturbation to the atomic oscillator, because the frequency between ground and excited state is ω_0 , no matter whether the atom undergoes a collision or not. In this limit, which is often realized in collisions with rare gases, in this limit, you do not have any phase interruption by the collision. And this is the prerequisite for Dicke narrowing. OK. I think we have the full picture now.

And that means we can move on to two-photon excitations. So whenever you start a new chapter, you should motivate it. And the question is, why should you be interested? Well, I can also turn it around and said, why not? Because those things happen. And if those things happen, we want to learn about them. And they actually happen very naturally.

The moment you go to higher laser power and you go beyond the lowest order of perturbation theory, you may actually excite two photons at the same time. And actually, what we just discussed about the emission spectrum of an atom, with the sidebands, with the broadening, these are actually examples where we really have to think about it in a two-photon picture, not a single-photon picture.

The second motivation is very practical. We may want to excite an atom from a low-lying level to a high-lying level. But the only laser we can buy, or the only laser we have in the laboratory has lower frequency. That doesn't mean that the case is lost by just stacking two photons on top of each other. We can bridge the gap and still excite the atom, which has transitions only far in the UV. And all we have is these

visible lasers.

There's another thing which changes. Maybe we want to excite hydrogen from the 1s to the 2s level. 1s and 2s have the same parity. And if we try to excite it with one photon, because of the dipole operator which has odd parity, this is not allowed. So there are maybe selection rules where we can go to desirable final state only with two photons and not with one photon.

Finally, I will show you on Wednesday-- remember, there is no class on Monday, because the whole CUA will go crazy with the NSF site visit on Monday, Tuesday. So the next class on Wednesday.

And on Wednesday, I will actually show you if the two photons come from counter-propagating beams, the net momentum transfer to the atom is 0. And net momentum transfer of 0 means zero Doppler shift. So two photons give us an opportunity, which doesn't exist with the single photon. We can excite the atom without transferring momentum. And this is the basis of a Doppler free spectroscopy technique.

And finally, for purely conceptual reasons, we've talked so much about excitation of an atom, and then ask, what is the rate of light scattering? So we have actually used two-photon processes in the course many, many times without actually adequately addressing them.

I didn't tell you anything wrong. I was sort of choosing my words carefully, that everything I told you about light emission and absorption is comparable with the correct picture. And the correct picture is if a photon goes in and a photon goes out, you have to describe it as one process involving two photons. So I think we have enough reasons why we should be interested.

And if I take this picture, that I start in one state and have a two-photon absorption process, I actually want to redraw it, because in the dipole approximation, which I want to use here, we have the dipole operator, and second quantization is a plus a dagger. The electric field is a plus a dagger. So each application of the operator

takes us from one state to the next with a single photon.

So therefore, if you talk about two photons, we can only absorb them if we have an intermediate state, which I will call f . So we should really think about two-photon absorption as a two-step process which involves an intermediate state. And it has to be like that if we want to use, as the operator for atom light interaction, the dipole operator, because a dipole operator is creating and-- the dipole operator involves electric field. And that creates or annihilates one photon at a time.

Let me just make a side remark, but this will be addressed in 8.422. If you use the description of the interaction with the electromagnetic field, which is the p minus e Hamiltonian, and you square it, if you square that, you get an A squared term. The A squared term is actually the product of a plus a dagger squared. So you can actually, with the A squared term, scatter two photons by going from one state to another one.

But I'm not discussing it here. And I leave a detailed comparison and discussion of that to 8.422. Now we are strictly adhering to the dipole approximation.

Oh, actually, let me make a comment about it, because I just realized if I say something and don't say it fully, I may confuse people. So what I want to say is, we will show, in 8.422 that the two pictures, the p minus a Hamiltonian and the dipole Hamiltonian are equivalent. They're connected by canonical transformation.

So therefore, it is not a fundamental aspect of nature, whether you can scatter only one photon when you go from one state to the other. There are two equivalent descriptions. In one description, you sometimes scatter two photons. In the other description, you scatter only one photon. And you will all get a PSET in 8.422 where you show for one example that the results are the same when you sum up over all possibilities.

So therefore, I can maybe here should rather take the position, the generic description of light scattering uses the dipole approximation. And in the dipole approximation, we are describing the atoms, that they only exchange one photon

when they go from one state to the next. And this is a full description but not the only description.

So therefore, we need an intermediate state. But assuming that the first laser is not resonant with the transition to the intermediate state, we need this dashed line. This dashed line is sometimes called a virtual state. And in the following discussion, we will really learn what is the nature of the virtual state and what does it mean. But we need a stepping stone for a two-photon process in the form of intermediate states.

So maybe I should just use three minutes and show you how you would calculate it. I use pre-written slides here, because I'm getting a little bit bored of just writing down the same or similar perturbative expressions. And I just step you through.

So I said we use the dipole operator. But now-- and this is a new thing-- the electric field has not only one component and one frequency. We want to look at two photons, so therefore, it has two frequencies. So therefore, our perturbation Hamiltonian is what we had so far for monochromatic field. But then we have an additional term where we just change the index from one to two to have the second laser field described.

So let me now introduce matrix elements. We have to go from state a to some intermediate state. And then we go to state b. So therefore, we need matrix elements which take us from state a to state f. And again, we have two such possibilities-- one at ω_1 , one at ω_2 , one driven by field 1, the other one driven by field 2. And yeah, we all have the counter-rotating and co-rotating terms.

What happens is we have, of course, a complex conjugate. But I told you already several times that there is an $e^{-i\omega_1 t}$, which is responsible for absorption. The $e^{+i\omega_1 t}$ does emission. And let me just-- I want to do the rotating wave approximation, only keep the relevant term. So I look at two-photon absorption. And that means I only keep forward the minus terms.

You could, if you want, duplicate the lengths of each formula and carry forward the counter-rotating term. You're not learning anything new. You get additional Bloch-

Siegerts and AC stark shifts. This is not the new feature I want to implement.

So let's focus on the new aspect. And this is the following. We do first-order perturbation theory, which takes us in the first step to the intermediate state. And you have seen this expression many, many times. The only thing is in addition to what we had for one laser beam, we have to add a second possibility, which comes from the second field.

And now we are at the intermediate step. And we want to take the second step to the final state. So what we are now doing is we are writing down Schrodinger's equation, derivative of the wave function is Hamiltonian times the wave function. And we are especially interested in how do we accumulate probability or amplitude in the final state.

But what we are doing is on the right-hand side, because we can't start in the ground state, we've gone to the intermediate state, we are now plugging in the previous result, the first-order result for the intermediate state. So therefore, by using second-order perturbation theory, we want the second-order perturbation theory, so we integrate this equation with respect to time. But on the right side, we use our first-order result, which we had derived earlier.

And then we just write down the integral. Everything is just exponential function. And we get the result here. So now with two steps, we have obtained an expression for-- I know time is over, but let me just finish the argument.

So we have now an expression in second-order perturbation theory. What is the amplitude in the final state? And things look a little bit messy, because we have four terms. And we should get four terms, because we have two interactions.

We can take one photon of one laser beam-- first step, ω_1 ; second step, ω_2 . We can switch the order of photons. And of course, if you write down everything correctly, nothing is forbidding the atom of taking both photons out of the same laser beam. Yeah, let me stop here.