[SQUEAKING]

[RUSTLING]

[CLICKING]

JACK HARE: So I'm going to begin by giving a recap of the last lecture, just very briefly. And then we'll keep going with where we left off before. So you remember we talked about B-dot probes. These were little loops of wire, like this. They've got some area, A. They've got some magnetic field pointing through them. We're only sensitive to the component of the magnetic field that is parallel to the normal of this loop.

And what we want to do is measure some voltage across this loop. And we showed that the flux through this loop, phi, was going to be equal to the integral of B dot ds over this area. And the voltage that we were going to get out was going to be equal to the time rate of change of the flux, which is equal to B dot A, like that. So by measuring the voltage on this loop and then integrating it up, we can work out what the magnetic field is doing as a function of time here.

We also talked about-- these are B-dots. We also talked about Rogowski coils, which is effectively a set of B-dots arranged around the circle. So we have a circle like this. But now, we have a load of B-dots arranged all the way around. And we join them all together.

And we measure the voltage again here. And this tells us something about the current which is flowing through this loop here. And we found that the voltage that we get out here is equal to the number of turns per unit length. So for example, this might be three turns per millimeter.

The area of each of these little loops here times by I dot-- and there was an optional factor of mu r here. This is different from 1. Or if it's different from 1 because you've chosen a material that saturates, like steel or something like that. So these Rogowski coils are used for measuring current in a very similar way to B-dots. We get out of voltage, which we can then integrate up and find the current.

So then we talked about how we could use these simple devices to measure the plasma conductivity, which we represented by this symbol, sigma. And the bar here is just some sort of average conductivity here. We made lots of assumptions, such as the fact that we're in steady state, which allows us to ditch a load of terms.

And we then ended up with a power balance where we balanced the ohmic dissipation within the plasma, which is the integral of the current density squared over this plasma conductivity with the Poynting flux, which is the energy we're pushing into the plasma in the form of a loop voltage driven by our transformer.

In the case of something like a tokamak, this could be the voltage inside a Z-pinch as well, and times by the current flowing in that same direction here. And so this was a balance between the ohmic heating and the external applied power here. And we found that, for something like a tokamak-- I'm just giving you half a donut here-- if we put a Rogowski around the cross-section here, we could measure I phi.

And if we put a voltage loop around the tokamak like this, which will effectively sense the same voltage that's being driven inside the tokamak by our transformer, we measure V phi. So we can measure those two things. And with a little bit of rearrangement here, we ended up with a formula for our conductivity, which was equal to I phi over V phi with some geometric terms 2R over A squared, where R is the major radius and A is the minor radius of our tokamak.

And this was particularly interesting because we have good models for plasma resistivity. And we know that this conductivity is proportional to a load of constants times Te to the 3/2, like that. And so by measuring the plasma conductivity, we could effectively measure the temperature inside our plasma just using two loops of wire, which is pretty cool.

Finally, we started talking about pressure balance. So we looked at our MHD equation-- J cross B minus gradient P equals 0. This is steady state MHD in completely general three dimensions. And the idea is we want to be able to measure B and use that to infer the pressure.

And so then the pressure is an important quantity because the density and the temperature of the tokamak, or whatever device you're trying to measure, gives you the fusion power output here. And we decided that trying to solve this for like a three dimensional tokamak or a stellarator was very hard. And so we were just going to focus, just to start with, on a so-called straight tokamak, which is just where we take the torus and unfold it into a cylinder, like this. So we went for cylindrical geometry.

And we also said we would focus on just the first two modes, the m equals 0 and the m equals 1 modes. And this was the idea that if we put some probe out here measuring the theta, the magnetic field going around our little cylinder here, we could decompose that B theta, as we can always do, using a Fourier decomposition.

So would have some Bc component-- C0. And then we'd have a sum over these series of m modes from 1 to infinity with coefficients like Cm, [? pulse, ?] m theta-- I'm running out of space-- plus Sm phi m theta, just about made it. OK, cool.

And the idea is really, most of the time, we're just going to be interested in these lower order modes. But of course, if you have lots and lots of B-dots and digitizers, you might want to go after some of these high order modes. But what we're going to show in today's class is that the size of the m equals 0 mode tells you something about energy confinement. And the size of the m equals 1 mode tells you something about the displacement of your plasma, both of which are useful things to be able to measure. So I'm going to pause there and let you ask any questions on the material that we covered last lecture. Mm-hmm?

STUDENT: [INAUDIBLE]

JACK HARE: Yes.

STUDENT: [INAUDIBLE]

JACK HARE: I haven't drawn the return loop there. If you want to put the return loop on, which is a good idea for canceling out some of the stray fields, when you get to this end, you simply wind it back here and put it out this side. So one way you can do this simply is you take coaxial cable, like a BNC cable, you strip off the outer conductor.

And so you just have the dielectric and the inner conductor. You take some magnet wire. And you back-wind it along the section that you've stripped, and then solder it back onto the outer sheet. And so that makes this configuration very easily. Other questions? Anyone from Colombia? Yeah?

STUDENT: [INAUDIBLE]

JACK HARE: Yeah, so the question is, does the Rogowski have to form a complete loop? Or can you go halfway? And the answer is, no, you definitely have to form a complete loop. When you're deriving this, you're using Ampere's law, where you have a loop integral. And that requires you to do a full circuit around the surface that you're integrating around, but yeah. Matthew?

STUDENT: Yeah, I think you mentioned this last class, but I just want to make sure--

JACK HARE: I can't hear you. And that's probably my fault. So just give me a second. Can you try to say something now?

STUDENT: Test, test.

JACK HARE: Yeah, OK. I can hear you. Can you say your question again, please?

STUDENT: Yeah, so I think you mentioned this last class, but I just want to make sure. So another assumption in your computation of the conductivity is that all the current in the plasma is chemically driven, right? You can't have ECCD or neutral beam-driven current or anything like that. Otherwise, that V phi measurement wouldn't be accurate, right?

JACK HARE: Yeah, absolutely true. So in case anyone didn't hear properly, here, we're assuming that all the things that's driving the current here is to do with a transformer that's driving the current.

STUDENT: Right.

JACK HARE: It's not to do with waves or something else like that. If you have a tokamak where you're driving waves, this very simple treatment won't work anymore. But many tokamaks, for a long time, they were just inductively driven.

STUDENT: Right, right, right. All right, thank you.

JACK HARE: OK. Right. So yeah, we started with this derivation with the Poynting flux. But I quickly simplified it to this circuit model. But you're right, if I kept the Poynting flux in there, then obviously I'm capturing all of the power I'm putting into the plasma. But we wouldn't be able to do it in this very simple way with two loops anymore because these two loops are only measuring the contributions to the Poynting flux from the transformer. These loops won't capture your lower hybrid current.

STUDENT: [INAUDIBLE]

JACK HARE: Yeah. So if you know that power, if you know how much power you're coupling--

STUDENT: [INAUDIBLE]

JACK HARE: Right, exactly. Yeah, so you might know you're putting in 10 megawatts. But you may not be coupling 10 megawatts. So yeah, any other questions? All right, let's keep going.

So I'm just going to write that. Actually, I don't think I need it. I'm not going to write the decomposition again. Hopefully, you've remembered this. I'll write it again if I need it sometime later. So we're going to focus, first of all, on the m equals 0 term from that decomposition. This is often called the diamagnetic term.

We'll see, in fact, that the plasma does not have to be diamagnetic. It just turns out, for many plasmas we study, like tokamaks and stellarators, they are diamagnetic. And so this is a reasonable thing to call them. But we'll talk about paramagnetism briefly as well.

And the point of studying this, just so you don't lose all hope as we plow through all this mathematics-- the thing we're going for is it allows us to measure the energy confinement time inside our plasma, tau e. And you probably recall that the Lawson criterion depends on the density and on this tau e quantity.

And if you want to put a T on there and call it the triple product, you can do that as well. And pretty clearly, right now, we're able to measure T very roughly using two loops. And we're going to see it only takes three loops of wire to measure tau e. It takes significantly more than that to measure n, unfortunately. But we can get a pretty decent way to being able to measure the Lawson parameter for our plasma with just some little loops of wire, which is pretty cool.

So let's have a look at this. So the m equals 0 term, as we discussed last week, is a term that says something about the size of the plasma. This is something that tells us-- has no information whatsoever about the azimuthal profile of our plasma. We've thrown all of that away. So we're only interested in things that move in the radial coordinate.

So can, for example, take my MHD equation and dot it with the radial unit vector and say, this is equal to 0. Previously, the whole equation was equal to 0. But now, I don't care about all the other terms. They'll be 0 by symmetry. I'm just going to focus on this one. And that allows us to write down a simplified form of this MHD equation, which is minus dP/dR plus dBphi/dR times B phi upon mu 0 plus 1 upon RdR/d theta dR times B theta over mu 0.

And you'll recall that all of these mu 0's are popping out because we've replaced J with the curl of B over mu 0. And you'll also remember that, although we don't really have a theta in our analysis because we're dealing with our straight cylinder, we agreed that we would still continue to call this direction along the cylinder phi, like that, because, if we do bend this back around into a torus, that's what it'll be. And of course, the other direction that's important here is the theta direction, which wraps around [? but ?] your poloidal angle.

OK. If you take this equation 1 and you spend a little bit of time fiddling with it-- and this is where I'm not going to do the whole derivation in class-- you can take this equation 1 and you can multiply it by R squared. And then you can integrate it up-- equation 1 times R squared dR. And you can integrate it between 0 and A, where A is the edge of your plasma, here.

And if you do all of that plus some rearrangement-- and I suggest you go away in your own time and give this a go, especially if you're studying for quals at MSE, because this is the sort of question we love to ask-- you will end up being able to rearrange it into a dimensionless parameter called beta theta.

And beta theta is defined as 2 times mu 0 over B theta A squared. So this is the magnetic field at the edge of the plasma. And this is multiplied by the average pressure. So this is not necessarily the standard definition of beta that you've seen. It's a definition of beta that uses an average of the magnetic pressure-- sorry, an average of the thermal pressure over the whole plasma cross-section. And we're comparing that to the magnetic pressure at the edge. So it's a little bit subtle because it's not actually a local dimensionless parameter. This is an average dimensionless parameter here.

You've taken this. You've integrated it up. You've rearranged it so that you have-- on your left-hand side of your equation, you have this, which is what you're going after. On the right-hand side of your equation, you find out you have 1 minus B phi at the edge-- so B phi A squared minus the average of B phi squared across the plasma cross-section over B theta A squared.

So again, these two are at the edge. And so you could measure them just using B-dot probes, which are sitting at the edge, not inside your plasma. This one is going to give you more trouble measuring it. Now, if I asked you to measure B phi, you'd say, that's not a big problem. I'm just going to stick a B-dot all the way around my plasma. And you see how this is going to average the average B phi all the way over the area. And you say, job done.

But I didn't ask you for B phi. I asked you for B squared phi. I also didn't ask you for B phi squared, which you could do just by squaring that. So clearly, this is going to cause us some problems. And we don't have a diagnostic that just measures B squared phi. So we can't just get this out easily.

And so to make some progress on this, we're going to have to make some assumptions. And this is where we're making even more assumptions in order to be able to make progress. Before we keep going, does anyone have any questions about this? Yes?

STUDENT: [INAUDIBLE]

JACK HARE: All I need is for-- the question was, can I put a B-dot probe at A and have it survive? In fact, all I need is for the plasma pressure to be 0 where I place my probe. And it could have been 0 at A over 2, much further in. As long as it's still 0 where I put my probe, this will all work. So it needs to be a vacuum measurement. There cannot be any plasma at a larger radius than my probe. But yeah, OK, other questions? Yes?

STUDENT: [INAUDIBLE]

JACK HARE: It's free. It's a free agent. This is a fraction. The 1 is by itself. Yes, cool. I can't remember what equation this is in Hutchinson, but this is definitely one of his equations. Yeah, other questions? Yeah?

STUDENT: [INAUDIBLE]

JACK HARE: Yeah, so we're bothering with this definition of beta. Actually, the beta, it doesn't matter. What we've done is-- just ignore that if you don't like it. I've shown that I can write the average pressure in terms of some quantities. And we can certainly measure these quantities. I could multiply this up by the way, right?

I could put this over on the right-hand side. Then I'd have a B theta A squared minus this blah, blah, blah. So far, it's looking good apart from we can't measure this one. And we're going to explain how to measure that. So if you don't like thinking about B theta, just think, hey, you've measured the average thermal pressure inside the plasma. That's already pretty good. So questions from Columbia? All right. Yeah, another question here?

STUDENT: [INAUDIBLE]

JACK HARE: Apart from the last one, the squared term, yes, yeah. OK, good. So to make progress, we are going to make some assumptions. This happens quite a lot. Some people are like, is that assumption good? Well, it doesn't matter. It lets me keep going with the mathematics. So it's good enough for now. If you want to do it properly, you'll often have to resort to numerics as opposed to analytical results. So we want to make analytical progress here.

So we're going to assume three things-- first of all, that our toroidal magnetic field, B phi, as a function of R is roughly constant, as in constant spatially, it doesn't change much in R. We're going to assume that B phi is much, much larger than B theta. And we're also going to assume that beta phi-- and we haven't actually defined that yet. But if you squint over here, you can probably work out what it's going to look like.

Oh no, I'll define it here. OK, 2 mu 0 B phi at the edge, B phi A squared times by the average total pressure is much, much less than 1. Can anyone tell me a system for which this is a reasonable set of assumptions? Yeah?

STUDENT: [INAUDIBLE]

JACK HARE: Yeah, high aspect ratio tokamak, high aspect ratio stellarator as well. So some sort of tokamak or stellarator-- these are machines-- stellarator, OK. These are machines where we are deliberately operating at low beta to avoid MHD instabilities, where we have a very strong toroidal magnetic field to prevent MHD instabilities, like the kink.

And if you're at high aspect ratio, then your magnetic field as a function of R is going to be dropping as roughly 1 over R. So I just put my circular cross-section far enough out that this is roughly constant across the circular cross-section. So these are actually not terrible assumptions to make for certain plasmas. But they are violated for other ones. This won't work for reversed field pinch where this is not true. But it does work quite well for these.

OK. So using these assumptions, what can we do? We can say, again, that B phi is roughly constant, which means it's also roughly equal to the average value of B phi. But to take into account the fact that B phi probably does vary just a little bit, I'm going to write this as B phi at the edge, B phi A, plus some small quantity delta B phi, which will vary as a function of R.

If I take the square of this then, B phi squared, I'm going to get something that's roughly equal to B phi A squared plus 2 delta B E phi A plus a term, which is quadratic in this small quantity, delta B phi. And we will say that we are going to drop the quadratic terms. And we're just going to keep the linear terms here. So this is a standard perturbation theory trick.

And if you squint at this for a little bit, you'll realize you can then write B phi A squared minus the average of B phi squared is equal to 2 B phi A B phi A minus the average of B phi. Conveniently, this term here is the same as this term here. But we've now rewritten it in terms of things which we can measure. So we can now finally make some progress.

I will say, I find this derivation slightly hand-wavy. And it's not in Hutchinson's book. He skips straight to the result. So you might want to sit down and work it out yourself and convince yourself that it's really true. We're definitely playing fast and loose with our averages as we do this. So this suddenly randomly including these angle bracket signs is maybe a little bit dodgy. So give it a go. You might find it interesting.

OK. So that means we can now write our colloidal beta as roughly equal to 1 plus 2 B phi A by A minus the average of B phi over B theta A squared. And if I go back and I have a look at my little straight tokamak, I can put in a B-dot here to measure B theta A.

I can put it in a loop that's aligned with phi. And I can measure B phi A. And as I've already said, I can put a big loop that goes around the entire plasma and measure the average value across the plasma of B phi. If I number those 1, 2, and 3, you can see that we have now measured all the things that we wanted.

Where's 2? This one, this one, this one. Great, so it's pretty clear that we can actually measure the average pressure inside our plasma using just three loops and a lot of assumption. But these assumptions are reasonable for the sorts of plasmas that many of us work on. Any questions on that before we keep going and I show you how to get the [INAUDIBLE]?

Let's briefly return to our definition for beta phi because this tells us something about the [? damn ?] magnetism that I hinted at earlier. So beta phi we can write very simply as beta theta. But we need to swap out the magnetic field theta in beta theta for a beta phi field. And so we just multiply this by beta theta A squared-- sorry, not beta, B theta A squared over B phi A squared.

And if you remember what that looked like, you're just swapping out the B thetas for the B phis. And so then you'll have your definition of B phi. And if we pop in this B theta definition, we'd end up with B theta squared A over B phi squared at the edge plus 2 times 1 minus the average value of B phi over B phi at the edge.

I can't go down low enough. But if you look at this, you multiply it by that ratio of B theta over B phi. The 1 picks up that term. The B theta conveniently cancels with the denominator on this. We end up with something like this.

We can go further than this for the assumptions that we've made because we already assumed that the poloidal magnetic field is weak compared to the toroidal magnetic field. So we'll drop this. Now, can anyone tell me any bounds on beta?

STUDENT: [INAUDIBLE]

JACK HARE: It's positive. Why is it positive?

STUDENT: [INAUDIBLE]

JACK HARE: So the definition is 2 mu 0 over B phi A squared average pressure. Every single one of those terms is positive. So beta phi is positive. It has to be. There might be other limits to MHD instabilities. But this is the most basic of them. That also implies that this term is positive because that's just the same as beta phi.

So from that, we can see that B phi at the edge here is always going to be larger than the average value of B phi. If that were not true, then this could be a negative number, which would be bad. And what this is effectively telling us is something that you already know, but is nice to show just using some MHD, which is-- if you think about your tokamak as a function of major radius, you've got your magnetic field, which is falling off as 1 upon R, like that. The cross-section of your tokamak is something like this.

We have decreed that, in fact, B phi at the edge is the same on either side. So maybe I need to go a little bit further out until this line flattens off. But what this is saying is that the average magnetic field has to be less than the magnet field of the edges, which means that this magnetic field has a lovely little dip inside it. So the magnetic field is less than you would predict from the 1 over R. And it is, indeed, that dip that the thermal pressure sits inside, there.

So the plasma reduces the magnetic field. And that is the definition of a diamagnetic substance. Any questions on that? Yes?

STUDENT: [INAUDIBLE] magnetic [INAUDIBLE]?

JACK HARE: If there was no pressure gradient, there'd be no plasma. And so it wouldn't happen. There would just be a vacuum. You can't have a plasma with pressure gradient because you need to have a plasma which has 0 pressure A because otherwise you don't have the edge of a plasma. And if you have no gradients, then the only other pressure available to the plasma throughout the entire volume is 0, yeah.

STUDENT: [INAUDIBLE]

JACK HARE: It would, under these strong assumptions that we've made. And we're going to talk about relaxing these assumptions and finding a plasma which is paramagnetic, which actually enhances the magnetic field, in a moment. But for something like a tokamak, then this is always paramagnetic-- diamagnetic, sorry, yeah. Other questions? Hm?

STUDENT: [INAUDIBLE]

JACK HARE: I mean, this is a pretty good model. You can go to Hutchinson's textbook and see measurements he made in 1973 when I assume his mustache was absolutely brilliant. And he does a very good job there. The error bars are quite tight on measuring the energy confinement time.

So this was possible in the '70s. So I imagine it is still possible now. We're going to go on a little bit about how difficult it is because effectively you're trying to measure-- it's a very small dip. I've made it look very big here. But of course, the actual difference between the toroidal magnetic field at the edge and the average toroidal magnetic field is very small because the beta is so low. So this is a very difficult measurement to make. But as Hutchinson said in his book, if you do it very carefully, you can get reasonable results from it. So yeah, I think it is a real technique that one can use. Yeah?

STUDENT: [INAUDIBLE] so this A [INAUDIBLE] first where [INAUDIBLE]?

JACK HARE: Yes, the A refers to a point where pressure goes to 0. I'm saying at the edge of the vacuum chamber because it definitely should go to 0 there. But if you want to put A further out, you can also do that, yeah.

STUDENT: Does it matter [INAUDIBLE]?

JACK HARE: Oh, it absolutely matters that it's outside the last closed flux surface. You cannot have any plasma further out than this probe because that plasma could carry current. And that would violate all of the assumptions we've made so far.

STUDENT: So if there's [INAUDIBLE]?

JACK HARE: I don't understand. I would never put a B-dot probe inside [? the ?] [? last-- ?] even anywhere close to the last closed vector. You'd have them very far away because of reasons mentioned earlier, that they would melt.

STUDENT: Right, OK.

JACK HARE: So let me give you an example. If this is your tokamak with a divertor, we've got our core plasma and we've got all of this stuff here. You can just put your B-dots out here, somewhere where the pressure is equal to 0. The pressure is not equal to 0 in the last closed flux service. And it's not equal to 0 in the core.

And that effectively reflects that there is some plasma there. That plasma could carry some toroidal current density. And that would mess with of the magnetic field calculations we've been doing where we've carefully assumed that there is no further plasma carrying current outside. So we certainly couldn't suddenly introduce a blob of plasma out here. That would ruin our day, yeah.

STUDENT: [INAUDIBLE]

JACK HARE: I mean, the pressure of the plasma has to go to 0 somewhere. Otherwise, it has no edge, right? And you can define where that goes, like you say, by sticking a limiter in, like that. And so you can put it anywhere that is behind the limiter, yeah.

STUDENT: OK, cool. Anywhere?

JACK HARE: Anywhere where the plasma pressure is 0, yes. Yeah, I see a question from online.

STUDENT: Hi there. So my question is about the B phi at the edge on both sides. You said at one point that we are like confining them to be the same on either edge. And so it's the dip?

JACK HARE: Yeah, so we made an assumption here that B phi was roughly constant. So all I'm saying-- probably I shouldn't have drawn the diagram exactly like this. In the model that we're doing at the moment, what we have is a plasma like this. And we're saying that the magnetic field up to the edge of our plasma is some value B phi A.

And inside that, it has some dip. We don't know the exact shape of the dip. That's going to be due to the current profile. So the dip could look like this, for example. But I'm just-- some sort of dip. Does that make sense?

STUDENT: Yeah, that makes sense. And then the thermal energy is-- the thermal-- what did you say about thermal? I missed that part. You said there was a thermal energy--

JACK HARE: Well, we'll go through it in a little bit more detail in a moment. But the thermal pressure here is just the thermal energy density. So you can measure the thermal energy inside the plasma now that you know the average thermal energy density.

STUDENT: OK.

JACK HARE: Yeah, we'll get to that in a second, yeah. OK, more questions here? Yeah?

STUDENT: [INAUDIBLE]

JACK HARE: We haven't got there yet. If you'll wait five minutes, I will get there. The big reveal at the end. OK, other questions?

STUDENT: [INAUDIBLE]

JACK HARE: Well, the way we've done this is we have taken such a large aspect ratio that our 1 over R is effectively constant where we've chosen this. This is our unwrapping of our torus into a cylinder. It's not true. And we will talk a little bit about the not truthfulness of it in a moment as well. Any other questions?

Cool, I'm just going to raise these answers to questions. And then I'll keep going. So I want to briefly touch on times when we have paramagnetism instead.

So in the case where our poloidal magnetic field is not necessarily much smaller than our toroidal magnetic field, that doesn't necessarily mean it's bigger than the toroidal magnetic field. It might just be that it's not negligible. If you go and you work through this equation and stare at it for a little bit, it's actually pretty obvious that you can end up with a situation in which B phi is greater than B phi A, where the average B phi is greater than B phi A. So that would be a situation which looks like that. Anyone got any idea what's going on here, why that could happen? Yeah?

STUDENT: [INAUDIBLE]

JACK HARE: That'd be cool. Dynamo is forbidden in axisymmetric systems by Cowling's theorem. So I think we can't have that. But OK, it'd be a good idea. Yeah?

STUDENT: [INAUDIBLE]

JACK HARE: Good. OK, so let's draw that out. So let's start with a big plasma that has got a uniform magnetic field across it. But this plasma also has some current going in this direction.

And that current is generating a poloidal magnetic field, B theta. And this poloidal magnetic field is allowed inside the plasma as well, of course. This configuration gives rise to J cross B forces, which will always pinch inwards. These are pinching forces. This is like two wires attracting that you know about from electromagnetism. So this is your J cross B force.

Now, the magnetic field inside a sufficiently conducting plasma is frozen to it. So as this plasma compresses, the magnetic field inside will also be compressed. So we'll end up with a situation where the plasma is now smaller, the magnetic field on the outside is still pinned, but the magnetic field on the inside can have some pattern which, again, is going to depend on some details of how the current is distributed such that this condition is fulfilled.

And it turns out, if you go through mathematically, whether you get ferromagnetism or diamagnetism depends not on beta phi, but it depends on beta theta. And if you have beta theta less than 1, it's diamagnetic. And if you have beta theta greater than 1, it's paramagnetic.

Don't ask me what happens if it's equal to 1, probably nothing. So cool, and so that situation, can anyone tell me some sort of device where this could happen? So this was relevant for stellarators and tokamaks. Any sorts of devices where this could occur?

STUDENT: [INAUDIBLE]

JACK HARE: Yeah, Z-pinches don't actually have a toroidal magnetic field. So it wouldn't work, but a screw pinch would, a generic screw pinch, Yeah. Something like a reversed field pinch is like the classic example. So reversed field pinches, which look topologically like tokamaks, fulfill this condition here. In fact, they have a region where the toroidal magnetic field goes to 0 and reverses, hence reversed field pinch. And they definitely see this paramagnetic effect. Yeah?

STUDENT: [INAUDIBLE]

JACK HARE: Yeah, I think you probably would, yes. Sorry, for Columbia people, I think you would see this in a [? sparmac, ?] but I don't know. Any questions from Colombia or anyone else in the room? Yeah?

STUDENT: [INAUDIBLE]

JACK HARE: Ah, that's a good question. I guess the way we're coming at it here is what we can measure and so what we can learn about our plasma rather than what we can do to our plasma. So I'm not sure Hutchinson goes into much detail on this. I'm not sure the answer to your question there. Yeah, I don't have a good answer.

STUDENT: I have a question.

JACK HARE: Yes, please.

STUDENT: So how does this relate to the diamagnetic drift? Often in a tokamak, if it's diamagnetic, there will be some poloidal drift associated with it. But does that drift [? switch ?] direction in the paramagnetic case?

JACK HARE: Yeah, that's a great question. So if people didn't hear, how does this relate to diamagnetic drifts that we talk about from single particle motion? The answer is this is equivalent. So you can get these same results by following single particles as they orbit.

And you can also get it from MHD, which is quite nice because you can't always get everything single particle from MHD. But in this case, MHD retains that information. And so think, in the paramagnetic case, you're right. The drifts in the opposite direction would cancel the magnetic field, yeah.

STUDENT: [INAUDIBLE] theta theta that [INAUDIBLE] definition?

JACK HARE: No, it's still this one. So it's at that beta theta, yeah. All right, let's measure the energy confinement time.

So all of this has led up to us being able to measure the volume averaged pressure. And we found that we could do that using B theta at the edge, B phi at the edge, and the average of B phi. The reason we might want to do that is that this average thermal pressure is just equal to the density times the temperature. Using temperature in energy units, we swallow the Boltzmann constant.

And the stored thermal energy is just equal to this average energy density times the volume of our plasma times a factor of 3/2, which comes from equipartition in three dimensions And you've probably seen from ideal gases. And I'm not going to go through it. But just, it's there.

Cool. So the trouble is measuring this is still difficult because-- where's my board with it on-- the thing that we're measuring in order to get this pressure is the ratio of these two things. And it turns out that this dip is going to be very, very small. And so it's going to be very, very hard to measure this.

So I'm going to put that back now. The thing we're trying to measure, all of the errors are going to come from the average value of B phi over B phi at the edge. And this is going to be-- it's going to vary within about 10%. So it's not really the right use of approximately equal sign, so 10% variation.

So this is difficult to measure accurately. But you can still do it. It is hard, but possible. And as I said, you should have a look in the textbook. There's a very nice figure from a tokamak, I think in Australia. And they still had tokamaks where they did this.

And the reason you might want to do that is because you might want to calculate the energy confinement time. So can anyone tell me a useful definition for the energy confinement time? And you can do it in words, if you don't want to do it in formula. Mm-hmm?

STUDENT: [INAUDIBLE]

JACK HARE: Perfect. Thank you very much. So it is defined as the total stored energy W over the input power. And I'm writing that as a big P because it's not the same as our little p. It's the power going in. So as I said, this is simply-- well, I'm not going to write it out again. It's there.

And this, or an ohmically driven, inductively driven tokamak, doesn't include all the other heating terms you might want to do is going to be I phi squared RP where this is the plasma resistivity or plasma resistance. So we're treating the entire plasma, the entire donut here, as some sort of resistor with resistance RP, like that.

OK. And if you crank the handle on all of this, you remember the volume of your torus is pi A squared times 2 pi r, at least for large aspect ratio torus so that we have a simple formula. And you remember that your toroidal make-- your toroidal current from Ampere's law is going to be 2 pi A B theta at the edge over mu 0 squared.

You can write all of this triumphantly as 3/8 mu 0 beta theta R over RP. I personally dislike the way that we use R here because these are completely different R's. It makes it look like it's a ratio of things with the same quantities. But this is the R for your torus, the major radius. And this, again as we said, is the plasma resistive, which is going to be-- let's see-- da, da, da-- it should go up with length.

So 2 pi r dotted with the resistivity over pi A squared-- and that resistive we can just replace with a 1 over the conductivity. Huzzah. So we've measured that already. Remember, we used two loops to measure that. We used a Rogowski and we measured B phi. And then we also need to know this.

And again, we can measure that using three loops, which are B theta A B phi A and average of phi. Three loops-- so the conclusion of all this is 5 loops gets you tau E with a fair few assumptions along the way. But that's not bad. Questions? Yeah?

STUDENT: [INAUDIBLE]

JACK HARE: Yeah, I guess you can also extract rather more noisily the density at this point. So your argument is that from the conductivity we've also got temperature. And because we've got the average pressure, which is presumably going to be the average of n times t, we're going to be able to get out an estimate for n with lots of exciting averaging going on.

But again, for five loops, that's not bad. So you're going to argue that we can also get n. And so we can get the Lawson criterion as well. I'll talk about that. That sounds like fun. Someone should try that. OK. Yeah, I see there's a question online.

STUDENT: Oh, hi. Would you just quickly mind showing where you would place all of these--

JACK HARE: Thank you, that's a great question. I'm just going to erase some of this because I think you guys have seen it. And I want to have this space, so I can draw nice and big what's going on here. So I'm going to draw my tokamak or my toroidal device in cross-section, like this.

So I need to have a Rogowski. I'm going to put that around some poloidal cross-section. I'm going to back-wind it to keep everyone happy. But that's I. I am going to stick in a probe which measures the toroidal magnetic field at the edge. I'm going to stick in another probe.

I hope you can appreciate the beautiful perspective I've put into this circle here. And this is going to be measuring B theta at the edge. So remember, that probe should be lined up with the long axis of the torus here. So if you want me to, I could just draw it like this, I guess. There we go. Cool.

I also want to have something that measures the average magnetic field B phi. Note how this looks a little bit like the Rogowski, but it measures something very, very different. And then finally, I need to have my flux loop, which was measuring B phi like this.

Now, remember what we were saying is we're talking about here an ohmically driven, inductively driven device. And so that transformer, which will fit around the cross-section here, is driving some voltage that drives the voltage in the plasma. But exactly the same voltage is being driven across this loop. And so we can measure the voltage that's in the plasma using this loop here. OK. Hopefully, that is clear. Any more questions?

STUDENT: Yeah, thank you.

JACK HARE: You're welcome. OK, let's do m equals 1 [? after ?] m equals 0.

So as we discussed last week, m equals 1 is a displacement mode. And so what we're doing is we're talking about a cylindrical vacuum vessel here that has inside it some sort of cylindrical plasma. And the center of that plasma is offset from the geometric center of our vacuum vessel by a distance that we're going to call delta, capital delta.

And we'll say that the vacuum vessel has a radius A here. And around the edge of this vacuum vessel, we're going to place a series of B-dot probes. And these B-dot probes are going to be lined up to measure the B theta.

So if we look at what B theta is going to be as a function of theta-- and remember, each of these probes is at a different angle-- theta 1, theta 2, so on. So effectively, we are measuring this function sparsely at each B-dot probe here. We're going to see that this is equal to mu 0 over 2 pi A times the current inside this plasma.

For Ampere's law, we don't care about how the current is distributed. It's just the total current here. And this is going to be modified by a term that looks like sine squared theta plus cos theta minus delta upon A [INAUDIBLE] vacuum brackets squared square rooted.

This is just geometry. This term here is just R, where R is the distance from the center of the plasma to some probe, which is now different depending on your angle that you're going at, which is why I'm just going to clarify it. This now has a theta dependence to it.

And we're going to make an assumption again that delta upon A is small. Is this a good assumption? What would happen if it wasn't true? Yeah?

STUDENT: [INAUDIBLE]

JACK HARE: Plasma's gone. It's hit the wall already, exactly. So we're going to make it. And it's a very good assumption because otherwise you have already failed the [INAUDIBLE]. So having made that assumption, we can approximate this as mu 0 I upon 2 pi A.

And we get out 1 plus delta upon A cos theta, the Taylor expansion here. And you can see that there's going to be a second order term [INAUDIBLE] A that disappears. And there's going to be a sine squared and a cos squared. And they're going to come together to make a 1. So this all looks like it's going to work.

Remember here, in this case, we have specified that delta is in the x direction. So this is delta x, like that. If it was in the y direction, this would be 1 plus delta y upon A sine theta. And if you remember the Fourier decomposition that we said we were going to do at the start, for this magnetic field pattern there are only two terms.

There's the B-- well, this is the C0 coefficient. And these two are the C1 coefficients. And so we can say that we can directly measure the displacement x as 2A C1 over C0 or the displacement in y as 2A S1 over C0. So this is C1. This is S1.

So by doing a Fourier decomposition of our signal and therefore finding the coefficients C1 and S1, we can immediately find the displacement of our plasma. And that's pretty fast to do. And so we can see if our plasma is starting to move. And we can do something about it, like active feedback, to stop it hitting the wall. So that's pretty useful. Any questions on this? Yeah?

STUDENT: [INAUDIBLE] as you're varying theta [INAUDIBLE]?

JACK HARE: Oh, yes.

STUDENT: [INAUDIBLE]

JACK HARE: Yeah, so if your plasma is in the center, you would have-- B theta would be like this. And you might reconstruct a flat profile. If your plasma moves towards theta 1, you might have a profile-- let me use a different color. Yeah, that goes back to there.

And you'll notice that, in this case, I've got four sensors. And I can just about reconstruct the cosinusoidal variation there. And you'll find that very interesting when you come to the P set because this sort of question is fundamental to how many B-dot probes do you need to measure some displacement.

So thinking in terms of Nyquist-Shannon sampling theorem, things like that might get you a long way. Cool. Other questions? Yeah?

STUDENT: [INAUDIBLE] for measuring [INAUDIBLE] 0 [INAUDIBLE]?

JACK HARE: How many did we need to measure? We just needed 1.

STUDENT: But now we're [INAUDIBLE]?

JACK HARE: No, it's still here. You see the C0 term when you do your Fourier decomposition. No matter how many probes you have, you can always Fourier decompose it to C0 upon 2, and then the sum to infinity of Cm cos m theta Sm sine m theta.

And what we're saying is if you want to measure, for example, Te and all that beta stuff we did before, you just take this one. This is the m equals 0 term. If you want to measure displacements of the plasma, you take these first two terms, C1.

If you want to measure more exciting higher order modes where your plasma gets pancaked or it comes into increasingly unlikely shapes, then you want the higher order ones of these. But John made a very good point, which is that you'll have to think about how many-- what m you need to get to will set a requirement on how many probes that you have in order to properly do this Fourier decomposition. And this, again, is related to Nyquist theorem. So there's a bit to think about. Any other questions on this? Yeah?

STUDENT: What [INAUDIBLE]?

JACK HARE: Yeah, so the question was, at what point do you need to start worrying about eddy currents in the vacuum vessel messing up this measurement. These would definitely be affected by eddy currents. So that's a higher order correction that you need to do to understand what the eddy current does, yeah.

So in this assumption, we have assumed there are no eddy currents. We have a nonconducting vacuum vessel or something else. Yeah, exactly. Any other questions?

So just very briefly, Hutchinson spends a long time on this in the book. And I think it's very interesting. But it's very hard to go through in class because it's very mathematical. So I just want to give you a little taster of what would happen for a toroidal system.

So remember, we've really been assuming all along here that we've got a cylinder, a very high aspect ratio tokamak, so that we don't have to worry about toroidal curvature. But what happens in a toroidal system is that we have a very complicated force balance. Remember, our MHD equations, which we simplified massively, we're going to have some more terms.

But there's a complex force balance. And that shifts. If this is our axis of symmetry of our tokamak and this is the vacuum vessel, that shifts the magnetic flux surfaces so that they're no longer concentric. So we may have a flux surface like this, and a flux surface like this, a flux surface like this, and a flux surface like this.

And again, I'm not explaining where you get to from this. This is like three or four pages of algebra. But you start ending up with systems that are much more complicated. We have kind of been assuming a lot of the time here that we have concentric flux surfaces.

It's actually not important for most of the things we're driving. But we certainly would have no reason to believe the flux surfaces are not concentric because our system is azimuthally symmetric. There'd be nothing pushing them from one side to the other.

So if you start doing this analysis properly, you get to equation 2.2.25 in Hutchinson's book, just to let you know where I've pulled this away from. And you find out that your magnetic field, the m equals 1 term, is going to be equal to the m equals 0 term times a correction for where you are inside the plasma.

So this coordinate R measures radially outwards from the center of our poloidal cross-section. And this R0 is the distance from the axis the center of our poloidal cross-section. And we find-- sorry, this isn't very clear. And all of this is going to be multiplied by beta theta plus a term, Li upon 2, which I'll explain in a moment, minus 1. Again, none of this is derived. I'm just throwing this equation at you to show you some of the things that you can do in the toroidal geometry.

Now, beta theta is something that we can measure from our m equals 0 diagnostics. They're still going to work just fine. What's interesting is that this shift here results in a term that is to do with the inductance-- not quite the inductance, but it's similar to the inductance.

And we talked before about how the inductance is a measure of the geometry of your current, so how the current is flowing in the system. And so by measuring this very precisely, you may be able to determine whether your system has a current profile that is nicely peaked or whether it has a profile which has a hollow in the center, something like that.

And so you actually can end up getting more information out of your system when you have a toroidal geometry. But the price you pay for that is significantly more complexity with all the equations. So again, that's all I'm going to go into in terms of this. If you work with magnetic diagnostics for tokamaks, or stellarators, or other devices like that, you will definitely come across some of this very complicated stuff. So any questions?

OK. Now, we're going to stick the probe in the plasma, see what happens. So now, we're discussing internal probes.

So again, we'll have some vacuum chamber with some sort of clever re-entrance port. And we can stick into it our B-dot probe. We'll probably put a bit of shielding insulation on the B-dot probe, like that. And we'll talk a little bit more about that in the context of Langmuir probes next lecture. And there is some plasma here. And there will be some magnetic fields, which may be three dimensional in nature.

Now, despite what it looks like, this is actually not very perturbative. Despite the fact you would stick a probe right in your plasma, this is not as bad as it might seem. What do we mean by perturbative? We mean that it changes the measurement that we're trying to make.

This turns out not to be very perturbative. The main reason is that the magnetic fields at a given point here are not just set by the currents locally. They're set by currents over here, and over here, and over here. So B is set by global currents. And so disturbing just a little bit of the plasma doesn't actually change the magnetic field at a point that much.

Again, we're going to put an insulator around it. So we would set I into the probe equal to 0. And we're going to make the probe small. We want it to be small to reduce the perturbation on the plasma, but also to make delta B over the probe small, as we discussed before, which makes our analysis easier because it means the magnetic field is constant over the probe here.

So you wouldn't want to stick one of these inside a very hot and dense plasma. Well, you might. If went on to stick one of these inside a tokamak because it would just melt and it would fill your plasma with impurities. And your plasma would crash. That would be bad. If you're doing, for example, a pulse power-driven experiment, you might want to stick these inside a plasma. But you might not expect it to last more than one experiment because of the heat flux on it.

Now, what people do on some machines is they have sticks of probes that they can position. The stick has a load of probes wound like this to measure, for example, Bx. It also has a load of probes wound like this in the same location to measure By. And it'll have a load of probes wound like this to measure Bz. So you'll be able to get all three of the axes here.

And the nice thing about that-- you put these all together. You can find out what the current is locally inside your plasma because the current is simply 1 over mu 0 curl of B. Now, you're taking the gradient of the signal. And the signal might be noisy. So you've got to be a little bit careful here because this is going to be a noisy measurement itself.

But the fact is you can still have some go at measuring the current internally to a plasma, which is a very, very hard thing to do. I'll just make a note here that this is noisy because we have delta B over delta x. Delta x is your spacing of your probes here. So the smaller you can make delta x, the better you're going to be measuring this gradient. And this looks very Nyquisty again all of a sudden. Sorry, we can't escape from them.

The nice thing about measuring J is that then we would have J and B. And we can use that to measure the pressure. So we can actually measure the pressure inside our plasma using these closely spaced probes. Or at least we can measure the pressure along the line of the probe.

This turns out to be hard for low beta plasmas because grad P is small. So you won't see much change in J cross B. But when beta is reasonable, you might be able to do this. A very nice trick is if you have a steady plasma or a plasma which is highly reproducible, so you can do this experiment over and over again.

You can actually take this probe and say you've got a plasma that looks like this. And you can step the probe across the plasma. And either the plasma is steady or you're doing the same experiment again and again and again. And you can build up a complete map of B as a function of space and time and therefore J as a function of space and time as well. And people do this.

One great experiment is the large area plasma device at UCLA, LAPD. And on this device where they do astrophysics relevance experiments, they can map out all of this. They will do a shot every second. There'll be a servo stepper motor that moves this across. And they just press go after programming it.

And it will run for 24 hours a day for a week. And you'll have a huge amount of data with a full three-dimensional map of the magnetic fields. So this is the sort of thing you can do with high rep rate devices. It's very, very cool. OK, any questions so far? Yes?

STUDENT: [INAUDIBLE]. But when you're entering [INAUDIBLE]?

JACK HARE: It's the local current. So this is a local equation. This is B of x. So you get J of x wherever your measurement is being taken. So you would get J at the locations of each of these coils or, depending on your differencing scheme, halfway between the locations of each coil, yeah.

STUDENT: [INAUDIBLE]

JACK HARE: Yeah, so the thing is that this equation here does not imply causality. You don't have magnetic fields telling currents where to go or currents going where-- telling magnetic fields where to go. It's just true simultaneously that a global magnetic field also constrains the global current. And so you don't end up having too much perturbation here.

There will be some perturbation. You'll be cooling down the plasma and things like that. But if you're on small length scales compared to the plasma dynamics, the current will still be roughly the same. Yeah?

STUDENT: Can you repeat why [INAUDIBLE] theta [INAUDIBLE]?

JACK HARE: It's just that, in that case, you have P much less than B squared over 2 mu 0, which means that this J cross term is just going to look like this. And so this is x and this is B squared over 2 mu 0. So this is only going to vary by, say, 1% or so.

But you're trying to measure the pressure, which is going to be the opposite of this, down at about a 1% level, which means you need to have 1% accuracy on this reconstruction. So you need to have very low signal to noise because if your signal was noisy, you would just reconstruct a completely noisy pressure signal. And it wouldn't make any sense. Okey doke.

Some other cool tricks you can do with probes-- if you can't move the probe inside the plasma because the plasma is moving too quickly or it's not steady enough, you can just set your probe up and let the plasma flow over it. So for example, we could have a B-dot probe like this. And we can have a plasma that is moving towards the probe with some velocity, v.

And we're assuming here that this plasma is taking with it some magnetic field. And maybe you know which direction the magnetic field is pointing in prior to the experiment. And you set up your B-dot probe so that it's aligned with this field. And if you have a set of these B-dot probes-- for example, if you have one here and then you have one up here and they've got some separation delta x, like this, the signal that you'll measure on the two probes will be delayed.

So if one probe goes like this and gives you this magnetic field as a function of time, well, that one is 2, 1. And the next one goes like this. You can do some rather neat tricks and identify similar features in these two signals-- so the start of the magnetic field, the start of this foot, the peak of the magnetic field here, and the fall of this foot here.

And you can calculate the time differences here. So you'll have a set of time differences like delta T1, delta T2, delta T3, and so on like that. And from this, you can put it all together. And you can get out an estimate for the flow velocity of your plasma as a function of time at a specific point, so at the location of your probe, just from doing very simple time of flight differencing like this.

It turns out this formula is a bit simplified. You actually have to think in a Lagrangian sense rather than this is an Eulerian sense. And one of my grad students, [? Rachaad ?] [? Datta, ?] wrote a nice paper on how to do this properly last year, which is in *RSI.* So if you want to know any more details, [? Datta ?] et al.*RSI* 2022.

Another thing that [? Rachaad ?] then did, which is the final thing I want to talk about today, is he asked what happens if the plasma, which is flowing here, is flowing supersonically, has a Mach number greater than 1. Anyone know what will happen then? Yeah?

STUDENT: [INAUDIBLE]

JACK HARE: Yeah, so you'll get a bow shock forming around your probe. Anyone taking a course in shock physics and know anything interesting about the shape of the bow shock?

STUDENT: [INAUDIBLE]

JACK HARE: Indeed. So the opening angle of this shock, which we call mu, from simple shock theory, we get that the sine of this opening angle is equal to 1 over the Mach number. So if we have a probe and a camera taking a picture of this bow shock and we can measure the opening angle of it, we can get out the Mach number. What's the Mach number equal to? The dimensionless parameter?

STUDENT: [INAUDIBLE]

JACK HARE: Speed over the sound speed. Great. Well, we've already measured the speed using this technique. We've just measured the Mach number using this technique. So now, we have an estimate of the sound speed. And in a plasma, the sound speed is going to be something like gamma average ionization times the electron temperature over the ion mass all [? halved, ?] like this.

Now, some of these coefficients we don't know very well. We should know what our plasma is made out of. We should know the ion mass. It turns out that this gamma for high density plasma where you're going to see shocks is about 1.1. E would be 5/3 in an ideal gas. But it's reduced by the effect of ionization.

And so all of this together gives you a measurement of the ionization times the electron temperature, which is kind of cool because it means that, just by using a single B-dot and a camera, you can get out the electron temperature of your plasma, which is the sort of thing that you normally need to do optical Thomson scattering for.

We're going to spend four lectures on optical Thomson scattering. It's not an easy thing to understand, whereas this is a relatively simple diagnostic and relatively cheap as well. So this is kind of a nice way of being able to measure the temperature of your plasma just by taking pictures of bow shocks and measuring the magnetic field flowing over them. So I see there were some questions.

STUDENT: [INAUDIBLE]

JACK HARE: So the question is, does the heating of the probe cause problems? Definitely, these probes get blown up. And normally, you see the voltage signal go haywire. And then you stop trusting it after that point. So you integrate up until time of death. And then you leave it, yeah.

STUDENT: [INAUDIBLE]

JACK HARE: The thing is that heat transport is actually very slow on these timescales. These are like nanosecond timescale experiments here. So you're not too worried about heat transport. So yeah, ablation is actually like photoionization on the surface and heating through as an ablation [INAUDIBLE] problem, yeah.

STUDENT: [INAUDIBLE]

JACK HARE: So the question is, are you looking at optical wavelengths, are you looking at X-ray wavelengths? Yes, all of those wavelengths will work. In general, the brightness will be to do with the density of the plasma. It has a stronger dependence on density than anything else. So if the plasma is hot enough to be emitting, then you will see the density jump. And you'll see that as a dark region, and then a bright region where the shock is.

So measuring this is uncertain. You want to measure mu as far back as possible, where you've got to a weak shock. But where you've got the weak shock, you have a small density jump, so you can't see it. So what you actually do is you measure up here, and then eyeball it. And you go uh, maybe like that. And then you measure that.

So that contributes a lot to the errors. And then, of course, because you're only measuring the sound speed with temperature inside a square root, when you square the sound speed to get the temperature, that amplifies those errors. So we measured like 22 plus or minus 9 electron volts, which is not a very precise measurement. But it's still better than you can do with [INAUDIBLE]. Yeah, any other questions? Yes?

STUDENT: [INAUDIBLE]

JACK HARE: Z is the average ionization of the plasma. So for hydrogen, that would just be 1. For fully stripped carbon, that would be 6. But in the plasmas I work with, your atom can have some electrons attached to it. So it's still an ion. But it doesn't have all of the electrons removed. And we'll talk a lot more about ionization states later on. But just throwing that in there, if you're a tokamak person you don't like thinking about z, you can just make that 1.

But then you probably won't get any radiation because it's fully stripped. So you probably won't be able to see this very easily. So you need to have something going on, yeah. Any other questions? Anything from Colombia? All right, well, we're at time or past time. Thank you very much. See you next week.