GEORGE BARBASTATHIS: OK, good morning. So I'd like to pick up from where we left off last time. So we were discussing the various properties of the Fourier transform and what they mean in terms of the Fraunhofer diffraction patterns.

So what we're doing last time when we ran out of time was we were proving the convolution theorem. So this is the proof here. And just to remind you very briefly, we start with an expression that looks like a convolution. We have an output function equals an input function times a kernel. Then what we did is we wrote each one of those, the input as well as the kernel, we wrote those as Fourier transforms. Here they are respectively.

Then we rearranged the integrals and the order of integration. And we noticed that this expression here is also known as a delta function. The integral of an exponential, of a complex exponential from minus infinity to infinity is a delta function. Oh, too bright. OK, I think that's a little bit better.

I can rewrite this. Let me do it one step at a time. So I can rewrite now this big integral. Unfortunately, the black marker has run out of steam, so I'll switch color. So since we have a double integral with respect to the two frequency variables and a delta function inside, the delta function will knock out one of the integrals. And we simply replace the-- well, let's pick-- we can pick whichever variable we want to integrate, so let's pick u1. So we'll simply get du1 times G sub in of u1 H of u1 one e to the i 2 pi u1 x prime.

And now we realize and remember what this is. If you go back to the beginning of the derivation, this is the outcome of the convolution. So this is actually-- let me write it out again. This is G sub out of x prime. What we've got now is a Fourier transform. Because, you see, here is the Fourier transform kernel. Actually-- I'm sorry-- what we've got is a Fourier integral. We've got to write the outcome of the convolution as a Fourier integral, where this is the Fourier integral kernel or the inverse Fourier transform kernel. The two terms mean the same thing. And this is the actual inverse Fourier transform.

So, therefore, this result is equivalent to G sub out of u equals G sub in of u times H of u. So this is, then, the proof of the convolution theorem, which says that if two functions are related as a convolution with a kernel, then the equivalent relationship in the Fourier domain is actually a product of the Fourier transform of the input function times the Fourier transform of the kernel.

So this mathematical-- this analytical result is what you've seen in the simulations that I have on this screen here, where the two-- I will do this in one dimension. The simulations are obviously 2D, but it's a little bit easier to do everything in 1D here so we don't spend too much time writing.

So the two functions that we have here are actually-- one of them is a sinusoid. So if I have a sinusoid-- I keep using the blue marker, [INAUDIBLE]. So the sinusoid is-- let me just write it with full contrast to save some writing here. This is a sinusoid. Then it's Fourier transform will actually-- as we have said several times, it will consist of three delta functions.

Then we have another function which looks like a rectangle. That's on the right-hand side, top right. And we've already seen this one. So the rectangle is a rectangular function. If it has a width, let's say, a, then the Fourier transform will look like this. And we call this the sinc function. So the height of this function is a by virtue of the scaling theorem. And then the nulls, they go like 1/a, 2/a, and so on. And then, symmetrically, on the negative side, minus 1/a, minus 2/a, and so on and so forth.

So these are two Fourier transform relationships. Now, what we see on the bottom side is basically a grating which is truncated. So I take my infinite grating. It continues on and on. And then I truncate it. That is, I multiply-- I multiply it by a rectangular function which sets the side-- which sets the aperture of the truncation.

So the way I've written it, the width of this boxcar would be a, and the period of the grating itself would be uppercase lambda. So, of course, we can go ahead and compute this Fourier transform analytically, but that would be quite painful. We'd have to do a lot of changes of variable and so on. However, by virtue of the convolution theorem, because this is a product, this Fourier transform will be a convolution.

So, basically, what we have to do in order-- when we multiply these two results in this domain, in this domain we have to convolve them. So what is now-- so the assumption here is that the size of the boxcar is actually bigger by the size of the period. It does not have to be much bigger, but it's nice if it is bigger for what I'm about to draw to be correct.

So what we get if we can convolve this pattern with the three delta functions, each delta will produce a replica of this function centered at the location of the delta function. So, therefore, we'll get three replicas of this sinc pattern centered at the corresponding locations. So let me draw this a little bit carefully here. So I'll try to be as accurate as I can.

OK. So here are the two-- well, there's three delta functions. One is at the origin. One is at minus 1 up on lambda, and the other is at plus 1 up on lambda. And then, on each one of those, I'm going to center one of the sinc functions. So here is one sinc function, and here is another sinc function, and another one.

So I tried to do it carefully. Why is this taller? Because, if you recall, the delta function that was at the origin is actually twice the size of the other two delta functions. So the size of this one would be a/2 actually, because, if you recall, we're also picking up one factor of a from the sinc itself. Then the size of this would be a/4, and the size of this would also be a/4. When I say the size, I mean the height.

And then where are the nulls located? Well, these nulls are simple. They're still at 1/a, minus 1/a. Where is this null? For example, this null would be at 1 over lambda plus 1/a. And then I have another null over here at 1 over lambda minus 1/a, and so on and so forth. And then I have more nulls at 1 over lambda minus 2/a, minus 3/a, and so on and so forth.

And now my assumption that I have several periods of the grating with the boxcar, you can see it is quite convenient because I can draw these sinc functions and I can kind of ignore what happens in between over here. As you can imagine, something complicated will happen over here because the sincs are actually adding coherently, so something funky will happen. But if they're pretty far apart, you can see that the envelope of the sinc function actually decays quite fast. So I can more or less ignore what is happening over here, and I can simply draw the sinc function. Of course, this is because I'm doing it by hand. If I use a computational tool, such as MATLAB or Mathematical, it will do it for me.

So this is the one-dimensional calculation, that is easier to do by hand. As a bonus, I think you guys have already computed-- or, if you haven't already, you will compute some time between Tuesday midnight and Wednesday at 8:00 AM. Some time you will compute the same convolution for the two-dimensional case. And you can see the result here. So, in effect, this saves you from-- well, at least you know if your result is correct or wrong. If you derive something that looks like this, then you know you've done correctly.

So you see what's happened. If you had the infinite grating, the infinite grating gives you three delta functions whose axis is kind of perpendicular to the fringes of the grating, and the one that the origin is stronger. Of course, the rectangular function produces a sinc pattern. And then when you convolve them, when you have a tilted grating now and multiply it by a rect, then you see that you get three sinc patterns oriented along the same axis, which is perpendicular to the fringes of the grating. But the sinc pattern itself is actually perpendicular to the rectangular aperture. So this is what the convolution theorem says.

And, of course, you can turn it around. You can actually turn the grating around and rotate the aperture. In this case now, again, you will see the sincs oriented parallel to the fringes again, but now the rectangular pattern itself is rotated, and therefore the sinc patterns themselves are rotated. So the reason we use these properties is because-- again, as you can imagine, if you were to write this as an explicit integral, it can be done and you would-- of course you would still get the correct result, but it would be quite painful to compute.

So now that we have all these results, we can actually apply them. So far, these were mathematically results. I could pretend I'm teaching 18085, or whatever it is that you learn those things. But this is also-- of course, they actually become Fraunhofer diffraction patterns if you simply perform a scale-- a coordinate change. So the Fourier transforms, we compute them with respect to the frequency variables, u and v we call them. Goodman calls them F sub x, F sub y, but it is actually the same variable.

Well, if you substitute a special variable and you multiply by the scaling factor lambda z, that we saw the derivation in the previous lecture-- I'm not going to do it again-- then you actually get the Fraunhofer diffraction pattern. So it's a simple scaling argument. And you can tell that-- well, at least it is plausibly correct. Here, the units are correct, because u and v are frequency variables. So, therefore, their units are inverse meters. And then multiply it by meter squared. The wavelength is meter, the distance is meter. So, therefore, I get meters. So the units are correct.

And this is one example that we did. Then here's another one where I shrink the rectangular aperture. And, of course, the Fraunhofer diffraction pattern will grow. We call this a similarity or scaling theorem. Then I can, for example, use the shift theorem to calculate the Fraunhofer diffraction pattern from three rectangles. As we discussed the last time, this will give rise to a sinusoidal modulation in the Fraunhofer domain. And you can also do the convolution theorem in this case with a truncated aperture, and so on and so forth.

So this basically concludes the discussion of the Fraunhofer diffraction pattern. But since we've been discussing Fourier transforms, there's another very basic property of Fourier transforms that I would like to introduce here, and then we will see it in full glory for the next two lectures, and that is spatial filtering. So spatial filtering is basically the following. It says if you go to this Fraunhofer domain, or, in general, in the transform domain-- which, we will see a little bit later, that we don't need to have to go very far actually. By using a lens, we can produce a Fraunhofer diffraction pattern at the back focal plane of the lens. That's very convenient.

But if I go here and I do some modification, and then take another Fourier transform, then, of course, the signal I reconstruct is not identical to my original signal, but it will have changed because I've modified the frequency spectrum. So this is called spatial filtering. So here's an example that I have constructed. So in this case, I've contacted the signal in the space domain that looks like three sinusoids.

Now, you cannot tell very clearly from this pattern that you have three sinusoids, but if you take the Fourier transform, then you see three spots here, three dots. Actually, you see six. But you recall that each sinusoid corresponds to two dots. So the conjugate dots here, this one and this one, they are one sinusoid. This and this one, another sinusoid. This and this one are yet another sinusoid.

So spatial filtering, a very simple occurrence of spatial filtering is what happens if, for example, you go in with some black marker or some opaque screen in the case of optics and you remove one of these dots. If you do that in the transform domain, then you will see it. Now, watch as I transition the slide. You will see that the spatial pattern also changes.

OK. So now it becomes kind of horizontal, and it is horizontal because the two dominant-- the two dominant sinusoids are actually along the horizontal axis. So, therefore, your grooves are-- I'm sorry. Your grooves are vertical because the two dots here are on the horizontal axis. But there's also a weaker sinusoid that gives rise to these weak diagonal fringes that you see over here.

But, basically, you can see that one of the three spatial frequencies has vanished here. So this is the simplest case of spatial filtering. And, of course, you can generalize it. Here's again the Red Sox-- or I should say the GO SOX pattern on the Boston high rise that I showed last time. And this is the spatial frequency representation or the Fourier transform of this pattern. And then we can apply various filters. For example, if I go with a filter and I block all the high frequencies, then you can see that my pattern appears blurred. In fact, it is more than blurred. The windows have kind of disappeared of the building. And that is because the windows, if I go back, you will see that the-- hi, Colin. You're back.

COLIN: Yes.

GEORGE BARBASTATHIS: Welcome.

COLIN: Sorry I'm late.

GEORGE BARBASTATHIS: Oh, no problem. I thought you were still in Poland.

COLIN: No, I got back.

GEORGE BARBASTATHIS: Oh. Oh, OK, welcome back. And so the windows, if you recall, the windows are kind of periodic in this high rise here. So they correspond to these dots in the frequency domain, kind of like delta functions. And because, in this case, I have blocked the dots, you see that the windows disappear from the high rise. And, of course, you can do other funky kind of filters. This is called a band pass filter. And, in this case, the windows reappear because now I center this doughnut, this annulus. I centered it so that some of these dots in the frequency domain, they survive.

And you can see that, of course, the-- well, it is not fully reconstructed, the original building, of course, because there's still spatial frequencies missing. But you can see that the pattern of windows of the high rise has kind of reappeared. Now, what happened to the sign GO SOX? It vanished, and it vanished because, in this case, I have blocked the lower frequencies.

The GO SOX sign, it has survived the low pass filter because this is a relatively slow-varying signal, right? So its frequency content, you expect it to be centered-- I'm sorry-- to be concentrated near the center of the Fourier domain where the frequencies are low. When we do the band pass filter, the GO SOX vanishes. And that is because I blocked the low frequencies where this signal was represented.

So you can see that you can do quite interesting manipulations on images using this concept of spatial frequency. And, actually, the GO SOX signal has not quite vanished. If you look carefully, there is a little bit of evidence of it here, but it's quite hard to see. And, of course, there's a little bit of evidence because there is a little bit of the frequency content leaking into the intermediate frequencies. So, therefore, some of it has survived, but mostly it is gone. So that is the-- so it is not a perfect filter, but it works quite well.

Of course, the other thing that vanished is the average. You can see that the sky, that used to be kind of an average gray, it is gone also. Because the average-- of course, the average is presented at the zero spatial frequency, and I blocked it. So, therefore, the average is gone. So this is called a spatial filter.

OK. So we're still one lecture behind. So this is what I was supposed to have done last Wednesday. And if you look ahead, there is some discussion of the transfer function of a Fresnel propagation, and then something called the Talbot effect. So I will not do this right now. I will postpone it, if I may, for next week. What I would like to do is I would like to switch to the lecture that I posted today online, and that is Lecture 10A.

The reason I'm doing that-- I will go back and talk about the Talbot effect. Don't worry. But the reason I'm doing that now is because I would like to press on with the concept of spatial frequencies and spatial filtering. Because it is quite an important one, and I think the sooner you learn it, the better. Talbot effect, well you can learn later. But this business of spatial filtering, in my experience, it takes quite a bit of time to digest, so I would like to do it sooner rather than later.

So I already alluded to that. I said that this Fraunhofer diffraction pattern is a Fourier transform, but we don't have to go to infinity to watch the Fraunhofer diffraction pattern if we want to generate a Fourier transform optically. We can also do it by using a lens. So this is what I will do today. For the rest of the lecture, I will show you how a lens can produce a spatial Fourier transform, and what can we do with it.

So, very briefly, to remind you, from geometrical optics, this is-- we did this some time ago, maybe about a month ago. So, to remind you, a lens is a device that looks like a, well, at least one curved glass surface, typically more than one. And it is a device that we can use to focus or collimate light. So, for example, if you illuminate a lens with a plane wave, then the lens will focus that plane wave at one focal distance to the right. On the other hand, if you place a point source at one focal distance to the left, the lens will collimate it, will produce a plane wave, which we also refer to as an image at infinity.

And, of course, what I am discussing here is for the case of a positive spherical lens. There's other lenses that we saw, negative lenses that would do something slightly different. But I don't want to do a full review of lenses here, just to remind you what is relevant to our discussion here.

And, of course, the other thing that lenses do is they can produce images as finite conjugates. If you place an object at some distance s sub o, then the lens will form an image at a distance s sub i, which is related to s sub o by the lens law. So we did the stuff to that when we did geometrical optics, so I don't want to produce recurrent nightmares to you by repeating it here.

So what I will do now is I will describe the lens in the context of wave optics. So, of course, in the context of wave optics, we have to describe the lens as some kind of a transparency, as some kind of a phase function that is applied to the optical field. So I don't want to go into the details of this one. Is there a question, or-- someone has a microphone on, so I can hear you shifting on your seat. Anyway, it doesn't matter, though. If you have a question, please interrupt me. Of course, this applies always.

So what is this now? So we will do a very crude approximation here. We will actually neglect the thickness of the lens. We did this again when we did geometrical optics. And that is, of course, because it is not, strictly speaking, correct, but the results that we get are pretty good, and it makes our mathematics pretty simple. So the combination of the two is a good justification to make an approximation. If you have one of the two reasons, it is not good to make an approximation.

For example, if it makes your math simple but the answer is wrong, then you don't do the approximation. If you get the correct answer but the math is not simplified, again we don't make the approximation. You might as well go with the accurate calculation. But, in this case, we get two bonuses. And if we have both bonuses, then we do the approximation.

So what is happening here, if you take a field-- imagine like Huygens wavelets impinging on the lens from the left. The wavelet that impinges in the center will actually see the thickest part of the lens, so it will sustain the longer phase delay because it propagates through glass. If you take a Huygens wavelet that actually impinges away from the axis, it will see a thinner portion of the lens. Therefore, it will have less phase delay because it propagates a shorter distance in glass. It will still propagate some distance in air on the left and the right of the lens.

So if you compute that now, the difference is, of course, it is given by the spherical calculation. I don't want to go through this. You can go back and do it yourselves. It's a very simple geometrical calculation, with the addition of the paraxial approximation. So even if you glance at this here, you see that-- actually, this I copied from Goodman, so the equations are verbatim from Goodman's book. It's a scan, actually. I'm sorry. It is not a scan. It is a scan from my own notes from last year. But, anyway, it is verbatim copied from Goodman.

And you can see that I replaced the square root with a Taylor expansion. So it is a paraxial approximation. And the result that you get, which is what I really wanted to do, is something that looks like this. OK. This is what we get. So we'll get the complex amplitude transmittance of the lens if you express it in wave optics. It is actually a quadratic phase delay. That's if.

And in this quadratic phase delay, a magical distance happens. A magical distance appears, which we recognize to be the focal length. This, if we recall from geometrical optics, we used to call this the lensmaker's formula. So, basically, we recover the expression for the focal length of the lens, but now we have a wave-- I don't want to say wave function. That sounds like quantum mechanics. But now we have actually-- we have a complex amplitude transmittance. That's what we have associated with that one.

So now let's see why this is the same result as we had before. So the trick here is that we replace the lens-- when we have a situation like this one, we replace the lens with its amplitude-- complex amplitude transmittance. So forget about the curvature. Forget about the glass and everything else. We just replace it with a thin transparency.

And then we illuminate it with something. Let's start by choosing this something to be a plane wave. So here's the wave vector of this plane wave. And because it is propagating at an angle, we write it as g sub in. Actually, we write it g sub t sub minus, because it is the field immediately to the left of the transparency, of x comma y. It is a plane wave. Let's call this angle theta-- what did they call it? Theta 0.

So since it is a plane wave, the proper expression for it is e to i 2 pi sine theta 0 up on lambda x plus cosine theta 0 up on lambda z. And I'm going to do two things here. First of all, I'm going to place the transparency at z equal 0. So that knocks out this factor, because equal 0. And then I'm going to make the paraxial approximation, so that knocks out the sine. So, basically, the field incident upon the transparency, upon the lens, that this, is simply e to the i 2 pi theta 0 x up on lambda.

So what I would like to do now is to compute the field after the transparency, g sub t plus. So the rule that we described when we did thin transparencies is that we multiply. I'm sorry. I'm using a slightly different notation on the whiteboard than in the notes. You don't have a minus and a plus, but it's basically the same thing. So g sub t plus multiplied by the transparency itself. So the plus stands for after, the minus stands for before, and the nothing stands for the transparency itself.

Oh, and another thing that I did in the notes is I defined-- I defined this quantity u0 equals theta 0 up on lambda. And this now we recognize as a spatial frequency because it has units of inverse meters. So what do we get now? What do we do is a little bit of algebra, but it actually results in an physically intuitive, physically meaningful result. So that justifies the algebra, I suppose.

So let me write it out. Let me write out this product over here. OK, that's it. So now I have to do something that-- let me leave it here so you can see. We have to do something that you may remember from horror, from your high school or elementary school. I don't know where you learned these things. It's called to complete the square.

What is the square that I want to complete? If you look at the exponents, you have an expression that looks like this. I'll knock out the minus sign. x square over lambda f minus 2 u0 x. Can you see that? I've neglected the y's and everything else, and the pi's, and so on. If I can do that, then I can take care of the rest.

So how can I complete the sign? Well, I tend to get confused with this, so let me knock out the 1 over lambda f also. Now it looks better. Now I can do it. Basically, to complete the sign, I have to add and subtract the square of this business here. And now I can write it as--

OK, now I can go back and substitute into my original expression. I'm done with manipulating the exponent. And my original expression was this one. So I can rewrite out now, g sub t plus of x comma y equals-- first of all, I have this constant term. That is constant. It means it does not depend on x, the spatial variable. So I'll just take it out. I should not forget my pi's. So there's a pi over here.

So-- no, I don't like the way this came out. I was expecting to divide, but I had multiplied. OK, that looks better. So I'm doing OK here, because what are the units? No units. I have a spatial frequency squared times distance squared, so no units. And what I have left is something that looks like this now. OK.

So the first part, I don't have to worry too much about. This is just a constant factor, as I said. But this one, now I recognize that's a spherical wave. It is a spherical wave because it contains quadratic phases in the exponent. It is a converging spherical wave because of the minus sign here. And it's not quite its origin but its sink. The location where the spherical wave becomes a point is actually shifted by this factor over here. This is displacement. OK.

So, basically, this is what I've got. I've got a spherical wave which converges. Oh, and where does it converge? Well, the distance that a spherical wave converges is what multiplies the wavelength in the denominator. So this is where it converges. So this is basically what you see here. The spherical wave after the lens converges to a distance u0 lambda f. If you substitute the definition for u0, it is theta 0 times f away from the axis. And the distance between the lens and the focus is one focal distance.

So this is not news. We knew this from geometrical optics. We just rederived it using the thin transparency. So this, I guess, gives us conviction that our approach is correct, because we rederived something from geometrical optics. I will not do the next one. The next one, I'll let you do by yourselves. You can repeat a similar procedure of completing squares in order to see what happens to a diverging spherical wave placed at one focal distance to the left of the lens. And you can convince yourselves easily that this becomes a plane wave propagating at an angle equal to the ratio of the displacement over the focal length. So this is again something that we saw in geometrical optics. It is not new. OK.

The real result that I want to derive here-- and I will try to do it carefully in the time that we have left-- is the Fourier transform property which I will do for a special case. Actually, I will not do it for a special case. I'm going to-- I take it back. I will do it for the general case of a lens-- I'm sorry-- of a thin transparency placed at a distance z to the left of the lens. Goodman does three cases. First, he does the case z equals 0. Then he does the case z equals f, and then another case. It doesn't matter. We'll just do it for the general case, and we're covered.

So what-- first of all, let me do the derivation in one variable. So don't write too much. So we'll basically skip y. All of the derivations will be just with x. So I have a thin transparency g of x. And then what I will do is I will propagate it distance z to the lens. Now, on the lens, my coordinate is x prime.

And since I'm propagating a field-- also, I forgot to say-- that's quite important-- the implicit assumption here is that the illumination is an on-axis plane wave coming from the left. So that, if you recall, we said a couple of times, that is simply-- its complex amplitude is 1, because I can choose that-- a constant phase. And there's no x variations with this one.

So the Fresnel propagation kernel, if I go from z-- So g is-- let me maintain my notation consistent here. So to the left of the lens, L minus of x prime-- so that is the field to the left of the lens-- is going to be given by a Fresnel diffraction integral.

And, actually, in my derivation, I skipped the constant. And what is the constant that I skipped? It is this one. And this constant should be there, but it is not doing anything significant for us in this case, so that's why I skipped it. To save writing, basically. So from now on, we will basically neglect this. Even though it is there, we will simply neglect.

Now, the field after the lens equals the field before the lens times the lens itself. And the lens itself is something of the form e to the minus i pi x squared up on lambda f. And, finally, I have this field, and I have to propagate it. How long? Now I have to do this part, which means I have to propagate by distance f until I reach the back focal plane. And that is another Fresnel integral. I will call it g sub f, I guess. Again, there is a factor here which I will neglect. OK.

So now let's put everything together. I have two Fresnel convolution integrals-- one with respect to the input coordinates, one with respect to the lens coordinates. And what is left inside, I will simply substitute all the rest. I have the input itself. Then I have the propagation kernel from the input to the lens. Then I have the lens. And then I have the propagation kernel from the lens to the back focal plane. OK. That's what it is.

This looks a little scary, but part of the purpose of this class is to teach you how to not be scared by this kind of integrals. So the way you know this, you don't get scared by this sort of integral is you manipulate the exponents here. And you try to-- the first thing you do when you reach an integral of this kind is you try to collect terms. So I'll write the exponents here. I will expand the exponents and collect terms then.

So if I expand the exponents, I will get x prime squared plus x squared minus 2x x prime over lambda z. This came from this term over here. Then I have minus x square over lambda f. And then I have plus-- and I think I missed a prime here. That should have been x prime. Yes, correct. The lens should also be with an x prime. Thank you. It is very fortunate that you corrected me, because if you hadn't I would be kind of stuck here. OK.

So the thing that you notice first is that some of these exponents get knocked out. This one kills this one. That's very pleasant. Now what do I have to do? I still need to make an integration with respect to x prime. So x prime appears here and here. And I have to make an integration with respect to x. Well, here is x. Here is x. OK.

So what-- any ideas? Anybody want to speculate on what I should do here? Let me write the integral. That's a little bit confusing the way it is. Let me rewrite it so you can see what the integral looks like.

Let me do it carefully. So I have e to the minus i 2 pi x prime up on lambda. That's common in the two exponents. Inside, I have x up on z plus x double prime up on f. So what should I do next? Is there any glaring sort of integral that popped up here? What is-- here, we have an integral of another exponential, right? And they-- let me rewrite it like this.

OK. So the glaring integral that I was referring to before is this one. That's a Fourier transform. Whose Fourier transform? The Fourier transform of whomever appears in this location over here. And where is the Fourier transform computed? Well, it is computed in this spatial frequency, right? It is computed in whatever spatial frequency multiplies the dummy variable in the exponent.

Now, what is this Fourier transform? We don't know, but we have our notes, or we have Goodman, or we have the tables of formulas. So switching to Lecture 9B. This is our Fourier transform pairs. I recognize this integral, recognize this Fourier transform. It is the second row from the bottom. If you look at this expression and the expression over here, it is actually the same Fourier transform. It is the Fourier trans-- what you see here is the Fourier transform of the quadratic phase in the exponential.

So we have the answer. The answer is right here. Again, I will neglect-- actually, this constant is quite important, but I will be neglect it nevertheless. So, basically, the way to get a one-to-one correspondence is to simply substitute what would-- what in the table is denoted as a square is actually identical to 1 over lambda z in our case.

So I can write out now the answer. This thing equals-- first of all, before I do any further, we recognize that this does not play in the integration. This has the output plane coordinate, so I will simply take it outside. And then I will write out the-- in one shot, I will write out the outcome of this Fourier transform.

So that a square that I have in the original function, it will go inverse in the other one. So we'll get, then, e to the what? I will get an extra minus sign. If I have plus j here, I have minus j here. So this will then become e to the minus-- we'll have all the pi's and so on-- minus i pi lambda z. And then I will get the square of the spatial frequency. So it will be 1 over lambda square x up on z plus x double prime up on f squared. OK, that's it.

So now I can manipulate it a little bit further. And now let me write out all these exponents that come out of this square. So I will get-- the first one will be x square up on z square. So one z will be killed. One lambda has already been killed. So I will get e to the minus i pi x square over lambda z, right?

Then I will get this term. That will be e to the minus i pi-- this is tricky-- z x double prime square over lambda f square. This came out of the square of this one. And I will also get the cross term. So that will be e to the minus i. And then we'll do it carefully. I will get 2 pi. And what is left here-- one z will cancel. I will get x x double prime up on lambda f. And now, happily, we see that this additional quadratic that was very annoying over here, this one, it got killed by this one. This one is not playing in the integration either, so I can actually take it out of here.

And this is the result that I was after. You see that I actually got another Fourier transform. This is well recognizable as a Fourier transform kernel. So what I have in this part over here, it is actually the Fourier transform of the transparency calculated at these coordinates, that is the argument of the integral.

Now, something funny happened here, and this doesn't look quite right to me. That should be f, right? OK. I don't know how this became z. Oh, yes, yes, yes. OK, I know now, yes. That should be f, yes. Somewhere in my notes I converted this to f, but, thankfully, not the physics. So this should not be z. This should be f. OK, so now it looks right. OK.