GUEST SPEAKER: Good morning. Today we're going to start with a demo, and we're going to show experimentally some of the principles we've learned about 4F systems in the past classes, and actually we're going to see a comparison of an incoherent versus a coherent 4F system when imaging resolution target. So but since now we have all the tools to understand, finally, the set up here, let me just describe it. I've been showing some of these components over and over the past demos, but I guess now we can understand it, so let me try to focus the camera.

So let me start with this is a green laser. So the first component that we have here is [? anti ?] filter that its job-- thank you-- so this component, its job is to control the intensity of light so we can attenuate it in this case for experiment. This is an interesting component called a spatial filter that is composed of a microscope objective and a pinhole, and now that we know what Fourier transforms are, the microscope objective is essentially taking the Fourier transform of the laser beam, which ideally is a plane wave, so it should be a delta function. However, since this plane wave is not perfect, it has high frequency noise.

So then the pinhole, what it does is it does a low-pass spatial filtering. So it basically removes all these high frequency signal, and after these we have a nice spherical wave coming out. This component here is an iris, control the diameter of this spherical wave, and this element here is my collimating lens that will transform this spherical wave into a plane wave. In this case, this is just a beam splitter that will allow me to couple from this side, where I have a white light source. In this white light source, we have a green filter, so then we only have green light coming out.

We have an iris again to control the diameter. Right now I'm blocking it, so it is a beam block. And these are collimating lens. So here it should be the planar equivalent of the incoherent side. This is the planar equivalent of the coherent side. This mirror here bends the light 90 degrees to send it to the side from here. This is the one line of the incoherent line, and then should go into the common path of the coherent line. Then they basically bend again 90 degrees, and this is the 4F system section.

So here we have first an element, which is the object that we want to image. The object is called the resolution target, and it's a clear substrate, so a piece of glass that has patterned chrome lines of smaller and smaller sizes that we're going to see. Then we have our first lens of the 4F system. This is a non-magnifying 4F system, so these two lenses are exactly the same focal length.

So first lens. Then this is an aperture in the Fourier plane because we're going to start with low-pass filtering. The second lens, and then we're imaging this into our detector, which uses CCD seamless array.

So now we're going to switch to the camera and start with a coherent case. So the first case, it's like what we've seen in class. We're going to illuminate this object with an inline plane wave. OK, let me just zoom in a little bit.

So this is the resolution target, and it's in imaging condition, right? So the object is in the front focal plane of the first lens, and the detector is in the back focal plane of the second lens. So what you can see here is we are actually imaging all the way to groups two and three. This resolution target has lines even larger, which are one and two that we don't see here. And the first thing we see is the following. We see some fringes, which are produced by the interference from the glass, so it's basically the reflection of the plane waves interfering with each other.

So this is very typical in coherent imaging. We see other rings that are produced by dust particles maybe along the way. So all these fringes right now look destructive to the image, and perhaps they are for certain applications. However, for those of you doing digital holography, you notice that these fringes contain a lot of information that is encoded about the amplitude and the phase of the signal that we can exploit.

So just for comparison, I'm going to zoom in and see what is the maximum that we can see here. So you can see that we can have a resolution now that we know what the optimal resolution. With good accuracy we see maybe down to this side of the group number five.

All right, so in the homework we learned about the low-pass filtering in Fourier domain. So just to remind you, essentially what we're going to do is that the first lens is taking the Fourier transform, and it happens in the Fourier plane. So then in that plane we put an aperture, centered, and then what we do is we're going to start reducing the diameter of the aperture to try to low-pass filter the signal more and more. So what do we expect to see here? Well, you did it in the p-set, so we expect to see some blurring, progressive blurring.

So I'm going to do it, and hopefully you can see it from there. So right now it's all open, and I'm going to blur it. So I'm going to do it again. This is clear. Blurry. So I'm controlling with this iris. And you can see how the signal that contains the higher spatial frequency in this area goes first, then this one and this one. But all of them now-- in this blurry condition, if I try to zoom back again, now we see how the signal is really blurry. If I open the aperture, it becomes clear again.

So now let's do the comparison of the exact same thing. What would you think it would happen if I switch illumination now to incoherent light? So I'm going to block the coherent light here, and then the first thing that we notice is that we don't have those fringes anymore, right? So it looks like a very clean image. So again, in this case, you can think about that for microscopy, this is very nice. However, from the digital holographic point of view, we don't have that extra information that we can use in encoding of the signal. Now let me zoom in to compare the resolution in this case.

So you can see it's hard to tell if we don't do a quantitative study to see exactly what is the smallest line that we can see in each case, but we can see the comparison again versus the resolution in the coherent and incoherent cases. And now let's do the low-pass filtering. So again, this aperture right now is fully open, and then I started closing it. And I'm closing and closing and closing it. Open again. Close it.

So we see two things. Remember that the incoherent system is basically regulated by the OTF, as opposed to the ATF in the coherent system. So now that you can think of it-- the easy terms is it's a [? combo ?] because the OTF is a normalized cross correlation of the ATF, which means that when I close this aperture, you can see how the signal becomes weaker and weaker because essentially I'm restricting the DC component that happens to be in all the signal there, but also gets blurry. It's hard to see, but it's getting blurry. It is getting blurry at a different rate as the other system.

So now let's go to high-pass filter, and I'm going to switch back to the coherent case. So this is the coherent case. So right now what I'm going to do is that I'm going to remove the aperture in the Fourier plane, and I just did a very improvised high-pass filter, which basically is a microscope slide that has a little piece of tape in the center just to block the basic component, or kind of, but anyways, it works.

Somewhere around here-- let me crank up the light. And we can see how basically we remove the DC term. And you see the high-pass filtering effect, which is essentially accentuating all the edges of the signal.

Now I'm going to put the incoherent case and the coherent superimposed just to show edge detection. So this is a normal thing. We can accentuate the edges. So what is interesting also here is that if you move the filter in different places, you can, say, just accentuate the edges in the horizontal direction or in the vertical direction, et cetera, et cetera. We see that the incoherent case, we don't see the same effect, which is as expected. Again, it's a normalized cross correlation. But with a coherent case, we do see these drastic change.

GEORGE BARBASTATHIS: Thank you. Any questions about the demo or about coherence, incoherence?

AUDIENCE: How do you tell whether something's partially incoherent or completely incoherent?

GEORGE BARBASTATHIS: OK, so the test is with one of these interferometers that Se Baek showed last time. You basically measure the fringe contrast. For example, if you're worried about spatial coherence, you set up a [INAUDIBLE] interferometer [INAUDIBLE], and the fringe visibility will tell you the degree of coherence. If it is full contrast, it is fully coherent. And of course, you never see full contrast. You never see-- it's always somewhere in between.

OK, so today we'll have a flashback of geometrical optics. I will go back to something we saw in geometrical optics, and that's the single lens imaging system. Oh, and before I forget, on Wednesday, it is the project presentation for the graduate version of the class, so I wanted to remind you. I will also send an email from the website, but if you haven't done so already, you should have a meeting with your project mentors, finalize your presentations, and then on Wednesday, what we'll do is it'll be a conference style. Each group will have about 15 minutes, approximately 12 for the presentation, approximately 3 for questions. How many groups? I think we have six or seven groups. So typically, these things run over, so we'll take probably the full two hours with the presentations.

So for those of you in the 2.71, you don't have to attend, but I think it would be very useful if you do because it's sort of interesting. It is more advanced material. The graduate students will present some topics from current literature, and you'll get an idea not only over the basic optics that we'll do here but also what kind of things happen in the research front. For example, last week several of us-- Pat, Yuan, and Sean-- were at the conference in Vancouver, so some of the topics that were represented on Wednesday were actually big items in the conference. The conference was actually called Advances in Imaging, this spring congress of the optical society. Anyway.

So going back to the lecture, this is something that we've seen before, but with geometrical optics. Now we try to understand how the same system works but with the wave optics. And if I could have the tablet again, please. So the way we analyze the system, of course, is we'll basically have to propagate using Fresnel diffraction to propagate the field from one element to the next until we reach the end of the system.

So for example, at the input we have the input field. I will write it in one [INAUDIBLE] so we don't write too much. The input field is simply the product of the illumination field times the complex transmittivity of the input transparency. I will do everything for a spatially coherent illumination, and we know now that if the illumination is spatially incoherent, if we have already solved the coherent case, we're in good shape because we can simply take the modular square of the point spread function and the autocorrelation of the coherent amplitude transfer function, and then we can compute the correspondent quantities for the incoherent case. So the coherent is the one that we'll always start with.

OK, so doing that now, we can calculate, for example, the field immediately to the left of the lens. So let's call it g sub L minus for left, and I think I still use x double prime as my coordinate. So this is now a Fresnel propagated version of the input field, so we'll write out the big, bad Fresnel integral. e the i 2 pi z1 over lambda i lambda z1. So this is the [? prequel, ?] and then we write our big integral here, g sub in of x, and then the quadratic.

And of course, two things happen at the lens plane. First of all is the lens itself. That will impose an additional quadratic phase delay upon the signal, but now it's a little bit different than the geometrical optics case. I forgot to mention earlier that I have added a thin transparency in front of the lens.

So this is again a pupil mask. As we will see in a moment, it does pretty much the same job as the pupil mask did at the Fourier plane of the 4F system, but in this case, we actually put it sort of in contact with the lens. In fact, we don't have to. We can put the mask anywhere we like, but anyway, it makes the math a bit simpler in this case. And in most implementations, for practical reasons, they also do it this way because it is much better to have everything stacked together, as opposed to having different elements sort of floating around in an optical system.

OK, so this means then that the field to the right hand side of the lens, so the plus now means to the right, still x double prime coordinates, is going to equal to the field that was on the left multiplied by 2. Two functions now. One is the thin transparency itself, which I denote there as g sub pm. pm, of course, stands for pupil mask, like in the previous cases. And then times another complex quadratic exponential that corresponds to the lens itself. So this would be e to the minus i pi x double prime square upon lambda f, where f is the focal length of the lens. And finally, the field at the output.

Now the output coordinate is x prime. It's going to equal another Fresnel propagation from the pupil mask to the output plain, so I have to write another propagation term here. This would be z2 this time. And what I will put now is gL plus because this is the field to the right of the right of the composite lens and the pupil mask, so that would be x double prime and then another quadratic exponential.

OK, so this is it, and when you plug in all this math-- this is something we've done already, so I will not do it in the class again, but there is a supplement. There is a scan of my handwritten notes in the website where you will see a little bit more detail on how I did the derivation here, but I will skip some of the immediate steps, and I will show the result here. So the way you handle this kind of-- as you imagine, when you combine all of these integrals, you will get a sort of integral that involves all of these functions, g sub in, g sub pm and so on, but also you will have a bunch of quadratic exponentials.

So what you do is you expand the exponents of these exponentials, and you rearrange terms, and you write it out this way. So these are the output coordinates. They don't participate in the integration, so we can knock them out of the integral. And then of course, we cannot do much about this thing because, well, we don't know what these functions are yet, but we can certainly try to see what we can do with the remainder.

So what is sort of glaring here is that in the exponent, you have quadratic terms, and you also have linear terms. The linear terms, we know from experience that they lead into Fourier transform, so they're nice. We'll know what to do with them. The quadratic terms, they usually correspond to nasty things like the focus, but in this case, you can see that the quadratic term is multiplied by a coefficient here, which allows us to knock it out if we select the distance as z1 and z2 judiciously, and of course that is the result of the lens law. This is the same equation that we derived using geometrical optics, and we'll call it the imaging condition. So you can see here that when you satisfied this imaging condition, then you essentially knock out the focus term from the diffraction integral. So that is gratifying because we actually derived a result from geometrical optics, but this time it is perhaps more rigorous because it is based on electromagnetics.

What is the rest now? Well, we'll deal with the rest in a moment, but you see that there's another quadratic here that we cannot quite so easily get rid of, and this extra quadratic that popped up there is interesting because it is not predicted by geometrical optics. So this is something that-- well, it will happen only with spatially coherent illumination, and geometrical optics cannot predict it and actually cannot deal with it. And in fact, not so long ago, about the time I was born in the early '70s, this was a topic of research, what to do about this term. In fact, Professor Goodman, who wrote your textbook, one of his early results as a young scientist, he wrote the famous paper about this term, what this term means. And he wrote it-- I believe he wrote it in 1971, exactly the year I was born. I'm not that old, so I guess that you can still call it a relatively recent development.

So again, we're talking about this term over here, what to do about this one. And then of course, there's an entire section in Goodman that talks about this term, and I would encourage you to read the section. I will show a summary here. Something got messed up that with my animations here. This was not supposed to be here, but anyway. So Goodman actually goes over three methods that you can use to eliminate this unwanted quadratic. One of them is you lay out the input transparency as a sphere, or in a spherical surface, instead of the typical planar surface. That may sound a little bit strange, and it's probably very difficult to implement in practice, but conceptually you can see what's going on because now these points in the transparency start with an original phase delay, you can work out the curvature right here. In fact, the curvature would have to be exactly z1, the areas of curvature, so that it cancels this unwanted quadratic. So that's one way.

The second way is to actually stick a lens in front of the transparency. This is actually very often done in geometrical optics as well. This is called the condenser lens. So for example, we don't have this anymore, but actually, we do have one in the classroom. We used to use these projectors. Basically, if you look at the top surface of this projector-- I don't know if I can take it out here-- but if you look at it, it's actually a lens. After class, come over and look. There is grooves on it, and basically, sort of the grooves implement equivalent of a lens surface. If you do the operation modulo 2 pi on the surface of the lens, instead of having this surface that we typically think of as a lens, if you do a modulo 2 pi operation, you will get something like this, of course.

OK, so that's a way to make a thin lens. That is called a condenser. The reason we put it in a projector is not to eliminate this quadratic factor, it is simply because it makes the illumination uniform. So when you project the slide on-- I don't know if you guys are maybe too young to have ever seen anything like this in operation, but the next time if I remember, I'll bring an old fashioned transparency, and I will show you how people of my age used to give talks at conferences when we didn't have PowerPoint. But anyway, this is done actually very commonly, to stick a lens in front of-- just attached to a thin object. But there are some cases you cannot do that. For example, in the microscope, it is very difficult to imagine sticking a lens behind the glass slide. So in some cases, this is applicable, and in some it is not.

Anyway, a big application of this type of lens is of course in overhead projectors. Does anybody know another application where people use this kind of lenses, or what might motivate you to make this kind of lens? Can you think? These lenses are commonly used-- actually, you know?

AUDIENCE: Do they sometimes stick them on the backs of vans or cars to make it wide angle?

GEORGE BARBASTATHIS: That's right, yeah. That's another application. Exactly, to improve the rear view of the driver. And that's a good one. I wasn't thinking of that.

Yet another one is in very high power illuminators, which are most commonly used in photography in shooting movies because the glass, as you know, is not a very good conductor, so it would hit up very badly if you illuminate with very high power. They're called, actually, babies, these really big light bulbs that they use when they shoot movies. So they actually use these. It is called a Fresnel lens, the same Fresnel lens in Fresnel propagation. So they use the Fresnel lenses in these so-called babies. All right. And lighthouses also, yeah.

So again, every time that you have a really huge light source. Of course, it's a matter of practicality also, right? If you have a major lens, as you can imagine, if the diameter is big, then also this size would also be very big, right? So in a lighthouse, you might have a lens that is maybe 1 or 2 meters in diameter, then it would have to bulge over another half meter or so. That's not very practical. It is heavy, fragile, blah blah blah. So they make the Fresnel lenses to get around this problem.

A good thing to remember about those is that they generally give very poor image quality. So if you try to use this as an imaging lens, you typically get a very bad blur and a very bad image quality, but obviously in the lighthouse you don't care. In an illuminator, you don't care. And also the projector, it is not used as an imaging lens. The imaging lens is a real lens, this one. This is simply collimating the illumination so you get a uniform image at the image plane.

So for jobs like this one where you basically want to take a light bulb that is highly un-uniform. It has filaments, all kinds of crap. So this is a very nonuniform field. This type of lens is very good to uniformize the intensity, and then of course if you need to image a sharp object, then you need the real lens. You cannot get around that. OK, that's a bit about Fresnel lenses. So that's a bit of a detour, but I wanted to point out some practical ways that you might be able to do this kind of thing if you need to get rid of the lens.

And finally, what John Goodman proved back when he was young is that there's a condition that you can neglect this nasty quadratic. Again, we're talking about the quadratic that you see at the top of the slide. So it turns out if the field of the object-- when we say the field, sometimes in geometrical optics we mean the size of this. So if the field is smaller than about a quarter of the imaging lens, then it turns out again that the effects of this term are negligible, and therefore we can get rid of it in clear conscience.

So I will not go into the details. There is discussion in the book and also reference 303 from the book is an article that Goodman wrote back then, and he goes into this in a little bit more detail. It's a pretty interesting article. I'm going to leave a post on the website if you like. It's pretty interesting to read.

So assuming then that will get rid of this extra quadratic-- let me go back one. So we got rid of one quadratic from the lens law, and I spent a lot of time arguing about this quadratic, that maybe you can get rid of that then as well, assuming you can get rid of that as well. Well, anyway, even if we don't get rid of that, what is left in here is actually Fourier transform. I think I pointed that out before. A Fourier transform-- why? Because you have linear now, linear exponents times some coefficient, and then here you have the pupil mask. So again you recognize something familiar.

The same thing happened in the case of the 4F system. We again got a term like this one, but it was slightly different. Instead of z1 z2 in the denominators here, it had f1 and f2, the two focal lengths. Here we have the two distances, but it's pretty much the same, and again the pupil mask appeared in the kernel here. So this is a Fourier transform, and of course we get rid of this term, and let me jump through these animations now.

OK, so if we're doing that, then what we finally get is this expression here, where, again, this is the Fourier transform of the pupil mask. I wrote it in a form that looks a bit more like a Fourier transform, and of course u and v are now arbitrary arguments. These are the spatial frequencies in the Fourier transform, but what actually goes in there is these parameters over here. So basically we get again a term that looks like a convolution. So we see that the image, the complex optical field at the image is a convolution of the input field times something that again we will call the point spread function. And this point spread function, just as in the case of the 4F system, again here the point spread function is given as the Fourier transform over the pupil of the pupil mask scaled with appropriate coordinates.

OK, and there is various scaling forms and factors here. For example, this one again ensures energy conservation, which is kind of important, but very often in optics we don't care so much about the absolute intensities, which means that we don't really care about this term. What we really care about is about the spatial distribution, like the type we're solving, whether features make it to the image or they get lost due to spatial filtering and so on and so forth. So very often these multiplicative constants, we just drop them. That's sort of convention.

Now, to get a little more insight out of this, imagine for a moment that the point spread function becomes extremely narrow. It becomes a delta function. So if you substitute the delta function into the expression we had before, then you get this expression, which is of course again very gratifying because now the output really looks like, again, some multiplicative term, but then it really looks like the input, except in the input now, the coordinates have been scaled, so therefore you can immediately see that this term is the magnification. In fact, the lateral magnification of the system is given by this expression. So we basically rederived the same result that we got from geometrical optics-- namely, the lens law that is the imaging condition, the magnification and so on and so forth.

Now we've got them back from our wave optics approach, but we had to do some approximations. For example, we had to pretend that the point spread function is a delta function in order to get the geometrical optics result. If it's not, which is never the case, right? Delta function is a very good approximation. So if it's not a delta function, then the result is not exactly like this, but it is a convolution. That is, basically this will become blurry. Like [? Pepe ?] showed earlier, you will get a blurred vision-- or in general, an especially filtered version of the original input will actually survive through to the output.

What I would like to point out now, this is something that I believe Se Baek also showed last time, the same slide. What I would like to do is basically pay a little bit of attention to the scaling factors here. These scaling factors are actually very important. So I mentioned that to drop multiplicative factors that appear in the front here. This, it is OK to drop, but these scaling factors that go inside the argument, these are very important because they determine the amount of spatial filtering that goes on inside the system.

So what I will do actually for a while is I will compare the 4F system with a single lens similar system. So the scale factors, again, they are different. You see f1 is inside the transfer function, and z1 in the case of the single lens. So what is the effect of that one? Also in the incoherent case. Of course, in the incoherent case, the transfer function is the autocorrelation of the coherent transfer function, and since the coherent transfer function is basically proportional to the pupil mask, then what we call the optical transfer function, the incoherent case, is an autocorrelation of the pupil mask itself. This is again a very basic result, but again, be careful. When we compute these autocorrelations, we have to apply different scaling factors in the argument.

OK. So I will show in a second how this works, and I will show it in two cases. One is what we called before the Zernike phase mask. Actually, this is not exactly accurate. Zernike did not invent exactly this. He invented another one that has pi over 2 phase shift at the edges. It's like a ring. But anyway, the function is the same, so everybody refers to all of these types of masks as Zernike nowadays.

So this is familiar. We saw a few examples in the past when it was in the middle in the focal plane of the 4F system, and this is the case when we put it as a pupil mask in a single lens system. It's the same mask. It will do a very similar thing, but there's a subtle difference, so this is what I want to point out. So this is the subtle difference. What you see here is, of course, the schematics, so you see it is opaque outside some amplitude. Then inside, you have this little extra phase delay near the optical access in a small region.

And OK, this is the mathematical expression that I don't want to dwell upon. I'd rather focus on these plots here. So these two plots are the magnitude and the phase of the mask. So the magnitude is the top. The blue plot it goes from 0 to 1. So it is 1 within the amplitude, so let's see if that's correct. The amplitude has a size of 1 centimeter, so indeed this is 1 between minus 0.5 and 0.5. And the phase, well, the phase, nobody knows what the phase is. Where the magnitude is zero, we can not define the phase for a complex number that equals zero. Nevertheless, let's set it equal to zero.

But what I want to emphasize is that within the past band of the system, the phase is 0, except for a small chunk of width of 0.2 centimeters where the phase jumps to pi over 2. Pi over 2 is of course i. And this is the real and imaginary part. These two are actually equivalent. So the real part, it is 1, except at the little extra phase protrusion there where the real part goes to 0, and the imaginary part becomes 1 because of course when the phase equals pi over 2, then the actual complex amplitude equals i. So that's what you see here.

So these two pairs describe both of the phase mask. So the [INAUDIBLE] transfer functions of course are the same. I only plotted them here as a real and imaginary part, but they have the same shape. What is really different-- and unfortunately, my plot is a little bit too small here. You cannot see very well-- is that they actually have different sizes. This one, if you do the scaling factors here for the numbers that I used, I think I used-- so I used lambda equals 1 micron. Lambda equals 1 micron, f1 equals 10 centimeters, f2 equals 1 centimeter for the 4F system. And then z1 equals 11 centimeters, z2 equals 1.1 centimeters, and f1 equals still 10 centimeters for the single lens. And I worked those out so that in both cases, you get the demagnification of a factor of 10. That's why the numbers are a little bit strange there.

So in both cases, you get a demagnification factor of 10, and I did that deliberately because the two systems give you the same effect kind of, but you can see in terms of geometrical optics, they are identical. They both give you a demagnification by a factor of 10, but you see here that in terms of wave optics, they're slightly different because one of them, the 4F system is scaled by the focal length. Its spatial frequency extends, goes from minus 50 to 50 inverse millimeters. This is what you see here. Whereas in this case, it goes from approximately minus 48 to 48 inverse millimeters.

So the single lens actually does a slightly more severe spatial filter than the case of the 4F system. That is to be expected, of course, because the distance z1-- you can see from here-- the distance z1 that you subtend from the object to the aperture is longer, so therefore this system is cutting off more angles, or more spatial frequencies than this system. We will see that in more detail in a second.

The other thing I want to point out to you is that if you take that autocorrelation of this function, which is the coherent transfer function or the ATF, if you take its autocorrelation, you obtain the OTF, which is, as Se Baek demonstrated last time, it is the transfer function for spatially incoherent illumination. So again, this is a little bit of an exercise that I will not do here, but it is in the notes. I have posted another set of practice problems.

So it is in pages 16 and 17 of the last set of practice problems. I have gone ahead and derived this one, and I would encourage you to go through. It's a little bit of algebra, but you can think of it as mental fitness right because you might say, why do I need to do all this algebra? When would I ever have to do so much algebra? I can just plug it into Matlab, and I get it. Well, I'm sure all of you do some kind of fitness. You go to the gym, right? When you do bench presses, what is the probability that you need to do a bench press motion in real life? It's actually very small. So you do bench presses in order to keep fit. So the reason we do these kind of calculations is it's the equivalent of mental bench presses, right, so I would encourage you to do it. Because presumably you are at MIT because you have you want to have mental fitness as well as physical fitness.

Let me talk a little bit more about these scaling factors. So in this case, this is just a clear aperture that I put here, and you can see that in the case of the clear aperture, it's actually very easy to do sort of the ray diagrams, and you can see also here that-- well, let me back up for a second. So you remember from geometrical optics, we had the definition of the numerical aperture. So we defined the numerical aperture as the angle that we subtend towards the optical system if we place a point source on axis. So therefore, the numerical aperture is limited by one of the physical apertures that are in the system.

So in the case of the 4F system, that would be the pupil mask. At some point in any physical system, the pupil mask cannot be infinite. It has to be finite, so the physical diameter of the pupil mask, assuming the lenses themselves are large enough, the pupil mask will become the limiting factor. And so we can compute the numerical aperture then. It is the ratio of the radius of the mask over the focal length, f1. You can see it very easily from this triangle over here. And of course this is an approximate expression. In reality, the numerical apertures would be the sine over the inverse tangent of this quantity, but of course we're doing [INAUDIBLE] approximations here, so we can drop the trigonometric functions.

In the case of the single lens the numerical aperture is-- again, you can get it from this triangle over here, but now the one of the orthogonal sides is z1, not f1. So you can see that the numerical aperture is r over z1, and of course if this system is supposed to produce a real image, as opposed to a virtual image, then z1 must be longer than f1. You can see very clearly from the imaging condition, 1 over z1 plus 1 over z2 equals 1 over f. If you want z2 to be positive-- so for z2 to be positive, it means I have a real image-- for this, it is required that z1 is bigger than f. So because of that, you can see that if you have the same size pupil mask, the single lens system would have a smaller numerical aperture. That is sort of an unfortunate fact of life.

So another thing I want to point out here is that the numerical aperture actually changes when you go to the second leg of the system because, of course, the system has angular magnification. So what started as a numerical aperture at the input is of course an angle. If you take the marginal ray, it is propagating at a given angle with respect to the optical axis. By the time it comes out, this angle will have changed. By how much? By the angular magnification of the system. So therefore, the numerical aperture at the output is not the same. It equals the numerical aperture at the input times the angular magnification of the system. So this is not a cause for confusion. We can use either one. We'll get actually the same conclusions whether we use the numerical aperture or the input or the numerical aperture at the output, but we have to be a little bit careful that we don't confuse things, as long as we remember this simple relationship that the two are connected by the angular magnification.

So the reason the numerical aperture is so important is because if you consider a circular-- of course, you cannot see the circle here. This is just the projection. But imagine that I have a circular pupil. Then recall that the point spread function of the system is actually the Fourier transform of that circular pupil because this is a clear aperture now. And we learned sometime ago-- I don't remember when-- but we learned that the Fourier transform of a circular function is this crazy [? zinc ?] so-called function that is given by a ratio of a Bessel function to its argument, and it's probably better to think of it as a plot like this one that has a main lobe and then a smaller side lobe. And actually, it has a lot of lobes. It continues on forever, but the size of the lobes decays away. How fast it decays away is 1 over the argument.

So what I want to point out here is that in both cases, the point spread the function of the system looks like this [? zinc ?] functions. Sometimes it's also known as an Airy disk. Not an Airy function, by the way. Airy function is different. This one is referred to as Airy disk. If you're curious what Airy function, is you can open the table of formulas. Abramowitz and [INAUDIBLE], and you'll see a monstrous thing. That's called the Airy function, and this turns out to be a special case of an Airy function.

But anyway, this is the Airy disk, and what I want to emphasize is that in both cases, you get the same shape of Airy disk, but different in size. And you'll get different size because in the case of the 4F system, the size of the aperture is the same, but because the scaling factor is smaller-- it is f1-- then you actually get a bigger ATF. The size of the disk in the frequency domain is bigger. Therefore, it will give you a slightly narrower point spread function. If you work out these ratios over here, and you work out these coefficients, you will see that very clearly, but I want you to get it sort of intuitively by using the scaling theorem of Fourier transform.

In both cases, you start with the same physical aperture, but what matters is not the physical size of the aperture, but the numerical aperture. So in one case, you have this numerical aperture. This is a f1. This is the physical size, and numerical aperture is R over f1. In the other case, you have again the same aperture, but now you have a longer distance here. This is still R. So now the numerical aperture would be R over z1. So the same physical aperture in the two cases will actually give you a different size in the ATF. This will have a bigger ATF.

So the size of the ATF, if you work it out, it will be proportional to actually 1 over lambda f1, and the size of the ATF here-- I'm exaggerating, of course-- will be proportional to 1 over lambda z1. OK, so since we've got the smaller size of the ATF in this case, it means that we'll get a broader PSF. So which system is better? Well, obviously this one because this one will give you a narrower PSF. Therefore, it will give you a smaller blur.

So of course there's a caveat. I sort of took it for granted that the desirable in an optical system is to minimize the blur, which in most cases it's true. If for some reason someone asked you deliberately to produce a system that causes a lot of blur, then of course you would go for this one, but in most cases, we try to minimize blur. So this means, given our resources-- that is, given our physical size of the numerical aperture-- we should try to maximize the numerical aperture, and that is what the photo system is doing. So there are a few comments about the resolution that are in the rest of the slides. I don't want to destroy the rest of your morning.

Because the PSF is finite, you can imagine that if you have two point sources that are at distance. So now this is like an experiment. I'm moving two point sources together. There will come a point where the two point sources will merge. And if this is your image now, you don't really know whether you had one point source or two. Now, of course you say, well, you also get twice the intensity. That's true, but very often you don't know the intensity you started with.

For example, if you're looking at the sky, and you're looking at the bright dot and you're trying to decide, is it one star or two stars? Well, so far we cannot yet go to the stars and measure their brightness, right? We can only measure the total brightness that we receive here. So therefore, if you're looking at a telescope, and you see this, you don't know if there's one star or two stars that are too close to be resolved by your telescope. So actually, this sort of situation was dealt with by a fellow called Rayleigh, and of course, the size of the PSF, as I worked out before, it depends on the numerical aperture. So I will go over this in more detail later, but this is just a preview that the numerical aperture actually gives you an idea-- I actually believe also Professor [INAUDIBLE] mentioned it in one of the past lectures.

The numerical aperture gives you an idea about the capability of your system to resolve point images. So imagine that each one of these lobes here corresponds to the point spread function of one point object. So in this case, they were resolved because I spaced them so that the diameter of the main lobe falls exactly on one null of the other lobe. So if you work it out, and you take into care into where the zeros of the Bessel function are located and so on, you come up with an expression that looks like this. The space in between the two sources is 1.22 times the wavelength divided by the numerical aperture. And of course this would be at the input plane. So this would be the spacing at the input required so that the two sources can be resolved.

But when you look at the output, the distance of course would be given by a similar expression, but the numerical aperture at the output would appear in this case. And of course this is sort of the most common definition. Some books define the resolution-- instead of the diameter, they use the radius, so they come up with an expression that is exactly 1/2. Which one is correct, it actually depends on your requirements. If you require a very crisp image, or if you're dealing with a very noisy situation, then you go for this one because as you can imagine if you organize noise into this one, then it becomes progressively more difficult to resolve the two spots.

So in a very noisy environment, when you have a very little light, then you go for this definition. If you have plenty of light, then you can possibly resolve two sources that are very close like this, right. Here you see you have very poor contrast. You basically have to rely-- in order to resolve the two sources, you have to rely on this little dip in density. So if your signal is noisy, this may be lost, and then you cannot resolve anymore. OK, so that's a preview of resolution. We'll talk about this a little bit more.