

[SQUEAKING]

[RUSTLING]

[CLICKING]

JACK HARE: Right, so today, we are going to start a series of several lectures on refractive index diagnostics. So these are diagnostics which use the fact that the refractive index of a plasma is not 1 in order to make measurements about the plasma. So first of all, what we're going to do is a very quick recap of electromagnetic waves in a plasma.

And you have doubtless seen this before in some courses that you've done, like it's in Chen. There's also a very good explanation of it in Hutchinson. And I'm going to give a little taster of how this derivation goes, just to remind folks of how this result actually comes about. And I'm going to do it in a rather restrictive way here so that we can make some rapid progress. And we'll go back and add in more bits of theory later on as we need it.

So for our electromagnetic waves in a plasma, we have two fields. We have E and B, like this. And we have Maxwell's equations. So we have that the curl of E is equal to minus B dot, and we have C squared times the curl of B equals E dot plus J over epsilon 0. You might be more used to seeing this equation with a mu 0. I've just moved the C squared over the other side because it makes the math a little bit easier on the next step here.

And so these equations are just true. They're true in any medium. And so what we want to do is try and reduce them down a little bit and then put some plasma physics in. So the first thing we normally do when we're deriving electromagnetic waves is we say, what if all of our vectors had some sort of time and space variation that looked like exponential of i, some wave vector k dot, some position vector x minus omega the frequency times time, like that.

And so this is a little bit like Fourier transforming our equations here. And we end up with k cross E equals i omega B, and we have C squared k cross B equals minus i omega E plus J upon epsilon 0. And the astute amongst you have noticed I made a sign error in the first equation here. So this is plus i omega B, not minus i omega B. Very good.

And so we could call these equations 1 and 2. And we can note that if we do k cross with equation 1, that is equal to i omega upon C squared of equation 2. So what we're doing here is we're just looking at this term and this term and being like, hey, they both got B in them. We can probably make them look the same.

So then we can equate. We can do this calculation and we can equate the two sides of the equation. And with a little bit of vector magic, calculus magic, we would end up with something that looks like k k dot E plus k squared E is equal to i omega J over epsilon 0 C squared plus omega squared upon C squared E, like that.

Now, we are searching for transverse waves here, electromagnetic waves that tend to be transverse, and that means that k dot E is 0. So we're just going to look for transverse wave solutions here. We can drop that. And that means that our equation can be rewritten as omega squared minus C squared k squared times E equals minus i omega J upon epsilon 0, like that.

And I just want to point out, there's absolutely no plasma physics in this at the moment. All we've done is manipulate Maxwell's equations. We haven't said anything about the plasma.

So if I take a standard limit here, say I let the current equal 0, which is what it would be in the vacuum of space where there's no particles to carry any current, then we would simply end up with an equation for light in free space and you would have a dispersion relationship $\omega^2 = c^2 k^2$, like that. So those are light waves. So life is good.

Any questions on that before we put some plasma physics in? Hopefully, you're dredging up your memories of Griffiths and Jackson and all sorts of wonderful things like this, and this is all making sense. So let's keep going.

The next thing we want to do then is add in some plasma physics. And we're going to add in some plasma physics with some serious assumptions which let us make significant progress quickly. And these assumptions can all be justified for most plasmas. And if you can't justify it for your plasma, you may want to revisit these a little bit.

So the first assumption we're going to make is that we're using high-frequency waves here. And high frequency, this is kind of a wishy-washy term. We want to make that more precise by having a dimensionless parameter.

And so we're going to say that ω is much, much larger than ω_{pi} here. So the frequency of our waves is much higher than the ion plasma frequency. What does that condition physically respond to? What are we saying about the ions and their interaction with this wave by putting this condition in?

STUDENT: You're saying they essentially don't interact at that point.

JACK HARE: Yeah. So the ions are frozen in place. And they are not going to participate in any of the physics that we're interested in here. And again, you can work out the ion plasma frequency for some density that you're interested in, and you'll find out it's pretty low. So this is pretty reasonable, but if you start using very low-frequency waves, it won't be reasonable.

So another thing that we're going to do is make the cold plasma approximation. And again, we can't just use the word cold. We have to say what we mean by cold. And we're going to say the thermal of the electrons is much, much less than the speed of light here. And this condition is equivalent to us not worrying about the Maxwellian distribution of the electrons.

So all the electrons are just going to be moving with no-- they will be moving, unlike the ions, which are frozen, but they will be moving all at the same speeds. There's no spread of velocities here. So we have a delta function of velocities, and we do not have our Maxwellian distribution that we might normally think about.

And the final condition we're going to write down is unmagnetized. Now, this is the one which I think is most complicated. Because, in fact, there are a dozen different ways you can write down a dimensionless parameter for unmagnetized, and they all mean slightly different things. We could think about collisionality. We could think about pressure balance, all sorts of things like that.

For the purposes of this derivation, unmagnetized means ω is much, much larger than ω_{ce} . I know it's confusing. And I'll put a little subscript e here because we've frozen the ions. But this is the gyro motion of the electrons. So the electrons may be gyrating around field lines.

There may be some magnetic field. It doesn't have to be 0. But on the time scale that the wave goes by and does its stuff, the electrons do not move appreciably around their gyro orbit. So we don't have to care about their gyro motion.

This is one of the places we will definitely have to relax later on when we want to do Faraday rotation imaging which relies on that gyro motion to give us the effect. So this is the final condition. Any questions on those three assumptions?

OK, so then, we can write down that the current inside our plasma is simply going to be equal to the charge on the electrons, the number of electrons, and how fast they're moving. And so we've simply transferred our lack of knowledge about J into our lack of knowledge about V . So we better do something about that. And we're going to do that using the electron equation of motion.

And that electron equation of motion looks like $m \frac{dV}{dt}$ is equal to minus e the charge times capital E , the vector electric field, like that. So we could rearrange that so that we have V is equal to e vector E over $im\omega$. We've done the same Fourier transform trick that we did with all the other quantities. We assume it's going to be oscillating in some way.

Substitute that back in to our equation for E , and we'll get ω^2 minus $C^2 k^2$ E , as we had before, is equal to nee^2 over ϵ_0 and E , capital E . Sorry, was there a question? I heard someone speak.

STUDENT: No. That was an accident. Sorry.

JACK HARE: No worries. All good. And just to be clear here, this equation of motion doesn't have the $V \times B$ term in it, which would be like the magnetic field term, because we've dropped it because we're making this unmagnetized approximation. But that's where you put it back in. You put in a little term that was like, $V \times B$ in here and put little brackets around. But for now, we're just setting that equal to 0. OK. Good.

And the solution to this equation that we've got now is ω^2 equals ω_p^2 plus $C^2 k^2$. So this looks an awful lot like what we had before, which was just these two terms. But now, we've got this additional term, which includes some plasma physics. And I could write that this is ω_{pe}^2 and make it clear that it's the electrons, but remember, we've dropped the arm motion already. So there's only one plasma frequency that's very interesting, the electron plasma frequency for this derivation here.

And then, we can go ahead and do all the standard things we do with one of these functions, which is to try and write down the phase velocity. And the phase velocity squared is just equal to ω^2 upon k^2 . And so that means that our phase velocity is C upon 1 minus ω_p^2 upon ω^2 . It's a half.

And then, we can write down the refractive index, because our refractive index, capital N , is just equal to C over the phase velocity. And that is going to be equal to 1 minus ω_p^2 over ω^2 . And we often write this not in terms of frequencies, but in terms of densities, because the plasma frequency has inside it an electron density. We write 1 minus n_e over $n_{critical}$, like this, where $n_{critical}$ is some critical density.

And we'll get on to why it's so critical in a moment. But if you want to approximate it, then n critical in centimeters to the minus 3, so particles per cubic centimeter, is roughly 10 to the 21 over λ when λ is in microns here. So if you're using a laser beam at 1 micron, that would be a very standard laser wavelength from a neodymium YAG or a neodymium glass laser, then you have a critical density of about 10 to the 21 .

Which depending on which field you're working in is either hilariously high and unreachable or crazy low and happens all the time. So again, this is one of the exciting things about doing plasma diagnostics course where we span 16 orders of magnitude in density. Any questions on that?

STUDENT: Yeah, what's the critical density point again?

JACK HARE: As in, what is its physical significance?

STUDENT: Yeah. Like, why did you choose that number?

JACK HARE: We're going to look into that in a moment. It's a good question. It's a solid question. But just to be clear, the reason I've got there is I've taken this ω_p^2 and this ω^2 and I've noticed that ω_p^2 has inside it the electron density, and I've rewritten all the other terms. So this n critical now has inside it things like ω^2 .

I made a critical mistake in my equation up here. This is λ^2 here for approximating this in terms of simple quantities. So the n critical now depends on the laser wavelength or the frequency of your probing electromagnetic radiation. So it's different for every frequency.

But that's the only thing it depends on. The rest of it is all like, fundamental constants, like E and the electron mass that don't change. So this is a parameter there's a critical density for every single laser wavelength or electromagnetic wave frequency.

I'm going to talk about lasers a lot. Of course, this also applies to microwaves and things like that as well. But some sort of source of radiation. So I'm going to get on to what n critical is in just a moment. But any other questions on this before we keep going?

OK, so we just said that n is equal to the square root of $1 - n_e / n_{critical}$. If we work in a regime where n_e is much, much less than $n_{critical}$, so we work far from the critical density, we can do a Taylor expansion, which is what we often end up doing, and we write $1 - n_e / 2 n_{critical}$, like this. So my questions for you now-- and we will try and work out what $n_{critical}$ is doing together-- for n , for a density which is greater than the critical density, what happens to n , the refractive index?

STUDENT: The wave becomes evanescent, right?

JACK HARE: Yeah. Well, that's the result. So what happens to n itself?

STUDENT: Big N ? It's imaginary.

JACK HARE: Right. Exactly. Yeah. Yeah. You've got the right answer. I just wanted to take in a few more steps. So n is imaginary because it's going to be the square root of a negative number, which we can see here.

And so that means our wave becomes evanescent. So it's going to have properties which decay in time. So it's going to look like $e^{-\alpha x} e^{-\gamma t}$, like this. So it dies off.

So that means that we can't propagate a wave at densities greater than the critical density. What happens to all the energy, then? Because this wave is carrying energy.

STUDENT: It's absorbed by the plasma.

JACK HARE: No absorption mechanism in our equations, actually.

STUDENT: Yeah, so that-- I mean, from the math we have, the only option's reflection, right?

JACK HARE: Yeah. So reflection is indeed the answer. So we will get reflection of this wave. It will bounce off the critical surface and go somewhere else.

We'd only be able to have absorption if we put, for example, collisions into the equation of motion, all the way back there. And then, we could have what's called inverse Bremsstrahlung, which is, effectively, the electrons get oscillated by the wave, and then they collide with some ions and transfer the energy to the ions. So that's a damping mechanism.

There could be other damping mechanisms like λ , but in the equations we've got so far, we don't actually have those. So reflection is the only thing that can happen. Now, we also had this equation for the phase velocity, which was that v_p is equal to c over 1 minus ωp squared upon ω squared to the $1/2$. Does anyone want to comment on what happens to the phase velocity as we go above the critical density?

STUDENT: Does it go to infinity?

JACK HARE: I mean, we'll do that eventually, yes. Yes. Right there, at that point--

STUDENT: It diverges?

JACK HARE: Extremely large. Yes. Yeah. So this is going to start doing very silly things here. And those things don't seem very physical, of course, because we don't want things traveling faster than the speed of light.

Fortunately, this isn't a big problem because the phase velocity doesn't carry any information, and there are limitations on things going faster than the speed of light have to do with information. And the information is encoded in a quantity called the group velocity, and the group velocity just looks like this, c times the square root of 1 minus ωp squared upon ω squared. And so as we get close to the critical density here, all that happens is the group velocity goes to 0 , and effectively, we transmit no information through that evanescent region. So we don't have to worry about the fact that phase velocity. It's superluminal because our group velocity is still subluminal, as we'd like. Sean?

STUDENT: So if the wave's being reflected here, and the group velocity is going to 0 , to me, that seems like the wave information is sort of stagnating at the reflection point. How do we see from these equations that the wave information is actually being reflected back out.

JACK HARE: Very good question.

STUDENT: I think-- excuse me. I think you need to impose boundary conditions to do that. Because there would be some sort of discontinuity, right?

JACK HARE: Yeah. I suspect the model I'm presenting is a little bit too simplistic to handle this stuff.

STUDENT: So we would need to have some more information.

JACK HARE: I think so. There's certainly-- as you get very, very slow group velocities, you're going to start-- we've been making some assumptions about the homogeneity here, and so, reflectively, there's going to be some length scale in here, which is going to be like k , the size of the wave vector of the light at a given point, and we're going to be comparing that to the length scale associated with how quickly the electron density changes, this gradient length scale here.

And as the group velocity gets very, very low that k is going to get very, very long, and we're going to start violating our assumption that k is-- let's see. Maybe I can write this as λ . That λ is going to be much, much less than this change in the gradient.

So for any realistic system, our density has to ramp up. It can't just immediately get up to the critical density. This could be the critical density here. And we're going to find out that there's some region where some of the approximations we've implicitly been making break down. And then, you need to start doing wkb, and all that sort of stuff and doing everything properly.

So I think, effectively, this simple model breaks down. But if you do it properly-- and I think Hutchinson does this in the reflectometry section. So we may end up doing it. You can get the answer about the reflection there. It's a good question. This is very hand-wavy at this point. I agree.

STUDENT: Thanks.

JACK HARE: Cool. Any other questions on this? Because this equation here, we are now going to use an awful lot. So I'd like you to agree that it's valid within the assumptions that we've made. And if you don't agree, we should have a chat about it. OK, good. So let's keep going.

So there is a series of different measurements that we can make. And these are the refractive index diagnostics. So I'm going to just call these n measurements or n diagnostics. Because they rely on the change in refractive index here.

So one type is when the refractive index is not equal to 1. There's a refractive index inside our plasma is not equal to 1, which is true anytime there's any density inside the plasma. This sort of diagnostic causes a phase shift. So the plasma ends up-- the laser beam going through or the electromagnetic radiation going through the plasma ends up with a different phase than it would have done in the absence of the plasma. And we can measure that phase shift using a technique called interferometry.

And with interferometry, we can therefore say something about the density inside the plasma. Another technique is when the gradient of the refractive index is not equal to 0. So this is when there is any change in refractive index. And in a plasma, that corresponds very clearly to just changes in the electron density, so gradients in the electron density.

But of course, in general, this technique can be used for any medium where the refractive index changes. So air, if you heat it up, the refractive index changes, and so you could use these techniques. These are not specific to plasma physics.

And these diagnostics tend to be called refraction diagnostics, because the light refracts and it bends. And we end up doing techniques such as schlieren and sonography. And then, the final type that I'm going to talk about are ones where, actually, the polarization, the medium is birefringent. It treats different polarizations differently.

And so we can have polarizations of light which are circular. We can have the left-handed and the right-handed polarizations, which we sometimes refer to as plus and minus. And here, we would say that the refractive index for the plus wave is not equal to the refractive index to the minus wave here.

And so here, we measure the polarization. And this is using a technique called Faraday or Faraday rotation. Which we briefly discussed in the context of magnetic field measurements using Verdet glass. And in fact, it turns out that you need to have magnetic fields that are non-zero, and we also need to relax our assumption that the plasma is unmagnetized in the sense that the light frequency is much larger than the electron cyclotron frequency.

So those are three different types of refractive index diagnostic, and we're going to start with what I think is conceptually the simplest, but still often causes us lots of problems, which are the refraction diagnostics here. I see some people writing, so I'm just going to pause on this slide for a moment. Okey doke. Now, I just want to have a little aside. And this is on conceptual models for electromagnetic propagation.

Because I'm going to be switching quite a lot between different ways of thinking about electromagnetic radiation and how it moves through a plasma. Because sometimes, some models are easier to work with than others. Sometimes, models are simplifications and they throw away physics, but they make the intuition much simpler. So I just want to show you two different models that we're going to be using in these next few lectures so that you have an idea of what's going on.

One model would be a model of wavefronts. So this is based on the idea that, as we said, our electric and magnetic fields can be written as just a single Fourier component. So there's some strength and polarization of the electric field here, and this is multiplied by the exponential of what we call the phase factor. So $i \mathbf{k} \cdot \mathbf{x} - \omega t$, like this, which we could write as $E_0 \exp(i \text{ times some scalar quantity, which is the phase here.}$

And so if we think of our electromagnetic wave as having a phase, and the electromagnetic wave still exists in all places, all points in time, blah, blah, blah, blah, blah, but that's a very difficult thing for me to sketch on my iPad here or on the board in front of you. So what I'll probably end up sketching are what we call isophase contours. So these are contours along which the phase is constant.

And so, for example, it could be at some integer multiple of π , right? Yes. And belonging to z . So this might look like some waves like this. This would be an electromagnetic wave which is diverging here. Actually, it could be an electromagnetic wave which is converging to the left. But at the moment, it looks like it's diverging to the right.

So that's one way I could draw a wave here. Another way I could do it is with a ray model. And this gets into a topic which is called geometric optics. And it turns out what you can do, if you have some isophase contours like the ones I just drew, say, these contours here, they're doing something slightly strange, but perhaps there's a plasma there, which is like moving the phase contours around.

If there's a change of refractive index, will affect the phase. The rays that we draw here, I can just take these phase contours and I can draw rays such that they are everywhere normal to the isophase contours. So this ray would look like this. This one would look like this. And this one would look like that. So they are perpendicular to the wavefronts.

Conveniently, they are also parallel to the Poynting vector. At least, I'm pretty convinced they are. If someone knows more about geometric optics and thinks I'm wrong, please, shout out, because this was a very hard fact to check in like, 10 minutes before the lecture. But I'm pretty certain they represent the direction of the energy flux in electromagnetic waves. So they're quite conceptually useful as well. They tell us where the power is flowing.

Now, when we think about these rays here, we can start thinking a little bit like it's a particle trajectory. And I put particle here in speech marks. I don't think you really need to think about these as photons, but you can think about them as little point particles that move around inside a plasma.

And we'll find out some rules for how they move inside the plasma in a moment. And if you track their trajectory, that's where the ray's gone on. And then, you also know some places where you have lines which are normal to the wavefronts. So maybe you could reconstruct the wavefronts later on. But it's important that when we're doing this, we ignore the wave effects.

So we no longer track the phase of each particle. It's now just a little billiard ball. And billiard balls don't have phase. And so we're going to get rid of effects like interference and diffraction, and we're going to keep effects only like refraction here.

So no interference. No diffraction. Just refraction. So this is our ray model. So does anyone have any questions on these models before we start trying to use them?

STUDENT: I was just wondering, so if the Poynting vector right is $E \times B$, Does that mean that if we have any parallel electric field to k -- I'm just wondering, your point about the Poynting vector, would that break if there was like a parallel e to the main background magnetic field, or is that just the oscillating B there? If that question makes sense.

JACK HARE: It does, actually. And I know the answer to it. That's good. So if you had some background magnetic field, like in a tokamak. And then E was parallel to it.

Well, let's put it this way. The Poynting flux oscillates. And so, when you're averaging it, time averaging it, that's what gives you the actual power that's moving. If you've got a static magnetic field, your average power will go to 0.

STUDENT: OK. Yeah. Thank you.

JACK HARE: So you'll only get power flow from oscillating components here because that's what's transporting the electromagnetic energy. But it's a really good question.

STUDENT: Thank you.

JACK HARE: Like I said, I'm not completely 100% sure that rays follow the trajectory of the Poynting vector, but I'm pretty certain, after thinking about it for about 10 minutes, that they do. So if someone finds out that's wrong, please let me know and I'll take it out. OK, so everyone is going to be pretty happy if I start drawing ray diagrams.

And they'll understand that these ray diagrams represent the trajectory of little beamlets of light, and you can also reconstruct the wavefronts from them, and therefore, you could reconstruct visually what the entire electromagnetic field looks like. And we're implicitly assuming everywhere here that our magnetic field is perpendicular to our electric field, which is a pretty good approximation to the assumptions we've made so far. OK, so now, let's try putting an electromagnetic wave through a plasma.

And it's not going to be any old plasma here. I'm going to choose a slab of plasma like this. And this slab is going to be much denser at the top than it is at the bottom, which I've tried to really clumsily do with some shading here. So it's going to have a gradient of electron density going up, like that.

And of course, you remember our formula, $1 - n_e / n_{\text{critical}}$ for our refractive index. We're going to work in this regime where the density is much less than the critical density, so we don't have to worry about what happens if we get close to the critical density. And so you can see, then, that if the gradient in the electron density is in this direction, then the gradient in the refractive index is in the opposite direction, like that.

And we're going to start by putting through some phase fronts. And we're going to start with a plane wave. So this is a wave in which the phase fronts are flat and parallel and uniformly spaced. So those are my wavefronts. I'll put a little coordinate system in here.

I'm going to tend to put the z-coordinate in the direction waves are going. And so there'll be two transverse coordinates, y and x. And I'll probably just write y on most of these. I'll try and do things in a one-dimensional sense. But everything I say you can imagine could be applied to a three-dimensional picture here.

I'm going to say that this plasma slab has some length L, and it's homogeneous within that length apart from the gradient in the density here. Does anyone know what happens to the phase fronts as they emerge out from this plasma? The wavefronts or the rays. I don't mind.

STUDENT: They bend in the up or down, right?

JACK HARE: Sorry? Yeah, Daniel?

STUDENT: Oh. Yeah. They're bent downward, right? Because you've got a lower refractive index in the upper half.

JACK HARE: Yeah. You're absolutely right. So these rays will emerge or these wavefronts will emerge bent, like this. And so the rays-- which I didn't draw on before, but I meant to. So here are some rays for you here. You see how they're all normal to the wavefronts. Here are some rays for you here.

And they're going to be bent by some angle, which we'll call theta here. And it turns out, if you go and look how to do this, theta is going to be equal to $d \text{ phase} / dy \times \lambda / 2\pi$. And so we can actually put that all together and we can say it's going to be equal to $d \text{ phase} / dz \times \lambda / 2\pi$, like that.

Which, for our plasma, is $-1 / 2 n_{\text{critical}} \int \text{gradient of the electron density} dz$. And this dz here is going to be running from 0 to L. Now, I don't know how clear this is to everyone that the rays should bend or that they bend downwards or why they bend. There's lots of different ways of thinking about it.

You can go and just solve a load of equations, if you want to. I like to think of it-- and you may laugh at me for this-- as a bunch of soldiers marching arm in arm through some mixed terrain here. So here's my soldiers. I'm looking at them from above. You can see how I'm lining them up nicely with the wavefronts very suggestively.

And maybe some of the soldiers over here have got some sort of marsh that they've got to walk through, and these soldiers are going to fall behind. And because they've all linked arms, they've still got to stay in a straight line with each other. And so, as they go, they turn more and more round like this, and this is what leads to our bending here.

And you can make this a bit more rigorous if you start thinking about the rays as particles, and you think about their velocity, and you think about the speed that they're going at inside the plasma, and you realize that they're actually going slower in the denser regions, and that's going to start giving you a twist. So they're going faster in the denser regions. They're going slower in the regions with high refractive index. And that's what gives you the bending here. So this is just like a little mental model to think about when you're trying to work out why it is that the rays of light are turning. But there's many, many different ways to get this.

STUDENT: Is there a great density over all you've got use this? Hi, is there a gradient and density along the z-axis? The way it's drawn it looks like it's only within the y-axis.

JACK HARE: It isn't only in the y-axis. Yes.

STUDENT: So an integral of the gradient of density along z from that constant, then, was it-- is there no-- is there a density change in z, so what I'm asking.

JACK HARE: Not in this really simple model I'm proposing. Of course, in general, there can be a density change. Really, this should read gradient of density dot dl, where L is an infinitesimal. No. Sorry. Ignore that.

Yeah, there is no gradient in density in z. And we don't need one to get any bending. And in fact, if there was a gradient of density in z, it wouldn't have any effect on the light. It would just go forwards at a different speed, but it wouldn't get bent.

STUDENT: All right. So then that integral is just gradient of ne times L.

JACK HARE: Yes. For this very simple model, you're absolutely right. I'm just introducing the generality because we may have something different. But you're quite right. We could write this as minus $2 \frac{nc}{L} \times L$. And maybe I'll put a subscript z so that I know that it's my length scale z and times by the gradient in ne.

And if this is some simple density ramp, so I would have any $0 \times 1 \text{ minus } x \text{ upon } L_y$, or something like that, I can simply put this in and say that the entire beam is now twisted by a nice linear angle, which has an L_z inside it, an ne_0 , and $\text{minus } 1 \text{ upon } L_y$, like this. So this, if I give you some analytical result, you can then go and work out what the angle would be. And that's a super useful thing to be able to do.

As I segue perfectly into my next remark, which is to do with the first problem that this causes. So this is issues with deflection. The first issue is if you've got some plasma and you're trying to put some electromagnetic radiation in it, you want to collect that radiation. You want to put it onto a detector.

Maybe that detector is a camera, or it might be a waveguide that you're collecting microwaves with. And so that camera has some physical size. And so maybe the camera is represented by this lens here. It's got some physical size D.

And if your rays get deflected-- that was a terrible straight line. OK. If your rays get deflected by an angle greater than θ_{\max} , where $\tan \theta_{\max}$ is equal to $d/2L$ -- I forgot to put in this L here. There we go. Then your ray is going to be lost.

So for $\theta > \theta_{\max}$, you lose your rays. So that means you can't collect them. You can't detect them anymore. So this causes big problems because it means that we're going to start losing light here. And for most situations, we can use the small angle paraxial approximation and just replace the $\tan \theta$ with θ here.

So that means you want to keep your deflection angle θ , which is equal to, as we said, $d \sin \theta$ of the integral of $n dl$. That wants to be less than θ_{\max} . So there's a few things that you can do to try and do this. You can have a nice big lens.

You can have a close lens. You can put it nice and close in. Or you can use a shorter wavelength. Because if you go to a shorter wavelength, you get a smaller deflection angle, which you can see if you go back to maybe this formula here. A shorter wavelength corresponds to a larger n_{critical} , and so you'll get a smaller angle.

Now, not all of these things are possible in a standard experiment. If you've got a tokamak, you may have a limit on how big your detector can be, because it's got to fit in a gap between some magnets. You'll certainly have a limit on how close you can put it to the plasma because you don't want to stick it right inside.

And you may not be able to choose whatever wavelength you want. Perhaps you're looking at electron cyclotron emission and you've got no choice but to use the wavelength that's emitted at. So this can cause big problems. And so if you're doing some electromagnetic probing of your plasma, one of the first things you should probably do is check whether the density gradients are going to make it hard to actually measure anything. Any questions on this?

All right, so we're now going to plunge in to our first diagnostic. And the point I want to make here is although deflection can be frustrating, it can also be useful. Because we can use it to measure something about the plasma.

The first thing we're going to talk about is schlieren imaging. This word, schlieren, people often assume refers to a person. It does not. So it doesn't have a capital, despite what Overleaf will tell you.

And it's actually after a German word, schlieren, which is like streaks, because this was first used for looking at small imperfections in optics. And so looking at these little streaks here. So it's a way of imaging things which would otherwise be impossible to see because they cause small gradients in refractive index.

So let's have a little example, building up towards schlieren imaging. This first thing I'm going to show you is not schlieren imaging. This is just imaging. But my impression is that some folks need a refresher with some optics.

So we're going to start with a solid object. It's going to be this nice little chalice here. And we're going to put in some rays of light. Like this.

Now, this object is solid, and so it blocks any rays of light which hit it, these two centers ones, and allows through rays of light going past it. Allows through rays of light going past it. Very good.

And what we would probably do here if we're doing a standard imaging system is we would have a lens. So this is how you form an image. So we'll put our lens here. It's going to have a focal length F . And we're going to place it at a distance, which is $2F$ away from the object we're trying to image. I'll just put that F up there.

Now, behind this lens, if we're doing a standard $2F$ imaging system, we're going to have a focal point, and that's going to be at F away. And then, we're going to have an object plane, which is also at F away. Sorry, an image plane. This is the object plane.

And this is the lens with focal length F . So hopefully, some of you have seen this sort of thing before. You know that the rays will pass through the focal point here. He says, drawing them carefully. And what we'll end up with-- can't do this on a chalkboard-- is a copy of our image, of our object here.

But it's going to be inverted. And you can tell that because you can see the rays have changed place. So this is a nice 1 to 1 image. It's at magnification 1, and it is inverted.

So this is the simplest-- I think, the simplest possible imaging system you could possibly develop. It simply takes whatever is at the object plane and puts it at the image plane some distance away. This could be a microscope. This could be a camera. All sorts of things like that. Yes, Vincent?

STUDENT: I think I missed it. What was F again?

JACK HARE: I beg your pardon.

STUDENT: What was F , like, in the diagram?

JACK HARE: The focal length of the lens.

STUDENT: Oh. Thank you.

JACK HARE: Cool. Any other questions? OK, let's make this more interesting. Let's put a plasma here instead. And we're still going to have our lens. There still can be a focal point. And there's still going to be an image plane here.

But the plasma doesn't block the rays of light, as long as we've got, for example, not much, much less than the critical density here so that the rays can pass through easily. Instead, what we're going to have is rays that come in. And then, they're going to be deflected slightly inside this plasma.

So I haven't drawn the density gradient. We can imagine, we've just got a whole range of exciting density gradients that cause some deflections. So they deflect this ray slightly downwards. They deflect this ray slightly upwards. They deflect this bottom ray-- what have I done to this one? Let's have this one go straight.

For some reason, there's no density gradient exactly there, so the ray just goes through. And this bottom one gets deflected downwards as well. And let's say it just about makes it onto the lens. And I'll move the lens downwards to make that true. Can't do that on a chalkboard either. OK, good. So what will the lens do to these rays now?

Well, it's still going to reflect the rays, and it's going to reflect them-- let's start with this one that actually didn't get deflected at all. So its angle hasn't changed. It's going to go straight through the focal point, as you'd expect. This one that was deflected upwards is going to be deflected down, but it's going to slightly miss the focal point. It's going to be slightly above it, like that.

This one that was deflected downwards is going to be the opposite way. It's going to be slightly below the focal point. And this one was deflected downwards. It's also going to be slightly below the focal point. And that was a mistake. There we go.

There we go. It's not a very good image compared to the one I was hoping to draw, but there we go. Nothing quite works out. So we should have an image of our plasma here. This image of the plasma should still be 1 to 1 mag 1 and inverted.

The fact that I haven't quite managed to get it to work is probably just a flaw with how I managed to draw the rays this time round. Not quite sure what went wrong there. Looks good on my notes, anyway. This stuff gets a little bit tricky.

The point is, although the rays here look like they've all gone upwards slightly, they should actually still end up in the same places that they did before. And the fact that I can't get it to work right now just means that I've made a mistake while drawing it live. OK, good.

STUDENT: Jack, this is a very basic question, but what's the point of having all the rays go through the trouble of going through the lens when we could just have them go straight through and hit our image plane? Guess I missed that. You know what I mean?

JACK HARE: Yeah, absolutely. So in the top case, the solid object, the rays could go-- if the rays went straight through and hit the image plane, they will be deflected slightly at the edges. And so you'll end up with something fuzzy, so it'll be out of focus. So you need a lens to bring it to focus, which effectively is mapping the rays from where they came from back to the same place on the object plane.

If in the case of the plasma, if you don't have the lens and you just put-- if you don't have a lens and you just let the rays propagate to a screen, that's a technique called shadowgraphy, which we'll talk about next lecture, which I actually think is more difficult even though it's simpler to draw. And so I want to talk about it after this one.

So we haven't done anything here at the moment. And in fact, if you do this with a plasma, you won't see anything at all. Because all of the rays are mapped back to where they started from. And that means that you are going to end up with just the same laser beam that you originally started with or the same microwaves you originally started with. So this will be invisible. So the only way we can make this visible is to notice that the rays do not all pass through the focal point.

Now, you saw in the case where we did the imaging that all the rays did, indeed, still pass through a focal point. But here, some of them have gone above and some of them gone below. And in fact, the distance they've gone above and the distance they've gone below is directly proportional to the angle with which they exited the plasma here. And so we can learn something about the angle they can select exited the plasma by placing a filter at this focal plane. And this filter maybe looks a little bit like this.

This filter, for example, here is like a little aperture. And it lets through these two rays. Let me color code them. This one and this one. And it blocks off these two rays. And so light is coming from that bit of the plasma, where the density gradients were large will be blocked and it will no longer appear on our final image. And this is what Schlieren imaging is.

So we place a stop at the focal plane and we filter by angle. I'm going to do some exhaustive examples of this to try and build some intuition for what's going on if you're a little bit confused right now. Any questions on this before we keep going?

STUDENT: I have sort of an overall conceptual question. I feel like the highest gradients, a lot of the time, or-- well, maybe this isn't quite true. But I can imagine, in a lot of cases, this is going to be affected most by the edges where it's entering and leaving the plasma because you're-- yeah, depending on how uniform things are.

So just curious what you actually get an image of. I mean, if you-- especially if you don't have a great sense of where the gradients are or if you have gradients inside you don't know about or something.

JACK HARE: Yeah, absolutely. These images are difficult to interpret. So this is not a generic diagnostic technique that will immediately tell you what's going on. You need to know something about your plasma. Maybe you have a simulation and you do a synthetic diagnostic on it. Or maybe you've set up your experiment such that it's particularly simple.

We'll talk a little bit about some simple distributions and what patterns they make and that will give us an idea for what sorts of things we might be able to measure with this. Turbulence in the edge of a tokamak or something like that, this is maybe not the ideal diagnostic for it.

STUDENT: Fair enough.

JACK HARE: Yeah. I want to make very clear because I don't know if it came across when we were talking about it before. But we only get deflections from density gradients, which are perpendicular to the direction here. So if our ray is going in this direction, in the z direction, we sense dd_x of n_e and dx sub n_e , but we do not sense dd_z of any.

So if there's density gradients in the direction the ray is propagating, the ray will slow down or speed up, but it won't actually deflect from that. And so that helps you a little bit. You're only sensitive to gradients perpendicular to the probing direction. That might also tell you if you've got a plasma, which you think has some geometry. There may be a good direction to send the probing beam through there maybe a bad direction. So you want to think about that a little bit.

So let's have a talk about some of these stops, right? So I've said that we can place these stops here. Let's have a chat about what sort of different stops are available to us. So types of stop.

The first type is to decide whether our stop is going to be dark field or light field. So I'm putting **darken** in brackets because it will save me writing in a moment. The difference between dark field and light field is that the dark field blocks undeflected rays. So ones that do pass through the focal point.

And the light field blocks deflected rays, ones which do not pass through the focal point. So you can either look for regions where there are density gradients or where there aren't density gradients. We can also choose the shape of our stop because our stop is a two dimensional plane at the focal plane here. So we can have a circular stop.

And that doesn't care what direction the ray is deflected in, it only cares on the size of the angle. So the size of theta. So that is basically, are there any large density gradients? Or we can have what's called a knife edge, which is linear like this. And that is sensitive to density gradients in only one direction and it still cares about the size of the density gradient. That is like x hat here.

So we could, for example, have a stop at the focal plane. No, it's not going to do it. OK, fine. I can't get a nice, round circle. We can have a stop which is an opening inside an opaque sheet of material here. And this opening could be positioned such that the focal spot in the absence of any plasma sits inside it. What sort of stop would this be?

STUDENT: Light field?

JACK HARE: So this is a light field stop. And what shape is it?

[INTERPOSING VOICES]

Yes, OK, circle. Thank you. We could also have a stop that looks like this. And we can position it such that the focal spot is actually within the opaque region. And what sort of stop would this be?

STUDENT: Dark field.

STUDENT: Dark field.

JACK HARE: OK, dark field. And what shape is it?

STUDENT: Linear.

JACK HARE: So the knife edge here. Yeah. We call it a knife edge because actually using a razor blade is a pretty good thing to have because you get a very nice, sharp, uniform edge to it. OK, and so depending on these stops you can think of as filters in angle space, right? So they allow through certain angles. You can think of arbitrarily complicated versions of this.

There's a technique called angular-- not fringe. Angular filter refractometry, which has a set of nested annuli, which let through light which has been deflected by certain specific angles. So the world is your oyster. You can come up with all sorts of exciting different stops if you want to.

One thing I will note is that the dynamic range of your diagnostic, which we'll talk about more later, depends a great deal on your focal spot size. So I've shown these focal spots to be relatively small here. But that focal spot size, at least in a diffraction limited sense, it covers an angle, which is equal to the wavelength of your light over the size of your lens, the diameter of your lens here.

And so you might end up having focal spots which are not small, but actually could be rather large. And then, part of the focal spot could be obscured and part of the focal spot could be clear for some given deflection angle. And in general, when we've got a plasma here, we have, say, our small focal spot before we put any plasma in the way.

When the plasma is gone in the way, different rays of light have been deflected by different amounts. So this thing may take on some complicated shape here. And this is the shape that you're filtering. You might be filtering it with your knife edge like this, or you might be filtering it with your circular stop like this.

So we're basically filtering the rays based on how far they've been deflected. At the focal plane, there's no information about where the rays came from inside their plasma. So their spatial information has been lost. The only thing we know about them is their angular position.

So rays, which are deflected by an angle θ_1 from at the top of the plasma, are rays which are deflected by the same angle from the bottom of plasma. It ends up at the same place in the focal plane, even though they came from different parts of the plasma. This is the magic of geometric optics.

So any questions on this or should we do a little example? OK, let's do our example. So let us consider a very simple plasma. This plasma-- we'll have the coordinate system y vertically and we'll have z in the direction of propagation as we discussed before. We'll have rays. Well, I'll draw the plasma first.

So the plasma is going to have a density distribution that sort of looks Gaussian ish, some sort of nice peaked function. So this is density, n_e . So you can think about this, for example, as like a cylinder. So you've got a cylinder of plasma, like a z pinch, and you're probing down the axis of this z pinch and it's got a Gaussian distribution of density to it. Anything like that.

And then, we'll have the rays of light coming through. So we'll have rays of light, which are sampling very small density gradients at the edges here. And these rays will just go straight through.

There's also be a ray that goes through the center, which also sees a very small density gradient here, right? At the center of this distribution, the density gradient is zero. But the edges here where the density gradient is large will have some deflection.

And if we place our lens, as we did before, some distance away, it's going to focus those rays onto our focal plane. And in our focal plane, we're going to put some sort of stop here. So the undeflected rays are just going to go straight through the focal point.

But the deflected rays are not going to go through the focal point. The deflected rays are going to go above and below. So now we can put a series of stops inside here. So we can have a stop on that blue dashed line that looks like a light field knife edge. We can have a stop that looks like dark field knife edge. We can have a stop that looks like a light field circle.

And we could have a stop that looks like a dark field circle. And what we're going to do is sketch out what we expect the intensity to look like from each of these different knife edges here. So if I plot this one-- I'll just draw the density distribution again like that. So that's our density. Now we want to know what our intensity distribution looks like. So this is now intensity.

We have our initial intensity 1 and we have 0 intensity corresponding to all the rays being blocked. So where for the light field knife edge do we see no intensity?

STUDENT: In the upper side where there is the largest density gradient.

JACK HARE: I'm sorry, I didn't hear that properly. Can you say it again?

STUDENT: I think the upper portion where there is the largest density gradient.

JACK HARE: So I think what you said was in the upper side where there's the largest density gradient, right? So this is where the density gradient is very large. So we expect outside of that region our intensity would be pretty much constant. But inside that region, we'd expect the intensity would drop to 0. Because we're blocking the rays which have a large deflection angle in one direction upwards and the rays which have a large deflection angle in one direction upwards corresponds to that specific density gradient there.

OK, anyone want to have a go at telling me what happens with this one?

What's happening at the edges? What's happening out in these regions here where the density gradients are small?

STUDENT: They're cropped out because they mostly go through the focal spot.

JACK HARE: Right. Yeah, so these are going to be 0, right? So 0, 0, is there any region where it's not zero?

STUDENT: I think it's not 0 for the high gradient region that's lower in y. Because it's deflected.

JACK HARE: So you think it's not 0 for this region here?

STUDENT: Yes.

JACK HARE: Does anyone agree with Sara or does anyone disagree?

STUDENT: I think that looks right.

JACK HARE: OK, so we're talking about this ray here. Yeah. So this ray is, indeed, passing low down compared to the knife edge. So we'd expect this to show some intensity here. But not for the upper one, which passes higher up.

OK, good. We're getting there. What about for the circular light field? Anyone want to tell me what the intensity looks like here?

STUDENT: You'd probably get three peaks. So on the edges where the density isn't really-- where the density is just low and then right in the middle where the density is high, but not really changing too much.

JACK HARE: Yeah. So you say three peaks, I'm going to think about them as two notches, but I agree with what you're saying. So these notches here correspond to the points where the density gradient is largest here. And so those rays got a big deflection angle there being blocked out. What about for the final one here, dark field circular aperture?

STUDENT: Well, there you're basically notching out all the minimally deflected rays, so you lose the ones on the edges and center.

JACK HARE: Yeah, so we'd have something that looks like this, right? So we would just see light where there was a significant deflection angle. OK, this may still seem a little bit abstract, so we're going to try one more thing. And I hope that you don't hate me for spending so much time on Schlieren, but I absolutely love it, so that's your loss.

Which is going to be a two dimensional example here. I'll just draw on this distribution function. There we go, density function now on that last one just so you've got it. So this actually goes back to a sketch that I did during the $b \cdot$ lecture where we stuck a little $b \cdot$ inside the plasma. And we have plasma flow coming from left to right. And because it collides with this, we get some sort of bow shock like this.

And as we all know from our shop physics, at the bow shock, we have strong density gradients like this where the density jumps across the shock here. So if we look at this system and we have our probing laser coming towards us, so our laser is looking towards us like this, we are the camera. What would we see about this bow shock? What could we tell about it?

So maybe the first thing we could do is ask for a light field circular aperture, what do we think we would see here? And if you're not quite sure, at each of these places where I've drawn this little arrow, you could say that the density maybe as a simple model looks like a sort of hyperbolic tangent type thing like that. It's got some region where the density is n_{e0} some region where the density is n_{e1} , and then some region where the density changes rapidly.

That's not true for a shock, but it's a model just to get us thinking about what this looks like. So what would I see on my image? I've got this expanded laser beam going through this shock and I've decided to use a circular light field stop.

STUDENT: I might be doing this backwards, but from-- it looks like what we drew above. That would mean that you're not getting your steep gradient sections, so it should be dark where the bow shock is.

JACK HARE: So you'd actually have an image. Yeah, you'd have an image, which if I was on a chalkboard, I could do as an eraser, but it's actually quite hard here. If you imagine this is filled, a green laser beam beautiful nice laser beam image. Then you would have a region where it-- no, the arrays is not very good. If I can make the eraser smaller that would work much better you would have a region where there was no light whatsoever and there was just darkness, right? So you just have this dark region here corresponding to the bow shock.

And if you did this with a dark field circular aperture, then you'd have the opposite. You'd have complete darkness and you would just have a region where the bow shock shows up very nicely like this. And so this is actually typically what we use for shock measurements is dark field circular aperture.

If you do happen to do something like dark field with a knife edge and you put your knife edge like this, so you were measuring density gradients that were in this direction, you would end up seeing something a little bit like just one half of the bow shock like that. So if you set-- your probe was sitting here, you would just see a little bit of the bow shock. You wouldn't see this section because the density gradients wouldn't be in the correct direction. So you wouldn't be able to observe those.

But the knife edge might be a better choice for some shock geometries if you're only interested in gradients in one direction. Aidan, I see your hand.

STUDENT: Yeah, I'm curious how distinct the actual image gradients would be given that this is like a cylindrical phenomenon, right? The shock. So I assume there's some density gradients that's they slowly become parallel to your propagation direction as you rotate to the--

JACK HARE:

Yeah, so you won't see those density gradients, but you will see the ones on the edges of the shock. Yeah, it depends on the exact shock morphology, whether it's some extended like a bow shock that's extended like this or whether it's a bow shock that's sort of rotated like the tip of a nose cone on an aircraft or something like that. So it will make a difference. But this is just to try and get a feel for what this looks like in terms of imaging here.

I can send around some papers later, which have some very nice images of Schlieren of shock structures that show what sort of quality of data you can get out from this. So any other questions on Schlieren? I will then just summarize exactly how to use it and when not to use it and things like that. But if anyone has any questions on these sort of worked examples we've done, please go ahead and shout out.

So Schlieren is good for visualizing density gradients, right? And in particular, it's good for visualizing strong density gradients. So in particular, it's good at visualizing shocks. So if you're looking at shocks in plasmas, this is a nice diagnostic. And it's particularly nice because it's very simple to set up.

As I showed you, it's just got at minimum a single optic. You need to form a focal point, so you need to have an optic. You need a focal point and you need a stop. So this is a very simple thing to set up. You can do it very, very quickly. The trouble is, although it's simple, it's very difficult to be quantitative.

We can say there is some sort of density gradient there. We think the density gradient has to be larger than some certain limiting value. But other than that, we can't really say much more than that. So it's useful for seeing where shocks are and their morphology, but is not useful for measuring gradient of any. So we can't measure gradient N_e . We can measure shape and location of the density gradients.

Now, there are some ways where you can make this more quantitative. So imagine you've got a little beam of light coming in like this. This is going in this direction. And you have a knife edge that looks like this. That knife edge is going to be entirely blocking what's coming through, we'll have 0 light coming through.

But if we have a beam that comes through and it's lined up like this, you can see that now about half the light is going to get through. And if we have a beam coming through that's lined up above the knife edge, we'll have all of the light coming through. And it turns out that if you have a large enough spot that you can actually see this partial obscuration of the focal point, you end up in a regime where the intensity that you see on your detector eye is, in fact, proportional directly to the gradient of the electron density.

I'm putting it in brackets because you need a large uniform focal spot. And I'll talk in a moment about why that's actually extremely difficult using a laser, which is what we normally use here.

OK, if you can't guarantee that you have a nice large uniform focal spot, you could also use a graded neutral density filter. So a neutral density filter is sort of smoky glass that blocks light. And you can change the amount of impurities in it to block more light. So we could carefully fabricate for ourselves a neutral density filter that has a changing absorption. So we can have a gradient in, say, we'll call the absorption α here.

And then, the idea is that different rays of light at different points of your stop will either come through not at all or very attenuated. Or they will come through partially attenuated or they will come through lazy not attenuated at all. And so if you have a really good graded neutral density filter, you can, again, end up in this regime where you actually get some sensitivity.

To the position of the beam. And that can get you back into this nice regime up here where the intensity of your signal is actually proportional to the gradient of the electron density. So that'd be a very nice place to be, but this is hard to make. And you still need a uniform beam to start with, which is also hard to make.

So these techniques, a lot of Schlieren was actually developed using light sources, which are very different from the light sources that we have to end up using in our experiments with plasmas. And so you end up, yeah, so just looking at the beautiful pictures that Matthew put in the chat here. These pictures here are absolutely glorious. And you can see there's a huge amount of detail on them.

And this detail is due to the fact that actually I think for these ones they're using the sun as the back lighter. I may have forgotten this exactly, but light source is the sun. And the sun is actually quite large. It's an extended object, you may have noticed, in the sky. And this gives you really nice Schlieren imaging. But we don't tend to use nice large objects for our experiments with plasmas. We tend to use things like lasers.

And lasers are not a good Schlieren source. And if you want to read in more detail why lasers are actually a really bad idea, you should go read the book that's listed in the bibliography by Settles, which is an absolutely cracking book. Has lots of lovely pictures, like the ones which were just posted in the chat. And also, a very detailed description of this stuff.

But can anyone tell me why-- if lasers aren't very good, why do we end up using lasers as the light source for Schlieren imaging in plasma physics? What property of a laser is it that we particularly want?

STUDENT: Monochromatic light.

STUDENT: Monochromatic.

JACK HARE: So monochromatic is somewhat useful. Actually, it turns out you can do really cool Schlieren techniques with a broadband light source as well and filters because the different wavelengths will be reflected by different amounts. And so if you have multiple cameras with different filters, you can really carefully reconstruct the deflection angles.

So monochromatic is a good guess. But actually, we'll be OK with broadband. What else-- when you think laser, what do you think?

STUDENT: Coherence.

JACK HARE: Coherence. Actually, coherence isn't good for this. And we've been using a picture and the coherence screws up our nice simple picture. It's going to cause problems. So we really want coherence for interferometry and we absolutely hate it for Schlieren. And that's another thing that Settles says in his book, so coherence is a bad thing. So I'm going to put coherence. No, we actually don't want it. What else do we think about lasers?

STUDENT: Well, they tend to produce nice round spots. So it's also easy to deflect them without changing them that.

JACK HARE: Nice beam. Yeah, you can do that with other light sources. But fine, it is a nice beam. Anyone ever shone a laser pointer in their eye? If not, why not?

STUDENT: High energy density.

JACK HARE: They're very bright. Should we put it that way? OK, so lasers are extremely bright. Why do we want a bright light source when we're dealing with plasmas?

STUDENT: The plasmas glow.

JACK HARE: Yes, so we need to overcome the plasma glow, should we call it, or should I say emission here. So when you're working with jet planes flying around and you're taking Schlieren images of them using the sun as your background, you don't actually have to worry about the plasma or the shocked air around the plane glowing and ruining your measurements.

Whereas, a plasma, you really do. It makes an awful lot of light. So you need an incredibly bright light source. And so despite the fact lasers are terrible for Schlieren it's the brightest light source that we have available so we end up using that. But the trouble with lasers is that they have a very small focal spot.

So there are, in fact, incredibly easy to focus down to a point. And we don't want that. We want a nice large focal spot. And it's actually very difficult to get a laser to decohere enough to get a nice big focal spot. But maybe there are some techniques we can do.

And because we've got a small focal spot, laser Schlieren is effectively binary. And I'll explain what I mean by that when I remember how to spell Schlieren. What I mean by that is the spot is so small, the laser Schlieren, it is either completely blocked by the knife edge or it is completely visible, unblocked, by the knife edge.

And so when you do a laser Schlieren image, you get an image which is either dark or light, 0 or 1, no intensity or full intensity. And so you don't get those nice images that we were just looking at that were linked. And we also don't have the ability to use the change in intensity to measure the gradients in the electron density because the intensity is either 0 or 1.

So we have very, very limited dynamic range. You can think about it that way. We have no dynamic range. We even know the density gradient is larger than the number or it's smaller than a number and that's it. So it makes very good shock pictures, but you can't do all the beautiful techniques that people use Schlieren for in standard fluid dynamics where they have access to different sources.

So someone came up with a nice, bright, incoherent large area laser that we could use. That would be absolutely great. I have a few ideas along this direction, but haven't had a chance to try them out yet. So yeah, they have some serious limitations. So Schlieren is a lovely technique for plasmas. It's very limited. It's obviously useful when you have shocks, so you need to be having fast moving plasmas.

And so this is not really a technique that we would normally use inside say a tokamak or a magnetically confined device, but it is very good when you have larger densities. All right, so any questions on this? Nigel, yeah.

STUDENT: So did we ever say how exactly the critical density was determined? Is that for a later lecture?

JACK HARE: Do you mean like, what is a numeric-- what is an analytical result for the critical density?

STUDENT: Yeah, when we had like $1 - n/n_{critical}$ critical. Where that was ever determined? Or is it just kind of guess and check.

JACK HARE: No, absolutely not. You can write it down. It's a function. And I neglected to write it down here. And I am not going to try and remember what it is. But it is a function which has inside it the frequency of the laser light, epsilon naught, e and me and the-- not the density. Definitely doesn't have the density inside it.

So this effectively comes from rearranging the plasma frequency here. And you can find this in Hutchinson's book, and I really should have written in my notes, but I just didn't, OK? And I'm not going to try and bullshit you and look it up now.

So yeah, the critical density is a number that we can find. It is uniquely defined for every electromagnetic wavelength and that is the density, which you cannot go above. And if you're sensible, you'll work well away from the critical density. Because if you get anywhere close to it, it really screws up most of these calculations.

So most of the time, we're relying on this approximation that the density we're using is much less than the critical density. And we'll come across that approximation very strongly when we deal with interferometry in a little bit. But we did also use it already when we were deriving the Schlieren angle. When we got to this point here, we used this linear approximation where the density is much less than the critical density. Any other questions? Yeah? No?

STUDENT: So can you hear me?

JACK HARE: I can hear you. Yeah.

STUDENT: OK, cool. So would it not be possible to use the quote unquote glow of the plasma itself to use in some other part of the plasma that isn't glowing itself or that it's glowing at some other frequency? Say, once we're putting--

JACK HARE: I thought about this. I think it's a really cool idea. You definitely have to arrange your plasma so there's a bit that's glowing and a bit that isn't. It's a bit you want to measure as also glowing that's going to make things very, very difficult. So yeah, you could potentially have something like that, like a ICF hotspot backlighting the whole plasma. That sort of thing could work. Yeah. But I don't know of anyone who's done it.

STUDENT: OK, Thank you.

JACK HARE: All right, we're past the hour now. So I think we'll leave it here. I will be back in the classroom next Tuesday. I look forward to seeing all of you there and all of the Columbia folks online. And yeah, enjoy your weekend and bye for now.