

[SQUEAKING]

[RUSTLING]

[CLICKING]

JACK HARE: So once again, we are going to be discussing current scattering. And we're going to derive Thomson scattering for the third time. Huzzah! Huzzah. Thank you. Good. Well done. A grade.

Coherent scattering-- so just as a quick review of what we discussed at the end of the last lecture, we are looking at scattering when the wavelength of whatever the mode is that we're scattering of is greater than the Debye length here. So this wavelength is 2π upon the size of k , where k is k_s minus k_i , like this.

So we're not talking about the wavelength of our laser. Though, this wavelength may be similar to the wavelength of our laser, but we're talking specifically about the wavelength that we get from drawing this diagram where we have some laser beam coming in with a wave vector k_i . We have some scatterer at a vector k_s . And then we go, OK, k_s minus k_i . Aha, this is the k that we're talking about. We're talking about the wavelength of this.

And the reason we're interested in things larger than the Debye length is on scales larger than the Debye length, the wavelength of this scattering vector is going to be sampling the collective motion of the plasma, or the coherent motion of plasma. So you might see this called collective Thomson scattering, as well as coherent Thomson scattering.

So modes which exist in the plasma that fulfill this condition, they have a frequency, which is much, much less than our laser frequency in general. And this is a consequence of the fact that we're doing relatively low energy scattering. So our photon energy is much, much less than the kinetic energy of the particles in our plasma.

But the size of the k vector may well be on the same order as we see in this sort of vector diagram as the k vector for our free space mode. And so we still have modes which carry a lot of momentum, even if they carry relatively little energy. And these two facts together mean that our phase velocity, which is defined as ω upon k , is much, much less than the speed of light.

And the modes that we identified that fulfill this were the ion acoustic waves. And these have a dispersion relationship, ω equals square root of z_{te} plus t_i over m_i -- Could you mute your microphone, please? Times k . And we also have the electron plasma waves. And these have a dispersion relationship of ω squared equals ω_p squared plus $3b_{te} k$ squared.

So these are the sorts of modes we're going to be scattering off. This is a very low frequency mode. And this is merely a low frequency mode. So when I say scatter off, it would be reasonable to ask why am I so sure that these modes exist in my plasma. Is there something in my plasma that is launching ion acoustic waves or electron plasma waves? And the answer is my plasma, just by virtue of sitting here, is a bath of these different modes, these fluctuations going on all the time. We have in every single direction ion acoustic waves flying off from every part of the plasma to every part of the plasma.

So if you try and fire your laser beam through here, you are certain within the region of plasma you're interested in to find a wave going in the direction you want here. And you will get scattering. And we will be scattering off, as I've drawn several times now-- we'll be scattering off ion acoustic waves going in this direction at the sound speed here, but we will be observing them over here with our spectrometer, looking along k_s .

So I think this is one of the things that confuses people when they first start learning about collective Thomson scattering. They're wondering, why are these waves here to start with? The fact is there is always a possibility of there being a wave.

And as you scatter off it, you transfer more energy and momentum to that wave. So you actually transfer energy from your probe laser to the wave. This is sort of like a Landau damping process. And you will grow that mode. And so that mode will definitely be there for you to scatter off.

And the intensity of the scattering-- so again, we can write this intensity as the amount of power scattered into a solid angle scattered into a certain frequency will be proportional to the fluctuation strength for the number of modes-- fluctuation strength. Or if you want to think of this in a quasiparticle picture, you can think of it as like the number of phonons, or something like that, which are available.

So the spectrum that we expect to get if we do all the mathematics correctly-- this is ω . This is what our spectrometer measures. If we expect to see some peaks at very low frequencies, there's 0 in the middle here, corresponding to the ion acoustic waves. And we'd also expect to see some peaks at slightly higher frequencies, corresponding to the electron plasma wave, the EPWs.

And just remember I'm using this ω here, which is the difference between the scattered light and the incident light. So if I didn't have any plasma and I scattered off a lump of metal, I would just have a very big spectra right here in the center, ω equals 0. When there's plasma, I scatter off these modes inside the plasma.

And there'll be two possibilities. I can scatter off a mode, which is going this way, or I can scatter off a mode that's going this way. And they both fulfill that condition that $\omega = \omega_s - \omega_i$ and $k = k_s - k_i$. So I'll see both of these peaks here.

Now, in reality, they're not going to show up as little delta functions, even though if you look at the dispersion relationship, it looks like there's only one solution for a given ω and a given k . In reality, they'll be thermal effects, which broaden these. So what we will see is some sort of peak like this. It's good, because we don't like singularities. They don't really occur in nature.

This heuristically is what we expect to get. But now we have to go out and do the mathematics and actually go and get it all. So any questions before we launch into that? Yes.

AUDIENCE: [INAUDIBLE].

JACK HARE: Scattering gets chosen by where we put our detector, our spectrometer. So if you would like to think of it as a fiber-coupled spectrometer, I would have a little fiber optic bundle here. And that would go to my spectrometer, which disperses the light in the frequency space like this. There'll be some dispersive element.

And so for example, if I have a lens here, I would collect light from this region here. And this spectrometer, therefore, defines the k vector, k_s . Now, you bring up a good point, because for any realistic system, you actually collect a range of k_s . You can't just collect along a single vector, because different angles of light will be collected by your lens. And that will also contribute to this broadening as well.

But fundamentally, you can make that very, very small. And you still wouldn't get these delta functions. You're still going to get some broadening due to thermal effects.

AUDIENCE: [INAUDIBLE]

JACK HARE: Well, actually, you don't really choose the magnitude of k . You only choose the direction of k . The magnitude is set by the dispersion relationship and the frequency relationship here. So when I see a wave with a certain frequency, I can use the dispersion relationship to work out what k was going to be at.

So really, this is just setting the angle. So it's setting the vector relationship here. It's the size of k . Remember, we had this nice expression for the size of k , which was the size of k is equal to $2k_i \sin \theta$ upon 2 . And we got that from saying that we're doing elastic scattering. So the size of k_i is roughly the size of k_s .

So what you do is once you've chosen your angle and you know what your incident wavelength is, you have effectively chosen the size of this k vector straight away. And then when you choose the angle, you've chosen the other part of the vector, which is its angle. So this is the magnitude. And this equation determines θ and this is the magnitude.

And then you've set everything about k . And so if you look at a certain frequency and you see a wave show up at that frequency, you know what your k is. And therefore, from the frequency you measure, you can work out what this term has to be.

So for these peaks here, I can work out that this frequency shift is going to be proportional to square root of z times t_i -- I don't know why I made a temperature vector there-- t_i upon m_i . Other questions? Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: At this point, this is like, yes, I guess this is for the k in the center of our lens here. And then one side, we're going to have k 's, which are a slightly larger angle, and the other side, we'll have k 's at a slightly smaller angle. And you can see from this that corresponds to slightly bigger and smaller magnitudes of k , which in turn will, for the same temperature in the plasma, correspond to slightly bigger and smaller frequencies.

So this corresponds, for example, to θ_0 . And on this side, this is θ_0 minus $\Delta\theta$. And this is θ_0 plus $\Delta\theta$, where we're talking about this being the angle θ . And then this being the angle $\Delta\theta$ between the sort of main angle of your thing, and the spread of angles there.

AUDIENCE: [INAUDIBLE]

JACK HARE: No. It's the multiple effects that are going on. But the angle is one of them that would also stop you getting singularities as well. And you do need to think about it, because a good trick-- we talked about how few photons we have. So it makes sense to have a bigger solid angle to collect more scattered light.

The trouble is that each of the photons are coming at a different angle. They actually have a different spectrum. So it'll also broaden your spectrum, which will make it hard to measure things. Other questions? Anyone online who has a question? Let's do some maths.

This is a two-step process. Our first step is calculate the scattered power spectrum, which, again, we're writing as $d^2 p / d\omega d\nu$. And we want to know how this relates to density fluctuations. So fluctuations in the density, δn_e of r and t . Then we're going to quickly Fourier transform these and actually be looking at them in terms of a spectrum of density fluctuations with certain k vector and ω , because everything is going to be Fourier transformed here.

So we relate this to this quantity, a relatively simple formula. And then the more complicated bit is calculate what these density fluctuations are, δn_e -- well, we're just going to go call them any k and ω in a plasma.

So this goes back to my previous point that there are always these fluctuations in a plasma. So we just have to work out the size of these fluctuations. And we're going to do that using a test particle treatment and a little bit of Landau damping along the way.

And one thing I want to say is that this treatment here is going to reproduce the incoherent case as a limit. So our treatment now-- sort of derivation includes what we've done previously. It just includes it as a limit. And I'll try and point out where in our derivation we've exceeded what we did previously. But you will be able to get back all of the Thomson scattering spectrum, not just the collective part, from the equations at the end here.

So we want to recall a couple of things. We had our Fourier transformed scattered electric field. So it's in terms of frequency, which was equal to the classical electron radius over the distance to the scattering volume. And we were doing a Fourier transform on this. And we had this $\pi \cdot E_i / 0$, which was just the thing that gave us the shape of our scattering. So it's got some angular dependence.

Actually, I can just write this as E_i in terms of r and t prime. And we were Fourier transforming it. So we had an exponential of minus $i\omega s$ and e . And we replaced our t with the t prime, because that's what the electric field was. And when we replaced this t with a t prime, this becomes t prime. And we get a second term that was minus $k \cdot r$ at here. And now our integration is respect to d prime instead of t . So this was from a single particle what our scattered electric field was going to be.

So now we want to know the scattering from a distribution of particles. And we're going to take as our distribution function the Klimontovich distribution. So I'm going to use a capital F here. And this is a distribution function in space and velocity and time. And in some sense, it's a very dumb distribution function, because we just say it's the sum over the position of every particle, which we write as a delta function, $r - r_j$ of t , and another delta function, which is the velocity of every particle.

So sort of like an obviously true statement, but the Klimontovich distribution is not usually very useful to work with. But we'll get back to why we do it later on. You may have seen this in particle kinetic theory.

And we're going to say what we want now is the scattered electric field. We want the scattered electric field from every particle in this Klimontovich distribution. So we're going to call this scattered electric field E_s total. And we're going to do an integral over all of the plasma, all of the particle velocities, our Klimontovich distribution, and our electric field from a single particle here. So this is the total field.

The key difference in what we did in the incoherent case, I believe, is that now in the incoherent case, we kind of handwaved our way over this integral over space. And I think we effectively assume that all particles are in the incoherent case. All the particles j were at $r = 0$. And so now we're going to do the coherent case, and we're not going to make that assumption. So $r = r_j$. That's what the delta function up here does.

And that means that if we have our spectrometer, our observer over here, and our origin here, and we have two different particles at positions R_1 and R_2 , the electric fields coming from these, because they're scattering waves in this direction, they may be a phase difference. And so we may get interference from them.

Previously, because we assumed that all the particles were in the same place, we couldn't get a path difference. We couldn't get any interference. And so we just summed up their contribution.

Now, when we do this integral, we're going to very carefully take into account the difference in position of each of the particles inside Fe. And so when all the waves gather up together, we're going to be interested in how they interfere with each other. So this is the key difference in this derivation.

A couple of things to note, of course, is that our electron density, n_e is simply the integral of our distribution function over all velocities. That's still true. Just to be clear, this is n_e as a position of space and time. So we've lost the velocity component. We've integrated over it.

And so-- do we have space to write this? No. This total scattered electric field, where has it gone? That one up there. It's going to be equal to r_e on R . There's going to be this boring phase factor out front, $i k \cdot R$, which will vanish very soon. This is the phase that all of the waves pick up by virtue-- this is kind of like the average phase. It's the phase that they pick up from going from this point to here.

So it's going to be a very boring phase that will cancel out. But we're still interested in the additional phase from the exact locations of each particle. There's going to be a factor of this π dotted into the polarization. Again, it's just a shape function. But now, we're going to be an integral over π and an integral over volume of the density of R and T prime. So I've carried out the velocity integral and replaced this with the density.

And now I have the factor from the electric field. I didn't write it down again, did I? OK. I shall put it here. Our incident electric field is equal to $E_i = 0$. I'm going to do it this way. $\cosine \exp(i k \cdot r - \omega t)$ prime. This is the incident electric field.

So if I now substitute in the electric field, I'm going to get that exponential term here. But that's going to be exponential of $i a \cdot r - \omega t$ prime. And then I'm also going to have this exponential here. I'm just going to write directly below, exponential of $-i \omega s t$ prime minus $k s \cdot r$.

So far this looks a great deal like what we did before. But if you want to evaluate this electric field carefully, and then work out the scattered power very carefully, you have to pay an awful lot of attention to exactly-- sorry, I've forgotten to put in the limits here of this. $d t \text{ prime } d^3 r$. So integrating over time and space.

To do this properly so that we can take into account the interference requires you to very carefully go through this and pick out some important terms. And I had it in my notes from last year. And it was not a very educational exercise to do on the board.

So what I'm going to say is if you are interested in seeing the mathematics behind this-- this is Hutchinson's 7.3.8 to 7.3.12. And I'm simply going to quote the result now instead, which is that the total scattered power into some solid angle, into some frequency here is equal to the classical electron radius squared, the power of the laser over the cross-sectional area of the laser, this shape function $\pi \cdot \hat{i}^2$ times n_e times the volume that we're scattering from trying to find all of this times something called s of k ω .

This is the important quantity called the spectral density function. And it's important because all of the information about this spectrum is held inside there. This is where the k and ω dependence are that tell us what the spectrum looks like. These are all just scaling constants that tell us the intensity. And often, because we don't have absolutely calibrated spectrometers, we can't really measure this. So we don't really care. But we can measure the spectrum, the relative intensity. And that's all hidden within this function.

And this function is equal to 1 upon the time of our integration, the volume of our integration, and then the average of the Fourier transformed electron density squared over the average electron density. So these are the fluctuating term, and this is the sort of steady state n_e electron. I should really make that square very clear. It's very important. Question? Yeah.

AUDIENCE: So why are we the r instead of [INAUDIBLE]?

JACK HARE: Yeah. This is a generic Klimontovich distribution function. We'll be using it. Yeah, you're right. When it goes into here, I did put t prime in a Klimontovich distribution. So you're right. That's just the Klimontovich distribution. But you're right. We're evaluating a Klimontovich distribution at time t prime. So this should be t prime here. And this should be function of r d and t prime here.

I haven't generally been using primes on the r 's. Though, Hutchinson does to remind us that it's r evaluated at t prime. So he writes r prime is equal to r at t prime. I tend to drop the prime because I remember. But maybe I shouldn't drop it. Other people don't use that notation. So just if you note, if you go into Hutchinson.

In the syllabus, I recommended a book by Froula and Sheffield called *Plasma Scattering of Electromagnetic Radiation*, which is very good and very thorough. And I've kind of blended their treatment with Hutchinson's treatments to get some of this. So not all of this looks exactly like how Hutchinson does it, but I do think that Hutchinson's treatment is pretty good for most of this stuff. Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: This one?

AUDIENCE: [INAUDIBLE]

JACK HARE: So it comes in as part of this term. So let's be clear. So this, it actually should be back here as well. I have missed it out. So there's a times e to the i as dot r factor in there as well. And where it comes from is originally this is exponential of minus i ω s t dt .

And when we substitute out t equals t prime with a load of other bits, one of the other bits becomes k s dot r , and the other bit comes k s dot capital R . And we can take that outside. But yeah, I forgot to write it down here. It should have been in here as well.

And what that represents, that relatively boring term, is just the fact that all of the electromagnetic fields from all of the plasma have to travel some distance r to our detector. But some of them travel r . So if this is the r vector here, some of them travel r plus δR or minus δR , or whatever else like that. So we're kind of accounting for the overall phase that they all pick up as the waves travel. And then we're trying to calculate these very small differences. And those are the ones which interfere and give us the interesting coherent effects. Other questions? Yes.

AUDIENCE: [INAUDIBLE]

JACK HARE: That's the relativistic treatment, the relativistic treatment where we keep terms in β upon c . And we've dropped those. So this is the non-relativistic treatment. Good point. Any other questions?

So we've now done step one. We have calculated, to some extent, the power spectrum in terms of density fluctuations. So now the question is, step two, what are these density fluctuations?

What we're going to do is we're going to consider the response of the plasma to a test particle. So we imagine we have some plasma. And we put some test particle through it. And then we asked ourselves, what does the rest of the plasma do in response to this test particle?

The test particle, in this case, what we're going to do is we're going to pick each electron in the plasma and each ion in the plasma, work out what the rest of the plasma does in response to each electron and each ion moving. And then we'll sum all of those responses together.

So at the moment, this could be an electron or an ion. Remember, we're not interested in the scattering of the ions. But if the ion motion perturbs the cloud of electrons, and then we can scatter off that cloud of electrons, we will see that light here. So this cloud here is a sort of Debye cloud here. These are the particles which are shielding our test particle.

If our test particle is an electron, then the shielding particles are an absence of electrons. The ions can't move out the way fast enough as an electron goes through, but the electrons will move out of the way. And you can imagine that as this particle goes through, and the electrons move out of the way, it will cause a perturbation to the density. It'll be a negative perturbation. So it'll actually scatter less.

But that means that these electrons have to end up somewhere else. The density will be higher. They'll be scattering off those. If the test particle is an ion, then it turns out that half of the particles, or half of the shielding that's made up that shields this ion charge are due to electrons being attracted. And the other half of the shield is by ions being propelled.

So it's slightly different from the argument for the electrons, where only electrons can shield electrons from ions. Some of the shielding comes from drawing electrons in. The rest of the shielding comes from pushing ions out. And that can happen because the ions can move as fast as an ion. And so we can see that shielding here.

So what we get is a perturbed charge density, ρ_0 , at some position r and t , due to our test particle, which is $q \delta(r - vt)$. We're just saying our test particle perturbed density looks like this. As it moves through, it's at some position r , moving at some velocity v in a straight line. And that means that its Fourier transformed contribution to the charge density is equal to $q^2 \pi \delta(a \cdot b - \omega)$. So this is the Fourier transform.

Now, we're going to take some equations from electromagnetism and talk about this in terms of a polarization and a polarization charge density. So it'll be free and bound charges. So we have a polarization vector, which is equal to the susceptibility times the permittivity of free space times the electric field. And this is also equal to π upon ϵ , which is just the permittivity, and the d field.

And the e and d fields are defined as divergence of the electric field is equal to ρ_0 . And divergence of the field is equal to ρ_e . So this is the free charge. And these are the bound charges. And in our case, our bound charges are our shielded charges. I ran out of space to write shielding. There we go. You get the idea.

And we can solve these equations to work out what the size of our shielding charge is. And our shielding charge, ρ_e , is equal to minus χ_e upon ϵ times by this Fourier transformed test particle charge $q^2 \pi \delta^3(\mathbf{r})$ minus ω .

What this is effectively saying is as this test charge moves through or as it exists as a mode in Fourier space, it induces some bound charge to shield it. And that bound charge has a charge density, which is to do with the susceptibility of the permittivity here.

Then we can-- this board-- say that our electron density, n_e of $\mathbf{k}$ and ω , is going to be equal to a density, which is due to the electrons, and the density, which is due to the ions, which have been perturbed in this system. So for the electrons, we sum over j electrons.

And we have a 1. This is the test particle charge. So whenever we have an electron test particle, it's going to contribute 1 to the electron density here. And we're going to subtract off it the number of electrons which were repelled by that test particle, π upon ϵ . And both of these are $2 \pi \delta^3(\mathbf{k})$ for the electrons minus ω .

So this is the electrons. This is our test particle electron. And these are our shielding electrons. And then we also have a term that's for the ions. We sum over l ions. And here, we simply have a term π upon ϵ here. The ions have a charge of z $2 \pi \delta^3(\mathbf{k})$ velocity of the ions, v_l , minus ω . This is the ion contribution here.

Just to be very clear, this is the contribution of the electrons shielding the ions. We're only talking about the electron density here because we're only interested in electron density. So this is the contribution of the electron test particle and the electron shielding the electron. And this is the contribution of the ion test particle. So we don't get a 1, because we don't add.

The ion density of the test particle doesn't go into the electron density overall. So we merely have the response of the shielding electrons, these ones, which are following the ions around here. Again, we're only interested in ion. So this, in case you're wondering where we got to in Hutchinson's book, is his equation 7.3.19.

And we said that our spectral density function, s of $\mathbf{k}$ and ω , was equal to 1 upon the average electron density, the time and the volume that we're integrating over, and then the ensemble average of $\mathbf{k}$ and ω squared like that. We can plug this in. We note that we've got two terms.

So we're going to have a cross product, but the cross between these two terms is going to be a correlation between the electron and the ions. And there is no correlation. So those cross terms vanish. So we simply get this term squared and this term squared. Let me write some of this down.

So we would get the electron contribution squared plus 2 ion times electron plus ion contribution squared. This cross term goes to 0, because we have uncorrelated electrons and ions.

And we also use our dodgy identity from last week that the square of a delta function in frequency space is equal to π upon 2π of that delta function. So when we take the square of this term and this term, and we have a square of a delta function, we're just going to replace it with itself, with a π over 2π in front here.

And so all of this gives us 2π upon the electron density times the volume summed upon j electrons $1 - \pi e$ upon $\epsilon^2 \delta(k \cdot b_j - \omega)$. So this is the electron term, summing over j electrons, plus a term that looks like $\sum$ upon the ions, πe upon $\epsilon^2 \delta(k \cdot v_l - \omega)$, the ions minus ω .

And I need to put some big ensemble average brackets around this. And I actually want to go back and put in some absolute value signs inside here, which is why we end up with these absolute value signs. We're about to convert this into something that makes more sense, but any questions at this point?

The size of the density fluctuations in Fourier space, these can be complex quantities, because they've got a phase attached to them. So this is any times any star. It's the same as that. So we're taking the-- the power that we measure depends directly on s of k and ω . The power we measure has to be a real quantity. We can't measure an imaginary power. And so therefore, this needs to be a real quantity as well. So we're just making sure we've only got the magnitude of it, not its phase. Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: What we're really doing is the cross terms would involve delta functions of ions and electrons times by each other. So if you can imagine, it's very unlikely that an ion electron is in the same place at the same time. These are delta functions that are very, very sharp. So really, what we're saying is not so much they're uncorrelated, but they don't occupy the same space. So they're different locations.

So if I did a little diagram of some x -coordinate here, I might have an ion here and I might have an electron here. But those delta functions will not overlap in general. Any other questions? Yes, we'll do this.

This still does not look very useful because we're summing over all like 10 to the 23 or whatever electrons in our system. So this is a difficult sum to do. But you must have spotted this and thought, OK, if I'm doing a sum over a load of particles $\delta(k \cdot d - \omega)$, this is looking an awful lot like my Klimontovich distribution again. And in fact, this looks like a sort of integral over space and over velocity of my Klimontovich distribution, r and v and t , times by some delta function $k \cdot v - \omega$.

Remember, this consists of a load of delta functions-- $\delta(r - r_i) \delta(v - v_i)$. And so if I integrate this up over this delta function, this is going to pick out all the particles. So I can replace this by my Klimontovich distribution. And I'm going to end up with a term that looks like the volume from our d^3v integral.

There's a $1/k$ that I'm going to swap out $\omega = kv$. I'm going to swap out my v integral, the v for-- when I do that, I'm going to pick up a $1/k$ factor. I'm going to have an average over my Klimontovich distribution function, ω upon k .

And here, this is like an ensemble average of lots of particles. The Klimontovich distribution function is very spiky. And we don't really want to work with it. So I wave my hands and say, this is $V_{1/k}$, some nice, smooth distribution function that you can actually work with.

So I'm not going to claim that's a particularly rigorous argument. But you can have a look in Froula's book and in Hutchinson's book if you want to see maybe a slightly better argument for this.

So we're going to basically replace these sums over a large number of particles-- this should have been v_i here-- with the distribution function, which is kind of, obviously, what it is. If you have a sum over a load of particles, it's going to be some distribution function. But this distribution function is now being evaluated in terms of ω and k . So I'll write that down, and then we can talk about what it is. Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: Well, our test particle treatment we talk about a single particle moving through. And that's what starts introducing all of these delta functions. And so I think that that is why we do it, because when we've done the test particle treatment, we end up with all of these, and it's easier to get rid of them by going through a Klimontovich. So I think we started with a Klimontovich, knowing that this would happen later. Then we take out the Klimontovich later on.

There may be other ways to do it. And it's not like I'm claiming that what I've done on the board is particularly rigorous. So there may be an equally rigorous way to do it without the Klimontovich distribution function. But I think if you want to do it rigorously, you have to put this.

So if we do this substitution, we then end up with the very, very key formula for the spectral density function in terms of things that we actually stand a chance of evaluating. So s , k , and ω , 2π upon k times the average electron density. Then we have the brackets, the first term, 1 minus the susceptibility of the electrons over the permittivity of the entire system squared.

Now, the distribution function of the electrons in the direction of k , so this is our incoherent scattering thing again, where we're only taking a slice of the distribution function in one direction. So a one-dimensional distribution function here. Plus a term that looks like susceptibility of electrons over the permittivity of the entire system squared.

And then, formally here, we have to sum over every single ion species. So now we're doing a sum over different ion species. So if you've got carbon and hydrogen, things like that, in a quasi neutral plasma, you need to sum over the charge of each of these ion species and the distribution function for each of these ion species.

But again, if everything is Maxwellian, this becomes much simpler. But this is the fully general formula here. And if you didn't want to write this down, this is Hutchinson's 7.3.22, or you can check in case I've made a typo here.

So I'm going to talk through what each of these terms physically means, physically motivated. So this here, the distribution function at ω k is the number of electrons with a velocity v equals ω upon k . So the number of electrons in this distribution whose velocity is equal to the phase velocity of the mode that we're scattering from. Remember, we've been setting ω and k all this time.

So this tells you how many scatterers there are which are electrons. And this tells you the response of the plasma to each of these electrons. So each of these-- there may be more than one. Each of these electrons are flying around as a test particle. And the plasma is going whoa and moving out of the way in order to accommodate them.

So this is the response of the plasma. We've got a contribution from those electrons. That's the one. And then the minus is all the electrons getting out the way of those electrons. So that's the total amount of electrons you have to scatter off. We have the density of those electrons, but we also have a depletion of density of other electrons getting out of their way.

This over here is the number of ions with a velocity ω upon k . And this is that response again. But note that we don't have a 1, because we don't scatter off ions. Their scattering is tiny. So all we get to see when the ion is moving-- the scattering we get is just the scattering of the electrons, which are pulled into the ion.

Note there's a difference in sign here. Well, I guess-- yeah, it doesn't matter. When we square it, we get rid of it anyway. But these are the ions which are drawn into the region where the-- sorry-- these are the electrons which are drawn into the region where the ion is.

And so we're scattering off those. So we only look at the response of electrons to these ions here. This is what your total scattering spectral density function looks like. Where does this function have peaks? Where am I looking in this for that spectrum what I drew of r here?

Where do we get peaks in s of k and ω ? There will be a singularity, shall I say? I'll try and make it very clear. Where will I get singularities in s of k and ω , which will look like nice, sharp, defined resonances scattering off specific waves?

We will get peaks, singularities for ϵ equals 0 in the denominator. So whenever ϵ tends to 0, this function will be extremely large. ϵ is also-- by the way, this is χ of k and ω . This is also ϵ of k and ω . So there's k 's and ω 's hiding everywhere inside here, not just within these fusion functions.

These depend on the speed of the test particle as well. And so if the test particles are going at a certain speed, they will induce a certain response to plasma. That's a wave. So this is equivalent to saying, ϵ equals 0 is equivalent to finding the normal modes of the plasma, just like you do with determinants of matrices and stuff like that.

So are we done? Can you plot me a Thomson scattering spectra from this? You can? You were nodding. We don't know anything about χ and ϵ . So now we're going to have to go do some [INAUDIBLE]. Don't worry. We're almost there.

But in principle, if you did know these for some reason and you did know your distribution functions, you could work out what the scattering spectra was. But we need to know what these χ 's are. And just-- yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: So again, still, we are never scattering off the ions. But people will talk about scattering off the ion feature. So these are actually-- these two terms tend to be large at different frequencies. This term is large, around the ion acoustic waves. This term is large around the electron plasma waves. And so people call this the ion feature. But just, it's important to remember that no ions were scattered off in the making of this spectrum.

Any other questions on this before we keep going? And I don't think I said it explicitly anywhere before. But your permittivity is always equal to 1 plus all of the susceptibilities in your system. This is just true for some electromagnetic system. It could be dielectrics and stuff like that. So in this case, our permittivity is equal to 1 plus the electron susceptibility plus the ion susceptibility.

And technically, because I've written this in terms of a load of ion species, it's the sum over all of the ion susceptibilities. So what we're heading out to find is the electron susceptibility and the ion susceptibility. And that will give us epsilon straight away. And we'll also find that when epsilon equals 0. Remember, epsilon is some function of k and ω .

So this will be some function of k and ω equals 0. And all of a sudden, out will pop ω equals c/k like that. So that will be a solution to epsilon equals 0. And it'll be the dispersion relationship for the ion acoustic waves, or the electron plasma waves, or whatever wave you've got in your system. Questions?

Run out of tea. Oh, no. There's a bit left.

Well-- yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: Please.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yes.

AUDIENCE: So when we get a measurement, we get a direct [INAUDIBLE].

JACK HARE: Yeah. We effectively do. Do I have it anywhere? What we measure is this quantity. We measure watts into some solid angle into-- at some frequency range here. That depends on the cross-section, the intensity of our laser, how well we focused it, if we put our spectrometer in the right place to actually see the scattered light, how many scatterers there are, how big the volume is I'm scattering off, and s .

This is the only thing that has a shape. This is just a number. All of these together for a spectrometer in a certain place, this is like two. Whereas this s is actually this function that looks like it has peaks at different frequencies and stuff like that. So that's ω . This is s , like that. So this is the only interesting thing.

If we had absolutely calibrated our spectrometer so that we could calculate this value, we can, in principle, get the density of our plasma as well. It turns out the absolute calibration of a spectrometer is very hard. So very few people use it to get this density. It also turns out that the electron plasma waves carry information about the density. So you're usually better off just measuring the frequency shift of these.

So this is like the spectrum, what we used to think [INAUDIBLE] all of our lines and our k-shell, and all that sort of stuff. And this is just some constant factor that we probably won't be able to do anything with. So it does tell us that if we want to get more light onto our spectrometer, we should use a more powerful laser. We should focus it more closely. We should put our spectrometer in the right place. We should make our plasma denser. And we should collect the light from a larger region.

So it's sort of useful information to know how to do Thomson scattering better. But this is the sort of thing that you actually spend your time fitting to. So you alter the temperature and density of your model plasma until your model σ matches the data that you've recorded. Yeah. Cool. But good question, yeah. Thank you.

AUDIENCE: [INAUDIBLE]

JACK HARE: We have not assumed it's Maxwellian yet.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yes. If we did assume it was Maxwellian, then we could get a temperature. We'll go through it. We will derive-- I can't even remember which board we're on now. We would derive these susceptibilities from an arbitrary distribution function. And then we would discuss it in the case of a Maxwellian distribution function, which is a pretty common case.

Any other questions? We're getting there. What shall we erase?

The refractive index of a medium is simply equal to its permittivity to the half. So if we were dealing-- and for example, for a plasma, a cold plasma, where the thermal velocity was much less than the phase velocity, we derived, for example, [INAUDIBLE] mode, this was something like $\omega_p^2 / (\omega^2 - \omega_p^2)$. So I'll write that as n^2 .

So if we were dealing with a plasma like this, it would be very easy to find the permittivity. But we're not. Remember, we're dealing with plasmas where the dispersion relationships we're interested in look like $\omega^2 = \omega_p^2 + k^2 v_{th}^2$. So that means that our phase velocity now is on the order of the thermal velocity.

So we can no longer do a cold plasma treatment. We need to do a warm plasma. And because we're doing a warm plasma, that means we're going to have to think about Landau [INAUDIBLE].

This is because our distribution function, something like this, previously, for this case, for the cold plasma, our wave was all the way up here-- v_{ph} much, much larger than v_{th} . There were no particles which had a velocity close to the phase velocity. And so therefore, from the point of view of this wave, the plasma was just cold and stationary.

What we're doing now is we're looking at the phase on the order of v_{th} . And so we have to consider how the wave interacts with and damps on these particles. So you will recall that we start with the Vlasov equation, where we're looking at the time derivative of some distribution function, plus the velocity dotted into the spatial derivative of a distribution function, plus the force-- the acceleration from the electric field, $q \cdot \nabla f$, the change in the distribution function in velocity space. And we're going to set this equal to 0. So this is explicitly a collisionless system.

Now, it turns out you can derive Thomson scattering for collisional systems, but it also turns out that most Thomson scattering systems are collisionless. I'm going to spend 30 seconds on this, and I hope I don't confuse you too much. You may be doing Thomson scattering in a system, which is collisional. A collisional system is one in which the length scale is larger than the mean free path.

But when you're doing Thomson scattering, the length scale we're interested in is λ . Remember, this λ is 2π upon k . And that λ may be on the order of the Debye length. If we're down towards the incoherent limit, this is just the Debye length here.

And on these length scales here, the particles never collide. The mean free path is huge. And so these waves never see a particle collision. And so it's still legitimate to use this collisionless treatment, where λ mean free path is much, much larger than our scattering wavelength. So even in a very collisional plasma, I can still do collisionless Thomson scattering.

And that's good. If I wanted to do collisional, I'd have to put some collision operator in here. And if any of you have done plasma kinetic theory, it's a pain. But people do it. There's some really nice theory out there on Thomson scattering in collisional plasmas.

We are going to assume that we have some distribution function that looks like some background distribution function plus a perturbed distribution function, which is oscillating, so some Fourier decomposition here, so e to the $i\mathbf{k} \cdot \mathbf{x} - \omega t$. When we substitute that back into this equation, we get $-\mathbf{i}\omega \mathbf{f}' + \mathbf{k} \cdot \mathbf{v} \mathbf{f}' + q \int \mathbf{m} \mathbf{e} \cdot d\mathbf{f} d\mathbf{v} = 0$.

And we can define a current, which is equal to this perturbed distribution function times the charge of the particles we're talking about times the velocity integrated d^3v . And we can rearrange this to be an equation for $\mathbf{f}'$, substituting into this. And we get q^2 upon $\mathbf{i}$ times the mass, integral of velocity dotted into-- sorry-- velocity times the electric field dotted into the $d\mathbf{f} d\mathbf{v}$ over $\omega - \mathbf{k} \cdot \mathbf{v}$ $d\mathbf{v}$ like this.

We then say this current here, $\mathbf{j}$, is equal to some conductivity times the electric field. And if we have a conductivity, we know that the susceptibility is equal to the conductivity over $-\mathbf{i}\omega \epsilon_0$. So this is just from electromagnetism here.

We have taken all of the complicated plasma physics. And we're going to chuck it inside this σ , calculate it, and then we're going to say from electromagnetism, we know the link between the conductivity and the susceptibility here. So this χ is equal to q^2 squared over $\epsilon_0 m k$ the integral of $\partial \mathbf{f} / \partial \mathbf{v}$ over $\omega - \mathbf{a} \cdot \mathbf{v} d\mathbf{v}$. So now for some distribution function $\mathbf{f}$, we can calculate this χ here.

And I need to make this $\mathbf{f}$ in the direction of our $\mathbf{k}$ vector here, because along the way, I've stopped this from being a tensor. And I've just made it into a scalar. And we're taking the longitudinal component because these waves are all longitudinal waves. So this is the zz component of that tensor, which would otherwise be a 3 by 3 tensor.

So we're just interested in the component of this along the electric field, which is causing all of this in the first place. When you do this, you remember that there is a singularity in the denominator here. And you have to be careful. There's a pole at $\mathbf{v} = \omega / \mathbf{k}$.

And then you do your integral in the complex plane. And you close it so that you don't go around this pole. This is called the Landau [INAUDIBLE]. I assume that you have all seen this at some point. And I'm just reminding you exactly what's going on here.

So now we have-- we do this integral properly using the Landau contour, we have χ . And then we also have ϵ is equal to $1 + \chi_e + \chi_i$, where these subscripts here are going to come from us putting in the distribution function for these. So this is χ_j . And this is f in the k direction for the j -th species. Any questions on that?

Hopefully you've seen this in more detail.

So now let's consider a Maxwellian distribution.

So there's an arbitrary distribution function. If your plasma is not Maxwellian, you'll get some complicated looking spectrum you can, in principle, calculate. But if we specify a Maxwellian such that our f for the j -th species in the k direction-- so it's a one-dimensional distribution function-- is a function of velocity is equal to whatever density of particles, some constants to make sure that this is normalized so we get density out when we integrate over velocity.

Thermal velocity of the j -th particle electrons or ions exponential of minus v squared upon that thermal velocity squared. So this is just the Maxwellian. Then we find that χ , in this case, susceptibility for the j -th particle looks like ω_p^2 the j -th particle upon k^2 thermal velocity the j -th particle squared.

I've still got this $1/\sqrt{2\pi}$ hanging around. Then we have this integral of minus b upon bt_j exponential of minus v upon bt_j word over $\omega - kv$ -- sorry-- ω upon k minus vbv .

Now, finally, we get back to α because we recognize that this term here is just the Debye length for the j -th particle squared. So we can say that χ_j is equal to $1/\sqrt{k^2 \lambda_{Dj}^2}$. That is effectively $1/\alpha^2$, our coupling parameter. So this is why this has come back in and why I didn't find [INAUDIBLE] 2π is, because at the moment, it's [INAUDIBLE].

And then out in front here, we're going to have a term that looks like the charge of the j -th species. For electrons, this would just be minus 1. But for ions, it could be some bigger number-- the density of these species times the electron temperature over the electron temperature times the temperature of the j -th species. And we've got a squared there.

We've pulled a load of terms out because what I'm about to write next is a different function, but I just want to point out if you stare at this a little while, this is equal to 1 for electrons. For electrons, this term just vanishes here. It's kind of obvious for one electron density, electron temperature, electron density, electron temperature. This all cancels out.

But for ions, it will be a different number. And that will be mostly due to the fact that the ions may have different temperatures than the electrons and due to the fact that the ions will have different charges. And then the thing that is left is this function, which we call w . And it's a function of some dimensionless parameter squiggle of j .

This squiggle parameter-- I don't remember how to pronounce it, so I may as well be honest-- is equal to the phase velocity over the thermal velocity. Remember, the phase velocity is ω upon k . So this is one when we're talking about particles or thermal velocities that are same as the phase velocity. And it's less than 1 or greater.

This function here is the-- I think it's the Wolf-Z function. And it's simply a function, which is defined to be-- to involve the solution to this integral in a dimensionless way. And this function w of squiggle is equal to $1 - 2$ squiggle e to the minus squiggle squared integral from 0 to squiggle e to the x squared ex plus a term, which is the residue due to the Landau contour, $i\pi$ to the $1/2$ squiggle e to the minus squiggle squared. Obvious-- take it home to meet your parents.

This is a very ugly function, but because we use Maxwellian so often, this is a function which can be tabulated. So we can just calculate it for various values of epsilon. We can make a table of-- sorry, not epsilons-- of squiggles from very small numbers to very large numbers. And we can just calculate this up.

Remember, it's a complex number, a complex value here. But we can have a lookup table here. And there is something called the plasma dispersion function. I think it's normally written as z . That is related to this function here. And it's related. It's sort of like-- I can't remember what it is, but the z is equal to squiggle to the $1/2$ w of squiggle, or something like that.

I can't remember exactly what it is right now. I prefer to work with this thing, because this is a thing that you'll find in Python, and Julia, and Matlab, because it shows up in other contexts as well. And so, people, there are very nice fast implementations of this function in all of these languages and many other languages. Fortran as well, no doubt.

So it's almost always better if you're writing your own code to write it in terms of this function, rather than the plasma dispersion function. But you may have come across the plasma dispersion function in other Landau damping contexts. And there they're very, very similar.

This is also related to the error function, the complex error function, and a related family of functions that crop up in all sorts of other interesting contexts as well. The link between them all is that they have these fun integrals and e to the something squared inside them all.