[SQUEAKING]

[RUSTLING]

[CLICKING]

JACK HARE: So we talked about interferometry. We had some sort of plasma like this. And we had a probing laser beam that went through the plasma. And we split off a fraction of that probing laser beam and sent it around the plasma, which meant that when the two beams recombined, we got some interference effect between them. And we call these beams the probe and the reference.

And we said-- or we found that the phase difference between these two beams, delta phi, which is the phase accumulated by the probe beam minus the phase accumulated by the reference beam, that was going to be minus omega over 2c times the critical density, which is itself a function of the wavelength of the laser times the integral of the line integrated electron density. So the electron density inside this plasma, and e over some plasma length L like that.

And so this is the phase difference between the probe and the reference. We want to measure that. We came up with a simple system where the intensity on our detector was just going to be equal to 1 plus the cosine of this phase here.

And we realized very quickly that this causes us some problems, because if we have a signal on our detector of 1 plus cos delta phi, and that signal looks like something like this, so some sort of 1 plus cosine of some constant a times t like that, something that's oscillating, then we have a lot of phase ambiguity in the sense that we don't know whether the phase is going up and down. And we also can't measure the phase more than modulo 2 pi.

And for this example here, we had a look at what possible paths we could take in our little delta phi space that would give us exactly the same signal. We said we could have one that ramps up like this. We could also have a delta phi that goes down like that. And then every time we get to some multiple of pi, we lose track of whether we are going up or down.

And so we start having these multiple branching pathways and any possible path through this space is valid. It will produce the same signal on our detector. And we can't tell the difference between them. So this we then rechristened a homodyne technique. And then we looked at heterodyne techniques instead.

So when we started working with heterodyne technique, we borrowed some tricks from FM radio transmission. And we now have going through our plasma some radiation source, which has got a frequency omega 1. And now our reference beam has a frequency omega 2. So they've got some frequency shift between them. We talked about techniques for doing that.

We put in some recombining beam splitter here. And we put it through to our detector. And we said, if omega 1 is equal to omega 2, we just get back our homodyne system, which is kind of obvious, because here we split it. So they will have the same frequency.

But in the more interesting case where omega 1 is not equal to omega 2, we'll end up with two frequencies present in our final signal. We'll have a signal that on our detector-- even the absence of any plasma. So if we just leave the system running, we'll have a signal that looks a little bit like this.

And there will be two frequencies inside this. There will be this envelope frequency called the beat frequency, which oscillates at the difference between the two, omega 1 minus omega 2. And then there will also be this fast frequency within it at omega 1 plus omega 2.

Now, these frequencies for any radiation we're likely to use are very high. So it's very hard to get a detector to work at these frequencies. So we won't actually see this at all. Our detector will just average this out. And what we'll see is this slow beat frequency instead. And we can detune omega 1 from omega 2 to get a beat frequency that nicely falls within our detector range.

So far so good, but there's no plasma physics in here. What we'll end up then is measuring something that looks like i/i0 equal to 1 plus cosine of omega 1 minus omega 2 t. That's the beat frequency. But we will also have an additional phase term, delta phi, as we had before, because that's what represents the phase of going through this plasma here, delta phi.

And when we looked at this, we said, huh, that's interesting. We've got frequencies times time, and now we've just got this other term, delta phi, which means that if delta phi changes in time, we have a change in delta phi with respect to time. That's going to look like an effective frequency inside here.

And then we can rewrite this as 1 plus cosine of omega 1 minus omega 2 plus partial of delta phi partial t times time like that. And what we actually measure on our detector at the end of the day is a signal, which is oscillating with some frequency, which I'll call omega prime, which is the sum of these three different frequencies here, which means that if we see a change in the frequency omega prime, we know that change is due to a change in delta phi in the phase. And so now we can measure the temporal change in phase, which is the temporal change in the electron density.

The nice thing about this technique is that we found that now we can distinguish between the phase going up in time from the phase going down in time, which we weren't able to do with this technique. And that's why we got this ambiguity every time we went past pi, whether we were going up or down.

With this technique, it was resolved, and we saw that we resolved that by looking at this in Fourier space. So in the case of the homodyne technique, we were detecting, effectively, this frequency here at some positive omega. So that's d phi, partial delta phi, partial t. But that gave exactly the same result as some negative frequency here. So we couldn't tell the difference between the negative version of this and the positive version.

When we moved the heterodyne technique, because we are effectively encoding our fluctuation quantity around some frequency omega 1 minus omega 2. Now we can tell the difference between whether we're in the negative side of that or the positive side of that. So we've shifted this frequency a small distance by changing the phase by having some time changing electron density.

So this enabled us to resolve the ambiguity between these two and, indeed, this technique in general helps resolve a lot of the ambiguities associated with homodyne interferometry. Though as I mentioned in the problem set, you'll come across a couple of other techniques which can do this in a slightly cheaper, but more ambiguous way.

So we got through all of that. And I just want to pause here and see if there are any questions on that material before we go on and finish off the spatially heterodyne version where we do this as an imaging technique. So questions? Yeah.

AUDIENCE: So I understand that the detector is not able to resolve the omega 1 plus omega 2 component. But does the time change in the phase difference contribute to that oscillation as well if you were able to resolve it?

JACK HARE: Yes. But it would, if we look at it in a Fourier domain, I'm on another board. I'll just use this one briefly. I won't exaggerate this a lot.

Let's say we've got our two frequencies omega 1 and omega 2 close together, omega 1, omega 2 like that. The difference between them is down here. This is the beat frequency, omega 1 minus omega 2. And indeed, we have some shift to higher or lower frequencies. And that shift is due to the change in the phase in time.

The sum frequency is up here. This is omega 1 plus omega 2. And indeed, it will also be shifted by the same amount here. But it will be at such a high frequency still that there will be no way to measure it. If you happen to have a system where your phase changes so much that you can shift this one down into a range you can measure, then I don't think you need a heterodyne technique at all at that point.

But anyway, certainly, if that happens, you should make omega 1 and omega 2 larger until it doesn't happen, because you remember that we have this condition, but you only really get good results of this if delta phi delta t is much, much less than omega 1 minus omega 2, which necessarily means it's much, much less than omega 1 plus omega 2. You don't want your shifted frequency getting anywhere close to 0, because then you [INAUDIBLE] again. And you won't be able to measure it.

AUDIENCE: Thank you.

JACK HARE: Good question. Any other questions?

Well, I've either taught it very well or you're going to find the homework very hard. Let's go on to [? spatial ?] heterodyne techniques.

So here is the idea that we actually put an expanded beam through our plasma. So we take our laser beam, which maybe is initially relatively small, and through some beam expander, we get out a large laser beam. And we passed that. We've expanded our beam enough that our plasma is maybe slightly smaller than the beam diameter. So there are some regions outside the plasma that we can still image where there won't be any phase shifts. This turns out to be useful for zeroing our system. We'll talk about that a little bit later on.

And what we said is this beam, of course, consists of some wavefronts like this. And these wavefronts, as they go through the plasma, are going to advance, because the phase speed of the plasma is faster than the speed of light. And so the phase actually advances inside the plasma. And so the wavefront that comes out is also going to be advanced in respect to the plasma.

We want to measure the change in this wavefront. So one thing we could do is we could interfere it with a set of plane waves that we derive from the same laser beam. We'll put a beam splitter in somewhere down here. Send these around the plasma. Send it back in here. And then we'll have another beam splitter as we normally do that recombines these, and then we would have this nice flat phase front. And we would now have from the probe beam some phase front, which is advanced a little bit.

And if we put all of this image onto a detector-- here's our camera like this-- we get a series of constructive and destructive interference fringes that maybe look like this. Destroying a set of nested contours here, we assume we've got some sort of peaked central structure here.

Now, the trouble with this is this is still a homodyne system in the sense that we can't tell whether the density contours are going up or down. If I draw a random line out across this, my density could look like a peak structure. Maybe I've got a prior that that's true. But I can't prove to you that my density doesn't look like that instead, or any number of other different paths through phase space that will give us the same fringe pattern.

And so this is still problematic. So what we want is a spatially heterodyne version of our temporary heterodyne system that we had out there. And we do our spatial heterodyning by tilting these fringes. That literally means slightly adjusting this mirror here so that the fringes come through at an absolutely tiny angle. So we're not talking about 10 degrees here. We're talking about much less than a degree.

And that means that our phase fronts are now coming at some slight tilt here. And as opposed to having omega 1 minus omega 2, we now have an inbuilt phase pattern that looks like k1 minus k2. And these are vectors in the xy plane here. We don't care about the z components, where this is x and this is y. And by convention, we usually have the z coordinates going the direction of our rays here.

So we're obviously putting a camera here. We don't measure anything in z. We measure in y and x. And we're interested in the misalignment of these wavefronts in x and y. So in the absence of any plasma, this misalignment is simply going to give you a series of straight fringes evenly spaced like this.

And so that is like the signal that you have, your beat signal here, the green line. In the absence of any plasma, this signal just goes on and on and on. But when we introduce the plasma into our system, each of these fringes will get distorted. And they'll get distorted by an amount that corresponds to their line integrated electron density.

And by looking at the shift between where the fringe was before we added the plasma and where the fringe is afterwards, we can then calculate the amount of density that's being added because we know that the fringe shift is linear functional to the density here. And once again, because we've got a heterodyne technique here, we have avoided this ambiguity.

Even if these fringes overlap, even if we have a distortion so large that this fringe goes above this one, when we go out to the edge of the plasma, where the fringe shift is 0 out here, we can still uniquely identify each fringe with its background fringe. And so we can track them along, and we can say, aha, this one's done two fringe shifts or four fringe shifts.

And there is a Fourier transform way to think about this as well. But now we need to have a two-dimensional Fourier transform. And what we're looking at here are now kx and ky. And so, originally, your k1 minus k2 beat frequency is maybe up here. And by symmetry, because our signal [INAUDIBLE], down here. So just at the negative number as well.

And now we've distorted these fringes. It's taken this initial beat frequency. And maybe we now have Fourier components that look a little bit like this, or they could look like that. And moving around in Fourier space changes what the shape of your background changes look like. This here, where my components are roughly equal in kx and ky, that will correspond to fringes, which are 45 degrees.

So if I had my beat frequency down here, this would correspond to fringes with the k vector in that direction. And that would be like that. So you can choose what your carrier frequency is. And there was a good question in the last lecture about your sensitivity to different density gradients in different directions. And indeed, you have more sensitivity in the direction perpendicular to where your carrier frequency, spatial frequency is.

There's a lot going on there. I actually have a load of slides with pictures of real data on this that might make a little bit clearer. But before we get on to that, are there any questions? Yeah, uh, John.

AUDIENCE: You might have answered this. So by tilting the wavefronts of the reference beam, what in essence we're doing is we're changing the wave vector component that is interfering with what's coming through the plasma. I mean, we're trying to create some beam interference here. So I guess, technically, we're not changing the magnitude of k.

JACK HARE: No, because in free space, the magnitude of k is fixed-- precisely-- because the dispersion relationship for a wave in free space, not in a plasma, is omega equals ck.

AUDIENCE: And so what we've done by tilting the wave is we've changed the distribution of that magnitude in either dimension. So now the component of k that is interfering with the wavefronts is slightly different.

JACK HARE: Yeah. So k2 here is, for example, the reference beam here. And k1, it's been coming through, for example.

AUDIENCE: So K1 is exclusively an x component, I guess. I think about that if we align our--

JACK HARE: If we aligned in this way, but if I align the fringes-- if I rotate the fringes, I can measure other components of it as well. And it turns out, as I'll show you, you still get to measure some of the y components and things like that, even if you're in this setup, which is most sensitive to the x component. But you're right, yeah. Was there another question? Yeah.

AUDIENCE: So as the probe light propagates through the plasma, it'll refract and bend around in all this. Why don't you get heterodyning for free from these changes to the wave propagation as it transits the plasma?

JACK HARE: So those-- so the question was, why do you get heterodyning for free as you get k changing within the plasma here? I mean, that's effectively what you're measuring here with these phase contours. It's just that they're still ambiguous.

You need to shift them so that they're-- as we did in frequency space for the temporally heterodyne version, you need to shift them into such a direction that they're completely unambiguous, whether your phase shift is up or down. And here at the moment, even with this, you're going to get these ambiguous fringes. I'll show you some pictures that maybe will make it a bit clearer in a moment. Were there any questions online?

AUDIENCE: Hi, yeah.

JACK HARE: Yeah.

AUDIENCE: So how fine can you get for the amount of x components or y components in your scan? How detailed can you get? Because I'm assuming-- because the wavefront is continuous, but I'm assuming you can't get perfectly granular understanding of the plasma density just from this [INAUDIBLE].

JACK HARE: Did everyone in the room hear the question? It sort of limits the resolution in a way. So our spatial resolution is set by our fringe spacing. So usually, we can say, this is a dark destructive interference, and this is a light constructive interference. And theoretically, you could identify the gray point halfway between light and dark, but it starts getting a bit ambiguous there. Dark and light are pretty obvious.

And that means that your spatial resolution is set by the spacing between your fringes. And if I choose to have my fringes closer together-- so I choose a k1 minus k2 which is a larger number, so a higher beat frequency, then I'd have my fringes closer together like this. I would gain spatial resolution. And I'd be able to keep playing that game down to the resolution of my camera, where I need a certain number of pixels to be able to tell the difference between light and dark.

The trouble is, as I shrink this down, I'm gaining spatial resolution, but I am losing resolution of the density, because these fringe shifts are now smaller and smaller. The fringe is now only moving a very small distance. Maybe it's only moving two pixels or one pixel. Now I've got a 50% error, because I don't know whether it's two or one pixels.

So there is a tradeoff directly between the density resolution of this diagnostic and the spatial resolution of it. So that's a very good question. And that is the same. It's always the same, because mathematically it's the same. It's the same for the temporal heterodyne version as well.

So you can have time resolution, or you can have density resolution. But you can't have both. They trade off against each other.

AUDIENCE: I see. That makes sense. All right, thank you.

JACK HARE: Thank you. Yeah.

AUDIENCE: Can this be used for temporal measurements as well? You can just keep [INAUDIBLE], imaging--

JACK HARE: So can this be used for temporal measurements as well? Yes, if you have a fast camera. You can do this-- for example, I know a guy who used a CW laser beam, so basically a continuous wave laser beam at like 30 watts or something terrifying. And he had a fast camera. And so he took-- the fast camera could take 12 pictures, one every 5 nanoseconds. And he was able to make a little movie.

And depending on the speed of your plasma, if you don't need every five nanoseconds, but you're working with a plasma where the timescale is milliseconds, then you can actually just have a continuous camera. So you need a nice, bright light source that is continuous enough and you need a camera which is fast enough. And that's what sets the resolution. So you can make a 2D movie of the density evolution in time.

AUDIENCE: So the limitation here is just technological?

JACK HARE: Yes. Yeah, yeah, yeah. This technique is-- I mean, that is time resolved, spatially heterodyned interferometry. I don't think you can do temporally and spatially heterodyne interferometry in the same diagnostic. But if you work out a way of doing it, let me know.

That sounds hard, maybe not necessary, because you've already got around the ambiguity in one way. So I don't know if you need both. And if you had a homodyne system, in one sense, I think you'd be able to use lack of phase ambiguity from the heterodyne part of the system to get over that. But I haven't really thought about it that much-- interesting question. It could be a fun diagnostic-- very expensive.

[LAUGHTER]

So yeah, cool. Any other questions? Yeah.

AUDIENCE: So in this case, it's not just we're measuring the k factor of how we're measuring the wave. We're actually deflecting the probe away a little.

JACK HARE: So the question is, is the probe wave actually being deflected as it goes through the plasma? Well done for spotting that. I was going to have that as a question later on when we looked at some data. But as you pointed out, you've ruined the game.

So I told you earlier that rays are always perpendicular to the phase front. And so as I'm drawing this, the rays are like this. Fine. The phase fronts are flat. But you can see here that if I drew those lines like this, I would start to get deflection. And so you will have shadowgraphy effects overlaid on top of your interferometry signal.

It turns out the modulation from interferometry is usually stronger. And so you see that more strongly. But if these are very big, you may actually just lose the light, because your lens will be this big. And your rays will exit it.

Now, in general, interferometry is so sensitive that I've drawn this in a very exaggerated way. The phase fronts can be almost perfectly planar still, and I'll still get really nice interferometry patterns. But it won't be so distorted that the light will all be spreading outwards and I won't be able to do anything about it.

So yeah, you're quite right. In this picture, we should have shadowgraphy and [INAUDIBLE] and all sorts of things like that as well. Any other questions? I'll show you some pictures in a moment otherwise.

AUDIENCE: I have a question.

JACK HARE: Yes, please.

AUDIENCE: Because of the shifting of the wavefronts, is there a possibility for interference within the same wavefront? So would k1-- could it interfere with itself if there's enough of a shift in density, the plasma?

JACK HARE: In fact, that's when we were talking about shadowgraphy and I said, we don't really want a coherent light source for shadowgraphy. So even if you take out the reference beam, and the question is, can this shift so much that it actually interferes with itself? Yes, that happens.

And you can see that in shadowgraphy. And it's bad, because it's really hard to interpret. But there is a technique-- I should have read up on this more before saying this, but it is a technique called phase contrast imaging, which is used with X-rays. And that actually exploits the interference of the X-ray, which is just radiation, like all this other stuff, with itself to make very, very precise measurements of sharp density gradients.

So in general, you want to avoid coherence in shadowgraphy, because it messes up your data. And it's hard to interpret. But if you can do it very precisely, you can do some very nice techniques with it. So it's not always a curse.

AUDIENCE: Is there--

JACK HARE: In general, all of these effects will be overlaid on top of each other. I've just been presenting them one at a time. But they're all present in the same system.

So this is a very biased sample of interferograms. And it's biased in the sense that I just went through quite a lot of papers I've written and tried to grab them because I was in a hurry. But hopefully, some of these will be informative.

I tried very hard to find some temporary heterodyned interferometry. And it's actually quite hard to do-- find one where they show the raw signal, because we have such good electronics for this stuff these days that mostly you just do the signal processing on the chip and give the output of it. So you don't really digitize the raw signal.

So this is my best attempt so far. This was a HeNE laser beam. So that's a green laser beam on [INAUDIBLE]. And this is from a paper from 2017, so it's relatively recent. They had-- the HeNe is obviously green, but they used the heterodyne technique to produce a probe at 40.

That's the beat frequency there. And that effectively sets the temporal resolution of this. And they actually did something even more complicated than we've discussed where they heterodyned the system, and then interfered the heterodyne probe with the heterodyne for reference, which is very weird. And they did it with a quadrature system where they shifted one of the signals out of phase by 90 degrees. And you'll learn about quadrature in the problem set.

And then they digitized those two signals. And so what they saw was these signals here. And if you look at this carefully, which you can't do right now, but I'll put the slides up later, you'll find out that these signals are actually 90 degrees out of phase, which is really, really cool. And then they were able to process those together, and they could get out of phase shift.

And you can see the time on the bottom here is sort of in millisecond-ish time scale. So this is pretty fast for a tokamak. And you can see that the phase shift is going multiples of 2 pi. So they've resolved that ambiguity. So they're saying, look, the density went up, and then it came down. And it went up again. And then it went down. So they have some confidence that this is real. So this was quite a nice example.

Another example is from the pulsed power world. This was on the Z machine at Sandia, where they actually showed the raw data-- not quite the raw data. What they show is a spectrogram. So this is like if you do a very short time window Fourier transform on your temporal signal.

And you plot what frequency components are of single time. So if I take a slice at a certain time here, I can see a dominant frequency component down here. And you can see that dominant frequency component change in time. And they've shifted this so that the beat frequency is at 0. But that would actually be gigahertz or something like that. And this would be a shift from that beat frequency of gigahertz.

And they've done a technique where they shift the beat frequency in each window a different amount. And this gives them a much higher dynamic range. But effectively, this is looking at an increase in phase that's chirped in time. And the time scale on the bottom here is nanoseconds. So over about 100 nanoseconds, they've measured a significant phase shift.

Now, what they're doing with this technique is not actually measuring a plasma. They're measuring the motion of a conductor. So this is photonic Doppler velocimetry for the [INAUDIBLE] kids who have heard about that before.

And they're doing that to measure all sorts of cool squishing metal type things, but exactly the same physics at play, because that moving conductor just gives you a phase shift. And that phase shift could be density, or it could be some moving conductor. So it's up to you afterwards to interpret what the data looks like.

And they've run out of bandwidth up here. So 25 gigahertz, they can't sample any faster, because that's already a very expensive digitizer, which is why they have this clever technique which effectively [INAUDIBLE] the signal. So it goes up on one, and the beat frequency appears to go down on the other. And then when it hits this point here, it starts going back up.

And they do the same trick several times. And by sort of appropriately flipping and splicing these signals together, they'll actually get a signal that just keeps going up and up and up, and then they can measure this motion of this conductor over a very long time scale. So there's some very cool, advanced techniques in electronics involved in all of this.

But this, at least, is closest to the raw data. And again, for the p set, you'll be making your own raw data. So you can see what it looks like there. Any questions on these two temporally resolved techniques? Hmm?

AUDIENCE: In the [INAUDIBLE] example, did they have any kind of spatial resolution [INAUDIBLE]?

JACK HARE: No, it's just a cord. So it's a laser beam through the plasma. And that's pretty typical for tokamak plasmas. One reason for that, which doesn't apply here, is that you often want to use microwaves because the density is more appropriate for microwaves than for lasers. And so that means it's actually quite hard to do imaging with microwaves. We tend to just have an antenna which launches microwaves, an antenna which collects them. So you tend to just have a line.

With a [INAUDIBLE], that's not a limitation, but, obviously, they probably couldn't do a camera that is resolving on this time scale. Maybe they don't want to. They certainly couldn't have a camera that covered the entire tokamak cross-section on that time scale. So I think they went for this sort of time resolved, but just one point in space technique instead. We'll talk a little bit about how many cords like that you need in order to do some sort of reconstruction later on in the lecture. Other questions? Anything online?

So these are examples of spatially heterodyned interferograms. This is the case with no plasma. You see you've got these nice, uniformly spaced fringes here. And some of them are light. That's constructive interference. Some of them are dark. That's destructive interference.

So the probe beam is going straight into the page like this. And the reference beam is tilted at a tiny angle upwards. And we've chosen that angle to give us this nice fringe pattern here, because when we put a plasma in the way, we have plasma flows coming from the left and the right. And we can see that all these fringes are distorted. You can see most prominently the fringes all tick up in the center here. And that corresponds to an increase in the line-integrated electron density.

You can also see there are regions where there's quite large fringe shift distortions around here. These are actually for plasma sources on either side here. And you can see that some of the distortions are so large we've actually formed closed fringes again. So in that place, we have violated the condition that k1 minus k2 has to be much, much larger than the spatial derivative of the phase.

We've effectively recovered by accident the homodyne system, because we weren't able to keep our fringe spacing close enough together. We made this fringe space even closer. These closed fringes would go away, and we'd lose that ambiguity. But we'd also be sacrificing our dynamic resolution of the electron density.

So for an interferogram like this, I have very strong priors that the density is going to be higher here than here. And so I can just, when I'm doing the processing on this data, make sure that the density goes up here instead of down, effectively making choices on that decision tree that we had before.

And if you spend a little while processing these interferograms, this is the raw data, and this is the line-integrated electron density here. Now, for this, the electron density is in units of 10 to the 18th per centimeter cubed. So technically, the thing that you get out of this is line-integrated electron density. So that's per centimeter squared.

In this system, we have a lot of symmetry in this out-of-plane direction. And we knew how long the plasma was in that direction. So we just divided by that length to get the line-averaged plasma density.

And you can see, again, although we did have these homodyne regions here where we have some ambiguity about phase, because we had strong priors, we were able to assign the correct electron density. And if we decided incorrectly, if we said it's going down, we would see a weird hole here that wasn't on the other side. So it also helps to have a bit of symmetry in your system as well to check that you're assigning things correctly.

So I won't go into the details of how you process these, though, some vaguely involved techniques, but you can take data like that and get out some really nice pictures of the electron density in your plasma. So that was quite a nice one. You can see all the fringes are still roughly parallel. They only move a little bit.

Here's some interferograms with slightly more twisted fringes. You can see the background fringes are like this on this image, and they're like this on this image. It's just how it was set up in the two experiments. And you can see that this is an example I used earlier of a B-dot probe sticking in a plasma and a bow shock is forming around it.

And you can see strong distortions of the interference fringes, especially very close to the bow shock, where, in fact, the reflections of the rays are so large from the density gradients that they're lost from our imaging system. And so we don't have interference fringes here, because they've been refracted out of our system. We no longer can do interferometry. So when the density gradient gets too large, it's very hard to do interferometry.

And again, these images were processed, and you get nice pictures of the bow shock in these two cases here. And this paper was comparing bow shocks with magnetic fields aligned with the field of view and perpendicular the field of view. And we have very different geometries there.

So those are quite complicated. But probably the most complicated one I've ever seen traced was this one by George Swadling, 2013. So this is a 32-wire imploding aluminum Z-pinch. There's a scale bar up there, which is actually wrong. That's a millimeter. It says a centimeter. It should be a millimeter. No one's ever noticed that before.

And so these that are positioned of the 32 wires here, and there's plasma flowing inwards. And as the plasma flows inwards, it collides with the plasma flows from adjacent streams. And it forms a network of oblique shocks.

So this is two wires here. The first oblique shock is out of the field of view and forms a plasma here. And then these two plasmas interact, and they form another oblique shock structure. And then these two interact, and they form another oblique shock structure. And there may even be like a fourth or fifth generation in here.

So this is a complete mess. This is extremely complicated. But because the interferometry is very high quality, with a great deal of patience, you can follow each interference fringe all the way around. And you can work out its displacement. And you get out this rather nice map of the electron density.

And we see that there is still sufficient spatial resolution, despite the fact that we're using interference fringes, which limit our spatial resolution. There's still sufficient spatial resolution to resolve these very sharp shock features here. So this is a nice piece of work.

And then my final example for this batch here was actually showing something we've already discussed. This is where we had plasma flows from the left and the right colliding here. And as opposed to seeing interference fringes, we just see this dark void. And you can see the fringes are beginning to bend downwards here, which will indicate enhanced electron density.

But because the density gradients are too large, the probe beam has been refracted out of our collection optics, but we don't get to see anything. Now, you could say, well, perhaps it's because the plasma is too dense. If we got to the critical density, the laser beam going through the plasma would be reflected.

And so that could be the case. It's just that density is really, really high and very hard to reach, whereas we know that in this system, that density gradient is very easy to reach. And so we're pretty certain that in these experiments, it was the density gradient rather than the absolute density that caused us to lose our probing in the center here.

And there's not really much you can do about that. You can go to a shorter wavelength if you've got one, so that your beam doesn't get deflected so much. You can use a bigger lens so that you collect more light. But there's only so big a lens they'll sell you.

And so sometimes you just have to deal with the fact that your data has got holes in the middle. And when this image was processed in his paper, he just masked this region of the data off. And he said, we don't have any data there-- really the only thing you can do. Any questions on spatial heterodyne interferometry? Yes.

AUDIENCE: So I think as you alluded to, just to make sure that I understand, in order to extract useful information from one of these pictures, you need to be able to trace each fringe from edge to edge of the picture. And if you lose that, you're in trouble. So then how do you-- I mean, is this all done with computer image processing for your ability to, say, continually trace all of these fringes?

JACK HARE: So you don't have to be able to trace each fringe from side to side. But you do-- ideally, you'd be able to assign-- say you numbered each of the fringes from the bottom to the top of the image here. You'd like to be able to assign numbers to the fringes in the image without a plasma.

So you need a reference interferogram. So you always have to have two pictures, because that reference interferogram gives you the background signal that you're effectively modulating here. So in temporal heterodyne interferometry, you'll have to measure the beat frequency for some time before the plasma arrives.

The trouble comes that, actually, you don't need to be able to uniquely allocate each of these fringes to a fringe in the reference interferogram. You'd like to. If you can't do that, then there's some offset constant of density that you can't get rid of that's ambiguous. So you're like-- you can say that my density is going to change from here to here by 10 to 18, but it may also be like 10 to the 18 plus 10 to the 18 or 10 to the 17 plus 10 to the 18, something like that. So there's some ambiguity there.

If the fringes are broken-- so in this case, some of the fringes-- well, actually, this one is simpler to see. In this case, some of the fringes on this side here, you can't actually follow them through on the other side. But you can make some pretty good guesses in the absence of the plasma, because they'll be nice straight lines. And you can trace them across like that.

So for these complicated interferograms, the best process we've found is grad students, but lots of people say, I'm going to write an image processing algorithm. And indeed, they'll send students from the PSFC who were doing a machine learning course, who tried to do this.

The trouble is humans have incredible visual processing. So when you look at this, you can work out what all these lines are perfectly. Every single algorithm I've seen to try to do this automatically, it starts getting hung up on the little fuzziness on this line here. And it's like, oh, I think that's really important, so I'm going to spend all my time trying to fit that perfectly.

And so there may be techniques which can do it automatically. But at the end of the day, it really requires a human to look at this region where there's no fringes and say, ah, we've lost the fringes because I know that density is really high there.

Or in fact, rounding these regions here, it's hard to see, but there's actually some very strong shadowgraphy effects. There's some brightness in this region here. And that looks like additional interference fringes.

But if you've looked at these enough, you'll know that it's through shadowgraphy. So it seems to be very hard to train a computer to do it. So I'm not saying it's impossible. I just haven't seen a realistic program yet.

If you have an interferogram where all the fringe shifts are relatively small and well behaved, you can do this using Fourier transforms. So there are techniques which are Fourier-transform based, which basically do a wavelet transform. So it's like a small region Fourier transformation and looks at the local frequencies there.

And that's like those spectrograms that I showed you. That's the 2D equivalent of the spectrograms I showed you here, where we have the spectra each time. You want the k spectra at each position.

And those do an OK job. But as soon as there starts being any ambiguous feature, or even some relatively large distortions, they also fall over really badly. So it seems like a hard problem to automate. So that was a long answer to your question. I'm sorry. Other questions? Yeah.

AUDIENCE: So the image processing, it seems like choosing which areas to mask and which areas to not mask is important.

JACK HARE: Yes.

AUDIENCE: Is there any weird, like cut and dry rules for that, or is it all sort of related to intuition?

JACK HARE: At least the way that I do it, it seems to be very intuitive. Yes, exactly. You sort of have to know what you expect to see, and then work with that. Other questions? This side of the room is much more questioning than this side of the room. Next one.

AUDIENCE: From a practical standpoint, how many time points can you resolve with diagnostics such as these?

JACK HARE: Yeah. So if your plasma is-- the question was, what's the sort of temporal resolution of something like this? If your plasma is only lasting for a few hundred nanoseconds, then it depends whether you can afford a camera that can take more than one picture in a few hundred nanoseconds.

These were taken with off-the-shelf canon DSLR cameras bought in 2006. And there, the shutter was actually open for one second, but the laser pulse is only a nanosecond long. And that sets the time resolution.

So you can get one picture in an experiment as this. And then you do the experiment again, and you hope it's reproducible enough. And you move the laser later in time. And you take another picture. And you keep doing that. Yes, it's hard.

AUDIENCE: This will be a silly question. But what are those circular fringes that appear?

JACK HARE: Yeah. So this is what I was saying, where we've effectively violated our--

AUDIENCE: They're very light in the background.

JACK HARE: These are diffraction patterns of dust spots.

AUDIENCE: Dust spots.

JACK HARE: There's dust on an optic somewhere. It creates a diffraction pattern. It's out of focus. It modulates the intensity of the laser beam. It's another thing that makes it hard for automated algorithms to work. We tend to normalize those out, but I'm showing you-- this is the actual raw data from a camera. I haven't done anything to it.

But you can do tricks to get rid of those, because it's like a slow moving, slow changing effect. You can do like a low pass filter.

AUDIENCE: In these sorts of papers, how is uncertainty communicated with the result?

JACK HARE: Yeah. So we tend to estimate something like your uncertainty and density is going to be about a quarter of a fringe shift. And we'll talk about-- I'll talk about what that means actually in the next bit of the chalkboard talk. So you can estimate uncertainty by saying how certain are you that the fringe has shifted up this far versus this far.

So there's sort of like a pixel uncertainty, and also how good you are at assigning, this is the lightest part of the fringe, this is the darkest part, because, effectively, you're looking for the light parts and the dark parts, but there's several pixels which will be equally light, because it's near a maxima or a minima.

You tend to whack a relatively high uncertainty on it and just call it a day. In this field, if we get measurements right within about 20%, we're pretty happy. So this is very different from other parts of plasma physics.

Where there any questions online? I'm sorry. I can't see hands online at the moment, because I hid that little bar, and I don't know how to get it back. So if you put your hand up or something like that, I can't see it. I assume I'm still on Zoom somewhere. No idea how to get back there.

AUDIENCE: Escape, I think. [INAUDIBLE].

JACK HARE: Ah, OK. There was something in the chat. No questions here. All right, well, that was easy. Thank you.

[LAUGHTER]

I think we'll go back to the chalkboard for a moment then. I've got a few more pictures, depending on how we do for time.

I guess I should look down here for the remote control at some point.

So maybe just a little bit more practical stuff-- one thing that really matters is your choice of your probe wavelength. I've been talking a lot about frequency. It turns out that a lot of the time, people quote their frequencies in terms of wavelengths, and obviously, they're very intimately linked. So if you remember, we had our phase shift was minus omega over 2c [INAUDIBLE] like that.

We often define a quantity called a fringe shift, which I'm going to write as capital F. And a fringe shift is just a shift of an intensity maxima or minima by an amount in time or space that makes it look like another intensity maxima or minima. Having said that out loud, I realize it's pretty incomprehensible. So let me draw the picture.

Let's say that this is space or time. It doesn't really matter. And you've got some intensity here and some background fringe pattern like this. So this is the beat frequency that you're measuring, either in space or time in the absence of any plasma.

And then, say, that you have some plasma signal. I'm going to draw this wrong if I don't look at my notes. Give me a moment. So in the presence of a plasma, your fringe pattern has been distorted. Why have I got that one twice?

There we go. I did it. So this fringe here, you would think, should line up with this one. But in fact, it's been shifted all the way so it lines up with this one instead. So this is the case with plasma, and this is the case with no plasma. And that is the definition of one fringe shift. So it's effectively delta pi over 2 pi. We're just counting the motion of minima and maxima here.

And we can write that in practical units as 4.5 and 10 to the 16 lambda times the line integrated electron density. So people often write this quantity here like this, I think, because it looks nicer on one line in an equation, because you don't have the big integral sign. But effectively, it just means the average electron density average over some distance L like that. And all of these units here are SI.

And so that means you can then work out what your line-integrated electron density is in terms of fringes. It's 2.2 times 10 to the-- that was a minus 16 here, 10 to the plus 15 over lambda times the number of fringe shifts. And that's in units of per meter squared.

And so now I'm just going to give you for different lambda what this number actually is so we can have a look at some different sources. So I have a little table where I have wavelength of the source, and then I have any L or F equals 1. That's in units of meters squared.

I'm just going to go down a list of sources. So if, for example, we're in the microwave range here, this might be a wavelength of 3.3 millimeters. So that would be a 90 gigahertz source. So this is relatively low density plasmas here. And that density here will be 6.7 times 10 to the 17.

We could jump quite a bit and go to a CO2 laser. This is a nice infrared laser. And you can make very powerful CO2 lasers. So they're quite popular for some diagnostics. [INAUDIBLE] had a CO2 interferometer. So this is 10.6 micrometers here. So I've dropped a couple of orders of magnitude from the microwaves.

And you can see that the densities that we're measuring here have gone up quite a bit by similar amounts. Ah, that should be 20. And then something like a neodymium YAG laser, this is the sort of thing I use. If we use the second harmonic, that would be 532 nanometers. That will make those beautiful green images that we looked at. And that is 3.2 times 10 to the 21.

So if I see a fringe shift of one, and on those images that I showed you first, the fringe shift was maybe two or three fringes, each of those corresponds to an electron density of 4.2 times 10 to the 21. So just by looking at the image and by eyeballing it, you can start estimating the line-integrated electron density. And then if you have some idea of how long your plasma is, you can then get a rough estimate of the electron density itself.

I think what's very interesting about this, which I worked this out just before the lecture, and I hope I'm right, because I was very surprised by it. If we take, for example, a wavelength of 1,064 nanometers. So this is an infrared Nd:YAG laser. When we take a length of 10 to the minus 2 meters, so a centimeter-- so this is the sort of experiment that I might do-- one fringe would then correspond to a density of 2 times 10 to the 23 per meter squared. Not that interesting so far-- meter cubed.

But then you ask yourself, what is the critical density here? This turns out to be 10 to the 28 per meter cubed. So what is the refractive index? We've been saying that it's 1 minus Ne over 2Nc. This is a very small number. This is now-- well, it's almost close. It's 1 minus 10 to the minus 5. So all of these interferometry effects we're looking at are to do with changes in the refractive index on the order of 10 to the minus 5 or so.

I kind of find that remarkable. We're able to measure very, very, very small changes in the refractive index. It's not like N ever gets close to 0, or 2, or something bizarre like that. Anyway, I thought that was interesting.

So you will pick your source to match your plasma. If you're doing low density plasmas, then if you use a 5.2-nanometer interferometer, you won't see any fringe shift. The fringes won't move at all. You won't be able to measure any plasma. So you need to use a long wavelength source that is more sensitive to those lower densities.

And conversely, if you try to use a long wavelength source on a nice dense plasma, first of all, the beam may just get absorbed, because it'll hit the critical density, or it might get refracted out. And even if it doesn't do any of those, you'll have such huge phase shifts you won't be able to meet that heterodyne criterion. And you'll just have very complicated fringe patterns and no chance of processing.

So you've got to pick very carefully what sort of source you have. And there are other ones out here as well. But I just sort of picked a range that might be relevant to some of the people in this room. Any questions on that before we move on?

So I want to talk about a few extensions to this technique. And the first one we're going to talk about is called two-color interferometry. There are two reasons to do two-color interferometry, conveniently. One is to handle vibrations, and the other one is to handle neutrals.

So let's have a little talk about vibrations, first of all. Your system is made up of lots of mirrors and other optics. And there are vibrations everywhere. And so all of these mirrors and optics will be vibrating slightly, which means their path length will be changing by a small amount.

How big a deal is this? Well, if we imagine that you've got some mirror, for example, here, and we bounce our beam off it like that, and the mirror is oscillating with an amplitude little l here, we're going to get a phase change, very simply, just by looking at the distance that this moves on the order of 2 pi little l upon lambda.

So if the phase-- if the amplitude of these vibrations is on the order of the wavelength, you're going to get a phase shift of 2 pi, which is actually already pretty huge. That's one fringe shift. But this is a tiny number. I mean, if I'm working with green light, this means I'm sensitive to vibrations on the order of 532 nanometers.

So this is extremely hard to avoid. You can't get rid of vibrations after all very easily. So you're going to have big problems with these vibrations. And if your whole tokamak is vibrating, so you've got all these cryopumps and neutral beams and exciting things going on, this is going to be an absolute nightmare.

It turns out not to be a huge nightmare for my stuff because although these are very sensitive to vibrations, the timescale over which our experiment takes place, the nanoseconds of a laser pulse, the vibrations are not in gigahertz. You don't have mechanical vibrations at those frequencies. So we can ignore it. But if you have vibrations at kilohertz or even hertz from people walking around, this will ruin your nice steady state experiment.

And what we want to notice is that this phase shift in particular, as I've just sort of alluded to, is large or small wavelengths. And this already suggests the beginnings of a scheme that we're going to use, especially because I called it two-color interferometry to deal with this.

So the solution is we run two interferometers down the same line of sight with two different wavelengths. One of these wavelengths is short. And so that could be something like on a tokamak a laser. So it could just be a HeNe laser beam. It goes straight through the plasma. It doesn't see it at all because it's not dense enough.

But that short wavelength is very sensitive to vibrations. The vibrations are proportional [INAUDIBLE], should we call it, a vibration [INAUDIBLE] over lambda. So this will be a very good diagnostic, not of plasma, because it won't see the plasma, but it will be a very good diagnostic of vibrations.

And the other one, the long wavelength here, which might be on a tokamak something like a microwave source, so many, many, many differences in wavelength by two or three orders of magnitude. That long wavelength will be very sensitive to the plasma, because my plasma is proportional to lambda.

And so what you can do, when you measure the overall phi with these two devices-- well, I'm not going to do it explicitly-- you're going to have two sources of phase. One of them is going to be the standard plasma here, NeL. And the other one is going to be the vibrations due to L.

So you'll have two unknowns. And now you've got two measurements. And so you can use a system of equations to solve for that. And if you're very, very quick, you can even use the short wavelength to feed back into your mirrors with the PA side and stabilize the mirrors.

So you can do vibration stabilization feedback mirrors. So you can use this very fast, short wavelength interferometer to vibration stabilize all of your mirrors. I don't know if anyone's actually done this. It's in Hutchinson's book, so presumably someone tried it. It sounds like a lot of work. But I guess it could be very, very effective.

So this is maybe more of a question mark, rather than something that everyone does. It's pretty clear that if you've digitized both these signals, you should be able to work out what was the vibration and what was the plasma. Of course, if the vibrations are huge, it might still ruin your measurements. So it might be worth doing this feedback system.

So that's one use for two-color interferometry, just in case you were trying to work it out. The two colors is because we have two wavelengths, and we tend to associate wavelength with color. Any questions on that?

The second thing I want to deal with is neutrals. So far our refractive index has been derived assuming a fully ionized plasma. And so in that fully ionized plasma, we just have ions. And we have electrons.

Now, these have associated frequencies with them, the ion frequency, which is much, much less than the electron frequency. And so when we write down the refractive index, we just have 1-- this is refractive index squared-- 1 minus omega squared over omega [INAUDIBLE] squared. I need the chalk there.

There's technically another term in here, 1 minus omega squared omega Pi squared. But because the arms are so much more massive, we always just ignore this term. And we subtly drop the subscript on the plasma frequency here. So this is the refractive index we've been using so far.

However, when you've also got neutrals, you've got some density of these as well. And I'm going to write that N alpha, because-- no, not alpha, Na. The a is sort of being atoms. So let's go for that.

Now, neutrals are much more complicated, in fact, because they have atomic transitions inside them. And those atomic transitions change the refractive index. If you're closer to an atomic transition, you have very different effects, like absorption, than you do if you're far from an atomic transition. And so your spectra or your plot of refractive index for your neutrals here against frequency will look like some sort of spiky minefield of lines, something like that. It will depend exactly on the atomic physics.

And in general, we can write down this refractive index as equal to 1-- always a good start-- plus 2 pi e squared upon me. Ignore that. It's just some constants that they're normalizing by. And then we sum over every single atomic transition in this neutral gas.

So first of all, we want to sum over all of the atoms in state i here. So for example, there may be atoms which are partially ionized. They've lost one electron. And so then they have a different refractive index here. So all of the atoms in a certain state i.

And then all of the transitions between that state i and some other state k as this is divided by the transition frequency, which is one of these lines here between i and k minus omega squared here. So in this case, Fik is the strength of one of these transitions, which determines how likely it is to happen, and so how strongly it shows up. And this is the frequency of one of these transitions.

I misspoke earlier. I spoke about this being through ionization. It's just that it was excitation. So it's to do with whether your atom is in its ground state, or some other state, or something else like that.

Now, this formula is intensely complicated. And you can spend a very long time doing quantum mechanics calculations to try and work out both of these two terms. And of course, as soon as you go to something above hydrogen, it becomes very complicated. Even for hydrogen, it's pretty complicated. But above hydrogen, it's extremely complicated, because there are multiple electrons interacting here. So you don't actually stand a chance of solving this directly.

Your best bet is the fact that if you look at some parts of the spectrum or some part of the refractive index that's away from one of these lines. So this is for omega not equal to any of these transitions here. There's a general formula that works pretty well, which says the refractive index is just equal to 1 plus 2 pi alpha Na, where this Na here is still the number density of neutrals. And this new alpha is a quantity called the polarizability.

[INAUDIBLE], which is easy to calculate, and also very easy to measure. So you can measure this for your different gases. And so for example, this little table of different gases, if we have it for helium, or hydrogen, or argon here, here's the 2 pi alpha in units of meters cubed. The 2 pi is just some normalization constant that comes from somewhere else in the theory. So we always quote it with the alpha. I'm just going to quote you 2 pi alpha here.

But this number is like 10 to the minus 30 for helium 5 times 10 to the minus 30 for hydrogen, and 1 times 10 to minus 29. I wrote this as argon, because I missed the i in my notes. It's actually air. So there we go.

And then just as a little calculation here, air at standard temperature and pressure, the number density is around about 2.5 times 10 to the 25 per meter cubed. And so therefore, the refractive index of air is about 1 plus 2.5 times 10 to the minus 4.

So what began to change in the refractive index is very small compared to 1 for the neutrals. It's on the same sort of order as the change of refractive index you get for a similar sort of plasma. But crucially, the refractive index is always greater than 1 for neutrals. And it is always less than 1 for a plasma. This is our first hint at how we're going to use two-color interferometry here.

So let's have a look at how to use two-color interferometry to determine both the number density of the neutrals and the number density of the electrons. And you might come across this scenario quite a lot. If you're doing low temperature plasmas, you always come across this scenario.

Even if you're doing something in a tokamak, maybe there's a region at the edge where there's a large number of neutrals, and your beam has to go through that region at the edge. And you want to ignore it. You just want to measure the core. But you're still actually picking up a big phase shift from these neutrals at the edge. So this is a big problem. And this is a hard one to solve. So let's have a look at this.

So remember that when we derived the phase shift, delta phi, I did it, first of all, just in terms of a generic refractive index before specifying it to be a plasma. And that refractive index was just n minus 1 dl. There's 2 pi lambda at the front here. So this is just affecting the change in the path length, effective path length that the probing beam sees.

And we also said that NA is greater than 1 and dNa d lambda is equal to 0, which is mathematically saying that in this polarizability model here, we assume that it doesn't depend on wavelength, as long as we don't go too close to one of these transitions. So the polarizability here is the same as the polarizability here is the same as the polarizability here.

So then we'll end up with a total fringe shift on our interferometer of minus 4.5 times 10 to minus 16 times lambda times the integral of the electron density. That's the plasma component here. And we will also have a term, which is plus-- notice the difference in the sign here-- plus 2 pi alpha upon lambda integral of Na here.

So they will cause shift and fringe in different directions. So to a lower effective spatial or temporal frequency, but they also have a different dependence on lambda. And this is key, because, again, as we saw with the vibrations, we have different dependencies on lambda. We can use a two-color technique to get around this.

So phi plasma is big for large wavelengths. And phi neutral, phase shifted for neutrals, is big for small wavelengths. And again, we've got two unknowns. We've got the density of the electrons. Oh, sorry. We don't see the neutrals.

And if we have a two-color technique, we have two equations. And so we can solve all of that. And I'm not going to write down the algebra now. It's quite boring. And you can work out uniquely, apparently, what the electron and neutral densities are. And I'll show you that this doesn't generally work in practice.

So in practice, you often end up with negative predictions for your density of both the neutrals and the electrons. And this tends to be, as far as I can tell from reading the literature, that when we did this approximation of n is about 1 plus 2 pi alpha here, we have assumed that alpha is constant.

But it doesn't have to be constant. It could change with wavelength. And if we're using two different wavelengths, and there are two different values of alpha, then that will cause chaos with your two equations, two unknowns, because alpha is actually relatively hard to pin down in some sources. You can very rarely find it for the exact wavelength you're working on.

You may also have horrifically ended up using one of your wavelengths, like here, halfway up one of these resonances, or even worse, at the peak of one of the resonances. And if you did that, your whole model is completely off.

And I think this is what causes this. And I'll show you some data I took where we predicted negative densities. And I'll talk a little bit about that as well. Obviously, negative densities are unphysical. So we thought that was probably wrong, but we published it anyway, because other people were doing the technique and not pointing out they had negative numbers. And we thought it'd be nice to point out that we knew that it was wrong.

So that was a quick roundup of two-color interferometry. And I have some quick slides after this showing some examples. But I'll just pause here and see if there are any questions. Yes.

AUDIENCE: Does this technique work better for certain levels of [INAUDIBLE]? Is there a spectrum where it works better than others?

JACK HARE: The question was, does this-- sorry. Go on.

AUDIENCE: If you're using a really weakly ionized plasma, it's really not a good option as opposed to a more [INAUDIBLE].

JACK HARE: Right. Really good question. So the question was, does this work better for different levels of ionization in a plasma? So you might be thinking to yourself, if I have a very, very weakly ionized plasma, then it may be very hard to measure the electrons over the overwhelming change in refractive index from the neutrals. And that's true. It's going to be hard.

But if you look at this equation here, this is the thing you're measuring, the fringe shift. You can choose your two wavelengths to optimize the wavelength sensitivity of one of them to the electrons and the sensitivity of the other one to the neutrals. So you're going to need to have some widely spaced wavelengths.

So if you try and do it with just two different frequencies from the same laser, that's going to be really hard. But if you have a microwave interferometer and a HeNe green laser beam, like they do on the tokamaks for vibration stabilization, that will work much better.

The difficulty there is then you have two completely different detection techniques. And so it's not easy to compare these two. But that's what you probably want to do if you're dealing with 1% ionization or something like that. You might have to do this technique. Were there other questions? Anything online?

AUDIENCE: What physically is the polarizability? Is that the electric polarizability of the medium?

JACK HARE: Yeah. The question was, what physically is the polarizability? This polarizability is very strongly related to how the electron wave functions are distorted by the electric fields of the electromagnetic wave.

AUDIENCE: OK.

JACK HARE: Yes. Which is why when you get close to a transition and the frequency of the wave is now resonant with some atomic transition, this polarizability changes dramatically. As opposed to just going through the medium, the electromagnetic wave is absorbed. I'm saying it in a very classical way, but of course, you need to start doing quantum if you want to have absorption.

AUDIENCE: It's like the wave field is inducing a small dipole moment.

JACK HARE: Yeah, absolutely. The wave is inducing a dipole moment, and that is slowing down the wave. Just slowing down the phase of the wave. In a plasma, remember, it always speeds up the phase of the wave. Any other questions?

Maybe I should have saved all my pictures till the end and avoided having to find this thing twice. Can you see this online?

AUDIENCE: Yes.

JACK HARE: OK, perfect. And it's showing up slowly here.

So this was a set of experiments we did with a very sexily named, but boring device called a plasma gun, which is actually just a bit of coax cable where you've chopped off the end of it. And you pulse it with some current. And the current flows up the inner conductor across the chopped off plastic insulator and back down through the outer conductor.

And as it flows across here, it sends plumes of plasma outwards. And they're moving 10 kilometers a second, but this only works in a vacuum. So it's not a very good gun.

Anyway, this was a fun object to study, because there was a grad student using it for his PhD thesis. And we put it on our experiments. And we did two-color interferometry.

And these interferograms were made using an Nd:YAG laser, neodymium YAG. We used the second harmonic at 532 nanometers, which shows up as green here, to do one of the measurements. And simultaneously, along the same line of sight through the same bit of plasma, we used the third harmonic, which is 355 nanometers. So that's in the ultraviolet. So you can't see it by eye.

You might be asking, why does it show up as orange? And this is when you remove the ultraviolet filter from your off-the-shelf Canon DSLR camera. The pixels get confused and think that this is orange. Obviously, it can't-- it can see ultraviolet, but what we can render it as on the screen. Anyway, this is 355.

And these are the fringes before the plasma was there. And these are the fringes after the plasma was there. And you can see the fringe shift is very small. There's just tiny little shifts here and here. So this wasn't a very high plasma density.

But what we were able to do is to infer the phase shift for these two interferograms like this. And then we combined these two together using the simultaneous equations. We also did a novel inversion, which I'll talk about in the next lecture. And we got out the electron density.

And this looks quite reasonable. We get up to about 10 to the 18 here. And it falls off nicely in various directions. But we also got a prediction of the neutral density here. And these red regions are fine. These are positive numbers.

But then right in the middle here, there's a big negative number. And in fact, it's so negative we're predicting many more absences of neutrals than we had electrons. So it's like clearly completely nonsense.

And so we went back to some of the textbooks that explained this technique for measuring neutrals, and we found that the example data they were showing also had negative numbers in. It's just they didn't bother to mention that this was a huge problem. We think this is a huge problem.

So I'm a little bit baffled by the fact that people will in a textbook say that this technique can be used to measure neutrals. But in reality, it seems to be really, really tricky to do it properly. And I think the problem is the quality of the polarizability data that we have.

So we were trying to use the polarizability data, assuming that it worked at 532 and at 355 nanometers. But it was derived in the lab by some group in the '80s who published a paper on it. And they did it like at 10.6 microns in the infrared.

So there's no really good reason to believe the polarizability is the same. But it's really hard to get hold of this data in a consistent fashion. So if you're going to try and use this technique, I think you should be very skeptical about the results, especially if you start seeing negative numbers.

So we've only got a couple more minutes. I'll just take any questions on this, and then we'll do [INAUDIBLE] inversion in the next lecture. So any questions? Yes.

AUDIENCE: How do you go about measuring [INAUDIBLE] parameters to get rid of any disruption?

JACK HARE: You could use interferometry to measure the polarizability in the absence of any electrons. Then you'd know exactly what you were measuring. So if you were able to puff some gas in-- you need some measurement of the number density as well.

So that's-- you can imagine if you've got a gas cell at a certain pressure and it's at room temperature, then from the ideal gas law, you know the number density inside that gas cell. And you know the size of the volume.

And you could do interferometry on that volume, for example, and that would give you a measurement of the phase shift. And then you could back out what the polarizability must be. Maybe we should have done that here, but we didn't.

[LAUGHTER]

Yeah.

AUDIENCE: That's just like an additional calibration step?

JACK HARE: Yes. Yes, exactly. So I think it's a doable calibration. It's just quite hard. Whereas for the electron density, it's like that is all in terms of fundamental parameters, like the electron charge, and the electron mass. And you're like, OK, we know what those are.

So then when you're applying those formulas, you have absolute confidence when you measure the phase change what the electron density is. It's just for this, the theory is a little bit more wobbly. Other questions?

AUDIENCE: If you have good data on individual species, then you know you have a certain ratio in your plasma. Is there any reason to think that you couldn't just do an average?

JACK HARE: If you think beforehand you somehow know for some reason the number of-- the electron density and the neutral density, or--

AUDIENCE: Oh, sorry, like two different neutral species. And you have good--

JACK HARE: Oh, we didn't even get into that. That's a nightmare, because then you'll have two different polarizabilities, and even more transitions that you're trying to miss. So we're already like, oh god, let's stay well away from any of these transitions. But if you have two species, you'll have an even harder time finding a region with no transition. So it'll be very hard to find on its source.

Any questions online? All right, thank you very much, everyone. See you on Thursday.