

[SQUEAKING]

[RUSTLING]

[CLICKING]

JACK HARE: So today, we are starting a new topic, and we're going to be looking at radiation or self-emission from the plasmas. In a way, this is the natural language of the plasma. Previously, we have surrounded it with magnetic probes. We've plotted it with Langmuir probes. We've fired various beams of lasers and microwaves through it, but now we're looking at what the plasma produces itself as a way of trying to diagnose what's going on inside.

And it's important to note, of course, plasma is not special in this sense. So all objects will emit electromagnetic radiation. Humans, of course, mostly in the IR. Pokers when they get red hot, famously, glow different colors. And of course, the universe, in the form of things like stars and black holes. But even in the completely empty parts of the universe, we still have the cosmic microwave background left over from the Big Bang rattling around.

And so even something that is as cold as outer space still has some radiation associated with it. In general, if we have an object which is hotter, we get higher energy photons out. So we could say higher frequency or energy, which is just, of course, h times the frequency. Photons out of it.

And so we might expect in the case of plasmas, which are usually pretty hot, that we would be able to get some pretty high energy photons. And indeed, plasmas tend to span the gamut all the way from the radiofrequency up to X-rays and even gamma rays.

So we're going to be trying to talk about radiation that spans different energy bands varying by many, many orders of magnitude. And just like when we talked about the different plasmas we worked on at the start of this class, we saw there was a huge diversity. Here, we're also dealing with a huge diversity, so we're going to try and use a framework that treats all of these in a similar way. But of course, there will be some subtle details.

And at the end of the day, some things are going to look more like radio waves where maybe you need to think more about the wave equation because the wavelength of the radiation is on the order of the size of your plasma. And some things like X-rays and gamma rays are going to be up where the wavelength is so small compared to the size of our plasma that we can use a ray treatment instead. So just keep in mind that we have different conceptual frameworks for dealing with radiation and we may need to modify it a little bit, depending on what we're dealing with.

Now, the fact that plasmas are hot and emit all of these different types of radiation has two consequences. First of all, this can be a significant cooling term. So for example, the radiation may cool down our plasma. You saw this particularly if you took the fusion energy class. And we remember that our Bremsstrahlung cooling there was a significant loss mechanism in our Ohmic power balance.

But as well as being maybe pesky in the sense that they cool our nice, hot plasmas down, of course they also give us information. So if we can study the spectrum and the location and the temporal variation of this electromagnetic radiation, we can get information on our plasma. And this, of course, is our focus in this course because this is a diagnostic course. So in this course, we like radiation. Maybe in the previous course, you decided you didn't like radiation because you wanted to do fusion. But here, we like the radiation.

Now in general, the way that radiation moves around inside the plasma, or indeed, inside any fluid is not trivial. And we want to think about how that radiation is transported. And this is a whole topic in itself that we'll only cover briefly here-- the topic of radiation transport.

The basic idea in radiation transport is that you have some sort of plasma. There'll be radiation emitted in one place, and that will locally cool the plasma. And that radiation may be absorbed in another place, which will locally heat the plasma. But then that bit of plasma will also re-emit, and so the cycle continues. And you can see straight away that in three dimensions in a highly inhomogeneous system, this is going to be very, very complicated. So this is, effectively, the transport of energy.

And this transport can be highly non-local. So you may be used to thinking about transport in a system where we've got some sort of heat, thermal conductivity. And we look at the diffusion of the heat through the material. That's a very local process. The amount of heat that's traveling just depends on the local temperature gradients.

But here, we could have a region which is emitting, and it gets absorbed a very, very long way away. And because that radiation is traveling close to the speed of light, this can be a very, very fast process. So this non-locality makes solving the full radiation transport problem extremely difficult. In reality, if we want to solve radiation transport, we often assume that this is a diffusive process. And there are reasons why that assumption might be valid, but there are also good reasons why it may not be valid in general.

So just to give you an overview that radiation transfer is complicated. And we're just doing a simplified version of it here. But radiation transport is extremely important for understanding how plasmas work-- for example, the significant cooling. But it's also very important for understanding how we get information from the plasmas because if you just have a camera sitting out here looking at the plasma and you just see some light coming out, you really want to have some model that tells you where that light came from inside your plasma.

If you just think it's all coming from the surface, you'll get one result. If you think it's coming from modestly within the plasma, you'll get a different result. So it's very important to understand this radiation transport conception.

So we're going to start by having a look at radiation transport before we even start to talk about what the radiation is or how it's produced. So even without knowing any of the details of what's making all these radio waves and X-rays and gamma rays and things like that, we're going to look at radiation transport framework, which we'll then apply for all of the different wavelengths that we're working with. OK. That's very high level. Any questions on that so far?

AUDIENCE: Professor?

JACK HARE: Yes.

AUDIENCE: When you say assume diffuse, what does that exactly mean?

JACK HARE: Sorry, I was talking at the same time. So the idea was that most of the time, when we're trying to solve radiation transport, we can't use this non-local model because it's very, very complicated. And so we often make assumptions that our transport is diffusive. I'm not really going into this in a great deal of detail here, but there are-- I can give some references later if you want.

AUDIENCE: Yeah, that'd be great. Thank you.

JACK HARE: The other thing I'll say is that a lot of this radiation transport stuff, as we'll find out, is important when you have some sort of opacity in your system for sufficiently high energy X-rays in a sufficiently sparse plasma like a tokamak. Then radiation transport is not necessarily very important. And so you might think, oh, I'm a tokamak person. I don't need to know this. But we'll find out very, very quickly that radiation transport is still incredibly important for the lower frequency waves, like the electron cyclotron emission in a tokamak. So people still need to pay attention even if they think they're too good for radiation transport.

So let's have a look at what's going on here. Again, we're going to have some sort of plasma. And we're going to have some radiation coming into the plasma. Maybe this is radiation from another part of the plasma. Maybe we have generated a beam of X-rays or lasers or microwaves that we're using to shine through the plasma. It doesn't really matter here.

This radiation is going to have an intensity, I , and we're going to parameterize the path of the radiation with this parameter, s . So we're going to call the point where the radiation enters the plasma s_1 . And then there is some path, s , through the plasma. And we're interested in the properties of the radiation at point s_2 . That's I at s_2 when the radiation has exited the plasma here.

So what we want to know is, how does I change along this path, s ? Now, in general, this path, s , could be curved because we've talked extensively about the fact that when you have changes in the refractive index, you're going to have refraction of your ray. Now, although I have drawn it curved here, a lot of the time I'm going to assume it's a straight line.

So just watch out for that. But if you want to solve the full thing, you need to take this curvature into account. OK. And this quantity, I , that we're dealing with here, this is a quantity which we can formally define as the spectral radiance.

I've used the symbol, I , which is often used for intensity because this is a symbol that's commonly used. Although people often refer to the spectral radiance as the intensity, they also use the word intensity to mean lots of different things. And so if you go on the radiosity Wikipedia page, there's an incredible table that has very, very niche words for all sorts of different quantities.

And this is the table I always go to when I want to work out exactly what I'm talking about. And the reason is that this word, spectral radiance, here corresponds to one exact set of units, which is the watts, the joules per unit time being emitted through an area, meters squared here-- through a solid angle per some spectral unit here.

So the spectral unit could be something like hertz. It could be something like joules. It could be electron volts. It could be meters. The difference in this being depends how you're resolving your radiation spectrum.

So in terms of hertz, that's where we're using angular frequency. Joules and EV, that's where we're using energy. And meters, that's where we're measuring it in wavelength. So if we're talking about the spectral radiance in terms of, I don't know, gigahertz or per EV or per meters here. Of course, if you put meters in this one, then it starts to get complicated because this becomes a cubed and stuff like that.

And so people will actually write this as watts per meter squared per steradian per meter, just to remind themselves that they haven't accidentally folded this in. So what this means is that there is some amount of radiation that's going through some unit area-- that's the meter squared.

We've got watts going in this direction, and that watts is subtended by some solid angle that's measured in steradians. And we're also resolving it in terms of electronvolts or joules. This is a complicated quantity. But it will correspond to your intuition of what intensity of light is like if you think about it a little bit.

OK. So let's just have an example here. Maybe our spectral radiance-- and I may start from this intensity every now and again-- initially has a spectrum. And I shall do this in hertz here. This is what Hutchinson uses in his book. Maybe our spectrum initially has some sort of complicated set of features-- spectral lines in the background, stuff like that-- like this.

And by the time that it has traveled through the plasma, there will be some emission inside the plasma. The plasma will make more light which adds to our beam of light, and there'll be some absorption inside the plasma that gets rid of these. So for example, we might have absorption at these low frequencies. They've been completely wiped out.

And then we will have some additional emission from inside your plasma. So here, we would have a region where there was absorption. And here we have a region where there was emission.

So in general, we want to come up with some formulas that will tell us how, for a given amount of absorption and a given amount of emission in different parts of the plasma-- how the spectra will change from one point to the other. And this is the radiation transport equation, which is the very fundamental equation in this field. And the radiation transport equation says that the change in spectral radiance along our path that we've parametrized by s is simply equal to this quantity, j , which is the emissivity minus the intensity of the spectral radiance times by another quantity, α , which is the opacity.

The important thing about the equation as we've written it here is that each of these quantities are functions of frequency. And effectively, this equation is linear in frequency, in the sense that we can solve the equations separately for each frequency. So whatever is happening down in the gigahertz range has nothing to do with what's happening up in the X-ray range. This is not necessarily true. In order to make this assumption, we have to neglect interesting processes like fluorescence, where we might excite part of the plasma at one frequency and get back lines at different frequencies.

But in order to be able to solve this equation easily, we're going to neglect that. So this is an assumption that we've made. And I'll just tell you what the units of all these things are. You can work them out yourself, but the units of opacity-- very clearly, it has units of inverse length scale. And the units of emissivity are still watts per steradian per hertz.

But now it's not per meter squared, but per meter cubed. And this reflects the fact that the emissivity is a measure of the light coming out from a small volume in every direction like this. The first region means that if you integrate over 4π steradians, you get the total radiated power per cubic meter per hertz.

AUDIENCE: Emissivity and opacity are also functions of s , right?

JACK HARE: Yes. I've suppressed the s in here, but you're quite right. As we go through the plasma, these will in general change. Yeah. These are properties in the material.

It could be a plasma or it could be a lump of cloudy glass. We haven't actually decided or said anything about them. And we won't learn how to calculate these for some time. You just have to take my word for it that you can calculate the emissivity and the opacity.

You've come across some of this stuff before if you've done blackbody radiation. But we're just going to leave it like that at the moment. OK. Any questions? Yes, Sean?

AUDIENCE: Does this equation-- if you include temperature dependence of the emissivity and the opacity in this equation to be used to understand heating of the plasma, if you're thinking about maybe you're injecting waves or something.

JACK HARE: So the question was, can this equation be used to understand heating of the plasma? Yeah, effectively this term is the amount of energy lost from the beam per unit length. And so you could add that energy into your plasma. So you could couple this to an energy equation in your plasma. But this itself is not an energy equation for the plasma. This is an energy equation effectively for the energy in the radiation.

AUDIENCE: I see.

JACK HARE: Yeah. But if you wanted to do full radiation transport, you would couple this to your energy equation, yeah. Other questions? OK. It's very convenient to introduce a slightly mysterious but important quantity called the optical depth.

And the optical depth, which is still a function of frequency-- and I'm just going to keep trying to write these as long as possible to remind you of all functions of frequency. This is a quantity which is defined as the integral of the opacity along some section of path. And just written like this, this is an indefinite integral. We will come up with some limits to it later on.

So at the moment, it doesn't have a huge meaning. But the important thing about it is that it depends on the actual path, s , and frequency. So the optical depth can be very different for X-rays-- good for microwaves-- so again, very important to realize that. The reason we introduced this is that if you stare at this definition and this equation, you might be able to convince yourself that we can then rewrite the equation as $dI/d\tau$ is equal to j upon α minus I .

And this is rather a nice equation because we use the equations of the form $dI/d\tau$ equals minus I . These are exponential growth, exponential decay equations. And we're also used to adding an extra term in here, which is like a source term.

So what this equation says is that the intensity is going to drop off along the path parametrized by s , which is now sitting inside τ . But the intensity is going to increase with a parameter that looks like J upon α . And I'll show you how we're going to use that in a second. Let's label explicitly. This, then, looks like a source function.

So now we can go back to this problem and we can solve this equation. So we can solve for s_1 to s_2 , like that. And we're going to solve it by integrating up with respect to τ here. And we're going to find that the intensity at s_2 on the other side of the plasma is going to be the intensity that we started out with at s_1 times by a factor of τ_1 minus τ_2 .

So this is an attenuation that takes into account any absorption along the path. And then along with the initial radiation we started with, we're going to have a second term which corresponds to the radiation we've picked up through the plasma due to the emissivity. And that's going to be the integral from s_1 to s_2 .

Our emissivity, now explicitly parameterized as a function of s . As someone asked already, all of the J 's and α 's are functions of s . But now I'm making explicit here e into the $\tau_1 - \tau_2$ ds . What this exponential here is doing is saying, yeah, OK. You're adding radiation, but that radiation may also be absorbed.

If that radiation is emitted very early on, near s_1 , it's going to be absorbed a lot as it goes through the plasma. If the radiation is emitted late along the path, close to s_2 , it's not going to be absorbed very much. So it matters where along the path the radiation is actually emitted.

So let me just label these terms here. So this is the original. I'm going to call it spectral radiance so I don't accidentally say intensity. So the original spectral radiance-- this is the optical depth from s_1 to s_2 . This is the emission along the path, s . And this is the depth from s to s_2 , where s is wherever along the path this radiation has been emitted.

This is a complicated equation, and I'm now going to give you two simple examples to try and build up your intuition for what this equation is saying. If you don't think it's complicated-- well, never mind. I think it's quite a complicated equation, especially what these τ 's are doing here. Any questions on it, though, before we keep going. Yeah?

AUDIENCE: So for the emission along s term, if we're wondering about what's appearing at the point, s_2 , I'm just confused why we don't care about the e angle that it is emitted from, right? Because if it's emitted back along the pathway that you completed, it's not heading towards s_2 . I guess where is that dependence in there?

JACK HARE: Right. So the question was, where is the direction of the emission featuring into this equation? And the question is because when we're talking about the emissivity, we said that it is in some solid angle. So we could imagine that this emissivity has an angle with respect to the magnetic field. And that would be a perfectly reasonable thing to think about. And indeed, this is where I've started slipping into a straight line kind of picture and also an isotropic emission kind of picture.

If you want to do this properly, all of these things that are scalars start to become vector, or even worse, some sort of horrific tensor and things like that. And you have to account for all of that properly. So just to try and simplify this, I'm going to do that. But you're right. If you're dealing with an emission that is anisotropic like electron cyclotron emission, that might be very important. Yeah. Another question from Sean.

AUDIENCE: Can you define τ as an integral over the path?

JACK HARE: Yeah.

AUDIENCE: I'm a little confused how-- is tau inside the exponential inside the integral? Is this a sub integral to be evaluated before you evaluate the true integral?

JACK HARE: I think there needs to be some primes and things like that in order to do this properly. You're certainly going to evaluate this tau, and it's going to be evaluated based on whereabouts in this integral, s , you are. So as you're incrementing s here, this tau is going to be evaluated from, for example, whatever is halfway between $s_{1.5}$ to s of 2, where the endpoint of this tau is fixed already.

I don't think I gave a very good explanation of that. I think probably if I went back and did this more rigorously, I'd put some more primes on this, and we'd have integrals within integrals. OK. Other questions? Yeah, I see Nicola.

AUDIENCE: The path that the radiation takes through the plasma is going to depend on the radiation itself, right? So it's going to depend on what kind of radiation it is.

JACK HARE: Yeah, could be. So if you're a different frequency, you might have a different refractive index, so you'd have a different bending. Yes, exactly.

AUDIENCE: But in this equation, the path is predetermined. So it doesn't-- we're saying that if we know that the path is this, this is how we would calculate the--

JACK HARE: Right. But the path of the radiation, although it does depend on the frequency, doesn't depend on the intensity, necessarily. And so for a certain frequency, you can trace out a ray through your plasma. You'll know what trajectory it is. And then you can go back and solve the radiation transport equation along that path.

AUDIENCE: And all the emissivity and opacity figures for that particular frequency that we are looking at.

JACK HARE: Yes. At the moment, we're splitting it up so you could solve this for any different frequencies you want. And they don't interact in this simple model here. Yeah. OK.

AUDIENCE: When you were answering Sean's question-- I'm a little confused about how the $s_{1.5}$ works into the bounds of the tau integral. If you're providing it with some estimate up to, what's the starting s , or if you're providing it, starting, that's with your ending s in that tau integral?

AUDIENCE: So in this integral is an indefinite integral. So we're not providing any limits whatsoever at all on this. The limits only come when you start specifying the endpoint. So for example here, we completely specified both the start and the end, τ_1 and τ_2 . And that integral is going to look like-- this is going to just look like e to the integral of αds from s_1 to s_2 .

This is the equivalent. These two are equivalent here. In this case of τ_2 , at the moment what we've got here-- we don't really have the lower bound. So we have e to the integral of s to s_2 of αds , like that, where the actual s that we put in the bottom here depends on whereabouts in evaluating this integral we actually are.

Maybe that also might answer Sean's question. I don't know. OK, good. Thank you for teasing that out. OK. Just any questions online? OK, Nicola.

AUDIENCE: Is it possible that radiation emitted from outside of that line would end up joining that same rate?

JACK HARE: Yes, it could do.

AUDIENCE: And they're not accounted for?

JACK HARE: In this case, no. What we're going to do in an awful lot of this is assume that the radiation is moving in straight lines, in which case it won't cross. The reason is because this is already complicated enough to solve without that. If you need to solve it with that, then you do. If you're dealing with something like an X-ray going through a tokamak plasma, we are so far away from the critical density with an X-ray that the refractive index is unity, and it doesn't change.

And so therefore, they do go in straight lines. But you're absolutely right. You could imagine all sorts of radiation coming through, crossing through this point, going back to an earlier question. That radiation could heat up a bit of the plasma, and then that plasma could then emit in a way that you didn't expect before by changing, obviously.

So radiation transport is extremely complicated, even without magnetic fields. And then when you put the magnetic fields in, the whole thing becomes much worse. OK. Any other questions? I will give some examples of this equation, so hopefully it will begin to make more sense.

OK. Let's keep going onto that. Good. So let's consider a really simple system where we have some radiation coming in here. We have a plasma, which I'm just going to split up into two regions here. These regions are actually going to be identical, but I want to show you how the intensity changes in each of these regions. And then we've got some radiation coming back out here.

So in this plasma, we're going to have an opacity, which is just simply some amount of opacity as a single frequency, ν_1 . And we're going to have some emissivity, which is simply some amount of emission at a different frequency, ν_2 , like this. And my initial radiation that I'm going to put through here is just going to have some amount of emission at ν_1 here.

So I've really simplified this, where effectively we are solving the radiation transport for all frequencies. It's just that I've come up with this problem where there are two frequencies. So this is much easier. You can come up with something arbitrarily more complicated than this if you want to.

I just want to point out, you're going to find out soon that these two, emissivity and opacities, are not compatible. In fact, there's a really strong thermodynamic link between the two. So don't shout at me there. This is just to make life easier for you right now. But it turns out that you couldn't physically have a pattern like this. We'll talk about that in a little bit.

OK. So what does this look like as we go through the plasma? So at this first step here, we're going to absorb some amounts of the initial radiation, which was at ν_1 . So our radiation is going to go down. But we're going to pick up some amount of radiation at ν_2 because the plasma is emitting.

And this next step here, this process continues. ν_1 goes down, ν_2 goes up. And then finally, as we exit the plasma, we might end up with a system where ν_1 is very, very small and ν_2 is very, very large.

So effectively, as we've gone through the plasma, we have been absorbing that ν_1 , the size of that line has gone down, and we've been emitting that ν_2 . The size of that line has gone up here. So that's a cartoon example. It turns out that for this approximation, that the plasma is homogeneous, then we can solve this equation analytically as well.

AUDIENCE: Professor?

JACK HARE: Yeah.

AUDIENCE: So just to clarify, the absorption is throughout both sections of the plasma.

JACK HARE: Yeah. I split it into two sections so I could just draw the intensity in two places. But the plasma has uniform properties throughout. This applies to both sections, and this applies to both sections like that. Yeah, good question.

AUDIENCE: Thank you.

JACK HARE: OK. Any other questions? OK. At these times, I wish I'd more boards. I really want that equation.

Yeah?

AUDIENCE: Can we assume anything about energy conservation in this equation? Can your additions have more energy out than what caused the emission? Or I guess--

JACK HARE: So the question was, are we conserving energy in this equation? So explicitly, we are not conserving energy in this equation. The plasma is the source and the sink of the energy. Any radiation we lose is heating the plasma, any radiation we gain goes to cooling the plasma. That's why you would need to couple an energy equation into this to do it properly.

AUDIENCE: OK.

JACK HARE: Yeah. I'll put that equation there. So if we have a system, which is now homogeneous so that along our path, s , this quantity, J upon α doesn't change. So again, this is just the homogeneous condition.

We can solve this analytically, and we get that the intensity at point s_2 is simply equal to the intensity of point s_1 , attenuated by a factor $e^{-\alpha \int_{s_1}^{s_2} ds}$, which is, as we said before, the integral from s_1 to s_2 of αds . So this is just an exponential damping factor on the intensity. This is what you would expect for some sort of constant opacity through your system. The further you go, the more that initial signal is going to be damped. And then we're also going to have a term, J upon α $1 - e^{-\alpha \int_{s_1}^{s_2} ds}$ -- same again here.

OK. And when we look at this equation, we can see that there's a strong dependence on τ . And so we want to identify two limits which have names, one of which is this $\tau \ll 1$ the optical depth is much, much less than 1. And the other one is $\tau \gg 1$ is much, much more than 1. And these are called optically thin and optically thick.

In these limits, this equation reduces to either I_2 is equal to I_1 or I_2 is equal to J on α here. Optically thin corresponds to the radiation streaming through without being absorbed. Optically thick corresponds to the radiation being very, very strongly absorbed.

So the optically thick case, we have no information about the initial intensity. The optically thin case, we only have information about the initial intensity. We don't have any added emission here. So people may also call these transparent and opaque.

Now, remember, this is a strong function of frequency. And so you can be optically thin or transparent to X-rays and optically thick or opaque to the first and second harmonics of the electron cyclotron emission and tokamak. And these are not contradictory. We have solved these equations separately for every different frequency that we're interested in in our system.

Now, the interesting thing about this optically thick case is that this corresponds to a black body, which is a thermodynamic system that I'm sure you've studied in which, again, the radiation is in equilibrium with the temperature inside the plasma-- perfect thermodynamic equilibrium. And for a black body, we know a second expression for the intensity here. For a black body, I black body-- or, yeah. I as a function of frequency is equal to a function that's often called B. I'm guessing the black body.

And this is equal to u^2 / C^2 , Planck's constant $h\nu$ over the exponential of $h\nu / T$ minus 1. Or in many cases, it suffices to use the approximation $\nu^2 T / C^2$. And this is valid for $h\nu$ much, much less than T . And this is the classical limit.

So this is the limit we had in classical physics for some time. And it was the violation of this, the ultraviolet catastrophe implied by the fact that, as ν keeps getting larger, you keep emitting more radiation. That catastrophe is well known too. In some part, the development of quantum mechanics-- and this quantum mechanically correct correction.

It just turns out that this is pretty good with [INAUDIBLE] law. It's a pretty good approximation for low frequencies, where the frequency is low compared to the temperature. What is interesting, then, is we now have two expressions for the optically thick case. We have an expression for the black body radiation, and we have an expression to with J upon α .

And these expressions must be equal to each other. And so this means that the emissivity of the opacity is equal to-- I'm just going to use the classical limit-- $\nu^2 T / C^2$ times the temperature. And this is called Kirchhoff's law, along with all the other things that people cotransport.

Kirchhoff's law is very profound. Although we derived it in the limits of an optically thick body, it still has to apply into an optically thin body. And effectively, what this says is if you know the emissivity, then you automatically know the opacity. These two quantities are intimately related. This is why I said that this wasn't actually valid because this does not obey Kirchhoff's law.

So the nice thing about that is you only have to calculate the emissivity for a system, and then you automatically get the opacity. So if you calculate J , you get α for free. That's rather convenient. That's nice, but maybe the more profound thing about Kirchhoff's law is the fact that where you have high emissivity, we also have high opacity.

What this means is that regions which strongly emit also strongly absorb. And that also means that regions which weakly emit do not-- sorry. Regions which weakly absorb do not strongly emit. They also weakly emit. It's confusing.

This is where this comes in. Some of you been staring at this being like, wait, is this obvious? The point is if you have very, very low absorption, you are also going to have very, very low emission. And so you're not going to see significant self-emission from the plasma being added to your initial beam of radiation. And so you might think, oh, perhaps I would get the initial beam plus some extra term. Well, you would get an extra term, but that term will be proportional to τ to 1, and we've already said the brightness is much less than 1. So that extra radiation wouldn't be very significant compared to the initial beam that you put through. OK. Questions?

AUDIENCE: Does the built-in thermodynamic equilibrium assumption or something-- like you always have the same temperature everywhere along your path or something like that? Or is it just--

JACK HARE: No. I think the power of Kirchhoff's law is although you derive it using all these black body assumptions, it then works everywhere else as well. Because these two have to obey this principle for a black body, but it means it also pins it even when you're not in non-local thermodynamic equilibrium. Yeah. So I think that's why it's a powerful result. It doesn't just apply to bodies. Yeah, another question.

AUDIENCE: You said that you can calculate the emissivity and get absorption. That also requires you knowing the temperature, though, right?

JACK HARE: Well, if you're going to calculate the emissivity, you need to know the temperature of your plasma. We haven't got there yet, but the emissivity is always going to be a very strong function of the temperature.

AUDIENCE: OK.

JACK HARE: Yeah. So to calculate the emissivity, you're going to need to know the density and the temperature of your plasma. And if you want to do it anisotropically, you'll need to know magnetic field direction, all sorts of fun things like that. But at the very basic, we're going to do Bremsstrahlung in a little bit-- Bremsstrahlung, density squared, temperature to the half. So you need to know those two things in order to be able to get the Bremsstrahlung out.

AUDIENCE: So we assume that you divide σ of J over α on this [INAUDIBLE]. Are we-- do we also have to assume α is constant?

JACK HARE: Maybe this is more powerful than just simply a homogeneous plasma. Yeah, it looks like they both go up and down together. Yeah. I don't think we need to assume that they're actually homogeneous in order to get this result. So maybe you're right. Yeah, maybe. As long as J and α increase and decrease together.

I'm trying to work out, is that effectively a statement that the temperature is constant? I think it is. So the temperature is constant because of this result. But I think that means the density could change. So it's homogeneous in temperature only. OK. Any questions online? Yeah?

AUDIENCE: Maybe I just missed this in the second [INAUDIBLE] question. But this differential equation in $\gamma \alpha$ equals zero. Why do we need Kirchhoff's law to conclude that I_J and I_α are the same, that they increase and decrease in lockstep? And what that differential equation says.

JACK HARE: This is an assumption. This equation here is an assumption that we've made in order to derive this equation, which is particularly simple to understand the optically thin and optically thick limits because τ just linearly increases with distance rather than changing rapidly over different bits of plasma. Once we've done that, we then find out for the optical case that this is true.

And then, I don't know. I don't think it's obvious from the start that we were going to end up with this result. And this result is still true for inhomogeneous plasmas as well. It just happens to be true for-- we've proved it in the case of a homogeneous plasma. Yeah. I think I see your question. But I don't think we've baked in our result in our assumptions, if that's what you're asking.

AUDIENCE: Yeah.

JACK HARE: I don't think so, but yeah, OK. OK, other questions? Anything online, anything in the room? All right. So now what we probably want to do is calculate J for a variety of different cases.

So we could calculate J , the emissivity from free electrons. What sort of radiation do we get from free electrons? Brems, OK. Anything else?

AUDIENCE: Synchrotron?

JACK HARE: Synchrotron? You're not going to do synchrotron.

AUDIENCE: Thomson scattering?

JACK HARE: Thomson scattering.

AUDIENCE: Oh, no, Compton.

JACK HARE: Compton scattering? Compton scattering is just relativistic Thomson scattering. So we're not going to consider that because that is scattering of radiation from external radiation by the plasma. So we're talking about radiation being produced by the plasma here. But this is a good point. We will talk about scattering extensively later on, but not in this case. Anything else?

AUDIENCE: Recombination.

JACK HARE: Recombination, yeah. It starts with three electronic beats, so it counts. Anything else?

AUDIENCE: Lamar radiation? Or is that just--

JACK HARE: Yeah, that's just cyclotron radiation. To be honest, if that's the same as synchrotron, that explains a lot. So I've always been confused about that. But we're definitely doing cyclotron. Is that the same as synchrotron?

AUDIENCE: I think they're slightly different.

JACK HARE: Isn't one the relativistic version of the other?

AUDIENCE: They're definitely related.

JACK HARE: Anyone?

AUDIENCE: With any three positrons? Do you think that?

JACK HARE: What would we get if we had three positrons?

AUDIENCE: [INAUDIBLE]

JACK HARE: Yeah, OK. So we're not going to do that, but that'll be great. No, no, no. That would be neat. We're not going to derive that. It's a little bit terminal to the [INAUDIBLE], isn't it? Yeah. That's all I've got here. And then we'll also do this for bound electrons. And really, for bound electrons, this is the whole zoo of light emission.

I think I'd be very inconsistent whether I've put two m's or one m in emission throughout this lecture. So you can tell how good my spelling is. OK. Now, in a previous version of this course, we went through all of these without talking at any point about what you'd actually do with it. And I found this very difficult to teach.

So what we're going to do is I'm going to cover each of these topics in turn, and then interject with here's how a bolometer works or here's how a pinhole camera works or here's how a spectrometer works. But the point is that all of the techniques I'll discuss are applicable to all of these different types of radiation. I'm just trying to break it up so it's not learn every type of radiation and then learn every type of diagnosing radiation. But we will start, as Hutchinson does, with cyclotron.

But before that, I have a little question for you about something that I thought about in the shower the other day, and thought, I wonder if this is obvious or not. And you can tell me whether or not it's obvious. And my question for you is, we've encountered two places in which radiation ceases to go through a plasma. We have found a limit where the optical depth is much, much greater than 1. So that's particularly thick.

And we've also talked about a cut-off. That's where the refractive index is less than 0. In the context of a plasma, remember, we might have n is 1 minus n_e over 2 and C . And so when we get close to n critical, the wave is cut off. And that also appeared to stop the radiation going. So what I want to ask you is, are these two things the same? And if not, why are they not?

AUDIENCE: But they're not the same, the reason being that the optically thick case includes absorption, but not reflection, whereas the in less than 0 case, [INAUDIBLE] reflection.

JACK HARE: OK. So let's start writing down some differences. I agree with you. I don't think they are the same. So here, we have reflection. There's no energy absorbed, at least in our WKB picture of this. There might be reality. Here, we just have absorption.

What happens to the wave in these two cases? What happens to the electric field in these two cases? OK. So in both cases, what happens to the electric field? I guess we have a critical surface here. Our wave is coming in. What happens to the electric field past a critical surface in the case of a cut-off?

AUDIENCE: It can't propagate?

JACK HARE: It can't propagate. So what does it do instead?

AUDIENCE: Reflects?

JACK HARE: Is there no electric field?

AUDIENCE: Evanesence.

JACK HARE: Evanesence. So what does it do?

AUDIENCE: Decays exponentially?

JACK HARE: Decays exponentially.

AUDIENCE: It's like an Airy function somewhere in there.

JACK HARE: OK. So this is because our refractive index is less than 0. Our refractive index is defined as k^2 over $\omega^2 C^2$, like this. ω and C are positive, which means that k^2 is less than 0, which means we can define some quantity.

I'm trying to write this very exaggerated κ to try to make it look different from my k 's-- which is equal to k , like this. And then our electric field now goes as exponential to the minus κx , like this. The electric field itself decays. It does not oscillate. What about in the case where optically thick? Does the wave oscillate or not?

AUDIENCE: Can you just attenuate the amplitude?

JACK HARE: Yeah. So in the case where we're optically thick, the amplitude just goes down. So we still have a wave which is oscillating. So E is going as exponential of $-\kappa x$ minus ωt .

It's just that this envelope, I , which is proportional to E times its complex conjugate-- it's not proportional. This is E times its complex conjugate. That is dropping like exponential of minus τ . But the wave is still oscillating. Any other differences about this?

AUDIENCE: Here's a clarifying question. These cases are actually the same, right? It's just that in one case, you've entirely gone to an imaginary k , and in the other case, you put, in principle, a complex k or something.

JACK HARE: That actually feeds into my final point. So yeah, keep going with that thought.

AUDIENCE: I guess just to say one case could become the other, assuming you lost all of your real part.

JACK HARE: Yeah. So people online, the question was these may actually be the same, in some sense. This is a slow decay. This is a very rapid decay. The difference I see here is that you can have this slow decay for many different values of the optical depth. This happens suddenly and only when we get to the critical density.

Before that, the wave knows nothing about it. In reality, if the density is ramping up, the wavelength will get longer. But the amplitude of the wave packet in our [? wkp ?] approximation that we use in the derivable equations-- that amplitude doesn't change. So I guess the way I summarize it in my notes is that this is a gradual process. And this is a sudden process.

In a sense, this is an overdamped oscillator. The amplitude just drops without oscillating. And this is an underdamped or maybe critically damped oscillator, where the wave keeps oscillating, but it goes down slowly.

And this will happen. There will always be some absorption in some plasma. So you will always have very gradual decrease, whereas you don't actually have to have this happen in any plasma at all. It doesn't have to be a critical density.

But there will always be some finite opacity, even if it's very, very small. Anyway, I don't know if that's profound or not. I just thought it was interesting that these two phenomena look quite similar, but they're actually very different. OK. More questions, yes.

AUDIENCE: Weird follow-up. We talked a little bit about not being able to see [INAUDIBLE] for some things.

JACK HARE: Usually reflectometry, yes.

AUDIENCE: Yeah. So if you have a case where your decay is longer than the length scale over which your critical density is maintained, do you get weird things where it becomes a wave on the other side?

JACK HARE: So the question was, if you have a region where the density drops below the critical density-- for example here. So any less than critical density, what happens to the wave? Has anyone done this experiment? I did it in undergrad. It was incredible, life changing. Yeah?

AUDIENCE: Keep oscillating?

JACK HARE: It will start oscillating again. So you will couple an evanescent wave. You can couple an oscillating wave through an evanescent gap where the wave itself does not propagate, and energy will start coming out the other side. We did it with wax blocks and microwaves in an underground lab-- these big blocks of wax and a microwave generator.

And as you move the wax blocks apart, you can generate a wave with increasingly small amplitude. But the wave is still there, and it bridged the gap. And then the remaining energy-- because now your wave is just oscillating on the other side, much smaller. That remaining energy is reflected instead.

But this is a really cool example of electromagnetism. I assume the same thing happens in a tokamak or in any plasma. Does it, people who do waves?

AUDIENCE: [INAUDIBLE] lower hybrid.

JACK HARE: Right. OK. So you can get it evanescently crossing a bit of the plasma and then coming back to life on the other side.

AUDIENCE: Yeah. That's what you want to put your launcher as close to the plasma [INAUDIBLE] as you can. Obviously, that gives you plasma surface problems.

JACK HARE: Just for the people on line, Grant is saying that you want to put your launcher very, very close to the last closed flux surface because otherwise it's evanescently decaying in free space. So you want it very close to the plasma, where it's actually in oscillating mode instead.

OK. Any other questions on this? And then we will go on to deriving radiation from free charges. No screams. None of you've done Jackson before.

The good news is we're not going to do the full Jackson treatment on this. It's extremely boring, and you've hopefully got it before. And if you haven't, there's no way I'm going to teach you it in a couple of lectures. We're going to quote some of the main results. If they look completely perplexing, then it might be worth going to have a look at something that deals with radiation from free charges.

But we're not going to do the whole thing. If you want to see it in very rigorous detail, you should look at Hutchinson's book. It really does go into this with a lot of rigor. But what I want to start with is I want to start with a very simple picture of why charges radiate.

And I have not been able to find this in any textbook. But it was taught to me in undergrad, and I thought it was a rather nice physical picture, so I will teach it to you. And it may or may not be helpful to you.

Radiation is measured by the Poynting flux. The Poynting flux is equal to the electric field crossed with the magnetic field over a factor of μ_0 , or probably c , if you're using CGS. I don't know. The point is that radiation moves in a direction which is perpendicular to the electric and the magnetic fields in the system. And we're going to use this simple formula and sketch a few different moving charges. And we're going to use that to see whether those charges radiate or not.

So let us start with a very simple system-- an electron at rest. What are the electric and magnetic fields in this? Yes?

AUDIENCE: A radial electric field? It's a known [INAUDIBLE].

JACK HARE: So it's a radial electric field like this. Electric field drops off as $1/r^2$, Gauss's law in the r half direction. The magnetic field is 0 because there are no moving charges, so there's no currents. And so therefore, the Poynting vector is 0. Our stationary charge is not radiated.

OK, next one. Now we've got a particle, and it's traveling at a constant velocity. For example, it's traveling in this direction. At a snapshot in time, I'm looking at this particle from my lab frame. What are the electric and the magnetic fields here? Not you again. We'll have someone else. But I'm glad you know.

AUDIENCE: The magnetic field, you can do by the right-hand rule so it's a moving charge.

JACK HARE: Yeah. So there'll be some sort of-- there's current in this direction, so there'll be a magnetic field surrounding it. I've drawn this tilted, out of the page. Otherwise, I'd just have to draw it straight up and down, which would be hard to do. OK. This is the theta. What about the electric fields? Not a trick question.

AUDIENCE: Radially.

JACK HARE: Yes. There was still the same electric fields radially in.

AUDIENCE: Can I ask the annoying question? We're observing from close enough that we're able to see what's moved and all that good stuff?

JACK HARE: Yes. This is definitely not rigorous, but it does get the right answer. So please bear with me, even if you're like, oh, no. OK. So then we've got electric fields, which go as $1/r^2$, r hat. Magnetic fields, which are going to go as $1/r$ theta hat.

This is assuming that we've got this electron as a current carrying wire. Clearly, it's not a current carrying wire. That's a result for a current carrying wire. But it's dropping off in this fashion. And so that means what is s ? How does it scale and what direction is it pointed?

AUDIENCE: It's pointing in the bi-direction or whatever depends on your unit.

JACK HARE: Yeah. Well, the system's a bit screwy here. Let's say I'm using cylindrical coordinates, where I've got z direction here, a radial direction here, and I've got some theta angle like this. I know that doesn't make sense with the definition of r that we had earlier, but again. Yeah.

AUDIENCE: That'd be 1 over r cubed mu 0.

JACK HARE: Yeah. I'm going to drop the mu 0. Treat myself. And what direction is it pointed?

AUDIENCE: Z.

JACK HARE: In the z direction, yes. What is this Poynting flux doing? The Poynting flux is the transport of electromagnetic energy. What's it doing? Why is it pointing in the z direction? Is it radiated?

AUDIENCE: It's following the particle.

JACK HARE: It's just following the particle. This Poynting flux is simply moving the electric and magnetic fields that the particle has with the particle. And it drops off as 1 over r cubed. It drops off very, very quickly, away from the particle.

If I draw a surface with surface area r squared around it, and I ask, how much total power do I have going through that surface? That amount total power will drop off as I make my surface bigger. This is not a propagating wave. You can't observe this.

Again, you need to be close enough to see the electric magnetic fields. But this is not a radiation you can see from far away. OK. So this moves the n energy with the electron like that. This is the point of it. I can add-- sorry, now that we've decided it's in the z direction.

OK. Now finally, we're going to look at a system where the velocity is not constant. And I'm going to look at a very, very specific velocity profile here. And you'll see why I've chosen this in a moment.

This is a profile in which the velocity of our particle in this sort of z direction is initially at some value, and then suddenly, instantaneously drops to 0. A function of time, which means that the x-coordinate is going to go up and then flatten off. The particle will then be at rest, and we'll define this to be $e = 0$ here.

What am I going to get? I've changed something just before the lecture to make it clearer. Now I'm not sure it's consistent. Give it to me. I'm doing it the other way around. Let's see. The particle is initially at rest, and then it suddenly accelerates to some velocity. I think this will make it work.

So the particle is initially at rest. Let's say it's at rest here. And then all of a sudden, it's over, moving in this direction like this. Now, information about the electric and magnetic fields of this particle can only propagate at the speed of light. And so that means if I'm an observer some distance away-- and let's say that distance is this circle, like that.

The information I have about the electric and magnetic fields is the same information I had-- from my point of view, the particle is still at rest, which means that outside of this circle, all I can say is that the electric fields are pointing inwards here. Inside the circle, I now know that the electric fields are doing something quite different. They're now looking like this. And we have some sort of magnetic field like this.

Now, the electric field lines cannot be broken. These electric field lines therefore, in this very fictitious scenario, have to be joined up in this way. So I'm also going to draw some electric field lines like that, electric field lines like that. And these ones are outside like this. And now again on the inside, I've got some magnetic field lines, which is still in this poloidal direction.

So what's happening now? So while pointing flux or electric fields or magnetic fields, let's go through the magnetic field and say that at this interface here-- so at this interface, we still have a magnetic field, which opposes upon our theta hat. We've got an electric field.

OK, I think this is the bit where I just have to tell you to believe me. If I'm right, I can prove this straight away. This electric field turns out to have a scaling which goes as 1 upon r . I can't remember where that comes from. Not 1 upon r squared, but it's very important for the conclusion.

And also here, if we think about this direction, we're going to have something which is tangential to this theta direction. So I'm going to call it some sort of [INAUDIBLE] where I'm subtly switching into a spherical coordinate system. And all of this together means that our Poynting vector now goes as 1 upon r squared. And it has a direction, r hat.

And this means that finally, we've got a system which is radiated. So this is moving radiation radially outwards in every direction. And this is the result that you probably remember is that accelerated charged particles radiate.

These disappointing lectures are not simply shuffling the energy along with the particle. It's actually taking energy from the particle and putting it into electromagnetic waves. And like I said, it'd be much more convincing if I can remember the argument for this fact. And I decided not to write it down in my notes, so I can't remember it. But it's clearly a very different circumstance than what we had here.

The reason we've derived in this very hand-wavy way is I'm about to write down the actual correct equation. And you'll see that it has almost no intuition involved in it whatsoever, whereas I like to think that this gives you a little bit of intuition about what's going on and why, moving charges right here. But questions? Yeah.

AUDIENCE: So in this situation, you've drawn in your third situation with the non-constant B and C lightcone of the particle as it begins to move. That's where the energy is located?

JACK HARE: Yeah. So this distance here, this is where an observer outside of this would still believe the particle is located here. Inside this lightcone, the observer can see that the particle has moved, so of course, the particle is actually going to be moving in several discrete-- I've drawn these electric field lines as straight. They should actually be curved inside here. I've just drawn them as straight from the current position of the particle at this point.

And so there must be a tangential discontinuity of the electric field at this surface to go from one radial vector pointing at one point to a new radial vector pointing at a new point. And it's that tangential discontinuity which generates this electric field perpendicular to the magnetic field that gives us Rossby Poynting vector. And that goes radially outwards.

AUDIENCE: And so for the situation where we have basically a step function in the velocity, the energy is infinitely localized to this shell?

JACK HARE: Yeah. In this case, the particle is only accelerating at this time. The acceleration just looks like a delta function. And so therefore, there's only radiation here, propagating outwards as a spherical wave. And from then onwards, the particle is just in this situation again where it's shuffling its fields along. And so this is acceleration. Yeah.

AUDIENCE: So does that mean like if you have somehow the spatial distribution of energy as a function of your acceleration, [INAUDIBLE].

JACK HARE: Absolutely. And we will see that in the long and complicated formula in a moment, yes. You can do this stuff rigorously as well. Any other questions? Yeah.

AUDIENCE: Just my vague memory of how the $1/r$ thing is because you're-- I think it's like to have your voltage make some big nonsense, you have to integrate along that path length that it's getting extended across.

JACK HARE: OK. So there's a suggestion this $1/r$ comes from trying to keep the voltage to a reasonable value and your integration along the path.

AUDIENCE: Yeah. There's some weird--

JACK HARE: Have you seen this before, then?

AUDIENCE: Yeah.

JACK HARE: OK, good. I'm glad I'm not just crazy. I didn't come up with this. It was underground, and I'll try to reconstruct it. OK. Any other questions before we move on? Any questions online? I see Matt's hand. Yeah.

AUDIENCE: Yeah. Maybe you said this, but I just want to make sure. The Poynting flux direction in the second picture. Yeah. That's only valid in the plane perpendicular to the motion of the particle, right? Because the electric field direction would be ρ , not-- it would be in cylindrical coordinates, right?

JACK HARE: I was saying that the fact that I decided to use cylindrical coordinates here is not actually very useful because the electric field isn't just in that direction or in this direction. It is in every direction. It's just the only place where $\mathbf{e} \times \mathbf{B}$ is significant is really in this plane here.

AUDIENCE: OK. Sure, sure.

JACK HARE: [INAUDIBLE]. And it will drop off. If you go back to this point here, $\mathbf{e} \times \mathbf{B}$ -- because they're not very well anti-aligned, it will drop off. And so you have a ring moving the electric fields and magnetic fields with the particle.

AUDIENCE: Yeah, that makes sense. OK, thanks.

JACK HARE: OK. Any other questions? OK. Let's see the full version.

So if you do this properly, you get out of the equation for the radiation from a single moving charge. Like I said, if you haven't seen this or can't remember how to do this, I recommend you go take a look at Jackson or some other sufficiently advanced text.

And this equation is the electric field seen by the observer who is some distance away from the charge. And this electric field has got the standard $q / (4\pi\epsilon_0)$ that we know and love from Gauss's law. And then it has two terms corresponding to radiation, which we call the near field and the far field.

So the first term, which is part of the near field, is a $1/k^2$. This k is not the normal k that we've been using before the moment. And I've just realized this is a cube. $k^3 r^2$. And then you have a term that looks like $\hat{r} \cdot \mathbf{v}$ minus the particle velocity on C . $1 - \mathbf{v} \cdot \mathbf{v} / C^2$. We're going to see lots of velocities showing up normalized to C because obviously the speed of light is pretty important for this sort of stuff.

Here, I've defined two new terms. This k takes account of the fact that we're looking at the particle at some time. We're looking at the particle and we're seeing where that particle is, but the particle has already moved. So this is helping us with our retarded time, which is looking at the earlier time that we're at.

So this k is defined as $1 / (1 - \mathbf{v} \cdot \hat{\mathbf{r}} / C)$. And I haven't defined $\mathbf{r}$ either, so I should do that. This $\mathbf{r}$ is a vector that is effectively the vector that joins the observer to the particle. So this is defined as the position of the observer minus the position of the particle at time, t .

So if I draw a little diagram, we've got some origin to our coordinate system. There's some vector, $\mathbf{r}$ of t , which is our electron. And then there's some vector, $\mathbf{x}$, which defines our observer. And so this is the vector $\mathbf{r}$ here.

We could get rid of this by just assuming either the observer or the particle is at the origin of the coordinate system. But as the particle is moving, that's not necessarily particularly useful. And so we're doing it in this more generalized way here. So we'll talk about what happens to this term in a moment.

The second term inside here is $1/C^2 k^3 \hat{\mathbf{r}} \cdot \mathbf{v}$. And we're just going to find $\hat{\mathbf{r}}$ is equal to the $\mathbf{R}$ vector over the size of the $\mathbf{R}$ vector-- so just a normalized vector. And that $\hat{\mathbf{r}}$ is crossed with $\mathbf{v}$ minus the C , which is, in turn, crossed with $\mathbf{v} \cdot \mathbf{v} / C$. And $\mathbf{v} \cdot \mathbf{v}$ is going to be extremely important in a moment. And I'm going to close my brackets.

So we have two terms here. We have a term which we call the near field and a term that we call the far field. We're not going to explicitly derive the magnetic field here because fortunately, these are properties of electromagnetic waves in a vacuum.

And we're still talking about vacuum waves here. $\mathbf{B}$ is simply equal to $1/C \mathbf{e} \times \hat{\mathbf{r}}$, like that. So the magnetic field is perpendicular to the electric field, as we expect.

And this means that for the near-field case, we have a Poynting vector, $\mathbf{s}$, which again is equal to $\mathbf{e} \times \mathbf{B}$ with some constants. And that $\mathbf{s}$ is going to go as $1/r^4$. And so the far field case. That $\mathbf{s}$ is going to go like R . And $\mathbf{s}$ is going to go as $1/R^2$. Here we had r^2 of this.

So this straight away tells us that if we consider the power through a sphere of some radius, r , the total power, $\mathbf{s}$, integrated over the surface of the sphere will drop off as R^2 for the near-field. But it will stay constant for the far field. So only the far field is actual propagating radiation. That's the only thing we're actually going to observe.

So we don't actually need to have this near-field term in the equations we're looking at. It just drops out when you solve this equation. We figured [INAUDIBLE]. They're propagating $\mathbf{e}$ and $\mathbf{m}$, $\mathbf{e}$ [INAUDIBLE]

The only reason to show you this is, first of all, to show-- as many of you know, with the radiation even from a single moving charge is very complicated. And then of course, we have to ask, what is it that we're trying to achieve in a plasma?

But the first thing we want to do in a plasma is we notice that our near field far field term has lots of things that we might know, like where we're sitting and how long it's been. But there are several things inside here that we don't know, such as v and $v \cdot$. So that tells us that we're going to need to solve the equation of motion for a particle, something that you've done several times in a plasma already without a magnetic field, with a magnetic field.

You know this is the thing that tells you that the particles are spiraling around field lines. And because they're spiraling, they're continuously being accelerated. And so we know straight away that these accelerated particles are going to be emitting propagating radiation. And that's the thing that we want to detect.

So first of all, we'll have to solve for the equation of motion. But the second thing we'll have to do is integrate overall distribution function, f of e e^3v because we will have then-- once we've solved the equation of motion, we'll have it for a single particle. But the particles have different velocities depending on where they fall in this distribution function.

And they may have different velocities in different directions. They may have one velocity along the magnetic field and another velocity perpendicular to the magnetic field. And so these two steps are very non-trivial to do properly. And these are the steps that we're going to skip in this class.

We're just going to give the results from this. But if you want to do this properly, if you want to work with electron cyclotron emission, you could probably go back and check that you can actually do all these intermediate steps. OK. We are well over time. I'm happy to take questions. But otherwise, I will see you on Thursday.