

[SQUEAKING]

[RUSTLING]

[CLICKING]

JACK HARE: So just briefly, there was some discussion last time about Lyman sequence and the Balmer sequence and how these correspond to the K shell and the L shell. And I just looked this up, and I thought for your amusement and entertainment I would explain it.

So the idea here is that we are looking at hydrogen. So the Lyman, the Balmer series only apply to hydrogen. And we have levels of principal quantum numbers, 1, 2, 3, 4, and things like that. And there were two series which are historically seen. There, first of all, was a set of lines which turned out to correspond to transitions down to-- sorry-- n equals 2. And these are the Balmer series.

These are transitions to n equals 2 from some other level like this. Confusingly, the symbol that we use to describe these is h . So we have, for example, h alpha, which is n goes from 3 to 2. h beta is n goes from 4 to 2, and so on. And the reason that these were discovered first is that h alpha is actually in the red. It's in the visible.

And the other lines are increasingly high energy, as you can tell, because we're doing a transition from a higher and higher energy. And in fact, after h beta, these go off into the ultraviolet. So they're very hard to see in early experiments.

The other sequence of lines which seem more fundamental, but were discovered later are the transitions down to n equals 1. These are the Lyman series. So this is n goes towards 1. And these are usually called Lyman alpha, Lyman beta, and so on. So that's transitions from n equals 2 to 1, and 3 to 1, and so on.

And even the lowest energy of the Lyman alpha series is deep in the ultraviolet. So it took much longer for people to spot this. And so this is the Balmer series came first, and the Lyman series came later. So you're looking at this and you're saying, OK, this does look at an awful lot like what you told us was the K shell and what you told us was the L shell here.

And apparently, the pedant's answer is when we defined the K shell, the way we were thinking about it involved some other electrons here. And so you need to have greater than three electrons left for both of these definitions. And so that means that you have to have at least lithium-like atoms.

But I would say, frankly, these look almost a pretty obvious continuation of the L shell and the K shell stuff. There's also a series of lines for helium lines as well, which are for two electron atoms. These are for one electron.

So my short answer would be, yes, the Lyman series corresponds to the K-shell emission and the Balmer series corresponds to the L-shell. And then, of course, you have all the fun things like L shell, Lyman, H K. Spectroscopy is awful for that sort of stuff. So it's just nomenclature. Any questions on that before we move on to line broadening, the reason we're here?

So we're going to be talking about line broadening. So far we've mostly been considering line emission, which is at a specific frequency. We've said that we have some lower energy level and some upper energy level. And we have some transition of electron from an upper to a lower. It could be stimulated, or it could be spontaneous emission, and emit a photon with energy $\hbar\omega_{upper\ lower}$.

So what we'd expect to see in our spectrum, because we have lots of different levels available, is a spectrum that has a series of very sharp lines. Their spacing is determined by these energies. And the height of them is determined by the probability of transition. And in our plasma also, the occupation of the upper energy level. So how many atoms are in that excited state. And you with an electron up here have a chance to decay down.

Now, in reality, when we look at it, our spectra does not consist of these delta functions. It has lines, some of which may be very broad, some of which may be very narrow. And so there are several different mechanisms, which we'll go through one after the other today, which broaden out these very narrow, single energy peaks into a range of different energies here.

So the first of these is the one that we absolutely can never escape. It's always going to be present in our plasma. And this is called natural broadening.

The idea of natural broadening is that our electron does not immediately transition down into a lower energy level. It has some finite dwell time that it will stay up here. So the transition is not instantaneous here. So we have what we call a finite excited state lifetime.

And we assume this is going to be some sort of random process. It's not like all of the excited energy levels decay simultaneously after 1 millisecond. They're going to decay at various different times. So it looks a little bit like, for example, radioactive decay, where as a function of time, the probability of being in the upper state is going to decay away. And that probability is going to be proportional to exponential of minus time over some time scale τ .

And that τ here is what we refer to as the finite excited state lifetime. So different states, some may have a short τ . They may decay very quickly. Some may have a long τ . They may decay very slowly.

So we want to have some handle on what this τ parameter is. And so we make a very hand-wavy appeal to a sort of uncertainty principle here, where we say the energy of our line is $\hbar\nu_{upper\ lower}$. Again, I'm switching between this in hertz and this in radians per second. But the same thing here.

And we say there's going to be some uncertainty in energy, which is going to go like \hbar over 2π this excited state lifetime. So this is effectively a bit like having $\Delta E \Delta t$ is greater than \hbar upon 2π , so greater than \hbar like that. So this is an extension-- this sort of is an uncertainty principle type argument.

You actually don't need to invoke the uncertainty principle in order to get this. But some people this way of looking at it. So this means there's some uncertainty in the energy of the line related to the excited state lifetime. And so some of the photons will just come out with a slightly different energy. So that should be enough to explain the broadening. But we'll do it a little bit more rigorously in a moment.

The other thing I want to say is, yes, how to actually calculate this τ here. This is one way. Another way to do it is to say that the intensity of the emission i is going to go like the probability of the decay squared. And that's going to give us a factor of 2 upon τ , which is equal to the sum over all of the spontaneous emission processes that take us from the upper level to the lower level, because, of course, what we're interested in is not just decay from this upper level to a specific lower level.

What do I mean by that? Because I cannot work out what this means in my notes, I'm going to delete this here. Possibly, there's a decay chain that includes some other processes, and we want to sum over those.

But in general, we're just going to be looking at, for example, in this case, spontaneous emission here. And this has units of 1 over seconds. So we can quickly identify the excited state lifetime with 1 over this. The factor of 2 just comes from the fact that when we're dealing with intensities, we square the probability. So that goes into the exponential here. And that's what this comes out here.

Now, we know that this A_{ul} has some relationship-- could someone close the doors, please? Thank you-- to the black body. This is how we were able to calculate this before. And this black body has a relationship like ν cubed.

So what we're saying is for the higher energy lines here, we're going to have a shorter excited state lifetime. And we're going to have a larger energy uncertainty. So ΔE equals $h \Delta \nu$. It gives a short excited state lifetime, and that gives a large range of energies. So in general, if you have i energy lines, they're going to have larger natural broadening.

So as I said, this is a slightly unsatisfactory way of getting relationship between energy and time. And we can do it more precisely by looking at Fourier transformations of this time series here.

So if we have a system, which is decaying in time-- no, I guess I've drawn it already. We have an intensity of time that has a proportionality exponential minus $2t$ upon τ . And we can Fourier transform this time-varying signal. And we will get out the spectrum that corresponds to that.

And that spectrum, does anyone know what the Fourier transform of an exponential decay is? It's got a name. It looks like whatever the intensity is of the unshifted line here divided by $1 + \nu^2 \tau^2$. Where's the squared go? Here.

This has a shape which is symmetric in frequency, because we're just squaring the difference between the frequencies here. And it looks like this. Does anyone know the name for this function? This is Lorentzian. Sorry? Yeah, this is Lorentzian.

What this is saying is that there is some width through the distribution of frequencies and the width in the frequency space, which is obviously also the energy space. The full width at half maximum is equal to $1/\tau$, which you can get by just setting what is ν when i is half of its peak value here, so making this denominator equal to a half.

And so what we find is that this full width at half maximum, which is proportional to $1/\tau$, which is proportional to ν^3 from our argument over here, means that as we said before, if we have higher energy photons, they're going to have more and more natural broadening here. So this we can't get rid of. This is just due to the uncertainty in when the electron is going to decay from the excited state to the ground state. And that naturally just leads to some broadening here.

And so you're always going to have some Lorentzian broadening to your system. And then you'll also have other broadening mechanisms on top of it. But this means that when you're trying to fit, for example, like trying to fit these lines of something like a Gaussian function, because you look at it, and you go, ah, that looks like a Gaussian.

Well, Lorentzians are slightly different from Gaussians. And your fitting procedure won't work properly. So you need to be careful when you're fitting spectral lines not to just assume that the spread is Gaussian distributed. There are some other functions, like Lorentzian, which are more fundamental. Questions on that first natural broadening mechanism? Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: Do you want me to give you intuition for quantum mechanics?

[LAUGHTER]

I don't think I've got a better answer to you than this mathematical derivation. I think the energy levels are not as exact. Because of this uncertainty principle, we can never be quite sure exactly what the momentum, and therefore the energy of the electron wave function is. There's some uncertainty in exactly what the energy of the function is. And that translates to this uncertainty in the frequency.

AUDIENCE: [INAUDIBLE]

JACK HARE: I feel like that's above my pay grade.

AUDIENCE: How much [INAUDIBLE]?

JACK HARE: Thank you. But if I can't answer this, then that's very embarrassing.

AUDIENCE: [INAUDIBLE] the probability of [INAUDIBLE]? What's the state we're talking about here? Is that the end state?

JACK HARE: No, the up-- sorry. So let's say this is 1 here at time t equals 0. We have whatever process has just promoted the electron up to here. This took place at time t equals 0. And then the down process takes place at time t greater than 0.

AUDIENCE: Or up [INAUDIBLE]

JACK HARE: Sorry. Yeah, so the probability that is still in that upper state. Yes, exactly. Yeah, I mean, in this-- if you're like, but it could decay to a different state, or it could go up again, then that wouldn't affect the shape of the line, because we're looking at a very specific line, which is the line with the energy $h\nu_{ul}$. And if you want to-- u and l-- you want to know the broadening for some other state, then you have to do this calculation with whatever other state you're going for.

Any questions online or other questions in the room? Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: OK.

AUDIENCE: [INAUDIBLE]. That doesn't get [INAUDIBLE], does it?

JACK HARE: Right. Yes, I agree. So we are dealing with an ensemble of different [INAUDIBLE]. Yes, thank you. That's actually a very good point.

So the point that's being made here is that, in fact, in reality, any given photon decay, the intensity looks like this. It's a delta function at a particular time. That's when the photon was emitted.

This treatment here is talking about an ensemble of particles that we've prepared into the upper state. And then we're watching the intensity as one after the other, different ones decay. More of them will decay early in time. A few of them will decay later in time.

And so when we take the Fourier transform of that, we get this. You're right. Each individual photon has a specific energy. And this is made up of lots of different photons coming from particles. The ones here decay at t around 0. And the ones out here decayed at t a great deal larger than 0. Thank you. That's a better interpretation.

OK, Doppler broadening. Doppler broadening, we have some intuitive understanding of it. It's to do with whether the particle is moving away from us or towards us. And we get a shift in the frequency of the light that we receive. To keep things simple here, I'm going to just stick in a non-relativistic limit.

And then this limit, the shift in frequency of our light, which is defined as whatever frequency we see minus the initial frequency, the sort of, if I was in the same frame as the emitter frequency. And this is just a sign convention here. So I'm choosing it so that $\Delta \nu$ is greater than 1 when the light has been upshifted in frequency.

And for a non-relativistic system, this is just the velocity over the speed of light times whatever the initial frequency of the light was. So I can rewrite this in terms of velocity and I get ν over ν_0 minus 1 times the speed of light.

So the intensity of light which I see at frequency ν is going to be proportional to the number of particles traveling at speed v , which cause a Doppler shift, which gives a frequency ν here. So this is the distribution of particles frequency ν over ν_0 minus 1 mc^2 -- not enough brackets.

So if we're dealing with something like a Maxwellian, where F of v is proportional to the exponential of minus $1/2$ mv^2 upon the temperature like this, then we will end up with an intensity, which is proportional to exponential ν over ν_0 minus 1 squared mc^2 over $2t$. Have I missed the minus sign? It seems likely.

So this is the broadening for Maxwellian. Now we're looking in frequency space here. So if we initially had some line at ν_0 , and I'm plotting i of ν against ν with my spectrometer, what can we say about this function, any facts about it?

Is it symmetric? Close to symmetric, good answer. So for $\Delta \nu / \nu_0$ much less than 1, so for small shifts, this is a very symmetric function. And it's peaked at ν_0 here.

And it does indeed look like a Gaussian. So your normal approach when you see some function that is symmetric and peaked is to fit it with a Gaussian. And that would be a reasonable approach here. It just turns out in the wings, where probably you can't measure it, because signal to noise is so bad, this may be slightly asymmetric. This is symmetric if you do the Taylor expansion around 0.

And what we find here is that the full width at half maximum total $\Delta \nu$ at $1/2$ is equal to $2 \nu_0 \sqrt{2t / mc^2}$, which the punchline of this is that the broadening is proportional to t for the $1/2$ here, which is the thermal velocity, as you might expect. So this means if you measure the width of a line, then you may be able to infer the temperature of it.

Now, you noticed here I've used $\Delta \nu$ to the half. I like using the full width at half maximum for these, because this is a relatively unambiguous thing to define. If you start writing down σ , then you're intrinsically thinking about Gaussians, really. And that means that when you use this full width at half maximum here-- again, this I could also write as $\Delta \nu$ to the $1/2$ -- it makes it harder to compare. This picks up loads of extra factors of $\log 2$ and stuff like that. So I like working with full width at half maximum. It's easy to find here.

So this result is not really very surprising. Effectively, the line traces out the distribution function here. There are some particles which are moving-- so a few number of particles which are moving very fast away from you, redshifted, a few particles which are moving very far towards you, blueshifted. But most of the particles are not moving towards or away from you. So we have most of our mission at the initial line frequency here. Any questions on Doppler broadening?

AUDIENCE: Professor.

JACK HARE: Yeah.

AUDIENCE: So I'm assuming all these happen in laboratory plasma settings. But which one plays more of a role? Do they have equal amounts or is it one or the other that has more of an effect?

JACK HARE: Yeah. So we'll go through all of these and maybe discuss that in a bit more detail. But all of these have an effect. But some of them may be negligible. And so it will depend a great deal on your plasma, and you'll want to calculate it. So for example, we did an experiment recently where all of the broadening we saw was due to natural broadening and the response of our spectrometer, and not to do with Doppler broadening. Though we thought it was going to be Doppler broadening originally.

But you can calculate it. If you think, I know what the temperature roughly is, I can calculate the amount of Doppler broadening. And then you can compare that to the natural broadening. And remember that you can calculate the natural broadening if you have access to a database with these spontaneous emission coefficients. You can quickly estimate what the lifetime of the state is. And so then you can estimate the natural broadening. And you can compare that to the Doppler broadening you expect.

And depending on your plasma-- remember, the natural broadening doesn't care about the temperature. So for a colder plasma, the natural broadening will probably dominate. And for a hotter plasma, the Doppler broadening might dominate. Or there might be other mechanisms that dominate as well.

I briefly mentioned it then, the response function of your spectrometer is also important, because your spectrometer has some finite size detector to it, which means that even if you give it a nice delta function here, it'll be broadened in some way. And if you're lucky, this will be Gaussian. And it'll be easy to understand.

But of course, your spectrometer doesn't have to have a nice response function. It could have a response function like this due to some weird reflections inside it. So you also need to understand the response function of your spectrometer before you try and use, for example, the broadness of a line to estimate the temperature. It might just not be possible.

AUDIENCE: I see. That makes sense. Thank you.

JACK HARE: Yeah. I'll talk a bit more about that with the other three mechanisms as well. Other questions?

AUDIENCE: [INAUDIBLE]

JACK HARE: Yeah. We've taken a one-dimensional slice of the distribution function here. And because we're doing non-relativistic, we don't really care about the motion perpendicular to us, which in the relativistic case does happen. So yeah, this is sufficient for a three-dimensional system, because the three components of the velocity are-- no, this is OK. We can use this one-dimensional approximation as long as the distribution function is isotropic. If the distribution is anisotropic, then you would see different broadening in different directions.

So the next mechanism we're going to talk about is called Stark broadening. Now, the only way to do this properly is with quantum. And you can go have a look at Professor Hutchinson's book. And you can see that there's a very long derivation for this. And even in his derivation, he admits that he doesn't cover all of the possible cases.

So Stark broadening in general is extremely complicated. So what I'm going to give you is a very hand-wavy motivation for why Stark broadening might be happening. And then you can either accept that, or if you need to go and do spectroscopy Stark broadening, you can go do the calculations properly.

So effectively, the idea of this is we have some ion, or some atom, I guess, we've been calling them. And if this atom is alone in the entire universe and it has few enough electrons, we can just about calculate the energy levels. And then we can do all of our nice quantum mechanics that we talked about. These have wave functions that possibly we can calculate. And so we can work out what the spontaneous emission coefficient is, because it goes lower upper upper, all that sort of stuff.

But this is great. But as soon as we put anything else in the universe, this treatment will stop working very well. And the reason is, say, we've got some sort of electron here flying around. At some snapshot in time, this electron has a radial electric field like that. And that radial electric field will distort the potential around the atom. And it, for example, will distort these energy levels. And maybe we don't know exactly how they're being distorted, but maybe some of them go up. Some of them go down. It changes all of these energy levels.

And that means now when we have transitions, their energy, $h \nu$ prime l prime-- this is the lower state. This is the upper state-- does not have to be the same as the original energy level $h \nu$ l . So the photons will come out differently.

So effectively, what Stark broadening is is the effect of the electric fields of other particles on the energy levels, shifting the energy levels, giving us different energies. And we can imagine as these electrons move around, they're going to create some sort net effect, because there's not just one. There's more of them.

And all of these are going to shift the energy levels in different ways. And we might imagine there's some smearing out at this point, because some atoms are going to see some electric field, and some atoms are going to see different electric field configuration, which may move the lines apart or closer together and give us that broadening.

So that's the first level of hand-wavy approximation. The next level is to try and put a little bit of math to it. And we say that the change in the frequency is going to be proportional to the electric field felt by the atom. Now, I've used this symbol to mean energy a lot. In my notes, I started writing the electric field like this, simply to make it clear that there's a difference here. So I'll do that. This big curly E is the E field.

I want to know-- I kind of just stated this. You might be able to motivate it by thinking about the energy of an electron in an electric field and how that energy shifts. But it turns out that, quantum mechanically, sometimes you get $\Delta \nu$ goes as the electric field squared. This is called the quadratic Stark effect for obvious reasons. This is the linear Stark effect, which is the one I'm going to use as my toy model today.

And the only way to really find out whether you have quadratic or linear is to go do the quantum theory properly for this. My apologies. So this electric field is going to be related through the distance between-- say there's one electron that's closest that's doing the most distortion. And so that is inducing electric field by Coulomb's law that goes as $1/r^2$. There's some distance r between these. And so we'll get some shift that's related to distance to the closest electron.

Well, how close are these electrons? Well, we have a density n like this. And so that density is going to give us an average spacing of $1/n$. So we can write that the other way around and say that r goes as $n^{-1/3}$ like this.

But this is asking for some density n , what is the average spacing between particles? So if I just plot them around in this volume, I can see as n goes up, r goes down. If I put more dots here, I get a smaller r . This is a very crude estimate. So that means that the electric field that the atoms are feeling is going to scale as $n^{2/3}$ like this. And so therefore, that's also our scaling with $\Delta \nu$.

Obviously, what I've not got is the coefficient out the front here. You need that coefficient to actually link the frequency shift back to the density. But in general, the idea is, as we had before, that we have some line at ν_0 and that line is broadened and the width of that broadening at full width at half maximum is proportional to $n^{2/3}$ of the density.

So potentially, by measuring this broadening, as long as you've got this coefficient from doing all the quantum mechanics right, you can infer the density here. Doppler broadening got us the temperature. Stark broadening might get us the density. Not very satisfactory, I appreciate it.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yeah. These are the electrons which are inside your Debye cloud. So they are the Debye shielding electrons, but some of those electrons may still get very, very close to the ion.

AUDIENCE: [INAUDIBLE]

JACK HARE: I think that that's the electric field outside of the Debye cloud, but I might be wrong. I think that's the electric field, which drops off far away from the atom, and then goes to 0 very, very quickly. But I think if you imagine you have your ion and you have this cloud of electrons which are shielding the charge, there's still a possibility of one electron going very close.

And that's sort of averaged closest approach thing is what we're estimating here. So you still will get some perturbation to the energy levels here and some broadening from it. Again, I think my picture with individual electrons here is not the right way to do this. To do it properly, you have to do it with a wave function.

But people do this. People use Stark broadening to measure density. And so what you'd want in order for this broadening mechanism to dominate is to have a dense plasma at relatively low temperature so that you don't get the Doppler broadening, which is proportional to temperature to the half. Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: OK. Maybe that's a better answer.

[LAUGHTER]

Any other questions on Stark broadening? Sometimes called collisional broadening as well. [INAUDIBLE] Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: I never worked out what the motional Stark effect is. It's like it's the electric field that the particle feels in its lab frame. So it's actually related to the magnetic field. So you use the motional Stark effect to make local measurements of the magnetic field. But they call it Stark because it's effectively the electric field. But I don't think it's Stark in terms of density. I think it's to do with that lab, like you transform into the co-moving frame with the particle, and it fills an electric field instead of a magnetic field. Never quite got it.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yeah. So the question is, if you have both Docker and Spark, are the line shapes different in such a way that you might be able to pluck out both the density and temperature? I think, in general, because all of the line shapes that we get are Gaussian-like functions-- they look pretty Gaussian. They are symmetric, and their peak's in the middle. It's quite hard to deconvolve both of those.

For the natural broadening, Though we have a Lorentzian, you end up-- if you have a Lorentzian, or if you have the natural broadening and the thermal broadening, the Doppler broadening, you end up with something called a Voigt profile, which is a well known function. And then you can fit a Voigt profile to your line. And then you can get out the thermal broadening component, which is what you actually want to do.

Presumably, there is some function that we can come up with, which is a convolution of whatever the Stark broadening is and whatever the thermal broadening is. And maybe you could fit that. But as always, you're really at the mercy of noise. And so if you're not measuring this line very precisely, it's going to be really hard to determine the difference.

By the way, if you look in Hutchinson's book, there's some really weird Stark broadening features, some of which actually have a little dip in the middle like that. So Stark broadening, depending on the quantum mechanics of the quadratic and the linear can give you really funky line shapes as well. So that could be a good way of trying to deconvolve, because then this function here looks quite different from a peaked function. Any other questions? Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yeah. I think you're right. I think any particle coming near will distort the electric field. And it could be another ion as well. I was just using electrons [INAUDIBLE].

The next mechanism is Zeeman splitting. I'm going to explain this. And you're going to say, this isn't broadening. But it is. I promise you.

So Zeeman splitting is what happens in the presence of a magnetic field. So this is effectively the B field splits energy levels. So if we think of a system that has, for example, a shell with p angular momentum and s angular momentum-- if you haven't seen this in quantum. Don't worry about it. But these are letters which refer to specific amounts of orbital angular momentum.

Then we can have a transition from the p level down to the s level. And this is allowed because we have a change in angular momentum of plus or minus 1, which is the selection rule that's enforced by our dipole operator here that converts from the s energy level or the p energy level down here.

When we have a magnetic field, we see that this line actually splits. So the s does not split. But the p, for example, will split into three energy levels here. And these energy levels will now have different angular momentums and they will also-- we're going to give them a label of sigma minus, sigma plus, and pi. The reason for these are just obscure spectroscopy notation. But this is what they've been called.

The energy is different because the magnetic field is coupled to the magnetic dipole of the electrons. So we have a change in energy-- do I want to draw that as ΔE ? No-- of ΔE . And this change in energy, ΔE , is equal to $-\mu_B$, the magnetic moment of the electron, dotted with \mathbf{e} like this.

And this μ_B is a combination of the orbital angular momentum of the electron and the spin of the electron. If you've ever seen this in quantum mechanics, I apologize for butchering it like this. I'm just trying very quickly to get to the result. But effectively, you can have electrons which are spin up and electrons which are spin down. And so some of them are going to have their energy shifted up and some of them are going to have their energy shifted down.

And that means that we're going to be looking out for transitions. Now, as opposed to just having a single transition down, we're going to have three different transitions, all of which are allowed, and all of which have slightly different energies here.

So if I now look at my initial frequency on an intensity plot here, as opposed to initially having a single line, I may now have three lines. I've labeled them with their labels. And see the sigma minus has lower energy.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yeah.

AUDIENCE: Is there anything like [INAUDIBLE]?

JACK HARE: No. So this is this s.

AUDIENCE: Ah, no, no. [INAUDIBLE]

JACK HARE: This doesn't, because it's a singlet state. But that's quantum mechanics. And if you haven't seen it, I'm not going to explain it to you now. You can, in general, have a splitting of the other state, but then often transitions in those states is forbidden by the selection rules. So you won't see those lines.

And also, this can split into two, or five, or all sorts of other numbers like that. I'm just doing a case where it split into three. Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: It's to do with a number of different ways you can add up the angular momentum vector and the s vector. So it is finite. There's only a finite number of them. But oh god, you've really taken me back now here. But we have some vector B here. It depends on how we can add up i and s.

And because s is the spin of the electron, this can only take finite values. And so this is $1l$ plus s projected onto b. But we can also have s like this. And so this would be another l plus s projected onto b. So it does take discrete energy levels, and it depends on the angular momentum. And this s is actually the total spin of the electrons in this energy level. So these electrons, it doesn't have to just be one. There can be three here. And then there's lots of different ways you can get s out of that.

We're really not going to go into quantum mechanics in this. You should look it up if you're interested. I realize as I try to explain it we're getting deeper and deeper. The fact is magnetic fields, I think the heuristic motivation here is that the magnetic field will shift these energy levels, and it will create new energy levels, which are-- either have their spin aligned with the magnetic field or against the magnetic field. And those energy levels will have different energies. It's like a little bar magnet rotating to align itself with the magnet.

And because those energy levels have different energies, the lines will have different energies. The light that we get has different energies. And so if you're in a state where your broadening is very narrow-- so there's not very much natural broadening, or not very much thermal broadening-- you will get out these maybe three very distinct lines.

And then the cool thing that you can do there is you can measure $\Delta\nu$. And you can get out the magnetic field magnitude. So that's pretty neat. In many plasmas, we don't know the magnetic fields. And so Zeeman splitting is a good technique. In something like a tokamak, we tend to know the magnetic field very well. So Zeeman splitting is not that useful a technique.

The trouble comes that actually a lot of the time, these lines are relatively small splitting. So the magnetic field may not be very large. And so these lines may be very close together, or they may have large thermal broadening. So if I draw this now with each of these lines having significant thermal broadening, we can see that we won't see this well resolved triplet anymore. We will instead see some big blob of lines.

And so this is effectively another broadening mechanism, because we split these lines, but we don't resolve them. So the total overall shape here, $\Delta\nu_{1/2}$ is greater than $\Delta\nu_{1/2}$ of the Doppler. So if you interpret this line as just being Doppler broadened, you'll infer like some ridiculously high temperature. But in fact, that only looks so high because you've got some of this Doppler-- this Zeeman splitting broadening built in.

Questions on this? I have another point, but I just want to pause here. Yeah. [INAUDIBLE].

AUDIENCE: [INAUDIBLE]

JACK HARE: Oh, sorry. Yes. That's from-- so you know μ because it's actually very constrained by quantum mechanics. And so if you're seeing this line, you know what the energy is. And so you know what transition that corresponds to. So you know the electronic configuration of the upper state. And then you can go do your quantum mechanics and be like, there's three levels, or there's five levels, and the spacing of those levels is this or that. So that's OK.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yes.

AUDIENCE: [INAUDIBLE]

JACK HARE: Oh, yeah. Yeah, yeah.

AUDIENCE: [INAUDIBLE] the bottom [INAUDIBLE].

JACK HARE: Yes. That will smear things out. As always with spectroscopy, if you're line integrating through a region that has different parameters, then there will be some smearing there.

AUDIENCE: [INAUDIBLE]

JACK HARE: Oh, we've got time. It's a rather famous work by the group at the Weizmann Institute in Israel, where what they do is they do Zeeman splitting spectroscopy, but they do Zeeman splitting on a range of different lines. And what they have in their Z-pinch is a temperature profile.

Maybe something that looks like this, the temperature of the Z-pinch is hottest in the center. And so they make the Z-pinch out of oxygen. And the reason they do that is that as you look at different shells of the Z-pinch, different radial bins like this, for example, in this bin, it's dominated by oxygen 2. And then this bin is dominated by oxygen 3. Maybe the center is dominated by oxygen 4.

And so these different ions of oxygen-- so this is oxygen that's once ionized, twice ionized, three times ionized-- has different energy levels available to it. So it has different excited states available to it. And so it has different lines available to it. And so here, for example, there's a well-defined line at this frequency. For this one, there's a well-defined line here. And for four, there's a well-defined line here going up in energy.

And so they know when they see this line that it must correspond to the outside of the plasma. And they know when they see this higher energy line, it must correspond to the inside of the plasma. And so when they look at the Zeeman splitting of these lines, they can effectively localize where the magnetic field is in space, because they're able to localize where those lines come from in space.

This looks a little bit like with the electron cyclotron emission all over again. So you need to have some idea of what the temperature is inside here. And this, effectively, then looks like an [INAUDIBLE] inversion problem. So if we look at the Z-pinch down from above, there are these different rings with oxygen 4, oxygen 3, and oxygen 2.

And you can do spectroscopy on a series of parallel lines of sight like that, and do an [INAUDIBLE] inversion on it, and work out what regions emit oxygen four lines, three lines, two lines, and so on. So a rather nice technique that you can use. But in general, you're right. If you're looking at some astrophysical plasma where you can't do an [INAUDIBLE] inversion easily, then this may not be the best technique.

A couple more things on Zeeman splitting-- interestingly, these transitions here from the sigma plus and minus and the pi levels have different polarizations. And that can be used to actually do things like measuring the field direction. So trying to find somewhere where I've got enough space to draw this.

If we've got some magnetic field B like this and a plasma, which is emitting Zeeman split lines, then if we observe at 90 degrees, then the light has two polarizations, depending on whether it comes from the sigma plus or minus levels, or the pi levels. And so in this polarization here, we get the sigma plus or minus. And in this polarization, we get the pi.

And so by using a spectrometer with a polarizer in front of it and measuring the ratio of these two lines at different polarizations, we can back out the angle of the magnetic field here. And this can actually be very useful in a tokamak, where we know what the size of B is. So we don't need the Zeeman splitting. But we do want to know what the twist of the field lines is because that tells us about how big the poloidal field, particularly a useful diagnostic on devices with large poloidal fields, like reversed field pinches.

It turns out if you look along the magnetic field, by the way, then you just see the sigma plus and minus lines. But you don't see the pi line. And I did at one point understand why that was the case. And it probably comes into quantum here as well.

I don't think that was a chat noise from Zoom. I think it was something else. All right, any questions? Yes.

No, not for quantum mechanics. What are you talking about? Are you talking about this gyration?

AUDIENCE: Right. [INAUDIBLE].

JACK HARE: Is it?

AUDIENCE: [INAUDIBLE]

JACK HARE: The electrons are not gyrating. The ion is gyrating. It has electrons which are bound to it, which have wave functions moving around it. All the electrons are gyrating with the ions. I don't think that's going to be a big perturbation to their energy.

So when we talk about hyperfine, we talk about coupling with the nucleus. And we're-- I am not talking about the hyperfine coupling here, because that's very hard to measure. So here, this is the spin angular momentum. And this is the orbital.

I don't know whether someone has done the quantum mechanics on how the gyrating particle affects the energy levels. I imagine it would be small. Can we come up with an argument that shows that it is small? Or is it the case that because-- does it not change the energy levels? Because if you move into the frame of the particle, the ion-- all the energy levels are going to be shifted by the same amount. I don't know-- interesting question.

Any other questions on-- sorry. I don't know the answer to that. As I said, it's been like a decade since I had to really do this math. So any other questions? Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yeah, I guess it depends on the exposure time of your spectrometer. You're capturing this spectrum on a camera. And that camera has a shutter, which is open for some time. So if the shutter time is shorter than the time over which parameters change, you effectively got a snapshot of the plasma. So that doesn't matter.

But if the plasma is changing a lot and you just have a spectrum that is collecting for all time, then it will be a problem. So for example, if you have a spectrometer that's using a bit of film without a shutter, and you're just like give me all the light over the entire course of the plasma experiment, then the time changing properties are important. And you would be better off using a camera that had a shutter. Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: So I don't know where your statement comes from. You might be right, but I can't verify that.

[LAUGHTER]

Is it really smaller? I mean, it feels [INAUDIBLE]. So the field strength here, we're talking about relatively large magnetic fields. You see Zeeman splitting. And the nice thing about Zeeman splitting is it does just keep increasing with magnetic fields. So you can get very large splittings. I think it is a pretty large effect in general.

The final mechanism we're going to talk about is blackbody absorption. So we talked a little bit about this before. We had this idea that if we are an optically thick plasma, where tau is much greater than 1, then we end up having an emission profile that looks like i of ν is equal to ν squared t upon c squared. This is the blackbody spectrum.

Now, it turns out from thermodynamics, you can also show that for whatever plasma you have-- optically thin or optically thick-- the emission i of ν is strictly going to be less than or equal to this blackbody spectrum. You can never exceed that.

And so that means if you've got a system with a very bright line, like this line here, and it's got some width to it because of all the other broadening mechanisms that we've discussed, then the intensity of this line is going to be capped by the blackbody emission spectrum. And so for example, if the blackbody emission, we're dealing with a small range of ν here. So this is roughly flat, even though it's got a ν squared dependence on it.

Then our line would end up actually having a flat top to it. And if the temperature was even lower, then our line would have a broader flat top like this. So this T_1 , T_2 , but T_2 is greater than T_1 here.

And again, because I'm-- although there's a ν square dependence here, I'm only looking at a very small range of ν . So the ν is roughly constant over this. And so we get these little flat tops for the lines. And this looks like another form of broadening, because, previously, we were defining $\Delta\nu$ to the $1/2$, which is the full width at half maximum, which is where you take the peak value and go halfway down and define it like this.

But now when you take-- if you just naively apply this formula, you're going to end up with a broader looking line. You may observe some broadening, but it's actually just using this blackbody effect. So this is only really a problem for very, very bright lines. This is usually a very large value for your plasma. But those lines will have a cut off to them where the intensity can't go any higher than the blackbody level. Any questions on that? Yeah.

So I was saying, for an optically thick plasma, you get the blackbody spectrum. But from thermodynamic arguments, you also find that the spectrum from any plasma for a given frequency cannot exceed that blackbody value. If you did that, you could work out how to break the second law or something like that. So not allowed. Other questions?

So maybe a couple of applications here for the spectroscopy, starting with this Doppler broadening. There's the obvious use of Doppler broadening to measure the thermal velocity. Less obviously is when the Doppler broadening actually stops you making a measurement of the overall Doppler shift.

So if we have a plasma which has some bulk motion, all the particles have some bulk velocity towards you or away from you. Then you're going to have a shift of wavelength, $\Delta\lambda/\lambda$, which is approximately the same as $\Delta\omega/\omega$. And that $\Delta\omega$ is going to be whatever the k of your wave is dotted with v here.

This is just the classical non-relativistic Doppler shift that we talked about back here. So this Doppler shift has a size $\Delta\omega = v \cdot k$, and the broadening, as we've already said, has a $\Delta\omega$ that is about the ion temperature square rooted.

So if you have a system in which your Doppler broadening is very large, then you will not be able to measure your Doppler shift. So for example, if your Doppler shift is just a tiny bit over, it's going to be very hard to measure that on top of a very broad profile.

On the other hand, if you have relatively narrow Doppler broadening, then you should be able to measure that Doppler shift very easily. So this kind of puts a limit, but the interesting thing is we can exploit the fact that this square root of temperature is really square root of temperature over the ion mass, because it's related to the ion thermal velocity here, which means that we can actually choose to look at ions with very heavy masses.

So these would be impurities in a tokamak. So this would be, for example, your hydrogen, it's got a very low mass. So it has very large thermal Doppler broadening. And so we can't measure the Doppler shift very easily. But if you've injected some tungsten into your plasma from the walls, for example, then you're going to have very, very small Doppler broadening. And you're going to be able to measure that Doppler shift very accurately. So this is basically using an impurity as a tracer of the flow.

And if you are APS this year, and you made it to the Tuesday morning plenary, where John Rice was talking about this, people have built their whole careers out of this idea that you can use the impurity ions, which are easy to measure Doppler shift for, as a proxy for the motion of the whole plasma. And it turns out it is more complicated than that, but this is, at least, the basic idea is it's much easier to measure Doppler shifts for heavy ions than for light ions. Questions on that? Yeah.

AUDIENCE: [INAUDIBLE] Doppler broadening [INAUDIBLE].

JACK HARE: So sorry. Let me be very clear then. So Doppler broadening is when you-- the Doppler shift is to do with the entire plasma movement. So imagine a tokamak. This is a toroidal velocity of the plasma moving round and round in a circle. The broadening is due to within that moving plasma, some particles going slightly faster and slightly slower. So it's sort of like they're all going like this, and some of them are also going like this as well.

AUDIENCE: Broadening [INAUDIBLE].

JACK HARE: Yeah, exactly. Yes, yes, yes. So in order to get the broadening, you need to know your distribution function, like a Maxwellian like this. In order to get the shift, you just need to know the velocity of the entire fluid. Yeah, that's a good way of looking at it. Any other questions? Yeah.

AUDIENCE: [INAUDIBLE] broadening [INAUDIBLE]?

JACK HARE: Yeah. In general, with all of those diagnostics, the answer is, in theory, yes, but the noise floor of your diagnostic will probably limit you, because there's very few of those fast particles, and so they're not making very much light. And so any noise is going to make things hard.

So this is true with any diagnostic that promises to map out the distribution function. The problem is that it's usually proportional to the number of particles at that velocity, and there's not very many of those particles because our distributions are close to Maxwell. I'll talk about another one of these diagnostics that you might be able to use as well in a moment.

I want to finish up with a couple of diagnostics which are spectroscopy adjacent. And these are often called active spectroscopy diagnostics. So all the other diagnostics we've been talking about, we just look at the light coming from the plasma. With these active diagnostics, we actually do something to the plasma, and then look at the light coming out of it. So it gives us a little bit more control.

Active spectroscopy-- so the main one I want to talk about here is called laser-induced fluorescence-- LIF. The idea of laser-induced fluorescence is that we have a system with maybe three energy levels like this, lower, upper, and intermediate. We actively pump the atoms in our plasma up to this level u . And then we look for light emitted when the system drops from level u back down to some intermediate level i like this.

So we're stimulating the absorption. And then we're looking for the fluorescence coming out. And the reason this is a powerful technique is if you have a plasma like this and you focus your laser beam through it, you can now have your viewing cord perpendicular to the laser and collect the lights like this on your spectrometer.

And you have chosen this transition to be very unlikely to happen. So you're not going to randomly get that many atoms in the upper state. So you're not going to spontaneously get that many photons with this wavelength being emitted. So that means when you see any light there, you know it came from the region of the plasma, which is the intersection between where you are pumping the plasma with your laser, and where you are collecting with your spectrometer.

And you remember previously the trouble with spectroscopy is you collected all the way along the line of sight. And you didn't have any localization. But this technique gives you that localization.

That's very powerful. You can also-- because of this, if you have your spectrometer absolutely calibrated and you know exactly how many photons you've counted, and how many photons you've put in with the laser, you may be able to get out a measurement of the density here, of the atom density.

So this can be a way of directly measuring the number of atoms in this volume. And indeed, people use this even for neutral gas, or measuring neutral gas density. An extremely clever thing you can do on top of this is to scan this frequency here.

So the laser frequency, you can get a laser, which you can gently tune up and down by a very small amount. And the neat thing about that is that you can then choose which atoms are absorbing the light inside here, whether the atoms are moving towards you or away from you due to their blue shift here.

So if we-- this is a little bit exciting to draw. So we'll see how it goes. If I scan my laser over a range of frequencies around ω_0 , which will correspond to this transition in a stationary atom, then for these frequencies down here, the red-- the lower frequencies here, they will be resonant with the redshifted or with the-- they will be resonant with the particles which are moving away from them. They'll be resonant with these particles here.

If instead I'm using a laser light, which is upshifted, then it will be resonant with the particles which are moving towards me here. So as I scan my laser, I'm exciting particles, some of which are moving towards me. Some of which are moving away from me. And then I'm going to get out a spectrum at the upper to intermediate transition. Here is the upper, the lower transition here.

And the amount of light I get out is going to tell me how many particles there were that were resonant with this frequency here, how many particles there were that resonated with this frequency. And so I can actually map out the entire distribution function using this, because I'm very carefully tuning the light. And so I know when I meet the resonance condition how fast these particles are going with respect to me.

You need to have a steady plasma to do this, because what you're actually doing is ramping the laser frequency in time from some initial value. You will ramp it around that, and then you'll probably sawtooth it like this.

So your plasma has to be nice and steady during this time now that you're ramping it. So this doesn't work on very short lived plasmas, but it does work very nicely on like low temperature plasmas that people use for various chemical processes. Any questions on that? Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: So it can perform arbitrarily well if your plasma is very, very stable, because you can just do this many times, or you can have a very long integration time. So if you run it for twice as long, you're going to get twice as many photons, and then the error goes to the square root of n kind of thing. So you can make this very good.

The trouble is if your plasma is only short lived, then you can't do this many, many times. But if you've got like a glow discharge that just sits in the chamber for 24 hours, then you can just do this 24 hours, and get very good statistics. And you really can measure the tail very, very carefully. Any other questions?

There's a variant on this, which is called two-photon absorption LIF, Laser-Induced Fluorescence. And so this has the very compounded acronym TALIF. TALIF is used because in many of the interesting transitions we've talked about-- and you'll remember I talked about the Balmer series and the Lyman series earlier. And I said for the Lyman series, even the lowest one of these transitions here requires an ultraviolet photon, maybe something the Lyman series at 121 nanometers or higher-- lower, shorter wavelength.

So if you're trying to do this LIF technique with hydrogen, which you might very well want to do, then you're not going to be able to do it, because these photons don't propagate through free space. And so you can't just set up your laser outside the vacuum chamber. The air will absorb it very strongly. So tau is much, much greater than 1 for air.

That's your big problem with all sorts of vacuum spectroscopy is that a lot of the wavelengths of vacuum light gets absorbed in air. So you either have to put your spectrometer inside the vacuum chamber or something else. You can't put a laser inside the vacuum chamber. That just simply doesn't work.

So what you do instead is you take advantage of a very, very strong laser, because if you have a very strong laser, there can be a virtual energy level here, where we have a photon coming in with h lower upper divided by 2, half the energy. This will excite the electron into this virtual energy level, which absolutely does not exist.

But if quickly enough another photon comes along, also with the right energy level, then we can successfully promote that electron all the way up into the other energy level. Because this frequency only has to be half of this transition frequency, we can go from having something at 121 nanometers, which is in the vacuum ultraviolet, or maybe 240 nanometers, which is just in the near ultraviolet. And we can propagate through a vacuum here.

So this means you can still do TALIF, even if the energy levels of the atoms you want to work with aren't really playing ball. But what you do need is a very intense laser, because at this point, if you don't have another photon come along, then the electron will realize that this is not a real energy state. And it will just return back down to here. And the original photon will just go on as if nothing has ever happened.

So you need to have a very high probability of two photons coming in a very short period of time. So we need a very intense [INAUDIBLE]. And you can still do exactly the same tricks we talked about here with scanning your laser in order to slowly map out the distribution function of the particles in your plasma. Any questions on that? Yes.

AUDIENCE: [INAUDIBLE]

JACK HARE: Not really. I mean, almost going back to what I'd erased with the natural broadening, any energy state in a particle does exist for a very, very short amount of time. The stable energy states are the solutions to Schrodinger's equation.

But we can have unstable states, which have very, very short lifetimes. And so we can make an electron into this state here. It's just that this is extremely short lived. So it may be picoseconds, or something like that, before we need another photon to come along. Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: No, these are not measurable light. This is why you need a very intense laser. This is not something that you can-
- I don't believe that you can measure the virtual state lifetime. Oh, sorry. [INAUDIBLE] first, and then.

AUDIENCE: [INAUDIBLE]

JACK HARE: This is better if you have the same energy, because then it's more likely the second photon that comes along will have the right energy, and also half as many lasers, which is always good. Yes, I believe you're right. This scheme could work with various other things. I'm sure you'll get \$10 million in venture capital with no problem. Good.

AUDIENCE: [INAUDIBLE]

JACK HARE: And I think there's even a way of looking at this where you're like, ooh, there's a nonlinearity in the system, and the medium itself does the second harmonic generation, and then absorbs that photon. So if you prefer to think of it in terms of second harmonic generation within a medium, any medium is nonlinear if you drive it hard enough. So again, a very intense laser-- you can get out the second harmonic generation from anything. So you can think about it like that, if you prefer.

And then all it is is the medium converts these two photons into one photon, and then it immediately absorbs that photon. So maybe that's a nicer way to think about it. Yeah.

AUDIENCE: Thank you.

JACK HARE: Oh, OK, cool, great. Any other questions? Okie doke. See you on Thursday.