PROFESSOR: Ready to roll? OK, good afternoon. Count down four lectures and you know everything you have to know about atomic and optical physics, at least for those of you who take part two-- for those of you who've taken part one and part two-- for the other ones, well, there is one more semester which is the following spring.

So we have discussed light forces, first with the optical Bloch equations, and then we discussed the stimulated light force using the dressed atom picture. So this is sort of finishing the presentation and derivation of they light forces. Today, I want to continue to discuss for the aspects. We've discussed dipole traps, different ways how we can understand why there is a dipole potential from the harmonic oscillator to the dressed atom to the refraction of light by small spheres. So in all cases, we realize red detuned light traps.

Today I want to discuss what is at work here, electrical magnetic forces, and then I want to address what several of you have asked, where is energy conservation? Where does the energy go when we cool with a stimulated force? But this will only take 10 minutes. The main part of today's lecture will be techniques to go to very low temperature to ultra-low temperature. This is sub-Doppler cooling, sub-recoil cooling, evaporative cooling, and magnetic trapping as a technique. So this is the menu today. Again, a feature-picked menu.

Now talking about electric and magnetic forces, you may think it's a trick question, but it's rather subtle. All optical forces were derived from the electric dipole Hamiltonian D dot E. This is what we used as a starting point for light forces. So my question is, what about the Lorentz force? The Lorentz force comes from the magnetic part of the electromagnetic field. Now, is there Lorentz force on the atoms? And, if yes, does it contribute to the light forces or is it negligibly small?

Actually, the question I'm asking you, I don't think it's addressed in any textbook and I think you can go to conferences and ask some of your colleagues. You may get a wide range of answers. So but anyway, has anybody of you thought about it? What about, is there Lorentz force on neutral atoms? If no, you're done. If yes, does it contribute, is it negligible, or is it even dominant?

AUDIENCE: [INAUDIBLE] temperature or is equivalent to the T minus A picture, and the T minus A picture has the Lorentz force, so [INAUDIBLE]. So.

PROFESSOR: So it should be included because we are talking about two different representations. Collin?

AUDIENCE: Yeah.

PROFESSOR: OK, so yes. The Lorentz force is included. Also, when you ride on D dot E, you think it's the electric force of an electric dipole. We come to that in a moment. The next question is, do you have any idea if the Lorentz force is important or whether it's negligible?

AUDIENCE: For neutral, or?

PROFESSOR: For neutral atoms.

AUDIENCE: [INAUDIBLE]

PROFESSOR: OK, who of you is working on optical lattices or wants to work on optical lattices? OK, just to catch your attention, when the atom goes up and down the optical lattice, if feels the lattice potential. The force is 100% the Lorentz force. There is no contribution from the electric force. So that's what I want to show you now.

So we have two fundamental forces, the electric force and-- if you have two charges separated in the field, two different electric fields-- we derive from the Coulomb force, the dipole force. And the force on the electric dipole is-- that potential is-- D dot E. The gradient of it is D times gradient E-- with correct Victorian notation-- whereas the Lorentz force is V cross B. So these are the two forces we have to

consider.

Let me show two configurations where we use dipole forces, the stimulated light force. One is here where we focus a laser beam, and let's just assume for the sake of the discussion that the laser beam-- it's propagating here and the linear polarization goes up and down. So now we have a dipole moment of the atom, which oscillates up and down, which is parallel to the electric field. So now we have D dot E, and indeed if you take the atom and move it into the laser beam, it will experience an attractive force and this attractive force is purely electric.

But let's now come to the situation which many of you work on in the laboratory, would that we have two plane waves-- so the laser beams are infinity extended-- we have two plane waves and they form an optical lattice. Now, the dipole, the oscillating dipole, the electric dipole, is again driven by linearly polarized light. It points up, but the gradient of the electric field is in the direction of the interference of the lattice, so the gradient of the electric field is perpendicular to the dipole moment. So therefore, D dot E is exactly 0, and therefore, the electric part of the potential-- the Coulomb force which is microscopically behind it-- will not contribute anything.

So I could stop here and say, well, OK, since we have only the Coulomb force and the Lorentz force, I've proven to you it's not the Coulomb force, so therefore it's a Lorentz force. But let me just tell you why. So what I'm really telling you is it is actually remarkable that if you go along the Z direction and we take the Z derivative of the AC Stock's shift potential, and the AC Stock's shift potential is one half, alpha is now the polarizability-- we use alpha in many places-- times E square. So this is the AC Stock's shift potential, and the spatial derivative of the AC Stock's shift potential is a force, but it's the Lorentz force.

Actually, I was amazed. I had to work it out myself because I haven't seen it discussed anywhere, but you can simply take the electric field of a standing wave, perform this derivative, and what you find is you get the Lorentz force. If you wonder how you get it without going through all the Victorian notation, the Lorentz force is V cross B. The velocity of the charge is related to the derivative of the dipole moment.

The dipole moment is charge times position, and the derivative of the dipole moment is charge times velocity.

So don't let yourself fool yourself. When you have an oscillating field, there are charges which move their currents inside the atom. So our V cross B, the V is pretty much D dot. The D dot is polarizability times E dot, this is alpha polarizability-- frequency times the electric field. But, as Cody said, we can go form P dot A to D dot E-- or, to be precise-- the magnetic field is the curl of the vector potential, and the curl gives us the [INAUDIBLE]-- I'm suppressing Victorian notation, here-- and you know that E, the electric field, is the temporal derivative of A, so therefore we get that.

So if you now multiply D dot with the curl of A, you get something which is K alpha E square. K is for a standing wave, the spatial derivative. So I've shown you-- at least you can see by dimensional analysis, or by just pointing to the different terms, that if you take the AC Stock's shift potential through the derivative, you get everything which you have in the Lorentz force. It would be a 10 minute, easy homework assignment, but we're running out of homework assignments-- we're at the end of the semester-- to show that, explicitly. So it is really the Lorentz force which provides the trapping potential in an optical lattice. Questions? [INAUDIBLE]

AUDIENCE: So in the simple picture you have [INAUDIBLE] the lattice, you're saying that the transverse confinement is provided by the electric force, but what about the perpendicular direction? [INAUDIBLE] electric field like this, but the other direction?

PROFESSOR: OK, I can discuss with you two simple geometries. In one case, in one direction it's the pure Coulomb force at work. In the other direction, it's the pure Lorentz force. But in general, when you go through an AC Stock's shift potential, you have both. So you may wonder that the Coulomb force, when you integrate it, would not be integrable. It does not give rise to potential. Also, the Lorentz force does not give rise to potential. But if you add the Coulomb force and the Lorentz force, you can integrate it up to the optical dipole potential. I find it remarkable, but it is what it is. Questions?

So energy conservation, yes. That's a good thing to discuss now. Let me first discuss I want to have-- energy conservation has to be discussed with two different perspectives. One is a transient phenomenon-- which I'll do in a few minutes-- but let me first discuss something, which is CW, and this is the following. You have atoms which move over a standing wave, and I use the dressed atom picture to show you that-- for a strong, blue detuned standing wave-- we have a friction force. So the atom-- this was the Sisyphus cooling-- that the atom is climbing up hills, there's some transition, climbs up the next hill, but on average the atom climbs up more hills than it scoots down, and therefore there is work done. There is a net force, and this net force provides friction. It slows down the velocity of the atom.

My question to you is, where has this energy gone? The kinetic energy lost by atom, where does this energy appear? Which fields?

AUDIENCE: [INAUDIBLE] the laser beam?

PROFESSOR: Well, we have two fields, the laser fields, and we have maybe a spontaneous emission? But the force is only a stimulated force due to-- yeah? Which fields?

AUDIENCE: The laser field? Because when you stimulate into a laser field, your [INAUDIBLE] modulator [INAUDIBLE].

PROFESSOR: Well, I have a little bit problem with that because when the atoms-- the force comes from distributing photons from one standing wave to the other one, and these are photons of equal frequency. So in that sense, the net effect is that those two laser beams-- actually, the atom moves, here. So what should happen to the laser beams? Photons are not disappearing. I will later tell you, in the transient picture-- and this, maybe you've already done your homework assignment number 10-- there is something happening to the laser beam. But this would only happen when atoms fly in and fly out of a laser beam. Then we have a different situation. But here, if the atoms stay in the laser beam the laser beam-- so to speak-- sees a constant index of refraction medium. It just goes through and nothing has happened to the photons except that-- occasionally-- photons are spontaneously emitted.

Well, OK. If it's not the laser beams, what remains now?

AUDIENCE: Spontaneous [INAUDIBLE].

PROFESSOR: Spontaneous emission. Now how can spontaneous emission carry the energy? What is the spectrum of spontaneous emission?

AUDIENCE: Triplets?

PROFESSOR: Yes, the mono-triplet. We have the central carrier, the central carrier is not carrying away any energy because it is at the laser frequency. Now we have two side bands. Can the side bands carry away energy?

AUDIENCE: Interlacing them.

PROFESSOR: Huh?

AUDIENCE: Interlacing with the [INAUDIBLE].

PROFESSOR: Asymmetry in what sense?

AUDIENCE: If the side bands are [INAUDIBLE] escape.

PROFESSOR: Well, but didn't we discuss, with a radiative cascade, that the blue detuned side bend comes when you go from the upper dressed level to the dressed level? And in the lower dressed level, you can only emit a red detuned side band, not a blue detuned side bands. What we said in the radiative cascade after one blue detuned photon, there is a red detuned photon. So the number of photons emitted on the red side band and on the blue side band are exactly equal. [INAUDIBLE]?

AUDIENCE: But then you have the [INAUDIBLE] cooling picture, where the red detuned side then was emitted in the lower field region?

PROFESSOR: Yes. So the situation is the following. The atom emits an equal number red, blue. Red detuned and blue detuned, red detuned and blue detuned. But when it emits the blue detuned, the side bend spacing is larger because it emits the blue detuned photon when it's on top of the hill. And the red detuned photon is emitted when the

atom is more in the dark, and therefore, the side bend spacing is lower. So it alternates blue, red, blue, red. But the blue detuned photon is emitted when the generalized RB frequency is larger, when the atoms is at an anti-node And the red detuned is more preferably emitted when the photon is at a node of the standing wave.

So it is the difference of these modulations, the side bands are modulated, and there is a preference to emit the blue detuned side band when the side band spacing is larger than in the red detuned case. And this is where the energy goes.

AUDIENCE: You could make the argument and that would give a simple picture for the maximum cooling rate?

PROFESSOR: Yes.

AUDIENCE: Where you have omega times your [? scattering ?] rate?

PROFESSOR: Yes. That's exactly how, last class, I calculated the cooling rate. The cooling rate, the energy removed from the system, is, well, yeah, I said it is the height of the lattice, but the height of the lattice is also the amount by which the-- the lattice is nothing else than the generalized RB frequency. This is the energy level of one crest level, and therefore the lattice is the increment by which we modulate blue and red side bands. Exactly. This can be quantitatively worked out. If I had another half hour, I could do it for you. It's in the references, but this is the physical picture for cooling.

OK, but so you realize, this is one of the subtleties. We can understand the force simply from a simple picture, energy levels, Sisyphus cooling and such, but in order to understand where the energy comes from, we need the mono-triplet. We need spontaneous emission. So the force comes from the redistribution of photons. It's the stimulated redistribution of photons which is responsible for the force, but the energy balance comes from, what can be neglected for the momentum balance, namely the few spontaneous emission events, and especially those which involve the blue and red side bands. And it beautifully works together and all the equation of

physics are obeyed. OK, but that tells you now that the cooling effect hinges on spontaneous emission. And yes, spontaneous emission is dissipation. And whenever you want the friction coefficient-- you want dissipation-- you need an open system and the open system is spontaneous emission into those nodes.

So let's now talk about an opposite situation where we do not have any spontaneous emission. And I think almost all the people who work with optical lattices use infrared lasers for them, or [INAUDIBLE] detuned laser where spontaneous emission happens only once every 20, 30 seconds or 100 seconds per atom, so we can really neglect it. So now the question is the following. Assume you have a focused laser beam and you load a Bose-Einstein condensate at the edge of the cloud, and now the condensate is accelerated, sucked into the laser beam. So now the atom has kinetic energy. The question is, where does this kinetic energy come from? And this time, we do not have spontaneous emission as we have when we do frictional cooling. So the question is now, where does the energy of the atom come from?

AUDIENCE: But if we were talking about a dielectric medium, it's coming from those light waves that are reflected or not reflected.

PROFESSOR: If comes from those light rays, but-- the atom is a dielectric medium, that's good-- but the reflection and deflection of light rays-- at least in leading order-- you would say photons of the same energy are deflected--

AUDIENCE: Wouldn't the [INAUDIBLE] of the atom just get absorbed [INAUDIBLE] and the [INAUDIBLE]?

PROFESSOR: I would hate to use the picture of absorption and emission. What really happens is it's just scattering. It's scattering. you scatter a photon which has one frequency into a different direction. It's really light scattering. Well, you have a homework assignment on that, and I don't need to go into details because it can easily be worked out, but what happens is the following. The atom is a dielectric medium. It has an index of refraction. And if you put an index-- if you suddenly put an index of refraction into a laser beam-- then you change the phase of the laser beam. The

atoms act as an electro-optical modulator, and they change the frequency of the transmitted photons.

You have a homework assignment on that. So if you increase the number of atoms in your laser beam, because they move in, this is-- the phase of the transmitted photons is shifted. If the phase of the photons is shifted during a certain time, phase over time is frequency. So therefore you will see, when the atoms move in and gain kinetic energy, the transmitted photons of the laser beam have lost some energy. They are slightly shifted to the red. So in other words, it's funny but if you have a dipole trap and the laser beam is focused and atoms slosh through the trap, if you would carefully analyze the frequency of the transmitted beam, you will find that it's a little bit blue, a little bit red, a little bit blue, a little bit red, and this compensates for the kinetic energy. Again, I don't think you'll find it in any textbook. It confused me for a while until I could work it out. But the homework assignment gives you an exact dissoluble model where you can exactly show what I told you.

So these are the two situations, here. It's either spontaneous emission, which is responsible for cooling, but here we have a conservative potential. And in this conservative potential, the only player is the laser beam which is transmitted. And, indeed, it is the photons which have been transmitted, and usually you wouldn't put that into your picture. Remember, the dipole force can be explained by eliminating the coherent field of the laser field-- for canonical transformation-- and then we have a C number. We have a classical electric field. This is your Hamiltonian, and from this Hamiltonian, you derive that you have a conservative potential. But if you now ask where the energy goes, the energy-- really, in this situation-- goes into a frequency shift of the photons. Also, in your Hamiltonian, you have assumed that the electric field is an external classical field, E naught times cosine omega T.

Well, if you want my advice, whenever you use the AC Stock's shift for trapping, just use the AC Stock's shift and regard it as a classic potential, and just completely forget that there are photons behind because if you want to account what really happens to the photons-- what really happens to the energy in this system-- it really gets ugly and it can get really confusing. Questions?

OK, we are done with the stimulated force. We are now in preparation for next week when I want to tell you about degenerate Fermi gases and degenerate Bose gases. Those gases require nanokelvin temperatures, so today, in the next hour, I want to tell you what are the techniques to create such ultra-low temperatures, and the techniques I want to discuss is one addition to laser cooling. We have so far discussed laser cooling to the Doppler limit, but now I have to tell you that that's not where laser cooling stops. You can go sub-Doppler and sub-recoil.

And, at least when the Nobel Prize was given to Bill Phillips, Steve Chu, and Claude Cohen-Tannoudji, well, if you go to the Nobel website and read the report of the Nobel Committee-- well, they should have given it to those three people for the many, many wonderful things they have done, and I've made frequent references to their work. Bill Phillips invented Zeeman slower, Claude Cohen-Tannoudji invented the dressed atom picture, they have many, many contributions, but the Nobel Committee tried to justify their choice with a coherent story, and the coherent story was Steve Chu cooled to the Doppler limit in molasses, Bill Phillips discovered sub-Doppler cooling, and Claude Cohen-Tannoudji realized sub-recoil cooling. So the storyline was cold, colder, the coldest. And that's what we want to talk today.

So sub-Doppler cooling. Well, I could have spent-- or, 10 years ago, I spent a whole lecture on sub-Doppler cooling, polarization gradient cooling. Wonderful! The kind of epiphany of elegance in [INAUDIBLE] a mechanical description of an atom. But I have to say-- and then I explained, in another lecture, sub-recoil cooling, how you can even cool below the recoil limit. Well, this year, 2013, I spent 15 minute on it. The reason is the following. When I joined atomic physics and was a post-doc in the '90s and assistant professor, the conference's [INAUDIBLE] of [INAUDIBLE], [INAUDIBLE] was full on intents to find new ways to laser cool. Ideas of how to get to lower and lower temperatures. There were many, many different-- sub-Doppler, sub-recoil techniques were discussed. This was the main topic.

But then suddenly, in 1995, evaporative cooling-- an intellectually boring cooling technique, just have atoms collide and evaporate-- this led to the lowest temperature ever, and it was almost a sudden transition. Within a few months there

was no research, no papers anymore, on advanced methods of laser cooling. Evaporative cooling it just completely wiped out this area over atomic physics.

The reason being because all cooling schemes which have been discussed had, in the end, some problems. At some point, photons heat. Heat because they excite an atom and if an excited atom collides with another atom, there is a heating mechanism and such. So even all the optimistic proposals for laser cooling reached very low temperature, but never at high density. And evaporative cooling just did everything for you. But anyway, I think it-- not just because a Nobel Prize was given for sub-Doppler and sub-recoil cooling, this is really an accomplishment to understand how can you go through conventional cooling limits, and at least every graduate student who happens to graduate in AMO physics at MIT should know what are the concepts behind sub-Doppler and sub-recoil cooling.

So what I need in order to introduce those cooling methods for you is I have to remind you that optical pumping is a cooling scheme, and one could actually say, to some extent, every cooling you do with lasers is based on optical pumping. And so let me explain that. Optical pumping happens that-- let's say you have a level structure with a circularly polarized light where you can go maybe from M equals 1 to M equals zero. And then there can be spontaneous emission back to M equals 1 and back to M equals minus 1, but the laser-- because it has angular momentum-- cannot excite this state to any excited state.

So what happens is, of course, pretty clear. You have whatever distribution you have in state one and two. You switch on your laser, and after a few cycles everything has fallen into the dark state. You have 100% population in stage two. So this is the simplest example of optical pumping using the three-level scheme. Well, you may ask now, what has that to do with cooling? Well, temperature is a Boltzmann factor, and if you-- we can define a temperature by saying the population between two levels is given by a Boltzmann factor. Let's just introduce an energy-splitting delta E-- which is somewhat arbitrary here-- but in any event, you see if you completely pump out a level, if you have all the population in one state, this corresponds to zero temperature.

So if you optically pump the atoms, preferentially, into certain states, you have lowered the entropy of the system, the atoms are no longer distributed over as many states as before, and this corresponds to lower temperature. And the message I want to give you is that you can understand laser cooling as optical pumping in translation space in velocity space. When you laser cool, you excite the atoms at high velocity and then because of the mechanisms we discussed, spontaneous emission leads to lower velocities, and this is one form of optical pumping.

I also like the word-- optical pumping is, in some sense, a spontaneous Raman process. You go up with one photon, you go down with-- you go up with the laser photon, you go down with the spontaneously emitted photon. So this is a spontaneous Raman process, and laser cooling-- even if you have just a single ground state-- can be regarded as a Raman process where the initial and final state differ in velocity or momentum. You have the same internal state, but you have a different external state. So you can see laser cooling is spontaneous Raman scattering between different momentum states.

OK, now how can we achieve sub-Doppler cooling? Sub-Doppler cooling was discovered in 1988, and I remember I was at the [? IKA ?] Conference in Paris when Bill Phillips' group said, we've carefully measured the temperature and in sodium we measure a temperature which is lower than the Doppler limit, which was rigorously derived. I derived for you the Doppler limit of molasses. So what went wrong here? Well, it became clear the only loophole-- if you have a theory which predicts something and then you find a violation, you carefully have to check the assumption. And the assumption which was made-- and which we have made in this course-- is that you had a two-level system.

Now, atoms have hyperfine structure-- and for pedagogical reasons, because I want to explain to you why in sub-Doppler cooling scheme I assume this hyperfine structure with F equals one half, F equals three half-- but let me first tell you what the novel feature is about a multi-level atom which has hyperfine structure. Now, the

one thing which happens is that-- instead of just going between our single ground state and our single excited state-- we now have transitions. You may want to call them Raman transitions, or optical pumping, between the hyperfine atoms.

So there are now transitions not only up to an excited state which rapidly decays, there are now transitions-- Raman transitions-- between ground states which have a very low widths. The width of the excited state is gamma. Well, what is the width of the ground state? Well, if you don't have a laser beam, the width is zero, but if you have laser beams the atom doesn't stay there forever. There will be a time, which is estimated here-- depending on the RB frequency and the detuning-- this is the time, this is simply nothing else than the scattering rate which we derived before. So that's the rate at which you scatter photons. And then, depending what the bunching ratio is, this is also the time in which you may optically pump to another ground state.

So therefore, what we have to take into account now is that we have narrower widths in our system, and those widths correspond to another two-level system which is driven by two photons in which connects the two ground states. I will not give you any derivation how this exactly leads to a lower temperature, but I want to give you two pictures which I have already introduced to you. One is, remember, when we discussed Doppler cooling, the final temperature was proportionate to gamma, the width of the transition. So now I wave my hands and say, if you have transitions between ground states which have a much, much smaller width-- which is the rate of optical pumping-- and for low laser power this width can become very, very narrow. But you have at least one ingredient which can lead to a lower temperature.

OK, in all truth in advertising, for the Sisyphus-- for the polarization gradient cooling scheme-- the final temperature is not given by gamma prime, the rate of optical pumping. It is a proportionality factor, but there is another factor. If you want to learn about it, you have to read some of the classic papers. But there is another picture which I can also use to tell you-- and that's actually related-- why optical pumping at least gives the possibility for lower temperature. I explained to you with the

stimulated force in [? blue ?] molasses-- that the only reason why we have cooling-- is that the atom has a lag time. It cannot instantaneously adjust itself to the light field. It lags behind. And you can say, in Doppler cooling, the lag time is the spontaneous emission time, and the inverse of it is gamma.

Here, in this case, the lag time can be the very long time to optically pump. So therefore, if the atom is in one hyperfine state and now moves into an area where the polarization of the laser beams is different, it may take a long time to adjust. And it is this lag time to which the friction coefficient was proportional. So this is a new feature, long delay times, narrow widths, and resonances between ground state levels.

OK, a little bit show and tell now. Nobody thought about polarization creating cooling, nobody thought about Doppler cooling, but it was discovered when people simply did Doppler molasses and it's one of those big violations of Murphy's law where cooling worked much better than everybody had thought, and the result is the following. That-- when I drew for you the blue curve, which is force versus velocity-- it's pretty much the subtraction of two [? Lorentzian ?]. What happened is-- for sodium atoms, or for any alkali atom-- it's the red curve. For large velocity you have the Doppler cooling mechanism, but for smaller velocities you have a steepening of the slope, and the steeper the slope, the larger is your friction coefficient, alpha. And this is the mechanism of polarization gradient cooling.

Let me just take this famous paper and show you one way how polarization gradient cooling works, which is the most famous form of sub-Doppler cooling, and this goes as follows. If you have molasses with two laser beams and the two laser beams are lin-perp-lin, linear polarization perpendicular to the other linear polarization. So you have two laser beams with these polarization. When these two polarizations overlap and they have the same phase, you get light at 45 degree, but if these two polarization have 90 degree out of phase, you get circularly polarized light. So as these two laser beams counter-propagate, you periodically go from linear polarized sigma minus, linear polarized sigma plus.

So you have spatial-- at any given point you have a polarization, but the polarization changes. You have a-- you would say, naively, lin-perp-lin don't interfere, don't form a standing wave. Well, they do not form a standing wave in intensity. They form a standing wave in polarization. So now what happens is the following. If you have a multi-level atom and I use this simple scheme here, the different polarizations, linear polarization drives, the pi transition, sigma plus and sigma minus [INAUDIBLE] different transitions, and those transitions have different strengths.

So an atom here-- when it experience this polarization-- reaches cycle, but when it experiences the other circular polarization, it will be optically pumped over here and then it cycles. So therefore, an atom which from sigma plus polarization and was here, it flies over to an area where you've sigma minus polarization, it will actually be pumped over. So the atom will constantly be pumped back and forth between those hyperfine states. And this actually gives rise to a beautiful form of Sisyphus cooling that the atom experiences sigma plus light in one ground state. It climbs up the hill, it sees the AC Stock's shift potential-- and the AC Stock's shift potential, because of the, for instance, for one hyperfine state, sigma plus drives the strongest transition-- and therefore we have sigma plus light, one ground state has the deepest potential. Where we have sigma minus light, it is the other ground state which has the deepest potential.

And so what happens is that the atom is in one hyperfine state. It climbs up the hill, and then it's optically pumped to the other hyperfine state, and we have exactly the same kind of Sisyphus cooling. I just show you pictures and you sort of match it with what you know. This is more complicated because this involves optical pumping between hyperfine states. The different hyperfine states have different [INAUDIBLE] coefficients and such. It's not really complicated but more complex. What I explained to you is how Sisyphus cooling works in the dressed atom picture just for a two-level system. And here you find a more subtle form of Sisyphus cooling, but this form is more important because whenever you operate a magneto-optic trap, you get this cooling mechanism for free.

OK, so that's all I want to tell you about sub-Doppler cooling. OK, sub-recoil cooling,

we can quickly deal with it because I want to prove to you here that sub-recoil cooling is impossible. And unless you tell me what this wrong in my derivation, I don't need to discuss sub-recoil cooling because I've convinced you that it's impossible .

OK, so let's assume our atom has an initial kinetic energy and then we absorb one photon. That would mean the momentum has now changed by the laser photon by the recoil-- and then it emits a photon, so the momentum gets reduced by the immediate photon-- and all I do is I ask, what is the difference between the initial kinetic energy and the final kinetic energy after two photons-- a laser photon and the spontaneously photon-- have been exchanged? And what you find is-- and this is just exact-- that the change in energy is two times the recoil energy. The recoil energy is h-bar square K square over 2m, plus-- and this is now the cross-term-- plus KL-- the K vector of the laser-- minus K spontaneous emission times the velocity. And now, of course, we assume-- which is correct-- that spontaneous emission goes randomly in all directions. Therefore, if we average over many cycles, this does not contribute.

So therefore we find now that the average energy exchanged by an absorption and emission event is two times the recoil energy plus h-bar K laser times [? width ?], and we want to make it negative because we want to cool. Well, we make it negative by arranging the laser beam-- surprise, surprise-- counter-propagating to the velocity. This is how we can get the best cooling. Well, surprise, that's what we expected. But you find, of course, that when this velocity is a very small then K V is smaller than K squared. So therefore, once the velocity is smaller than the recoil velocity of a single photon, you cannot choose-- you cannot make this equation or this expression go negative.

In other words, what I've shown to you is if there is an atom which has a velocity which is smaller than the recoil velocity of a photon, whenever this atom scatters a single photon, it will be hotter than it was before. So is it clear, what I've shown you, in energy? When an atom has a velocity which is smaller than the recoil velocity, any further photon scattering will not cool. It will lead to an energy transfer delta E,

which is large, which is possible.

So that shows you that sub-recoil cooling is impossible. Any idea how we can sub-recoil cool? Any idea why Claude Cohen-Tannoudji got the Nobel Prize?

AUDIENCE: Further down in the lattice, you can have an effective mass that is much heavier than [? bare ?] mass?

PROFESSOR: OK, great idea. We put the atom in a lattice, a lattice has band structure, in a band structure we have an effective mass, and heavier atoms can be cooled to lower and lower temperature. Actually, it's well known that, in Doppler cooling and sub-Doppler cooling-- especially in sub-Doppler cooling-- cesium an rubidium reach a microkelvin, sodium reaches only 25 microkelvin. So making the atom heavier is OK, but I have to say you're only rescaling your recoil limit, you're not breaking through the recoil limit which is now defined with a heavier effective mass.

AUDIENCE: [INAUDIBLE] emission direction--

PROFESSOR: Yes?

AUDIENCE: --means that the atom must have [INAUDIBLE] responding is [INAUDIBLE]?

PROFESSOR: Yes.

AUDIENCE: And [INAUDIBLE]?

PROFESSOR: That was the assumption.

AUDIENCE: [INAUDIBLE] at least for this cooling, it needs the time much larger than the Doppler cooling [INAUDIBLE] spontaneous emission time?

PROFESSOR: Well, that was the assumption here. We scatter several times, a spontaneously emitted photon do not contribute the recoil of the spontaneously emitted photons cancels out because spontaneous emission-- and this is correct, this is not wrong-- spontaneous emission goes equally probable in the plus X and the minus X direction.

Well, I thought I want to show you a demonstration how-- mechanical demonstration-- how you can sub-recoil cool. It's not a demonstration in velocity space, it's a demonstration in position space. What I have here is a Plexiglas tray and it has a little hole in the middle. And I have a bunch of ball bearings and I put those ball bearings in. Let me translate. I can blindfold myself and all I do is I shake the tray. So whenever I shake it, I kick the atoms randomly in position space from here to here. So the amount of position control I have over the atoms is on the order of this size.

But the question is, can I-- without having any control of the transfer of position I give to them-- can I steer all the atoms into an area here into the hole in the middle, which is very, very narrow in position space? So the translation is, if I randomly scatter photons and they kick the atoms around with h-bar K, I don't have any control about momentum transfer smaller than h-bar K, is it possible to localize the atoms in momentum space to a momentum around zero, which is much, much smaller than h-bar K?

Well, let's do the experiment. I just close my eyes and I shake it, and I shake it, and I just shake it for a while and I continue shaking it. And, well, zero temperature! So without controlling the motion on the scale I'm interested in, I manage to cool into a target region which is much, much narrower. So in other words, what you need is-- all you need is-- you need some dark state which is velocity selective.

If you scatter light but you create a situation where once the atoms are close to zero velocity, they're not excited. You're not steering the atoms in a deterministic way, you just wait until-- by random chance-- one photon is emitted, and by the random chance the atom hits the hole. And then, it will never be re-excited again. This was the idea of this demonstration. And I could, yes, in the old days I may have spent two hours on teaching it, but by using Raman resonances between ground states-- which are terribly narrow-- or by using VSCPT-- velocity selective coherent population trapping-- you can create such narrow dark resonances which have the effect that the atoms scatter, scatter, scatter, but the moment they reach a very narrow region around zero velocity, they stop scattering.

So what was wrong in my proof is that here we have a situation where we stop the laser cooling exactly at the time when the atom, by random chance, happens to be at low velocity, at very low velocity. And then you don't need any control, you just accumulate in the same way as I accumulated the ball bearings.

OK, this was sub-recoil cooling in five minutes. [? Timo ?]?

AUDIENCE: So I wasn't here last class. But just to summarize, the Doppler cooling rate gets us to a temperature on the order of gamma, which is natural alignment.

PROFESSOR: Yes.

AUDIENCE: And the recoil limit is the recoil energy which is usually tens of kilograms. But if you have a really, really narrow line, would you use the Doppler cooling limit? You could, in principle, beat the recoil limit with Doppler cooling, no? If you have a--

PROFESSOR: No, you can't, because if you do Doppler cooling-- [? Timo's ?] question is, if you simply do molasses with a very, very narrow line, we derived that the limit of Doppler cooling is Kt equals gamma, and what you are saying now, if you use narrower and narrower lines, can we reach [? arbitrary ?] low temperature? This is not the case. In our derivation of Doppler cooling, we made a continuum assumptions when we plotted the force versus velocity and we had the friction force, we assumed that an atom, when it scatters a photon, stays, let's say, within the linear part of the force versus velocity. So we had a hidden assumption which required that the recoil energy is smaller than h-bar gamma. But there is a lot of literature where people looked into it, and it turns out that when you have a very narrow line, you can go to the recoil limit, but you can't go beyond.

If you want to go-- and I think this is what I showed you in the last five minutes-- if you want to go below the recoil limit, you need some velocity-selective dark state, and you would not have that in this situation. On the other hand, if you have a very narrow resonance, you can probably engineer a dark state that there is destructive interference in some excitation in a very, very narrow velocity class, and this narrow velocity class is simply selected by the Doppler effect. So you can use the narrow

line to engineer sub-recoil cooling, but the simple arrangement of just having molasses with narrower and narrower line will not work. [? Collin ?]?

AUDIENCE: I think the group that laser cooled to BEC just had-- their dark state was just a Stock's shift, [INAUDIBLE] and that was their-- and it shook. There were resonances in kilohertz or something.

PROFESSOR: Thy used a very-- did they go below the recoil limit in laser cooling? I'm not sure. They may have had a situation where it was enough to go to the recoil limit. I have to check out the paper, but--

AUDIENCE: I remember they did have some sort of dark state. Maybe it was to increase density or-- they had the high--

PROFESSOR: Anyway, let me just be clear. Sub-recoil cooling requires that you have a dark region in velocity space. For that, you need some narrow line widths. That's necessary. But I think you just don't get it by having counter-propagating laser beams, you have to do something more. Other question?

OK, so I've told you about sub-Doppler cooling, sub-recoil cooling, but what really got atomic physicists to nanokelvin temperature were two other techniques. Laser cooling was used as pre-cooling, but then the final trapping and cooling was done by magnetic trapping and evaporative cooling. So I want to give you now, in the last half hour, a quick overview of magnetic trapping and evaporative cooling.

So I think any discussions of magnetic trapping starts with a theorem, a theorem which tells you that not everything is possible you would like to do. And this is the following, then if you have a region without charges and currents without-- through an empty space, and what this theorem says-- in empty space, you cannot have a local minimum-- you cannot have a local maximum of your electric and magnetic fields. So if you take the strengths of the electric field or the strengths of the magnetic field, you cannot have a maximum. You can only have a minimum.

Well, and this is important when it comes to magnetic traps. Depending on the magnetic moment-- spin up, spin down-- we could create a trapping potential

around a maximum of the magnetic field or the trapping potential around the minimum of the magnetic field. You want to be at the minimum of the potential, but if you have a minus sign, you may want to be at the maximum of a magnetic field. But [? Bing's ?] theorem says that only one of them is possible. The proof really goes in two lines. You really show, if you assume that there is a maximum, that would mean that if you add a field to it, you add a field-- and no matter in which direction you add a field-- the total field strengths become smaller. But that requires a violation of the Laplace equation.

It's easier for you-- you can probably read it at home-- but you simply assume you have this situation, so you add-- and you show you have an incremental field, delta E. This field, delta E, you are at the maximum of an electric field, and then delta E over R is the small difference of the electric field. And you can show that this expression is only negative. That means you have a maximum of the electric field when the electric field would not fulfill the Laplace equation.

So the gist of the argument is the following. Each component of the electric and magnetic field fulfills the Laplace equation. And the Laplace equation-- I don't know if you've probably heard about it-- if something fulfills the Laplace equation, it says at any given point the function which fulfills the Laplace equation is equal to the average of function, averaged over a small sphere around it. And that means-- if something fulfills the Laplace equation-- you cannot have a local maximum or a local minimum because if the value here equals the average, that would mean in one direction the value gets higher, in the other direction the value gets lower. So if something fulfills the Laplace equation, you cannot have a local maximum or local minimum.

But now, we're not asking for maximum or minimum in one component of the electric or magnetic field. We ask for a maximum or minimum in the total value, or the square of the electric field. But then you show, when you assume that you have a maximum, the field points along the Z direction, that then the Z component of the field cannot fulfill the Laplace equation. So maxima are not possible. Minima are possible. Good!

So the magnetic trapping potential comes because we have a magnetic dipole moment which interacts with a magnetic field. Magnetic traps are actually classical. You can have a classical model for a magnetic trap, and in that case, you say this potentially is mu times B times cosine theta. [? Crendo ?] mechanically, of course, the angle cosine theta is quantized and we have the different M F levels, the different orientations of the spin, relative to the Z-axis.

OK, the fact that the magnetic field can only have minima and not maxima means that we want to make sure that the magnetic moment times the G factor times M F is negative, or cosine theta is negative in the classical picture. And if we absorb whatever we choose for cosine theta-- or what we call mu Bohr G M F, the magnetic moment-- we want that the magnetic moment of the particle is anti-parallel to the magnetic field. Then we have a magnetic trapping potential, which is-- I've taken care of the sin now-- mu times B, and if B has a minimum this potential has a minimum, and that is a magnetic trap.

But so the consequence of [? Bing's ?] theorem is that we can only trap particles which are anti-parallel with a magnetic field. And they can always lower their energy by flipping to the other state. So therefore, a magnetic trap does not allow us to trap particles in the absolute ground state. There is always the possibility of spin-flip collisions which lead to an anti-trap state which is expelled from the trapping region.

Now, spin-flip collisions can happen when you have particles at sufficiently high density and they collide. For the experts, there can be spin relaxation, there can be dipolar relaxation, they are two different kinds of spin-flip collisions. Fortunately, in magnetic traps, they only become relevant when you pick the wrong spin state. But for suitable choices, magnetic traps are very long-lived.

Dipolar relaxation is often-- can be a limiting process for Bose-Einstein condensates, for instance, for high [? dose ?] Bose-Einstein condensation, dipolar relaxation was-- limited the number and density of atoms in the Bose-Einstein condensate. But I don't want to talk about cold collisions, here. You should just know, based on the fundamental theorem, you cannot trap in the ground state. You

have to trap into a state which has more energy than other states, and, in principle, it's possible to flip the spin and go to a lower state.

Now, what we have assumed here is-- when I wrote the trapping potential like this-- I assumed that the atom stays in a given hyperfine state. So we put the atom in one state and it experiences potential. Classically it means that the angle of the dipole, with respect to the magnetic field, stays constant. And this is the case due to rapid precession. Classically, the magnetic dipole precesses around the magnetic field, and when the direction of the magnetic field changes, the rapid precession keeps the dipole aligned. Or, in other words, for slow changes of the magnetic field direction, cosine theta is an adiabatic invariant. So the word adiabatic is important, either classically or you stay in a quantum state as long as you're adiabatic.

Well if an atom would rapidly go through a region of very weak magnetic field, then the precession-- the [INAUDIBLE] precession-- of the atom is very slow, and this adiabatic condition can be violated. And this violation of an adiabaticity condition is called Majorana flops. So if you operate a magnetic trap at very low magnetic field, you may destabilized the trap because you violate adiabaticity. OK, so these are the consequences of the fact that we cannot create maxima, we can only create minima of the magnetic field in free space, and therefore we have to deal with those two possible loss processes, but we have learned what kind of magnetic trap, what kind of atom to pick, and in general these are not big problems.

OK, so we have to-- so this shows here the typical hyperfine structure of an alkali atom. It could be rubidium 87 or sodium 23. The fact that we have only a local minimum of the magnetic field means we can only trap hyperfine states-- which I've marked here in green-- where the slope is positive, and these are the atoms where the magnetic moment is inter-parallel with the magnetic field. And, well, for stability reasons we usually pick the highest state of the lower manifold, or the highest state of the upper manifold.

In principle, it would be also possible to trap here-- you have the correct slope-- but often those states suffer form spin relaxation and collision loss. [? Mickey ?]?

AUDIENCE: [INAUDIBLE] the picture you showed before, is it actually possible to make a magnetic trap with state number four from the top so that they expel from the center? From [INAUDIBLE] a--

PROFESSOR: Yeah, there is a peculiarity here, you have the transition from the weak field to the strong field region. Whenever the magnetic moment is constant, you need a local minimum of the magnetic field. But if the magnetic moment would change and you have a gradient of the magnetic field, you could actually trap here. So in this case, you have a magnetic field gradient but you have a spatial variation of the magnetic moment. This has been discussed in one special paper in the literature, but it has not really found any [? lucent ?] realization.

AUDIENCE: Wouldn't it help to put the Majorana flops in the center? Because they're repelled from the center.

PROFESSOR: Well, yeah, but these atoms would undergo spin-flip collisions so it's not a good choice. It would help against Majorana flops, but we have many solutions against Majorana flops. And here, you would solve the Majorana flops problem but you would solve it with another problem.

OK, so we need minima of the magnetic field. This provides magnetic trapping potential, and there are two kinds of possible trap configuration. One is where the minima is at zero magnetic field, and one is where the minima is at finite magnetic field. Now, the zero magnetic field minimum one can be simply created with anti-Helmholtz coil. What happens is the magnetic field, as a function of position, would actually cross through zero. It's just a field gradient which crosses through zero, but, of course, since we are interested in the absolute value of the magnetic field, we get the V-shape potential.

This V-shape potential was what was used for the first demonstration for Bose-Einstein condensation because a V-shape potential is much, much more confining in this cusp than in harmonic oscillator potential. And this had advantages for tight confinement and rapid evaporative cooling. So those quadrupoled traps with the V-shape potential are the best confinement you can get, the best confinement for the

buck, and the buck here is your power supply and your coils.

However, they have a problem when the magnetic field is zero. We violate adiabaticity because the different spin configurations become degenerate, and in degenerates you can't have adiabaticity. So therefore, the two first demonstrations of Bose-Einstein condensation avoided this cusp in two different ways. One way was to use a rotating magnetic field, and to use a perpendicular field but which was time-dependent, and the MIT solution was to use a blue detuned laser beam, use the optical dipole trap-- optical dipole potential-- to push the atoms away from the dangerous [INAUDIBLE] field region.

I should say rotating-- none of those-- most traps which are now used are the other kind of trap, the trap which has an harmonic oscillator potential, which doesn't have the cusp and where the minimum is at a finite magnetic field. Let me just make one comment. The rotating trap was very popular because it led to the first BEC and a lot of people built that, but to the best of my knowledge, it's only used in different places. And the only real application it has is-- since you have some rotating magnetic field-- by some modification you can actually make of rotating potential. And this is nice if you want to study Bose-Einstein condensation on a rotating frame, create vortices and things like that.

But for simply creating a Bose-Einstein condensate, this trap is used, by far, most frequently. But there is actually a renaissance of this trap. One reason why people use this trap and why my own research group immediately switched to this trap-- and we had the idea how to built it-- is, well, harmonic potential is nice. Every physicist loves an harmonic potential. You can solve, immediately, thermodynamics in a harmonic potential. Who wants to deal with that potential? Also, if the laser beam just drifts by a micrometer, this symmetric potential becomes asymmetric. So you characterize your potential today and a few hours later you have a different potential. Whereas a magnetic potential, harmonic potential, once it's characterized-- as long as you don't change the current in your power supplies-- it's the same week after week, months after months.

Well, there is a renaissance now because a lot of groups now doing experiments in optical lattices, they don't care what shape the potential is because they don't do physics in this potential. They just take the Bose-Einstein condensate or the cooled Fermi gases and transfer them to an optical lattice. So in that case, it doesn't play any role, and then the advantage of this potential is you get more bang for the buck.

AUDIENCE: [INAUDIBLE] you were saying before [INAUDIBLE] more rapid evaporation in V-shape versus the [INAUDIBLE]. Where is the differences in evap time?

PROFESSOR: It depends. I may run out of time today. I needed-- I can probably-- I wanted to go through evaporative cooling, and I will talk to you that, for evaporative cooling, there could be a threshold in confinement where you go into a runaway regime that evaporative cooling is speeding up. So a little bit of confinement can make the difference without never getting into a runaway regime or being in that runaway regime. So confinement, extra confinement, can make the difference between getting BEC and not getting a BEC. So it really depends. It's highly nonlinear, but I will show you later on that when you put a system together, the kind of threshold density at which the cloud can evaporate to BEC is much lower here than it is there.

But-- and this is the next thing I want to explain-- it's not really that this is a linear potential and this is a [INAUDIBLE] potential because those finite magnetic field-- these traps which have a minimum at a finite magnetic field-- are usually done in the following way. Those traps are called Ioffe-Pritchard trap. Ioffe actually suggested such a magnetic field configuration for confinement of plasma, and Dave Pritchard was the first to point out that such a configuration would be a good choice for neutral atoms.

So I sometimes joke and I say, well, plasma physics is the study of very hot matter. Cold atoms is the study of very cold matter. But when matter is either too hot or too cold that you cannot put it in ordinary container, you want it confined with magnetic fields. And it happens that, for plasma, you need a minimum of the magnetic field, and for neutral atoms you need a minimum of the magnetic field. So therefore here is something which ultra-cold atoms have in common with plasma physics. I don't

think there is a lot the two fields have in common, but when it comes to magnetic field configurations, yes. A similar magnetic field configuration can confine a plasma and can confine neutral atoms.

OK, so the generic way of how these magnetic fields are generated is you want to have two coils, called pinch coils. You can say each of them creates a magnetic field which decays, a magnetic field which decays, and now in the middle you have the parabolic minimum. But then you have to add the so-called four Ioffe bars which create a linear potential, and this is shown here. So the pinch coil is simply creating a local minimum along the Z-axis, and if you only want to trap in one dimension, you would be done, But you want to trap in 3 dimensions. And so what happens is you add now an anti-Helmholtz-- or quadrupold-- field in X and Y, and this blue field is done via the green bars, by the Ioffe bars.

So you add together a harmonic quadratic field in the Z direction with the linear quadrupoled field in X and Y. And what matters for the atoms is-- and what matters for the Zeeman energy is-- the absolute value of the magnetic field. So you add those things in quadrature, and because you add it in quadrature, you have now an harmonic oscillator potential in X, Y, and Z, so you have an harmonic trap in three directions. However, if this field is not larger than this-- if these field becomes larger than this field, so then when you add it in quadrature, you actually get a linear potential.

So in other words, the Ioffe-Pritchard trap is quadratic for smaller values of X and Y, but if you go out to large X and Y, you get the linear potential. So in the end-- and this is maybe in response to Matt's question-- for a high temperature, the Ioffe-Pritchard trap is actually not harmonic in three dimensions. It's linear in two dimensions and harmonic in the third dimension. So, therefore the hit you take in evaporative cooling is not as large as we'd initially assumed. And the moment we realized that, we were building a Ioffe-Pritchard trap. And this has been the traps we' have been using at MIT ever since.

Well, just as a warning, this is sort of the simplified description, but if you have

curvature, the curvature has to fulfill Maxwell's equation. And you cannot have a curvature only along Z. Maxwell's equations are three-dimensional, so you get sort of all sorts of curvature terms. If you really build a magnetic trap, you should understand those. If you just want to understand why magnetic traps work, the previous slide is sufficient.

The different ways to build those traps, this is the design we invented at MIT. It's called a clover-leaf trap. So you have one coil package here, to another coil package here, and what you see is the pinch coils, and when you have two of them you create this parabolic minimum. We don't like Ioffe bars because if you have a vacuum chamber you either have to put the Ioffe bars into the vector vacuum chamber or you drill holes through your vacuum chamber to put the Ioffe bars back and forth.

Don't laugh, some people have done it. Of course, they put little tubes around it so it was a very highly-engineered vacuum chamber where they could string Ioffe bars through the chamber. But we realized that the same field as Ioffe bars can be generated by taking the Ioffe bars and flipping them out, and after flipping them out, they had the shape of clover leaves. It's just a variant of creating the same field geometry.

Yes, that's all I want to tell you about magnetic trapping. I think I should not start with evaporative cooling. This will just take 10 or 15 minutes, but I'll do that on Monday. So since we're on time, do you have any questions about magnetic trapping? Different forms of magnetic trapping? OK. Final announcement, next week is the due date for the term paper. Yes, and the term paper is due on the date of the last class, which is Friday. Any questions about that? OK. Good.