[SQUEAKING]

[RUSTLING]

[CLICKING]

JACK HARE: So maybe we'll start with the Schlieren photography stuff. You're thinking to yourself, this is interesting but it doesn't really apply to my research. I can do tokamaks, I can do low temperature plasmas. We don't really have high enough density to use visible radiation like lasers, and so we're not really going to be able to see things like shocks, even if there were shocks, and there probably aren't shocks inside your tokamak plasma. Everything is nicely subsonic.

So that's OK. Maybe you tuned out a little bit. But today we have a topic, which is applicable, really, to every form of plasma, which is interferometry. Because we can do interferometry all the way from very, very low densities all the way up to-- it is possible, or at least applicable to very high densities, as long as we have some probing radiation which is sufficient to penetrate your plasma.

So this is interferometry. And I spend an awful lot of my time thinking about interferograms and so I really like this stuff. And hopefully, we'll have a good time learning about it as well.

So the basic principle of interferometry is that previously, for the Schlierenography. Or the Schlieren, we were interested in gradients in the refractive index, which were, of course, gradients in the electron density. For interferometry, we're interested in refractive indices which are varying in some way. So they're not just equal to one, which is the refractive index of the vacuum.

So with interferometry, we're looking at this refractive index capital N again, which, as we've written a couple of times in a cold high frequency unmagnetized plasma, looks like one minus omega squared over omega p squared. And then we wrote that as square root of 1 minus the density over the critical density.

And I promised I would define the critical density because I forgot to do it in one of the earlier lectures. So here's critical density. It's defined as epsilon 0 me over e squared times omega squared. So you can see there's a different critical density for every probing frequency and the rest of it just depends on these fundamental constants. So that doesn't change at all here.

And one of the things we want is to be able to recreate this refractive index so we can work out how-- it relates to the frequency and the k vector of our wave, because those are the things that go into the phase of our wave. So just here, I'm just going to write that N is actually defined as the speed of light over the phase velocity of the wave, which is equal to the k vector of the wave and the speed of light over the frequency of the wave. So we're going to be using the link between these two here, which means that if we can measure something as a k vector, or the frequency of the wave, we can measure something to do with the electron density side of the plasma.

And the setup that we'll have, we'll have our plasma like this. And maybe it's got some length, capital L like this. We'll put our probing radiation through the plasma and it's got some electron density associated with it as well.

And that electron density can vary. It can just be a function of position or of time. We might be interested in measuring the variation in space or in time here. And we call this beam that's gone through the plasma the probe beam. And so far, this looks an awful lot like Schlierenography.

The trick is we also put a beam that does not go through the plasma, it goes around it. It can go through the same vacuum chamber in a region where there is no plasma, it can go around the vacuum chamber, doesn't really matter. And we'll call this beam the reference beam, or ref for short here.

And what we're interested in knowing is what phase the probe beam and the reference beam acquire. So the probe is going to acquire a phase by probe. Now remember, the phase is the thing, when we have our oscillating electric and magnetic fields, we have this exponential i of a.x minus omega t, like this, adding this quantity here, which is the phase.

So if we want to work this out and we just want to know what it is at an instant in time, we can drop the time variation here. And so this phase that this probe beam picks up is going to be equal to the integral of k.dl, where dl is some infinitesimal distance along the probing line of sight. And that k.dl is going to be equal to using this definition here, the integral of N omega upon c dl, like that.

On the reference beam, the situation is very similar. So we have bireference. The reference beam is going through a medium where the refractive index is just unity, and so we can just write this as the integral is omega 1 c dl.

And that means that there's a phase shift between these two quantities delta phi, which is equal to the phase of the beam going through the plasma minus the phase of the beam going around the plasma. And that is going to be equal to the integral of n minus 1 omega upon c dl, like this.

Now the electromagnetic radiation, omega is going to be constant in this system. The frequency of the wave doesn't change as it goes through the plasma. The speed of light is constant. So really, this is just an integral over n minus 1 here.

And using our Taylor expansion here, we assume that Ne is much, much less than Ne critical, we can get out an expression here where delta phi is roughly minus omega over 2 c and critical integral of Ne dl.

So we now have a phase which is linearly proportional to this quantity, the line integrated electron density. So say, for example, our dl was roughly in the z-direction, we would have Ne dl as a function of x and y, for example.

And there's probably still some time variation inside this as well. This could be a function of time as well. So we could be trying to probe how this phase changes in time or how this phase changes in space, and we'll look at both types of interferometers. So the goal, clearly now, is to try and measure this phase so that we can get out the electron density. Any questions so far? Yes.

AUDIENCE: So in the lines where you're defining capital N as the square root of 1 minus omega on omega p squared, and then right below you say that N is important to k, c, and omega. K is a vector, and so N is also a vector. How do we get the vectorial information from that first line defining N?

JACK HARE: So the question was this k here is a vector and so this N should be a vector. And so where have I lost the vector nature of it? N definitely is a tensor, I think, in reality. So I've not even started that. That complicated even further, where does the tensor nature of it go away.

I think at this level, I'm just going to say that this plasma is isotropic and so it doesn't matter. When we bring in the magnetic field later on, we'll find out that N isn't isotropic anymore and you might have to think more carefully about the direction your probe is going in. It's going to have something to do with the direction of k here.

So I just want to point out that I'm actually deriving all of this, at least in many places just using capital N, but if that means that later on, you want to go and work out what-- this is for a magnetized plasma or tokamak, where this really matters whether you're doing x-mode, or o-mode, all that good stuff. You can go back and rederive this. I just want to try and keep this as simple as possible to start with. And I don't know exactly where we're sort of losing the vector nature.

So I'm being a bit wishy-washy here. Any other questions? Any questions from online?

AUDIENCE: Yeah. Maybe this is what was just asked, but couldn't-- depending on the type of wave, couldn't the like omega change when it goes in the plasma?

JACK HARE: So the frequency of the wave doesn't change, but the wave number does, the wavelength does. But the frequency of a wave won't change as it propagates through a plasma. But this omega doesn't work.

Sorry, this k definitely works the wave number. So the wavelength of the plasma, the frequency of the irradiation doesn't change. Any other questions?

OK. And so we'll talk about a few different schematics for actually interfering this reference and probe beam in a moment. But for now, just imagine that we've somehow managed to get the reference and the probe into the same place. And so these two beams of radiation are interfering.

And then calculate what we're going to see on our detector. And our detector could be some sort of diode that resolves the signal in time, or it could be some sort of camera that resolves a signal in space. It doesn't really matter for our purposes here.

So we're going to have an electric field due to the probe beam, which is equal to whatever the strength of the probe beam is times this exponential of ik.x minus omega t plus delta phi. So this is the key bit. This wave is now ahead or behind in phase by an amount, delta phi, and that delta phi is this quantity here that we're trying to measure.

The reference beam is much simpler. It just, again, has a strength of the reference beam exponential of ik.x minus omega t. There's no additional term here.

Where we actually start the 0 of the phase is arbitrary, and I've chosen to put the 0 of the phase such that all of the phase shift is in the probe beam. I can do the same thing [INAUDIBLE]. It doesn't matter now.

And then we will have a total electric field that's hitting the detector, which is equal to the probe electric field plus the reference electric field. That is going to be equal to the strength of the probe beam times the exponential of i delta pi plus the strength of the reference beam. And all of this it's going to be timed by our old friend, exponential of ik.x minus omega t, like that.

Now we don't sense electric fields with our detector because these electric fields are probably oscillating far too fast for our detector to see. So even if you're using visible light, it's very hard to detect if it's fast enough. You see that? Even if you're going down into your gigahertz, even 10 gigahertz detectors are expensive.

So we don't resolve the actual oscillations here. What we resolve is a time average power hitting our detector. And so we're averaging over some cycle. And so this intensity is equal to these brackets here, which are going to indicate time averaging and then the total electric field times the complex conjugate of the total electric field.

That's going to give us a real value here, which is proportional to the square of the absolute value here. That, when you work through all of this, you get out the strength of the probe beam squared plus the strength of the reference beam squared, all of that times 1 plus-- I'm going to run out of space. It's going to squash [INAUDIBLE].

Strength of probe squared plus strength of reference squared times 1 plus 2 strength of probe strength of reference over sum of their squares, all of this times cosine of delta phi, like that. And during this averaging process, we've effectively averaged this operating function times its complex conjugate. And when we average over multiple periods, this goes to 1, which is why we don't see this anymore.

So when we're measuring the signal here, what we're seeing is the signal depends on delta. It doesn't depend on any of these little details because these are all fluctuating far too fast for us to see. So this is a function which has intensity in it.

And that intensity is linked to delta phi here. Anyone know-- if I want to measure delta phi very nicely, I clearly want to maximize this term with respect to this term. Do you know what values of the probe field strength and the reference field strength I need in order to maximize that?

AUDIENCE: Do you want it to be equal?

JACK HARE: Yeah. So if I set this equal to this, then I will maximize this. And then, in fact, I can write this as i over i0, identifying this as i0 here. And this will be equal to 1 plus the cosine of delta phi, like that.

And so this is quite neat. It's actually relatively easy to balance these. You can just put some attenuators in until they're about the right strength. Give them the super sensitive function, so if you don't get it perfectly right, you'll still get nice balance.

But this function clearly goes between 0 and 2, like that. So you can have constructive interference, where you have twice as much intensity as you had initially, and destructive interference where you have half as much intensity. And you come across presumably constructive and destructive interference. Any questions on that? Yes.

AUDIENCE: Maybe I'm asking this question prematurely, but it seems like a major limitation of this is that you have some cyclic pattern where a phase offset of 0--

JACK HARE: Yeah, we're going-- yeah. So the fact we can only measure delta pi modulo 2 pi is a problem and we're about to address it. But yeah, it's a great question. So I see a question online. Jacob?

AUDIENCE: Hi, yeah. Does the length of the time you send out the probe matter? Because if you send the probe for too short of a time, then are you not going to be able to average it over a certain-- is there a certain amount of time you have to average it over, or does it--

JACK HARE: I mean, you want to average it over a few cycles. Really strictly speaking, you only need to average it over one cycle, 1 over omega. But in reality, it's actually really hard to make a probe beam that's that short. And so I think people who do femtosecond lasers, and femtosecond interferometry might begin to have to worry about that, but most of us don't have to worry about that.

Yeah, so in this case here, this is just a statement that your detector can't possibly catch up and so it is going to average anyway. Except for if you're doing very low frequency radio interferometry where you could actually track the waves independently. Yeah. Other questions? I saw one over there.

So this is your result. I now just want to take a little detour and show you some ways in which you would set up this so that we can measure delta phi and then we'll talk about some of the limitations and ways to overcome those limitations.

So again, our goal is to measure delta phi, the phase shift within the [INAUDIBLE]. So one configuration that we can use which is quite popular is something called a Mach-Zehnder interferometer. So we have our plasma here.

We send our probe beam towards it. We put a beam splitter in, beam splitter. And we send half the beam through the plasma and half the beam around the plasma. I always have to remember to get the beam splitters the right orientation by the very stupid beam splitter. And then this goes through to our detector here.

So the nice thing about this setup-- I should give the name for it, sorry. It's a Mach-Zehnder. The nice thing about a Mach-Zehnder is you set it up so that in the absence of a plasma, this length, l probe, is equal to this length l reference.

That's pretty useful if you're using pulsed radiation. If these two parts are different lengths and the pulses arrive at different times at the detector, you won't get interference. So you want them to arrive at the same time, so you'd like to have it balanced here.

It's also-- it's a nice setup here because it's very flexible. These are very easy to set up. You don't have to have this beautiful square configuration.

If, for some reason, your lab is a weird shape and you have to do like that, you can do something like that, as long as you keep the lengths about the same. So that's quite nice. The trouble with these things is that you do need to get the alignment very good. I'm putting a negative sign here because this is a negative of this set up.

When you're positioning all of these mirrors, you need to get them within what's called the coherence length of your laser beam. And that coherence length for a laser could be maybe a centimeter or so if it's like a nanosecond pulse. But if you're dealing with picosecond or femtosecond pulses, that distance could be only microns, and then you've got to get the mirrors in the right place within a micron.

That's very, very challenging. So this is a problem with alignment and coherence for your laser. So they're easy to set up, but they may be very tricky, depending on the pulse length that you're working with.

The next one that you have probably heard about in other contexts in physics is the Michelson. So a Michelson interferometer has a beam coming in and it's got a beam splitter. It's got a beam going up and we put a plasma in one arm of this beam.

And the beam goes up and then it bounces back down. At the same time, because this is a beam splitter, some of the light is transmitted through, reflects off another mirror, comes back, reflects off, and these two go down to your detector. This is a Michelson, of the famous Michelson Morely Experiment.

This is very sensitive. The reason it's so sensitive is we have a double pass here. So if your plasma density isn't very high, then you can increase your signal delta phi by going through twice.

That just comes about, because when you evaluate this dl layer, you're obviously doing the integral through the plasma and then back through the plasma the other way. So this enhances your sensitivity. [INAUDIBLE].

But of course, there is a problem here that you don't want your plasma to change while you're doing this batch. Your plasma dynamics are such that the density changes, as you bounce through and come back on the other way, then you'll be measuring, effectively, a different plasma. And so this approximation of the density as constant as you go through in time won't be very good. So you do have some issues from time resolution.

AUDIENCE: [INAUDIBLE]

JACK HARE: Yeah?

AUDIENCE: Does that really come up that much in practice, though? Because you're talking transit time over reasonable size experiment. It's got to be nanoseconds at longest.

JACK HARE: I do nanosecond experiments, yeah.

AUDIENCE: OK.

JACK HARE: But it might not be a problem for the experiment you're thinking of. So if you're thinking of a tokamak or something, we don't tend to look at dynamics on nanosecond timescales there. But they might be a problem for some people.

So it's worth thinking about. Yeah. Any other questions on these two as we pause for a second? I've got one third one to show you. Yes?

AUDIENCE: With the one on the left, [INAUDIBLE].

JACK HARE: Mach-Zehnder?

AUDIENCE: Yes. You said that you just need the [INAUDIBLE] to be reasonably close together, but it means high frequencies--I would imagine that the air between them needs to be approximately like a wavelength or a couple wavelengths?

JACK HARE: No. No, it doesn't matter. It doesn't matter as long-- because we can only measure the phase difference modulo 2 pi. So if your reference beam goes another 100 phase units, radians, but we can only measure the phase difference modulo 2 pi, in the absence of any plasma, we'll just see the same result as if you went 0 phase units.

AUDIENCE: Oh, OK.

JACK HARE: Yeah. So they can be different by many wavelengths, but they have to be the same within the coherence length of the source, which defines the length scale over which you can get interference. If they're not coherent, they can't interfere.

Which is, by the way, one reason why we use the same source and we use beam splitters here, if anyone was wondering. You might think, oh, I could just use two lasers. But if you have two lasers, even if by the same manufacturer, sitting in the same laser bay at the same temperature, because of the very slight differences in how those sources are made, they will not be coherent with each other. And so you are better off splitting the radiation and interfering the beam with itself. Yes?

AUDIENCE: So if I'm understanding correctly, that means the lengths of the two arms can be different, but they have to be different by 2 pi multiple.

JACK HARE: And if they weren't different by 2 pi multiple, that would also be OK. So if I-- we'll get onto this in a little bit, but anyway, if I plot the intensity, which is 1 plus cosine of delta phi here, and maybe plot it in time, and think about a case where they've got these two beams with exactly the same length.

So if they're going at exactly the same length, when they interfere, they're going to interfere constructively. So we'll get an intensity of 2 up here. So this is where my initial delta phi is 0. And I'm not going to get any plasma in here, so just watch my diode and it reads 2, and it just reads 2 for all time.

Say I now change this so that the difference in length between these two gives me [INAUDIBLE]. So this could be 0, 2 pi, 4 pi, something like that. What happens if I have a delta phi of pi or 3 pi, something like that? I chose my lengths so that my phase differences pi or 3 pi.

AUDIENCE: 0.

JACK HARE: 0. My detector will just read 0 forever. And if I choose something in the middle, it will read 1 forever.

When I put the plasma in, the signal will start to deviate from this. And that's what I'll be measuring. But the background signal, my baseline, I can just measure that in advance. So it doesn't actually matter in reality.

AUDIENCE: So it's like a [INAUDIBLE] probe where we can only measure the time change.

JACK HARE: And when we start talking about temporally heterodyne interferometry in a moment, we will see a very deep link between what you're talking about there. We're really measuring time derivatives of phase, but it doesn't really look like it to start with. But we'll get on to the point where we are measuring actually changes in-- yeah, we're only sensitive to changes in phase. I guess that's the main thing. Yeah, whether that's a time derivative or spatial derivative doesn't matter. Other questions?

OK. So then the final type of interferometer that I want to talk about is very simple to set up. You have your probing radiation coming through. You've got your plasma like this.

And then on the other side of your plasma, you have-- what's this called? I believe this is called a Wollaston prism. This special device called a Wollaston prism.

And the cool thing about a Wollaston prison is that it splits your light into two orthogonal polarizations and separates them by angle. So this is a birefringent material. It sends out some of the light up here and that light is going to have, maybe, polarization in this direction, and the other little light that's going to come down here. That light is going to have polarization in this direction.

Then we put a lens here. Start sending these two back together. And if you know something about interferometry, you're thinking, well, this is useless, because these two are not actually going to interfere, because they've got orthogonal polarizations.

You can prove that to yourself, if you go through this and do it in vector notation instead, you eventually end up with an ep.er. And so if they're orthogonal to each other, you don't get an interference. The beams don't see each other. But the really clever thing here is you put in a 45-degree polarizer.

So again, this beam is coming in polarized in this plane. This beam's coming in polarized out of the plane. But this 45-degree polarizer picks out the 45-degree component of this, the 45-degree component of this. So by the time they hit your detector, they've both got the same polarization. And so then they interfere with each other.

You're effectively creating your reference and probe beam from the same beam. So the idea here is that in fact, I've drawn it very well, you have an expanded beam that goes-- some of it goes around the plasma. And you effectively end up interfering this bit that's going around the plasma.

There's the bit that's gone through the plasma on the inside. So this is called a Nomarski. Can someone close the doors [INAUDIBLE]?

The nice thing about Nomarski is very easy to align. Because we talked before about how if you have femtosecond beams, trying to get all these optics in the right places very hard. Well here, all of the splitting takes place within the same optics. So it's always in the right place, like that.

But because you're interfering a bit of the plasma-- a bit of the beam that's outside the plasma, there's a bit of beam inside the plasma, you have a limited field of view. So you can only image something that's as big as the Wollaston prism you can afford to buy. Those things are expensive. But they've got very small fields in the view, so [INAUDIBLE].

And also, we'll talk about spatial heterodyne interferometry in a little bit, it's to do with the angle between these two beams here. And that angle is set by your Wollaston prism. And so your fringe pattern, which we'll talk more about later, is fixed. You don't have any ability to modify it, which you do with the other two diagnostics here.

So this is the sort of thing you'd use on a femtosecond laser experiment, or picosecond laser experiment where you really, really need to have an ultrastable interferometer. But you wouldn't necessarily use this on many other sorts of assets. Any questions on that-- oh, yes. Many, OK.

AUDIENCE: So in the Nomarski setup, it seems like you could only get out of spatially averaged quantities, or about the electron temperature, since you have to basically encapsulate the whole plasma with your beam.

JACK HARE: Sorry, what did you say about the electron temperature?

AUDIENCE: I mean, sorry, the electron density.

JACK HARE: Oh, OK. Cool. Yes. Yeah.

So the question was, do you only get spatially average quantities effectively if your beam is like-- if you're looking into the beam, your beam is like this, and your plasma is about this size within it. You will be able to resolve all of the line integrated electron density within this region. So you'll get a map of Ne dl that's a function of x and y here.

As long as you can fit your entire plasma into the rest of the-- if you imagine translating if it fits inside the [INAUDIBLE], you'll still get a full image of the [INAUDIBLE]. So it's not averaged over x and y, it's obviously averaged in the probing direction, z in this case.

This might make more sense when we talk about the spatial heterodyne interferometry, and I realize now it was a mistake to introduce some of this stuff before we talked about some of the advanced techniques. So maybe it'll make more sense later and we can talk about this. I saw another question there. Yeah?

AUDIENCE: Yeah. I Wanted to mention another geometry that hasn't been used a lot is the self-mixing interferometry, which has-- it uses diode lasers with a photodiode behind the laser so that the plasma beam just goes through-- the probing goes through the plasma and reflects off a mirror, and comes back into the same diode and mixes with the ray that goes to the other side of the laser diode.

JACK HARE: So you're talking about a system where you've got a laser like this. It sends out a beam and then it comes back. And where's the detector?

AUDIENCE: The detector is behind the laser [INAUDIBLE].

JACK HARE: Yeah. And then how do you get-- what's this interfering with?

AUDIENCE: That's interfering with the light coming from the resonator backwards.

JACK HARE: OK.

AUDIENCE: Right

JACK HARE: Yeah.

AUDIENCE: So it's a very simple geometry--

JACK HARE: This is a type of Michelson.

AUDIENCE: Kind of, yeah.

JACK HARE: It's a variation. I agree it's different, but it's a variation on the Michelson.

AUDIENCE: Yeah.

JACK HARE: Yeah. But you're right. I'm not saying that these are the only three interferometers. People have made a lot of them and I'm just-- this is detailing three popular ones. But yeah, you're absolutely right. That's something is interesting. I've not come across it before, so thank you. I saw a question online. Nigel?

AUDIENCE: Which two beams are actually interfering with each other in the Nomarski setup?

JACK HARE: Yeah. Again, I think I should have introduced this later on. So bear with me, it will make more sense later on. I'll try and circle back to the Nomarski once we've done a bit of the lecture that explains what's going on. But thank you. Yeah?

AUDIENCE: What does fringe fixed mean again?

JACK HARE: Again, we're going to talk about what fringes do, or what fringes are later on. But effectively, you are fixed in a single configuration. And so you can only measure-- yeah, you don't have very much flexibility to change how your interferometer is set up.

So really, fixed fringes is meant to be in comparison to the flexibility of the Mach-Zehnder. And to be honest, the Michelson is pretty flexible as well. But we'll get on to what fringes are and how we change the fringe pattern.

And in fact, that's a big part of the problem set as well, is learning about what these look like by doing some synthetic diagnostics for them. Were there any other questions? Yes.

AUDIENCE: What's the name of the polarizing material?

JACK HARE: Oh, I think this is-- I'm almost certain that this is a Wollaston prism. And I think it is two slices of some birefringent material like calcite that have been cut and then rotated and glued back together in such a way that they do this splitting.

It actually was a basis for how the Vikings navigated during cloudy days across the planet. But we'll talk about that when we do Faraday notation. Any other questions?

All right. Let us get into some of the ambiguity that has already been referenced. So I think this has come up a few times.

You probably know what the answer is to this. But say I'm your advisor, or maybe a diagnostician of D3D or something like that, and come and give you a time trace on the interferometer. And I go. What's the density? Can you tell me?

Well, what's the phase shift maybe is a more reasonable question. So is it obvious? So let's draw out a sort of a tree diagram of different phases it could be.

So this is time here. This is delta phi. I'm going to be nice. I'm going to say, let's assume we started at 0.

So they're going back, this signal was like that. And then the plasma started here at t equals 0. And then this signal started wobbling.

So this looks like, to the best of my drawing ability, some sort of cosine function. And it seems to be a cosine. And again, to the best of my drawing ability, maybe these are all evenly spaced peaks.

And so maybe this is just a drawing of cosine delta phi where delta phi is just equal to some constant times time. It's just linearly ramping up, like that. Certainly one possible trajectory. Any other possible trajectories?

AUDIENCE: Could be negative.

JACK HARE: Could be negative. I have no way of telling the difference between those two. Now let's put-- let's start making this interesting.

Let's say this is pi and this is 2 pi here because it's minus pi. This is minus 2 pi. Any idea is for other trajectories now? Yeah.

AUDIENCE: I think if you add some portable sawtooth thing, you're going down.

JACK HARE: Yeah, it could start going down. Or no, it's not going out. What happens here? Do I just continue? What else could I do?

AUDIENCE: [INAUDIBLE]

JACK HARE: I could go down. So here, I have a choice to go up like this. Here, I've got a choice to go here, like this.

Here, I've got a choice to go like this or like that. I also have a choice coming from the other way to go like this like this, and it's now like driving in Boston. It's getting extremely complicated.

And I can just keep playing this game. I don't have to stop any time. I am going to insist on causality. I don't think a plasma goes back in time. So we have to worry about that.

But in general, this gets very complicated. You cannot tell me what delta phi, and therefore the density, is doing. You're quite right. A solution to this is a linearly ramping up density. Another solution is a sawtooth. It could be something that goes like this.

Hard to get negative density, but something like that where you start with some background density here, and then density drops and goes back up. There's lots and lots of different solutions. And so the problem is that for all of these five t tracks, we end up with the same intensity signal.

And so this is the inverse problem. We want to go from the data on our diagnostic back to what the plasma is doing. And we can't do it, because it's not well-posed. There are multiple different solutions.

This is not completely hopeless, right? We are physicists, and so we have some intuition or priors about the world. So we might be able to use our-- I'll call them priors because that makes it sound like I'm some Bayesian person, but think of them just as intuition.

So you might say, well, a reasonable thing for any plasma to do would be to not exist, exist, and then not exist again, like this. Some sort of shape that comes from simulations. Something like that.

And so if you can't, the fringes here, eventually they will stop. I mean, we assume that plasma won't go on forever. And then maybe we could pair up all of with these, and count to make sure we got an even number, an odd number of that, and then we can reconstruct this. But of course, there's no guarantee that in the middle, it doesn't do that. So it may have a more complicated shape.

Again, we could just do a simulation. And from our simulation, we could get a synthetic i of t. And then we can go, hey, my simulation says this, my data, the split looks the same. Therefore, it's the same as a simulation.

But you haven't proved it. You just said it's consistent with your simulation. You can't prove anything using this technique or what we're going to do in a moment, which are advanced techniques.

And in the problem set, you will encounter at least two other advanced techniques which we'll not be covering in class, which help to try and avoid this phase ambiguity. We're going to jump straight to the gold standard in my opinion. But those other two techniques are significantly cheaper [INAUDIBLE].

And so if you can get away with doing them, you would do those instead. So any questions on the phase ambiguity before we jump into the really good stuff?

AUDIENCE: Quick question.

JACK HARE: Yeah?

AUDIENCE: So is there not information about the phase in the incoming light at all, then? But isn't there usually-- can't you find some phase by polarizations and such in the light?

JACK HARE: I don't think the phase is linked to the polarization, no. I can have any arbitrary polarization with any arbitrary phase. They're not linked quantities.

AUDIENCE: OK.

JACK HARE: But if we had a detector that was fast enough, then we could actually measure the phase kx minus omega t. But we can't do that. We can only measure the phase shift with respect to our reference beam. And that is, with this setup, ambiguous.

AUDIENCE: OK, I see.

JACK HARE: Was there another question? All right. Probably this.

Rishi looked into getting one of these fancy AI cameras that would rotate and follow me around the room so I could use all of these boards and still broadcast. But the reviews were not very good. Seems like it should be easy, right?

Isn't there something that can just point at me all the time and I can wander around? But anyway, anyone needs a start-up idea and thinking of dropping out, that's one for free. Really, really huge market for egotistical academics who want to be filmed all the time.

AUDIENCE: Probably need it to work [INAUDIBLE].

JACK HARE: Yeah. The motion capture [INAUDIBLE]. Oh well, it can't be more ridiculous than what I normally wear.

So what we're going to talk about is a technique called heterodyning. Heterodyning. Or we might call this a heterodyne system.

And heterodyning, or a heterodyne system, this is very similar to how an FM radio works. So FM radio, anyone what the F and the M stand for? Radio is this thing we had before podcasts. Yes?

AUDIENCE: Frequency modulated.

JACK HARE: Frequency modulated. OK, cool. And so what we're going to be doing is using some tricks to separate the signal that we want, delta phi, from some larger background signal, which is the reason we use FM. But in this case, we're also using this trick because it helps us distinguish between the phase going down and the phase going up, and also the phase rolling over 2 pi. And so these techniques are extremely powerful.

And the way that we do that is we notice that we've got this phase here. And remember, this phase shows up in a function that looks like exponential of ik.x minus omega t plus delta phi.

By the way, I'm probably going to drop the delta in most of this because it's-- oh, no. It looks like I've kept it. I've kept the delta. OK, ignore me.

So you notice when you look at this, delta phi, which should be inside these brackets because we multiply it by the imaginary unit. It looks an awful lot like a frequency times a time or a k times an x. And this means that we can say, what if we thought that we had some sort of frequency which was equal to the change in phase in time, or some sort of k that was proportional or equal to the gradient of the phase in space.

And then if we put these into this equation, they would actually end up looking a little bit like an effective k vector, or an effective frequency. So this is the basis for it. Hold that thought. And we're going to go see how that works and why this is a really, really good idea.

So let's have a think about a system now where we put some radiation through a plasma, it's got some frequency, omega 1 here, and we mix it with another beam which has a frequency omega 2. So you can think about these as the probe and the reference beam, but they've no longer got the same frequency. We'll talk about how you get that in the end.

And that means that the electric field that you have on the other side here has got some default electric field strength. But then it's got a cosine of omega 1t plus delta phi, that's the one that's gone through the plasma, and the cosine of omega 2t. That's the one that's gone around the plasma. So one's gone through the plasma, has picked up our standard phase shift. But we're also keeping track of the fact that in the time it's taken for them to go through this plasma, we've had a different number of oscillations because one of them is oscillating at omega 1, that everyone's oscillating at omega 2.

We do the same thing again and we work out the intensity here. And we get out cosine squared of omega 1t plus delta phi plus cosine squared of omega 2t. So again, I'm just squaring these brackets here. It's not particularly complicated.

And then we have the cross-term here, 2 cosine of omega 1t plus delta phi cosine of omega 2t. And then we squint at this for a little bit and we realize we know some very clever trig identities, and we end up with this formula, which looks like cosine of squared of omega 1t. Sorry, just cosine.

Just cosine. Cosine of-- I'm going to write this on the next line up [INAUDIBLE]. Cosine of a term that looks like omega 1 minus omega 2t plus delta phi and a second term that looks like cosine of omega 1 plus omega 2t plus delta phi.

And probably a lot of you have seen this before when you've looked at wave mixing because you see that we have a term, which is the difference term and we have a term, which is oscillating at the sum frequency. And sometimes this difference term is called the beat frequency because this is the sort of slow, modulating beat where if you play two notes slightly out of tune, that's the beat frequency that you hear. And in fact, that analogy is quite good because in general, we can detect the beat frequency.

Our detectors are fast enough to detect omega 1 minus omega 2, assuming that omega 1 and omega 2 are quite close. But our detectors are still not fast enough to detect this sum frequency. And so when we do averaging over a few periods, this sum frequency on any reasonable detector is just going to be averaged out and it's just going to be 0. And this is the term that we're going to be able to measure here.

Let's just do a little check for our sanity. If we say that omega 1 is equal to omega 2, we've gone back to one of our standard homodyne-- not heterodyne, homodyne interferometers. And our homodyne interferometer is going to have i is equal to 1 plus cosine phi. Don't ask me where the 1 came from. It was spotted [INAUDIBLE].

I'll have to check that later. The 1 shows up somewhere, I'm confident. I'm pretty certain it doesn't show up here because the average of a cosine is just 0.

I'm wondering if it shows up somewhere between these steps as just the 1 at the front there. I suspect it does. But I can't prove that. Suspect you'll end up with a cos squared omega t plus a sine squared omega t, and that will give you a 1.

But in the case that we're interested in, because we don't want to just simply reproduce what we derive before with much more algebra, we would get out-- for different frequencies, we would have i is equal to 1 plus-- there should be a delta phi up there-- cosine of omega 1 minus omega 2, and then a term that we're going to say looks like the time rate of change of the face, which we said is a frequency-like quantity, it's certainly got the right units, all times time, like this.

So this detector, when the phase is 0-- when d phi dt is just 0, so say this is a phi like this in time, this is intensity. A detector is just going to chug along, recording a signal which is oscillating at the beat frequency omega 1 minus omega 2.

But when that phase begins to change, this frequency itself is going to begin to change. And what we tend to do is we set our beat frequency, omega 1 minus omega 2, to be much larger than the effective frequency we get in the time rate of change of the phase. And now let me draw you a sketch of what might be happening here and then we can have a chat about what's going on.

So let's imagine that we've got some time trace of density and it's just something peaked like this. And then we think, what are we going to see on our interferometer here? Like I said, in the absence of any plasma here, our interferometer would just have a signal looks like this. So this is where Ne equals 0, so therefore delta phi is 0.

But when we add the plasma in, we're going to see something very interesting. We're going to start off when the gradients are small with exactly the same signal. But as we get to the point where the gradients are large here, this frequency is going to start increasing. We're going to have a faster signal.

Now, around the point here where dn dt is 0, we should end up with the same frequency, which is pretty hard to draw, but I'm trying to draw the inverse of this one here. And then what happens here is dn dt is now negative. And instead of having fast waves here, we're going to have a more slowly oscillating system. So it's going to look like that before eventually following the Ne equals 0 signal here.

So here, we get oscillations, which are closer, so higher omega. And here we get oscillations which are further apart, corresponding to lower omega. And I may have screwed this up by getting my sine wave wrong, which I have, because the phase is negative of the electron density.

So I apologize for that. But you get the idea. If I redrew this and I had the further spaced fringes here as the density ramps up, and the farther apart-- the closer spaced ones here [INAUDIBLE]. And you have a chance to get the sign right when you do this in your problem set as well.

And I will say one more thing, and then I'll take questions, of which I'm sure there are many. We now have an ambiguous signal in the sense that the signal you get from d phi dt are less than 0 is not equal to the signal we get when d phi dt is greater than 0. And if you remember, that was the problem that we had back here, where we had this ambiguity where we couldn't tell where the fire was going up or down when we reached these inflection points.

And we've now solved it. So this is great. OK, questions. Yes?

AUDIENCE: So if you're writing d phi dt, is that actually d delta phi by dt?

JACK HARE: Yeah, sorry. I think I did drop it somewhere. Let's put a little delta over here. Might have missed it in a few other places. There's one there. Sorry.

Should have been delta phi. So this is the change in the difference between the reference and the probe beam. So sort of few differences in deltas and stuff like that going on. Yeah, another question.

AUDIENCE: [INAUDIBLE]

JACK HARE: No, absolutely not. It works well with both, but it certainly works well with the CW system maybe better. You would like to probably have the length of your pulse be longer than the lifetime of the plasma, because you'd like to be able to see this increase-- at least the increase in density here. If you started your interferometer here-- sorry?

AUDIENCE: That's in time [INAUDIBLE]?

JACK HARE: This is in time, yes. I'm doing-- sorry, this is the temporary heterodyne version. We'll do the spatial one later. So temporally heterodyne interferometry. Fun to say. OK, other questions. Yeah.

AUDIENCE: If we have two different frequencies of light, how do we get around the coherence length effects that we were talking about before when we had only one frequency?

JACK HARE: The question was, if we have two different frequencies of light, how do we get around the interference effects that we had before, the coherence effects? We will show in the next page of notes how we derive two different frequencies from the same source.

And then they will have the same coherence. But we will-- I should be very clear that the difference between the frequencies, we can call it delta omega, is very small. So this is like detuning your laser by one part in 10 to the 8, or something like that. It's extraordinarily small frequency changes. You don't need very much to do this technique.

AUDIENCE: But still the same source?

JACK HARE: And so what we'll show is that you can use the same source and so we'll have the same coherence properties. But we will shift one of them in frequency very slightly. I see a question online.

AUDIENCE: Hi, yeah. Could you just reiterate over the relationship between the density and the intensity that you drew on the right? So the peak density right there, you have it-- it's as if there is no phase change?

JACK HARE: Yeah. So at this peak density, there is no phase change. And so the frequency of the wave will be the same as the frequency of the green wave, which is what drew as the background signal.

And because I've drawn the wrong number of oscillations, it's now got the opposite phase, but that doesn't matter. It's still got the same frequency. I tried to draw it so it sort of has the same wavelength here.

AUDIENCE: It's frequency--

JACK HARE: I say wavelength, but this is in time, so it's frequency. Yeah.

AUDIENCE: It's frequency, but it's inversed because--

JACK HARE: Yeah. It's got the-- I've tried to draw it so it has the same frequency. So in reality, if you actually fulfill this condition, omega 1 minus omega 2 is much greater than this, then I will show you from the data I managed to find online later, what your signal actually looks like is just like this, because it's oscillating incredibly quickly and there were just incredibly tiny changes to the fringe spacing.

But that's not very informative if I try and draw that on the board. So I've done a shitty version where I haven't actually fulfilled this criteria, but at least intuitively, you can see. And again, when you do the problem set, you'll be coding this up yourself. So you can have a play with changing this beat frequency compared to how much the phase is changing and you can see for yourself what criteria you need to get a good signal out and how this actually works. But it's pretty hard to draw it on the board, unfortunately.

But yeah, that's the idea, is-- what you're meant to get from this is the frequency is higher here, the same here, and lower here. And that's directly related to the gradients of the electron with respect to time. Yeah.

AUDIENCE: Got it. Thank you.

JACK HARE: You're welcome. Sorry, go ahead.

AUDIENCE: Sorry if I missed this earlier, is there a kind of limit on the frequency of the density changes in time that we can resolve?

JACK HARE: You can-- so if your density is changing rapidly, you need to have a higher beat frequency to be able to resolve it. So you need to keep this going. So if you have a very rapidly changing density, like on the nanosecond timescale, then you need to have a beat frequency which is larger than that.

And then your trouble is your digitizer. So if you want to resolve that beat frequency, we actually have to do here. Previously, I was like, hey, we don't care about the frequencies because we'll just average it out, but you do need to resolve the beat frequency because you do need to see this. Then you have to get a very expensive digitizer.

If you can get a 50 gigahertz digitizer, you're pretty happy these days, and that's going to cost you tens of thousands of dollars. So they do this at Sandia National Labs for a related technique. And there, they have 24 channels at 50 gigahertz and it costs them millions of dollars to do it. So this is completely-- those sort of timescales are completely out of reach of a university level.

If you're doing this on a tokamak, you can do it much more slowly and then you can use a megahertz scope, which you can get for like $200 off eBay. So then that would be acceptable. So yeah, when I'm teaching the course, it's meant to be our principles, so I don't talk about technology too much.

But every now and again, there are some really hard technological limits that means that these techniques will work in some regimes and not in other regimes as well. You also need a detector that can detect that frequency. So we have good detectors in the infrared, and in the gigahertz, and in the visible, but we don't have good detectors, really, for x-rays, and we don't have good detectors in the terahertz range at the moment. We're getting better. So there are some other technological limits. Yeah, another question.

AUDIENCE: Yeah. So you mentioned that this is an advanced technique. To me, this seems like resolving what is a pretty substantial flaw with the simpler techniques. So does anyone in reality implement the simpler techniques, or is this pretty much what everyone has to do?

JACK HARE: Well, the question is, does everyone do template heterodyne interferometry or do they do the homodyne interferometry. I don't think most people will do homodyne interferometry.

Did I lose it? I think it's [INAUDIBLE]. I guess this is a sample homodyne signal. You could imagine if you only had very small phase shifts less than 2 pi, then homodyne would be fine. Like if you knew off the top of your head or from other experiments that the signal just did that and never got to the ambiguous phase down here, then homodyne is fine.

If you don't have a lot of funding, then that's fine as well. No, I'm serious, because this heterodyne technique is very expensive and you've got to be able to generate two different, very closely and precisely tuned frequencies. So you've got to be very stable frequency sources because if they start drifting with respect to each other during the experiment, you can't tell whether that's due to plasma or due to drift.

And so this is a very expensive technique. And in the problem set, you'll come across two other techniques called quadrature and triature, which actually resolve this phase ambiguity. Quadrature almost resolves it but it has a pathological case, which hopefully you'll find.

Triature does resolve it, and that only requires two or three slow detectors as opposed to one very fast detector. So there's good reasons for using those simpler techniques. Yeah. Yeah?

AUDIENCE: This one doesn't really-- it's kind of a prior, but is in most cases, the shift from running upwards in density and then running downwards, is that going to be necessarily the point where the phase is exactly pi or 2 pi? So it's going to be-- you're going to be able to determine that here, my phase shift started going backwards from just the homodyne image. So you're not-- you can't tell by looking at it that at some point, you can guess that it's probably the point where the [INAUDIBLE] started going backwards.

JACK HARE: You said you can guess that it's probably-- so you're talking-- I agree, about price. So you can use this technique if you have strong priors and this is one of the reasons why these are actually relatively hard to automate. You often need a human to have a look at the signal.

But I completely agree, you can do a lot of stuff if you have some strong priors about what your plasma is doing. If you are trying to measure density fluctuations in the edge of a tokamak, I would argue that it's very hard to have good priors for what that is. And so this is a very difficult technique.

If you're trying to measure the afterglow of a plasma discharge where it exponentially decays, then yeah, it's fine. But yeah, I agree. You can get around these with a little bit of extra knowledge.

But if you don't have any good priors, if you're too stupid to have any intuition, you can spend a lot of money and resolve that ambiguity entirely. So that's always a nice thing in diagnostics. So any other questions? Yeah?

AUDIENCE: When you started this discussion [INAUDIBLE]. But why don't we write that down again? And how [INAUDIBLE] omega 1 and 2?

JACK HARE: It's over here. So we had-- if I skipped a line here, it was because this bracket is really omega 1 minus omega 2 times t. I think we can agree on that.

It's actually-- sorry, it's up here. So we definitely have omega 1 minus omega 2 times t. What's interesting is outside the times t, we still have this delta phi phase. But if delta phi is changing in time, the oscillations which it produces in i are going to look like the oscillations that we produced if we treated d phi dt a frequency. And I think that's kind of a difficult thing to get your head around.

But like if this changes in time, you're going to see changes in i that look like oscillations, and those oscillations are going to look like the same oscillations you'd get from having a frequency. So effectively, this forms a third frequency omega 3, this one here, which is mixed in with the other two frequencies.

AUDIENCE: So do we just drop k, or?

JACK HARE: Oh. That's spatially heterodyne interferometry. We'll talk about that next.

AUDIENCE: So if you think of this demonstration [INAUDIBLE]?

JACK HARE: Yeah. So what we're doing here is-- what we're doing here is we just have a single ray of light going through the plasma and the detector, which is not an imaging detector. It's just a diode, which is recording a signal in time.

If we want to use this and we want to measure gradients in phi, then we need to have an imaging detector like a camera. And it's likely that if we do have an imaging detector like a camera, we will not have the resolution on it to detect frequency changes. We'll get one picture or something like that.

And so we usually either do this temporary heterodyne version or the spatially heterodyne version. And when we do spatially heterodynes, we're not mixing in omega 1 and omega 2, we're mixing in k1 and k2. And we'll talk about that in a moment. They're mathematically identical, it's very beautiful.

This one is simpler because we don't have any dot products in it. So I'd like to start with this one. And this is also the one that the tokamak folks are more likely to use.

But in my research, we used the spatial heterodyne version like this, which gives you pictures of the plasma and the electron density within it. Any other questions? Let's keep going. Maybe I should have drawn this earlier.

So actually the way you can interpret this temporary heterodyne system is you're going to be doing Fourier transforms on this signal. And you'll be doing Fourier transforms of small windows so that you get the frequency within a small window because obviously, the frequency of your blue curve is changing in time and that's what you want to detect. So you can plot this, thinking in frequency space.

So we take our Fourier transform of i of t and we get out some frequency space intensity. For a homodyne system, we have omega here and 0 at the center here, we might, for example, have a signal here which corresponds to d phi dt.

So on some short time window, this is the frequency induced by the changing the phase. But we can't tell the difference between this signal and one at a negative frequency. So when we look at the Fourier picture, we see why there's this fundamental ambiguity. So this is for a homodyne system. And the equivalence of these two is our ambiguity.

When we have a heterodyne system, we have 0 here and omega here, we are mixing in our phase signal with this beat frequency. And so this beat frequency, omega 1 minus omega 2, means that if omega 1 minus omega 2 is greater than d phi dt, then the signal that we would get from plus d phi dt now occurs at a different frequency from the signal that we would get at minus d phi dt. And so you can clearly resolve whether you got plus d phi dt or minus d phi dt here. So this is the heterodyning system.

And this is where the link to FM radio comes in. In FM radio, we have a high frequency carrier signal that's equivalent to our beat frequency. And the audio is encoded as a modulation to that high frequency carrier signal.

For FM, it's so that we can filter out around that carrier signal, and those carrier signals transmit better. In this case, that does help with the signal to noise issues, but also more importantly, it means that, again, we can tell whether d phi dt is positive or negative, whereas in this case, it would give us, with the symmetry of the Fourier transform, the equivalent signal to if we had a negative frequency so that [INAUDIBLE].

So hopefully that helps those of you who like Fourier transforms to think about it. If not, maybe something to ponder [INAUDIBLE]. So just a quick note on practicalities.

So practicalities, in this case, is generating omega 1 and omega 2. So as we already said, we really want just one source so that we can keep coherence. So one thing you can do is you can put in-- you could split off your incoming laser at omega 1 with a beam splitter and you can reflect half the light off a mirror, which you are rapidly accelerating towards or away from the beam. And then we would have omega 2 is equal to omega 1 plus v upon c. We're able to get the mirror to go at v.

Obviously, this is only going to work for relatively low frequencies where we only want a relatively low frequency. We can't have this mirror going-- moving all velocities. But for small frequencies, the Doppler shift on that will be pretty significant and that works pretty well.

The trouble is, of course, presumably, your mirror has to stop somewhere. And so you have a limited amount of time you can accelerate the mirror for-- move the mirror in one direction before it stops, and so that will limit the length of your pulse.

If you want a continuous version of this, you have a reflecting wheel that is going at some frequency, v, like this, and you bounce your light off it. Omega 1 comes in, omega 2 comes out like that. And we find that omega 2 here is going to be equal to omega 1 1 plus b over c over 1 minus v over c, because it actually gets Doppler shifted twice once on the way in, once on the way out. This should probably be more like a grazing reflection in order to be able to use this. Otherwise, there'd be some geometric correction to it.

And then the final thing you can do is this really neat technique which is called an acoustoptical modulator. These are useful for very many different things, and effectively, it's a material that has strong sound waves inside it. And you put in one of your beams of light and when it reflects off, it's undergone an interaction with the sound waves here. So we have photon or phonon scattering.

And depending on the frequency of that phonon, we can get out a different frequency here. So those are three different ways that you can get out different frequencies that you can then use these heterodyne techniques, but all using just one source here.

So there would be a beam splitter here that takes off some of the omega 1's go through the plasma and the other part of the omega 1 reflects off the wheel, and then that gets interfered and becomes this one. Yeah, any questions on any of these? Yes?

AUDIENCE: The heterodyne system, do you determine the path length just by tuning where you're putting your beam and beam pointer, or is there some systematic way of--

JACK HARE: Oh, for a place where you want to get within the coherence length?

AUDIENCE: Right.

JACK HARE: Yeah, you'll have to tune it. I mean, what I would do is I'd measure it by hand very closely first and then I would move an optic back and forth until I got a good signal on my system. Or you can do a lot of this in fiber. Depending on if you want to use telecom wavelengths, 1,550 nanometers, you can use fiber optics and stuff like that, and then you know your fibers are the right length and so that's very convenient.

So if you're doing CW, the coherence length on CW beams, continuous wave beams, is extremely long. It could be kilometers on a telecom fiber. So you only have to get it right within a kilometer, which is usually easy in a university lab.

AUDIENCE: Yeah. OK.

JACK HARE: Other questions? All Right. Let's do some spatial heterodyning. I had some pictures to show but we haven't got far enough yet. So I'll have to show the pictures next time. There are some nice ones in Hutchinson.

So now we're not dealing with signals that vary in time. We're interested in signals at a single time, but now making an image of the plasma. And I warned you that I've been drawing lots of wave diagrams, and would eventually come back to drawing wavefront diagrams. And here I'm going to make good on my threat.

So here is our probe beam coming in. Now in reality, this is going to be a two-dimensional system. So if I look dead on to the plasma, I will see the plasma here, and I'll see some probing radiation source, which ideally would be bigger than the plasma for reasons we'll discuss later on.

But I'm just drawing it from the side and I'm drawing it in 1D. But you can imagine I could draw these as square phase fronts coming in. We're not going to do that, just wanted to give you an idea. We can do all of this just thinking about it as a 1D system.

And we've split some of this light and we have sent it like this, around the plasma. So this is coming in k2. This one is going k1.

The k's is here are vectors, and these k's refer to the direction the radiation is traveling in. They refer to the normals, the phase fronts. And so this is-- yeah, this is a vector system in this case here.

And then this is our beam splitter. And just to test you, I've drawn it wrong. And of course, it needs to reflect off like this. Maybe I should put these reference beams in a different color. Might make my life easier later on.

In the absence of any plasma, the beam is just going to go through and the phase fronts will remain completely flat. And then on our detector, which is somewhere back here, you'll get a nice overlap between the perfectly flat phase front that has not gone through any plasma, could have done, and a perfectly flat phase front that hasn't gone through any plasma because it's the reference one here. So this is the reference.

So my detector, I just have the same phase. Maybe they're in phase. We get constructive interference.

And so the image, if I expand this detector out like this, it would just be one uniform color. One uniform shade of gray. And that uniform shade of gray would be 2, in the sense that it's twice as bright as either-- as what you would expect if you didn't know anything about these things.

Now let's put the plasma in and let's see how the plasma distorts to bright. Remember that if there's any density inside here, our phase velocity is actually faster. Our refractive index is less than 1. N is equal to 1 minus Ne over 2Nc.

So when this phase front encounters the plasma, it will start to go faster. And in fact, it will start to distort outwards like this. And then when it exits the plasma, it will still be going faster like that.

Now when it reaches this detector, the reference beam is still going to be the same, but our probe beam is going to have a delta phi which is larger. And that delta phi is going to be different in different places.

There's going to be a delta phi here, which is relatively small, and a delta phi here, which is relatively large. And so what that might look like in our detector for a homodyne system, and this is still homodyne at this point, is we might have fringes like this, regions of lightness which are just concentric like that. And the trouble with these fringes is exactly the same problem we had with the homodyne, temporally heterodyne interferometer.

If you're trying to take a line out across here and you have some signal like this in intensity, and now instead of time, this is some spatial coordinate, you can't tell me whether the density goes like that or whether it goes like this. Because you can't tell me each time I get to a constructive interference fringe, whether that's because the phase has gone up by 2 pi or back down by 2 pi, or even worse, has stayed the same. So this, this is no plasma, this is the homodyne system, and now finally, we can consider the heterodyne system.

So in the heterodyne system, we tilt these fringes. And so we have a slightly offset A2. That means that we have tilted fringes coming into our detector.

So in the absence of any plasma, but this is now at third plasma, we will still have interference fringes because we've tilted one phase front with respect to the other. And we'll just have straight interference fringes like this.

But when we put the plasma in, these fringes will now be bent. So maybe they'll still be straight at the edges at the outsides of the beam where there is no plasma. But in the middle of the beam, like this.

And can you see once again that the fringes are getting further spaced apart and bunched up together again? And this corresponds to regions where there is high density gradient and regions where there is a low density gradient here.

And so this looks like if I take a line out in this direction here, I would end up with a signal that looks an awful lot like the one, and if I've still got it. That signal there that's hiding underneath.

So it might look like something that has big fringes and then very rapid fringes like that. So this would be intensity. And again, this would be some coordinate like y, not time.

But you can see there's a very deep mathematical link between the two of them. If I take a slice through any of my spacey heterodynes, interferograms, I'll get something that looks like a temporary heterodyne interferogram.

I am running over. I will take some questions, but we will pick this up next lecture, do a review of all of this. Hopefully, it will start to make sense in [INAUDIBLE]. But yeah, questions.

Anyone online? OK. We will pause there. You can spend the weekend thinking about heterodyne interferometry.

Come back fresh on Tuesday and it'll all make sense, I promise. OK, thank you very much. See you guys later.