COLIN SHEPPARD: OK, so for the ones that just arrived, I'm going to be giving the lecture today. But George is here. So if you've got questions, you can ask me.

I don't mind being interrupted at any time. But also, if you want George to answer something, he's sitting here as well. So he can also contribute.

So last time, George was looking at the Fourier transforming property of a lens. And he looks at the-- derived the 1D case, where you've just got a function of, x rather than a function of x, y. You've got this transparency a distance z in front of the lens. And then you've got a screen placed a distance f behind the lens.

And so the question is, what do you get on that screen? And I'm not going to go through all the algebra of that again because George did it last time. This was the calculus. This was the answer. What you find is that what you get in this plane here is actually the Fourier transform of the thin transparency.

So g of x is this transparency here. You get the Fourier transformer that. And multiplied by a parabolic phase factor, which depends on the position of this screen-- of the transparency.

What is the value of this z? You'll notice that you get this factor here, 1 minus z over f. So if z equals f, this is going to be 0. And this phase factor is going to disappear.

So for the special case where z equals f, all you see here is just the Fourier transform of this. Very simple-- very simple answer. And people actually have known this for quite a long time.

And at one time, it did look really attractive that you could actually imagine using this as a real way of doing Fourier transforms. Especially as we're now going to do, you can immediately generalize that to the 2D case. So you can do a 2D Fourier transform. And when you think that the light propagates through this at the speed of light, you can do a 2D Fourier transform at the speed of light.

And when people first realized this, which was back in the 1950s-- '50s or early '60s, computers were very slow then. So doing a 2D Fourier transform was not trivial. So a lot of effort was put into trying to design instruments to do that. Yes.

GEORGE BARBASTATHIS: Do you know who was the first one to realize this? Was it Vander Lugt?

COLIN SHEPPARD: The question is, who was the first to recognize that? Yeah Vander Lugt was at Michigan, wasn't he? And interestingly, the holography guys were also in the same lab.

What's the guy's name. The one-- Emmett Leith. And interestingly, Emmett Leith and Vander Lugt both published virtually the same paper in the same year. It's a very, very interesting historical thing there.

They're both from the same lab, and they both invented almost the same idea as far as I can see. And neither of them referred to the other guy. Very peculiar.

Anyway, I don't know who really came up with this idea first. But actually, I suspect that actually, even Ernst Abbe realized this. Ernst Abbe was the scientist who designed the Zeiss microscopes 120, 130 years ago.

And he came up with a theory for image formation in a microscope, which is actually what we're going to be getting onto pretty well, the next thing-- the imaging with a 4F F system. He understood that 130 years ago. And I don't think he's often given as much credit is he should get.

Anyway, you can use this as a way of doing a Fourier transform, and even a 2D Fourier transform. So you see, going from this stage to this stage is trivial, really. All we've done here is assuming that it's a separable function of x and y so you get exactly the same sort of integral for y as you get for x. So it now becomes a double integral.

This is the two-dimensional Fourier transform object of g of x, y. And you can see, you get again the phase factor. Here, because it's separable, you've got this e to the i pi x squared, and you got an e to the i pi y squared.

And of course, x squared plus y squared very nicely equals r squared. So this becomes a spherical type wave. So we've still got this 1 minus z over f So if z equals f, then again, this phase factor vanishes, and we end up with just the pure Fourier transform.

Yeah, as I say, so a lot of people thought that this was quite a neat way of doing Fourier transforms. There was a lot of work on optical computers all the way through the '60s and '70s. But the trouble is that, as fast as they could develop these systems, the computer people developed their computers even faster. So they were always one step behind, I think.

So I think in the end, people realized it was a waste of time, and that it was never really going to become a practical way of doing real calculations. But in the early days, it was. And Emmett Leith we were just mentioning, he first developed these systems for actually analyzing the images they took from a synthetic aperture

They'd fly an aeroplane over North Korea, taking pictures. And then the idea was that the plane takes many pictures as it's traveling along, and they're coherent with respect to each other. So you can actually combine them digitally afterwards and get an image, which has got amazing resolution. And you can see the guy in North Korea reading his newspaper or whatever. So they really did that sort of thing way back then.

Right, so that's the equation here then. And yeah, so this is just going through how that thing was derived. So we've got the screen here, the g of x. By the time we get here, it's convolved with some parabolic phase function.

And then after it goes through here, it's multiplied by another parabolic phase function. And then as it propagates here, it's convolved with another parabolic phase function. And if we have that magic condition, that this is f and this is f, most of these parabolic phases cancel out and we're left with this very simple final answer. So that's basically what's happening.

And George went through the various stages of doing this. Of course, you end up with a couple of integrals of products of three phase functions here. So the algebra process was quite complicated. Well, not really complicated, but just lengthy. But everything in the end cancels.

If you haven't got z equals 2f, though, then you can see that you have got this phase factor. So you don't actually get exactly the Fourier transform in this plane in that case. You get the Fourier transform multiplied by this phase factor. And that can be important in some cases. That's something that you need to take into account.

There's another special case, actually. That's the case when z equals 0. The case when z equals 0 is if this screen coincides with the lens-- immediately close up to the lens.

And I think Goodman does that one as well, doesn't he. And so if z equals 0, this goes. And so we're still left with this phase factor, e to the i pi x double dash squared over lambda f.

GEORGE BARBASTATHIS: And also, Goodman does the case of the transparency after the lens, which I did not do in the notes.

COLIN SHEPPARD: Right, OK. Yeah, so Goodman does this case of the transparency placed in here.

GEORGE BARBASTATHIS: That this is z negative.

COLIN SHEPPARD: Yeah, but it's very similar answers. OK, so this is just writing down finally, then. We're defining then the Fourier transform of little g is big G. So the final answer is that you've got this parabolic phase function times the Fourier transform of g.

OK. So this is pointing out that we've now looked at two different ways you can do a Fourier transform. One way is to just propagate the light a very large distance. You remember, if you have [? frame ?] of a diffraction, if your screen is a very long way away from the diffracting screen, object, whatever it was called-- transparency then again, you get this Fourier transform relationship.

But of course, here he's put a proportional side. There's some things out the front there. And in particular, there's going to be, again, a parabolic phase term, isn't there? Because in this case here, what you'd expect actually is that this Fourier transform is going to be actually on the surface of some curved surface, not of a plane surface.

So if the screen is plane, then it's going to be multiplied by some parabolic phase term. And now we've just looked at this other case where we don't have to have a very large lab in order to do this experiment. We can actually make the experiment quite small so that it fits on top of the table just by using a lens. And really what you can see, that that is actually doing, according to your geometrical optics ideas now, you can see what this is doing.

This is producing an image of this at infinity. So you can see that these two things are actually equivalent. Well, what it's doing is bringing what would have been at infinity to some finite plane.

So that's all it's doing there. But again, we've got this [? string ?] here, this proportional sign, which means there's some other constant factors out there, including in general this parabolic phase factor, which for this particular case, when you got f and f, it disappears. The parabolic phase function disappears when you've got this very, very special case with the two f's there.

OK, now this one. What this is really trying to show is that, really, there's two ways of looking at this problem. These two diagrams here are really trying to show the same case, but trying to get you thinking about a different way of viewing the problem.

So in this case here, you can think of any point on the screen on the transparency is going to act like a point source. So it's going to produce a spherical expanding wave. And then this is a distance f, so that spherical wave is going to be converted into a plane wave The angle of propagation of this plane wave depends on the distance at this point from the origin, of course. And so you can think, then, of what you see on the screen as being made up of what's called an angular spectrum of plane waves, where the strength of each of those and get a spectrum components comes about from the value of at that point.

So that's one way of thinking about it. And this way [INAUDIBLE] thinking about it the other way. This way, you're thinking of the object. It's going to diffract. If you imagine your transparency as being like a grating, it produces a diffraction. It's producing light grating orders

For example, if this was a very simple case of a sinusoidal grating, you'd expect it to produce some grating orders. And these grating orders correspond to plane waves traveling in particular angles. So this shows the angular spectrum of the diffraction of the light given by the transparency. And then the lens collects these different plane white components. All of them are going to be converted into a spherical converging wave that converges a distance f away.

It doesn't say that, but this must be [? there. ?] And the plane wave's moving at different angles are going to converge to a point which, of course, is proportional to the sine of the angle of the propagation. So here, at any point on the screen, you're going to sum up again these angular spectra components, each of which is going to give a spherical wave. So you can see, looking at that diagram again, they're like inverses of each other.

I think this is another example where, very often, you can think of a reciprocal type behavior. You could just flip this over, and it would look the same the other way. OK, so this way here, then, we're thinking of this as just like the Huygens sense that we spoke about when we were first doing diffraction [? there. ?]

OK, so this is giving some nice pictures. I'm going to steal this presentation afterwards, because there's so many good displays in this thing, that I'm going to use it in my own lectures in the future. So this has given me some good ideas. So we start off, then, this is a rectangle. Oh, I'm going to say something in a minute, though.

So this is a distance f. This is a distance f. So this length, then, as we just described, is going to produce a Fourier transform of that. So the Fourier transform of rect is a is a sinc. So you'd expect to get this a sinc in x times a sinc in y.

Now, the thing I was going to say though, George, this looks to me that this is longer than that. This is longer this way than that way. So this should be longer that way than this way, I think.

GEORGE BARBASTATHIS: That's the aspect ratio of the screen actually.

COLIN SHEPPARD: Yeah, the aspect ratio of the screen ought to be 90 degrees round.

GEORGE BARBASTATHIS: There's not a [? way of the ?] Fourier transform, I guess. [LAUGHS]

COLIN SHEPPARD: Anyway, so this conversation does bring it home, I think, and a very important point that we've mentioned before in the course. That if you make a function smaller, it's Fourier transform is wider. If you make the function wider, this Fourier transfer is smaller. And you'll notice, then, these products of these two sincs gives a bright spot in the center, a series of bright spots along the axes.

There are also some bright spots in this region here. But they're so weak that you can't actually see them. Really, you can only see the ones along the axes.

All right, so that's the first example. Then this one's showing how, if we make the transparency smaller, the rectangle smaller, so this pattern gets bigger, now actually we can just about manage to maybe see some ones other than along the axis.

And then this one is an interesting one. So now we've got not just one aperture, but three apertures. And so how can we work out what the Fourier transform of three apertures would be?

Anyone got some ideas? I'm sure [INAUDIBLE] will have an idea. Anyone? How would anyone think of doing this sort of thing?

Have you come across a thing called the convolution theorem? Probably you should have done somewhere, I would have thought. So there's the thing in Fourier transforms called convolution theorem.

So this is actually, like, three spikes convolved with a square, which really all it means is we place a square on each of the three spikes. And the convolution theorem says that the Fourier transform of a convolution is the product of the Fourier transforms. So what you'd expect is that the Fourier transform that you're going to see here is the Fourier transform of the rectangle, which is this, multiplied by something else, which is the Fourier transform of the three spikes.

And the three spikes is like a delta function and a plus or minus delta function. The plus or minus delta function together, e to the i sampling plus e to the minus i sampling gives cause. And the one in the center just gives a constant. So the Fourier transform of these three delta functions is 1 plus a cosine, or something like that. And so that's what we expect to see, then.

This thing, this pattern that we saw before, is going to be multiplied by a cosine function which comes about because of the separation of this. And of course, because this distance is actually much bigger than the width of the rectangle, the spacing of the cosine is much finer than the Fourier transform of the rectangle. And so that's why these fringes-- you see these fringes, very fine fringes, they become a result of the fact that these components here are quite well-separated. So that's another very nice example.

And so all these theorems that you learn in Fourier transforms can all be directly taken over into optics. And sometimes, they can save a lot of work, because you can just use the standard forms for things that you can look up in a book of tables. And it saves you a lot of algebra, or even worse, arithmetic.

Right, so this is another example. So now we're back to the bigger rectangle, or something close to the bigger rectangle. The Fourier transform of the rectangle, we saw before. But now this one, you see this is multiplied by some fringes.

So this is like the opposite of this one, really. So this is a convolution. This is a product. So we use the same convolution theorem again.

The fringes, of course, correspond just like we said before. The Fourier transform of the three delta functions was the fringes. So the Fourier transform with the fringes is like three delta functions. So what we'd expect, this is a product of the fringes and the rectangle. So we expect here to see the convolution of the Fourier transform of the fringes and the Fourier transform of the rectangle.

And that's what we see. So we see three of these diffraction patterns situated on these three spots corresponding to the Fourier transform of that grating. So a very nice display, that one. I like that one.

OK, so this is now introducing the idea of what's called the 4F system. So we're going to do this not just once, but do it twice. So this just shows what we've had so far. If we've got some particular transparency, we get some particular diffraction pattern in the back focal plane.

And then in the 4F system, we do the same again. We now have another lens which is going to do exactly the same. It's going to look at this pattern and produce it's Fourier transform.

So remember, if this is f and this is f and this is f and this is f, we won't have to worry anything about these parabolic phases. It's just straight Fourier transforms with nothing else. And so you can see what you get.

It's very simple. You start off with this object. So here is Fourier transform And if you take the Fourier transform of the Fourier transform, what do you get if you do a Fourier transform of a Fourier transform?

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Yeah, it's not quite, is it? You remember? You invert it. That's right, Yeah, so someone got that.

If you remember, if you do a Fourier transform of a Fourier transform, you get the original function inverted. You have to take the Fourier transform four times, and then you get back to where you start. And so that's what we can see anyway. So this is the inverted. Well, of course, inverted, that's just what we expect, isn't it? Because we know that when we have an imaging system, we get an inverted image.

So everything fits together. It gives exactly the same results that we expect from our geometrical optics. And so in this case, of course, this is a rather symmetrical object. So the fact that this has been inverted, I guess you can't really even see.

OK, so this is saying, then, that you can make a 4F system which produces, then, an inverted version of the original object. You might wonder what the point of that is. I mean, I guess in a way, there's no point in making an image which is much the same as the object, really.

But one reason why you might want-- well, there's two reasons. One is that it's actually produced this image in a different plane. So it's actually transported this distribution of light here from this plane to this plane. So that's one type of application where you might want to do something like this.

And the second thing, of course, is that although this is always called a 4F system, one we obviously have to have this and this [INAUDIBLE] both equal for this to work. We have to add this and this f to be equal for it to work. But this f doesn't have to be the same as that f.

So you can actually have a system like this, where you actually produce a magnified image or a demagnified image. And now, that, of course, is getting a lot more useful. So that's how you can produce an image of an object in another plane.

Now, the next thing to note is, of course, so far we've assumed that all this light gets through this optical system and gets here. So if that was the case, you get a perfect image. If all the information from the objects was getting to the image, then you wouldn't lose anything.

But in practice, of course, that's never going to be true. And, well, one way it's not going to be true, of course, is because in practice, these lenses are never going to be infinitely large. But also, we could, if we wanted to, actually put some sort of mask in this plane here to actually stop some of that light getting through to the image plane.

And so now we're going to look at what happens if we do that. This is much the same overall effect as George spoke about in previous lectures, actually. Not imaging case, but he spoke about-- you remember the results with the building with the windows that you could actually, just by doing digital processing on the Fourier transforms, you can manipulate the strength of the grating components in the final image. So the same is true here, except now we're doing it optically rather than digitally.

So what you'd expect, then, is if we put something in here-- so here, this is showing a passband. So let's say you put in here some-- a [? paint ?] screen with a small hole in. The hole is this passband, and the screen is the block band. So it only let's through the region in the center. According to how big that passband is, you'll get different answers in the image plane.

Yeah, of course, it says the transparency may be a grayscale. It doesn't have to be a binary type thing, black and white. Or it could even be a phase mask. It could be that this is changing the phase by, let's say, pi by 2 relative to this or pi relative to this, or whatever. And you get different answers according to what you choose, obviously, for the geometry of this what we call spatial filter.

OK, so yeah, this is going back to this idea of these two ways of thinking about the same problem. So this is the first way, where we think of our objects as being made up of lots of point scatterers Each of these scatterers gives rise to, like, Huygens' spherical wavelets. Each of these is collected by lens. The lens convert that to a plane wave.

The plane wave reaches the second lens. And because it's a plane wave, it's converted back to a spherical wave that converges at distance f away from that. And now we've actually introduced the fact that these f's can be different-- so f1, f2. This has to be equal to this, and this has to be equal to this. Otherwise, the phases won't cancel out properly.

And so you can see that what you'd expect, then, is that because in this diagram here, F1 is bigger than F2, this image is actually smaller than the original object. And note now, you can also see directly the geometry showing the inversion, the point above the axis has come to an image below the axis here. So that's the one way of thinking about it.

And then the other way is to think in terms of spatial frequencies. We think of this object as being like a grating, which is going to produce grating orders which travel in different directions. Each of these grating orders is like a plane wave. And when it reaches the lens, it's going to produce a converging spherical wave, which comes to a focus at distance f behind the lens.

And then after that, of course, it doesn't stop. But it carries on. It diverges again, reaches the second lens.

And then the second lens, because this is a spherical wave coming from a distance f away, is converted into a plane wave again. So again, you've got this nice symmetry between what goes in and what comes out. So there are these really different ways of thinking about how this is all going on. End of show.

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Right then.

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Right, OK. So all that was just a warm-up. That was last week's lecture-- last Monday's lecture we've done so far.

Right. Oh, yes. Lots of interesting things coming up later. OK, so we're going to carry on straight away with more of this 4F system and spatial filtering. So we're just getting started on this. And so there's lots of very interesting things you can do with these.

So this is back to the same thing again [INAUDIBLE] showing again, but a bit more detail now, because we're actually going to really look at more detail of where these orders are formed and what you can do with them. So we have a grating object. Let's imagine this is just a 1 plus a cosine, which means then it produces a zero order, and then a plus and minus 1 order. So three plane waves. And as we said, each of those plane waves is going to be converted to a converging spherical wave.

And so in this back focal plane, you'll see three spots corresponding to the three orders. And then carrying on, these plane waves then converted back into a spherical wave. And you get an image which is back to the grating again.

So at the moment, the optical system has transmitted all the information. And so what we really expect is that the image is exactly the same as the object, apart from the fact it's inverted and it might be magnified, or the opposite. But the resolution you'd expect would be the same. All the detail in here, you'd expect to see in there. So the next thing to look at is what happens if we now place something in this plane that changes the relative strength of these different orders. And there are lots of things that we could try.

The first one, the simplest one to look at, is what happens if we just put in something, a mask here that lets through the zero order and not the others. The plus 1 and the minus 1 order are obstructed by this screen. So all we let through is this single order, which is going to be converted to a plane wave.

And so on your screen here, all you'll see is the plane wave. So the image that you see-- well, I call it an image. But it's not really much of an image. It's completely featureless. It's just like a gray spot-- a gray region.

There's no detail there. So we've lost that structure. Actually, because we put this mask in here, the structure is beyond the resolution limit of the optical system. And so we can't see it.

OK, so this is what we might call a low-pass filter. So if you think of these as being like these grating orders basically represent different spatial frequencies. So just in analogy with what you have in electricity, where you've got frequencies and time, in space, we've got distance and we've got spatial frequencies. So this is why we could talk about low-pass filtering. We're filtering the spatial frequencies.

And just like in time, you remember when you've got your hi-fi system, if you turn down the treble, that's acting like a low-pass filter. It's cutting out the high-frequency sound and letting through the low-frequency sound only, isn't it? So that's what we call a low-pass filter.

And this is exactly the same. We've put in the mask. It's a bit different in a way from the electronic case, because in optics, we've got a lot more freedom as to what we can do.

We can actually put in something here which has got real sharp edges to it. In electronics, if you're trying to make a filter, there's always a limit to how much you can make the roll off of the filter. So you can't get something that suddenly cuts off like that very easily.

OK, so this is going through the maths of this. Our original transparency, then, is a 1 plus cosine. And the 1 plus cosine, we Fourier transform that to find what we call the object spectrum.

And so in this case, finding the Fourier transform of cosine is very straightforward, because we just expand cosine as e to the i something plus e on the minus i something over 2. And each of those e to the i type terms will correspond to a single frequency. So we end up now with three terms at different frequencies.

The one Fourier transforms to a delta function of u, which is a spike where u equals 0. And the two components of the cosine correspond to another 2 delta functions, each situated-- this one's a u equals some value u0. And this one is some value u equals minus u0. So you've got these three spikes that we've been talking about for the last few slides now.

So this is the object spectrum. It's the angular spectrum that we've got here. So actually, you can see that u here basically represents the angle that these plane waves are traveling in. And there's three different angles that the light can travel in.

And so yeah, that will propagate, then. And as we've said before, you can see that these rays are moving at an angle. And so as you go along a distance f, you get a shift in the x position which is given by, of course, the slope of this.

The more the slope, the more it deflects by the time you travel a distance f. So this then gives the shifts of these two things. So that's before the screen, then-- sorry, before the mask.

[SIDE CONVERSATION]

Yeah, so this is showing, then-- so these raise, then, as they come down here. We put in this pupil mask which is going to get rid of these two. And it just leaves the other one. So we just left them with our single beam, which as we said before gives an image which is not really an image.

Right, yeah, so after the pupil mask, we multiplied by the mask. Gets rid of those two. Just leaves the other one. And finally, then, the final image you'll see is going to be the Fourier transformer of that. And the Fourier transform of the delta function is a constant. And so you see something which has got no detail in it.

OK, so this gives another example, then. Here, we're looking at some binary amplitude grating. So this is a particular example-- another example.

So last time, we looked at the cosine type thing. So this one is now a binary grating, which means that the object is something which is either black or white. And we shine a light onto this. So it's like a square wave grating. And so unlike the previous case where cosine, of course, just produced those three grating components, in this case, there's going to be a whole load of grating components.

They're going to go on forever. But they get weaker and weaker. And so you can see here, it shows a few of them. The first, the second, so there'll be the third, the fourth, and so on. And then minus 1, and so on and so on.

And each of these is moving at a different angle. So when the plane wave gets to the pupil mask-- sorry, when the plane wave is focused to the pupil mask, it's going to reach there. At a point, the distance from the axis, of course, is going to depend on the angle of this plane wave component, the grating component.

So if you look here, you're going to see all these different grating orders all arrayed in this position here. And so by putting some sort of mask here, we can monitor and change the relative strength or phrase or whatever we like of those different components. And then they'll be added together, produce some nice effect that we might want to look at. So let's see what we're going to get in this case.

So in this particular case in the picture, the size of this aperture is enough to let through the central three components, the zero order and the plus and minus 1 order. But it stopped all the other ones. So what you will see here is the Fourier transform of this.

[INAUDIBLE] before. It's these three spikes again. The Fourier transform of the three spikes is just a 1 plus a cosine. The relative strength of these might change according to the properties of these-- of the transmission of the mask.

But anyway, what you see finally is an object which is like-- you see the 1 plus cosine type variation in amplitude. And so you notice that this thing here started with these very square edges. We no longer see the very square edges.

So this is an example of how we've lost resolution. The optical system has got a certain finite resolution. There's detail in this original object that we can't actually see in the final image. And this is because not all the light has got through the system.

OK, so we're going to look at a particular case now. This is the bit I was saying I wasn't very keen on. No, no, it's all right. But anyway, there's a load of maths here, which I might actually gloss over a bit, and let you look at it yourself later.

But basically, then, we're looking at a particular example. It's a binary grating with perfect contrast, m equals 1-- i.e., this goes from black to white. That's what that means. It's got a period of 10 microns. The spacing, the period, is the distance between the repeat on this grating.

And it's got a duty cycle of a third, which means that the white bit is a third of the black bit. Is that what that means?

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: A third of the whole.

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Yeah, a third of the period. OK, and then we're also told that the 4F system consists of two identical lenses of focal length f equals 20 centimeters.

So both f1 and f2 are 20 centimeters. And we put a pupil mask of diameter 3 centimeters here at the Fourier plane. What is the intensity observed at the output image plane? So that's the question.

OK, the secrets to solve this kind of problem, what you've got to do is to work out what you're going to get in this plane, which is basically the Fourier transform of this. And then you're going to have to multiply by the mask and see which of those orders get through. So the relative scaling of this diffraction pattern in the mask is obviously all important for that. If you don't get that right, then you'll get the wrong answer, obviously, because you might get orders getting through that shouldn't, or the other way around.

And so once we've done that, we know now what the grating orders, the Fourier components in the final image, are going to be. And we then do another Fourier transfer. And we've got the image amplitude.

There we are. Again, we go [? do ?] scaling to get it the right size. Normally, of course, at the very end, we would be interested in intensity, not amplitude, as well.

OK, so this is our example low-pass filtering, a binary amplitude grating. So this is this grating as we've chosen. And so you can see this, that the dimension's in microns. So this is quite a fine grating.

So you can see that it's not symmetrical, by the way. The duty cycle is not a half, right? So the width of the bright regions is smaller than the width of the dark regions. And that's important when it comes to looking at the Fourier transform, and therefore the spectrum of that.

And then the pupil mask that we're going to look at looks something like that. We let through the frequencies near the center. And we keep out-- reject the frequencies which are further from the center.

OK, so for the Fourier transformer of that, this is a repetitive function. So the Fourier transform just becomes a Fourier series. And this is an expression for the Fourier series of this. It basically consists of two parts. It's got a sort of envelope, which is a sinc.

Basically, this is going to consist of an infinite number of discrete frequencies. It's going to consist of a constant term, which is the zero. It is going to consist of a first harmonic and then a second harmonic and a third harmonic, and so on.

If the duty cycle has been chosen to be a half, i.e. it was a square wave, then by symmetry, it turns out that the second harmonic and the fourth harmonic and all those all vanish. So you'd only end up with the zero, the first, the third, the fifth, and so on. But in general, there will be all the orders there.

And the relative strength of those orders here, you can see the size depends on the value of this sinc function. Actually, that's a very easy way of seeing about the even orders canceling. Because what happens is, if the duty factor is a half, it turns out that the even orders when q equals even number coincide with the zeros of the sinc function. And that's why they vanish.

So where this comes from is that-- so we're saying that this function is a-- maybe it's easier if I go on to the next slide, I think, which is the Fourier transform. Yeah, so the Fourier transform, the strength of these components is what we've got to find out. And so this function Fourier transforms to something which is a series of delta functions corresponding to the different orders, different values, of q. And the strength of the different components is given by the value of the sync. So you can see again, if there are some values of alpha q which can be 0, and therefore would mean that those components [? will ?] be there.

OK, so when it gets to the pupil mask, all that's happening-- maybe we ought to go back to the-- we'll go back to this diagram. So you can see that what we need to know is, this is coming to a focus at this point here. But we need to know what that distance here is. But we need to know that relative scaling of where that diffraction order appears. Oh, that was well up here [INAUDIBLE].

OK, so this is this lambda f that comes in here, because it's going through a distance f. So that gives the scaling of this distribution that you get there. And so that's where u now is replaced by this x double dash over lambda f-- the scaling of that size of that diffraction pattern.

And then this is put in the answer, [INAUDIBLE], is that you get this-- oh, yeah, we're told the alpha is the juicy fat cycle. So that's the third. So alpha is a third in here. And where do we get the 1 centimeter from, George?

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Yeah, I think I saw something like this on the crib sheet. Which he will hide, by the way. So we've got-- all right here, we've got lambda equals 0.5 microns, f equals 20 centimeters, lambda period is equal to 10 microns. So that gives that lambda f over [INAUDIBLE] lambda is 0.5 times 20 over 10. No, that's not-- this is microns.

AUDIENCE: That's right. Lambda-- both lambdas are in microns.

COLIN SHEPPARD: Oh yeah. These are in microns. These are in microns. These are in centimeters.

So this is in centimeters. And this is equal to 1 centimeter. Right, OK, so that's what's going on there. So that's the 1 centimeter.

OK, so that's just plotting, then, these different orders. And so you can see here the zero order plus 1 plus 2 plus 3, so on. But the screen blocks out all the ones, except for these central ones. And so this thing, these relative strengths, must come from-- well, there's a factor of third at the front is this one. And then this is multiplied by the value of sinc of a third, which must give the rest of this thing here.

GEORGE BARBASTATHIS: And so I have posted in the website a supplement to this lecture. So in the supplement, I have worked out all these calculations-- how you compute the sinc of 1 over 1/3, sinc of 2 over 1/3, and so on. So you can see some more details in the website.

COLIN SHEPPARD: Yeah, OK. You remember a sinc, you've defined the sine of pi x over pi x. Remember that? So that's the way that people will usually define a sinc of x equals sine of pi x over pi x. All right, so it's defined in that way, because then lots of nasty pi's and things cancel out and everything else.

All right, so we've got to multiply that by the size of the mask. So as it shows here, the mask going from minus 1.5 centimeters to plus 1.5 centimeters, which means that this grating component, which is at 2 centimeters, doesn't get through. So that's what gets through the aperture. And then we Fourier transform that again. And so the delta function transforms to an exponential-- a shifted delta function, and corresponds to exponential.

And therefore, our output-- this f was the same in this case. The two f's of the two [INAUDIBLE] lens were the same. So we get 1 over lambda f again there.

AUDIENCE: What will happen if we make the width of this mask minus 2 to plus 2? Will there be any diffraction due to this second order?

COLIN SHEPPARD: Well, if it was exactly minus 2 to the plus 2, then I guess if it really was exactly there, you'd expect half of it to get through this order. I mean normally, we would assume it either lets it through or doesn't let it through. So if you made the size of the aperture 2.5, then of course this next order would get through, wouldn't it?

AUDIENCE: But if it is exact [INAUDIBLE], that means do we have to consider the diffraction of [INAUDIBLE]?

COLIN SHEPPARD: No, I think that effectively, I think you have to say that half of it got through.

GEORGE BARBASTATHIS: Actually, that's correct. Because in actuality, the grating is finite. So these are not delta functions. They're sinc patterns. So if you place the grating at exactly 2, then you allow 1/2 of the sinc to go through, right? So therefore, 1/2 is [INAUDIBLE].

COLIN SHEPPARD: Yeah, that's a very good way of thinking. Here, we started off by saying that this is a continuous repetitive grating. And so we can think of it in terms of a Fourier series.

But of course, it's not actually infinitely big. I guess it's not going to be more than a few light years across. So this is multiplied by another rect function, really. So the Fourier transform of this whole thing multiplied by rect is something convolved with a sinc. So each of these things is not really a delta function. But it's really a very narrow sinc, which depends on the size of our grating. .

So the point George is making is, if you think of it as a sinc, then it really would be letting through half of it. I think mathematically, it'd have to be like that as well. You get used to these when you do-- you must remember from solving partial differential equations, when you do the boundary conditions, you always have to do that sort of thing. It seems a bit of a cheat, but--

OK, so this is putting the numbers in, then. So this is the three components. These are these root 3 over 2 pi's we had there. And then finally, these two combined together to make a cosine. And our final amplitude of our image is going to be a constant plus of cosine. You notice that it's not-- the amplitude of this is not the same as the amplitude of that, of course, because of the fact that it was a this rectangular grating.

Yeah, and finally, though, as I just mentioned briefly a little while ago, what you actually see is not the amplitude. What you see is the intensity, which is the modulus square of the amplitude. So you remember, we said that we normally are going to call the intensity just the modulus square of the amplitude. There perhaps should be some halves and funny things out there if we were really doing it according to the Poynting vector.

But normally, we don't bother with that. We just say it's the modulus square. So the modulus square of this, we're going to get this squared and twice the product of these two, and then this one squared.

So we get this term in Cos squared. And then you have to remember your formulae that you learn in maths, Cos 2 theta equals 1 minus 2 cosine squared theta, which allows us the right Cos squared theta is 1 plus Cos theta over 2. So this Cos squared, we can now replace by a constant plus a cosine of twice the angle. And so this is the constant bit, and this is the cosine of twice the angle.

All right, so that gives a another harmonic in the image. So that's something that people are a bit confused by sometimes, first of all. Because we've gone through a lot of talk about how it's only the first time on it that's gone through. But actually now, when we look at the intensity, we've now got a second harmonic that's appearing in the final intensity.

So note, the second harmonic term intensity due to the magnitude square operation. This term explains the ringing in coherent low-pass filtering systems. So the ringing is the wiggles you get on the sharp edge.

Right, and that's a sketch of what it looks like. So it doesn't look quite like a cosine grating anymore, because when you square it, the low values get depressed more than the high values. So it makes it look more peaky, like that.

OK, so perhaps we'll start again. I was just waiting for one of our people to come back. Ah, well, there we are. Just in time.

OK, so yeah, before we carry on, then, there was a question just before the break. So we're going to have a question.

AUDIENCE: OK, actually, I have two questions. The first question regarding the physical meaning of the low-pass filter, just now in the previous slide.

COLIN SHEPPARD: Yeah, OK, so the question is about the low-pass filter. So you remember in-- perhaps I'll draw something from electronics. So you remember from your hi-fi amplifier.

So this is frequency. So these are high frequencies and these are low frequencies. So this corresponds to [? ah! ?] And this corresponds to [? uh. ?]

And if we multiply this by some function that looks like that, this is what we call a low-pass filter, because which we let through the low frequencies and we stop the high frequencies. So the same is true here. So when we say low-pass, we mean that we let through the spots which are close to the axis, and stop the ones that are far from the axis. So that's the first question. Then there's another one.

AUDIENCE: The second one is regarding the Abbe's rule because Abbe's rule tell us if we can [INAUDIBLE] two orders of the wave, we can form an image. But just [INAUDIBLE] even with three orders of the wave, we still get very [INAUDIBLE] compared to that original one. So I wanted to check out what is the reason.

COLIN SHEPPARD: Right, OK, so this is going back to Abbe's rule, which I don't know-- I've never really heard what he said, actually. But have you got that in a book or something? It's probably in Chinese, is it? But anyway, the Chinese Abbe came up with this rule that you must have at least-- well, we gave an example just now, didn't we, where you only had one beam. One order got through, and we didn't get an image, because you're obviously not going to get any interference if you've only got one beam.

So according to Abbe, you've got to have two beams interfering in order to get some variation in amplitude in the image. And so he says at least two, I guess. So if you think of these, you could let throw the-- last time we had the 0 and the plus 1 and the minus 1, which is three beams. And the point that was made in the question was that, although we got an image, it was wasn't a very good image. So I think that's probably right.

I guess Abbe's rule is all to do with resolution of an optical system-- the smallest detail you can see. But it doesn't necessarily mean that you're going to get a good image. And changing the relative strengths, we're going to see this in lots of examples now, actually. I guess the whole of the rest of this lecture will be about these-- more examples of stopping beams and so on, and changing beams and seeing what it does to the pattern.

And so very often, what you'll find, then, is you get something which is maybe not a good image of the original object. But it's an image, which is much better than nothing. OK, so does that answer the question?

OK, so we now looking at a bandpass filter. So going back to what I was saying here, the bandpass filter, of course, is something that looks like this. A bandpass filter stops the low frequencies and also the high frequencies. And so you can see here, the optical equivalent of this, it's now stopping the zero order, but allowing the first orders to get through. And then the later ones would also not get through.

Oh yeah, that was another point that George mentioned to me during the break that I hadn't appreciated fully. And that is that this particular example, which had this duty cycle of a third, actually, it turns out that the third order of this one vanishes. So I guess the third, the sixth, the ninth, all these.

And this is just like we said. That for a square wave grating, all the even orders cancel. So it's all to do with the zeros of that sinc function. We add sinc of alpha q, where alpha is the duty cycle. So for any duty cycle, there will be particular values of the particular integers for the order, which would vanish in the Fourier transform.

Right, so here, we're letting through the plus 1 and the minus 1. Did something [INAUDIBLE] there?

AUDIENCE: Yeah, [INAUDIBLE].

COLIN SHEPPARD: Ah. Right, yeah. OK, so the amplitude is going to be a 1 plus cosine. And intensity, we're going to get a second harmonic image, just like we had with the previous case.

0 and then now we're going to go through this in more detail. So it's the same grating as before. So we haven't got to work out all that stuff again.

You remember, 1 centimeter was where the first order was. So we've got this series of these different grating components. And these ones are 0, as I just described. The next one wouldn't be 0. And we're putting this masking, which is just going to let throw these two.

And so mathematically, we can write mask as being the sum of two rectangles-- rectangle, another rectangle. So then it lets through those two components and stops the others. So we've got these two delta functions. And then we just got to do a Fourier transform again. Each of the delta functions produces a plane wave-- [INAUDIBLE] complex exponential.

We're then going to add those together. And we've got to change the [? interdistance, ?] of course, and then combine these two complex exponentials to make a cosine. And then finally, we've got to square it again in order to get the intensity [? out. ?] So notice what we get now. You see we get the constant term.

You see, here, it's just cosine. There's no constant term anymore because the constant term is gone. So because when we square it, we get Cos squared. The Cos squared, we go into double angles again. We get a constant plus a cosine.

And so you see that actually, we get an image where the first harmonic doesn't appear at all, but the second harmonic does. So this comes back to your question in a way. You end up with getting an image which is actually really not like the object. Because you see that this has got twice the frequency of this, although Abbe says that you can get an image, he doesn't say that you can get a correct image.

And in this case, you could be very confused by this. You could look in your microscope at something and you could see this. And you'd be completely mistaken about the object that was actually producing this, because it's actually twice as big as what you expect. So--

AUDIENCE: I've got a question.

COLIN SHEPPARD: Oh, yeah. You frightened me.

AUDIENCE: If you grab just the zero order and say the first order, but not the minus 1, you get the same period as the original grating?

COLIN SHEPPARD: Yeah, I guess so. Would that be right, George? If you take this one and this one, so you've got-- and then you square it. You're going to get 1 plus an e to the i. You're going to-- yeah, he's working it out. He's working it out. We'll have a race, shall we?

You're going to have a 1 plus e to the [INAUDIBLE] i something or other. And then this module is squared. And then you can take out an e to the i something over 2.

You've got the answer? Yeah. Something like that. So this is cosine. So you're going to end up with cosine squared.

It's the same as this, isn't it? [INAUDIBLE]? It would just be the phase term. No, it can't be, because it's half as far apart.

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Yeah, you got the first harmonic, because it's the square of half of it. Yeah. OK, so it's double half of it. So you do get the first harmonic, yeah. You're quite correct.

But it might work for this object. But I guess there'd be other objects, which would also produce funny results. And of course, as we know-- well, some of us know. Most of you won't know.

But for example, [INAUDIBLE] will know that if you have a mask, which is actually metric like that, it's going to have some quite interesting properties. So I don't know whether you're going to go through these at some later point.

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Maybe in the exam. OK, so everyone [INAUDIBLE] up on asymmetric gratings.

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Yes, single-sideband modulation. Yeah. And what's the other word for it? What's this, when they measure the flow of the-- in aerodynamics? Schlieren. It's basically Schlieren.

OK, and now we're looking at a tilted illumination. So same sort of system as the last one, but we're now not illuminating it where the plane wave along the axis. We're going to tilt this. And so what we find if we've got a grating, and we illuminate it with a plane wave, what do we find happens if we tilt the plane wave?

Have you come across this-- you did this earlier on, did you? I don't think I was there for this one. But actually, all that happens is it just tilts the diffraction pattern. Very simple.

So you can see now, the zero order is still in this direction. And all of the other orders just move [INAUDIBLE]. They all rotate through the same angle. This is actually only true if you can think of this as being a thin grating.

And if you go into theory of volume holograms and things like that, then you'll find that things become much more complicated. The relative strengths of the gratings can change and so on. But this is a thin grating, so that you can use this property that you just multiply it by the transmit transmission of the grating, then that will be the case.

So you can see what that's going to do. All the gratings move up, which means that they all get moved up there as well. So you see the zero order grating, which was appearing down at the origin, is now tilted upwards and appears at this point here.

So with this bandpass filter that we've chosen, we now get the zero order through. So we're getting the 0 to through this one. And we're getting the minus 2 order through this one.

Right, so we illuminate the grating with an off-axis plane wave at an angle of 2.865 degrees, which is obviously very carefully designed, so that this grating this grating order moves through from here to here. So what happens? OK, so it's just describing what I just said.

Yeah, this is just describing why that happens, then, in terms of multiplying the-- this tilted wave is just going to have a complex exponential phase variation like this. So you multiply that by the grating. And then it's just [INAUDIBLE] this result of rotating the thing.

OK, so there are the grating orders. The zero order one is this one. And this is the plus and minus 1 and the plus and minus 2.

And so we let through those two. So we're now letting through the zero order and the minus 2 order. So it's a bit like the one in the question.

But now it's the minus 2 order, not the minus 1 order. So this is going to be another example where it doesn't give a very good image of it from that result. OK, so yeah, let's just [INAUDIBLE]. Here we are. So it's 1 centimeters. That magic angle was designed so that they're all shifted by 1 centimeter.

So after passing through the mask, we just get those two components. And so these are the two components. Ah, but [INAUDIBLE] no, it isn't. Yeah, that's right, because there's still-- it's a bit more complicated than where I was-- it's tricky, this.

So not only of course that we let through the zero order and the minus 2 order. But those orders have got this complex phase shift. Is that right?

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Yeah, exactly, exactly. Because of the phase, because of the-- yeah. OK, so that's what I'm saying. Although this is a zero order and a minus 2 order, they're actually shifted.

And so when you do the Fourier transform, they're just as though they were there plus 1 and the minus 1 order. And that's neat. Very interesting.

AUDIENCE: But with different amplitudes.

COLIN SHEPPARD: But with different amplitudes given by these two strengths here. And then-- yeah, so then we have to put the values of u in terms of the x. And then we combine the two exponentials.

Well, can't do that anymore. Can we? Oh yes, we can.

AUDIENCE: Well, you can, but [INAUDIBLE] sine and the cosine.

COLIN SHEPPARD: Yeah, exactly. So have you actually done that, then?

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Right, so you've got to do the modulus square of this. So you're going to have this squared plus this squared.

AUDIENCE: Plus [INAUDIBLE].

COLIN SHEPPARD: Oh yeah, OK, so [INAUDIBLE]. Yeah, of course you are. So then twice the products of these two. And-- OK. Is that right? No, sorry. That's not right, is it?

AUDIENCE: The magnitude of the first [INAUDIBLE] the magnitude of the second plus [INAUDIBLE].

COLIN SHEPPARD: Yeah, that's right.

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Yeah.

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: OK, and this is what it looks like. Ta-da. OK, so the field is the first harmonic. The intensity is the second harmonic because of the squaring. And the contrast is 0.7.

Right, so again we end up with a bad image as it's shifted, because it's double frequency. And also, the contrast is reduced in this case, because the strength of the two beams was not the same. This is calculating here the contrast.

This is the normal definition of contrast fringes-- the difference over the sum. And it comes to 0.7062. In the previous case, of course, the contrast was 1.

Right, another example-- a new pupil mask consisting of two holes, each of diameter 1 centimeter and centered further away from the axis plus or minus 2 centimeters from the optical axis respectively. What is the intensity observed at the output image plane? Right, so at least it's told us the 2 centimeters.

We know where that is. We don't have to worry about working out what 0.3479 degrees is. We know that the first order was at 1 centimeter.

So the second order is where the aperture is. So the second order is going to get through, the plus 2 and the minus 2. And did not appear, then, up there?

Right, so the second harmonic is going to get through. And when you do the square of that, you're going to get the fourth harmonic. So this is going through the mass of that, then. So we've got just these two components. So these have now got 2 centimeters in, rather than 1 centimeter. So that means that you get a 2 multiplied in here that wasn't there before.

So that's where the second harmonic comes in. And then when you do the square, again, you're only going to get the cosine and no constants. And then you do Cos squared and change that into double angles, or fourth angles, actually. And now we're going to see what it looks like.

And now, it's even worse an image of what we wanted. So I guess this is maybe telling us what we shouldn't do, rather than what we should do.

AUDIENCE: Got another question. Is this is still assuming f1 equals f2? And what happens if f2 is smaller than f1?

COLIN SHEPPARD: All it would do, if you changed the relative values of an f1 an f2, all it would do is change the size of the final image. If you think about this system-- let's go back to one where it shows that-- there we are. You can see that the relative distances from the axis only depend on this part of the optical system. And then after the pupil mask, we've already decided, then, what orders get into the final image. And then the final part, all that's going to do is just change the relative size of the final image you see.

Right, now we're going to do another bit more complicated one-- a phase pupil mask. And so what this is shown as is, you can see it's got an opaque part of the screen here. And then in the central part, it's transparent, but it's got some sort of structure there-- some diffractive dielectric element, which changes the relative phase of this order compared with the other orders.

So the question is, what's that going to do to your image? So the rest of the system is completely the same. So this is what we are considering putting in there.

It doesn't matter what the thickness of this is, because all that's going to do is change the relative phase of all of these. So it won't affect how they interfere with each. It will just change the phase of the final image, which you don't see anyway.

But what is important is this bit here. The thickness of this s equals 0.25 microns, and the wavelength was half of 0.5 microns. So that's lambda by 2. So that means it changes the phase by 180 degrees.

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Ah, n minus 1.

AUDIENCE: [INAUDIBLE] the next slide.

COLIN SHEPPARD: Yeah, n minus 1. So it's not going to do that. Take it back.

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: It's [INAUDIBLE]. OK, this is known as a phase pupil mask or a pupil phase mask, or a mask pupil phase.

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Yeah, that's right. So that must be in the next lecture, because it's not in this one. So let's get down to the real application. OK, so George said that I would have just failed the exam because I forgot the minus 1. Of course, this light goes through the glass. But this light goes through the glass and then the air.

So to work out the relative phases of these, you have to look at the relative change in the refractive index times the optical thickness. So this is 0.25 microns, and this is at 0.25 microns divided by the wavelength. And it's going to be multiplied by 1.5 minus 1.

So this is assuming it's glass. So this means that the phase change is pi over 2. So not pi, but pi over 2. OK, so exactly the same grating. These are all the orders.

And so you see, first of all, we get rid of these outer ones. We let through just these three. But of those, we're also going to change the relative phase of this one relative to the other, too.

That doesn't matter, of course, in your sum whether you change the phase of this or you change the phase of the other two. Certainly, the relative phase, which is the important thing. OK, so these are the three conditions, then. For the ones that are not don't get through at all, the strength of those is zero.

Then here, we're taking these ones to be of strength one, and this one in the center to be strength of e to the i phi. And we know that phi is pi over 2. So this is e to the i pi over 2, which is just i. Is that right? Why have we got this i minus 1 then, George?

GEORGE BARBASTATHIS: [INAUDIBLE] So one way to write this pupil function is as a rect that is as big as the opening.

COLIN SHEPPARD: Yeah, OK. Sorry, yeah. So you can think of this as being the whole thing minus-- well, this on top of the whole rectangle. Otherwise, you'd have to do this little rectangle and this bit a rectangle and this bit of rectangle as three things.

So you could do it either of those ways, I guess. But the way that George has done it here, it's the rectangle, of the whole [? lot, ?] it's all of this minus the central one, plus the central one with the phase change. So that's the minus the one that's not there plus the one that is there where the phase changed. And yeah, so now-- well of course, we still get all these delta functions. But now we've got some i's coming in. Sorry?

AUDIENCE: Just one i.

COLIN SHEPPARD: Just one i, yeah. And then these two combine to produce cosine still. I guess if you change the phase of one relative to the other, it'd be more complicated. But these two produce a cosine. So we've got an i plus a cosine.

And then we do the modulus square of that. And the modulus square of that, of course, is you're not going to get a cross product term are you, because these two are in phase quadrature. So when you do the modulus square, you're only going-- is that right? Yeah. You're going to get this squared plus this squared. And this squared, you can write as these two terms when you do the double angle.

And I guess the next one, we're going to see what it looks like. Ta-da, this is what it looks like. First harmonic in field, second harmonic in intensity because of the squaring. And again, the contrast is now much reduced, actually, in this case-- down to 0.25. So I guess this hasn't turned out to be a very [INAUDIBLE] strategy. This lecture seems to be mainly showing you what you shouldn't do when you make your microscope.

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: [LAUGHS]

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: OK, so now we're going to go on to think of this in terms of what we call a point spread function of the low-pass filter. So do you remember, going back to our analogies with electronics, we can think of-- any electronics you can either think of your hi-fi amplifier as being like a black box, don't you? And you put in something and you get out something.

And you can either think of what happens was inside there as being like an h of t, an impulse response. Or you can think of it as being a frequency response. So there are two different ways of thinking about this black box. The same is true in this case.

So up until now, we've been thinking of this in terms of the frequency response, what frequencies get into the image. And now we're going to think of what that means in terms of the point spread function. We've come across the point spread function before, of course. So the point spread function is basically the image of a point object. And our point objects, we make by having a screen with a very small hole in.

So this, then, is going to expand. It reached the lens. It's going to be collimated because we were seeing this on the axis.

This collimated beam is going to arrive at our pupil mask. Part of this collimated beam is going to go through the aperture. It's going to get to a lens. It's then going to be focused down and produce a converging wave to form a spot on the screen. So this is what we think of as-- the final spot is going to be what we call the point spread function h of x, y.

OK, so how do we work out this point spread function? So what we do is, we can again work through this system in terms of Fourier transforms. This is a delta function. So the Fourier transform of a delta function is just a constant.

We multiply the constant by the pupil mask. And then we've got another system here, which does another Fourier transform. So what we're going to get in the final plane, the image of a point object, is going to be the Fourier transform of the pupil mask. And yeah, so it makes the point here that we have to be careful, of course, how we scale things. And of course, the scaling is going to depend on the relative sizes of these f1 and f2 as well, of course.

So, example-- if we got the low pass filter, so 0 is in the center. It lets through the low frequencies, stops the outer frequencies. So this is the pupil mask.

What does the point spread function of that give? And so we're assuming that the input transparency is delta x on the axis. So this is illuminated on-axis.

The field after the pupil mask is just the same as the pupil mask. And then we Fourier transform that, and it's going to give a sinc. That the Fourier transform of rect is sinc. And here, we've got the relative sizes of this as well, so that we know how big it is.

And so that the amplitude point spread function, then, is this Fourier transform rescaled with this x dashed over lambda f. We replace the u by the x [INAUDIBLE] lambda f. And this finally, then, is our answer. Which three times is it saying there? This three, or is that three?

GEORGE BARBASTATHIS: Yeah.

COLIN SHEPPARD: Yeah, OK, because the energy has to be conserved by this thing. So you remember, there's this theorem in Fourier transforms again that says the area under the function squared is equal to the area under the Fourier transform squared. So it's conservation of energy.

And so this is the amplitude point spread function. It's sinc, so it goes negative. Sinc is sine pi x over pi x. So it goes negative.

So that's why it's plotted here. This is 0. It wiggles up and down above the 0 [? light. ?] And so of course, if you looked at the image of a point object, you'd see the modulus square of this. Just like when we looked at our examples of the grating, we finally look at the actual intensity, rather than the amplitude.

The intensity, of course, is always positive. So you'd actually be looking at the square of that. But if you're trying to think of this thing as being-- what we're going to carry on to say, of course, is in Fourier space, or in Fourier domain, the spatial frequencies of the object are multiplied by the mask.

In terms of real space, you're going to get a convolution. The product becomes a convolution. So we're going to convolve our objects with an amplitude point spread function. But we have to do it with the amplitude point spread function, find the final the total amplitude in the image, and then the modulus square to find the final the intensity.

OK, and this is an example of the point spread function of this phase filter we looked at. The pupil mask, this is the same as we had before. So this is its modulus. This is the phase.

And if you calculate the point spread function of this, it comes to this. This is the modulus. This is the phase.

So you can see that that is why we got a pretty horrible image. It's because it's got a very nasty point spread function. [INAUDIBLE] point spread function, for a start, it's got these very large side lobes, which normally we don't like.

We normally like to have our point spread function to be nicely narrow and smooth, and without the big side lobes. And then more than that, it's also got these horrible phase jumps, which I think are very likely to produce a pretty awful image, which is probably what we saw. So this goes through the sum of actually doing that. So this is the same as we had before, then-- the big rectangle minus the little rectangle with the phase.

And then we've got to do the Fourier transform of that. And so the way we've done it here, then, we just got the sum of two rectangles. Each of them becomes a sinc.

But you've got this i here you have to be careful of. And then finally, we work out the modulus square of the part of the point spread function and the phase of the point spread function right so this is the modulus square of the point spread function. So these side lobes are actually quite huge-- quite big.

OK, yeah, that's just looking at the difference between the two, then. So if we have our low-pass filter-- so this is what I was saying about these being huge. You see the low-pass filter, these side lobes are pretty small by comparison. I guess you could use some sort of shaded aperture, or maybe an aperture with some structure that would allow you to make those side lobes even smaller.

This is this process that's sometimes called apodization. Apodization means to cut off the feet. I think if George was giving the talk, he'd say something about the Greek for feet being something to do with pod.

Right, there we are. OK, we are over time, really. [INAUDIBLE] stop now?

AUDIENCE: [INAUDIBLE]

COLIN SHEPPARD: Yeah, maybe that's a good point. I think that's good. You've got the concept of the point spread function.

Now you've got a week to think about the point spread function. And then next time, it will be very clear. OK, so any questions from either side?