

[SQUEAKING]

[RUSTLING]

[CLICKING]

JACK HARE: Welcome, everyone, to the final lecture. We will be discussing some of the ramifications of the collective Thomson scattering spectrum, which we derived last week. So just to remind you, we derive first of all that the scattered power into a given solid angle into a given frequency seen by our spectrometer was equal to the classical electron radius squared, the intensity of our laser and its power over the area. There was some shape function like this that told us where we should look to see the scattered light. It was the polarization of the electric field.

It depends on the number of electrons we're scattering off and how many electrons are within our scattering volume, and then it depends on this interesting object, the spectral density function, which is what contains all of the actual frequency resolved information here. And this spectral density function, we found, depended on the time and the volume over which we were integrating. And it was an average over the absolute value of the Fourier transform of the electron density, squared-- this normalized by the electron density.

So after we derive that, we said, OK. Now we need to find out what these electron density fluctuations look like. And we did that using a test particle formalism, where we let some electrons or ions flow through our plasma. And we looked at the response of the plasma to these electrons' ions, with the electrons being repelled from other electrons and ions being attracted and all sorts of things like that.

And from that, we were able to write down a formula for the spectral density function for some arbitrary distribution of electrons and ions. And this was 2π upon k and E . And then there was a factor that looked like 1 minus πE upon ϵ squared. The electron distribution function in the k direction, evaluated at the phase velocity of the wave, ω upon k , and the component that looked like πE upon ϵ squared times πv .

Charge on the ions and the ion distribution function in the k direction, evaluated at ω upon k , like that. And we said that this was the number of electrons at ω on k , moving with that velocity, and that this was the response of the plasma to this number of electrons. And the same here-- we have the number of ions and the plasma response here.

And so these we call, respectively, the electron component and the ion component. And as I've stressed several times before, we are only ever doing scattering from electrons. This turns out to be Thomson scattering from electrons looking at other electrons being repelled by them. And this is Thomson scattering from electrons which were following the ion particles around. So we're scattering off electrons in both cases, but some of the electrons tell about other electrons, and some of the electrons tell us about ions.

And we said that if we specify this now so our distribution function is a Maxwellian so we have for the j species, the distribution function in the k direction is a function of velocity. That's going to be the number density of the j species from factors to do with normalization-- 1 over square root 2π , 1 over the thermal velocity of the j species, exponential of minus v squared upon the thermal velocity of the j species squared, where this thermal velocity is equal to the square root of $2T_j$ upon m_j . So if we assume that this is Maxwellian for the electrons and the ions, which may have different temperatures. So we're still allowing T_e to not equal the T_i here.

We found that these susceptibilities, χ , for the electrons was equal to α squared times this mysterious w function, squiggle e . I'm going to define all of these in a moment. And for the ions, was equal to α squared zT_e upon $T_i w$ squiggle of i , where if α is our collective scattering parameter and it's defined as 1 upon k lambda Debye.

And if you're curious about which Debye length it is, it's the Debye length of the electrons. And this function, w of squiggle is ω upon k times v_{Tj} . So squiggle for the j species looks like that. So it's the ratio of the phase velocity of our wave or mode over the thermal velocity of our distribution function.

And this w here of squiggle as a form that looks like 1 minus 2 squiggle e to the minus squiggle squared, integral from 0 to squiggle, e to the x squared. The x , that's the real part. And then there's an imaginary part, i square root of π squiggle a minus squiggle squared. And this, of course, is the imaginary part of the function.

And we often write this as equal to the real part of w of squiggle plus i times the imaginary part, w squiggle. And we're breaking it down into the imaginary part several times in what's to come. OK. The physical meaning of this squiggle parameter here is to do with what part of the distribution function you're interacting with. So if I draw a little plot here of velocity versus log of the distribution function here, and I say for Maxwellian, this is just going to look like minus v squared so it'll be something like that.

So this is the distribution function for the ions. Then if I'm looking at how modes with low phase velocity compared to the thermal velocity interact with this distribution function, maybe down here. So this is squiggle ion much less than 1 . Or I could look at high frequency modes up where there are very few particles with a velocity that matches the phase velocity. This would be squiggle greater than 1 .

And of course, if I'm somewhere in the middle, this is squiggle about 1 here, where we're sampling in the middle of this distribution function. And it's worth noting, of course, that for roughly equal temperatures, the electron distribution function at e is going to be much broader. And so that means that, in general, squiggle for e is going to be less than squiggle for i when the temperatures are equal.

So if you're asking, how do the ions interact with the mode of a certain frequency? Then the electrons will be interacting as if that mode had even lower frequency. Because the electrons-- this line, for example, here, which was high frequency for the ions is going to be low frequency for the electrons. And you'd have to go all the way out here to find a mode where squiggle e is much, much greater than 1 here.

So in general, the fact that the ions and electrons have very different masses will divide the spectra that we see into modes which interact with the ions, the modes that interact with the electrons. And those are the anacoustic waves and the electron plasma waves that we heuristically discussed already. So that is my summary. What we're going to do next is take this rather complicated looking formula and do some limits on it in order to recover the modes that we know are hidden within all of this hideous mathematics. But before we keep going, does anyone have any questions? Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: No. Well, yeah. You can solve it for a bit. I'm pretty certain it's positive. No. It's χ_e . So if you remember, again, we're just doing the electrons. The term here, the 1 here, is just the test particle itself. So each electron in the distribution function.

This term here is e minus repelled from the test particle. Over here, there would be a 1 if we had Thomson scattering off the ions, but we don't. So we don't see any scattered light from the ions themselves. So there's no 1 here. All we have is electrons attracted to the ions.

One sort of formula that maybe you remember, but I will write down just to be very, very clear, is that this ϵ is equal to $1 + \pi e + \pi I$, like that. So that's the permittivity, and it's related to the two susceptibilities like that. OK. Another question?

AUDIENCE: Wouldn't it be 1 over [INAUDIBLE]

JACK HARE: Yeah. The volumes can cancel very nicely here. I think the time comes from the integration of the incident power.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yeah. Once we rewrite it in terms of this, they've disappeared. I'd have to check my notes. I think it might be related to one of the conversions between f and density, yeah. No, this is like an average over an ensemble of articles.

AUDIENCE: I don't think it's a time average.

JACK HARE: Yeah. I can't remember off the top of my head. I can certainly check it out later. OK. Any other questions? All right. So now we are going to approximate this w here.

And just to rewrite it here, we have w of x . I'm just going to use x instead of squiggle because it's easier to write. It's equal to the real part of w plus the imaginary part of w . And for x less than 1, it turns out that the real part of w has the following expansion as an approximate solution-- $1 - 2x^2$ times $1 - 2x^3$ upon 3 plus a load of other terms that you could keep going to, but we won't in this class.

And then for x much, much greater than 1, we have this real part of x being approximately -1 upon $2x^2$, $1 + \frac{3}{2}x^2$, plus some other terms. And the way you get this is by carefully thinking about what's going on with this real part here. It's not completely trivial to get.

And just to write down again that the imaginary part for all x is just equal to square root of $\pi x e$ to the minus x squared. In general, we're very often going to try and take this imaginary part to 0 for small x . But when it goes to x , it goes to 0 for small x and for large x , which is the two regimes we're going to be working in.

There is one place later on where I'm going to talk about the actual value for this. But we tend to be working in these two limits because that's what makes these results analytically tractable. Of course, you can always just go and calculate the full s of k ω using this full tabulated function that you can find in Python and Julia and Matlab and all sorts of things like that.

So the first limit, we're going to pick the limit where α squared goes to 0. Yeah, I wish I had more boards. You can see the principal thing that this is going to do is that these susceptibilities both scale as α squared.

So for both of them, the susceptibilities are going to go to 0. So χ_e and χ_l are going to go to 0. Remember the limit where α squared goes to 0 is the same as saying that the wavelength of the fluctuations we are scattering off is less than the electron Debye length.

So this is scattering within the Debye cloud, and we expect that this scattering will therefore give us our incoherent spectra, which we derived a long time ago under much simpler circumstances. But I did promise that this new formulation would recover that. So as these two go to 0, our connectivity, which is just $1 + \epsilon + \pi l$ just goes to 1 like that-- quite boring. Because it goes to 1 for all frequencies, we don't have any waves here. The permittivity is never 0, so there are no waves.

And indeed, what we see is that s of k ω , which is proportional to $1 - \chi_e^2 / \epsilon^2 f_e(\omega) / k + l \epsilon^2 f_l(\omega) / k$. When we send these χ_e 's to 0, we just get that this is proportional to the electron distribution function, exactly as we found for the incoherent case.

So in this case, we've dropped the response of the plasma. We've lost all information about the plasma response. We've just kept the scattering off each of the individual electrons here. So no plasma response. Just individual electrons. Does that make sense? Any questions?

I just want to clarify when I say let's take α to 0. This is the thing that you have complete control over. So α is defined as being $1 / k \lambda_D$. This is the size of k . So this is this, which is equal to $2k \lambda_D \sin \theta / 2$. You've chosen the wavelength of your probe if you make that wavelength larger. So you go to a longer wavelength, you'll get a smaller α .

So larger-- actually, the other way around. Sorry. Yes. Larger. I'll write it in terms of wavelengths. So we tend to think of lasers like they're smaller, lower, α . Or we can also do the same thing with this. We can have a larger θ .

Taking this limit here is not some mathematical abstraction. This is a choice that you've made as an experimentalist. Now, of course you might be constrained in what your laser wavelength is. You might be constrained where you can place the detector. And then you'll go, OK, well, it turns out my α is very low. So I guess my Thomson scattering spectrum is incoherent. But perhaps there's a chance to place your detector somewhere else or use a different wavelength. And you could start looking at the coherent particle spectrum, which we'll discuss now. OK. So any questions on that before we keep going? Yes.

AUDIENCE: Is measuring the electron distribution function cytoplasm useful?

JACK HARE: Very. Is measuring the electron distribution function cytoplasm useful? Yes, extremely useful. Because this could include information about fast particles, like runaway electrons, particles which have been heated by IDRf. Non-thermal particles doing chemical reactions in a low temperature plasma-- this is incredibly useful, yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: It depends what you're trying to do. In general, it's very hard to get both in the same experiment, but if you can, you can do some really cool things. Yeah. OK. Other questions?

AUDIENCE: [INAUDIBLE]

JACK HARE: Oh. What I mean-- sorry. Yeah, there's no mode to the plasma. What we found when we looked at s of k . What we found when we looked at this before is we said, OK, if we do get ϵ going to 0, then there will be singularities in f of k ω . There'll be resonances. And so our spectrum will contain resonances which are related to modes in the plasma, waves in the plasma. But I'm talking about plasma waves here, not electromagnetic waves. I'm just being a little bit brief here. You're quite right.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yes. You just scatter off the individual particles. The wavelength of the scattering is smaller than the Debye length, and so we only see individual particles. We don't see these collective plasma effects do Debye shielding. And it's those collective plasma effects which give rise to the modes and the waves within the plasma. So that's a reasonable way of thinking about it, just too fast. Yeah.

Yes. What we're going to see very soon is that yes, of course, you get the electron distribution and the distribution function. But they're multiplied by this thing, which is a very complicated function of frequency. And so it's like your electron distribution function put through the mangle. And so you may not be able to work out what your original electron distribution function was from that-- again, because as we're about to show, these are present when ϵ goes to 0.

And so at those particular-- I should be clear. All of these here-- this is πe of ω upon k and ϵ of ω upon k . So what this means is that for a specific ω upon k which makes ϵ go to 0, this function will really blow up. So maybe before, you had scattering just off the electron distribution function, and it looked like this. But now as you raise α , the presence of these waves become more apparent.

And these waves are a singularity, so the scattering is absolutely massive. So now your spectrum looks like this. And so now you can't really see the electron distribution function at all. It's just some busy background noise. And all you're actually seeing is the scattering of the electron plasma waves on anacoustic waves.

Any other questions? But it sounds like we should go on to the next because we're getting ahead of ourselves. OK, cool. So now let's talk about arbitrary α squared. So α , not necessarily going to 0. And I'm going to cover two cases here.

The first case is the high frequency waves. So for the high frequency waves, we're going to say that ω is much, much greater than 1. These waves of frequencies way out, they have phase velocities much higher than the thermal velocity. So there's very few ions around here. And this effectively means that w of squiggle goes to 0, which is the same as saying that there's no ion response.

So just to be clear here, what we've done is we've taken this limit here as x gets very, very large, and we've just gone, well, the first order effect gets very, very large. Even this term vanishes. There's no just one term up front here, so we're just going to approximate the real part of this as 0, and the imaginary part will also go to 0, because for very large x , this exponential for minus x squared kills off this as well.

So we don't have to worry about what the ions are doing anymore. They're irrelevant. What you do find with the electrons is that you get s of $k \omega$ is equal to $2 \sqrt{\pi}$. And then you get the term which is quite interesting. And I'll put the $k v_e$ squared here.

You've got a term that looks like exponential of minus squiggle e squared upon something a bit more complicated. This is $1 + \alpha^2$ real part of w , evaluated for squiggle e . All of that squared plus α^2 squared imaginary part of w , the squiggly e , all squared. This is just by taking the previous expression I gave you and reducing the susceptibility down towards 0.

The reason this is interesting is it goes directly back to Grant's question. This here, this looks exactly like e to the minus v squared upon $e T_e$ squared. So this is equal to your Maxwellian. You've got a spectrum that knows about the Maxwellian, but the Maxwellian is being modified by this horrific denominator down here. I think there should be another closing square bracket there. Maybe I didn't need those, but there.

So you're not just getting back the electron distribution function. You're getting back the electron distribution function, and it is strongly modified. So now we want to go see what modifications this interesting denominator gives us. So we're going to be looking for resonances, as always, where the real part of the permittivity goes to 0.

We're going to be using, again, high frequency. So we're going to say that the electron squiggle factor is much, much greater than 1. Note we said the same thing for the ions. It's not as big as the ions. So squiggle for the ions is still much bigger than squiggle for the electrons, which is much bigger than 1.

So we're not going to be as brutal in cutting off all of the terms for electrons as we did for the ions. We're still going to take the imaginary part of w for squiggle e to be 0. We're going to ignore the damping here. But we're going to take the real part of w to your e , and we're going to keep the first two terms. And those first two terms are minus 1 upon 2 squiggly squared minus upon 3 upon 4 squiggly to the 4th.

And so if we're seeking solutions where this is equal to 0, we'll end up with something that looks like $1 + \alpha^2$ squared, real part of w , squiggle e is equal to 0. Because again, just looking at this denominator here, this is permittivity. And that gives us out a dispersion relationship, which is $\omega^2 \epsilon_p w$ is equal to $\omega^2 p_e$ squared plus $3 T_e$ on me k squared. So we have recovered from all of this complicated mathematics the electron plasma waves here.

And so our s of $k \omega$ -- I'm going to put it like this. I'm going to put some little crosses here to suggest that these are pretty high frequency compared to other ways we'll be talking about. It's going to have some little peaks in it. Remember, all of our analysis so far has been for high frequency, so we don't actually know what the spectrum looks like at these low frequencies here. But we do know what it looks like at large frequencies.

It's got these very sharp resonances. The reason that they are not actually singularities is because there will be a small amount of damping. We've neglected it in our treatments here just so that we can get the dispersion relationship. But if you allow a small amount of damping back in, that will prevent the denominator from going to 0. And what you find with these-- these are occurring at this frequency, the electron plasma wave, and at minus the frequency.

I think I'll cover that later on. And so the position of these depends on the plasma parameters. This depends on the density, and this of course, very, very clearly depends on the temperature. So the exact location of these peaks gives you information about the temperature and the density. And it turns out the broadening of the peaks-- so this sort of thickness in frequency space is related to the temperature of the electrons as well.

And you can get that out if you do this more thoroughly with the Landau damping term, this imaginary component. That imaginary component broadens these peaks, and the amount of broadening is to do with how hot the electrons are. OK. So that is the first mode we wanted to get.

And this is the mode related to the electron plasma waves, or Langmuir waves, if you like. And these are what we call high frequency waves. So of course, their frequency is still much less than the frequencies of the radiation you've heard. OK. Questions?

Yeah. You can rewrite this so it looks like something plus alpha squared. And then in the limit where alpha is really big so you're very collective, it turns out that it really just depends on the density. And this is the limit that lots of people work in. So you'll find papers where they're like, this frequency omega is proportional to density, and the width is proportional to temperature.

But that's only true for alpha much greater than 1. If your alpha is like 3-ish-- and I've done experiments in this regime. Then the position depends on both of these parameters. And so you have to go find the width. That gives you the temperature, and then you sub the temperature out from the displacement of the location of the peaks, and that will give you the density.

In reality, you don't do any of those things. You fit. You do a least squares fit to the full spectral density function. But if you don't want to do a least squares fit, you can still actually estimate all these parameters. And that's what my old Russian professor did. He would just go and point and go, ah, this is the frequency shift. And he would calculate temperature and density just from looking at the raw data.

You can also do it using computers as well. Any other questions? OK. We're going to go to the low frequency.

So just to remind you, we did arbitrary alpha, alpha squared, and high frequency. Now we're going to go looking for low frequency waves here.

So for these low frequency waves, we know that squiggle for the electrons is going to be much, much less than 1. And so we're just going to say that the real part of squiggle for the electrons is going to be about 1. Just to remind you, the small x. This is $1 - 2x^2$ plus a load of other terms. So for very small x, I can just drop this term and just get the 1 here.

OK. So now we find that our permittivity, which is equal, as always, to $1 + \frac{\omega_p^2}{\omega^2 - \omega_c^2}$, is equal to $1 + \frac{\alpha^2}{\omega^2}$, the imaginary part of ω for the electrons. I haven't dropped the imaginary part because they play an extremely important role in damping these waves. So I've got to keep this. Plus $\frac{z T_e}{T_i}$ times $\frac{\alpha^2}{\omega^2}$, the real part of squiggle ω plus ω times the imaginary part of squiggle ω .

And again, we seek resonances where the real part of the permittivity $\epsilon_r = 0$. And if you play around with this formula for long enough, you'll find that these resonances are damped. That means they don't appear. And they're damped for the condition that $\frac{z T_e}{T_i}$ is less than 3.

I haven't proved this at all. You can go and prove it to yourself by trying to solve this condition here. And it's in Froula's book, and it takes a few lines of algebra. So I'm not doing it on the board. So we don't get any resonances if we have this condition. So damped in this case means don't appear.

So let's just quickly take a look at this. We're very used to having plasmas where the ion electron temperatures are equal. But also, many of us are used to having plasmas where z is equal to 1. So in this case, this would be T_e is less than $3 T_i$ for a hydrogenic plasma. So this would be ions hotter than electrons.

So that's not a plasma you necessarily find very often. So this is a slightly uncommon thing to not have these resonances appear. In fact, most of the time, we're going to be in the regime where $\frac{z T_e}{T_i}$ is greater than 3. And this will be possible even with electron and ion temperatures being equal, if the charge on the ions was 3 or more.

So this could definitely happen if you have a carbon plasma. The electrons and ions are in thermal equilibrium, but because of the charge on the ions, this condition might be fulfilled. So this is relatively easy to find. And we're going to often-- this greater than 3, this is really the same as being much bigger than 1. Because as we all know, 3 is much larger than 1.

So if you fulfill this condition, then you find that the velocity of your anacoustic waves here, which is equal to $\frac{z T_e}{m_i}$ at $\frac{1}{2}$ must be much, much larger than the thermal velocity of your ions, which is equal to $\frac{T_i}{m_i}$ at $\frac{1}{2}$. Just by plonking this condition that it's much greater than 1 into here, we can see that the modes that we're looking for-- and these will be the anacoustic waves-- are going to be moving much faster than the thermal velocity of the ions themselves. And so that means that we can now write down that squiggle for the ions is going to be much greater than 1.

That means that the real part of ω or the ions is roughly minus $\frac{1}{2}$ squiggle ion squared minus $\frac{3}{4}$ squiggle ion to the 4th. Remember, for the electron plasma waves, this is the limiting factor that we use for the electrons. But now we're doing it for the ions here.

And we're again going to assume that there are no damping from the ions. And that's because these waves are velocities much faster than most of the ions in the system. And so there's no ions to damp up. So we're going to set the ion damping to 0.

Just remember, we have not set the electron damping to 0. So the reason that these resonances don't blow up is because they damp on the electrons. And then we can do the same thing we did before. We can substitute all of this in to the permittivity. We can find the resonances when this is equal to 0.

And we find that we have resonances now at the anacoustic wave frequency. And this is equal to k^2 . What follows may surprise you-- $\alpha^2 \frac{1}{1 + \alpha^2 \frac{z_{Te}}{m l} + 3 \frac{Tl}{m l}}$. Why do I write it like that, not just like this? Oh, because actually nobody asked that question.

So this looks a little bit like what you're expecting, the anacoustic waves. But it looks most like what you're expecting if α^2 is much, much greater than 1. This goes to $\frac{z_{Te}}{m l} + 3 \frac{Tl}{m l}$, because this factor, this goes to 1 at that point. And that looks a fair bit more like the anacoustic velocity that you're familiar with.

So now we know that this spectrum is going to have resonances which are appearing at this anacoustic frequency. So questions on that before I sketch what this spectrum looks like? Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: There's no [INAUDIBLE] or radii in here. No point to-- sorry. What does it mean physically? If α is small now, this term just goes away. OK. I'm going to show you what it means physically in a moment. Actually, no. That's not in these notes. That's something else.

α is small. If α is small, although these resonances will appear, they will disappear because a small α -- the ions and electrons don't respond anyway. And we get back to our incoherent scattering spectrum. So if I made α smaller here, although it would look like this term would drop out, and I'd get anacoustic resonances happening at $3 \frac{Tl}{m l}$, in reality those anacoustic resonances wouldn't occur because we'd just be doing incoherent scattering. Can I explain exactly why it's $\alpha^2 \frac{1}{1 + \alpha^2 \frac{z_{Te}}{m l} + 3 \frac{Tl}{m l}}$? No, not at the moment.

This comes about because we have modes. But α^2 . I don't know. We have derived this assuming that α is-- no, we haven't derived assuming α is not small because we've done it for arbitrary α .

So how would we get back the incoherent result? OK. So the interesting thing is that-- I think I might have erased it, sadly. Let me just go back.

χ_e is equal to $\alpha^2 \omega$ of squiggle e . And χ_l is equal to $\alpha^2 \frac{z_{Te}}{Tl}$ squiggle l , like that. So when we're doing the incoherent, we're just looking at this first term, and we're taking that to 0. So no matter what ω is doing-- and in this case here, ω has got all sorts of funky resonances in it.

Mathematically, at least, that will just not be important because whatever this response of the plasma is goes away because we are scattering with modes smaller than Debye length so that we can't see this part anyway. So it's like they're two separate arguments. We use the α going to α^2 going to 0 argument to get rid of this term to make it small, setting these equal to 0.

If we did that again-- I'm trying to think what you would end up with. But I have a feeling you're going to end up with a term that looks like α^2 the 4 or the strength of this resonance. And then that is still going to go to 0. So we could feed all of this back through.

Yeah. I don't know the full answer to it. I think we're talking about two separate things here. I think that the limit that we took to get the incoherent scattering is separate from what we've done here. So although this term would still exist in the coherent/incoherent limit, it would be inside this bit of the equations. And so whatever these are doing is irrelevant because this.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yeah. We took this very brutally to start with so we didn't even have to think about these. Once we restore the α being some arbitrary value, then we can go on about these things.

AUDIENCE: It's not like [INAUDIBLE]

JACK HARE: No, I don't believe so. There's nothing related to α inside w , yeah. w really has to do with the distribution function, whereas α is more fundamental than that. α can be evaluated regardless of distribution function. So probably if you're trying to work out what the Debye length is for a nonthermal plasma, you also have problems. So I'm not quite sure how this gets defined if you're very non-Maxwellian.

We're going to move on. if you happen to be in the case back here where zT_e upon T_i is less than-- we've got 3 here, but I'm going to say less than 1. It turns out you get a spectrum s of a ω . And this is for low frequencies. So we don't care about what the electron plasma wave is doing. You get a spectrum that actually looks like the ion distribution function.

Now, this is because the anacoustic waves are strongly damped in this limit here. So the anacoustic waves would be somewhere here. But they're Landau damping on the ions when we have this criteria fulfilled. And so the waves are completely damped.

As far as I know, I've never seen a paper where someone has done this. This is actually a pretty strange place to be, where you have really rather hot ions and cold electrons. I assume that it is doable, but I've just never seen it done in an experiment.

The next limit we could have is where zT_e upon T_i is roughly 1. So that's in between. The resonances are completely damped, or the resonances are not damped. And so what we end up with is a spectra that looks like a strange hybrid.

It almost looks like a flat top like this, which is effectively adding some anacoustic waves onto your distribution function. These anacoustic waves are still pretty damped, but they're not as strongly damped here. So there is an anacoustic wave contribution.

Finally, we have the case where zT_e upon T_i is much bigger than 1. This is the case we just derived, where we're going to have very strong anacoustic resonances. And in this case, your spectra-- although it does have some contribution from this f of I , as we've discussed before, the contribution from the resonances are much larger. You get something that looks a little bit like that. There we go.

And these occur at the anacoustic free. So now we have a large ionic contribution. And of course, if you write a code to implement all of these equations, and you change the value of z upon T_i , you will get these different shapes. I'm just describing to you hand-wavingly what shapes you get and why they end up looking a little bit like that. OK. Any questions on the anacoustic waves?

Wherever you want to be-- it depends what you want to measure. No. it's pretty irritating. I ended up there once. You'd much prefer to be here. You can learn an awful lot more from it. This is hard to fit properly.

And there's other things that can cause your spectrum to look like that as well. So there's probably a better place to be, but of course you don't get to choose this in your plasma. Unlike where we're choosing α by moving our spectrometer and changing the wavelength of light, I have no choice whatsoever on what my temperature of my plasma is. That's the thing I'm trying to measure. So sometimes you just end up here by accident.

OK. So we derived the anacoustic waves and the electron plasma waves. And a reasonable question is, in my spectrum, which of these is more important? Which do I see more light from? Because that might tell you which ones are going to be easier to detect.

So again, we might have a spectrum where α is much bigger than 1 and zT_e upon T_i is much bigger than 1. And we might go, OK, we're going to have some anacoustic peaks here. And then at some much higher frequency, because they really are very, very different frequencies. We're going to have some electron plasma waves, like this.

And which of these contribute more to our spectrum? So we get this by calculating what's called the total cross-section, which we call $s_{\text{sub } T}$, as a function of a , where our detector is. And this is just the frequency integrated cross-section of s of $k\omega$. So we can do this analytically for the ion and electron temperatures being equal.

So you have to make this assumption in order to make progress analytically. And this is called the saltpeter approximation. The saltpeter was the first person to come up with it when he was looking at Thomson scattering from the ionosphere. And he showed that our total scattered cross-section looks, therefore, like 2π upon $1 + \alpha^2$. This is the contribution from the electron feature.

So this region. And then there's another term, which looks like 2π upon $1 + \alpha^2$, $z\alpha$ to the fourth power, $1 + \alpha^2$ plus α^2 , the T_e upon T_i . You're going to say, we assume these are equal. We assume they are approximately equal.

There we go. Good. So this can still be slightly different than 1. And of course, the z makes a big difference here. And this is due to the scattering from the ion feature.

But now we look at this, and we go, OK. Well, depending on my α parameter, this seems to depend very, very strongly on α . So if I look at α^2 going to 0 here, it's very clear that this term, the aisle term, is just going to drop off. And we're just going to have electron feature only.

Which is what we've seen several times already. But when we go to incoherence, we're just scattering off the individual electrons. If we have α equals to 1-- so some intermediate α parameter-- then we're going to get something that looks like π times $1 + z$ over $2 + zT_e$ upon T_i .

And so there are two limits to this. There's the limit we already saw, zT_e upon T_i is much, much less than 1. Then we got both these terms. So again, this is the electron term. This is the ion term. This gives both e^- and ions.

And we actually end up with coherent electrons and incoherent ions-- or incoherent ions scattering from the ions. This is a strange situation, as I've said. It requires you to have hot ions and cold electrons, which we don't often encounter. The other one would be zT_e over T_i is much greater than 1. And then we get coherent electrons and ions. And finally, just so-- well, I'll create some space here.

We could look at when our alpha parameter is very large. Alpha much greater than 1. And then we'd have this $k \omega$ go to sk total go towards 0 for the electrons plus z over 1 plus zT_e over T_i , the ions. And so we just have ions dominating this.

If we have a very large alpha, your spectrum would just look like-- and then flat out to here. In some of these other cases, for example, we just have electron scattering. And we wouldn't have any contributions from the ions. And this is important because depending on what alpha you have, if you're like, I really want to go look at the electron distribution function directly, but then you calculate your alphas large.

But you know you're not going to see any of the electron feature at all. The electron feature is going to disappear. So you won't learn anything directly about your electrons. You'll only learn something about your ions.

And this is important because in general, the eTw frequency is much, much larger than the law frequency, as I alluded to here. So it's very hard to measure both on the same spectrometer, so we can't. The eTw 's and the law's on one spectrometer.

Because if you set up your spectrometer to have a nice, wide frequency wave to get the eTw 's, then the law's will just appear on a couple of pixels. And if you set it up to see the law's, then the eTw 's will be way off to the side. And so that means you generally have to choose which of these you're going to measure or you have to buy an alpha spectrometer.

So it is actually a pretty big practical concern. And so you want to know which of these features is going to be brightest so that you can set up your spectrometer to go look at that, because that's going to make life better. OK. Questions on this? Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yes. We haven't talked about magnetic fields at all yet. Oh, yes. If we have time, we'll get there. But the very, very, very brief answer is in certain conditions, the anacoustic feature gets a modulation so it wobbles up and down. And that wobble is the ion cyclotron frequency. And the electron feature gets a wobble, and that wobble is the electron.

But you have to be looking precisely perpendicular to the magnetic field, which is very hard. If you look even half a degree off, this effect washes out completely. So it's a very challenging measurement to make. I haven't seen anyone do it since 1972.

I didn't see it myself. But that's the last time someone published a paper actually measuring it in Thomson scattering. And I've not seen anyone do it. It's a really hard measurement.

So people say, oh, Thomson scattering can measure magnetic fields. Because you measure the cyclotron frequency, so you're like, OK, now I know the magnetic field. But in reality, it's really hard to do. Other questions? I will show you something in a moment that can help you infer the magnetic field, though.

OK. So now I'm going to talk about electron ion drift. And as you probably realize, if your electrons and ions are drifting with respect to each other, that means you have a net current in your system. And so if we can measure how much electrons and ions are drifting with respect to each other, we can measure the current locally within the plasma, which is pretty cool.

So we're going to say that there is some drift velocity which is equal to the difference between the electron velocity and the ion velocity, like this. Now, this is not the velocity inside your distribution function. This is the bulk velocity of your entire distribution.

So this is a big capital V here. So your entire distribution of ions is moving with some V_I , plus the thermal velocity on top of that. The entire distribution of electrons is moving with some V_e . And we define this drift velocity to be the difference between the flows of the electron ion species.

And we know from all the other Thomson scattering that we've done that we're only really sensitive to a component of this in the k direction. So we can only measure V_d dotted into k because everything we've done we only measure it along the wave vector of the scattering mode. And so that means that we can define a squiggle, which we call squiggle d , which is defined as V_d dotted into k over V_{de} .

Once again, this is analogous to the other squiggles, where we were using the phase velocity. Now we've got the component of the drift velocity along k . And that means we're going to replace the real part of the w function, which was normally in terms of squiggle.

And we're going to now evaluate it in terms of squiggle e minus squiggle d . So we're effectively looking at the damping from a different part of the electron distribution function. So this is our shifted Maxwellian.

The reason this is important is because it was the electrons which were damping the anacoustic waves. And both the anacoustic waves were damped by the same amount, and they had the same intensity. But if we shift the electrons, some of the anacoustic waves are going to see a different part of the electron distribution function than the ones going the opposite direction. So we've broken that symmetry.

Previously, for example, our electron distribution function was symmetric and our anacoustic waves were, say, here and here. And they both saw the same, df/dV , which is what causes the damping. If now we've shifted our distribution function, we can see that this wave now sees no df/dV because it's at the maxima.

This wave, these are different amounts. So we can imagine that the damping will be differently effective for the plus and minus frequency waves, and so they'll have different heights. So just to make that clear here, the damping of the anacoustic waves is proportional to the gradient of the electron velocity, e , at the phase velocity of the wave, ω upon k .

And that, it turns out, is proportional to this imaginary part of the w function that we've been carting around all this time. So it's proportional to the imaginary part of that at x . And that is proportional to x exponential minus x squared.

And you shouldn't be surprised that it's x exponential minus x squared because that looks like V exponential minus V squared, which would be the derivative of our Maxwellian distribution function. And so this is what we're saying, that the amount of damping we have is directly proportional to the derivative of the distribution function.

So if I now draw this again, but now I draw it in terms of the damping, I have axes with squiggle I down here. And up here, I have log damping. Previously, for example, I may have had the anacoustic resonances, these frequencies, minus omega law plus omega law. And the damping coming from the ions would have looked like this.

This is just the derivative of the Maxwellian here. It's just this function plotted. And the damping from the electrons would have looked like this. So the ions are feeding the ion acoustic waves because of the slope here. And the electrons are damping these waves.

If we shift the electrons now-- I put the anacoustic waves in the same place. I put the ion damping-- I'm going to try to make this one a bit more symmetric-- in the same place. But I now shift the electrons over so that they've got some drift velocity like this.

We can see that we have increased damping here and reduced damping here. And what that means, in terms of the spectrum, is that we go from something that looks like, in frequency space on our spectrometer, this. And we end up with something that looks like this instead. There's a very pronounced asymmetry between these two peaks here.

Now, unfortunately there's no simple formula I can give you for the ratio of these two peaks that gives you the drift velocity straight away. This is pretty complex. In general, as far as I know, you have to fit the full s of k omega, allowing there to be a drift of the electrons here. And it depends on the exact-- the ratio of these two peaks depends on exactly what the drift velocity is, what zT_e upon T_i is, what the alpha parameter is.

And so we have to fit the full s of k omega to our data. But once we've done that, we can get out this drift velocity here. And then we know that our current is equal to minus $e n V_d$ plus $z e$ and I VI. So this is simply equal to e and $e V_d$. So if we have some way of knowing the density-- for example, from interferometry or from scattering of electron plasma waves of a separate spectrometer-- by inferring the drift velocity from the spectrum, we can now get out the local current. Of course, it is the local current in the k direction.

So if your current is pointing upwards, but you've chosen your scattering geometry so that k is perpendicular to it, you won't measure anything. So this doesn't tell you the direction of k . It just tells you the size of it. If you know the direction of J from symmetry, and you have measurements of J at different positions-- for example, if I do Thomson scattering at several positions along x .

I've got my laser beam going through the plasma, and I collect the scattered light several places along here. And I calculate J at each of these points, then with some symmetry arguments, I may be able to solve $\text{curl of } B \text{ equals } \mu_0 J$ and locally find the magnetic field. But that requires a lot of assumptions about symmetry.

So it may not be completely general, but even getting to current is pretty cool. You get the magnetic field. That's really nice. OK. Questions on that? Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: No. We're just going to work in the ion frame. So this V_d here, I've basically set this equal to 0. And the ions we can assume just aren't drifting and [INAUDIBLE]. Just choose a frame, whether or not. So I can always choose a frame with the arms stationary. This only depends on the relative drift between the electrons and ions. It doesn't matter to the Thomson scattering which of them is actually true. Yes, exactly. I've never seen a case where the-- I don't think I've seen a case where the ions are drifting, and you're like-- that's unusual because the electrons tend to be the ones that move. Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yes.

AUDIENCE: Why do we have this?

JACK HARE: Oh, I'm just saying in terms of solving the equation. Yes, it definitely is. And the solution only depends on the difference between the two. Yeah.

So all the Thomson scattering will tell you is the difference between the two of them. It's then up to you to say, and I believe that the ions are stationary in my lab frame or the electrons are stationary in my lab frame or whatever else. But the drift velocity will be invariant, regardless of what frame you choose. Other questions?

We talked a little bit about what you can learn from the electron plasma waves before. Just in terms of what you can learn about these-- these anacoustic peaks here, their frequency depends on z times T_e . So it depends on the charge of the ions times the electron temperature. The width of these turns out to depend on the ion temperature here.

And to some extent, the ratio between the peaks and the center here is also sensitive to the ion temperature. So you can use both of these two things to get the ion temperature out. Because these are relatively low frequency modes, you can also see bulk shifts of the plasma. So if the ions are moving, your spectra will be shifted slightly.

And that shift here will be proportional to the ion bulk flow velocity. So you can measure the whole plasma moving. This ion velocity is the V in the mhg equation because the ions have much more momentum. And then finally, as we just discussed from the ratio between these two peaks-- if they are asymmetric, we can get $V_d \cdot k$ equals $J \cdot k$. That gives us $J \cdot k$.

So from the anacoustic spectrum, you can measure an awful lot. You can measure the electron temperature, the ion temperature, the ion flow velocity, and the drift between the electrons and the ions. The only thing you can't get directly from the spectra is the electron density. From that, you need the electron plasma waves. But still, there's a remarkably rich amount of information we get out of this by scattering off these anacoustic waves.

OK. So I have many pages left, and I knew I wasn't going to finish. So we discussed magnetic fields already, to the extent I said they were complicated. We vaguely discussed collisions.

The only thing I want to say about collisions here-- if you do include collisions, and you do this in the form of some-- sorry. Some collision operator acting on the distribution function inside your Vlasov equation. So it'll be about Vlasov Fokker-Planck equation. You include collisions here.

The main result I want to tell you about is that we end up usually damping coherent features. So we don't get these nice waves anymore. We just get the incoherent features. And it also further-- you might think, OK. We just get the coherent spectrum, incoherent spectrum where s of k ω is just related to the electron distribution function.

But it turns out that this is times another factor to do with collisions. So if you have collisions present, they will modify what you're measuring. And so if you don't include the collisions in your theoretical treatment, you will be interpreting the spectrum wrong. So collisions can be a problem.

As we discussed before, to be collision for Thomson scattering, you really need to have very short mean-free paths. And most of the time, people are not in the collisional regime, even if their plasma is collisional, in the sense of having a mean-free path less than the size of the plasma. OK. I like finite volume.

So a lot of the arguments I use here also apply to finite time measurements. So all of our measurements will take place in a finite volume of plasma or over a finite time scale. We'll have some plasma like this. We'll focus our laser beam through the plasma and we will collect the scattered light along the direction k_s . This is the collection volume of our spectrometer, and so we'll be looking at light scattered from this region here. So of course, one thing that could trip you up is that your f of V and r , and I guess time can have spatial variations.

So if, for example, the temperature on one side of your collection volume is different from the temperature on the other side of the collection volume, you're going to be collecting two different spectra, but they're both going to show up on top of each other on your spectrometer. So that can be quite hard to do. If it's not temperature, maybe it's something like velocity.

So if we have velocity shear, I have a plasma where across this collection volume, the velocity looks like this. And I look at the anacoustic waves, then the spectra coming from this point here would be Doppler shifted in one direction. The spectra coming from the middle of my collection volume would look like this. And the spectra coming from the other side would be Doppler shifted in the other way.

And so what you'd have to do-- what you see in your spectrometer is all of these averaged together. And maybe you end up with something like this, where the peaks are much wider than they were in your original case. If you didn't know that you had gradients in velocity within your collection volume, you would interpret this plasma as being hotter. So T looks larger than it is due to velocity gradient.

One way around this is actually by observing from multiple angles. So if you have k_{s1} but you also observe k_{s2} or maybe even k_{s3} as well, the spectra that you get out will not all be broadened in exactly the same way. And so you may be able to infer from these different angles what the true temperature is, but also what the true velocity gradient is.

But of course, that's expensive. Now you have to have multiple spectrometers. You might not want to do that. This problem gets much worse if you have a plasma which is turbulent. So instead of having a nice velocity gradient, you have a plasma where there are lots of little hydrodynamic eddies inside here.

If you collect your scattered light from over this turbulent region here, then you're going to be having all sorts of different velocity components inside. And what might start out as a nice anacoustic feature might turn out to be just some gigantic, featureless blur. And then it's very, very hard to infer anything from that.

So in general, you've got to have a nice, small collection volume so that you have smaller velocity gradients over the volume. And as you make the volume smaller and smaller, you get less scattered light. And so therefore, you get less signal, and then your signal to noise drops.

So you can't really win this. Then you'd have to get a more powerful laser. That's a problem. And we also discussed this, but I'll just mention it briefly, is our finite collection angle.

So again, if we've got some sort of plasma and we've got our laser going through it, we're generally going to be collecting the light that's coming out of here using a lens or something like that-- something with a finite solid angle. This is the ω we kept going on about it-- until it goes on to our spectrometer here. But that means that there's going to be several different $\mathbf{k}$ vectors. There's going to be $\mathbf{k}_0$ here, the one that goes to the center of our lens.

But we're also going to have light that's going to an angle, $\mathbf{k}_0 + \Delta \mathbf{k}$. And on the other side, $\mathbf{k}_0 - \Delta \mathbf{k}$. And these are all going to have slightly different vectors that will have slightly different spectra. So for example, you can imagine that one of your spectra looks like this, one of them is slightly more in the collective regime, an incoherent regime where α is smaller than the other spectra.

It looks like this. And one of them is even more in the collective regime, and it looks like this. When you sum those three together, you're also going to get a signal that looks a little bit like something you should be able to fit with your standard fitting techniques.

But you'll never be able to quite fit this. Or if you do fit it, you'll infer the wrong temperature and density. So the solution here is to make your lens as small as possible because if you make your lens nice and small, then most of the light is being scattered with the same $\mathbf{k}$ vector. And of course, if you make your lens small, then you collect less scattered light, and so then you have worse signal to noise, so then you need to buy a bigger laser. But there are no ways to win in Thomson scattering.

OK. I don't think I have anything else I really desperately need to tell you. Any questions on that?

AUDIENCE: [INAUDIBLE]

JACK HARE: Yes. If your laser gets more powerful, then the plasma starts to absorb the laser power through inverse Bremsstrahlung. And there are laser plasma instabilities, like filamentation, so there's an upper limits as well on the laser power, too, as well as upper limits on your budget for a laser.

Yeah, so it's not-- so a laser to do Thomson scattering is like a \$200,000 to \$500,000 object. And then bigger lasers are more expensive than that. So this already could end up costing more than the thing making the plasma along with [INAUDIBLE], yeah. Any other questions? Yes.

AUDIENCE: [INAUDIBLE]

JACK HARE: What's new at the moment? A lot of work on non-Maxwellian distribution functions. So I've erased it from the board, but for Maxwellian, we have this nice function, w , that people have worked out and they've tabulated it. But if you have an arbitrary non-Maxwellian function, you have to do Landau integration for every single ω over $\mathbf{k}$ -- so for every point in the complex plane.

And evaluating that function is extremely slow, so you can't really do fitting with it. So people are doing machine learning or they're doing other ways of speeding up those evaluations. People are also doing experiments when they observe the light from many, many, many angles.

And so the Thomson scattering experiment there looks like you've got a plasma, you've got some laser light going through. It gets scattered off in loads of different directions, and you put a convex mirror here. And that convex mirror focuses the light onto a spectrometer. And each different bit of the spectrometer corresponds to a different angle.

So then you're recovering the full three-dimensional Thomson scattering distribution function, or at least the two-dimensional Thomson scattering distribution function there using these very clever optics. So those are two things people are doing. It's a very active area of research.

OK. Good luck on your final, those of you who haven't submitted it yet. And as I said on Canvas, I'd be very grateful for any feedback from this evaluation.