[SQUEAKING]

[RUSTLING]

[CLICKING]

JACK HARE: So as you remember, in the last lecture, we were looking at waves in magnetized plasma. So I'm going to give a quick recap of what we learned, and then I'm going to go on and show you how to measure magnetic fields using the Faraday effect [INAUDIBLE] plasma.

So the geometry that we were using, we had our z-coordinates like this. You have our y-coordinate in it, and we have an x-coordinate. And we're looking at systems where there's some static background or slowly varying background magnetic field, B, which is in the z-direction. And we realized that we can just rotate our coordinate system so that the k vector for our wave lies in the z-y plane like this, and then it's only specified by this angle, theta, with respect to the z-axis and, therefore, with respect to the magnetic field. So we don't have to worry about the components in x.

And we churned our way-- well, we didn't actually churn our way through a lot of algebra, but we could have churned our way through a lot of algebra. And you can go have a look at it. And we found that we had to solve the determinant of this big matrix here, omega squared minus C squared k squared times the identity-- I'll write it with two arrows, OK-- minus the dyad formed by two k vectors plus I omega over epsilon 0 times this tensor, sigma, which contains all the information about the conductivity of our plasma. And it's a 3 by 3 matrix because the plasma is anisotropic due to this magnetic field.

So we want the determinant of this is equal to 0. And that will give us all of the modes or the eigenvalues of our waves, and we can substitute them back in and get out the eigenmodes. So in general, this is very, very complicated to solve. If you want to do it at some arbitrary angle theta, it's a lot of work. But we decided we would simplify down and focus on two angles, which were most likely that you're going to come across, and they demonstrate the physics very nicely.

So first of all, we looked at theta equals pi upon 2. This is a particularly good one if you're working, for example, with tokamaks because if you've got your cross-section of your tokamak like this with magnetic fields predominantly coming B toroidal out of the page, and you are probing from, say, the outside or the inside of your plasma, you tend to be in a situation where your k vector is perpendicular to your magnetic field. So this applies quite a lot of the time for tokamak diagnostics or stellarators or things like that.

And we found that in this case, there were two modes. There was a mode, which we called the O-mode, where O stands for Ordinary. And this had our standard dispersion relationship, N squared equals 1 minus omega p squared over omega squared. And we've seen this several times before from unmagnetized plasma. This corresponds to a mode where the electrons don't actually feel the magnetic field, and so we just get the unmagnetized plasma result.

And we also had the X-mode, which had a rather more complicated dispersion relationship, N squared equals 1 minus omega p squared over omega squared times 1 minus omega p squared over omega squared. And that term here is over 1 minus omega p squared over omega squared minus capital omega squared over lowercase omega squared, where that capital omega is the cyclotron frequency for the electrons.

And we said that if we want to choose one of these modes for doing, for example, interferometry we need to select the polarization because these two modes, when we substitute these eigenvalues back in, we get out different eigenmodes, and they've got different polarizations. And so you can select X or O with your polarization. If you don't know your polarization, you might get too confused, and that leads to some further ambiguity when you're doing something like interferometry.

So this is very interesting. We will come back to this a little bit when we discuss electron cyclotron emission. It doesn't actually help you measure magnetic fields necessarily, though you can see there is a bit of a magnetic field effect here. So maybe if you had an interferometer that measured both O-mode and X-mode, like two interferometers with cross polarizations, the difference in refractive indexes here might give you some measure of capital omega, and therefore the magnetic fields. I don't know whether people generally do that.

The one that people generally do is the other mode, which is where theta equals 0. So we're going to take theta equals 0 here. And here, we found that we have, again, two modes which we called plus and minus. And because they both have a very similar dispersion relationship, I'm just writing the two of them together, so N plus or N minus is equal to 1 minus omega p squared upon omega squared. All of that is over 1 plus or minus capital omega over lowercase omega, like that.

So in the case where capital omega goes to 0, where the magnetic field goes to 0, you just recover your unmagnetized dispersion relationship, the same as your O-mode. But otherwise, we see that these two modes in the plasma have different refractive indexes, and so they have different phase velocities. And the phase velocity also depend on the sign of the magnetic field. And so we would think that by measuring something to do with these waves, we should be able to get out the magnitude and the sine of that magnetic field.

Now, the neat thing about this dispersion relationship that we started talking about last lecture is it actually works for a wide range of angles, not just theta equals 0. But it works for quite a large number of angles where theta is less than pi upon 2. So this dispersion relationship is approximately correct all the way up until the point where this dispersion relationship is correct.

And the formal criteria is that capital omega over lowercase omega secant of theta must be much, much less than 1. But if you have a system where your cyclotron frequency is much less than your probing frequency, which for some plasmas is very easy to arrange, then you're going to end up fulfilling this criteria for a huge range of theta. It doesn't give you exactly this dispersion relationship. You have to modify it slightly.

Instead, here, this omega, is defined as eB over m e. That's our standard cyclotron frequency. We instead have to use B parallel, which is B in the direction of our probing laser, which is B0 cos theta, looking back at that geometry over there. And then the omega that we actually use is equal to e B parallel over m e.

So you get the same dispersion relationship, but you don't have omega which corresponds to the total magnetic field. You have an omega which corresponds to the projection of the magnetic field in the direction that your wave is propagating. And we substitute this back in, and we find out what the modes are here. And we find out that the modes are related so the polarization in the x-direction and the polarization in the y-direction is equal to plus or minus i.

So that means the electric field in x, the electric field and y is 90 degrees out of phase to each other, either leading or lagging. And we talked about using these things called Stokes vectors, which allow us to write the electric field in x and the electric field of y just as a compact vector here. And for the right-hand polarization, which maybe is the plus 1 here, this Stokes vector was 1 and i, And for the left-hand polarization, this was 1 and minus i. So this is plus and negative, like that.

Now, these are two basis vectors. But obviously, when we prepare laser beams or other permanent radiation, we tend to just have linearly polarized waves rather than circularly polarized ones. And these linearly polarized ones have Stokes vectors that look like, for example, x polarized linearly in the x-direction is 1, 0, and y polarized linearly in the y-direction is 0, 1. But we can just rewrite them in terms of these right- and left-hand circularly polarized waves as R plus L over 2 and R minus L over 2i, like that.

And if we have some arbitrary fertilization, so if we just are dealing with some arbitrary polarization P, we could write that as some constant times R and another constant times L. And equivalently, that would be a different constant times X plus another different constant times Y, which is just to say that we could write some arbitrary vector as a sum of some other basis vectors.

And it doesn't matter. We can switch between these basis vectors depending on which is more convenient for our calculations. And it turns out that we will start off using these X, Y basis vectors, but then we will write things in terms of R and L because those are the natural basis vectors for this dispersion relationship here. So that's where we got to last time. So any questions?

Now, it's often the case that we can make an approximation to this dispersion relationship. And the approximation we'd like to do is the same one we do for interferometry. We'd like to have this dispersion relationship be linear in the quantities that we're trying to measure here.

And so we can say, if we happen to be in the case where capital omega over lowercase omega is much less than 1 so that our gyrotron frequency is much less than our probing frequency, then we can say that our dispersion relationship is going to look like 1 minus omega p squared over omega squared 1 minus/plus capital omega over omega here, where I've just Taylor expanded this again.

Now, I want to be clear that this is not always true. You need to check whether this approximation works for your plasma. And for example, it may not work if you're using gigahertz radiation looking at a tokamak, because gigahertz radiation here, lowercase omega, in a tokamak, the cyclotron frequency is also in the gigahertz range. And so this approximation may not hold at all. This is a good approximation for the sorts of plasmas I work with. And it would be a good approximation if you used a higher frequency in order to make this measurement on a tokamak. But you'd have to think carefully about this when designing some sort of Faraday polarimetry system.

So what we see from this dispersion relationship is two things-- that this quantity, capital omega over lowercase omega, it has a small effect on the overall phase. So if you tried to measure this with interferometry, you're mostly measuring this term, and there'd be a very small change in this term. There'd be very, very small change to the overall refractive index caused by this small term here. So this is a very hard term to measure with interferometry. It's probably not going to work.

But on the other hand, this term turns out to have quite a big effect-- [INAUDIBLE] effect-- on the polarization. And we'll show that in a moment. And so although we can't measure it directly with interferometry, we can measure it in terms of polarization and so then we can still measure the magnetic field.

So let's have a look at what that looks like. We'll start by writing down the phase or a wave going through some plasma. So let me just draw a little diagram of what's going on here. We have some plasma like this. Inside it, there's some magnetic field, some density. And we've got our light going through. And the light could be right-hand polarized or it could be left-hand polarized.

And what we want to do is calculate delta phi that is the change in phase from a reference beam that has gone around the plasma. So this is the quantity we're going to calculate here. So delta phi R or L is going to be equal to the line integral of the wave vector along our probing path. That's just by definition. That's how much phase we pick up.

And that is also equal to omega upon c refractive index times dz, here. OK. And we have an expression for N. Here, we're going to be using the form where we've done the Taylor expansion just to keep things nice and linear. And so this is going to be approximately equal to omega upon c 1 minus omega p squared upon omega squared 1 minus/plus capital omega upon lowercase omega dz, by this.

And the approximate sign here is because we've made this approximation that capital omega upon lowercase omega is much less than 1. So we can see straight away that the phase of the right-hand and the left-hand polarized wave going through this plasma is going to be different. And it's going to be different by an amount that is to do with the magnetic field that we encounter inside this plasma.

So the right-hand wave, we take the minus sign. And for the left-hand wave, we take the plus sign. So the left-hand wave is going to have more phase than the right-hand wave. OK. Everyone following so far?

So this is where it's very convenient to go back and work with these Stokes vectors here instead, with X and Y and R and L. That's because, as we said before, if I prepare my wave so that it's polarized in x, so that it just has an initial electric field 1, 0, that is equal to R plus L over 2, after-- for example, this is what the wave looks like here. After the plasma, we want to know what the polarization state is, which I'll call X prime.

So after the plasma, this is going to go to a polarization state X prime, which is equal to still R and still L over 2, but I've less space here because the R and the L waves have picked up phase factors. They've picked up this phase here. So the R wave has picked up e to the i phase R, and the L wave has picked up e to the i phase L. And these two are not equivalent to each other.

So what is the consequence of the fact that these two are not equivalent to each other? What has happened to our polarization with respect to our initial polarization? Yeah?

AUDIENCE: It's rotated.

JACK HARE: It's rotated, yeah. But it's still linear because we can always write it as some linear polarization that's an instant in time. So you can just say that our initial polarization-- this is x, and this is y, and you're rather ill-advisedly staring at the laser beam as it comes towards you. This is your initial polarization X, right? And then you've now got a new polarization X prime. And that's being rotated by some angle alpha.

For some reason, people in the literature use alpha for this angle. Just go with it. Know that it's not theta. It's nothing to do with the angle between the probing direction and the magnetic field. This is a different angle.

OK. So then what is alpha? Can we actually work out what it is from X and X prime? Yeah, it's relatively simple geometry. We just have that alpha is equal to the arctangent of the y unit vector dotted into X prime over the x unit vector dotted into X prime. That's simply geometry here. We're calculating this length in y over this length in x and taking the arctangent of it.

So that means you have the arctangent of minus i e to the i phi R plus i e to the i phi L over e to the i phi R plus e to the i phi L. And you can verify that by just going through these definitions of R and L and thinking about the fact they've got x and y components and [INAUDIBLE] the x or the y components.

So this angle here looks quite complicated. Alpha is the thing that we're going to be able to measure. And that alpha is going to be related to the magnetic field, and it's related to the magnetic field through these phase factors. And these phase factors have picked up some term through the magnetic field inside here.

So it's not completely trivial to get the results here if you crank the algebra up a little bit. And I'm not going to do it. But the result is that alpha is equal to the integral of omega p squared over omega squared times all over 2c times capital omega squared dz here. We can write that in terms of some other quantities that we like a bit more, as e over 2 m e c. times the integral of ne over nc, the critical density, times the magnetic field in the probing direction, like this.

And we already talked about the fact that we're going to use the theta equals 0 dispersion relationship, even for theta 0 equal 0, as long as we replace omega with the omega that's due to the components of the magnetic field in our probing direction. And that's just what this represents here. This is exactly the same as having B dot k hat here.

I just want to note down here that this gives us a lambda squared dependence. Because it depends on the critical density, it matters what your probing wavelength is, and you will get different rotation angles depending on your probing wave. This is actually, let's say, yeah, nc is proportional to 1 over lambda squared. OK, questions? Yes? Oh, sorry.

AUDIENCE: Are we just rotating the perpendicular component of the polarization if we have some normal wave vector? Is the part that's parallel unchanged?

JACK HARE: So for these modes, there is no parallel component. These are purely transverse modes. There is no electric field in the direction of propagation.

AUDIENCE: OK. I thought that you said that you can have k not directly [INAUDIBLE].

JACK HARE: Ah. So this is the real subtlety that people trip up on. We have transverse modes, and we have longitudinal modes. We have perpendicular waves, and we have parallel waves. These words, although they sound like they're talking about the same thing, are talking about two different angles, right?

So here, when we're talking about this angle theta, and we're talking about components of the magnetic field, this is talking about whether our wave is parallel to the magnetic field or perpendicular to the magnetic field. But the electric field is going to be sticking out in some other direction like this. And it turns out for the X-mode, you actually have a small component of the electric field that is longitudinal that is along the k vector.

But when you solve the special relationship for theta equals 0, you don't pick up any of that. You don't have any electric field in that direction. It's purely like an electromagnetic wave in a vacuum. It's purely transverse. So there's only the electric field perpendicular to the direction of propagation. And so when I draw diagrams like this of the electric field, this is all of the electric field that there is, all completely perpendicular to the direction here, which I've chosen k here.

Now, if your question is, do you pick up a very small longitudinal component when you go away from theta equals 0, but you're still using the theta equals 0 dispersion relationship because you fulfill this condition, probably you do have a very small electric field coming up here. But we don't measure that on our detector. So we're only able to measure the transverse polarization.

And also, when the wave comes out of the plasma, it's going to have to couple back into the vacuum modes. And the vacuum modes don't have any longitudinal electric field. So actually, by the time it got to the detector, there won't be any longitudinal electric field, because that can't propagate in a vacuum.

AUDIENCE: So if there is a third one [INAUDIBLE] will it get deflected back into the plasma--

JACK HARE: Maybe.

AUDIENCE: --or something crazy.

JACK HARE: The plasma deals with it somehow, yes.

[INTERPOSING VOICES]

JACK HARE: [LAUGHS] It's a good-- that is a good question. I hadn't really thought about it before, but yeah. Certainly, by the time you get to the camera, you can only measure these R and L, which have no electric field in the direction of propagation. Yeah? There was another question?

JACK HARE: OK. Any questions online?

OK. So what are some implications of this? I mentioned one of them already.

So alpha is proportional to-- which looks like an alpha as well, I guess. I'll use this symbol instead-- lambda squared. So if we want to measure a large rotation angle, because presumably measuring bigger rotation angles is easier, then we want to use a longer wavelength.

Alpha also goes like the integral of B dot dl, like this. So we pick up the component of B along the probing direction. So it's B along. So we don't measure the vector B. We just measure a projection of it.

And that means that if the magnetic field is misaligned with our probe, we will still get some rotation, but we don't know necessarily what the orientation of the magnetic field is. OK. So that means that this, like all other line integrated measurements, requires a little bit of interpretation.

And the final thing is that alpha is still proportional to any dl, just like we had for interferometry. So we're not actually measuring just the magnetic field with our rotation angle. That's the thing we want. We're measuring the magnetic field weighted with the electron density. So we'll get a bigger contribution to the signal in regions where the density and the components of the magnetic field along z are high.

So you could imagine a system where the density is high and the magnetic field is low, and you might actually end up amplifying that low magnetic field due to the weighting with the high density here. All this to say that it's a line-integrated measurement, but it's sort of doubly line integrated. So it makes life quite hard.

At the very least, in order to have even a slight chance of measuring B, you need to have an inline interferometer. And that inline interferometer, its sole job is to measure this term so that you can divide it out of alpha and try and recover the magnetic field.

So let me just give you an example of numbers because one thing you might be thinking is, well, alpha is an angle, so that means we can probably only measure it modulo 2 pi, or even worse, maybe only modulo pi if we think about polarizations. So do we have to worry about alpha being an absolutely huge number?

So let me give you an example here. If we have a system where our density, electron density, is 10 to 19 per cubic centimeter, which is 10 to the 25 per meter cubed-- so this is pretty dense-- and we have a magnetic field of 10 tesla, which is not nothing, and we use a wavelength of 532 nanometers-- so this is a green laser, second harmonic, of neodymium YAG-- and we have a plasma that is 1 centimeter long-- so this describes some of the plasmas work quite well. When you get through all the maps here, you find that the rotation angle is a gigantic 3.7 degrees.

So even with this plasma, which is very dense and has a very strong magnetic field, we are not actually measuring a very large rotation angle here. So this measurement is extremely difficult. We're certainly not in any danger of having any ambiguity about alpha. It's unlikely that it's ever going to be getting up to pi. But if you are worried about the ambiguity, you should spend a little bit of time thinking about how you would heterodyne this. It's very cool.

So you can temporally and spatially heterodyne this. You have to find a way to modulate the polarization in space or time and carry your polarization signal on a modulated polarization. I don't think anyone's done it. It's almost done. I think know how to do it in time, but it's really hard to work out for it in space.

Anyway, so how do you do this practically? How do we actually make this measurement of alpha in a real system? Well, we'll have our setup with our plasma here. We'll have our probing radiation, which we'll prepare in some linear polarization state, let's say X polarized like this.

We will set up a standard Mach-Zehnder interferometer around this. That interferometer, its job, as I said, is just to measure any dl, like that. But now we've got probing radiation coming out of the plasma with polarization state X prime. And then we split that. And we send some of it through the beam splitter and through a polarizer, and the other beam also goes through a polarizer. And both of these beams go on to some sort of detector.

These detectors, again, we could be doing a temporarily resolved measurement, in which case you could think of a diode which is outputting a voltage trace onto an oscilloscope. Or we can have a spatially resolved measurement, in which case, you can think about a camera taking a snapshot in time. Or of course, you could try and combine those two if you had a fast enough camera. OK. So either of these could be functions of time or functions of x and y.

So these polarizers here, what you want to do to make a nice measurement of your system-- what we're doing here is a differential measurement. So I keep talking about differential measurements because they're very good ways to make good measurements. You set these polarizers up so that they're at a slight angle, beta. This is some angle in degrees.

You set these up so there's a slight angle beta to the extinction angle for X. That is to say that in the absence of any plasma, these two polarizers are perfectly crossed with the polarization of X. And so you see no signal whatsoever.

And you set one of the polarizers up to be a plus beta and the other polarizer up to be a negative beta. And so these two will see different things depending on the sign of magnetic field. So say the magnetic field points in this direction here. One of these detectors will see a signal that goes darker, and the other one will see a signal that goes brighter as the polarization is either rotated towards extinction or away from extinction.

If we reverse the sign of B here, we would reverse which detector goes brighter and which detector goes darker here. So this is a differential measurement. You see a brightening of one channel and a darkening on the other channel. And you know that's due to magnetic field, not due to anything else.

And so we call these signals intensity signal on the plus channel, intensity signal on the negative channel, like this. And these two intensities, I s plus or minus, are just equal to whatever our initial intensity was times sine squared of the rotation angle induced by the plasma and plus or minus beta, which is this slight rotation angle here.

So again, if there's no plasma in the way, alpha is equal to 0, there's no rotation. And these two detectors would just measure sine of beta. Beta is normally going to be a small angle on the order of a couple of degrees. So these are going to be measuring a signal very, very close to 0.

As alpha increases, on one of these detectors alpha minus beta is going to be close to 0, so that detector will measure nothing. It will measure darkness [INAUDIBLE] sine of 0 is 0. On the other detector, you'll have alpha plus beta. That will be some larger number. And so you'll measure a larger signal noise.

So you should see, for example, that this is a function of time-- actually, yes. If this is a function of time, and we, for example, have the magnetic field going up, one of these detectors is going to measure a signal that looks like this, and the one is going to measure a signal that looks like that. So you get brightening and darkening differently on the two detectors.

And the reason you do that is because, as well as the signal that you're measuring from your probing radiation like a laser beam, you also have some background radiation because your plasma is bright. And so you're always pushing out a copious amount of light in every direction, And this light is unpolarized. And because it's unpolarized, some of it will get through your polarizer. The amount that will get through your polarizer is, on average, 1/2. This just comes from thinking about what unpolarized light does. If you haven't come across this fact before, it's worth looking up.

Unpolarized light doesn't have a polarization, but it will-- still some of it get through a polarizer here. And so what you can do with these two signals, you can see that they differ in terms of how they treat the polarization of the radiation, but they're the same in terms of how they treat the self-emission. So this immediately suggests to you that you can combine them in some way to make a differential measurement of alpha, which becomes 1/2 of arcsine the intensity on the positive branch minus the intensity on the negative branch, and all of that times the tangent of beta over 2.

Again, I've skipped an awful lot of mathematics to get there. But effectively, when you take the difference of these two signals, you cancel out that self-emission that's unpolarized, and you get a signal-- you'll get an alpha which depends on the-- it has a sensitivity on whatever initial offset polarization angle we chose and on these two measurements of the intensity for the two different channels.

And if you want to know more about this, there's a good paper by George Swadling in *RSI* in 2014. And that explains the mathematics behind this and a nauseating amount of detail on sensitivity analysis as well. It turns out that [INAUDIBLE] propagation, this gets quite complicated.

So that's Faraday rotation effect. Any questions? Yeah?

AUDIENCE: I'm trying to understand exactly why we're implementing this beta scheme. I guess my initial guess when we were thinking about this was, oh, I'm going to have two polarizers which are 90 degrees out of phase that I can look at two different components. Is this because the expected angle shift or polarization is so small?

JACK HARE: Yeah, exactly. Right. So if you did the 90 degrees out of phase, you would be very well set up for measuring sort of 90 degree-style deviations here. But if you're trying to measure alpha plus or minus 90 degrees here, and alpha is very, very small, you're going to be measuring very similar intensities. You're going to be measuring something that's close to having an intensity of, like, 1 plus or minus a tiny amount. And so you're not going to have very much sensitivity.

Down here, you're very close to the null, where you've almost got no light, and so that means any change in intensity is very, very measurable. So Swadling's paper talks about exactly how you choose beta. But rule of thumb-- if beta is about the size of your expected signal but a bit larger, then it's good. The reason you want it to be a bit larger is if beta is smaller than your maximum value of alpha, then this line goes through I equals 0, only, of course, it doesn't it bounces back up, and then we have our phase ambiguity again. So you're trying to keep beta sort of just a little bit bigger, so beta is greater than alpha, like that.

Yeah. If you had 90 degrees, then you might want to think about a different scheme. But the main reason not to just use a single channel, because the other thing you might think of is, I'll just use a single channel, is the self-emission problem. And then you'll see brightness on your detector, and you'll think, ah, that must be magnetic field. But no, actually it's just the plasma glowing. And you have no way to tell the difference between those two.

If you think that the self-emission is really small because you've got an incredibly bright light source and it just overwhelms the self-emission, then maybe you don't need this stuff. But in general, it seems to be very helpful. Other questions? Yes, Sean.

AUDIENCE: I've heard of people using Faraday rotation to make astrophysical measurements. How does that work if you're not controlling the light source, if it's all unpolarized?

JACK HARE: Yeah. So Faraday rotation in astrophysical measurements, I believe you need to find a source that has polarized light, something that is emitting polarized light like synchrotron light, from some specific type of magnetized object. And then you look at the rotation of that through the intervening medium.

So you have to look for a source and then wait for that source to be between you and the object you're trying to study. So generally, people are studying the intergalactic medium or an interstellar medium, so it doesn't really matter where the source is. But it is the same effect. I mean, you still get a rotation that's to do with these two factors. I don't know in their case whether they are worried about phase ambiguity. Obviously, these can be quite long in the universe.

So I don't really know how they do that. I do think that they use multicolor techniques. So they can probe lots of different wavelengths, and so they get a measure of alpha for different wavelengths. And that helps reduce some of this ambiguity, like we talked about getting rid of vibrations and neutrals in interferometry, if you have multiple measurements of different wavelengths, that's also interesting. Yeah.

There was a polarimeter on C-Mod. I don't really know how it works. Like, I don't know whether they had problems of alpha being close to pi or something like that or exactly how it works, but yeah. OK. Any other questions? Any questions online? Yes, Nicholas.

AUDIENCE: Why would we do-- why would we measure B with a diagnostic such as this instead of just [INAUDIBLE]?

JACK HARE: Well, this measures the magnetic field inside the plasma. So it's very hard to stick a probe, a magnetic probe, inside many plasmas. Yeah, exactly.

And a second thing I'd say is that magnetic probes are perturbative. You stick them in the plasma, you will change the plasma. Whereas this technique shouldn't change the plasma. So there's two reasons.

The disadvantage is that this technique is line integrated. So we don't get the local magnetic field. And in fact, we talked about Abel inversion. If you have a cylindrical system, like a z pinch that has a poloidal magnetic field around it, there are very complicated formulas for doing Abel inversion of both the density and the magnetic field at the same time.

So you can still make progress if you know something about the symmetry of your system. But if you have no idea about the symmetry of your system, this is very, very hard to actually use in practice because you don't know whether B dot dz is changing, whether B is changing, whether ne is changing, along your line of sight. You have many different places where you don't know what's going on.

OK. Any more questions? Yes, go on.

AUDIENCE: [INAUDIBLE] we use magnetic diagnostic [INAUDIBLE] general question-- are we trying to actually map the whole [INAUDIBLE] field everywhere in space? Or is this more so you get boundary condition somewhere so then with a [INAUDIBLE] equation or something like that, we can figure out the rest? Are we [INAUDIBLE], or are we [INAUDIBLE] trying to map the field everywhere?

JACK HARE: So the question is, are we trying to measure the field everywhere, or are we just trying to measure it in some point for some boundary conditions, do some reconstruction? So on a tokamak experiment, for example, where, because we're in a low beta regime, our magnetic measurements are pretty good, we can measure at the boundary. and that's very close to what the measurement is inside the plasma.

I don't think you would need this to help with the Grad-Shafranov type measurements. This might be useful for getting higher temporal resolution in some cases, for example. On the experiments I do, obviously you're never going to get the three-dimensional field, because it's line integrated and you only measure a component of the magnetic field. But you can certainly measure the magnetic field as a picture. So you can take pictures of the magnetic field. You expand your laser beam up and take an image. That is pretty useful.

But of course, you can only interpret-- what you measure is the polarization. That's true. You know what that is. What you're trying to get out of the polarization angle alpha is the magnetic field. That requires some sort of model or some sort of intuition or some sort of assumption. So you're never going to get the full three-dimensional structure.

And particularly, like I said, the fact that you only get the components of the vector is not great. There are things like the motional Stark effect, which I've never fully understood, which does give you some way of measuring the local magnetic field. But never quite been able to work out what's going on with that.

Any other questions? Anything online? OK. We're going to go on to a new topic, refractometry-- or reflectometry. And then that'll be all.

So the idea of reflectometry is that we launch some radiation from one side of the plasma, and we assume that the plasma has some sort of density structure like this. So maybe this is a coordinate like r. Could be something like a tokamak. And this is density. And we know that somewhere, as this radiation propagates in, it's going to hit a surface where the density is the critical density.

And at the critical density, N-- and we'll call that point-- I guess I'm using x [INAUDIBLE] my note. We'll call that point xc. And so the density at xc is just critical density. And the refractive index at the critical density is 0. And so at that point, as we know, the wave is going to reflect.

So effectively, we have a system where the refractive index of our wave is 1 in the vacuum region. And then as it comes in, the refractive index goes down, goes to 0. And the wave would be evanescent propagating forwards. And so to in order to conserve energy, the wave has to reflect back out like that.

And if we have a measuring device, we can measure how long it takes for the wave to bounce back. Or we can measure the phase shift between the ingoing wave and outgoing wave, and maybe we can learn something about the position of this critical surface here. So this looks a little bit like radar, right? We send out a burst of radiation, we wait for it to bounce off our reflective object, the critical surface, and then we measure that radiation coming back. And from the phase lag, the time of flight, if you will, we get some information about where xc is.

Now unfortunately, this is not radar. Why is this more complicated than radar? Yeah?

AUDIENCE: With radar, you dealing with bouncing stuff off of solids, versus this is [INAUDIBLE] fluid.

JACK HARE: OK. So the answer was, in radar you're just bouncing stuff off solid, but this is a plasma. But in this picture here I've drawn, it looks quite similar, doesn't it? We've got some radiation coming in, we bounce off a reflector, we come back out. Seems like it should work. I mean, the plasma is acting as a perfect reflector.

So you're getting there, but that's not quite it. Yeah, Adam?

AUDIENCE: The speed of the wave should be changing as it goes through the plasma due to the [INAUDIBLE].

JACK HARE: Absolutely. So the answer was, the speed of the wave changes as it goes through the plasma. And so in radar, the speed of the wave is constant throughout the entire system because we've just got air, it goes at the speed of light, and it bounces back. In this system, the speed of the wave changes.

And so although it's going to reflect off this point, the details of what happens in the region before it reflects are very important. So delta phi depends on the density for x less than xc. So it depends on whatever the density is doing in this region here. Doesn't depend on what the density is doing afterwards, because we've reflected. But this region here, it is very important.

OK. So let's look at that in more detail. Let's say delta phi is equal to the phase of the wave at B, this point here, minus the phase of the wave A, where it enters the plasma. That, as before, is equal to omega upon c the integral of the refractive index dx A, B, like that.

This is just what we've done before. So now, we're just dealing with the phase of the wave going in. We will double this when we want to come back out again because we'll assume the plasma is stationary on the time scales of this light being reflected.

But there's a big problem with this equation I just wrote down here. Does anyone know? I've made a huge assumption in all of the work I've been doing that some of you picked up on before that isn't valid for some region of this plasma. Especially people working with RF waves might spot something here.

So we've been using a WKB approximation. But that WKB approximation does not work when the refractive index gets close to 0. And one maybe intuitive, maybe not way to think about that is as the refractive index gets to 0, the wavelength of your wave gets very, very large. And the WKB approximation assumes that the wavelength of your wave is much smaller than any gradient in your system.

And so as N gets to 0, the wavelength gets long, even this gentle gradient in density starts to look too large, and the formulas that we've been using don't work. So we cannot simply double this formula here. It turns out if you do all of the math appropriately, apparently you also think of a factor of minus pi upon 2 for this reflection. Don't ask me where that comes from.

OK. So actually, your total phase, your phase doing the bouncing, coming back out, is going to be equal to 2 times this-- times that. But we only pick up the pi over 2 once at the reflection here. OK. And that's because you can't really put your detector here. You're always going to have to do a double pass in this system. So it looks a little bit like a [? Michelsen ?] problem.

OK. And this integral-- I probably should have just written this on a new line. OK, I'm going to get rid of that. So let's just say that the phase is actually going to be 2 times omega upon c.

And the reason I wanted to rewrite this is I want to replace these limits, A and B, with a being where the wave started and xc being where the wave reflects, times the refractive index dx minus pi upon 2. And that pi upon 2 is due to reflection.

And the point to give here is that if we measure delta phi and the phase lag between a wave going into the plasma and the wave coming back from the plasma, it isn't just like equal to something times xc. xc appears in the argument of an integral-- sorry, in the limits of an integral, and the argument of an integral depends on all of the plasma before we get to xc. So this is mathematically a little bit why this isn't radar. If it was radar, you'd just have delta phi equals some constant times xc, and that constant would be a distance [INAUDIBLE].

OK. So there are ways around it. And the way around it is you use multiple frequencies, use many omegas, like this. So for example, you have measurements of delta phi, and you measure it for lots of different wavelengths of different frequencies. Maybe it looks something like this.

And this could be you're doing the measurements simultaneously with lots of different wavelengths, like a two-color, three-color, five-color measurement. Or we'll talk about some other techniques to do with chirping it later on, where you have a continuously ramping frequency and you listen for the reflected wave coming back, like that.

But the reason you do this is that after a little bit of mathematics that you can find in Hutchinson's book, you can find out how this phase changes with the frequency you put in. And you get an equation that looks like d phi d omega is equal to 2 times the integral between lambda and infinity dxc d lambda p-- and I'll define some of these in a moment-- lambda p d lambda p lambda p squared minus lambda squared all to 1/2.

This looks very odd. Don't worry. It's going to make slightly more sense soon. Lambda p is a slightly odd quantity. It is the wavelength that a wave with a frequency of omega p would have in free space. Obviously in a plasma, a wavelength frequency of omega p has infinite wavelength. But this is just the equivalent in free space, and so it's just 2 pi c upon omega p. And lambda is just 2 pi c upon omega.

So we've gone from this equation, which implicitly, of course, has N is equal to 1 minus omega p squared upon 2 omega squared-- all the way through this, I'm assuming the O-mode. All the other modes are very, very complicated to work with, but there are definitely a few different modes. And you can do it with other modes as well. So we've taken this equation. We've differentiated it with respect to omega, and then we have taken the bizarre step of, although we've got omega on the left-hand side, we've written everything here on the right-hand side in terms of lambdas.

Anyone have any idea why we did that? It's a mathematical trick we're about to employ, and you've seen it in this class.

AUDIENCE: An Abel inversion?

JACK HARE: It's an Abel inversion. Well done. Has no right to be, but it is an Abel inversion. So for those of you who are not following, you remember we had the Abel transformation. So this is the one where you go from your [INAUDIBLE] distributed function, F of r, to your line-integrated function, F of y. And that was equal to 2 times the integral from y to a of f of r r dr over r squared minus y squared to 1/2.

And so I can make the identification of this term with this term, of lambda p with r and lambda with y, like that. This means I can use all of the mathematical tricks that I had from the Abel inversion to do the inverse Abel inversion. Remember, this is what we've measured. This is what we want. So we can now do the inverse Abel inversion.

And again, a fair bit of mathematics follows. So if you don't follow all of this, it's because I'm not showing you every step. But then we have the location of the critical surface as a function of frequency is equal to a minus c upon pi the integral from 0 to the frequency that you want to know xc for, d phi d omega prime d omega prime over omega squared minus omega prime squared to 1/2.

I want to point out that as we change frequency, we obviously change location of our critical surface, right? If we use a higher frequency, our critical surface will be at a higher density. And so for a monotonic density profile, it will be further in.

What this equation is saying is that we can know the location of the critical surface for every frequency if we know the phase shift for every frequency omega prime, where omega prime is less than omega. So if we have measured the phase shift for every smaller frequency coming up to the critical frequency we're interested in, then we can know the location of the critical frequency.

So this is a very powerful technique because if we know the location of the critical frequency, we also know the density at that location because it's just the critical density. And so we can build up-- as a function of x, we can build up the density here. This is critical density at x1. And then we can build up the critical density at x2.

And you see soon enough, we can start building up the entire density profile as a function of x, which again, in a tokamak, could be something like the minor radius. So this is a way of measuring the actual density profile to see how much more powerful this is in interferometry, which only gives you line-integrated profiles, this gives you the accurate density profile. So this is a very, very cool technique. Any questions on this so far? Yeah, [INAUDIBLE].

AUDIENCE: [INAUDIBLE] use Abel inversion [INAUDIBLE] do we need to make assumptions about the symmetry-- do we need to say something [INAUDIBLE]?

JACK HARE: No, we made those assumptions in order to derive this formula. Now by analogy, we notice the formulas are the same. And so there is-- I don't know if there's something deeper under here that I don't understand, but there has not been any assumption about symmetry in our plasma.

I will say we've made an assumption that our plasma has a monotonically increasing density profile. Or rather, should we say that if your profile is like this, you can only measure up to here. You can't measure the other side.

And that's a pretty good approximation for a tokamak. But if you have some funky system which has a dip in the center, you won't be able to measure that. Hutchinson says you can't look over the hill, which I think is a nice intuitive way of doing it. You can only measure up to a maximum. So yeah.

Just to say that this is an analytical solution for the O-mode. If you happen to decide to do this with the X-mode because you want to use that polarization, you can, but you have to do all of this numerically instead. And it doesn't come out as nicely as this. But it's completely possible. You just have to know what your dispersion relationship is as encoded by the refractive index N.

Any other questions on this before we move on to a few practicalities and then finish up, actually, quite soon? Any questions online? Yeah?

AUDIENCE: You expressly drew this almost as it looks like a microwave horn. I guess this is a question I can answer myself in the sense of just computing the range of critical density of plasmas of interest to know what range of frequencies to use. But on machines like tokamaks, is this normally done in the microwave range?

JACK HARE: Yeah, yeah. So I did draw it as a microwave horn, and that is to suggest microwaves because this is a technique that is done on tokamaks. Because of some of the practicalities I'm going to talk about in a moment, for a system which is very short lived, it's very difficult to make a measurement at lots of different frequencies. And so for the plasmas I work with, this is completely impractical.

You would also, for the plasmas I've worked with, need to generate radiation that goes up to the critical density. And that would be sort of like soft X-ray lasers, which are also very, very hard to get hold of. Like, visible radiation, we're in that nice regime where any over nc is very, very small in the sort of plasmas I work with. So that's great.

But if I wanted to actually work near the critical density, I'd have to get something with much shorter wavelengths. And that technology doesn't exist. So although Hutchinson's book is very good at emphasizing the principles of diagnostics rather than the practicalities, a lot of the time the practicalities of what radiation sources we have and what detectors we have really dictate what techniques we can use in different plasmas. And then when new technology becomes available, that enables us to do new things that we weren't able to do before.

So OK, so let's talk about a practical system for making this measurement, which is the phase lag at a range of different frequencies. And can I just point out something that probably some of you have spotted already? I drew this very deliberately as a nice, discrete measurement that is a discrete set of omega. And so you're going to have numerical noise when you differentiate that.

And once again, that actually matters because I've written this little small d's here, but of course, in reality it's always going to be capital deltas. And so you're going to have to think about, can I fit this with some basis functions nice and smooth? Can I do all sorts of techniques that I normally do to avoid amplifying up this noise?

And of course, I'd be mostly amplifying noise very close to the frequency that I actually want to measure. So noise here is going to be very important for measuring the location of the critical surface, this frequency. This is like the Abel inversion, where noise near the center of the image was very, very important. So there are some nice analogies you can draw here. OK.

So a surprising number of plasma physics techniques actually started off being ionospheric measuring techniques. So if we've got the Earth down here, and we've got some sort of radar dish, we have, floating above it, the ionosphere, which is a low-density plasma, like this. And you know, Sputnik is up here having a good time flying around.

And so when people invented radar technology during World War II, they suddenly had this ability to produce intense bursts of quite specific frequency radiation that was being used for radar. And they, therefore, were able to use that technology for civilian purposes and start doing interesting measurements. It turns out that radar is well matched to the sorts of densities that you get inside the ionosphere.

And so if you want to do this reflectometry measurements and try and build up a picture of the density in the ionosphere, you can start sending out different pulses of radiation. And these different pulses will reflect off, and they need to be detected, I mean, in principle, by another detector, maybe something very close to the original one here. And you can build up a map of what the density looks like without even having to fly a satellite up there.

The ionosphere is relatively slow evolving, right? It doesn't change very much. So you can send a pulse at one frequency, wait for it to bounce back, and then send the next pulse. And because the ionosphere is evolving slowly, the properties of the ionosphere don't change during the seconds that you're doing this technique, where you're waiting for the reflection and analyzing the data from it. So this works very well in the ionosphere.

For a tokamak, this doesn't work very well. If you have to wait a significant amount of time, the plasma is going to have changed by the time you get around. So what they do in a tokamak is a slightly more complicated system.

So let's have our tokamak, nice D-shaped cross-section, diverted plasma like this. The density is constant on a flux surface, so these flux surfaces I'm drawing are also surfaces of constant density here. And again, I'll have some sort of hole and I'll be launching radiation into the plasma. With this, I'll also have some hole collecting radiation coming from the plasma.

But what I'll do is I'll have a source of radiation, and I will also split some of that radiation that's going into the plasma and send it to a heterodyne mixer, where it adds a small amount of frequency. Remember, this could be our rotating wheel or our moving mirror or our acousto-optic modulator or something like that. And then I mix that back in. I effectively do temporally heterodyne interferometry on the returning wave, where I have omega and omega plus delta omega, like this. But at the same time, I sweep this initial frequency here.

So for example, I have on this plot here, I think-- I can't remember which tokamak I got it from-- there's a system which sweeps from 40 to 70 gigahertz, and it does that in 15 microseconds. But that's kind of like your sweet period, and then it drops back down and does it again. So if you want to capture dynamics, you can capture dynamics on timescales which are greater than 15 microseconds. If there's something small and fluctuating on a shorter timescale, you won't be able to capture it with this system.

And here, we're not measuring the delay directly as we did with this system, where you're literally counting seconds between sending it out and coming back. Now instead, you're measuring the phase shift with this heterodyne system, as we did for interferometry.

So just a couple of final notes on this technique and some of its limitations. So one big hope of this technique will be to measure fluctuations. So fluctuations are very important in magnetic confinement devices because they are associated with turbulence, and turbulence is associated with anomalous transport and, therefore, bad energy confinement and, therefore, uneconomical reactors. So we'd like to understand what is going on with fluctuations inside here.

And so Hutchinson says, a lot of people were trying to use this technique for a very long time to measure fluctuations, but it is very, very challenging to do this. And one of the challenging things is that it's sensitive near the critical density, so near xc, the location of the critical density. And that's kind of what I was talking about here, where this formula has this singularity in it. And so it's going to be very sensitive to contributions very close to the critical density.

But there's also still contributions from the plasma, which is the x less than xc. So if it was just sensitive to fluctuations in xc, this would be great. You could think about it like a little vibrating mirror, and we could measure that position of that mirror really sensitively. But unfortunately, it's not. There's also these contributions from elsewhere. And Hutchinson says, rather dismissively, that most people forget about this second problem when analyzing the data and just assume that it's just measuring fluctuations in the location of the critical surface.

If you include the fact that there's these other contributions from elsewhere, you can imagine the fluctuations are, maybe, out of phase or doing something else, and it's going to make your measurement really hard to interpret. And also, a lot of the picture we've been using here has been one-dimensional. But reality is very three-dimensional, and so there might be waves which are scattering off at slightly odd angles. There might be reflections off the walls, things like that. And that makes this measurement very, very tricky in general. So as far as I know, people mostly use it just for measuring the sort of slowly evolving density profile, and not these small-scale fluctuations.

And the final thing I want to do is just contrast reflectometry with interferometry. So reflectometry, you measure near the cut-off. So you're measuring near this critical density here. And you can't measure behind the location of the critical surface. So this was the "looking over the hill" problem here. So you're choosing your radiation so that you are near the critical density.

That's completely different from interferometry, where you usually work in a regime where the densities you're measuring are much less than the critical density because that means you don't have to worry about refraction and all sorts of other effects. And so you choose your probing wavelength so that it's sensitive to these densities much lower than the critical density, and we make very small measurements there.

And so this means that we avoid the cut-off, but of course, our measurement is line integrated in interferometry in a way that, although it is complicated in reflectometry, you still can actually get the profile along the probing line of sight, which you don't do for interferometry without some symmetry arguments, which could lead to something like Abel inversion.

OK. Any questions? Yes?

AUDIENCE: How exactly does the [INAUDIBLE] receiving [INAUDIBLE]. Why would it not just bounce back into the [INAUDIBLE]?

JACK HARE: So apparently, you do need a two-horn system. People tried to do it with a one-horn system, but the reflections are very complicated to deal with. And so it's been found to be better to have a launcher and a receiver. There's a few more details on that in Hutchinson's book.

AUDIENCE: OK, thank you.

I will say-- so when did this class two years ago, there were several people in the group who worked on reflectometry. And they told me that there may be some more up-to-date stuff than what Hutchinson had in his book. But they weren't able to point me to any review papers or even any readable papers about this.

And this is a huge problem with diagnostics. So there's not a paper, I can't teach it. So if anyone is here and they work on reflectometry, and they're like, this is ancient, we haven't done this in 20 years, shout out. I'd like to learn. It's just like, I've been unable to find any resources on how this is done better.

I know, for example, on ASDEX, there is a reflectometry system that is trying to measure density fluctuations and correlate them with the temperature fluctuations measured using correlation ECE, which we'll talk about later. So that sounds like an incredible system. And it sounds like they are measuring density fluctuations.

But I have no idea how, because, again, Hutchinson, he says this is basically impossible. So there must be a way to do it, but don't know what it is. Was it a technological advancement? Was it a conceptual advancement? I'm not sure. So yeah.