[SQUEAKING]

[RUSTLING]

[CLICKING]

JACK HARE: So Thomson scattering-- as you remember, we set up the problem so that we had some scattering particle here. We had some incoming wave with a load of parallel wavefronts at some plane wave. It's got a wave vector Ki in the i hat direction here. We said that there was some origin to our system, and the particle was at a distance R from our origin.

And we were going to observe the scattered light in another direction with our spectrometer here. And that scattered wave is also going to be a plane wave that's going to have a wave vector Ks in the s hat direction. There's going to be two vectors, one from the particle and one from the origin, pointing towards our observer here. And we call those R and R prime.

But one of the first approximations we said is that imagine that plasma where all of our particles have some vector R is relatively small compared to the distance to our observer. And then we can say that this-- our vector is roughly the same as the R prime vector here. The position of our particle, R of t, is equal to whatever it started off plus its velocity times time. And so we're assuming that these particles are going in roughly straight lines.

And we also wanted to distinguish between the time t prime at which the particle is scattering the radiation and the time t at which we're observing the radiation. And we found out through a little bit of algebra that t prime is equal to 1 minus i hat dot beta over 1 minus s hat dot beta, where beta in this case is defined as the velocity over the speed of light.

And this was effectively the origin of our double Doppler shift. The particle sees a Doppler shift of the wave that's incident on it, and then the observer sees a Doppler shift of the emitted wave because the particle is moving.

And we found that the electric field-- if we said that we had some incident electric fields E sub i at R and t prime, which was equal to some electric field strength and some polarization. That's why this is the little vector. And we're going to assume it's a plane wave, so it's just cosine of Ki dot r minus omega i t prime, like that.

So it's this wave with the incident K vector and the incident wave frequency. Then our scattered radiation that we observe at R at time t, which is not the same as t prime, was going to be equal to the classical electron radius over the distance, because this thing is scattering into 4 pi steradians. And then there was a factor that looked like s hat cross s hat cross the incident radiation-- whoop-- at r prime and t prime like this.

So we scatter some of that incident radiation. We scatter it into a shape that's given by this s hat cross s hat cross. And later on I'll start referring to this, including that second cross product, as just a tensor capital pi, which transforms the incoming electric field here.

And we said that if the velocity of our particle-- this v here-- is equal to 0-- if we're just dealing with a stationary particle-- then our scattered electric field that we observe at our detector is simply going to have a very simple form, which has-- well, rE on Ri acting on the strength of the electric field.

That gives us that nice doughnut-shaped scattering pattern. And then we'll have a cosine factor here that simply looks like Ki dot r of 0-- r of 0 because the particle isn't moving, so that's actually r of all time-- minus omega i t plus a factor Ki dot r.

And this final Ki dot r factor here is for all our particles because we basically said they all pick up that phase as we go here. But actually all the particles are clustered in a little plasma around about here. So this was quite boring. We just get a scattered wave, which has the same frequency and K vector as the incident wave.

But if we then allow our particles to move around, we found out that we get Doppler shifts, as we would expect. And we're also dealing with a regime where we're looking at elastic collisions. So this elastic collisions means that our incident K vector has roughly the same size as our scattered K vector. So we don't change the momentum of our photon very much.

And this is-- bless you-- valid when our photon energy is much, much less than-- I've written it in terms of particle velocity-- interesting-- of me ve squared. But I imagine this is also true when it's non-relativistic as well, so I'll just write that down as well. If we violate this limit-- if we start scattering energy photons-- we get Compton scattering, not Thomson scattering. We have to do a full relativistic treatment.

So this elastic collisions constraint comes down to us scattering off fluctuations with a K vector, which is defined as the difference between the scattered K vector and the incident K vector and a frequency, which is equal to the difference in the scattered frequency and the incident frequency.

And this difference is also equal to the scattering K vector-- the difference between the scattered and the incident dotted into the velocity of the particle. So this looks like the classic Doppler shift that you're used to seeing.

So the punchline of all of this is if you input a wave with Ki omega i, you put your laser beam through the plasma, and you measure some scattering in some direction Ks-- which, remember, you have control over because you get to choose where you put your spectrometer around your plasma here-- and you measure it on a spectrometer with some frequency omega s, then you can infer that, in order for you to be able to see this scattered light from your original probe signal, there has to be something within the plasma capable of carrying away a momentum K and an energy omega, or giving a momentum K and an energy omega to your photon here.

And the strength of the signal, so the intensity at omega and K, is going to be proportional to the number of scattering things-- scatterers. And these scatterers could be, for example, particles, so electrons, or they could be waves, and I'll often refer to these waves as modes. And we will get on to-- and these modes in this case have an apparent velocity, which is equal to omega/K, so that's our phase velocity here.

So this was a very hand-wavy way of trying to derive all of this. We're going to make it more rigorous by going through some Fourier transformations of all of these equations. But that's just a recap of what we talked about last week, so-- questions? Yes?

AUDIENCE: [INAUDIBLE]

JACK HARE: So there is no vector omega. Omega is a scalar [AUDIO OUT]. That's OK. What I like to think of them as is the momentum-- if you're thinking of it as in a particle picture, then h bar K and h bar omega is the energy and momentum of the electron that scatters them.

If you're thinking about it as a wave, you don't have to times it by h bar to get that. You can just say at some wave with K and omega-- so like a sound wave-- omega equals sound speed times K, for example. And we'll talk a lot more about that later on, so.

But, yeah, what these represent is something that can scatter this light. If you have no plasma and you fire a laser beam through a chamber, you really shouldn't expect to see any scattered light at some random angle there because there's nothing there to scatter it. So if you see some scattering, you have to say, aha, there must be somewhere something in that chamber to scatter it, and it must be able to change the momentum of my photons by K and their energy by omega.

And that limits the sorts of things that can exist inside your plasma, so. Especially for the waves and modes, they'll only be scattering at very discrete values of omega and K given by the dispersion wave. Yeah?

AUDIENCE: [INAUDIBLE]

JACK HARE: Oh, you mean like these?

AUDIENCE: yeah

JACK HARE: We'll get on to that in a great deal of detail later on. So these four lectures are like a wave washing on the shore. So we're going to derive Thomson scattering three times. We did the first time last week, we're going to do the second time today, and we'll start on the third time as well.

And each one-- the previous derivation is just a simplification of the next derivation. And previously, we got this very, very rough picture that, if we have a single particle, we might be able to say something about the particle's properties, its velocity, if we see scattering at a certain frequency and wave vector.

Next, we're going to do incoherent scattering, which is where we scatter off the distribution function of electrons. But these electrons are not correlated. They're all moving incoherently. And so we basically get scattering of each of the electrons individually. And then, finally, we're going to do collective scattering, where we scatter off the collective motion of the electrons where they're all moving together, which is a wave, so yeah. Other questions? Anything online?

So before we get going with Fourier transforms, I'm going to write down some assumptions. And these assumptions don't necessarily have to be true for every plasma that we're dealing with. But they are assumptions that we've made in order to do this Thomson scattering derivation. If you want to violate these assumptions, you need to go and check where this assumption comes in and re-derive your Thomson scattering.

So first of all, we are scattering from some volume V, which has N electrons inside it. The electrons have a charge of e minus, and there are N/Z ions with a charge of Ze. This is pretty standard. But this is effectively a statement of quasi-neutrality.

Secondly, we are mostly going to drop relativistic corrections. So if we want to look at the fully non-relativistic case, we will drop terms on the order of V/c. So if we do some sort of expansion and we see a term with V/c, we'll just drop it. So we're only keeping the zeroth-order terms.

If we want to do this sort of quote marks "relativistic," then we will keep the order of V/c terms, but we'll drop the terms on the order of V squared upon c squared. So you can see this is like a Taylor expansion. We're going to claim that they're much less than 1.

If you want to do the really relativistic, because obviously this still neglects some relativistic terms, then you should just go and derive Compton scattering instead, which we're not going to do. I cannot stress this enough. There will be no quantum. I don't like it.

There's no scattering from ions. We discussed this already. Because of the mass difference, the ions scatter much, much less. We will no longer think about the scattering directly from the ions, but we will think a lot about the scattering from electrons which are dragged around by the ions.

We're going to assume that our observer is at a distance R which is much, much larger than the volume within which the plasma is contained to the third power. So this is a length scale now-- V to 1/3. So this is effectively the max value of this little r vector here. So the plasma is small, and it's far away.

And this in turn-- this length scale here-- is much, much larger than the wavelength of the incident light. So we're seeing a large number of charges within a wavelength. They're not scattering off single particles.

This is equivalent to saying, of course, that there are many electrons in V-- in the volume-- and that V is small. And this is what we needed to justify that our R prime is approximately R-- our argument here. These two vectors are roughly the same.

What are we up to? Five assumptions. Assumption 6-- we're going to assume that the frequency of our incident light-- the light which is being scattered-- is greater than-- significantly greater than-- the electron plasma frequency.

Don't get fooled. There are going to be lots of i's showing up here. That's always for "incident." I am not thinking about the ions anymore. They're too slow and too heavy. So if we make this approximation, I can approximate the refractive index of the plasma as basically being 1 so I don't have to think about those effects. In reality, you might want to put in a small correction for this, and it's actually quite easy to do.

We also have the assumption-- I think this should be number 7, frankly-- no probe attenuation. So the optical depth at tau evaluated at our probe frequency is much, much less than 1. So the beam doesn't get absorbed. This just makes it much easier to do the calculations. If your beam was getting absorbed, you'd get different amounts of scattering depending on where you are on the beam. It's a pain, so we're not going to do that.

And then, finally-- we touched on this earlier-- the velocity induced from the electric field of the incident radiation we are going to assume is much, much smaller than the thermal velocity here. And so this was saying that our particles are going in straight lines at some velocity, like the thermal velocity. And the electric field is just making them wobble ever so slightly along those lines here.

And this is equivalent because this equals this velocity here from the electric field-- you can work it out as being e times the strength of the electric field over m e omega i-- if you set that to be less than V thermal, then you find out that the power-- the intensity of your laser beam-- which is equal to c epsilon 0 strength of the scattering electric field squared-- it limits this power.

I won't derive it directly, but-- so you can't use-- in order for this assumption to be valid, which we'll use a lot in our derivation, our laser cannot be too strong. So if you end up using a very strong laser, you have to re-derive this using the full orbit of your particle, which includes this strongly perturbing electric field, and that's a pain.

So obviously, that's not a good reason not to use a very strong laser. You want to use a nice, strong laser in order to get this Thomson scattering effect. But you should be careful about whether your Thomson scattering formalism is still valid if you use a very strong laser. Any questions on any of these assumptions? Yes?

AUDIENCE: [INAUDIBLE]

We are just dealing with these particles as like points. So we're going to use a Klimontovich distribution function at some point, which has delta functions representing the position of the particles in space and in velocity space. And so we don't have to worry about uncertain things like that. And so we're not thinking about particles as waves. We're not thinking about a semi-classical treatment of electromagnetic radiation scattering particles. So this is a classical treatment.

AUDIENCE: [INAUDIBLE]

JACK HARE: I would have thought it's when you get to short wavelengths, so where you're getting to short wavelengths, which are getting close to the classical electron radius, so not when you're getting close to the Debye length.

We'll actually talk about the fact that the Debye length is really important to this. We can definitely have wavelengths which are above or below the Debye length. But I would have thought when you get down to wavelengths on the order of the classical electron radius, then you'd need a quantum treatment. Yeah. [INAUDIBLE]

Depends. If your plasma is very hot, the relativistic corrections are hard to avoid. And I'll talk a little bit about the "relativistic" corrections, which are normally used in a tokamak when you get to like 10 kiloelectron volts, because then your rest mass of your electrons is 500 kiloelectron volts, and 10 kiloelectron volts is not nothing, and you need to have a small correction.

I would have thought this one is often quite hard to obey, but it's also quite easy to fix up. You just need to make sure that the dispersion relationship for your incoming light and your scattered light obeys this equation as well, so that it obeys a dispersion relationship where N is not equal to 1.

That's really hard to achieve in a lot of plasmas if you have inverse Bremsstrahlung. Because your plasma is relatively cold and dense, then you have very strong probe attenuation that's been seen in lots of plasmas. Yeah. Quasi-neutrality is pretty good. Quantum's pretty good. No scattering from ions is pretty good. This one is pretty good. It's actually hard to get a laser that intense, so yeah. Other questions? Yeah?

AUDIENCE: [INAUDIBLE]

JACK HARE: Optical wavelengths, because they are nice to work with. [CHUCKLES] But people do X-ray Thomson scattering as well, which is really complicated, and we won't get into it. But that's where things get a bit more quantum. But, yeah, optical wavelengths are nice to work with. People have done Thomson scattering with CO_2 lasers, which are about 10 microns. So that's about 10 times longer than optical wavelength.

And then the first Thomson scattering measurements were actually done off the ionosphere by Salpeter, and he used radar. So he did Thomson scattering off the plasma that surrounds the Earth and measured its properties, but-- that's when the theory was first derived. But in a lab, we don't use radar because our plasmas are smaller, so-- yeah?

AUDIENCE: [INAUDIBLE]

JACK HARE: No, this was just-- the previous derivation was just to give you a heuristic feel for a problem-- yeah. The previous derivation was just a single particle. So none of these really talk about single particle-- well, no, they do. Look, we've put lots of electrons inside our volume, yeah.

AUDIENCE: The last [INAUDIBLE]

JACK HARE: Yes. Yeah. Yeah. We're going to use this as the orbit-- the "orbit"-- of the particle as its equation of motion. And if that's not true, then you need to solve. It will be very-- it's already quite non-linear. It'll be very non-linear at that point because it will depend on how the electric field scatters. The scattered electric field will also distort the orbit of the particle.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yeah, exactly. Any questions on that? So let's talk about scattering for multiple electrons. So imagine I have now my incident wave coming in, Ki, and I'm still observing along Ks. And now I have a collection of particles here, and they are all oscillating, and they are all scattering light. And so I'm getting a load of e scattered from each of these.

And so now the total scattered power, which is, as you remember, dP d omega s d-- oh, no, just dP d omega s OK, fine, so we're integrating over [INAUDIBLE]. Doesn't seem quite right. [VOCALIZING]-- I'll throw this back in. What did we do in the previous lecture? No. I'll leave it as that.

That is now going to be equal to R squared c epsilon 0. Previously, we had a term that looked like e scattered dot e scattered star. That was for a single particle. So now we replace this with a term that looks like the sum over j and l to N of ej dot e l star.

This is basically summing up all of these electric fields, and these electric fields now interfere with each other. I shouldn't have put these vertical arrows here. This is a time average, remember. There we go-- like that.

And it turns out we can split this into two terms. So we split this into the term-- they're both R squared c epsilon 0. And then we have a term that looks like N times es squared. And this comes from the sum where we have j equal to l. So this is effectively the scattered light of each of the electrons summed up-- and there is N electrons scattering-- plus a term that looks like N times N minus 1 ej dot e l. And this is for j not equal to l.

So we've split this sum up into the terms where we're just talking about the electric field from the same particle interfering with itself, which is obviously constructive, and the terms where these are not equal, so when the scattered electric field from one particle interferes with the scattered electric field from another particle.

You might notice that there's an N minus 1 here because once we've picked which particle is j, there's only N minus 1 particles to be L. But you may also realize that in a real plasma, the difference between N and N minus 1 is very small because there's a lot of particles. So we can just write this as N squared most of the time.

And so in this expression, this first term here is called the incoherent term because this just depends on each of the particles independently. And this second term is called the coherent term. And it's called the coherent term because, unless there is something correlating these particles so their electric fields have some correlation, you would expect them all to be random with respect to each other. And then when you dot them into each other and average them out, you would expect that randomness to average out to zero. So to be coherent, we need correlation, and we'll talk a lot about that later on.

Now, what's the smallest length scale in a plasma? Yeah? The Debye length-- yeah, exactly. That's what makes it a plasma. So if you are probing this plasma on length scales smaller than the Debye length, you're effectively going to be looking at something that isn't a plasma.

And by probing, I mean if our wavelength of our incoming light here, which is lambda i, which is 2 pi upon Ki-- if lambda i is smaller than the Debye length, then we end up with incoherent scattering. So this wavelength here is smaller than the length scale on which the particles can organize themselves and actually look like a plasma. So instead, we just end up looking like a random gas.

But if we end up in the opposite limit, where the Debye length is less than lambda i, we get coherent scattering. Because on this length scale here, we end up being able to see the plasma as a coherent object. So this coherence is expressed in the form of waves here.

And so we end up with a very important parameter that we call alpha, and we define alpha as 1 upon K lambda Debye-- yeah, kind of fibbed here. We'll get back to that in a moment. And if you have alpha less than 1, you're in the incoherent regime. If you have alpha greater than 1, you're in the coherent regime. This shouldn't have been lambda i. This should have just been lambda, and this lambda here is 2 pi upon K, where K is that Ks minus Ki.

So this is saying, if the mode that you're scattering off has a wavelength, which is less than the Debye length, then you're just scattering off individual particles. If the mode you're scattering off has a wavelength that's greater than the Debye length-- you're scattering off the full plasma-- you're scattering off the waves.

And you get to choose which regime you're in because-- well, you don't get to choose how dense and how hot your plasma is, which is what goes into Debye length. But you do get to choose this K by where we place our scatterer. Let's see. I had that nice formula for K previously.

Remember that K is equal to 2Ki sine theta upon 2, like that. So, by choosing my angle theta, I can make this Ki-- multiply by something arbitrarily small, I can make K arbitrarily small, and so I can make alpha arbitrarily large. So I get to choose by where I place my detector not just the scattering angle, but, also, I get to choose which regime I'm in.

Now, it may be that, for some plasmas, this theta is so very, very, very close to the direction that the laser beam's pointing in that we can't realistically do coherent scattering. But there is still a regime where the scattering will be coherent there.

And if this doesn't make sense, again, we're going to do it mathematically later on. But this is just to give you an idea of where we're headed. And the next thing that we're going to be focusing on is this incoherent regime. Yeah?

AUDIENCE: [INAUDIBLE]

JACK HARE: It is-- it comes from the sum of all of these together. Yeah, exactly. Effectively, the wave correlates the electron density fluctuations such that, when you scatter off those electron density fluctuations, all of the electric fields from them end up at your spectrometer in phase, and they constructively add up.

So if you've done X-ray spectroscopy and you've looked at Bragg scattering, some people like to look at this in terms of Bragg scattering. But I find most people these days haven't done X-ray spectroscopy and Bragg scattering, and so it's not a useful way to describe it. But if you've seen it like that, the wave makes a grating, and that grating, therefore, gives us constructive interference. Other question?

AUDIENCE: [INAUDIBLE]

JACK HARE: Just how it's defined. You'll see later on-- if I think we didn't put it in there, there'd be lots of 2 pi's showing up everywhere. So when we derive this more formally, you'll see this factor alpha show up in the equations. And it will literally be a parameter alpha that multiplies the incoherent term and like a 1/alpha over alpha that multiplies the coherent term-- or the other way around-- other way around.

And so, if alpha is largest, incoherent term gets suppressed, and the coherent term gets enhanced. And if I put it in with your lambda over lambda E, which is perfectly reasonable, I'd have like a 2 pi squared or something somewhere, which I don't want to do. So this is a better definition. Yeah. Yeah?

Yeah. So, again, we will get to that. I don't think anything I've said here reaches the level of rigor where we can really say that or not. So this is just like a guide to where we're going next. But, like I said, we're going to derive the equation where alpha shows up explicitly as a parameter.

And then we'll say, OK, depending on the value of alpha, you see more incoherent or coherent. And it's important to note that, when alpha is about 1, you have both of these spectra at the same time showing up-- both of these effects. So you're probably going to end up seeing some of both types of scattering. Yeah. Any questions on that? Yeah?

AUDIENCE: [INAUDIBLE]

JACK HARE: Oh, yeah, that's fine. Don't worry about it. What's 2 pi between friends? So this is-- again, this is like, what does it mean to be much less than 1? What does it mean to be much greater than 1? Does a factor of 6 matter if I tell you something is much less than 1-- if it's much greater than 1? Yeah. If that was the point, yes, fine, but-- yeah. [CHUCKLES] I can rewrite these, I guess. I can rewrite like a k for Debye length, but then that's weird because you don't know the formula for that, so yeah.

Let's keep going. I'm going to need a lot of board space. Now we're going to do the Fourier transforms. So we're looking at incoherent scattering. You won't even notice where I've made the assumption of incoherent scattering. You'll only spot it, probably, when we go back and do coherent scattering. But just to be clear, we're only going to get the incoherent component out of this treatment.

So our scattered electric field measured at our detector at time t is equal to r e over R-- classical electron radius-- over the distance to the spectrometer-- this tensor pi, which, remember, is equivalent to s hat cross s hat cross this, is just something that tells you the shape of the emitted field here. And this tensor is acting on the scattered-- or the incident electric field at position r and at time t prime.

So if I do the Fourier transformation of this and I want to know what the frequency is-- the scattered frequency-- and because-- [CHUCKLES]-- so we're going to be swapping again very freely between omega s mu s-- so, once again, factors of 2 pi flying everywhere.

But it-- so I'm going to write this as a function of mu s because this is what we measure, and this is what Hutchinson uses in his book. But what I'm doing with Fourier transforms-- I'm going to be using omega s. So just feel free to keep adding 2 pi's and getting rid of them as you see fit.

So this is just a standard definition of a Fourier transform. We have the electric field in time times by e to the i omega s c t dt. Now, we already mentioned we have a discrepancy between t prime, which is going to be inside this, and t out here. So t prime is equal to t minus R minus s hat dot lowercase r upon c.

And we can also say that the dispersion relationship for our scattered wave is just the vacuum dispersion relationship. So K is equal to omega upon c s hat. By the way, this is one place where, if you want to modify this to take into account the O mode dispersion in a plasma, that you can modify this here. But it doesn't actually make a big difference to the derivation.

This is going to allow us to replace this s hat here like that. We're also going to say that dt is equal to dt prime times a factor of 1 minus s hat dot beta, like this. And we said that anything on the first order in v/c is much less than 1-- we drop it.

So, in fact, we're going to just switch between t, e t, and dt prime. We can't do the same for t prime and t. They are not the same time coordinate, but the amount they change by is the same because we're not dealing with relativistic time dilation in this treatment here.

So now I can go back to this, and I can start replacing all of my t's with t primes here. And I'm going to end up with r e upon R-- am I really going to make that much of a jump?

[VOCALIZING]-- I'm not going to make that much of a jump. I just need to go back up here and remind you that this Ei has a factor that looks like cosine of-- oh, this isn't going to work. I don't want to do it like that. And a factor of cosine Ki dot r minus omega i t prime.

If I make the replacement of t prime with t, then this becomes cosine of Ks r, which is the shared phase for all of the particles. This is just boring-- just the phase you pick up as the wave travels to the spectrometer-- minus omega s of t minus Ks minus Ki dot r of 0 that we saw last time.

So when we go back down here-- I'm going to bring this boring phase term out the front. I'm going to turn everything into exponentials. So I'm just going to replace this cosine with an exponential i times this phase factor here. I'm going to bring this phase factor exponential of iKs dot r out of the front here.

But what I'm going to leave inside the integral is the integral of the vector-- the tensor pi dotted into the incident electric field at r and t prime. And all of this is multiplied by this factor-- the Fourier transform factor. But now I've taken t, and I've replaced it with t prime using this equation. And so this exponential now has a term, which looks like i omega s t prime minus Ks dot r. And this Fourier transform is now in terms of dt prime instead of dt.

I then remember that this electric field here has this dependence-- Ki minus omega i t prime. If I take that as an exponential, it will add with all these exponential factors here. And we will end up with an equation that looks like r e upon R, exponential of this boring phase factor term iKs dot r, integral of i dot Ei0.

So this is just the strength of the electric field and its polarization-- exponential of i omega t prime minus K dot r dt prime, where, again, I'm using those definitions of omega and K that we had previously. So this omega is omega s minus omega i, and this K is Ks minus Ki.

So I've now got something where all of my t parameters are t prime parameters. So I can now do this integral. I couldn't do it up here because I had a mixture of t primes inside the scattered electric field and t's for where I really am. So you've really got to take care of doing all of these Fourier transforms carefully.

This pi dot Ei0-- that can go outside the integral as well. It's just constant in time. Remember, that's just s cross s cross something like that. That's set by where our spectrometer is. And we're assuming we've got an infinite plane wave but intensive electric field change.

So what is this Fourier transform? Fourier transform of an oscillating sinusoid. Just a delta function. Thank you. There's a factor of 2 pi from the Fourier transform when we get the definition of the delta function.

And then we've still got this tensor dotted into the electric field vector to give us a scattering pattern. Now we have a delta function K dot v minus omega. And if you want to follow along this in Hutchinson's book, this derivation here is Hutchinson's 7.2.15-- like that.

So this is actually nothing new. This is-- once again, we have just shown that our scattering is going to be off-- the scattering of the particle is going to give us a scattered frequency omega, which is K dot v. But now we've done it in terms of Fourier transform. Previously, we did it in real space and real time, and we just hand-waved our way through it. Now, we've got the spectra actually in a much more-- slightly more rigorous way.

So this is still just a scattering off a single particle. If we want the scattered power off this single particle, then we again go to c epsilon 0 r squared. And now we need to do the time averaging. So to time average, we integrate between minus t/2 to t/2, and we divide by 1/t here.

We're going to take the limit of t going to infinity in a moment so it will start to look more like a Fourier transform. And this average power, again, is defined in terms of the time. So Es is a function of time. Es star is a function of time, integrated d time.

But if we take the limit as t goes to infinity, we can use something called Parseval's theorem, which says that this quantity here is equal to the integral Es in frequency space, dotted with Es in frequency space, integrated over frequency, still with a 1 over time out the front here. So this is a kind of-- [VOCALIZING], I'm not quite sure about the 1 over time out the front there-- Parseval's theorem. Dimensions don't work.

I think I should have paused earlier and asked for questions, because this is actually an excellent derivation. So let's just go back-- draw a line under this here-- see if anyone has any questions on how we got to here. OK, then, we'll go on. So now we're trying to calculate-- yeah, all right.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yeah. Yeah, sure, sure, sure. I will pause in a second [INAUDIBLE]. I just want to get through this section first.

AUDIENCE: [INAUDIBLE]

JACK HARE: No, this is just the incident-- i Ei0, which is just the strength and the direction of the electric field, without any of its-- so it's-- the incident electric field is Ei0 times by this cosine. So I put all of the fluctuations and interesting time and space variation in the cosine. But this is just our strength for our electric field here. So all this is saying here is that the scattered electric field is linear in the incident electric field. All of the funkiness with time and space has gone into the delta function.

So here, we want to be able to evaluate this function, but we don't actually have a great expression for Es a function of time. But we do now have a great expression for Es a function of frequency. So we use Parseval's theorem to convert between the time-power spectrum and the frequency-power spectrum.

And that means we can now write down our scattered power dP into some scattering volume or scattering solid angle, omega s, and some scattering frequency. And this is equal to c epsilon 0 r squared of K-- the limit of t going to infinity. So we're taking a long time slice. This allows us to use our Fourier transforms where the limits here are plus and minus infinity here.

And we do the Fourier transform of this-- sorry. Don't do the Fourier transform. We've taken d nu from this side here, and we've put it down here to form our scattered power per frequency bin here. So we're getting rid of this integral now. And now we're trying to evaluate this term, which looks like Es dot Es star instead. And we have that equation up here. So we want this term-- I'm going to take it from here.

So, first of all, we have a term that looks like r E squared upon R squared. Then this term will cancel out because we'll have it times by its complex conjugate, so we can get rid of it. Good riddance. There will be a term that looks like 4 pi squared here. We'll have a term that looks like pi dotted into the incident electric field squared. And then we'll have a term that looks like delta squared K dot v minus omega.

The trouble with this is that we have a delta function squared. It's a pretty odd object because a delta function is already a pretty odd object, and squaring it makes it just odder. So we're going to do the sort of thing that makes the pure mathematicians wince. And we're going to say, [VOCALIZING], this is a little bit like a sinc function. So a delta function-- it's sort of like a sinc. And if our sinc function of t-- or, in this case, I guess, nu-- as this nu gets very, very small.

So we have a function that looks like a sinc, and then we approximate the delta function as just a very, very, very narrow sinc. And if you do this, you can rewrite this delta squared as looking like t upon 2 pi, just times the delta function itself.

So delta function squared of frequency is equal to t upon 2 pi delta function of frequency. This is by far the least satisfactory part of this derivation, by the way. But Hutchinson doesn't even do it. He just skips over it. So I'm trying to at least demystify where this might come from. But I don't know how to do this very, very rigorously.

Let's just go with it for now. We're going to replace this delta function with just a delta function not squared because those are the sorts of things we can work with. Then we can run through all of this, and we can find out that we have r e squared c epsilon 0 limit of t going to infinity of 1/t times t upon 2 pi-- that's this factor here-- times by 4 pi squared. We still have this i dotted into the electric field strength squared.

And now we just have-- I'm running out of space-- this delta function of K dot v minus omega not squared. The convenient thing about doing this is we've generated a factor of t which can cancel out with this t. So we've canceled out an infinity or a 0 that we've accidentally introduced. And now we've also got a 2 pi, and that will cancel out some of these 2 pi's. The whole thing starts to look a little bit more tractable.

And so the result is-- I don't quite have space for it on this board, which is a real shame after we've come all this way. I shall write it here-- dP d omega d nu is equal to r e squared c epsilon 0 Ei0 squared. What the hell is this? I'm going to keep it as it was-- i dotted it into Ei squared times 2 pi delta K dot v minus omega.

The point is that this is what we should call our spectrum. We're looking at the power scattered in a certain direction into our spectrometer per unit frequency. And so we can now evaluate this. And once again, we can say, if we're looking at power per unit frequency-- or rather, dP d omega-- the power scattered for a certain solid angle-- this tells us that, from a single particle, our spectra will just have a little delta function then.

And that delta function will be at a frequency omega, which is equal to K dot v. So if our initial laser line was at some wavelength like this-- this is omega i-- then this is at omega s equals omega i plus omega-- and we say, aha. This is blueshifted. It's of higher frequency. And so means the particle is-- got some component of velocity towards us.

Again, this is all just single-particle scattering, but we're now working in the Fourier domain as opposed to in real space, which means we can actually write down equations for the spectrum, which we couldn't do when we were working in real space. We just had a hand-wavy argument for it.

Now I will pause. Oh, one more thing. Just so you can have contact with the book again, this is Hutchinson's equation 7.2.19. So in five equations, he does all of this. I have tried to spend a little bit more time doing it. But I will agree with anyone who says this isn't [INAUDIBLE]. Questions? Yeah?

AUDIENCE: Where did that 1/t [INAUDIBLE]?

JACK HARE: [VOCALIZING] Yeah. So what I have written in my notes is 1/t here. But if you set these two equal to each other, the dimensions don't add up because of the limits here, or because of what we're integrating over.

So in a panic, I changed it to be t because that made the dimensions work. But it makes the rest of the derivation not work. So probably, there's some reason why we have this 1/t in front of both of them that I've missed. And I'm very happy to go and have a look at the derivation again in Hutchinson's book-- try and work out what's going on.

But what we are trying to do is replace this integral here over the electric field in space with this integral over the electric field in frequency, but the dimensions don't quite work out. So that's where the 1/t was meant to come from, and that 1/t is the one that gets canceled by the t over 2 pi that we get from our shoddy treatment of the delta function. Yeah. Other question?

AUDIENCE: Yeah. So this whole [INAUDIBLE]

JACK HARE: If you're doing a Fourier transformation, you have to do it over infinite time. So this is just simply-- but if you put that infinity in straight away, then you start getting 0's. So we as a t, where it's a really big t. But then eventually, we cancel out, so it's like you don't actually care how big it is.

So to do this rigorously, you actually have to do Laplace transformation because the laser is not-- cannot be a plane wave. It's not traveling through the plasma for all time. It turns on at some time. And so you have to do Laplace transformations where you only take into account for t greater than 0. But the mathematics of that becomes harder, so I'm doing these Fourier transforms, and there, you need infinite time. Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yeah. Sorry?

AUDIENCE: [INAUDIBLE]

JACK HARE: This one?

AUDIENCE: Yeah. That represents the time average?

JACK HARE: Yeah. This is where-- the time where we're working out the average power because we're working out the instantaneous power, we're integrating it over some interval T, and we're dividing by T. Yeah.

AUDIENCE: And then we're [INAUDIBLE].

JACK HARE: And then afterwards being like, yeah, but-- how long do you need to do this for? Well, to do this for a plane wave, you should do it for a few cycles. But to be rigorous, you want to extend it out-- well, maybe not rigorous in this case. To be slightly more rigorous, you want to extend that integral to all time if you're really doing a Fourier transform type approach.

Yeah. And this is the requirement of using Parseval's theorem, which I clearly don't know as well as I thought I did because of this problem with the units here. Yeah. Yeah. Yeah, I'm confused, because this is definitely something that has dimensions of electric field squared. This does not have dimension of electric field squared. Any other questions?

The nice thing about our treatment is we said that, if we want to get the scattered spectrum from many particles, we are not interested in how the scattered electric fields in those particles interact with each other. That's the coherent part. This is the incoherent part.

So if we now have the scattering spectrum from a single particle and it's showing up at a frequency that is to do with that single particle's velocity, if we want the scattering spectrum from many particles, we simply sum the scattering from each of the particles V together. So this now, you're probably thinking, starts to look like integration over a distribution function. So let's go do that. Do I need to write that down? No, I'm happy with that.

Actually, I'll write it down again here so we've got it on the page as we look at it. So our scattered power into some solid angle into some frequency looks like r e squared times i dot-- why am I writing it -- put an i everywhere else. I'm just going to call it i again.

So I've taken the electric field strength and polarization-- I've now split up even further into a term that's the electric field strength and its polarization. You get the polarization inside here. I've taken out this factor of Ei squared because, when times that by c epsilon 0, I can put all of that together and write it instead as the incident power of the laser over the area of the laser.

So this is now an intensity in watts per meter squared here. And we often think about our lasers in terms of their intensity. This is nice units. And then, finally, the interesting bit, which is this 2 pi delta K dot v minus omega. So this is a scattering from one particle. And to be clear, it's one electron.

So, again, nothing particularly new here but just now in frequency space. So now we want to integrate over some distribution function f of v d3v d3r. And remember, when we have a distribution function, we say that the density is the integral of f of v over velocity space here. So this is the density as a function of r-- the zeroth moment of the distribution function.

What we're going to end up doing is we're going to be integrating this f of v times this delta function K dot v minus omega d3v. When we do that integral, this delta function is going to pick out certain things about this velocity distribution. It's only going to pick out particles which have a velocity which is equal to omega/K and it's only going to pick out particles which have a component of velocity parallel to K. That's the v dot K part as well.

So when we do this integral, we end up with a scattered power spectrum that looks like-- scattered power, again, into some solid angle-- some frequency-- is equal to 2 pi r e squared, intensity of the laser beam, a factor which accounts for where the scattering light goes to, and then the interesting bit. And the interesting bit is-- as opposed to having the full distribution function, we now have a distribution function f sub K evaluated at omega/K.

And then there's also a factor of 1/K out the front here. So let me unpack some of these bits here. This f of K-- imagine we have some full distribution function in three dimensions where we have some coordinate that is perpendicular to K and we have some coordinate, which is along K.

And what we're doing is the integral of this function, integrating over all the perpendicular velocities. So our f of K is simply a one-dimensional slice through that distribution function. You can think of-- maybe, for example, in-- just in two dimensions here, we could have the K direction and the perpendicular to K direction.

And our distribution function could look like-- and I'm drawing contours of f here. So, for example, this is f equals 1, 2, 3 type thing. These can be much bigger numbers, but this is just to give you an idea of some function which is peaked around about the center here, but it's not isotropic.

And what our Thomson scattering is picking up is just a slice through in one direction of that distribution function here. So if I take this cut, I get out the distribution function that looks like f of K versus vK, and maybe it looks like this. So I'm only sensitive to one-- if a particle's moving in one direction-- in the K direction here.

And this 1/K-- if you're wondering where that comes from, this is because dv is equal to d omega upon K. So when I change this integral from being in terms of velocity space to being in terms of frequency space, I'm going to get out this factor of 1/K here. So why did I write all of that? No, this is right. I don't need anything else. So, just to be clear, this is a 1D distribution function.

Now, if your system is isotropic, this doesn't bother you at all because you can-- if you think there's lots of collisions, there's no magnetic fields, my distribution is isotropic, then I can just make one measurement of the distribution function in one direction, and I've got all I need.

But if your distribution function is anisotropic because there's magnetic fields or electric fields, then you won't be able to pick up the full distribution with just one measurement. You'll need lots and lots of different measurements. For example, you'll need another K like this and another K like this. And you can effectively do a tomographic reconstruction in velocity space by taking lots of different slices at different angles here. I will just deal with a few consequences of this, and then I'll take some questions.

So our scattered spectrum-- which, again, we've been writing this scattered spectrum in our differential notation as d2P d omega d nu. But every time you see this, you think, that's what I measure.

I am measuring some power into some solid angle represented by my spectrometer, and I'm measuring it resolved in terms of frequency. This is what I measure here. This measures-- if I can spell "measures"-- the distribution function f of v parallel to that K scattering vector here.

And so, if I have multiple Ks-- I have multiple K-- I get f of v for multiple directions. So, again, I have a plasma like this. I put my scattered light through here. This is my laser beam. It's got some Ki like that.

And then I measure the scattering, for example, in this direction. We'll call that Ks1. Then my K1-- remember, K is equal to Ks minus Ki-- is equal to Ks minus Ki like that. So I measure scattering along this K1. If I measure instead at 90 degrees-- sorry, 180 degrees to that-- this is Ks2. Then I go Ks minus Ki. I end up there.

This will be measuring direction K2. And these two will be measuring different cuts through our velocity distribution. So if it was anisotropic, they could be different. By the way, if these are at 180 degrees, these two should be at 90 degrees. They're not, because I'm bad at drawing that diagram, but that's where it should be. So you could measure orthogonal components of a velocity.

And, of course, if it's fully three dimensional, you might want to scatter out of the page here. And then you would measure some scattering K that's also out of the page. So if you have these multiple lines of sight, you can reconstruct your three-dimensional f of v.

The other cool thing about this-- is you might be able to measure non-Maxwellian distributions, even in one dimension. So if I just write it as f of v without the vector-- so we're thinking about this one-dimensional cut here-- if we look at our spectrum nu, and this is dP d omega-- our scattered light-- if we look at this on a log scale, then this dP d omega is just about log of f.

On a log scale here, our Maxwellian would just look like a parabola. It would go as minus v squared, which is minus nu squared, like that. So if we had a Maxwellian distribution, we'd expect just to see a spectrum like this. But if we have a non-Maxwellian distribution, perhaps we've got scale, some tails coming out of the side here.

And so, in principle, using this incoherent scattering, because we measure the full distribution function, if that distribution is not Maxwellian, we might be able to measure non-Maxwellian tails out here. Now, in general, these tails are weak scattering. There's not very many particles out in the tails. It's hard to make plasmas which are very non-Maxwellian.

So it can be very, very hard to measure this. But if you are able to repeat your experiment lots of times and build up your signal like they did in these low-temperature plasma experiments, you may be able to measure the presence of fast electrons-- non-Maxwellian electrons out here. So this is two very cool uses for incoherent Thomson scattering.

I will pause there and take questions. Maybe I'll put this back down so you can see what we were talking about before. Any questions online? Yeah?

Yeah, so we saw that the scattering cross-section was like-- you expect 10 to the minus 8 of your particles-- photons-- to scatter. So the chance of it getting scattered twice would be 10 to the minus 16, which is very unlikely. So we don't worry about double scattering. I guess I could have put that on my list of assumptions, but yeah. So we don't worry about double scattering.

Just one quick note-- a relativistic correction, because if you-- generally, in tokamaks, they do incoherent scattering because it's very hard to get into the alpha greater than 1 coherent scattering regime. And in tokamaks-- if you've got a nice hot tokamak, you may well have a relativistic plasma. So just a quick note here on what the relativistic correction looks like.

First of all, remember, you may need to think about relativistic beaming. So this is where the radiation pattern of your scattered light beams forward. This is v/c much, much less than 1, and this is maybe v/c less than or around about 1.

So, first of all, you might want to think about this relativistic beaming. But the other thing you might want to think about is, how does this affect your spectrum? Well, there's a long derivation of this, but there's a short version of it, which is that your spectrum-- which, again, we're writing as d2P d capital omega d nu-- the relativistic version looks a lot like your non-relativistic version, which we have written several times over on the side here. But it has a correction factor to it, and that correction is 1 plus 3 omega upon 2 omega i.

This means that the whole spectrum gets shifted up. Remember that omega here is equal to K dot v. If you're wondering why this is relativistic, that's where the v is coming in. So this is a correction on the order of-- and this omega i, by the way, is just cK. And so this correction here is on the order of v/c. These are the ones we said we'd keep for our relativistic correction.

What this correction means is that we get a blueshift. So the spectrum is always blueshifted-- always ends up as being at a higher frequency. You can see that by the shape of this and the fact that it depends on the sign of omega here. And, in particular, this correction is for temperatures-- electron temperatures-- on the order of the keV range. So if you're working below that, you really don't need this correction.

But if you're working in a 10 keV plasma, you do need this correction. Because of this blueshift, it's important to measure both the positive shifted frequencies and the negative shifted frequencies. Because if you don't measure those, then you're expecting your spectrum for a Maxwellian to be nice and symmetric. But if you are doing the relativistic one, then your spectrum might look like this.

And if you only measure one half of your spectrum-- this is frequency-- we only measure the positive frequencies or negative frequencies-- we might misinterpret what's going on here. If you're starting to work with relativistic plasmas, you may be tempted with your spectrometer, because you only have a finite number of frequencies you can measure at, to only measure positive or negative frequencies because the spectrum is symmetric.

But in the relativistic case, it's not. This is like d2P. OK. And this is Hutchinson's equation 7.2.28. And he does the derivation of this properly. I am just giving you the result. Any questions on that? Just like a side note. We have just enough time to introduce coherent scattering, but not enough time to actually do all the fun mathematics of it, so we will leave that for next week.

So that was incoherent scattering. We derived the scattering off a single particle. Then we integrated over the distribution function and said that the scattering off each of these particles does not interfere with the scattering off any of the other particles. Or, if it does interfere, another particle's scattering will add up, and that sum over the scattering between different particles will eventually cancel out. That is not going to be the case in coherent scattering.

So, again, we're just going to sketch what's going on. Before we go into the mathematics, I want you to have a heuristic understanding of what's happening first so that you have some faith that, when we wade through the math, it'll all be worth it in the end. So this is just a simple sketch for you.

Again, we're talking about scattering off fluctuations which have a wavelength lambda, and that lambda is greater than the Debye length here. So these must be coherent fluctuations because we can no longer see individual particles. We can only see particles that are Debye shielding other particles.

And so these particles must be moving in concert with the other particles. And just a note-- as we said before, this is the condition that alpha is greater than 1. The coherent fluctuations are modes or waves. I'll use both words.

So let's consider what sort of modes and waves there are in a plasma that we can scatter off. So we consider modes with frequency omega, which is equal to, as we keep saying, omega s minus omega i. We're dealing with the difference between these two frequencies.

And, again, we're looking at Doppler shifts which are on the order of K dot v. And so this is going to be a relatively small quantity. So this omega s minus omega i-- our omega is much going to be much, much less than omega i, which is basically the same as omega s.

So we're looking for modes with a relatively low frequency-- much lower than the sorts of waves that we see propagating in free space, like our laser. But, despite the fact that these modes are low frequency, they may not actually have a low wave number.

Because, remember, our scattering diagram-- we have some incoming wave, we measure it in some scattering direction Ks like this, and we said, because this is elastic scattering, that the size of Ks is roughly the same as the size of Ki. They didn't have to be pointing in the same direction, though. And so that means that K, which is equal to Ks minus Ki-- again, we just look at this and do this as a vector Ks minus Ki like this.

And so the size of K is actually on the order of the size of Ks or the size of Ki. So, although the frequency of the mode is much lower than the frequency of a free space propagating plane wave, the momentum, or the wave number of the mode, is roughly the same here.

AUDIENCE: [INAUDIBLE]

JACK HARE: Nothing I've done here actually is talking about the Debye length yet.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yes, though, again, none of this mathematics so far has said anything about that. We're just-- what all of this means together is-- we're looking for a mode, which has v phase which is equal to omega/K. We know that this omega is much, much less than omega i, which is equal to CKi.

But we know that this K is on the order of Ki. So if we take the ratio of omega and K, we can see that this will be much, much less than the speed of light, which is-- actually, this would be the same if I replaced this with the electron velocity here, if I was doing incoherent scattering of individual electrons.

But this same condition now holds here. So what we need to do is we need to identify-- find modes with-- so what modes are there in a plasma that satisfy this? They've got a low phase velocity compared to the speed of light. They have relatively low frequencies, but they have high momentum. We're not dealing with any magnetic fields here. So we're not going to scatter off alpha [INAUDIBLE].

Sound waves? What is-- and can you give me a posh name for a sound wave? An acoustic wave. There we go. So what is the dispersion relationship for an ion acoustic wave? I'm just going to take the square root there. And what is cs?

AUDIENCE: [INAUDIBLE]

JACK HARE: Right. Well, we're working in a plasma that can be multiply ionized. So I'm going to put a z in there-- z in there-- and it turns out, in reality, it's z t e over t i over M i. Is this a wave which has low frequency and high momentum, which is the same as saying, does it have a phase velocity much less than the speed of light? What's its phase velocity? This isn't a trick question.

AUDIENCE: [INAUDIBLE]

JACK HARE: Why? OK. [CHUCKLES] So the phase velocity of it is just the sound speed. And that sound speed is going to be much less than the speed of light. And, yes, maybe it is because ions are slow. So you'll need, effectively-- this is a relatively heavy mass. You calculate this, and this number is going to be like 10 kilometers a second or something like that. It's going to be much, much less than 3 times 10 to the 8.

So these are great modes. We can definitely scatter off these modes. These are actually very low frequency modes. These are the lowest frequency modes in a plasma. What's a slightly faster wave? Yes?

You're cheating. You've done this before, OK. [CHUCKLES] Go on-- Langmuir waves. I'm going to call these electron plasma waves, but they are sometimes named after Langmuir. So what is the dispersion relationship for a Langmuir wave?

And we could go and work out the phase velocity for these. It's a little bit more difficult. But we'll find out that these do have a low phase velocity, but it's not quite as low as the ion acoustic waves. So these are just the low frequency waves. So if we're scattering off the ion acoustic waves, we expect to see small frequency shifts. If we're scattering off the electron plasma waves, we expect to see slightly larger frequency shifts.