COLIN SHEPPARD: Right. So you can see it's me again. So George can have a holiday again. But he's not really having a real holiday because he's actually sitting in. And I guess one reason why we decided that he had to do that was because I have to use his computer because he's written these talks in this strange application that doesn't run on my computer.

So now we going to-- today, we're just about getting into interesting stuff now. Seems rather a long time since the beginning of this course before we get into the real wave optics, which I'm sure everyone would agree is the most interesting part. So this lecture is about interferometry. And then the next lecture-- I don't know how much I'm going to get through today, but the next lecture's about diffraction.

So these are two techniques, two phenomena that are obviously very closely related to each other. So in this talk really two types of interferometer-- the Michelson interferometer and the Mach-Zehnder interferometer. So it took me a few minutes to understand this picture. But anyway, it looks a bit messy. But in the end, I-- George drew this, obviously, otherwise I wouldn't be saying that.

But what this is showing, then, is trying to show how a wave varies both in distance and in time. So the wave's propagating along this z-direction. So if you look along this z-axis, you will see this amplitude variation given by the black line. So a cosine-type variation. And then at any point, any value of z, you can see, plotted in red, how the amplitude will change with time along this axis here.

So you can see that you get a sinusoidal variation both in time and in space. And a reminder here, then, how you can express this wave. This is a sinusoidal-shaped wave, which is moving in the positive z-direction. And you can either express it in this form, which is what you would actually measure if you could really measure it-- except that, as we said before, you can't really measure it because it's moving too quickly.

But normally we do the manipulations on the expressions using this complex notation because it makes it easier to do the sums. Yeah. So this is showing, then-- this is a particular value of distance, then, and showing how it changes here the value of the phase of this cosine wave at the time t equals 0 according to the value of z. I guess what that's trying to show.

And then let's think now what happens if we've got some plane wave propagating along at an angle with the axis. So here's this wave coming along, which it's going along a line-- the direction of the propagation of the wave is along this direction, which is an angle theta with the axis. And so you can see that when this wave reaches this plane z equals fixed here-- so this is some plane where we might put a screen or detector or something-- you can see that the different parts of this phase front are going to strike this plane at different times.

Or equivalent, you can say that if you look at this plane here, this wave will experience-- exhibit-- different phase at different points along this line. And so this is what we're looking at here then. As it says here, the path delay increases linearly with x. So there'll be somewhere where this arrives at a phase of 0, which I guess is this point here. And then on one side, the phase is going to be positive and negative, according to which way you go.

So mathematically, this is how we express this. So this is our wave, the same as the previous slide, but now the wave is moving at this angle theta to the axis. So we get this term. This is our-- you remember, we expressed this before as k.r for a wave moving in general in a direction r.

And so this-- so we could write this then. This is what is in terms of our phasor representation. So as before, we write this as an e to the i type thing. And then we also-- this e to the minus i omega t, we just missed that out for clarity, just because we don't want to have to write that every time. So there's an e to the minus i omega t understood that we don't write.

And k, of course, is 2 pi over lambda. So k is related inversely to the wavelength of the light. And so this thing here, this complex number here, you could write this as e to the i phi, like this, where phi is the phase of the wave, which is a function of position. And it can be written like this.

So if we look at a particular value of z, then this z is going to be fixed. So this part, e to the i 2 pi over lambda z cos theta, is going to be fixed. And the phase variation is only going to be because of this other part here.

So this is another example. And so this shows spherical waves or rather distorted spherical waves they become because of the advances in modern technology that love to distort everything you draw and make it so that it's not the right shape anymore. But anyway, this is a point source which is giving out spherical waves. And you can see that here we put our screen again or whatever it is that z equals z0.

So you can see that these waves, as they travel out here, they're traveling a different distance before they get to the screen. So you can see that's going to do two things. One is, of course, the amplitude is going to decay as 1 over r. So the amplitude of the wave when it reaches out here is going to be smaller than the amplitude here.

But much more important is the phase effect because you can see the phase actually might change tremendously. Because the wavelength is so small compared with these distances, it means that you only have to go quite a small distance before you get a big change in the phase. So that's what we want to be able to quantify.

And so here we're riding our spherical wave, just like we did before. So George has kept this the same as he's said before. You'll notice that it's got this funny pi over 2 in there, which he introduced that before. It has the effect of making this like a sine, of course. Something minus pi by 2 is like sine.

This is actually like a sinc type function. And that's the reason why that pi over 2 is there. It's so that it doesn't blow up when you look at what happens at x equals y equals z equals t equals 0. If you didn't have the pi over 2 there, you'd have horrible infinities. So you avoid those by putting this pi over 2 there.

And yeah. So this is the radius. r is just the square root of x squared plus y squared plus z squared. So this is kr minus omega t. And this is our r at the bottom. And now we make an approximation. The approximation we're going to make is that we're not actually really looking at this geometry as shown in the picture, but we're looking at the geometry when, actually, we're only looking in a small region near to the axis.

What that really, in physical practice, that means is that this distance z0 becomes very large. We're a long way away from the source. And so we can think of this value z being a very large.

Sorry. You understand my funny English language, don't you, when I talk about "zeds"? What you call "zees." I just suddenly realized that. Anyway.

So we've already had this before in earlier lectures as well. This is a square root that we can expand by the binomial and assume that z is large much larger than x and y. So you take out a factor z, and the square root becomes 1 plus a half-something. And so eventually you get this sort of thing.

So that's what we do inside this phase term. But the thing at the bottom, we say we don't have to be so careful with that because it's not so sensitive to this amplitude because all it does is just changes the magnitude of the wave by a small amount. Because we've said x and y is small.

But up here, you have to be more careful because we're talking about phases. So anyway, at the bottom, we can just neglect the x and y. So that's where the z has come from there. And when we do this binomial thing on the top, we get this expression here then.

So the dependence on x and y, you can see, just becomes a quadratic type term-- x squared plus y squared over 2z. And so this is what this picture is drawing here. The path delay increases quadratically with x. So as you go further, this point here corresponds to the center of the screen. And then as you go further away, this increases with some parabolic variation like this.

So that's doing it in terms of the cosines. If we go in terms of the complex exponentials, much the same sort of approach. You can just expand, again, the square root inside the exponential now. Exactly the same, really. And you end up getting this expression here.

And so this is an example now, you can see, of the power of this complex exponential form though because this complex exponential of here two, terms or it might be more terms in general, you can just expand them, can't you? You can just say that multiplied together. Whereas if it was cosine, of course, you'd have to remember all these things-- cos of a plus b plus c. Horrible.

I don't know why I'm making funny noises. Maybe I shouldn't keep getting too near the speakers. Is that what it is? Yeah. So with the complex exponentials, we can break it up very simply. And this inside the braces, then, is our phi before, or i phi. So we can say this is e0 e to the i phi, where the phase phi has just got this parabolic term, plus some constant term, which corresponds to this, which corresponds, of course to this distance-- the phase that you get through traveling in that distance there.

So that's how you get a phase delay.

AUDIENCE: I've got a question.

COLIN SHEPPARD: Yeah.

AUDIENCE: What happened to the pi over 2 in the phasor?

COLIN SHEPPARD: What happened to the pi over 2 in the phasor? So we've got good-- it's there, isn't it? And then it's disappeared.

[INTERPOSING VOICES]

COLIN SHEPPARD: Ah, the i outside. There it is. It's that one. Yeah. OK. So it's an i outside. And yeah. That's an interesting-- I think George said something about that before because effectively it makes this wave 90 degrees out of phase. And I like to think of that as being related to some sort of resonance-type phenomenon.

If you excite something that resonates, then it's going to resonate 90 degrees out of phase, isn't it? So if you've got a point source and you excite it with a wave, then the scattered wave is going to be 90 degrees out of phase with the driving force, just because it's like a resonance.

And I guess that's also true in the Huygens' thing that we come on in the next lecture. But yeah. OK. Thanks for pointing that out. I'd forgotten that.

OK. Other than that it's fine, isn't it? I hope. So we're now going on to discuss something about the significance there. Yeah. So the point is, I mentioned that there is this modulus term, intensity-type term. But this is a very weak change. So you can imagine, if you're trying to do an experiment, to actually measure the properties of that wave field by looking at the intensity would be quite tough because we said that these two terms are very much smaller than that one.

So looking at the phase of the signal is a much more sensitive way to look at that wave front. And as it says, the phase can tell us all sorts of things about the wave, like where the wave has come from, how far it's traveled, the materials it's gone through, and generally the optical path, which you remember we had optical path before, integral ndz. So that will change the phase. And so if we can measure the phase, we can find out information about those things.

The problem, though, as we've said before, is that the phase is not something you can actually see directly, because it's so fast. So 10 to the minus-- 10 to the 15 hertz is typical. Very, very fast. Much faster than you can measure with any electronics.

So how do we do this? We do it by some form of interferometry, which is basically mixing it with a reference beam. And then you look at the intensity of the interference between the two waves. So this is how we do it.

So I said that we're going to look at two types of interferometer. The first one is the Michelson interferometer. So this is how you make a Michelson interferometer. You get light from some source and split it into two parts with a beam splitter. And one of those beams is what we call the signal arm. This is what we're going to try and measure.

So here, this shows this light going through some object, and then being reflected by a mirror, and then coming back to a detector. And then the other arm is the reference arm. So this is going to produce a reference signal. So effectively, we're going to be comparing the phases in those two arms.

So the reference beam comes down also. So these add together. So we can add them as phasors. This is for the reference beam. Ar e to the i phi r. For the signal beam we've got As e to the i phi s. We add the amplitude because we're assuming it's coherent. Maybe I ought to stress that a bit, of course. So this is assuming that this is a coherent-- these two beams a coherent with respect to each other.

So if you had light here from a laser, that would be true. And we'll probably come onto cases where it might not be true later on. Once we've got the amplitude, then we just have to find the modulus square to find the intensity And if you remember, there's actually a proportional sign here. There's some funny constants with epsilon a half and things, which we just don't usually worry about.

So these phases tell us the optical path that the light has gone through in these two arms in order to get to the detector. And so what we're measuring is the difference in the optical path in those two arms. And so this shows how that can work.

So this is an example where these two waves arrive in phase with each other. So then they'll add constructively and you'll get a big signal like that. On the other hand, you could imagine that they might not arrive like that. They might arrive in anti-phase like this, in which case you can see the two waves now are going to cancel out. And you're not going to see now any interference.

In practice, of course, it could be anything in between here. All right. So as you change-- you can change the relative phase of this light by changing the position of this mirror or the size or refractive index of this object that's in there, or the wavelength of the light, or whatever. It's going to depend on all of these things.

So example, measuring distance. So we've said that effectively what we're doing is measuring the optical path. So optical path depends on a few things, doesn't it? Optical path depends on the distance, it depends on the wavelength, and it depends on the refractive index.

So these are all things that you can measure using this sort of system. So here it's talking as an example. There's no object in this one anymore. So this is just free space. And so in this case, then, what you can measure is the difference in the optical path in this arm and this arm. And the optical path-- well, the refractive index is now the same in each, the wavelength's the same in each, so you could measure the distance.

So this is an interferometer for measuring distance. So where could you use this? Well, imagine, this might not actually be a flat mirror. It might actually have some structure there, in which case you could actually measure what that structure is, or if it was a mirror, you could actually use this to accurately measure the distance of this mirror. How far away it is from the instrument, so to speak.

So the interesting thing-- another interesting thing-- this is the Michelson interferometer. And the Michelson interferometer, as the name says, was invented by Michelson. Well, I think he probably invented it. He certainly has his name attached to it. But he mainly used it for measuring wavelength. So we'll come back to that after I've gone through the matter bit, actually. I think it's probably best.

Anyway. So what we've said is we add together these two phases, find the modulus square, and then we've got to expand that. So this is a square of a sum of two things. So we're going to get the square of this, the square of this, and then we're going to get some cross terms. And this is what it comes out to be.

The two cross terms, one is the complex conjugate of the other, which means that when you add them together, the final result is real. That has to be true, of course, because this is an expression for intensity, and intensity has to be real. And so we've written it in this nice form, Ar squared plus As squared plus 2 Ar As cos of something. And this thing here, the something, is the difference between the phases of the two beams.

So this next line puts it into a sort of general term, general form. It's of the forms a constant plus a cosine. And then these quantities, i0 m and delta phi, you can also write in terms of these other parameters we had earlier. So as it says, when both paths are clear, then these phases adjust directly proportional to the distance ls and lr. And so the phase difference is just k times the difference in the two paths.

So therefore, if you measure the signal, if you know what the wavelength is, you can measure what the difference in those two distances is. There's a bit of a proviso there, which is quite a big proviso, that you can only do this, of course, modulo 2 pi. So because this is a cosine, it's going to keep repeating. And you won't know which of the cycles you're on.

So it means that interferometry is usually very sensitive. You can measure very, very small distances, easily much smaller than the wavelength. And people sometimes talk about measuring hundredth of an Angstrom and amazing small quantities using interferometry. The LIGO project, I don't know what they're claiming. 10 to the minus 15 meters or something they're trying to measure to detect these gravitational waves.

Anyway. Yeah. So you only measure it modulo 2 pi. So that means that if you're trying to use this method to measure-- for example, I'm looking at the door at the back of the lecture theater. If I'm trying to use this to measure the distance of that door, I can measure the distance to an accuracy, maybe, of something smaller than a micron. But I wouldn't really know whether it had added to that 10 meters or whatever.

So that actually has been what-- a lot of people have put effort into trying to develop interferometric methods that you can get that extra information so that you can actually do what we call unwrapping, because this is like-- they sometimes call this phase wrapping. The phase wraps because of the cosine. And so you don't know which of these lobes you're on.

Yeah. Well, going back to Michelson. So Michelson did this the other way around. He knows this distance and uses this expression to find the wavelength or the k value or whatever. And that's, of course, what you do in an interferometer, isn't it? If you do a Fourier transform interferometer, you're using that as a way of measuring the wavelength of the light.

In the case of Michelson, he looked at various things. He didn't look at lasers because there weren't any lasers then. But he looked at light from discharge tubes and various different types of lamp and so on. And he looked at how the signal would vary as you move this mirror. And it turns out that if you do move the mirror and you know accurately how you're moving the mirror, you can actually work out from that what the spectral distribution of the source is.

That also suggests a way around this wrapping problem. The way around the wrap or a way around the wrapping problem is to use not just one wavelength but many wavelengths. So even two wavelengths is a big help, actually.

Yeah. It's quite interesting that Michelson-- I think I've spoken about this with some of my students before, that Michelson, he wrote quite a lot about measuring wavelengths using that interferometer. And then many, many years later, people came up with this bright idea of using what they call the low coherence interferometer or the white light interferometer to measure distances. But as far as I know, Michelson never suggested this. I've never found it in any of these papers.

Anyway. So this is another application. Here we're going to look at a real object that we put in there and try and find out something about it. So we put in this lump of material, length l, a certain refractive index. And you can see now then that the phase of this light when it gets there is going to be changed by the light traveling through that material.

And have we taken into account, George, that it's gone through twice here, or even on the previous one? There ought to be a 2 in there. There ought to be a 2 in here, oughtn't there? All right. So the phase difference is actually twice this because it has this amount of phase difference on the way in and on the way out.

So I hadn't noticed that before either. But it applies to this next one as well. So this one, the light goes through this object here. And so there's this extra phase delay that comes because of this object, which is this bit here. So again, there ought to be a factor of 2 for all of these terms.

And so there we are, though. We've got this extra optical path and so now we can measure the optical path of that object. Again, modulo 2 pi. So if this is many wavelengths, you might have trouble to know exactly how long it is, of course, because of the wrapping of the phase again. So this modulo 2 pi we can actually measure-- this is the optical path difference. n minus 1 times L.

The minus, of course, is because the other-- the reference beam has gone through air when this has gone through this material. So what we measure is n minus 1 times L. If you knew what n is, we can know what L is. If we knew what L is, we could know what n is. But as it says there, if you don't know either, then you can't find either. All you know is what n minus 1 L is.

Anyone got any ideas about how you might-- what might you do to get both of those if you wanted to know both of those? Any ideas how you could do that?

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Based on what, though? You'd have to have two measurements.

GEORGE BARBASTATHIS: Press the button, please.

AUDIENCE: [INAUDIBLE] Yeah. Someone says it's to use two wavelengths to set up simultaneous equations.

COLIN SHEPPARD: Exactly. Good one. Yeah. Who said it? Someone in the back said it.

AUDIENCE: I heard it somewhere.

COLIN SHEPPARD: You're not getting the full credit for that.

AUDIENCE: Yeah. Yeah.

COLIN SHEPPARD: Anyway, someone got it. So yeah. So if you measured this for two different wavelengths, then you could have a pair of simultaneous equations, and you could solve for both of those. That, of course, assuming that there is no dispersion. If n changes with wavelength, then you are still back to square one, unless you do some more wavelengths I guess. But anyway, yeah. So it just shows you, though, you can do a lot of these experiments very easily.

So now let's sort of-- these are used-- this Michelson interferometer is used a lot in metrology, industrial sort of metrology. And all countries have their-- in the US, it used to be called the National Bureau of Standards, didn't it? But I don't know that that exists anymore.

AUDIENCE: It's called NIST.

COLIN SHEPPARD: NIST. NIST. That's right. It changed. I don't know why it changed its name. But anyway, so all countries have got this. They have this standard measures. They have procedures for calibrating distances and so on. And the idea is that if a company like General Motors is making parts for a car, that those parts will fit as a spare for another car that's made by another factory.

So say you make some shaft that's supposed to be 1 inch in diameter, you want it to be accurately 1 inch diameter. Otherwise, the bit from one place wouldn't fit the bit from the other. There is actually a true story about that. A thing called the Vickers inch. Vickers is an engineering company in the UK.

And this Vickers inch was a big thing that came up during the Second World War, where lots of factories all over UK were converted to make munitions. And somewhere down the line, people realized that Vickers had a different inch from everyone else. So when they made a shell and tried putting it in the barrel of the gun, it didn't quite fit.

So eventually that was sorted out. They worked out what the problem was. And I guess it's that sort of problem that makes it so that the companies now-- obviously the countries keep these standards and so on.

So the sort of thing you might do is measure the height of this. Let's say you put in-- they have these things called gauge blocks. A lot of you are mechanical engineers so you must know about gauge blocks. So you can buy these standard lumps of metal that have got a certain thickness.

And so you put this lump of metal in here and measure the thickness of it using an interferometer. So these standards, of course, they have these what they call primary standards, and then they're handed down. I'm not quite sure how the length works anymore, but it's all defined in terms of the wavelength of something, isn't it?

And then all of the countries around the world, they all have to compare their meter to make sure that they've all got the same meter. And then within each country-- so in the US, NIST would then be measuring these gauge blocks that would be sent to them from other companies to make sure that they've got the right size. And then in the companies, then they would compare those gauge blocks against other gauge blocks and other things. So this is what you use-- one of the things you use this for.

So now, fringe visibility. Yeah. So we've said that you get these fringes. Perhaps I ought to go back and explain that a bit more. Back to this one. So you can see here we've got a constant term plus a sinusoidal-type term or cosinusoidal-type term. And you can see that this term m is called the contrast or fringe visibility.

So if m is 1, then you can see that they are going to be-- for particular values of delta phi, this intensity is going to go to 0. So that's when you get-- you've got very high fringe visibility. And you can see that this visibility depends on this, Ar 2 Ar As over ir squared plus As squared. And you can show quite easily that that is going to reach a maximum. It's going to be equal to 1 when Ar equals As.

So if you ensure that Ar equals As, then you're going to get your visibility of 1. I guess it can't be bigger than 1 because that would mean that the intensity would go negative, which is not possible. So this plots what it is-- what you would see as a function of delta phi and for a particular value of n.

So there's some background, sometimes called the bias. And then you've got some wiggles on top of that. And so this plot is plotted against delta phi. Delta phi could change by a number of different ways, of course. One way it could change is if you actually move the mirror.

Another way it could change is if you tilted the mirror. If you tilted the mirror, then there'd be different delta phis at different transverse positions on the mirror. And if we take Ar equal As, then you've got perfect contrast, m equals 1. You see it goes right down to zero. And that's the high visibility of the fringes that you can get.

If, on the other hand, you find that if you take, let's say, the reference beam to be much stronger than the signal beam, then you can see that m is going to tend to zero. The contrast is going to become very weak. Actually, sometimes people do do that. Contrast is not necessarily the most important thing because if you notice what we had before, you see that-- let's go back to the previous one.

Yeah. Let's go back to this one here. You can see that you've got these three terms. So what we're saying is if you make Ar much bigger than As, then you make the contrast is going to get smaller and smaller because obviously this term is going to become much bigger than this one. But nevertheless, it does actually make this term bigger. So if actually what you're trying to do is measure this signal in some background of noise, it's sometimes a good thing to make Ar bigger than As because as you increase Ar, it increases the strength of this interference term.

And that's the principle that's used in heterodyne interferometry. So not always true that you try and make the contrast as big as possible. Yeah.

GEORGE BARBASTATHIS: There is one more factor of 2 missing here. On the left-hand side bottom, the swing of the fringes should be 2m times i0 times-- yeah. There's a factor of 2 missing. That's quite unusual. Three slides in a row. They're all missing a factor of 2.

COLIN SHEPPARD: And they don't cancel out.

GEORGE BARBASTATHIS: Yeah. I guess.

COLIN SHEPPARD: OK. So that's the Michelson interferometer. So the Michelson interferometer is basically, you can see, a sort of reflection interferometer. And we're now going to go into the Mach-Zehnder interferometer, which is basically a transmission interferometer.

So the principle is very similar. You split the light into two parts. The two paths-- two light beams go through different path, and then you combine them onto a detector. So this works in transmission mode.

Here you're looking through an object. So now you're not going to get that factor of 2, of course, because the light is only going through it once. So the light goes through here. You notice that the way this has been drawn, it's been drawn so that this light is going like this, this light is going like this. So you can see this path and this path are roughly equal.

You might equally-- well, you could have thought of making this by doing it more like-- couldn't you? You could have done it like that. So where now you can see, if we'd have made it like that, this path is obviously longer than that path. So I guess nowadays if you did that experiment with a laser which has got a very long coherence length, something like a frequency double YAG laser, this would work fine.

If you tried doing this with an arc lamp or something like that it wouldn't work because this distance is much longer than the coherence length of the source. So you're all used to the idea that actually you don't see interference fringes all the time. In the normal, regular, everyday life you don't see interference fringes very much. And that's because in everyday life, normally the light is incoherent rather than coherent.

So if you make this system here or the Michelson interferometer we showed before and you make it so that these paths are accurately the same length, then you could get interference even with light from a-- you can get interference even from light from a tungsten lamp bulb. And they produce really nice, beautiful fringes with different colors. Maybe my students maybe have seen these in the lab.

But anyway. So here you can use this to measure the optical path of some object. And yeah. So this is now, again, looking at this idea of tilting something. I mentioned about tilting with regard to the Michelson interferometer. I said, if you tilt one of those mirrors, you would actually get the optical path varying linearly with transverse distance.

So if you actually looked at an image of that interference, what you would see would be a series of fringes. And the spacing between the fringes would increase as you increase the angle. So this is the same here. We're getting the offset between these two paths-- these two beams-- by rotating this mirror slightly. And then as you can see, these two beams arrive on a CCD camera, it shows here.

So you will actually-- in this case, you would get some variation in intensity on this detector, which you could actually measure with the detector. So it would look something like this. This would now be, then, a function of x rather than just being a function of the optical path difference. So this bit here, I guess, is explaining, is deriving that.

So we're adding two wave together. One is arriving normally. The other one is tilted through an angle theta. And so this one here is going to have a linear phase variation as a function of x, which will then appear inside this cosine.

So finally what we get here then, the value of this bit here, the z cos theta, is just this constant here. So all that does is it shifts these fringes this way. But normally, if you were doing this experiment, one thing you could do then very simply is to measure what this spatial wavelength of these fringes is.

And from that you could, for example, if you knew all the other things, you could measure what this rotation theta was, couldn't you? If you knew what the rotation theta is, then I don't know what you could measure. But you could measure something, probably.

So the Mach-Zehnder, I mentioned that this is like a transmission interferometer. The Michelson is like a reflection interferometer. Both of these sorts of interferometers are used in microscopy, actually. And so you can make an interference microscope based on this technique. And you could put in here, let's say, a biological slide with a bit of tissue or something on it. And you could use this to measure the obstacle path in going through that object.

It's not trivial to do that. You'd normally, if you want to do it in a microscope, you'd have to place some microscope objectives and things in here. And usually it's found that if you do that, that messes everything up. So you normally have to put the microscope objectives in this one as well. So you end up really having to have two microscopes, and the whole thing becomes very expensive.

So that sort of method has really gone out of fashion, although many years ago Zeiss and people used to make instruments like that. The same goes for a Michelson type interferometer. You can measure that for looking-- use that in a microscope for looking at surfaces-- surface height changes, and so on.

OK. So this shows, again, how this comes about, then. You've got this one wave, which is coming along this angle here. And so this represents the phases. This is a movie. Do I just click it again and it'll do something, will it? Yes. There we are. OK.

So it didn't go very far. Anyway. Probably it was quick enough to see. It was a wave that's moving this way. But the important thing, if I go back and do it again-- OK. So you've all got to watch this time. And there's two things I want you to watch. One is the wave moving this way. But more importantly, look what happens on this line here where you've got the camera placed.

GEORGE BARBASTATHIS: [INAUDIBLE]

COLIN SHEPPARD: Right. Yeah. Yeah. Yeah. Otherwise, I think people have to be very quick. Yes. Look at that. OK. So good. So this wave is moving out here. But you can see, if you look at where the detector is, it's as though the fringes-- the phase is actually moving continuously in one direction like that. So right.

GEORGE BARBASTATHIS: So the point is that the reference has a constant phase because it is normal on the vertical axis. So as the incident wave is passing through, the peaks are the locations where the intensity accumulates. So you end up getting a peak, whereas in the nulls, it means that as the wave is passing through, the phases are not lining up. So you end up getting a null. That's how you--

COLIN SHEPPARD: Yeah. Yeah. OK. Yeah. So I guess the reference beam, as you say, is coming along this way. But of course, that's got fringes as well. So--

GEORGE BARBASTATHIS: That's right.

COLIN SHEPPARD: --the peaks of that, when they match up with the peaks of the other, would give the maxima in the--

GEORGE BARBASTATHIS: There's locations where the peaks always arrive at the same time. But a little bit below, the peaks will be actually out of phase. So they end up getting a null.

COLIN SHEPPARD: End of show. So that's the end of that. That sounds like a good place for a break. So let's make a break. Unless anyone's got some questions, of course. Yeah, of course. [INAUDIBLE] always got questions.

AUDIENCE: In Michelson interferometer, the object would affect the intensity according to whatever phase delay it introduces. But in the Mach-Zehnder interferometer, it will affect the spacing of the fringes. Is that correct? It will change the spacing-- we added this constant slope in the signal beam in the Mach-Zehnder.

COLIN SHEPPARD: I think this tilting-- the tilting the one beam relative to another you can doing either sort of interferometer. That's not a fundamental difference between the two types.

AUDIENCE: Oh. So I was thinking, when you have a tilted plane, tilted mirror, then you would measure the spacing between fringes to find out the object's optical path length. But if you didn't have that, then you would just measure the change in intensity to get the wrapped phase and then apply unwrapping. Is that correct?

COLIN SHEPPARD: Well, it depends what you're trying to do.

AUDIENCE: To find out the phase.

COLIN SHEPPARD: Yeah. Yeah. Yeah. You're thinking of it varying. It's a spatially-varying quantity that you're trying to measure, isn't it?

AUDIENCE: Yeah. Yeah.

COLIN SHEPPARD: So I mean, very often what you do, of course, is you actually-- very often you want the fringes because they give you something to actually look at. So normally, if you don't do the tilting thing, what you'd see if the two paths were equal would be, let's say, completely featureless, bright, wouldn't it?

AUDIENCE: Yeah.

COLIN SHEPPARD: And so let's say you put some object in there which has got some phase change, you would just see a small change from that bright, which might not be very visible.

AUDIENCE: I see. Yeah.

COLIN SHEPPARD: But for example, then, if you tilted the system, then you would get some fringes. And the sort of thing you might see, of course-- let's say you put in a disk of material like this and you're trying to measure the thickness of this disk of material-- so what you would find is if it was nicely uniformly thick, you'd find that you've got these regular sort of fringes across here, whereas if it wasn't, then you might find that you get fringes that have got some shape that looks something like this.

So by looking at the geometry of those fringes, you can quite easily measure the changes and the deformations in the shape of it. And as you change the tilt, it would change the separation between these fringes. And so if you make the tilt very small, then you'd just effectively be looking at just a single fringe or no fringes, really. So then it would be actually very difficult to see the changes in the thickness unless you actually made very sensitive measurements.

OK. So everyone back again. Yeah. I thought of something else that perhaps I could mention following on from this diagram. So in the lectures, the system, as you remember, was arranged to do this sort of thing. So the idea here is that the two paths now-- this path and this path-- are roughly equal in size.

In the old days, when people were trying to do this with non-coherent sources, then you had to be really careful. And for example, this is a beam splitter, which is actually, of course, a lump of glass maybe with some coating on it on one side. But you can see this beam has gone through this glass but this beam hasn't gone through that glass. And so these are all things that you also have to take into account.

And so in order to correct for the optical effects of, say, for instance, that slab of glass there, you might actually end up putting a slab of glass in here-- some sort of compensating piece of glass-- so that this would produce the same dispersion in the two arms. So there was a lot of effort put into trying to design these systems so that they would work with the broadband sources and so on.

Nowadays, I guess, most of the time we're doing it with lasers and we don't have to worry about any of these things, but if you look in any of the standard books on interferometry, there are a number of textbooks on interferometry that will give you all the information if you want to go more into it.

So now we're going to change step slightly from interference to diffraction. And so what's the difference between interference and diffraction? The answer is really not a lot, except that normally interference we're talking about two beams interfering, whereas in diffraction we're normally talking about lots of beams interfering. So you can think of the two as really being effectively the same phenomenon, actually.

And so we're going to look at Huygens' principle. We're going to go on to the Young's interferometer. So this really does show, doesn't it, how interferometry is related to diffraction because this lecture is really about diffraction. And then we're going to go to a more complicated case of Fresnel diffraction.

So first of all Huygens' principle. So I think at the beginning of the lecture course, George said a bit about the history of how Huygens' came up with this idea quite a long time ago, wasn't it? It was even before Newton or around-- I think slightly before but around the same time as Newton. So when people still hadn't really decided whether light was a wave or a particle. And of course, we still haven't decided. So times don't change, really.

But anyway so probably now, looking back, we could say that Newton was the first quantum opticist. But anyway. So this was this was Huygens' idea. He was saying that if you think of a light beam-- a wave propagating in space, then at some moment in time this is a wavefront. The wavefront is a surface of constant phase. And he came up with this model for how light propagates where he assumed that all the points on this wavefront can be thought of as being a source of what he called secondary wavelets.

So each of these points radiates like a spherical wave. And then if you sum those all up and calculate what the effect is at some distance in front here, then that would then predict what that wavefront-- what the wave had done at a slightly later time. So that's the principle.

And so quite a simple idea, really. And you can come up with a very simple mathematical description of it. Unfortunately, it doesn't work marvelously well. And people put a lot of effort into trying to improve the theory. And they came up with some rather arbitrary-- they seem rather arbitrary-- factors that they had to include in order to get it to work properly.

So I think nowadays we probably look back over all this and think it was just history. And there are better ways of really tackling the problem. But anyway, at the very simple level though, it's still a nice simple theory for how you do that. So if you want to calculate what the field is at this point here, say, you look at all these waves that come from all these points on the wavefront, add them all together as phases, of course, and that will give you the field at that point.

And there it's showing how you do that. And then you can do it again and again and again and show how the wave propagates. So example, a hole in an opaque screen. So I guess this is pretty well the simplest example you could possibly think of. Well, no, I suppose the opaque screen with no hole in would be even simpler. Then you wouldn't see anything.

But what we do is we just have a very small hole in this. And so this is going to act like a source that's going to radiate like a spherical wave. And what we'd see on this screen that we place here, you'd see that is virtually what we've been looking at before anyway. We've done that before. So there we are. There's the spherical wave. And then we can-- are we going to get some more on that? Did something happen then, George?

GEORGE BARBASTATHIS: Yes. [INAUDIBLE]

COLIN SHEPPARD: Ah! e to the ir. Right. Sorry. I couldn't see anything happening. So this is the plane wave coming in. And now something's come up-- the hole is a delta function. Did that just come up?

GEORGE BARBASTATHIS: Yep.

COLIN SHEPPARD: And now something else has come up. This is the spherical wave now. And this is the same expression we had for the spherical wave in the previous lecture, then. So you remember this parabolic type approximation we made for a spherical wave. This is called the Fresnel approximation, named after the guy Fresnel who I guess must have come up with this in the first place.

And you remember it's true-- it's going to be approximately true if the x and y are small compared with the distance l. So you need that in order to apply the binomial. Right. So there it is. I'll just mention again this i that someone from MIT asked about earlier. So that i then at the bottom is saying that this spherical wave is actually 90 degrees out of phase with the wave that arrives at this point. So that's quite interesting. So you can think of it being like a resonance sort of effect, I think.

GEORGE BARBASTATHIS: I guess this is something that Huygens' principle cannot really explain, right? This is why you have to do more than Huygens'. But we take it for granted in the class, so we don't get into it.

COLIN SHEPPARD: Yeah. OK. Yeah. Yeah. So I think George is really saying that Huygens' didn't actually have that i there. So we're cleverer than he is.

GEORGE BARBASTATHIS: We don't know why.

COLIN SHEPPARD: So there we are then. And oh. Now we've got two holes. I knew something had happened. And the spherical wave is now hidden behind here partly.

GEORGE BARBASTATHIS: More will come.

COLIN SHEPPARD: More will come. So now we've got two holes. Each of these is a spherical wave. And of course, when we get two of them, there's going to be two spherical waves, which I think is going to-- oh, yes. Another spherical wave. Ah, and another equation. There we are. So these are the equations of what you see on the screen for each of these two spherical waves.

And if you look at them, you'll notice that these are distance x0 and minus x0 relative to the axis. And you can see the only difference in these expressions is this sign, plus or minus there. So we now want to add those together. And so here we're looking at some observation point. You've got to add those things together.

Yeah. So of course, as it explains here, if this distance-- we look at this distance and this distance, if the path difference between those is an exact number of wavelengths, then they're going to add up in phase. If they're exact number of half wavelengths, then you're going to get destructive interference. And you're going to get a black region.

So there they are. There's the distances d1 and d2. This is the expressions for those. And again, we can use this Fresnel approximation to expand the square root. And then d2 minus d1 is going to be this minus this. And you see the only difference-- there's a sine there, George. And so yeah. So there's going to be 4x0 x dashed over 2l. So one of the twos cancels and you get this.

So the optical path length, optical path difference, is this. 2x0 x dashed over l. So x0 is this distance here. x dashed is this distance here. So both of those are small quantities. And l, of course, is a big quantity. So this is actually, then, in terms of distance, of course. You know two distances up the top and a distance at the bottom. This has got the right dimensions of length for optical path.

Right. So as we said, that should give constructive and destructive interference. What you're going to see is just fringes like this. So there's going to be positions where these two add up in phase and positions where you get destructive interference.

So if the optical path difference is equal to lambda, then if you put lambda equal to this thing, then you get that x dashed is lambda l over 2x0. So this then tells you-- if you made this n lambda, of course, you'd have an n in here. But this distance here to-- this distance to there is given by this lambda l over 2x0. So you can see the smaller x0 is, the bigger x dashed is. The smaller you make this, the bigger you make this. What happened?

GEORGE BARBASTATHIS: [INAUDIBLE]

COLIN SHEPPARD: Ah. That's facing there. So this the spacing of the fringes, then. It's this lambda l over 2x0. OK. And so this is going through the maths of that. So these are the two waves we've added together. The only difference between them is this sign.

So we can take out the other things as factors. And so these are the two terms with the sign difference. So minus i something plus e to the i something gives cosine. And so that is our final answer. What you see, the intensity is the modulus square of this. These are all phase terms so we don't worry about them.

And I guess in this expression here we don't worry about any of these other things either, these constants. The only thing that's important is this term here. And we seem to have kept the 4. 2 squared. I guess-- yeah. Why have we kept that? Anyway. 4 cos squared something.

And then we say that cos squared of something is equal to 1 plus cos 2 theta. And so therefore our final result, then, this is then our expression. Very similar to what we had before, of course, for the interferometry. A background with a cosine fringes on top of it.

So here it is. It's very similar to what we had before. The m here is 1. And that of course is just like with the case of the interferometer that we described before. You remember we got m equals 1 when the strength of the object beam and the reference beam were equal. So in this case they're equal, aren't they, because the two slits in the screen are equal. If you made those two slightly different widths or something, then that would not be true. Then you might end up with something which didn't have this visibility of 1.

So that's the first example. So that one really was showing how interferometry and diffraction are really the same. Now we're going to go on to some more about diffraction and introduce the idea of a transparency.

So all we're doing here is we're introducing the concept of a transparency that can change the magnitude and phase of a wave by transmitting through it. And so this shows here our transparency. And you can see here it shows that the-- I don't know what this color, whether it means--

GEORGE BARBASTATHIS: [INAUDIBLE]

COLIN SHEPPARD: So this means the absorption of the screen. But it might, I guess, also change the phrase. It might have some refractive index as well. Yeah. And it changes in thickness as well. So you have to work out, just like before, the optical path as the light goes through this structure. And it might change both the amplitude and the phase of the light.

So we can think of this thin transparency as just having a multiplying effect if it's thin enough. So you can use your Huygens' type principle to say that when the light has gone through there, it produces some wavefront on the far side. So the assumptions-- the thickness I've just said. Also, the features on the transparency-- what's that mean? What's that mean? The features on the transparency are smaller than lambda.

GEORGE BARBASTATHIS: [INAUDIBLE]

COLIN SHEPPARD: Ah, right. OK. That's why I was confused. So this transparency can change the attenuation. So the black bits are going to absorb the light. But it could also change the phase. So phase delay. And both of these are going to be-- both of these effects are going to be dependent on the material properties of this transparency, of course, and also its thickness.

And as it points out here, in some cases, you might be able to think of these quantities-- the attenuation and the phase delay-- as being binary quantities. They might only have two different values. For example, you might have black and white, or they might be gray scale-type continuously varying-type quantities.

And then the phase also you can think of that-- maybe it's continuously varying phase change, or it might have just two values, binary values, or it might even have multiple values maybe, say eight values. You can code-- people often use transparencies like this to code binary numbers up to some basis of eight or something like that.

So if these assumptions are true, then according to Huygens' principle the propagation through this, all it's going to do basically is change the-- this is going to change to this. This is going to change to this. And you're not going to get this affecting this very much if this assumption of this thickness being very thin is valid, in which case all it does-- all this transparency does-- is multiplies the field going in by the transmission of the grating to give the amplitude of the wave as it comes out.

So g minus is the wave-- is the amplitude on the way in. g plus is the amplitude of the wave on the way out. And this transmittances, we call this the amplitude transmittance of the grating. And you see it's a function of position. And as we just described, it can be broken up into an absorption term and a phase term. And as we just said, those two might either be-- they could be binary or they could be continuously variable.

Yes. So this is saying how-- you can think of it as being very similar to the Young's interferometer, the two slit expert experiment. But now you can see here we're thinking of it as being many beams. So you're breaking this up into not just two beams but many beams. And each of these is going to propagate. But I'm not quite sure what this is trying to show.

Yeah. OK. So it's pretty well what I said before. So if this thing is thin enough, then this point here is only going to be affected by this one and not by the others. So then you will get this simple multiplicative sort of relationship. So you can think of diffraction then as being like interference with many, many beams.

Yeah. So this is saying, just like we said before, we are describing this interference of these many, many beams now. When we did the Young's experiment, each of the slits was described by a delta function at a particular position. And now we're saying that in general, we can think of this as not being just two slits but an integral over many, many apertures-- very small apertures-- placed at these points x1 y1. Sorry, placed at the points xy, I guess. And then we're integrating over all those points in order to get the total effect of all those spherical waves.

This is now we're saying that g-- so this is g minus after the-- it's just saying that g--

GEORGE BARBASTATHIS: [INAUDIBLE]

COLIN SHEPPARD: This is just the shifting theorem. And now we're going to state that in there, are we?

GEORGE BARBASTATHIS: [INAUDIBLE]

COLIN SHEPPARD: Yeah.

GEORGE BARBASTATHIS: [INAUDIBLE]

COLIN SHEPPARD: Yeah. Yeah. OK. Yeah. Yeah. Yeah. So this is what goes in then. So this is our single point on the wavefront going in. And this gives rise to a wave that comes out, a spherical wave that comes out, which is of this form. It's the same as this multiplied by the transmission of the source.

And then this then radiates like a spherical wave. And then you're going to get some distribution of amplitude on this observation plane, which is going to be exactly the same as we had before. And we've got this-- again, we've got this parabolic phase function that comes from the Fresnel approximation. So it assumes that x dashed y dashed are much smaller than z.

So that's for a single point. And now we do-- now we integrate over all of these. And the final result for the amplitude we'll see on the screen is given by this expression then. So this is the incident wave, this is the transmission of the screen, and this is the bit that describes the propagation of the spherical waves according to the Fresnel approximation.

So yeah. So this result is known as the Fresnel diffraction integral, or simply the Fresnel integral. And so I mentioned earlier that this approximation, replacing the square root by the squared terms, is called the Fresnel approximation. So all of these-- Fresnel gets the name of all of these things.

Yeah. This is just pointing out that this expression we've just written down, the form of this is exactly the form of a convolution. And so that's a very important result. A convolution, of course, means it's space invariant, doesn't it? It means that you can write this-- it's what goes in convolved with something which is independent of the coordinates.

And so if you think of this as a convolution, the thing you have to convolve with is this thing. So you can think of this as being like-- well, it's called the point spread function or the amplitude point spread function. So if you put in some amplitude and convolve it with this, then you'll get something out.

So this is explaining what I just said. So this is the sort of behavior that you think of when you're looking at a mechanical or electrical systems. You can think of a system of being like a black box, or in this case, a blue box. And something goes in and something comes out. And we don't have to worry too much about what's in here.

But you can express what goes in, of course, either in terms of some sort of impulse response function, or as we'll come to in later lectures, as some sort of transfer function-type behavior. But at the moment, we're doing in the spatial domain. So if we convolve this with this point spread function, then we'll get this output, like this.

Now we go on to do some examples of this. And this is probably the most important example-- comes up all the time-- because lots of optical components are circular. So lenses and so on, very often they've got this circular form. And so this comes up time and time again. And so all we've got to do now then is we put in-- we say that our amplitude after the light's gone through this circular aperture is 1 inside the circle and 0 outside the circle.

And so that Fresnel diffraction integral that we have to work out for this circular aperture is just 1 inside and 0 outside. So it's just-- you just have to do the integral over the inside. So this is just over the values of x and y that are inside the bright part here.

Now, so unfortunately, that integral is a bit hard to do. So we're not going to do it at the moment in general. But we will do a very simple example, a special case of what happens along the axis. So if we put x and y both zero, now we get-- we put x dashed equals 0, y dashed equals 0. The integral comes down to this. And now that is an integral we can do. And it's an integral that probably a lot of you have come across before.

The way to solve that is to go into polar coordinates. And so if you change into polar coordinates, you've got a theta and a phi-- a theta and a rho part. The theta part is just an integral of a constant, so that just becomes 2 pi. And the integral of this other thing you can now do. Integral of rho e to the i rho squared over something you can do because the rho outside is the differential of the bit inside.

So to actually do that properly, you can make a substitution then. Make a substitution that xi is equal to this thing inside here. And after you've made that substitution, you now just get integral of e to the i xi. So that's something you can just do. And you get an answer.

Yeah. Notice that you get-- so this final result here, this term here is just the phase term as a function of z. So that's just what you would get if it was just a wave moving in free space. And then these other two terms, notice, they're exactly the same form. This is an e to the i something, and this is a sine of exactly the same thing.

So that's what we get. As I say, if you go back to this, if you wanted to do this integral here, it's a bit more complicated. And so we're not going to do that at the moment. But I guess eventually we have to do something about it.

GEORGE BARBASTATHIS: [INAUDIBLE]

COLIN SHEPPARD: Right. So if we look at this amplitude along the axis, then, you can see it's got this sinusoidal sort of behavior. If you looked at the intensity, this would just go, and you just get sine squared of something. Sine squared as a function of a constant over z. So as it points out here, then, it's going to have peaks and nulls, maxima and minima, corresponding to these particular values of z. So you'll get a sinusoidal variation in amplitude along the axis.

The other thing to note is that-- where's it gone? Yeah. Notice that this z and this z canceled out. This z here appeared when we made this substitution. And so these z's finally cancel. And so the final result has got no z in it. Sorry, no z in the modulus part term.

So you've got this series of fringes which are of constant amplitude however far away from this aperture you go, which might seem rather surprising. You might think that as you got closer to the aperture, it would get brighter, might you? But it doesn't, it seems. I'm just trying to think quickly why that should be the case. But I'm sure it is true.

GEORGE BARBASTATHIS: Notice also the pi over 2 phase shift has disappeared. It got eaten up by the sine.

COLIN SHEPPARD: Yeah. That's true. Sorry. I didn't mention that. Yeah. OK. So when you put in the limits of the integration, then you've got this minus this. And this we replaced by sine. But of course we need to get an extra i to get the sine. Yeah. So the pi over 2 phase shift is also gone.

The other thing that perhaps is worth mentioning also is that when we got down to this integral of course, this is just the Fourier transform of 1, isn't it, within this region. It's the Fourier transform of a rect function. And even if we looked at more complicated example, if you add this circular symmetry, whatever this amplitude transmittance of this aperture had been, all we'd have to do is to do the Fourier transform of that amplitude transmittances of that aperture as a function of this coordinate xi. So that's a sort of trick you can use to calculate these things easily sometimes.

I think we're here for the movies now. So MATLAB has calculated these things for us. And when I press the button, it will go. But I'm not going to press it until you're ready, just in case it's very quick. So this is a big aperture. So what you'll see is how the wave propagates from that aperture. This is actually quite a big aperture. So you don't expect to see a lot of diffraction in this case.

When we come to this one, this is a small aperture. And as we said before, the smaller it is to start, the more it's going to spread. So you're going to see a lot more effect on that one. And so you have to look quite closely to see what happens as it propagates.

So this is a radius of 40 lambda. Now you can see some sort of fringes appearing, and it becomes like a bright halo around the edges. And if you look very closely, you'll see quite a lot of fine structure on this. And then the small one, this is 10 wavelengths. So this one changes a lot quicker. You see it's doing quite drastic things. By the time it gets to the end, you've even got a black spot in the center.

The fact that, of course-- we said before that the smaller the aperture, the more it spreads, but also the quicker it spreads. So you can see that actually this one has got a lot further in the development of the pattern than this one. This one still looks roughly similar to what it started as. This one is looking very, very different. It calls this the blinking spot. I've never heard of it called the blinking spot. But anyway.

GEORGE BARBASTATHIS: [INAUDIBLE]

COLIN SHEPPARD: Well. Yeah. I'm not even sure-- I always think of the Poisson spot as being what you get from an opaque object, rather than for-- but maybe that's the Arago spot.

GEORGE BARBASTATHIS: But Binet says the two are the same, right?

COLIN SHEPPARD: Yeah. Yeah. Yeah. OK. So now the rectangular aperture. The rectangular aperture-- actually it turns out you can't get a simple expression for what happens along the axis of the rectangular case, unlike the circular case. So that's probably why George hasn't done it.

And so you can get expressions in terms of Fresnel integrals for everywhere, in fact, for this case. So the square case, then. Let's see what that does. The big square. Again, you can see-- you can see some sort of fringes. And by the time it's got to the end here, you can see bright lines around the edges, and also bright spots in the four corners.

And then if we look at the small square, again, it's changing more quickly. And you'll notice that we've developed by then-- we've got quite strong sort of streaking along these two axes here. And in fact, you see this final pattern there that is the diffraction pattern of the square aperture. What does it say? Well, we haven't done Fraunhofer diffraction yet, have we?

GEORGE BARBASTATHIS: [INAUDIBLE]

COLIN SHEPPARD: This is introduction to it. yeah. This is an example of Fraunhofer diffraction then. If you're a long way away, you get what's called Fraunhofer diffraction. This is what it looks like. And as we're going to show later, this is very simply related to the Fourier transform of the structure, the object that we're going to look at.

OK. Yeah. So this case here, eventually this one would look like this. But we haven't gone anywhere near far enough for it to look like that. But eventually it would get like that. The distance have to go depends on how big the aperture is. End of show.