COLIN SHEPPARD: Right, everyone, we're going to start now. So my name is Colin Sheppard, and I'm going to be giving the lecture today. George is here to keep me in order, and I think, probably, he'll come up with some comments at times, maybe-- I hope, anyway.

So he started off by saying that there's no real announcements except that he's changed-- he's uploaded some revised notes, apparently.

GEORGE BARBASTATHIS: There were some minor error in the notes, so I have-- I have fixed them. And the current versions should be-- the one that's in the website now is the corrected ones.

COLIN SHEPPARD: All right, so only minor changes, I think. OK, so I'm going to try and take off from where George left it, which was with Maxwell's equations and the derivation of the wave equation. It's quite nice to see this derivation of the wave equation starting from Maxwell's equations, because it brings everything together and allows you to see how the-- these different areas of physics are interrelated.

But normally, of course, in optics, we don't usually go to the electromagnetic-type theory, so we usually get as far as the wave equation. We show the existence of simple forms like plane waves and spherical waves. We carry on then using those plane waves and spherical waves in a simplified form.

But anyway, so back to Maxwell's equations. These were the Maxwell's equations. George, before, explained these and tried to-- I'm not going to go through the meaning of them, again, because I think you've got that. But in the differential form, you've just got these four equations connecting the electric field and the magnetic field, right?

So you'll see that there's two E's there. Sorry, there's three E's, and then there's three B's. And then there's a charge row and a current that's CJ.

So how we get the wave equation from that is quite simple. It's just a case of doing a bit of vector manipulation. Curl E, you can see, is minus the Bdt from the Maxwell's equations. So if you take the curl of both sides of that, you get this equation. And curl B is another of the Maxwell's equations, so you can substitute in from the Maxwell's equation into this one, and then you'll get another equation.

What you get after that George hasn't actually written here, but basically, what you'd obviously get then, this is curl of this, so it's going to be curl of curl. So this is what you actually end up with, is curl, curl, curl E is going to be-- and then the other term is going to be something-- a second derivative in the E. So actually, that equation, maybe I'll write it down.

It's not actually written there, but after you've just done that substitution, you'll get something like-- is this thing working, or-- yeah, it is. You'd get something like curl of curl E. There we are. Now you can see it. A lot of people right cross product like that, as a little, sort of upside down V. Curl curl E plus mu naught epsilon naught, d2, edt squared equals naught, is I think what you get after you've just substituted that in there.

And then finally-- so this is-- this has still got this thing, curl of curl E, which is quite complicated if you try and work out what that is in spherical polars or something, especially if you're expressing E as something which is spatially varying or whatever. This could be quite a complicated thing to have to work out. But anyway, but the-- you can get it into a slightly more usable form-- often we do this-- by using this identity, curl curl E is equal to grad of div E minus Del squared E.

It calls it here an identity. Really, it's actually the definition of what this thing means. And this is not obvious, what this means, actually. But these other things, OK, they're well defined, what they mean. This thing is actually just what it says in Cartesian coordinates, so you know you've got Del is equal to ddx in the I direction, plus ddy in the J direction, plus dd zed in the K direction.

So Del squared is going to be this dot this, which is going to be d2 Del squared equals d2, dx squared, plus d2, dy squared, plus d2, d zed squared. And you'll notice that that's exactly the same form as the normal Laplacian, right? So the lower Laplacian you've had before, Del squared of a scalar, but this is different, because this is Del squared of a vector.

So what we're saying is that, for Cartesian coordinates, if we define Del squared of a vector by this thing, we find that it's got exactly the same form as Del squared of a scalar. Now, that doesn't actually work with any other coordinate system, so that's a bit of a warning. If you go into sphericals or cylindrical coordinates, you'll find that this Del squared of a vector is not the same as Del squared that you would use for a scalar.

But nevertheless, let's do that. We stick in this as being that. And then we look at this term here, and mostly, we're going to be interested in, for example, regions where there's no charges. And very often, also, we're going to be interested in isotropic media.

If these things are true, you know that div D equals row, so Div d is equal to naught if there's no charges, which means that Div E is equal to naught if there's no charges, and it's also isotropic. So I'm saying this all out in length because you have to be very careful about this, sometimes. Div E equals naught is only true if you've got no charges, and it's an isotropic medium.

So if you've got, for example, let's say a diffraction grating, where you've got some variations in permittivity, then in general, you can't actually assume this, because epsilon is changing. And therefore, with div D equals naught, div epsilon doesn't equal naught. But if you're in free space or some constant isotopic medium, this is true. So this term goes, and so this is just equal to minus Del squared. And so you can see then, this Dell squared then is just replaced by this, and we got a sign change there.

So this is the final expression. And you'll notice then that this is exactly the same form as the normal wave equation that we're very used to in this scalar form, except that the scalar is replaced by a vector. So that's really very nice.

Yeah, so comparing with that wave equation, this is the normal wave equation we've had, the scalar wave equation before, and you can see then they're exactly the same, except that the scalar is replaced by a vector. And the one over c squared, where c is the velocity of the wave, is replaced by this thing, so we can say then that 1 over c squared is equal to mu naught epsilon naught, and that gives us then c is 1 over the square root of mu naught epsilon naught. And so here, there are some figures put in, and we get the expression for the speed of light in vacuum.

So that's very neat. I guess Maxwell must have been really amazed when he did this sum. You don't know, of course, the actual history of the thing, but he played around with these equations, didn't he? And he came up with this form which looked nicely symmetrical to him.

And then he comes out with an expression, comes out with a value, for the speed of light that it predicts, which is what they already knew was true. So he must have felt really as though he was going to get the Nobel Prize or something. But probably, I don't know-- maybe it was before the Nobel Prize, so quite interesting.

And actually, Naveen was telling me that he'd been looking at the original paper by Maxwell. Maxwell's equations, if you look in the original, are horrible, because he doesn't-- this terminology for vectors wasn't invented then. So it's done all in terms of components, horribly complicated.

And I think Naveen said it was the-- Heaviside who actually really derived the Maxwell's equations in the form that we know and love so well nowadays. Heaviside you might remember, he's famous for a few things, but one of them is a thing called the Heaviside function, which is basically a step function. And the other thing is the Heaviside layer, which is an ionic layer in the ionosphere that he reflects radio waves from.

OK, so that's all about free space. So how can we now deal with matter? Well, it turns out actually there's lots of different ways you can deal with-- lots of levels you can deal with propagation of electromagnetic waves in matter. You can, for example, treat it really from a proper atomic point of view. You can think of the material being made up of atoms which have got nuclei and electrons clouds and so on.

Or you can think of it in terms of matter just being like an isotropic material, where you don't really go into the microscopic view of it, but just think of it in a macroscopic way. And actually, to some degree, I think that, very often, the second of those is just as well, because normally, we're not really interested in the actual atomic nature of where these properties come from.

But this is going back to how you can think of it, really, in terms of atoms. So as you know, atoms are made up of a nucleus with an electron cloud. And the nucleus is relatively fixed, of course, because it's heavy, but the electron cloud, because the electrons are much lighter, can be moved around by the field that's supplied. So you get the distortion of the electron cloud.

So this is what this is showing here. It's showing here the nucleus here, which is pretty well fixed, because it's so heavy. And then this electron cloud, you can think of it like being joined to the nucleus by a lot of springs. And then the electrostatic-- the forces on that electric cloud caused by the electric field will tend to, for example, displace the electron cloud relative to the nucleus.

So it gets sort of distorted. And the way that's described is in terms of a property called the polarization then. So this is showing how-- you can see that if you move the negative charges relative to the positive charges, you'll set up a sort of dipole moment. Separate the charges, you've got a plus and a minus. They're separated in distance, so that acts like a dipole. So that acts like a-- gives you a dipole moment, and then if you sum over all those dipole moments, you'll get the total polarization effects of applying that field.

So what that does then is introduces these charge variations, which distort the charge that you really have there, which is caused by the bound charges. So this is what we get. You can say that div D equals row, so div E is equal to row over epsilon naught, assuming the-- yeah, well, this is-- epsilon naught, of course, is really a constant.

And we are breaking this up into it's bound-and-free components, and we're saying that the free part-- we've got something wrong with this again, haven't we? Sorry, that should say-- that should say free there, shouldn't it? Yeah. The free part and the bound part is this part, and so this is it, the final result then. We can now take the epsilon naught the other side, take this the other side, and we get the div of epsilon naught E plus P is equal to the free charges, right?

And then this thing in brackets here is what we call D, the displacement in the medium. So it's like if you've got this medium, div D equals row is what we know is the normal expression that we write from Maxwell's equations. And now, we've derived that that D, in the medium, can be written in terms of the polarization of the molecules that make up the medium.

So how do these things connect? This shows us how we do this then. So we've got D equals epsilon naught E plus P. Now, in general, of course, we don't really know what this does at the moment. And in general, it can be very complicated, actually. It can be a non-linear relationship with the electric field.

But if I go back again to the diagram, you'll see that this is a bit like-- you know, you've got springs here. Hooke's law tells us, in mechanics, if a system obeys Hooke's law, then the extension is proportional to the-- the tension is-- what is it? Extension is proportional to the attention. Sorry, I'm not a mechanical engineer, so I don't know these things.

So you'd expect the same would be true-- could be the same-- could be true here. If you're only applying a very weak field on here, then you might think that this relationship might be linear. You'll get these springs behaving like Hooke's law, so you'll get a linear relationship between P and D.

Now, also, in analogy with the mechanical situation-- yeah?

STUDENT: So if we're in a mat-- like in a material, why is it the permitivity of free space and not just the permittivity of the material? Like why is it epsilon naught and not epsilon?

COLIN SHEPPARD: Why is it epsilon naught here that you can write this? This is either you can think of as epsilon naught-- epsilon E, or you can think of it as epsilon naught E plus P. So these are two alternative ways of looking at the same thing, right?

So we've got D equals epsilon naught E plus P, and you can think of this as being epsilon E. So we're just about to carry onto this on the next slide, actually. So that, you can see, tells us that there is a relationship now between epsilon and P, right? We're going to look what that relationship is.

But in general, it's actually quite complicated, because we don't know the relationship between P and E yet, right? So what I'm trying to establish at the moment is that if this thing-- if the fields are very weak, then you'd expect, maybe, a linear relationship. P will be proportional to E. So this term here will also be something times E, and so you'll be able to take out E as a factor from this.

Now, this is only going to be true if this electric field is weak compared with the-- you know, if the force due to the electric field is weak compared with the forces involved in the binding the electrons to the nucleus. And I suppose it was probably not until the invention of the laser that probably even people thought you could ever possibly get to a case where that might not be true. But now, of course, we know, if you've got a laser, you can make really big fields. You can get, very easily, into this rate regime where this polarization changes.

And so you'll get non-linear-type terms coming. You're not going to do this at all in the course. But there is this whole area, of course, of non-linear optics, which is when this linear approximation breaks down.

So here we are. So this is back to this, what I said, D equals epsilon E. OK, so that's really the answer to your question. And the D equals epsilon E, but D equals this other expression, epsilon naught E plus P. And now, we say that P is proportional to E.

This Chi is called the electrics susceptibility, which I think is a confusing term because it's often, in practice, just called susceptibility. People just neglect the electric bit. But of course, there is also a magnetic susceptibility which is very important when you're doing magnetic materials.

And unfortunately, people use Chi for that too, so that really gets people confused if you're not careful. So we're not going to be doing anything with magnetism, so maybe it's no problem for us. But anyway, this is the electric susceptibility.

So we can put in our P as Chi E, and therefore, D equals epsilon naught 1 plus Chi times E. And this is equal to epsilon E, as we've said. And you can see then, we have that epsilon must equal this thing, epsilon equals epsilon naught 1 plus Chi.

And this is a another expression we've had before, in terms of the refractive index, so you can write the permittivity, the relative permittivity, in terms of refractive index, right? So this is the refractive index coming in here. So under refractive index then is the square root of 1 plus Chi.

And so you see how all these things are connected. You could also say, one more, of course. Where are we? Yeah, we could say that we've got n squared equals 1 plus Chi, so Chi is equal to n squared minus 1. So this is another sort of form you might come across at times.

You can see then, we know in free space-- free space, n is 1, so this thing is 0. So there is no polarization of free space, which is what we know, because there's nothing in it. So there's nothing to polarize. OK?

So, yeah, this bottom bit just says that you can do the same for magnetism. You can do virtually the same sort of expansions for magnetism, but we're not going to deal with this, because we're not going to deal with magnetic materials at all. But there are some quite interesting, you know, magneto optic materials and so on, which are quite important actually, nowadays, in terms of things like optical data storage and so on. We're just going to be assuming that, from the magnetic point of view, the material just behaves exactly the same as free space, so B equals mu naught H. Yeah?

STUDENT: The electric polarization subtracts the field out of the original. Polarization adds the field, magnetization subtracts the field.

COLIN SHEPPARD: Yeah. [INAUDIBLE] has asked this question of why this is a plus, and this is a minus. And the answer is, I think it's all to do with all these horrible things you learn about in magnetism, dire magnetism and paramagnetism. And sometimes, it's whether they oppose the change to which they were due and all that sort of stuff. So I think it's just a case of whether-- if you define magnetization in the way that it's normally done, then it turns out there's a minus sign, but I think you could equally well have defined it as being the opposite side, actually. Yeah. I don't really know much about magnetism, so don't ask me too much about that.

STUDENT: Also, I think it has to do with the part of the stuff that-- B is actually the conjugate of the electric displacement.

COLIN SHEPPARD: Ah.

STUDENT: So it includes the effect of magnetization. So another way to put it is if you could please go back one, if you solve this, if you take M to the other side, then multiply it by mu naught, then you'll get exactly the question as before.

COLIN SHEPPARD: Of course that's right. If you just expand this out, you get B equals mu naught H plus M, so it would be exactly the same form then.

STUDENT: So the field is actually H, not B. The conjugate of the electric field, E, is the magnetic field, H, not B. B is the induction.

COLIN SHEPPARD: Yeah. OK. So now then, let's write down those Maxwell's equations for these various different cases. So this is how you can write them in vacuum. That's how we first came across them. And if you've got vacuum, and in addition, you've got no charges and no currents, then of course, this term, this will be 0, and this J will also be 0. And now, you can see that they look nicely symmetrical. These two are both equal to 0, and these things are both equal to a first-- a time derivative, but there is a difference of a sine-- but apart from that, nicely symmetrical.

And then the middle row here says, if you've got some matter, and it's still allowed to have free charges and currents, then we can write these things. So this is what we just said a minute ago. Of course, if you've got no free charges, then div D is 0.

And then if you've got a material, matter, without the free charges, then again, we put the rows in the J's, 0, and we end up with these quite nice, simple, and symmetrical sort of relationships. And in any of these cases, of course, you could derive the wave equation and write it in that same form we had before, OK? And as we said, this term here is related to the speed of light in the medium, and what we find is that the speed of light in the medium is equal to the speed of light, your free space, divided by the refractive index. All of those things are things we've had before.

So I'll just stress again, one thing that I did say earlier, this was originally, if you remember, curl curl E. And then when we got rid of the curl curl E, we actually had to get rid of a-- what was it-- grad of div E. And we assumed that grad div E is 0. I'm just reminding you, again, that is not always true, so beware. And you probably won't come across anything in this course, but in the general wide world, it's not always true. And the second thing to remember is that this thing, Del squared of a vector, that Del squared is not equal to the same as the Laplacian except if you're in Cartesian coordinates.

OK, so now we've got the wave equation, we go back just like we did for scalar waves. We could look at different solutions of that wave equation, and the simplest possible one is the equation of a plane wave. And so this is a picture of what it looks like.

I presume these pictures are taken from [INAUDIBLE], but to me, it seems very perverse to have the wave propagating in the x direction. He's a bit strange there, but anyway, virtually every book has the wave propagating in the zed direction. But nevermind, we won't worry too much. It obviously doesn't make any difference.

But this is the important thing, is we get the electric field vector and the magnetic field vector, E and B he's drawing here. It could equally well be B and H, because the two-- the B and H are going to be proportional to each other, or at right angles to each other. I'm just telling you the results first. I haven't proved this yet. And the wave is propagating in a direction which is at right angles to both of those.

And so these three form a triad, and it's a right-handed triad such that E cross B is in the direction the wave goes. So that's the way to always remember it. Well, that's the way I remember it, is that E cross B is in the direction the field travels, which means, of course, if you've got a right-hand coordinate system, normally, we would take E in the x direction and B in the y direction, and then the wave moves in the zed direction. But he's taken it with E in the y direction like this, in order to get it so that it's still a right-handed coordinate system.

And so how do we get that? Well, first of all, of course, we can say it's quite straightforward for the E. We can just solve this thing for the-- just as a-- it's a scalar and get an expression for E. E is going to be then polarized in a particular direction.

We've then got to work out how B is polarized for this particular value of E. Well, George has said something here which is-- I guess it's correct, but it's quite sort of condensed, and I went through how I would really derive it. So perhaps, I'll say a bit more about that.

Well, this first part is all right. We've got curl E is minus db dt. And we know that the time dependence is this e to the minus i omega t. So ddt of e to the minus i omega t is very straightforward. That's just-- the minus i omega comes out the front. So that's what we get. Curl e is equal to minus db, dt, is equal to i omega b.

And so we can say then that b is equal to minus i over omega curl e, all right? So if we know what e is, we can work out what b is from that. And so we take e as being this thing, all right?

So in general, this might not be traveling in the x direction or any other particular direction. It might be just pointing in some direction in space. So really, I guess, you ought to do it for that general case, which means that effectively what you've got to do is to work out what curl of the e is. And e is, effectively, as you can see, a scalar quantity.

Well, let's forget about the time dependence. We're only looking at the space now. There's this scalar quantity here. Time is a constant vector. So we then have to use another of our identities, which is the curl of a scalar times a vector.

So let's write it like that. And you might remember that that is equal to-- and I had to look it up because I never know these things-- 5 times curl a plus grad phi across a, all right? So that is a general vector identity which allows you to work out the curl of this product of two things.

And in our case, this is a plane wave. So this direction of the e vector is constant. So our a is a constant. So curl of a is going to be 0 because it's constant. So that goes. And we're then left with the curl of phi times a is grad phi cross a.

And phi is equal to-- phi equals e to the i k dot r. And a is equal to this direction of the e vector, all right? So if phi equals e to the i k dot r, then grad phi is equal to ik times e to the i k dot r. And so we stick that into here.

And after we've done that, we can derive this expression here. We find that finally the magnetic vector is equal to the cross products of these two. So that's why, of course, all these three things have to be right angles to each other, all right?

So the next thing, this is a picture of it moving along then. And we'll stress, very strongly, something. And that is that these two, the electric and magnetic field, are in phase with each other.

I don't know what it is, but I've often found that students don't seem to pick this up. And they come up with these funny ideas about them being in quadrature. And I think they get it mixed up with some other things.

But for a plane polarized wave, the electric field and the magnetic field, you can see, are in phase with each other. When the e is a maximum, the b is a maximum. And then they will have this wave form.

Here this is showing the wave as a function of distance. This is like a snapshot of what happens at a particular time. But you could equally well get a very similar sort of thing by looking at a particular position and seeing how it changes with time.

You'd also get sine waves like this. And you would also see that these two, the electric and magnetic fields, have to be in phase with each other for that case, too. OK? This is expressing-- showing you, again, this right-hand rule or whatever it is for-- I don't know how you do these things.

But anyway, I always like the-- I can never do these, partly because I'm left-handed, I think. It ruins things. But this corkscrew rule is the one that I always know that one. And my students will know it's because I'm very handy with a corkscrew. [LAUGHS]

So this is showing, then, how this electric field varies with space, the fixed time. And the distance, then, between these maxima in the electric field is the wavelength, all right? So this is going to propagate. It's a traveling wave. So this structure moves bodily, doesn't it, in time along this direction of propagation.

Ah, and now we get onto this thing that I was talking with George about on the way, the Poynting vector. And I don't think you actually really sort of prove it, do you, what it really does. But basically the Poynting vector, it turns out, is a measure of the energy flow. So it's a way of looking at the energy that's carried by a wave.

And this is the definition of it. So s is equal to 1 over mu 0, e cross b. Actually, you can see b equals-- actually, I think this should really be mu, shouldn't it, not mu 0, really, if it was in a general medium.

And b equals mu h, of course. So this is really cross h. So that's the way it's sometimes derived, sometimes defined. But anyway, but if you're in free space, then you can write it with mu 0 here like that.

And then, using our expression for the velocity of light, then we can write it like this. But I think this is also true for free space only. I think that's right. Presumably this is from Hecht again.

OK, so this is what our waves look like. Each of electric and magnetic fields is a cosine-type wave. So this represents, then, a wave which is moving in this direction, all right?

So this thing in here basically is the phase, isn't it? You got cos of a phase term. So you can see that as omega t varies, so the shape of the wave is going to sort of move in a particular direction.

And you'll notice that-- oh, this is a vector. This is a vector, but this is a constant vector. So this is a vector which tells you the amplitude of that wave in magnitude and direction.

And you'll notice that these cosine bits are exactly the same. That's because, as we said, the e and the b are in phase with each other. And so all that's very simple.

And so what we can say is that if we represent the fields like this, we can say, according to our definition of the Poynting vector, we can write it like this. And so we end up, then, cos times cos is cos squared.

So notice that, although cos of course can go negative, it's periodic, isn't it? So it's negative half as long as it's positive. Cos squared is always positive.

So the way this is defined, the Poynting vector is always going in the same direction. It's always, when you've got a plane wave, the energy is going in the direction of the direction of propagation of the wave.

And OK, so this shows us how we get all this. Sorry, let me-- this shows us-- can I go back? Aw, yeah. Oh, I've gone back too far now. Here we are.

This shows us how this Poynting vector varies in space and time, all right? So this Poynting vector has also got a sort of wave-type nature. But normally what we're interested in is not how it varies in space and time, but what the time-averaged form of that is going to look like.

And so we can do a time average of it. And so that's what this thing here is. This s with the lines on either side is the time average of s. And so we've got b is k times e over omega.

And so therefore we've got here e cross b. And we know what b is. So we can put that in. And we're going to get now e squared. All right, so finally what we get is that the Poynting vector is proportional to the time average of the electric field squared.

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Sorry?

AUDIENCE: Let's take the average. We haven't done the average.

COLIN SHEPPARD: Oh, we haven't done any average yet. Sorry, but this is the modulus. So--

AUDIENCE: Anyway [INAUDIBLE].

COLIN SHEPPARD: It's the--

AUDIENCE: So right now--

COLIN SHEPPARD: What do the two lines mean?

AUDIENCE: The magnitude of the vector.

COLIN SHEPPARD: OK, it's the magnitude of the vector.

AUDIENCE: A function of time.

COLIN SHEPPARD: It's still a function of time. OK, we'll take it as that, yeah. And yeah, OK. And I think, really, if it's in a medium, this is epsilon, not epsilon 0. Is that right?

AUDIENCE: Yeah.

COLIN SHEPPARD: Yeah. OK, so now we're going to do the time averaging. So this e, e is a time-varying quantity. e varies with both space and time.

And so this is the modular square of it, of the vector. So we put that in. And we get, then, the s is equal to something we cos squared. So it's not much different from what we had before.

OK, and then the next thing we do is to do the time averaging. So this is something which is positive but oscillating, time-varying. And the speed of these variations is very, very fast.

We've said before the frequency of light is around 10 to the 15 Hertz, which is faster than any detector that we have to measure the fields directly, all right? So you can't actually measure the electric field or the power of an optical wave because the frequency's so high.

So a normal detector, like a photo diode or whatever, would be measuring some time average of this because there's no way it can respond to this sort of frequency. So we do the time average. And so this is how we can do the time average, right?

The square things mean time average, the caret signs. And the time average, you integrate this over a long time and divide by the time to get the time average. And so this is called-- well, usually called intensity. I think that "intensity" is now frowned upon by purists. And they come up with new names every now and then.

So "irradiance" is a word that you often read nowadays in the literature. And yeah, so here it's measured in watts per square meter, which is obviously a unit of power density, all right?

There's one other thing that's worth mentioning here. And that is actually that normal detectors actually really don't measure the Poynting vector. They actually usually measure the electric energy density.

And there's a thing called Poynting's theorem, which, if you study electromagnetism, will show that the power flow out of a volume is equal to the rate of change of the stored energy. And the stored energy is made up of electric and magnetic energy. So all these things are related.

And it turns out that for all practical purposes normally it doesn't really matter whether you're measuring the Poynting vector or the electric energy density or whatever. It's going to give the same sort of answer. But there are cases where you have to be careful.

OK, and then finally we've got to do this time average of cos squared. And the time average of cos squared is just a half. You remember how you do that? Cos squared, you go into double angles. And then you get a cosine that cancels out. And you're just left with the constant bit.

And then, finally then, you put that half in. And we've now got an expression for the intensity, defined as being the time average of the Poynting vector as being equal to the amplitude of the electric field squared, multiplied by some constant, which, in probably 99 cases out of 100, we don't even bother to even think what that constant is, let alone thinking of whether there's a half there.

So you'll see loads of books or papers where they just say that the intensity is equal to e squared, which is strictly not true, I guess, but people do say. So how are we going? We got one hour, yeah.

OK, so that's how you can calculate the power flow for a plane wave. I guess the method is not only going to be true for plane waves. But you can apply it for other sorts of waves, too.

But the next thing is you might think back to what we were doing with phases. And there are a few things that are important to realize here, what you can do and what you can't do.

You remember this is what we said from our phases. We said that what we actually measure in the real world is a wave that's like this. It's a cosine dependence in space and time. And we say that we use this complex representative where we say that this is the real part of some complex exponential.

And if you remember, the reason for doing that is that these complex exponentials are much easier to manipulate. It means that you don't have to remember all those horrible cos of a plus b things and so on. So it's much easier to do the algebra using these.

So what you do normally, as you remember, is you do all your algebra with each of the i-somethings. And then right at the very end normally, you find the real part to find what the real field is in real space, as we're used to.

But here, now you've got to be very careful because you saw that the Poynting vector actually was cos squared, all right? So this is a nonlinear. We performed a nonlinear operation.

And if you do nonlinear operations, the same would be true for nonlinear optics, as well, actually. You can't do it in this complex notation directly because you just get the wrong answer. If you squared this, you get an e to the 2i kz minus 5, wouldn't you? Now the real part of that is not what we want.

So what you've really got to do is to work through, I guess-- well, the surest way of doing it is to work through in terms of the cosines and so on rather than to go into e to the i's. So this is saying a bit more about this.

So what we've got, then, is this is showing how you can do it in terms of the cosines. So you can see that, for example, here we've got two fields. And let's say that you see the point is that, although the fields are additive, the power of those two waves is not going to be equal to the sum of the powers of the two waves, all right?

So we can say that the e's would add coherently. And then we can work out the power flow of the sum of these two fields. And you can see of course, expanding this square, you're going to have the power of the two.

But you're also going to get this cross term, the cross product term, which depends-- you can see here-- on the relative phase of the two waves. If they're out of phase by 90 degrees, then that will go to 0. But if they're in phase, this of course is going to be strong.

And so that's one way of doing it. But in terms of phases, then, you can do it in terms of phases. I'm just saying you have to be careful how you do it.

And so you have to add the phases first and then find the modulus square. You can't actually do the squaring and then do the real power, all right?

So you add these two phases together, which gives the total field. And then the intensity is going to be the modulus square of this. And you expand it. And you can see we've got exactly the same answer.

So it doesn't matter which way you do it, as long as you do it one of those ways. What you don't want to do is to try squaring the e-to-the-i things and then taking the real power because you'll get something different then.

And that is the very last slide. And that is dead on time. So, well, whoa, we've got a couple of minutes for some questions. So any questions? Yes, Sharon?

AUDIENCE: [INAUDIBLE]. Hello? I'm quite curious about the Poynting vector. And you mentioned that it's actually a cross term, and it's similar like propagation and direction, like where the energy flows.

But what about the peculiar case in evanescent waves? Does it mean the energy is not propagating there?

COLIN SHEPPARD: Oh, oh, dear. Oh, yes. Evanescent waves, you're getting onto something really controversial--

AUDIENCE: No, I'm thinking--

COLIN SHEPPARD: --here, I think--

AUDIENCE: [LAUGHS]

COLIN SHEPPARD: --actually.

AUDIENCE: Well, where has the energy gone then?

COLIN SHEPPARD: Sorry?

AUDIENCE: I mean, if the Poynting vector is the direction of the energy flow, then it seems like there's no propagation in energy in the evanescent wave then--

COLIN SHEPPARD: Well, it is normally said that there is no energy propagation in an evanescent wave. But you know, in a way, there is. There's sort of transverse propagation of energy. But there's no energy flow across-- if you're doing total internal reflection and you're looking at the steady state, there's no energy flow across the boundary, I think is true to say. Is that true, or-- yeah, this is--

GEORGE BARBASTATHIS: Yes.

COLIN SHEPPARD: --getting a bit controversial, isn't it?

GEORGE BARBASTATHIS: Yes, you can-- if you have the interface-- I don't know if anyone can see, but if you have the interface of the TAR, if you compute the Poynting vector in the vertical direction, those have to be 0. The time average is 0.

COLIN SHEPPARD: Yeah, yeah.

GEORGE BARBASTATHIS: So you don't have energy flow. But you do have energy density.

COLIN SHEPPARD: Yeah, yeah. And the other thing is, you know, the question of if you've got-- and this is your total internal reflection. So here you get this evanescent wave.

Now, of course, all these questions of these evanescent waves and so on, they all rely on assuming infinite plane waves and infinite surfaces and things like that, which of course none of these things can really be strictly true in practice. But if we do this, then we'd expect to get something which is actually traveling in this direction.

And so in some sense, I think that there is some sort of energy flow in the transverse direction. But I'm just having a big argument about this, not for this case, but to do with some other project we're doing at the moment where it's to do with when the h vector is 0.

And so the question is, is there energy flow? Well, I guess there isn't. But I think there's also this question of whether Poynting vector really always means the energy flow.

In fact, I remember having a conversation with Emil Wolf once about this. And, well, he said, let me give you an example. Imagine you have a DC planar electric field in this direction, and you have a DC planar magnetic field in this direction. Is there any energy flow? No. But you can work out what e cross h is.

And you know, so I think that he put this forward as an example which makes you wonder whether it necessarily really always means energy flow. We were talking about this in the lab today because people have come up with these solutions where, for example, energy flow sometimes does this.

So you might get the lines of-- the Poynting vector might do something like this. So here you can see that there has to be some sort of 0 around here because this is going this way, and this is going this way. And here you can actually get these, sort of, closed eddies of Poynting vector that's going round in circles.

Now, whether this really represents something that's physical or not, I don't really know. You know, in water flow, I guess you could get this. There's no reason why you couldn't get eddies of water going in circles.

But you know, so this is what you get. This is what you get if you've just got a very simple example of a lens focusing the light. You get this sort of pattern if I look at this in some detail. You get this sort of behavior in this region here.

And there've been papers that have described this. But what it really means in practice-- well, I guess nothing really means anything unless you can measure it. So the question is, can you come up with a detector that measures Poynting vector? And I don't know the answer to that.

As I said earlier, I think normally detectors measure electric energy density. So people have looked at what happens in this focused region of a lens. And there are various ways you can look at that.

People have used near-field optics, you know, a tapered fiber to look at it. Or you can do some sort of tomography-type experiments with a detector that you rotate and things like that. But to actually measure the Poynting vector, I don't know. Right, any questions from over there?