[SQUEAKING]

[RUSTLING]

[CLICKING]

JACK HARE: So we're going to go on to a slightly new topic today. We're going to be studying Abel inversion and then talking about Faraday rotation imaging. Does anyone have any questions on all the interferometry stuff we've covered so far before we leave that part alone? Everyone seems very happy with interferometry.

All right. So interferometry-- this technique that we're going to be talking about today, Abel inversion, is actually quite general. And I'm introducing it in the context of interferometry because we often use it with this technique, but it could also be used with emission from plasma, from individuals. We'll talk a little bit about what Abel inversion is and what symmetry requirements we have, and you'll see quite quickly it's actually quite general. And it can be used in lots of different cases here.

So what we have from interferometry-- I'll just write as IF like that-- is we've got some integration of any dl. So we have some line integrated quantity. And we might have that line integrated quantity along a very specific chord. So we might be integrating along the z direction, and we might be at x equals x 0, y equals y 0.

And we might be resolving it as a function of time. That would be our temporally resolved interferometry. Or we might have an image of x and y, and we would have it at some specific time. This would be our spatially resolved interferometry.

So this is what we have. And of course, what we want-- that's what we can't have, which is the electron density as a function of position everywhere in space, preferably as a function of position of time. This is what you'd like if you want to compare a simulation or a theory or something like that. And instead, we have these line integrated measurements.

OK. So if you want to calculate this density from some of these limited reduced data sets, you could do a technique like tomography. So if you've ever gone for an MRI or something similar-- tomography-- you'll know that what they do is they'll take lots and lots of different images, slowly scanning around your head or the injured part of your body, and then they'll do some very fancy computer techniques to reconstruct the three-dimensional structure of whatever it is they're scanning. And this works extremely well because people have lifetimes of years. But plasmas only have lifetimes of seconds or milliseconds or nanoseconds, and so it's very hard to just slowly rotate your plasma in place when you've taken lots of pictures of it.

The alternative, if you want to do single-shot tomography, will be to surround your plasma with lots and lots of cameras and look at it from lots of different angles. So for example, if we've got a circular cross-section plasma like this, maybe this is some sort of automat type thing. We could just have lots of different lines of sight.

And we can do our tomographic reconstruction like this. But of course, lines of sight, LOS, are expensive. So we don't tend to be able to just have a very, very large number of them. But of course, for some applications, this might be justifiable on each so they've got 550 kilometer lines of sight for reconstructing the emission from it. So they made a choice to have a lot of imaging lines of sight. In general, for interferometry, this is too expensive, so we don't do it.

But what we can do is a version of this where we make strong arguments about the symmetry of our plasma. Because if our plasma has some underlying symmetry, it helps us need fewer lines of sight, and we can still start getting out an approximation of the full density profile. And the symmetry we're going to talk about today with Abel inversion is an assumption of cylindrical symmetry. So we're going to assume cylindrical symmetry.

And so that could be, again, something like this plasma here with a circular cross-section. And we're going to assume that maybe there's variation in this direction out of plane. That doesn't matter. We're only trying to measure it in one plane. But we're going to assume that our electron density, any of xy and z, is just a function, any of R and z, where R squared is equal to x squared plus y squared.

So that's the same as saying that the density is constant nested circular surfaces. And of course, if this is something like a tokamak, that's quite nice because we know that our flux services quantities tend to be constant. And this could also be true for a cylindrical z-pinch plasma and the sorts of experiments I do. And you can think of other situations where you think things are approximately symmetric.

And if you have a system like this, you could have a set of interferometers looking along parallel chords like this. Or if you are working with a tokamak and you can do imaging, you could have a camera and expanded laser beam, as we discussed before. And that camera is now measuring any of-- let's see, xy. That t moves to 0. Or these could be a series of chords which are measuring any at x equals x0 times like that.

So this is just two different ways to look at this problem. I can either build up my data from multiple time resolved interferometers, looking along parallel chords, or I can have an imaging system looking at a single time. In both these cases, we can do this thing called the Abel inversion. I'll just write down-- oh I wrote it down. OK, good.

Abel is one of these guys who's depressing when you read his biography. He was a Danish mathematician. He invented all sorts of wonderful things in mathematics, as well as the Abel transformation. And he died of consumption at the age of 26.

So I don't know how many of you are still younger than 26, but you've still got maybe a couple of years to make such groundbreaking discoveries. I'm already past it. I don't have a chance. So you look at his biography and you're like, damn, Abel got a lot of work done.

OK. Cool. So what we have, say, from either of these two systems is a map of the line-integrated electron density as a function of y. And I'm just going to draw this very suggestively as this very blocky setup here. So each of these densities could be the density that we've measured at a single pixel on our image or it could be the density measured by our n time-resolved interferometers here.

So we've got some density value at each of these points, like this here. And what we want is, of course, our plasma density as a function of R, not as a function of y, which is a coordinate. For example, it could be one of these two. I'm not really distinguishing between y and x here. As long as it's perpendicular to the probing direction, it doesn't matter.

But what we want is this as a function of R. And we'd like to have some nice, smooth function. And it doesn't necessarily have to have the same shape, of course, because it's clear if you look at this and you think about it for a little while that your profile in dl doesn't have to be the same as your profile in any like this.

So what we want is some mathematical formalism to allow us to take this data and produce this. It's clear how to go back the other way. You can certainly do it numerically very easily. You can just make this up. You can make up some profile like this, and then you can just calculate the line-integrated density along each line of sight.

What's less clear is how to go back the other way from the data we have to the data we want. So this is the setup of the problem. So does anyone have any questions about this so far? Any questions online?

This is what we measure. And I'm going to call this some function, F of y. And I'm going to call this some function f of R. And the reason is because it doesn't actually matter whether this is density or brightness or whatever else. The mathematics are all the same. So I'm just going to refer to them as these two different functions. And we're trying to convert one to the other.

So mathematically, we have our line-integrated function, F of y, which is equal to the integral of any exerted like this. And this is equal to the integral from minus a squared minus y squared square rooted to plus a squared plus y squared like this-- f of r. Exactly. [INAUDIBLE] very quickly.

OK. I had it right in my notes, besides the change on the fly. I hate it. We're going back to the notes. OK. And I'll draw you a diagram of the geometry here. It might be slightly different from the background of the geometry I had just previously here.

So we have some plasma which has an approximately circular shape. It's bounded at a, so we can say that the pressure at a is equal to 0. This just stops us having to integrate out to infinity, which is very inconvenient when you try and do it in reality. So we're going to stop our plasma at some distance. So we only need to make measurements up to the boundary, a, in order to solve this problem.

We've got some chords going through the plasma. We've got our coordinate system, where y is transverse to the direction that we're probing and the x now is the direction that we're probing. I was using z previously, but the way I've got it written is x. I got myself getting confused with how to do that.

And so any point inside the plasma, you can say that there is some distance, y, that it sits at from the origin here and some distance x along that it sits out. And so there's some distance, r. And this is the radial coordinate. We're assuming that we have symmetry in the azimuthal direction, in the angular direction. So we're only interested in the size of this radial coordinate here.

And then if you look at this and stare at it for long enough, you can see that this is indeed the procedure I was talking about before, where you can easily go from your-- if you come up with some calculated profile, some guess at what you think the distribution is, like a Gaussian-- how you go from that to your prediction of what you're actually going to get from your detector. So this is the easy direction. This is actually the Abel transformation. And what we're really interested in is the inverse Abel transform.

Now, we're technically requiring this condition here. I've written as pressure, but why don't we, now that we're talking in terms of these functions, I've written that I want f of a to be equal to 0. This isn't actually quite true. The rigorous requirement is that f of r as r tends to infinity. This needs to drop off or fall faster than 1 over r.

So you can get away with a Gaussian type function or something like that, as long as it falls sufficiently fast. You can't get away with something that's uniform across all space. But as long as your function falls off nice and quickly, you can use this technique here. But this is actually quite the Abel transformation.

What we do at this point here is we realize that this is a horrific mess of r's and y's and x's and things like that. And we decide that we want to rewrite everything. And we already said that r squared is equal to y squared plus x squared. And so we want to have a go at substituting this x out for something that's in terms of r instead.

And this gives us a real transformation, which is f of y is equal to 2 times the integral from y to a, f of r r dr, square root of r squared minus y squared. And if you stare at this for long enough, you can convince yourself that doing this substitution into integration will work out, and that we've correctly dealt with the limits here as well. And so this is the thing which is called the Abel transform.

OK. But we, as I mentioned, don't want the Abel transformation. That's relatively easy to do. What we have is f of y when we want f of r. So what we want to be able to do is the inverse. And I'm not going to derive this. I'm not even sure I know how to.

But if you stare at what I'm about to write down in this for long enough, you can convince yourself there's enough shared features to it that it's probably correct. And you can go look it up if you want to [INAUDIBLE] So this gives us that f of r is equal to minus 1 upon pi times the integral from r of a of dF dy. That's our capital F here-- dy y squared minus r squared. I'll just [INAUDIBLE].

Remember, this capital F is something we have as a function of y. So this is our line incentive measurement, and this is as we move to the symmetric radial dependence on the measurement. This is the thing that we're actually trying to get at.

OK. So this looks like a complete solution to the problem. If I have some measurement from my detector which is line integrated and I have enough samples in y, I should be able to work out what that involves. Can anyone spot any limitations for this procedure? There are two obvious ones.

AUDIENCE: It's not clear where the plasma ends all the time.

JACK HARE: Can you just say that again, please?

AUDIENCE: Yeah. It's not always clear what a you should choose.

JACK HARE: OK. So that's a reasonable one, actually. So where is a? That is actually a problem. I would definitely agree it's a problem. It's less of a problem as long as a drops off rapidly enough, which is related to this.

But you're right, the edge of our plasma is fuzzy. You know where the density definitely goes to 0 if you have a vacuum chamber-- the hard metal walls of your vacuum chamber. So maybe that would be good enough. But yeah, you certainly want to do this experiment to know where a is. Other limitations?

AUDIENCE: If you do the derivative. So that's going to be limited by detector resolution and then experimental noise.

JACK HARE: Yeah. So dF dy is noisy. If we go back to this very suggestive picture that I put in here deliberately like this, for any realistic system, you have discrete measurements at discrete locations. And we all know that doing derivatives of discrete data is a nightmare because you're taking something that's noisy, and you're dividing it by a small number, so any noise here gets really amplified up. So this straight away looks problematic. Third thing. Yes.

AUDIENCE: I don't know if we want to count this, but once you start getting close to the edge and do that correctly, it'll be able to have a much smaller number in the bottom.

JACK HARE: Close to the edge?

AUDIENCE: [INAUDIBLE] we're closer to. [INAUDIBLE] I was thinking--

JACK HARE: You're close, but you're not quite right. Anyone else know? What are we talking about here? We're talking about the fact that there's something interesting going on here. And your physicist eyes have seen this and gone, a-ha. Whenever we start having numbers minus other numbers in the denominator here, there's some chance that this thing will go to 0.

And it will actually go to 0 for y close to r equals 0. So near the center here, this thing will have a singularity. Now in reality, we won't have a singularity because we'll never have a detector bin that's exactly at that position. But what we will do is for the bins which are close to the center here, we'll have a very small number on the bottom. And so we will amplify the value of this big number.

So if this is noisy, if there's some noise near the center here, that noise will be massively amplified and will appear everywhere else in our solution. And it will cause problems for the rest of our solution. So I'll just write here that we've got a singularity near y equals 0, and that this amplifies the noise of these points. OK. Any questions on any of this? Yes.

AUDIENCE: Why are we collecting for the data where y is minus? y can only be positive in this picture. But numerical, you can have my y be negative.

JACK HARE: That's an excellent question. And does anyone know why we are neglecting the data that we have for y less than 0? We have assumed symmetry. In order to do this calculation, we have assumed as a mutual symmetry. And so the data must be identical for y less than 0 than y greater than 0.

Now, in reality it won't be. We never have a system which is perfectly symmetric. So the good way to present your data is to do the Abel inversion on one half of your data and the Abel inversion on the other half of your data separately, and then see whether those two match.

And if they match close enough with an experimental error, great. You've got a good inversion. If they don't match at all, then you shouldn't have used an Abel inversion in the first place. Your prior that you have this cylindrical symmetry is incorrect. So you can't use this method.

So it's a good check, actually, on your data. The other reason is because of this symmetry, if you're trying to save money, you might only put detectors in one half. If you really sure that you've got symmetry, then you don't need to check.

Maybe you do the experiment a few times with all your detectors spread out, and then you're like, hey, this is great. Now I can get higher resolution by moving half of my detectors to the first half. So there may be some reasons why you only need half the data here. But yeah, that's a really good point. Any other questions on this? Anything online?

OK. Like I said, this is very generic. This could be interferometry on a tokamak. This could be interferometry on the sorts of plasmas I work with. This could be used for unfolding a mission on a hotspot from an X-ray image or something like that. So we're just introducing it here because it's a useful technique to know.

OK. How do we actually do this in practice? Can anyone think of some way to overcome some of these limitations, particularly this one here? So we've got, again, our data, which is discrete and potentially noisy. Yeah.

AUDIENCE: We mentioned before, just take data exactly as you [INAUDIBLE] always going very close to [INAUDIBLE].

JACK HARE: Sure. That could be a problem, yes. So you could deliberately shift your data so that y equals 0 is on one side of one of your bins or something like that so it's close. Hard to get in practice because the plasma might move around, so you're probably not going to be able to do it. Yeah, other ideas.

AUDIENCE: I guess that also means that this idea might have issues. But if you have prior information about where high grades are in your plasma, you concentrate your measurements in those regions so that you're better resolving hydrating ingredients can go for it.

JACK HARE: Yeah. So certainly these high gradients are going to be dominating this integral. But of course, also any gradients near the center is going to be dominated. So you might want to put more measurements near the middle so you'd have higher fidelity there. Other techniques?

AUDIENCE: Is it possible to integrate it by parts? Therefore, you don't differentiate by inside your different [INAUDIBLE] the other part.

JACK HARE: I don't think you can do this by parts, but I haven't seen the analytical version of this, which looks like it does that. But I can't immediately check this and tell you that it doesn't. I suspect it doesn't work. Yeah. Any other ideas?

AUDIENCE: You might have some reasonable idea of the distribution ahead of time and form an expected interpolation.

JACK HARE: OK, yeah. So maybe we've got some priors about the distribution, and we fit to this. That would be good. A similar version is we could fit this noisy data with a set of basis functions that we think has the sort of information in it that describes this.

And those basis functions won't have any noise, and they'll be nice and differentiable because we'll use some nice analytical basis functions. So we could have a sum of m Gaussian functions, where the Gaussians have some position and sigma and width or something like that. And then we know what the derivatives of all of those are, so yeah.

AUDIENCE: So that's putting a set of basis functions to the brightness of emissivity? Big F or small?

JACK HARE: We'd still be fitting into this. This is the only thing we know.

AUDIENCE: OK.

JACK HARE: Yeah. So you could say F of y is equal to some weighted set of basis functions. And there are some basis functions that work really well because they've got analytical Abel transformations. Gaussians are one of them, unsurprisingly. But some functions have a nice analytical Abel version and some don't.

So you'd want to use a set of functions that have nice analytical algorithm versions. So that works pretty well. And that's what most people do. So if you go online, Python has a nice Abel inversion package, and they've got different basis functions.

And depending on your exact problem, you might want basis functions that have got more spiky features at the edge or smooth features in the middle. And so just like all of these sorts of problems, there's no one size fits all. You have to tailor it to what you're doing. Did I see another question?

AUDIENCE: Can we fold it for a low pass or something to get rid of 5 [INAUDIBLE]?

JACK HARE: Yeah. So as always with our data, we can smooth it out, but then we lose spatial resolution. So that could help us, but at some cost. And so you have to balance those things.

AUDIENCE: My first thought-- I actually think how to do this is the ideal would be get that derivative in analog directly.

JACK HARE: Right. So the idea was how to get this derivative in analog directly. I don't know how to do it either, but I think this is just a fundamental limitation when you're making one of these measurements that you can't overcome. There are some really interesting links between the Abel inversion; the radon transform, which is also used in tomography; and Fourier transform. They form some sort of weird, little cycle. If you do all of them in a row, you get back to where you started.

So there's really fun mathematics going on inside this as well, which I'm not intelligent enough to know about. But if you like that sort of thing, you should go look on the Wikipedia page. There's a lot of good stuff.

So any other questions or thoughts on Abel inversion? This is just a little aside. We're going to go on and do Faraday rotation after this, so we'll completely change topics. So if you've got any more questions, speak now.

AUDIENCE: So this condition of cylindrical symmetry is very strict. So for instance, if we had a highly shaped [INAUDIBLE] flask with diverters, et cetera, it technically doesn't-- it possesses a symmetry on flex surfaces, but those flex surfaces aren't cylindrical. So there would be no way to incorporate that information.

JACK HARE: Yeah. So the question was how do we deal with non-circular flux surfaces, like in most modern tokamaks? And yes, you can no longer use an Abel inversion in that case. But there is still sufficient symmetry, and there's a lot of symmetry. So what you would do for a tokamak is I believe you'd do a gradual runoff reconstruction of the flux surfaces from your magnetic diagnostics.

You would then know that the density is constant along the flux surface because they are surfaces of constant pressure, and there's enough motion in the toroidal direction to smooth out any density perturbations very quickly. This is roughly-- obviously there's fluctuations and stuff like that. But in general, there's this constant. And then you would use that as a prior.

And you wouldn't do this Abel inversion analytically or even semi-analytically, but you would have to feed into your tomographic reconstruction algorithm. And for any tomographic reconstruction algorithm, the more data you have, the better it is. So it's obvious. And it's obvious here as well-- if I only have four chords of interferometry, my data is so sparse that I'm going to have a bad Abel inversion.

And so you can have four chords of-- you can have your chords of interferometry crisscrossing the plasma like this, or you could still have them crossing the plasma like this. And even if your plasma was strongly shaped-- so if it was some classic single null no x point type thing like this, you could have your chords frothing in this fashion. And that would be good enough to be able to do some sort of inversion in the middle here.

And there's actually-- I think one of Anne White's students who's about to graduate who has been working on this for X-rays. And he came up with some cool ideas about what if you have some sparse chords of data and then in some section here, you have really, really fine chords? Can you combine those and measure very small turbulent fluctuations with this? And apparently, the answer is yes. So there's lots of cool things you can do with [INAUDIBLE]. Yeah.

AUDIENCE: Do people commonly [INAUDIBLE] integral formulation of the algorithm or is it more a problem to do a matrix formula where you know the width of your sight line or all of your bins in space, and then you can take the inverse of the matrix with some math massaging?

JACK HARE: I don't know what technique is more popular. When I've done this, I've tended to use fitting with basis functions and then do the inversion analytically to the fits of the basis function. But I imagine there might be good reasons for doing that matrix formulation, especially if you're trying to use this for real time control.

So the bolometers are going to be used for showing where the glowy bit of the plasma is. And we want that to stay in the middle. If it starts going somewhere else, then you want to feed back on that. And so then you need to do it quickly. And so having some sort of matrix technique will be beneficial compared to, oh, I'll carefully tune my fitting function. So you don't have time to do that. OK. Any questions online?

All right. Let's do some Faraday--

AUDIENCE: Are there ways of handling systems-- are there ways of handling systems that don't have-- I guess you already talked about this. But are there ways of handling systems that don't have any symmetry priors?

JACK HARE: Yeah, actually. I have a grad student of mine who's been working on this recently for tomographic reconstruction. So you can always do a reconstruction, but it's always poorly posed. You don't have enough information to fully reconstruct the [INAUDIBLE] points in space.

But if you have some priors, you can make some guesses. Even with just a single line of sight, you can make some guesses. And the more lines of sight you have, the more you can constrain it. And he was working on a technique, which I can't remember. The acronym is ART, and I can't remember what it stands for.

But with two orthogonal lines of sight and just a flat uniform prior-- so no information about what it looks like to start with-- he was able to reconstruct some relatively complicated shapes out of this. So there are some clever things that you can do. And I think there's a lot of cool stuff that we can take in plasma physics from other fields. So a lot of stuff in computer vision from tomography, from medical imaging, which can be used for understanding this stuff. And people do already, so there's lots of nice things out there that you can do.

At the end of the day, the more data you have, the easier this is. So if you have very little data, you have to provide that information from somewhere else, which is your intuition, your guesses about the plasma. So you can't win. You can't get just get this information for free. OK. I'm going to go on now so you have time to cover all of this.

So we're going to be looking at Faraday rotation. But we're actually just going to take a little side track for maybe most of the rest of this lecture to look again at waves in magnetized plasmas, of which Faraday rotation is one. The reason is although you hopefully have seen some of this in earlier plasma classes, I think there are lots of different ways of looking it.

And I think the way that Hutchinson has and that I've adapted is quite a nice way of looking at it. And we're also going to need lots of these results, not only for Faraday rotation, but also for reflectometry and electron cyclotron emission. So we need to know how waves propagate in a magnetized plasma. And so we may as well just review this quickly. So if you've seen this before and you're very confident, feel free to relax. And if you haven't seen it, maybe pay attention.

So remember, we had before some assumptions that our plasma was cold, and we quantified that by saying that the thermal velocity of the electrons is much less than the speed of light. We said that our frequency was high, and we quantified that by saying that our ion plasma frequency was much less than a frequency of the waves ascending through the plasma.

And that meant that the ions are effectively stationary. So we can just neglect them and just deal with the electrons. That made life much easier. And we also made this following restriction, that k dot e was equal to 0. And this was the restriction that the waves were transverse.

Now, this restriction-- the final one, the transverse waves-- made our life very simple, algebraically. But it turns out that if you try and find the waves in a magnetized plasma while keeping this restriction, you don't get all of the waves. You, in fact, have explicitly ruled out one of the most important waves-- the extraordinary mode, which is extraordinarily useful. And so we have to drop this restriction here and then deal with all the horrible consequences of that.

So now we have k dot e not equal to 0. And if you go back and you look through the derivation, and you start rederiving bits, you end up with an equation now that looks like omega squared minus c squared k squared. We had something like this before, but our equation previously was a scalar equation. Our equation now is going to be a matrix equation with these 3 by 3 matrices.

And this is just the identity matrix. And this is an even more odd looking object, which is kk, which is a dyad, which is also a 3 by 3 matrix here. You can look up more of the details of this in Hutchison's book if you haven't seen this in a while. All of this is now dotted with the electric field. And that is equal to minus I omega J over epsilon.

Now, once again, we realized that J is equal to minus e and eVe. And what we want to try and do is write this entire equation in terms of e. We want to get rid of J completely. And then, of course, this is an equation that looks like a matrix times a vector, 0.

And we know how to do this. This is what we've trained our whole lives for. We take the determinant of this. We find the mode to search by [INAUDIBLE]. That's great. So we really love these sorts of equations. And so we're trying to make this equation look like one of those.

So now we need an equation of motion for Ve. Previously, remember, we just looked at the response of the electrons to the electric field. But now we want to have the magnetic field as well. So we have any d Ve et is equal to minus e times the electric field plus V cross b. So this is just the Lorentz force, but we've now got the magnetic field in here.

And we're going to assume that our magnetic field is the first order, is just some static field. And we're going to point it in the z direction. Of course, I can point in any direction we want. I've chosen z in this case here. And then from this, you get out a series of equations for the velocity. And I think you've all seen this, so I'm actually not going to go through this line by line.

Well, I might go through. OK. I'm not going to write this equation out in terms of vector components, like I had in the notes. I will make the point that, as before, we're going to say that we're going to assume that V equals Ve 0 exponential of I k dot x minus omega t. And that allows us to replace d dt with minus I omega.

So we turn to this differential equation into an algebraic equation. And then we find that the velocity of our single electron here is going to have a structure with some nice symmetry due to the magnetic field. So dx and dy are going to look very similar. They're going to look like minus Ie over omega me times 1 over 1 minus omega squared over omega squared. It doesn't work so well when I say it like that. Capital omega squared over lowercase omega squared.

And this is going to be equal to the electric field in the x direction minus I capital omega over lowercase omega electric field in the y direction, where we've defined here capital omega as the cyclotron frequency, e B0 over me. So dy, we have all of these terms again. But in the brackets, we have something that looks similar, but slightly different. I capital omega over lowercase omega ex plus ey.

And when you look at this and you squint and you calculate dx squared plus dy squared, you find out that this is just-- when you just look at how these work, this is, of course, how particles are circulating around the magnetic field line. Remember, the field line is going in z direction. And so these two components are just the spiraling dx and dy here.

And finally, we have the z components. And that's very simple. That's minus Ie over omega me [INAUDIBLE] like that. OK. You've seen all this before.

You can then take this, make it into a nice vector, substitute it back into this equation for J, substitute J back into here. And you see very quickly that all you're going to have left are things like omegas and c's and k's and capital omegas and this electric field. So we're going to have some matrix equation dot e is equal to.

And that matrix-- well, we can write J is equal to minus e and b times the velocity. And we can write that in terms of some sort of conductivity tensor times the electric field, where that conductivity is this monster I, any e squared over me omega 1 over 1 minus capital omega squared lowercase omega squared times pi on a nice, big 3 by 3 matrix. If you want, down the diagonal.

[INAUDIBLE] here. And minus I omega in here, pointing to off-diagonal elements, then 0s elsewhere. If you're surprised that seeing something interesting involving the magnetic field show up in the zz component of this tensor, it's literally only there to cancel it out here. So in fact, it doesn't exist at all. It's just it's more convenient than writing this factor underneath all four of these terms. So this is just a symbol. So this is the conductivity for our particle.

Which means we can then write out in short form that omega squared minus c squared k squared times the identity matrix minus this dyad, ak, plus I omega over epsilon null sigma dot e equals [INAUDIBLE].

And I like writing things in terms of this conductivity term. But if you like things in terms of e-- I'm sorry, in terms of epsilon, the dielectric tensor, you could rewrite this. And you will get Hutchinson's equation, 4.1.2.

OK. So the magnetic field in the system breaks the symmetry. So we have to treat that differently from dx and dy. But we don't have to treat dx and dy the same. We can pick our orientation of our x and y axes to make our life simple in what's going to follow. And we're going to do that now. I've got just about enough space here.

So we're going to have a coordinate system like this, with z pointing upwards. And that's the direction of the magnetic field. We're going to have this coordinate, y, and a coordinate, x, like this. And I'm going to choose our k vector to always be in the yz plane, with some angle of theta to the z-axis here.

So we're effectively choosing our coordinate system such that we can write k is equal to the size of the vector, k, times 0 sine theta, cosine of theta. This makes our life much, much easier and more flawless. But even if you do that, when you go to solve this equation in full generality, you have all these sines and cosines. It's an ungodly mess. So this is absolutely horrible. So you still get something terrible, which is equation 4.1.24, which I think is called the Appleton-Hartree dispersion relationship.

And it's a mess, so it's very, very hard to work with. And the trick is that no one actually works with it most of the time. We just work in the case where we have the theta equal to 0 or theta equal to pi upon 2. So those are the two cases that I'm going to tell you about-- when we've got waves propagating along the magnetic field or perpendicular to the magnetic field. If you're ever in the unfortunate case of having to do something in between, you'll have to go back to this equation and work it out yourself. But at least in the limiting cases, the behavior is slightly easier to understand. So those are the cases that I'm going to go through now. So that was a bit of a whistlestop tour. Any questions?

Good. Everyone loves waves [INAUDIBLE].

So we're going to start with the theta equals 0 case. Am I? Why? No chance. Probably makes sense. That was a z equals pi over 2 case.

So this case is maybe particularly relevant if you're trying to diagnose something like a tokamak, which many of you are. So if we draw our tokamak looking from above, we have magnetic field lines going around like this. And remember, the toroidal magnetic field of the tokamak is very strong. And so really, although the magnetic field lines are slightly not like this, they really are very much just circles around the top.

And where do we put our diagnostics? Well, we can't put them on the inside, usually. We don't want to put them at some weird line of sight like this because that would mean integrating through things where we're not really sure about the symmetry. We're very likely to put our diagnostics like this on a line of sight, which is indeed perpendicular at 90 degrees at the local magnetic field. That might be because you've got to fit through some gaps between the magnets or simply because it's very easy.

And almost every tokamak I've seen has diagnostic designed to look along this line of sight. If you have something that looks at a weird angle, that's very unusual. So it's very relevant to ask how the waves propagate in a magnetized plasma like a tokamak perpendicular to the magnetic field. And what you get are two different modes here. We have one mode where the dispersion relationship is very familiar. n squared is equal to 1 minus omega p squared upon c squared-- oh, omega squared.

So this is just the wave that we found in unmagnetized plasmas. What's interesting is that although we've done all of that mathematics, there is one wave which propagates the perpendicular magnetic field, which looks like the magnetic field doesn't matter at all. And so we call this wave the O mode, where O stands for Order.

All right. So in the ordinary mode, remember, we can then go back. This is our eigenvalue. We can work out what the eigenmode is in terms of dx and dy and dz. And we find that dx is equal to dy is equal to 0 here. And so all of our electric field is in the z direction.

Just remember we have this diagram where we have z like this, y like this, and x like this. I've restricted my k vector to be in the yz plane. I've now set theta to be pi upon 2, and so therefore, k is pointing in this direction, which means that our electric field is pointing purely in the z direction.

And this actually gives us a hint why the mode dispersion relationship doesn't seem to know anything about the magnetic field. And that's because the electrons are traveling along in the z direction. And the electric field, as it oscillates up and down, is simply accelerating or decelerating them along the magnetic field.

And so it doesn't have any effect from the gyrating particle orbits. You just have particles which are going like this. And maybe they're being slightly accelerated or slightly decelerated, but it's all in the d direction, so it doesn't have an interaction with the magnetic field whatsoever.

So these are nice and easy. And the nice thing about these modes, actually, is that they are transverse. So k dot e is equal to 0. So although we relax that condition, we obviously didn't need it in order to get this mode, as you can see, because we already got the mode before when we did have that transverse condition. So this is our nice, easy wave. The next one is not ordered.

So the next one has the dispersion relationship that looks like this. Now, in Hutchinson's book, he rewrites some of these terms as x and y and things like that to make it more compact, which I think is great if you're going to write it a lot. But just in this case, I want to write it out in its full generality, in terms of things like the plasma frequency and stuff like that so that you can see where all of those terms come from. So it's going to look a little bit more complicated than what you get in Hutchinson's book, but I think it's more useful.

So it's 1 minus omega p squared over omega squared. Looks good, but then it's actually times 1 minus omega p squared over omega squared. And all of those-- not the 1. All the rest of these are over 1 minus omega p squared over omega squared minus capital omega squared over lowercase omega squared.

So this is a little bit Escher-esque. They're bits of the same thing repeated in a fractal pattern. And the more you stare at it, the more you think, I wonder what plasma is playing at going with that. It seems very complicated.

And remember, this is the simplified version, where we've already taken theta equals pi upon 2. If you want to put in all the cosines and sines, it becomes much more complicated. So straight away, we can see that this is more than ordinary. And so this is, of course, the x mode, and it is the extraordinary.

Now, you can probably guess just by looking at this that when we substitute this back in to our equation and we try and get out the eigenmodes of the system, they're not going to be quite as simple as this. And you are quite right. So what we find out is that dx over dy-- so the x and the y components of the electric field are related to each other. And the relationship between them is minus I times 1 minus omega p squared over omega squared minus capital omega squared over omega squared.

All of that is over omega p squared over omega squared, capital omega. Then I've got omega like this. Although that looks complicated, fortunately for us, the z component of the electric field is, in fact, 0. So that saves us a little bit of hassle.

So if I plot out now y and z, x like this, I still get my magnetic field in this direction. And I've still got my k vector in this direction. It's got exactly the same k.

Can anyone tell me what the electric field looks like here? Previously, we said the electric field was just pointing in the z direction because it was. But now it looks a little bit more complicated. Does anyone [INAUDIBLE]? Yes.

AUDIENCE: An ellipse?

JACK HARE: So the answer was an ellipse. Yes. Do you want to be more specific?

AUDIENCE: In the xy plane. They'll be out of phase by [INAUDIBLE].

JACK HARE: OK. So we're talking about an ellipse in the xy plane. When we have an ellipse, we have a minor axis and a major axis. What is the orientation of the major axis, with respect to this xy coordinate system? That's effectively asking you, is ex bigger than ey or is ey bigger than ex?

And we made some very strong assumptions in deriving this that will help you work that out. In particular, we assumed that this was a high frequency wave. Any answers online? You've got a 50/50 chance of being right. I hope everyone realizes that it's not going to be wildly at some random angle.

AUDIENCE: If we make the frequency really large in there, then it would be basically 1 divided by-- so ex is way bigger.

JACK HARE: Yeah, ex is way bigger. So if we have a high frequency, this term is going to be very, very small, which means ex is much bigger than ey. And so we have an ellipse that is extended like this.

So ex, much bigger than ey. And because there's this I inside here, they're actually related to each other in a complex fashion. So it means the electric field vector is going to be sweeping out this ellipse as the wave propagates.

OK. Now, the main thing to note here is that this wave is not transverse. So k dot e is not equal to 0 because of course, ky is parallel to our small but existent ey. So this is the wave we would not have got if we insisted on only looking for transverse waves, which is why you have to go back and rederive it all with this.

Of course, we also put in the equation of motion for particles in the magnetic field. But you can see why it's much easier to go and derive the ordinary mode in an unmagnetized plasma than it is to get these two modes in a magnetized plasma straight away. So this is a very fun way indeed.

There is a question in Hutchinson's book which caused me a lot of thought as a grad student because I didn't understand what the hell he was getting at. And so I'll put it to you now, and we shall see whether you can spot it straight away or not. The question is in the books.

You have set up an interferometer looking across the plasma like this. If you have accidentally set up your interferometer to measure the x mode rather than the o mode, what would the error be in your measurements? And you give some plasma parameters so you can calculate it.

So remember, in our interferometer, what we measure is not density, but we measure changes in refractive index. And so you're saying, if you set up your interferometer such it measures this refractive index, how different would your result be than if you were measuring this refractive index? And my question I was always asking-- how the hell do you even set it up to do one or the other? So perhaps we can work it out together. How can I get to choose whether I'm using the o mode or the x mode?

AUDIENCE: Well, for o mode, you have a z chord, so you could polarize your light coming in, right?

JACK HARE: Yeah. So if I'm injecting light into my interferometer-- and maybe it's collecting here and bounces back to that. So if I have this polarization, that's the o mode because the polarization, e, is parallel to b. And if I have this polarization out of the page like that, that's the x mode.

So you get to choose which refractive index you probe by your choice of the polarization that you're injecting. And that was the bit that I missed as a grad student. It took me until I was teaching this class for the first time to finally work out what's going on-- is that you actually do have a choice.

It's not like the plasma decides for you, which is why I was like, what will the plasma want to do? But that doesn't matter in this case. In this case, you have a choice. The plasma is not in control of you.

And in general, you could launch a wave through the plasma at some arbitrary polarization at 45 degrees. And then because you can always decompose your wave into various sets of modes-- but this is a good set of modes which is valid for theta equals pi upon 2. Then you would see that one of your polarizations would travel faster than the other polarization because it would have a different refractive index. And so you'd see all sorts of funky effects going on that may be very hard to interpret.

So it's definitely worth thinking about the polarization. And when I went looking recently for some papers on how people actually do this in tokamaks, they spend an awful lot of time thinking about the polarization. And often, they have the ability to switch it between x mode and o mode so they can do different measurements. Yes.

AUDIENCE: The x mode should be polarized.

JACK HARE: The x mode should be polarized. So yeah, out of the page. Because you're primarily going to be exciting this large ex. And the plasma will do the work of giving you the ey, and it will do that because of the electrons gyrating around the field lines.

So you don't have to worry about it. ey is truly fantastically small here. So it's OK. You don't have to give it that yourself, which would be impossible because in free space, waves are only transverse. You can't launch a wave in free space that has a polarization along the direction of propagation.

But it's OK. The plasma has got your back there. And these are the only two waves-- the only two electromagnetic high frequency cold magnetized waves which can propagate in a plasma. So if you hit the plasma of some wave, it will instantly convert into some mixture of these two, because they're the only modes which are supported by the plasma media. Yeah.

AUDIENCE: When that converting step is happening, are you at all worried about reflected power than being different-- because if you're injecting a wave into the plasma, that's a vacuum wave. And so [INAUDIBLE] x. I'd imagine there's some chance of that vacuum wave reflecting back on you. Would that hurt your signal to noise ratio in some way? Or is there a very small amount reflects back so it doesn't really [INAUDIBLE]?

JACK HARE: So the question is when you have this conversion from the vacuum wave into the plasma waves, is there any power which is reflected? And the answer is I don't know. I would imagine that there could be some power reflected in that interface there.

And yeah, also an important thing to notice is that on a tokamak, the magnetic field stays in the same direction. But if you're in something like an RFP-- Reversed Field Pinch, where the magnetic field rotates, your wave will invert. You may launch in o mode and then may convert yourself to x mode.

And intermediate, you'll have all those intermediate polarizations and k vectors, with respect to the magnetic field. And you'll have to go back and do the Appleton-Hartree formalism. And that's probably why people don't work on RFPs anymore, because they're extremely difficult to work with.

Now, one thing I will note is you might ask know, Jack, you work on z-pinch and they've got strong magnetic fields. And yet, you don't seem to be worried at all about the magnetic field or polarizing your beam or anything like that. And the reason for that is when you start looking at the frequency orderings in the sorts of plasmas I work with, this term here is very, very, very, very, very, very, very, very small. And so this cancels out that.

And huzzah, you just have two o modes, so you're basically back to the unmagnetized plasma. So depending on your plasma regime, you may not be sensitive to the field anyway. And basically, that's a requirement that the probing frequency of your wave is much, much higher than the gyro frequency, which is the unmagnetized condition we wrote down all the way back when we derived the unmagnetized waves.

So you may not be sensitive to this. It just turns out in the tokamak, you tend to be in a regime where this is important. So we'll talk a little bit about that when we talk about electron cyclotron emission. OK. Any questions on this?

AUDIENCE: You noted just a minute ago there's a number of assumptions that have gone into derivation, chief among them being low temperature. Most of today's modern tokamaks are operating 1 to 10 keV ion temperature. It does not seem to me like something that needs a low temperature requirement.

So what is the-- maximum sequence theory yet applied to a diagnostic system, is this valid? Do we still meet the thermal?

JACK HARE: Yeah, so when we say cold, we're just talking about compared to the speed of light. And that's pretty fast. So even when you've got-- so you can't neglect relativistic effects completely for electrons at 10 keV. They've got over 20%-- some decent fraction of their rest mass at 500 keV. It's not less than that.

So in some cases, relativistic effects will be important. And we'll come across those when we do cyclotron radiation. I believe that for interferometry, we don't have to worry about relativistic effects here, and this holds well enough. And the corrections that you get will be on the order of the ratio between the thermal velocity and the speed of light, and that is a small number. So I think there's a really good point. But this actually-- this holds very, very well.

We're not dealing with plasma waves, the waves inside a plasma which are being generated, where thermal effects are important, like Langmuir waves. We're dealing with these high frequency electromagnetic waves, which are traveling so fast that their phase velocity is close to the speed of light. Really, sorry-- I shouldn't say c here. This is the phase velocity. It's just the phase velocity is very close to the speed of light for these electromagnetic waves, so c is close enough.

So this is the actual condition. For the waves inside your plasma, where you deal with hot plasma effects and Landau damping and all sorts of fun things like that, that phase velocity is much lower. And so the phase velocity is close to the thermal velocity.

And then you have the interaction between the wave and the distribution function of plasma that gives you Landau damping, all that fun stuff. We're nowhere near that. But that's a really good question. OK. Any other questions?

OK. Now we'll do the interface. This has nothing to do with the diagram. But it's interesting, and we'll use it twice again. So now we're going to do the case where the wave is propagating along the magnetic field here.

OK. So this was maybe a harder case to get inside the tokamak, but it's a very easy case to get inside, for example, the z-pinch or many other systems like that. So for example, if you have a z-pinch, some wobbly plasma like this, it's got current like this. And then it's got magnetic field like this. If I fire a laser beam through this, there's going to be at least parts of the laser beam which are parallel or anti-parallel-- you get the same result-- to the magnetic field.

And in fact, we'll show that this result that we derived for theta equals 0 applies to almost every angle between 0 and pi over 2. And the theory only breaks down very, very close to pi over 2. So in fact, this still works even when, for example, here, my k is in the same direction but my magnetic field is bent. And there's quite a large theta between here. So this is very useful.

OK. We go and we plug theta equals 0 into the Appleton-Hartree relationship, and then we solve the determinant and we get out our eigenvalues and our eigenmodes. And we get, first of all-- well, we've got two nodes. And these nodes I tend to write with a plus and a minus and write the two of them together because they're very, very similar. And so this mode plus and minus-- that's the names of the two modes-- is equal to 1 minus omega p squared upon omega squared over 1 plus and minus-- this is the difference between the two-- capital omega over lowercase omega.

So the first thing that we notice, of course, is that for capital omega and the lowercase omega much less than 1, which was our unmagnetized condition, we just reduce back to our magnetized dispersion relation. So that's good. We haven't introduced anything funky in the math.

OK. And then when we solve to get the eigenmodes, we find out that we have ex over ey is equal to plus or minus I here. Let me just draw again the geometry of the system. That's upwards y like this, x like this. Our magnetic field is in the z direction.

And now our k vector, theta equals 0, is also in the z direction. And now we have the ex and ey lying in [? k. ?] And ez is equal to 0. So this is, again, a transverse wave. k dot equals 0, like that.

So does anyone know what this wave is? The electric and the two components of the electric field appear to be out of phase from each other by a factor of pi upon 2, plus or minus sign.

AUDIENCE: Circularly polarized?

JACK HARE: Circularly polarized, yes. Everyone in the room is doing this. I guess they wanted to say circularly polarized. OK. So what we can do is we can say ex is going to go to exponential of I k dot x. In reality, this is just [INAUDIBLE] because we know which way our [INAUDIBLE] here. Minus omega t and ey is exactly the same, plus an extra factor of pi upon 2 inside the brackets.

OK. So if we plot in either space or time-- it doesn't matter-- we can either fix ourselves in one place and watch a wave go by or we can attach ourselves to the wave and see how it changes in space. We're going to get at two oscillating fields. The electric field in x will look like this. The electric field in y will look like-- I'm trying very hard to do this properly-- that.

And if you then look at what in the xy plane the electric field is doing, and you stop at this point here and then at this point here and then at this point here and then at this point here, and ask, which direction is the electric field pointing? Well, first of all, it's going to be pointing entirely in the y direction. And then it's going to be pointing entirely in the x direction.

And so we can see that the electric field vector traces out a circle. And this is what we call circular polarization. And often, these two waves, which I've been calling plus or minus here-- you might call them the right hand and the left hand circularly polarized wave. And I can never remember the convention, but I think it's right hand rule with a k vector.

And it's like, is it going this way or is it going the other way like that? I never really mind too much about which way around it is. But there is a convention about whether it's right or left, and that's what this plus and negative sign really mean here. We've got two modes, one of which is going clockwise and one of which is going counterclockwise.

OK. Any questions on that? We're going to use that in a moment. There's not really anything particularly profound at this point, but we will get on to something profound in a moment. Any questions on that before?

OK. Just so I can keep that up on the board, we're going to get into this stuff. Forget about that until we get on to reflectometry.

So now, finally, we are fulfilling the promise we started the lecture with when we were talking about Faraday rotation. And the main point about Faraday rotation is that your magnetic fields cause a phenomena called birefringence. Has anyone come across birefringence studying optics before and can give us a concise definition? Yes. Either one of you. You can say it in unison.

AUDIENCE: 2, 3-- [INAUDIBLE] refraction [INAUDIBLE], which [INAUDIBLE].

JACK HARE: Its object's index of refraction is different, depending on the direction of propagation.

AUDIENCE: Yeah.

JACK HARE: No.

AUDIENCE: Darn. OK.

JACK HARE: Was that your guess, too?

AUDIENCE: No.

JACK HARE: That is a really interesting phenomena, and it definitely happens, but it's not this.

AUDIENCE: I'm trying to make sure these words don't mean the same thing. That the angle of deflection depends on your frequency, which [INAUDIBLE].

JACK HARE: OK. So that was the angle of deflection depends on the frequency. No, that's also not true. But that's not a good definition of birefringence. Anyone got any thoughts about what birefringence could be related to? Anything on this board here that makes you think? Anyone online?

AUDIENCE: I'll give it a try.

JACK HARE: Sorry?

AUDIENCE: I'll give it a try.

JACK HARE: Yeah, please.

AUDIENCE: Birefringence is that there are different refractive indexes for different directions-- not different directions, different components of the material. So it's like a tensor, and then the different components have different refractive indexes.

JACK HARE: Yeah, this is very, very close. So it's different refractive indexes for different polarizations for different directions of the electric field. So everyone had some thought about direction in there, and that was all good. It's not to do with frequency, though of course, this is frequency dependent. But of course, the refractive index for plasma is always frequency dependent.

So that's not a unique thing for this solution here. The unique thing about it is that these waves have got different speeds. So the other thing, the x mode and the o mode-- those are also birefringence as well. They're just not birefringence in a particularly useful way. This is birefringence in a way that we can exploit.

And if you ever want to go down a rabbit hole of interesting stuff, there were several Viking variants where they found, from these Viking warlords, buried with lumps of calcite. And people were like, why have they got calcite? It's not a particularly pretty crystal. It's transparent.

And you can't really make jewelry out of it. It's quartzy. And it wasn't carved in any case. It was just this lump of calcite. And there is a belief unproven amongst archeologists that this calcite, which is a birefringent material, could be used to navigate on a cloudy day.

So when you're sailing across the Atlantic or the North Sea to raid some poor monastery and it's all cloudy-- like, damn. I don't know how to get to this monastery. And they didn't have compasses in those days, at least the Vikings didn't. The Chinese did.

So with compasses, this birefringent crystal is very interesting, because you might know that the light from the sun is polarized by scattering. And so although that light isn't directly hitting your ship in the fog, some of that light is getting through. And although the fog is messing with that polarization, there's still going to be some overall polarization there.

And by rotating this birefringent crystal in the light that you're getting that's slightly polarized from the sun and looking at markings on a bit of stone, when you rotate the crystal just so, the light is going to come through and the two markings will line up. And then you will know roughly where north is, and then you can go and sail and raid your monastery. So I'm not saying it's true. You can go read some really cool papers on people trying to do this.

People have tried to go out in a boat on a cloudy day and navigate like this, and it did not go well. But they didn't have as much practice as the Vikings did because we have GPS these days. And so we don't worry about such things. But it's really cool. So if you like optics, go look up Viking sunstones, they call them. I always like rings.

OK. Back to plasmas. So we said we have these two modes inside the plasma. These are the right-hand and left-hand circularly polarized modes, which I'm going to refer to as plus and minus like this. So plus is the clockwise going mode, and minus is the anti-clockwise going mode.

If you're confused about why I'm confused about this convention, it's because is it clockwise as you look down the ray of light? Or is it clockwise as you look towards the ray of light? Those are different. And I can never remember which one the convention is for.

I think it probably should be as you look down the ray of light, but who knows? And so maybe I've got this the wrong way around. If I got the wrong way around, you just swap where I'm looking down the red light or at the red light or whether I'm looking down the red light. And then I'll be writing it like so.

Note, by the way, that this does depend on the direction of the magnetic field. This is not capital omega squared. It's just omega. So the direction of the magnetic field changes. Which of these modes propagates faster?

If you're propagating along the magnetic field, one of your modes is faster than the other. If you're propagating against the magnet field, the other mode is faster. That's very important. That's going to be what we use to measure both the magnetic field's amplitude and its direction.

And we can use this neat technique called Stokes vectors, where a Stokes vector is just a vector of the ex ey, normalized by the sum of ex squared plus ey squared. These Stokes vectors make our life very easy when we want to do the calculations-- same guy as Stokes' theorem.

And we can say, looking at this relation between these two, that the e plus is going to have something triangular called the vector r. And that's going to be 1 over I. Oh, I lied. I'm not going to normalize all of them. I can't be bothered to put a square root 2 in front of this. So just remember there should be square root 2. It doesn't really matter.

And then the left vector is going to be 1 minus I. And you can see that these have the same relationships between ex and ey. And so these are the Stokes vectors for the right and left circularly polarized light. We also can write some other polarizations in this Stokes vector notation.

So if we were polarized entirely in the x direction here so y equals 0, this would be the vector, x. And that would be equal to 1, 0. And if we have the x equal to 0, this would be the vector capital Y. This would be equal to 0, 1, like that.

Now, there are two modes inside the plasma. We can write any arbitrary polarization of our wave as a sum of these two basis vectors, effectively. And so we can switch basis vectors if we want to. So generally, when we're launching light, we don't launch it as a circular polarization. That's quite hard to come by.

We launch it with some linear polarization. But this linear polarization is made up of these two circular polarizations. So for example, this x linear polarization is r. What's l over 2? The y polarization is r minus l over 2.

I don't know why I got this again. Oh, well. OK. Maybe it's a good point. So just to draw our geometry another time like this, we've got some k vector, which is an angle, theta, to the magnetic field, b, like that. And again, for theta equals 0, this is our dispersion relationship.

I just want to point out one slightly funny-- I think it's funny-- thing about this. So if we have k in this direction, that means that this wave is a parallel propagation. But of course, it's still a transverse wind. I think people often get this very confused because the words parallel and perpendicular and transverse and longitudinal have similar meanings. So this is a parallel and transverse wave.

And there's one more word as well which means something similar. I can't remember which one it is. But yeah, there's lots of different ways. Yeah.

AUDIENCE: For the r minus l over 2, wouldn't that be the I in the bottom?

JACK HARE: Oh, yeah. I had meant to. There's meant to be an I in here. And it's going to be like that. Or maybe minus that, but you get the idea. No, it's that. Cool.

OK. It turns out I derived this for theta equals pi over 2. I didn't derive it. I just showed you it for pi over 2 and theta equals 0. But it actually turns out that the theta equals 0 case applies over a huge range of different conditions here.

So in fact, this theta equals 0, which we call the quasi-parallel case, it isn't good simply for theta equals 0. It's good for capital omega over omega. The secant of theta is much, much less than 1. And it turns out that for some reasonable values of omega here and here, this can apply for theta almost to pi upon 2.

So you can use this dispersion relationship as I said up here, even for relatively large angles. The waves will propagate as if they had this dispersion relationship, or at least extremely close to it. And that's very, very convenient. They will propagate without dispersion relationship as long as you write your b to be the component parallel to propagation. So you replace b with b 0 cosine of theta.

So now the dispersion relationship here, which has capital omega in, is to do with the projection of the magnetic field along your wave. So effectively, the wave fields components of the magnetic field in the direction it's traveling and ignores that other component until it gets very, very close to pi upon 2, when the components along the direction of travel is almost 0. And then we suddenly see this pi over 2 condition here. So this is very, very useful because it means that we only need to have-- for some cylindrical object like this, we can use this dispersion relationship for almost the entire region that we're probing for Faraday rotation.

OK. I'm well aware that I've gone over time, so I'm going to leave it here. And we will get on to exactly how we exploit this interesting mathematics and these Stokes vectors in order to measure magnetic fields in the next lecture. So thank you very much.

I'll see you-- oh, not on Tuesday because you all have a student holiday. So for the people in Columbia who are unfamiliar with this idea, we have Monday off. And then to recover from the three-day weekend, the students have a second day off from classes. So I will see you on Thursday.