[SQUEAKING]

[RUSTLING]

[CLICKING]

JACK HARE: Well, welcome we few, we happy few. Let's do a little recap on electron cyclotron emission, and then we will go on to a few other things.

So we had a look at the physics of electron cyclotron emission. We didn't actually derive the emissivity of it, but we gave ourselves a hand-wavy reason why there may be multiple peaks here. So if we have our emissivity as a function of frequency, we expect to have multiple different peaks, even from a single particle.

And these peaks are going to be occurring at frequencies-- I'm going to switch into angular frequency units, omega m, which were equal to this cyclotron frequency, which was just our normal gyro frequency modified by a relativistic parameter here, and times by this relativistic rest mass term, which always makes the frequency smaller, divided by a relativistic Doppler term, which might make the frequency smaller, might make it larger, and, most importantly, all this multiplied by some integer m for some natural number N.

And so, we have these peaks here. And we can say this is m equals 1, m equals 2, m equals 3. And they're evenly spaced. And this is just for a single particle.

And then, when we put this all together, first of all, we looked at what a distribution of particles would do, and we agreed that these peaks would, in general, be broadened in some asymmetric fashion. But the exact shape will depend on exactly how big these two terms are with respect to each other.

And we said the neat thing about this is that these frequencies depend only on the magnetic field. So this frequency here, eB over m e-- admittedly, there's a gamma factor inside there. But if we neglect that, the frequency depends only on the magnetic field. And so, if you see some emission at some certain frequency, then you know that it's been emitted by a region of plasma which has this magnetic field.

And that was particularly useful when we considered a tokamak because, if we have something like a tokamak--

--or a stellarator-- hello?

STUDENT: Oh, sorry, my bad.

JACK HARE: No worries-- which has some magnetic field that goes 1 upon R, then different regions of our toroidal device are going to be emitting electron cyclotron emission with different frequencies. And those frequencies are different because the magnetic field is dropping off nice and monotonically throughout our device here.

So we said that, for example, we might have three different regions, which we could label positions R1, R2, and R3. And each of these then have different magnetic fields, B1, B2, and B3. And each of these magnetic fields will then produce a spectra of lines at these different harmonics here.

Now, the neat thing was, if we consider just the lowest harmonics, so the m equals 1 and often the m equals 2 harmonics, we said that these tend to be in an optically thick regime. So that is the region of plasma that's emitting is emitting as a black body. And that means the black body spectrum I of nu, is equal to T nu squared upon c squared.

And what we would expect to get out if we put our little microwave horn and collected the different frequencies coming out from this plasma, is we'd expect some sort of spectrum where different parts of the spectrum correspond to different regions of this plasma.

So, for example, we might have a region down here corresponding to lowish frequencies at low magnetic fields. So I've got this the wrong way around. R3 is the lowest magnetic field. We might have another region where the frequencies correspond to the center of the plasma and a higher frequency part of our spectrum which corresponds to the high field part of the plasma.

And at each of these points, if we measure the intensity, we then know straight away what the temperature is corresponding to the black body spectrum for that frequency. And, therefore, we know what the temperature is inside our plasma.

So we can go back and we can say, OK, using these measurements, maybe we have a temperature profile that looks like this where, again, we can identify three specific points inside here that we've been using as part of this example. But, in general, you can get it all the way through the plasma.

And so, what we're doing here is we're saying that we've got some mapping from radius, the magnetic field, to frequency of maybe the first harmonic here. We have some intensity as a function of frequency, which is proportional just to temperature because of our black body case.

And then we can back this temperature out. First it's temperature as a function of frequency. Then we can say, well, that frequency corresponds to a certain magnetic field. And then we can say that magnetic field corresponds to a certain spatial coordinate.

And so, this is a technique which will give us, by looking at the spectrum for lowish frequencies, the first or maybe the second harmonics of this cyclotron emission, we can work out what the temperature is as a function of space. So this is why this is a powerful and popular diagnostic in MTF.

And we also discussed a little bit about accessibility. If you remember, this is when we got into talking about the O mode and the X mode. And I'm not going to recap that now. And Sean posted an interesting way of looking at using something called a CMA diagram on Piazza if you want an alternative view on how this works. So that was just a quick recap. Any questions on ECE before we go into correlation ECE?

OK. Ugh, a bad chalk. This is correlation ECE, often called CECE. So what is the problem we're trying to solve with correlation ECE? Why isn't this technique sufficient as it is?

Well, what we want to do, we want to measure very small temperature fluctuations. We want to measure temperature fluctuations within the plasma that are maybe on the order of 1% of the baseline temperature, delta T upon T about 0.01.

So if you have a small temperature fluctuation at some part in the plasma, that means you'll have a small intensity fluctuation. If the mean temperature is just Te, and we have some fluctuation around 1%, this intensity will also fluctuate by about 1%. And that 1% is actually extremely hard to measure. And this is because the noise is just too high on these systems.

There are lots of different contributions to the noise. But, in general, they all add up to make it very hard to measure these very small fluctuations. So ECE, as it stands, is very good for measuring the broad average temperature profile with some error bars. But if you want to measure small fluctuations within those error bars, you don't stand a chance with this system.

And you want to measure these fluctuations because these fluctuations are related to turbulence. And, as we've discussed before, turbulence is one of these key properties that we want to understand in plasmas so that we can build an economically viable fusion reactor. So it'd be very nice to be able to know.

Now, the fact that the noise is too high does seem like a big limitation. But there are some clever tricks that we play where we use correlations. And I'll talk now about what exactly these correlations are and how they provide us with information that allows us to get a signal out despite the overwhelming [INAUDIBLE].

So our setup here is borrowed from ASDEX Upgrade. And, in effect, I referred to Alex Creely's PhD thesis, which you can find online if you want more information, from 2019, which is a pretty good summary of this. A top tip, if you're ever looking for an accessible description of some of the physics or a diagnostic, you should go find someone's PhD thesis. Usually much better than any of the papers that have been written about it, because there's not really any base restraint. And, of course, the person writing the thesis is desperate to tell everyone about all the cool stuff they just spent the last six years of their life working on. So it can get very detailed.

So in ASDEX Upgrade, we don't really have a circular cross-section plasma, but I'm just drawing it like that, be a nice D-shaped plasma, got our plasma inside here. And our system, at first glance, looks an awful lot like the system that we sketched out up here. We're going to have some sort of special lens.

It turns out, you can make lenses for microwaves. I didn't know this-- but some sort of plastic HDPE lens that then couples into a waveguide here. And that lens is going to collect light from a region like this.

So I exaggerated it slightly, but there'll be some region over which we can collect light from relatively small volumes inside here. So we can have a series of volumes inside here. And, once again, our magnetic field is going as 1 upon R, all the standard stuff that we had previously.

Now, specifically for this ASDEX Upgrade system, but generically for C-mod and other devices where you might have this, the first thing that's done is a band pass filter. So this is applying some filter to it that in frequency space looks a little bit like this, some sort of top hat. It's centered around the frequency where most of the electron cyclotron emission is. It's around about 110 gigahertz in ASDEX Upgrade. And this has got a bandwidth of 10 gigahertz.

So we've cut out an awful lot of the radiation in bands that we're not interested in. We are no longer going to study those, any bremsstrahlung, any higher order things. This is going to capture all of the information in, say, the first harmonic within some relatively small window. And that's the important thing about correlation ECE. We're not trying to measure the temperature profile throughout the entire plasma. That's very challenging. We want to measure it inside some very small region.

And the reason is that our turbulent eddies, our little fluctuations inside here, these are very small as well. So the size of our turbulence of R turb, is on the order of 100 microns. That's the width of a human hair. So we are trying to measure turbulent eddies inside a tokamak on this scale of a human hair here. And so, that means, we're not trying to cover this entire frequency range. We're zoomed in on only quite a small frequency range.

Once we've got our bandpass filtered signal here, 110 kilohertz is still too fast for us to digitize. This would be an extremely expensive digitizer. And so, what we actually do is we downmix it with the signal at 100 gigahertz. And so, then we get out our beat signal, which we can digitize, which is at 10 gigahertz here. So 100 gigahertz mixing was this.

Our beat frequency is about 10 gigahertz plus or minus 5 gigahertz. So this is-- actually, I will write this as 0 to 10 gigahertz. This is the sort of signal that we can actually digitize now.

And what we do is we then amplify it. And when we split it off, as we did with our standard electron cyclotron diagnostic, and we split it through a series of different bandpass filters-- so there might be another low, a lower filter, medium filter, the high filter. And then we put these through to a set of detectors.

And so, each of these channels, now that we've split off, is looking at frequencies in a very narrow range here. These are about 100 megahertz in width now. And the spacing between these different bins is 125 megahertz.

It's important that these bins do not overlap in frequency space, which means that these volumes do not overlap in real space. They are each sampling a separate discrete part of the plasma inside here. So let me just write that down.

So each channel samples a non-overlapping region in frequency. And, therefore, in real space, because, again, we've got this very strong link between our magnetic field and our spatial position and the frequency at which we're sampling.

So each of these represents a measure of the power that's being emitted by a very small region. And these regions are clearly distinguished. And you can tell that these regions are very small because we're dealing with 100 megahertz bandwidth. And we originally started with 110 gigahertz here.

And so, you see we've gone down by a factor of 1,000. If 110 gigahertz was roughly enough to cover this entire region here, then we're now dealing with regions which are 1,000 times smaller than the radius of our tokamak. And that's how we're able to get down to this 100 microns or so.

And we digitize these signals. And, just for reference, the reason why we downsampled these is that digitized is now much less expensive. We're doing it something like 4 mega samples a second. And that is quite an affordable digitizer compared to the ones you would need to digitize this signal up at hundreds of gigasamples a second.

So I have not yet told you anything about how correlation ECE works. I'm just giving you an outline of exactly how these measurements are made with an example from ASDEX Upgrade. But there are other similar devices on other tokamaks. Any questions on this before we get to the meat of what we're trying to achieve here? Yeah.

This wave is too damn fast. We can't digitize it. We mix it with 100 gigahertz. We did this before when we were heterodyning our signals for interferometry. We get out two frequencies. We get out a frequency at omega 1 minus omega 2, and we get out a frequency at omega 1 plus omega 2. Omega 1 minus omega 2, that signal is about 0 to 10 gigahertz. We can digitize that.

This one, whatever, 200 gigahertz, we still can't. On the integration time of our detector, the sampling time here, that will have many, many oscillations. And so, it will average out to-- well, because it's power, it will average out to 1 or something like that. So it's some DC offset that we can subtract off after.

So all this is doing here is mixing the signal down so it gets to a regime that we can effectively digitize. Yeah, another question.

STUDENT: [INAUDIBLE]

JACK HARE: So if we were doing geometric optics, which we're not, then you would have a lens like this. That lens could collimate your beam, as in we'd have a load of rays coming out into our horn.

And if it collimated that beam, that would mean there'd be a focus point at some distance f away-- if this lens has a focal length f, then it will diverge afterwards. So, because of reciprocity, that means that, as opposed to launching rays this way and seeing where they focus, we have rays coming from this way.

Now, we don't ever actually get to infinitesimal point like we predict in geometric optics or Gaussian optics tells us that there's a beam waist. And our beam looks like this instead. And that was what was trying to draw up here.

STUDENT: [INAUDIBLE]

JACK HARE: Yes.

STUDENT: [INAUDIBLE] very long.

JACK HARE: Yes, exactly. Exactly, and you can choose properties of your lens that will give you a narrower waist or not. It depends on your size of your lens and the wavelength of light and all sorts of other [? goods ?] like that. So this is not exactly a picture-perfect sketch of it.

But the idea here is that there was some region over which, in the transverse direction perpendicular to your collection volume, you have a very narrow scale, which means that you can actually collect from a very small region on the order of 100 microns. Any other questions? Anything online? OK.

STUDENT: Professor?

JACK HARE: Yes.

STUDENT: I have a question about how you split the signal to each of the harmonics.

JACK HARE: You mean this splitter here physically or--

STUDENT: Yeah, what's the physical method of splitting it?

JACK HARE: I don't know how microwave splitters work. But there is something that you can buy from DigiKey or something like that that will do it for you. Anyone know how micro splitters work? Looking at you?

STUDENT: RF.

JACK HARE: The answer I got was RF, which I don't think is much more satisfying than my answer. So, yeah, there are circuits which will split microwaves.

STUDENT: OK.

JACK HARE: Thing called a rat race, but I don't think that's for gigahertz. Have you ever come across a rat race? If you're doing lower things, it has ports which are spaced by wavelengths of your microwave and then there's constructive and destructive interference in different places? Yeah.

STUDENT: [INAUDIBLE]

JACK HARE: I don't think that's how-- I think they are literally splitting it and then having separate band passes. But, you're right, you could come up with some clever analog way of doing the splitting.

What they were actually doing with this system is all of these were very tunable. And so, if they wanted to look at turbulence in the edge or turbulence near the center, they could actually tune all their bandpass filters to do that.

But this is getting way beyond what I was hoping to talk about on this. So let's get on to correlation ECE see what that does.

So now you've got these n different channels up here. And each of those channels is measuring some signal. And we're going to write this little tilde here to remind us that this signal is some oscillating quantity or time varying quantity.

And that signal has two components to it. It's got some fluctuation, which is due to temperature fluctuation. This is the thing that we want to measure. And then, it's also got some noise. This is the thing that we don't want to measure here.

And so, we are making the assumption here that our noise is much, much bigger than our temperature fluctuation signal. If it's not the case, then you don't have a problem and you don't need to do this technique. But it is the case in most of the applications that we're going to see it's used for.

And so, the signal, as a function of time, is going to-- what we want it to do is look like, again, our temperature signal, which might be some nice and smooth function like this corresponding to some turbulent eddy. But what it actually looks like is some messy noisy thing which, even with some very aggressive filtering, is still going to be extremely noisy instead.

So what we want to do is find some way of extracting this temperature from the noise here. And what we do is we pick two adjacent channels. So we'll call them S1 and S2. These channels are adjacent in frequency space, which means that they are measuring from adjacent parcels of plasma in real space or inside the tokamak.

So, again, if we have our tokamak cross-section here, we have the plasma, we zoom in right at the edge of the plasma here. No, that's not going to work. We would have two volumes very close together but not quite touching. And these volumes we could call 1 and 2, like that.

And what we do is we say, we arrange-- we have some simulation or some prior, which tells us that the size of a turbulent eddy, the size of one of these swirls, is larger than the separation between these two points. So we say, our term is greater than delta R. Which means that, if there's a temperature fluctuation associated with this little vortex, it is the same temperature fluctuation in S1 as in S2.

And, again-- oh, I don't need this to be much greater than. I just need it to be greater than or on the order of. And, again, these are really, really closely spaced, 100 microns or so apart. This is really, really tiny. But, the nice thing is now that these two signals are carrying the same components that we're trying to measure.

And this is where the correlation comes in. We are going to correlate these two signals. And we're going to find that the noise is uncorrelated at random, but these two signals will correlate together, and we'll be able to measure it.

So we have-- just to be clear, we have chosen the frequencies omega 1 and omega 2 such that delta R 1 2 is less than the length scale associated with turbulence. And we can choose these frequencies because we have control over all of these bandpass filters and funky things like that.

So you might need to do this experiment a few times to get it right. But once you get it right, your signal will leap out. Because what we find is that the temperature-- well, we find that the temperature is correlated. So if do S1, S2, and we do a correlation operation on these.

And there are a few different ways to do correlations, and I'm not going to go into them, but I will give you a citation at the moment. We get out a term here that looks like the temperature term squared, the thing we want, the cross-correlation between the two noise terms, plus two cross-terms between the noise and the temperature signal from channel 1 and the noise and the temperature signal from channel 2.

And because this noise is just random, when we do some sort of averaging, this could be in time or best to think of it is a short time integral, then these terms are all going to drop out. And we'll just be left with a correlation signal that is proportional to the temperature squared here.

And, again, I've been a bit hazy about what these angle brackets are doing here. You could think of it as equal to some sort of short time integration of S1 and S2 from T to T plus delta T, like that. That would be a time correlation.

But there are actually lots of different ways of doing this. And the citation here for you is a review paper by Watts, 2007. I don't know if this is an example of nominative determinism. Someone called Watts goes around making power measurements.

[LAUGHTER]

There we go. Yeah, and this technique is incredibly powerful because it's enabled people to measure, again, delta Te upon Te on the order of 1%. And they've done this at 13-- on ASDEX Upgrade-- 13 radial locations. So 13 positions they can measure these temperature fluctuations.

And they've done this with a time step of 100 kilohertz. So 100,000 times a second they've been able to measure these temperature fluctuations. So on very small scales, very fast time scales, we can now measure temperature fluctuations in a tokamak. And this is something that has revolutionized our understanding of turbulence and its importance in tokamaks, because now we can finally characterize it.

That was a lot. Any questions? We're going to move on from ECE [INAUDIBLE]. Yes.

STUDENT: [INAUDIBLE]

JACK HARE: I think the idea is that this is so noisy that if you try and just-- that's effectively a bandpass filter at that point. Yeah, I don't think you can get this out of an autocorrelation. So this is more sensitive than that. Yeah.

STUDENT: [INAUDIBLE]

JACK HARE: Yeah, so the idea is that we have positioned these two volumes, which are producing frequencies omega 1 and omega 2, we've chosen omega 1 and omega 2 so they're very close together. And we think that that distance is smaller than the size of our turbulent eddy. And within that turbulent eddy, the temperature should be the same going up and down.

And if you do this correlation and you get out nothing, that probably means your volumes are too far apart because there is-- this T correlation would just go T1 T2. And there's no good reason to believe those temperature fluctuations are correlated, because they'll be part of a different turbulent eddy. And so, that would just go to 0 as well. Yeah.

STUDENT: [INAUDIBLE] 13 times [INAUDIBLE].

JACK HARE: Yeah, so the question was, what is the actual spatial range of this? Do you have pairs and then a gap and pairs? I believe that, for this, if I remember correctly, they literally just had 14 channels side by side, and that gave them 13 points of correlation between each adjacent pair. And so, they were interested in measuring transport in the edge.

But I believe, in a separate shot, you can then tune all of this slightly differently and move slightly further in depending on where you think transport is important. But if you're looking at transport in the pedestal region, you have some idea where that is. And so you focus your measurements in the pedestal.

But, yeah, you're right, this does not get you a huge spatial range. It gets you about a millimeter. But that might be enough for your measurements.

STUDENT: [INAUDIBLE] would you have [INAUDIBLE] in this localized [INAUDIBLE]?

JACK HARE: Yeah, so the idea here is that your eddies are sort of like-- there are eddies like this, or maybe, at some other point, there is an eddy like this. And so, it will depend a little bit on how you do your time integration step here.

So this is part of the art of it as far as I understand, is that you need to pick pairs because they're the only ones that could possibly be close enough to correlate. If you try and pick this one and this one, it will go to 0.

But, as well, the eddies are going to be moving in time. So you may end up with a situation where, at some point, your eddy has moved across, and these two are no longer correlated because the next two are correlated. So it--yeah. I think we're reaching the limits of my knowledge. Yeah. [LAUGHS]

STUDENT: [INAUDIBLE] correlated, [INAUDIBLE] second and the third one are [INAUDIBLE].

JACK HARE: I mean, there's no reason to believe that each of the fluctuations is the same size. And there's no reason to believe they're stationary in space with time. So this is going to be a time-evolving system here. Yeah, exactly. Cool.

Any other questions? Anything online? Now we're going to go to bremsstrahlung.

Some people these days just call this breaking radiation, which is just the translation from the German, and that's a perfectly reasonable term to use. But if you haven't come across bremsstrahlung as a word before, that's what it means.

And so, what we're dealing with is heavy ions, and all ions are heavy compared to electrons. And we've got some electron whizzing by. As the electron whizzes by, it feels a force of attraction to the ion. And so, it is briefly an accelerating charge. And so, it is going to emit photons as it's deflected from its trajectory.

Now, this electron isn't then just going to sail off and never see an ion again. In fact, it's going to see one very, very soon. And so, effectively, our entire plasma is full of electrons which are gently being deflected and breaking and emitting these photons.

So the main thing we can say in a classical treatment is that the bremsstrahlung is going to be isotropic. It's going to be the same in every direction. That's actually quite different from electron cyclotron emission, even though we didn't really look at the anisotropy of that in any detail.

There are lots of different ways to deal with bremsstrahlung to do the actual calculation. Remember, we said we wanted to have v and v dot, that was our first thing, by solving the equation of motion. And then we want to integrate over the distribution function f of v e 3v That would give us the emissivity that is the input that we're looking for here.

So there's lots of different ways of doing this, and Hutchinson lists a few of them. There's a purely classical picture. That classical picture looks an awful lot like what I've drawn here with the ions and electrons as point particles. There's a semi-classical approach where you start bringing in some quantum physics and treat, I think, the electron as a wave. And then there's a full quantum approach.

What's remarkable about all of these approaches is they all give the same answer with just a very slightly different coefficient. So we get a small change in coefficient.

The scalings are the same. So, in some sense, although it's important to get the exact coefficient, it doesn't matter exactly which one of these techniques you use. And that's why I'm not going to go through any of them in class. They're also, I think, quite complicated derivations.

So we'll talk about it in a moment, but this j has some scaling, which is n squared T to the 1/2 E to the minus h nu upon T. And then there's a load of stuff out the front here. And these coefficients change from 1 to 2.

So it's not really a big difference. I mean, it makes a huge difference if you design your reactor. Because [? if it's ?] 2, you could be a very long way off from breakeven. But, from the point of view of this course, it makes no difference. Yeah, I think it's kind of remarkable that it doesn't make any difference.

So, again, if you want the full treatment, go have a look in Hutchinson. And there's also a long treatment in Jackson of this same problem. I'm just going to quote some results. I'm going to quote results at you. I kind of already spoiled it now. It's here.

For the Maxwellian average, because, again, we can have all sorts of different distribution functions, but our plasma tends towards a Maxwellian. So a Maxwellian averaged bremsstrahlung emissivity, this is equation 5.3.40 in Hutchinson. This looks like 4 pi j nu the emissivity. I'm just putting 4 pi here because it's isotropic.

So j is normally in terms of per steradian. But, because it's isotropic, I can just multiply. I can integrate over the solid angle, and I'll just get a 4 pi here, and it won't change the j at all. This is ne, ni, z squared, T to the minus 1/2 e to minus h nu upon T. And then, there's a factor called G bar, the Gaunt factor, which we'll talk about in a moment, and lots of other constants.

And these constants are things like e, the electron charge, and the electron mass, and epsilon 0, and h bar, and c, all arranged in some way to make all the dimensions work. So we're not going to go in too much detail. This, therefore, has units of watts per hertz per meter cubed. Yeah.

STUDENT: [INAUDIBLE]

JACK HARE: Yeah, it should be minus 1/2 of my notes. That should be.

So this Gaunt factor here, this is the bit where you can spend a lot of time refining your treatment, and you'll get different values of the Gaunt factor. And you'll also find this Gaunt factor varies very weakly with the temperature of the plasma here.

So the Gaunt factor G bar goes from about 0.3 to-- in a range where the photon energy compared to the plasma temperature goes from 0.1 to-- I'm not writing this very well-- where the photon energy compared to the plasma temperature goes from 0.1 to 10. So this is a pretty weakly varying function. We can change it by two orders of magnitude in this parameter here, but this only changes by one order of magnitude.

And so, in general, it's reasonable to just treat G as being a constant and then, for this calculations we're doing, where we're going to drop the absolute intensity, we can just drop G with all the other constants as well. So we can drop [INAUDIBLE].

But if you want to go back and do this properly for some measurement that you're doing, then you'd have to include this. And so, this emissivity, as a function of photon energy, is a very simple function because it just decays exponentially with photon energy. So that is the spectrum coming out of a plasma that's emitting bremsstrahlung radiation here.

And that is all you need to do Problem Set 3. Any questions about this? Yeah.

STUDENT: [INAUDIBLE]

JACK HARE: You mean you take the bremsstrahlung into account? Well, so the thing about bremsstrahlung is, it's always there. It is the irreducible minimum amount of emission from your plasma. Cyclotron emission is a very specific frequency. So although, when we were drawing it before, we drew it over long range. Inside your plasma, the cyclotron emission might be just at some very narrow window here.

But this is everywhere inside your plasma at all frequencies, like a black body kind of spectrum here. That goes down to very low frequencies and goes up to very high frequencies. We're going to talk about lots of other effects which produce emissivity which is higher than the bremsstrahlung.

But, often, when we're doing power balance calculations for a tokamak, we will just use this, and that represents our most optimistic take. So, I guess, you always need to worry about this, even if you managed to clear out all the impurities, you don't have any lines, you don't have any electron cyclotron emission, you don't have any recombination, which we'll talk about next. This is still going to be there.

STUDENT: So [INAUDIBLE]?

JACK HARE: It's going to be on top of this. You're going to have this on top of. Now, the nice thing is, over a short frequency range, this does not change very much.

So if I'm doing my ECE and I zoom in on this region here, I'm going to have i frequency here, this bremsstrahlung is going to be basically constant. And, on top of that, I'd have my ECE lines or, in reality, my black body ECE spectrum or something like that.

But this bremsstrahlung may be quite small. But even if it's not quite small, it is constant because it varies only slowly as a function of a frequency. So for a small frequency window, it looks constant. And then you can just subtract off some background intensity and look at the actual signal you're interested in, like you see there. Yeah.

STUDENT: [INAUDIBLE]

JACK HARE: Ooh, ha, yeah. I mean, it should couple into-- So I mean, so the important thing to realize is what mode it couples in actually depends on where you're looking from in some sense. So if this is being emitted isotropically, some of it's going along the magnetic field, some of it's going perpendicular to magnetic field.

If I'm observing perpendicular to magnetic field, I must be observing it coming out as O mode or X mode. I'm observing it along the magnetic field, I see it coming out as R mode or L mode, right or left-handed circularly polarized light. So, actually, in a way, as the observer, we choose what mode it goes into.

But this emission is isotropic. So you can imagine you have a blob of plasma that emits. And then the wave has to ask itself, well, what sort of wave am I? At this point, well, if I'm going perpendicular to the magnetic field, I'm going to be O mode or X mode. And, in reality, there'll be some emission in O mode, some emission in X mode.

And the exact coupling between those will be related to the polarization of the bremsstrahlung. And the polarization should be pretty random as well. So I would guess, without thinking about it very much, you get roughly equal into both of those. And if you're then-- if your wave that's being emitted along the magnetic field, you'll be roughly equal into the R and the L. But I haven't thought about that too much.

Yeah, and then there will be differences in speed of propagation and things like that. But the bremsstrahlung is being constantly produced, so you probably won't be able to notice that very [INAUDIBLE]. Yeah, good question. Other questions? Sorry?

STUDENT: [INAUDIBLE]

JACK HARE: If it's optically thick. So, yeah, we haven't really talked about that. So from this j, you can now calculate alpha, the opacity, because we had j upon alpha is equal to nu squared upon c squared T. That was Kirchhoff's law here-- that's a terrible j.

So I can now calculate alpha. And might look at this and I might go, ooh, interesting alpha is actually quite high for low frequencies. And so, that means for any reasonable distance of plasma that I'm looking through, that will be absorbed, and it will become black body.

And so, in reality, for some plasma, we might have a spectrum that looks like this, where in this region here tau is greater than 1, and in this region here, tau is less than 1 or much less than 1. So it goes back to being optically thin here for the high energy photons. But the low energy photons will get absorbed as well.

There's actually another effect, which is in Hutchinson's book, which I haven't covered here, which is, of course, there is no mode in a plasma which propagates below the plasma frequency. And so, the spectrum will be even further modified because that wave will be evanescent. But we're skipping over that this year. But it's in there if you're interested. Yeah.

STUDENT: [INAUDIBLE]

JACK HARE: We're not going to look at synchrotron, yes.

STUDENT: [INAUDIBLE]

JACK HARE: I think we had-- Sean's not here this time. Sean asked me that a couple of classes ago, I said I didn't know. And then he came and told me the next class. So cyclotron is when the particles are spiraling around the magnetic field. Synchrotron is when the particles are following the magnetic field in a curve.

So there's two types of magnetic field doing acceleration, but they are distinct ways of producing light. Yeah, I think someone told me that maybe for a high field tokamaks synchrotron could start being significant. But I actually have no idea how big a deal it is. And I don't know if people are using it as a diagnostic. I've not heard of someone using it.

But synchrotron light is used as a source of X-rays for diagnosing many other things. So it's interesting in its own right. But I don't know whether people use the synchrotron radiation from a high field tokamak if we've ever built one high enough for it to be a problem as a diagnostic of something.

And I don't want to be a diagnostic of, because I think it depends mostly on the magnetic field and the radius. And those are two things in a tokamak you already know. So you probably-- but maybe there's a really clever diagnostic you can do, like fast particles or something like that. So worth thinking about.

STUDENT: [INAUDIBLE]

JACK HARE: Right, but I'm not interested in [INAUDIBLE]. This is a diagnostics course. [LAUGHS]

Any questions online while we pause? So this is bremsstrahlung radiation, and often people call this free-free. And the idea here with free-free is that your particle starts off free, and it ends up free. So it hasn't changed.

This is in contrast to some of the other types we'll be dealing with, like bound-free, free-bound, and bound-bound, you get the idea. Electron cyclotron is, of course, also free-free. So there are multiple types that people often just call bremsstrahlung free-free radiation. Now we're going to do free-bound radiation or recombination radiation.

So recombination, otherwise known as free-bound. So our particle is going along, our electron, [? past ?] this ion here. And this particle is going along. And it's got some velocity to start with. And it's emitting, let's say, a single photon of energy h nu, like that. That entire acceleration yields one photon, one energy.

And so, initially, we have a kinetic energy, a 1/2 mv squared. Then, afterwards, we're going to have an energy 1/2 mv prime squared minus h nu. So I have e, that, e prime, that. This is the energy here, and this is the energy here, [INAUDIBLE].

--b prime.

Now, if h nu is less than the initial kinetic energy here, then we still have some kinetic energy left. So v prime is greater than 0, our particle is still free, and it can continue. That's the case we've just considered, bremsstrahlung.

But, of course, there's another case where this photon takes away so much energy that this starts, I guess it becomes imaginary. So it's sort of pointless to continue at this point. But it's clear that the electron no longer has any kinetic energy. And, indeed, it's going to have to pay back any debt it has in some other way. And it's going to do that by becoming bound.

And you can do that because we're going to switch to a slightly quantum model of the atom though, to be honest, the Bohr model works for this as well, where we have a range of different discrete energy levels that the electrons can occupy.

So these energy levels are labeled by the principal quantum number n. n equals 1, 2, 3, 4, and so on. And, up here, infinity, this is ionization. If your electron gets this much energy, it becomes free again. And the energies of these levels are given by this unit, Ry, which is the Rydberg z squared of our ion over n squared.

Just to be clear here, n is this principal quantum number. It is not the density anymore. We're dealing still with the single-atom picture, one electron, one ion. The ion has a charge z here.

Anyone remind me what a Rydberg is? What its value is? 13.6 electron volts-- good number to know. So that's the ionization potential of hydrogen. So that means that when we have an electron coming in, that's initially up here, if it wants to become a bound electron, it has to drop down to one of these energy levels. And these energy levels only have discrete values.

And so, what we're going to see in our spectrum is that this is only allowed if the electron energy fulfills this equation, which involves these discrete energy levels. And that equation is going to be the photon that is emitted.

And that's the thing that we're going to see with our spectrometer is going to be equal to 1/2 mv squared plus z squared upon n squared times the Rydberg. We're going to get out different photons.

Now, of course, these are not going to be at completely discrete lines because it's going to be broadened by our distribution function. There's a range of different electron velocities available. So each of these photons has some sort of variance around it. It's going to be some value, which is strictly greater than the Rydberg energy.

So what does this look like? If we integrate over a Maxwellian-- so, again, this is Maxwellian averaged-- we get out two terms here. So we have-- do I want to say two terms? Well, I'm going to write it as the emissivity per principal quantum number n, so the emissivity per each of these discrete levels. And then we can sum them all up in order to get the total spectrum.

So that just has our familiar bremsstrahlung type coefficients, electron density times the ion density times z squared e to the 1/2, e to the minus h nu upon T times a constant. But now we have an additional term, which is a new Gaunt factor for level n. But, don't worry, these Gaunt factors are all about 1 again, so it doesn't really matter that much.

And then we have a term z squared upon T Rydberg energy 2 upon n cubed. So the strength of this emission drops very fast with n. So it's going to be strongest for n equals 1 and less strong for the higher principal quantum numbers here. And then, there's going to be a factor of exponential z squared upon n squared Rydberg energy upon T, like that.

If you're wondering where the energies are canceling out here, the Rydberg is in electron volts, and we're putting T in energy units. So that can be in electron volts as well. These will cancel out quite nicely here.

So you end up something that looks a lot like the bremsstrahlung but with an additional emissivity for each of these principal quantum numbers. And then, to get the total j, you have j brems plus the sum over n of jn here.

And the spectrum, therefore, looks like we had something like this for bremsstrahlung before, and now we have a spectrum that contains a series of edges, like this. And this lowest edge here-- I should draw this as a straight line-- this is n equals 1, 2, 3, 4, and all the way back up here, where, again, this is our brems result, and this is our recombination.

STUDENT: Professor?

JACK HARE: Yeah.

STUDENT: So if the photons released are greater than the kinetic energy, but it's not equivalent to one of the discrete energy levels, what happens then? Does it just not get bound?

JACK HARE: Sorry, could you rephrase that?

STUDENT: Yeah, so earlier you mentioned that it needed to be equivalent to one of the discrete energy levels.

JACK HARE: It has to fulfill this equation here, where the photon that's emitted is going to have an energy equal to the initial kinetic energy and plus the energy that the photon gains by virtue of the electron occupying one of these energy levels. So if you're asking, what happens if you can't fulfill that, it just doesn't happen. Quantum mechanically, that would be a forbidden transition.

STUDENT: Oh, I see. So it either becomes-- it either doesn't recombine if it doesn't have that and it's still free, and then that's just bremsstrahlung radiation, or it does recombine at these certain locations?

JACK HARE: Yeah, so you could imagine that you-- I feel like for every electron there should be, at some velocity v, there should be some wavelength that it can emit at.

I mean, this is a spectrum of solutions, it's not a discrete number of solutions. If I put in some velocity v and I occupy some principle quantum number n, then I will get out a photon here. There's no forbidden solution to this.

I think what's being reflected here with these sharp edges is, these correspond the sharp edge here, corresponds to v equals 0, and this corresponds to v greater than 0. And if you get up to a point where your energy is so high that you can then access the next level, then that becomes very favorable and, in fact, there are some-- these ones down here, it might continue from the 2 level or it might start occupying this level instead.

STUDENT: Yeah, I think I understand. So I was just wondering, because it's not the condition that there are photons that are released with greater than 12 mv squared the kinetic energy. It has to be either it's not released with greater than, and it's not--

JACK HARE: So see, it's the photon energy that's-- yeah, maybe we're doing this backwards in some sense.

STUDENT: I see.

JACK HARE: So if you see a photon energy with less than the kinetic energy of an electron, then that photon will have been emitted and the electron will remain free. You see a photon energy, how you would know what the velocity of your electron is, I don't know.

But if you're in this single particle picture, if you had all the information, if you see a photon that's released with less than 1/2 mv squared, then the electron still has some residual kinetic energy left. And so, therefore, it must be free. But if you see a photon being emitted with more than 1/2 mv squared, then the electron has over-emitted. And the only way it could have done that is if it fell down into one of these principles.

STUDENT: Oh, I see. OK, we just did it in reverse. I understand.

JACK HARE: Yeah, I think maybe I motivated this the wrong way around. Yeah. Other questions? Yeah.

STUDENT: [INAUDIBLE]

JACK HARE: Yes, exactly. So this long tail here is due to the fact that we have electrons not just with v equal 0 but much more than that, yeah.

STUDENT: [INAUDIBLE]

JACK HARE: Hmm, it's a good question. Now, are they actually sharp or is it that we want f of distribution of v squared, the distribution of v squared, like that? So perhaps these are actually like this. I have to admit, in the book, it's drawn with very sharp lines. But maybe my eyesight is failing me and I should look more closely. I'll have a look into that. Good question. Any other questions? Yeah.

STUDENT: So [INAUDIBLE].

JACK HARE: [LAUGHS] It's quantum, it's worse than that.

[LAUGHTER]

STUDENT: [INAUDIBLE]

JACK HARE: Right, there was a probability that this process happens, and that is related to the overlap in the wave function between the wave function and the free electron, like a plane wave, and the wave function of a bound electron there.

So the probability of that happening is related to the overlap integral between those with a dipole operator, the dipole operator being the thing that emits the photon in between them. So you could write it as Psi prime dipole operator, like that. That's the probability of this thing occurring. Yeah, but--

STUDENT: [INAUDIBLE]

JACK HARE: I mean, this is where, again, our quantum picture and our classical picture are irreconcilable, because if we're treating the electron as a plane wave delocalized over the entire volume. So, yeah, this is-- you can't really-- [LAUGHS] even in a single particle picture, this drawing is not reconcilable with quantum [? theory. ?] So, yeah.

STUDENT: [INAUDIBLE]

JACK HARE: Yeah, it depends on the ion charge. So if you've got some helium with 2 times ionized helium, and you've got some hydrogen, then it will have different recombination lines here. And that's because, fundamentally, the energy levels are in different places inside here. And the energy levels shift by the ion charge here. Yeah.

Any other questions? We'll go into some ways of actually making some use out of all this stuff in a moment. Yeah.

STUDENT: So an electron gets bound by an ion, then, presumably, at a later time, [INAUDIBLE] that absorbs some [INAUDIBLE].

JACK HARE: Ooh.

STUDENT: Wouldn't this balance out, electricity going in and out of ions?

JACK HARE: We'll get to that. There's a very good question. We're going to get to that and the idea of detailed balance and thermodynamic equilibrium and things of that ilk. Yeah, I'm not going to skip ahead. But, yeah, we have a section where we talk about all these inverse processes.

But these are the two simplest ones, free-free and free-bound, that I wanted to discuss before we did some diagnostics. But, yeah, you're right, there are inverse processes for all of these that will affect the spectra that you get. Any other questions? Let's keep going.

So one use of this is a diagnostic called bolometry.

In fact, bolometry is, in some sense, a very simple diagnostic. It cares not at all about the detailed spectrum of what the emission is. It just wants to know how much power is being radiated by the plasma, so bolometry.

And to motivate this, let's go back to our 0D power balance from a fusion plasma, where we have alpha heating, we have some external heating, and in steady state this is balanced by conduction losses and what we, at the time, we assumed were bremsstrahlung losses, but in general could be any sort of radiative losses here.

So I'll write this as S rad here. And so, bolometry is focused on measuring the total power that's being emitted. And so, that is going to be the integral d nu of j of nu.

So we take our beautiful spectrum like this, and we integrate over all of this. We lose all of this detail. But, of course, the detail matters when we're doing the integral because you can see the recombination increases the amount of emission above that baseline bremsstrahlung value. So there'll be more emission here.

And the reason, of course, we want to measure this is we want to know how much power our plasma is losing by radiation. That's a quantity that, in general, we want to minimize. Maybe we want to know where it's losing that power. And so, we can say, OK, well, that part of the plasma is clearly radiating too much. What can we do about it? So this is what bolometry is trying to measure.

And a really simple way to do bolometry is we have some radiation coming out of the plasma. We'll call it P rad here because it's an integral over S rad over the distance of the plasma here, whatever. It's going to have different units. And we have a little sensor sitting at the wall of our vacuum vessel. And that sensor is just a resistor. And we'll call this resistor M. It has a resistance M. And that resistance M is a function of temperature here.

And we apply some voltage, V0, over this resistor. And we also put another resistor that is shielded from the plasma that we call R, and we measured the voltage over this resistor R. Effectively, what we're trying to do is measure this volt-- this resistance M of T.

We'll choose R to be less than or roughly equal to M so it doesn't dominate this measurement. And the voltage that we measure, VR, is equal to whatever voltage we used across all of this system R divided by M.

So by measuring this voltage VR, by knowing the voltage V0, by knowing the reference resistor R, we can measure the voltage M. This is a very simple resistive divider just drawn in a complicated way.

And the reason why this is interesting is that, in general, M is going to change with T. We're going to have some resistivity as a function of temperature. And so, as this resistor heats up from absorbing the radiation, its resistance is going to change. And so, we can have-- yeah.

STUDENT: [INAUDIBLE]

JACK HARE: Have I broken this? Yeah. Thank you. It was right in my notes and I decided to innovate-- it's a terrible mistake. Never do that. Thank you. Right, that looks better. Good. So it's a voltage divider. We're dividing V0 between M and R. And we're measuring the voltage across R, which then, of course, tells us the voltage across M, which then tells us what resistance M has here.

So if we're measuring the resistance M, we can then use some table of rho of T to get out the temperature. And, from that temperature, we can make an estimate of the radiation power incident upon it. And we'll talk a little bit more about how to do that in the next step.

But this, effectively, serves as a very, very simple system. And you can be digitizing this as a function of time. And so, this is also a function of time and temperature. And so, we could get out some graph of our fusion reactor's radiation power as a function of time like this just by looking at the resistance.

Do I want to introduce this equation yet? OK, I'm going to do this, but I might regret it. So the way that you link these back to each other is you say, for some radiation power incident on my resistor-- and I'll make it explicitly a function of time-- that's going to cause some change in temperature with time. And all of this is time some heat capacity.

And, in general, this resistor is going to be connected. It has its own heat capacity c, but it's also going to be connected to some heat sink, which is unavoidable. Even if you don't connect it to a heat sink, it'll be connected to a heat sink through these wires. And that heat sink will have some thermal transport time tau. There's actually a second term here, which is delta tau delta T upon tau, like that.

So there's two different ways that the radiation can affect this. And so, by measuring delta T up here, the change in temperature in our resistance from the change in resistance, we can invert this equation and we can back out the radiation power.

So this is just a really simple way of using a resistor to measure the power coming out of the plasma. And bolometry was one of the first diagnostics that we had on many MCF devices. Questions on this?

STUDENT: [INAUDIBLE]

JACK HARE: [INAUDIBLE], we'll talk about that in a moment. I worked on the bolometers for [INAUDIBLE], so neutrons were the bane of my life.

Now, this itself is a very simplistic system, and it does not work because it is very susceptible to noise. So we have to, as always, come up with a clever system, which is noise resistant. And think I'm just going to-- ugh-- get myself another board. What time are we on? Oh. OK.

So what we do in-- to make this measurement more precise is first we do two things. First of all, we do something called a Wheatstone bridge, which many of you have come across before. A Wheatstone bridge is an arrangement of four resistors. So there's a resistor here, another resistor here, a resistor here, and another resistor here.

And we connect these resistors up. And the clever thing, in a bolometer, is we allow two of the resistors to see the plasma. And these are the ones we call the measurement resistors M1 and M2.

So you can imagine these resistors just having a little window that allows them to see the plasma, and the other two resistors cannot see the plasma. They have a large lump of metal which gets in the way. And these are the reference resistors, R1 and R2.

I think I'm just going to call these R2 and R1. And, in a Wheatstone bridge, what we now do is we still have a potential drop, V0, over the entire bridge. But we now measure the potential difference across these pairs here. And that is called V bridge.

So these measurement resistors, they see the radiated power P rad, and their temperature is equal to the temperature of the reference resistors plus a change in temperature due to radiation, whereas the reference resistors just have a temperature which is equal to the reference resistors.

The point of this setup is your entire bolometer is going to heat up because the entire vacuum vessel is going to get hot. And so, what you want to know is not just how hot they are, but how hot they are relative to the vacuum vessel.

And so, this is why we have this system where we're trying to measure a very small change between R1 and R2, which would be identical, and M1 and M2. And it's this small quantity here, delta T, that's due to the radiated power that we're really trying to measure.

And we do this because-- or the way that we do this is we notice that V bridge is the difference between two potential dividers, one potential divider that goes on M1 R1. And so, the potential that this floats at here, Va, is going to be different in general from the potential that this floats at, Vb, where the two resistors are arranged in the opposite way around.

So V bridge is going to be equal to V0 R upon R plus M. That's just a potential at a here-- minus the potential at b, V0 M over R plus M.

And if you look at this, quite simply, you can see that it's proportional to R plus-- or R minus M. So it's a differential measurement. We're just measuring a voltage which is proportional to this small quantity. And that is entirely due to the change in temperature. So we're now able to isolate the change in temperature very, very precisely.

And, in general, we don't just use this Wheatstone bridge, we also use a phase locked loop measurement, which is a form of heterodyning again, my favorite technique. And so, we actually have V0 is oscillating at around 50 kilohertz or so.

And so, we can see that our signal, delta T, or the change in temperature as a function of time, is a slowly varying signal which is embedded on top of this 50 kilohertz. And then, using heterodyning techniques, we can measure it very cleanly without noise. This may be an experimental technique you're familiar with.

What the bolometer head actually looks like is we have some large lump of gold here. And by large, I mean it's about 100 microns thick and about 1 millimeter long. This is what faces the radiation. The radiation is coming in from this side.

This is often made out of gold. Gold is chosen because it absorbs all wavelengths relatively evenly. So you don't need to worry too much about the spectral response of this. This thin layer of gold has been deposited on top of a substrate.

And it's on the back of the substrate that you have your resistor. And your resistor is literally a little zig zag of gold deposited on the back side of this. So this is M or R. And depending on whether it is M or R, you either have this open to the plasma or you have a thick block some distance in front of it so it can't see the plasma.

It's the heat capacity of the absorber, c, that goes into our previous equation. And it's the heat transport kappa grad T that gives us the time constant for thermal conduction through the substrate from the absorber to the resistor.

And I'll just write that equation again because it's disappeared now-- that we can work out the radiated power by looking at how the temperature of the measurement resistor M changes in time. And that change in time is due to heating up and heating down from direct heating and then also a time lag phenomena delta T on tau.

And, effectively, this tau here sets the timescale at which we can measure. So the larger tau is the slower our measurement of the radiated power is going to be.

And if tau gets very large, because we've got a very thick substrate here, or it doesn't have very good heat transport, then we're going to have a very poor time resolution of our radiation power. And, of course, we'd like to have a nice time resolution of that.

I'm just going to finish up with a couple of details. Then I'll take some questions and we'll leave it there. So the trouble with these, as someone asked, when it comes to radiation, first of all, they do actually sense, basically, all energy coming out of the plasma or all power coming out of the plasma. It's not just radiated power, but they will also sense neutrons and particles.

So if you've got ions or electrons coming out and hitting these and being absorbed, that will also heat up the absorber. This will measure those as well. So it's very hard to tell the difference between those.

But, also, the neutron damage leads us to use much thicker substrates, which gives us a longer time response and so, therefore, a worse bolometer. So, ironically, the bolometers that will be used on [INAUDIBLE] are significantly worse than the bolometers used on existing devices, not because they couldn't work out how to make better ones, but because the radiation forces you into a regime where you can't use a good bolometer anymore.

And I'm sure that [? Spark ?] will have exactly the same problem. Almost everyone in the world uses this design of bolometer that they pioneered on ASDEX Upgrade in the '80s. No one has come up with a better system yet.