GEORGE BARBASTATHIS: There's some comments from the grader who looked at your homeworks. So a couple of points to remember that they seem to be sort of pervasive in several homeworks. One is that when we do Snell's law, we have to apply the rule-- the rule goes like n sine theta equals n prime sine theta prime. So remember that theta prime is measured with respect to the normal to the interface. So if this line is the interface, I have, for example, air and the glass here.

AUDIENCE: [INAUDIBLE]

GEORGE BARBASTATHIS: Oh.

AUDIENCE: Yes, that's good.

GEORGE BARBASTATHIS: I think my hand is confusing it. Maybe I should turn off the autofocus. OK. So if this is the interface, air and glass-- yes, I think my hand was confusing the-- and the ray is incident, for example, like this, then we apply Snell's law with respect to the normal. So this is the angle theta, not this one out here, not the other way around. And for this one, of course, the same goes to the other side of the-- whoops-- the other side of the interface. The ray would go in kind of like this. That would be theta prime. So that was one comment from the homeworks.

The other comment is about TIR. TIR actually happens for all angle of incidence that are larger than the critical angle. There seemed to be some confusion that somehow TIR happens only at the critical angle. That is not true. In fact, strictly speaking, at exactly the critical angle, TIR does not happen. What does happen is this surface wave that I mentioned in the notes. But for all angles that are greater than the critical angle, this is when we have the TIR condition for all those angles, from critical all the way to whatever it is, 90 degrees, I suppose.

The third comment-- and, actually, this is the last one. The third comment is about the speed of light. Some of you-- and, actually, this is also in relation to the streamer problem. There was some confusion where the speed of light is proportional to the index of refraction or inversely proportional. It is, in fact, inversely proportional. The speed of light goes like c over n, where n is the index of refraction.

And I don't know if there's a mnemonic for this, but one way to remember it is that c, the speed of light in vacuum, is the highest speed that anything can have. So, therefore, when light enters a material, it can only go slower. So, then, the index of refraction tells you the amount by which the speed of light becomes slower in a given dielectric. So these are some things that are useful to keep in mind as we move forward. These two dots don't mean anything. Just the ink from the markers that penetrated the piece of paper, so please ignore it.

So, last time, the topic was composite optical elements. So we started with an example of two lenses, two thin lenses that were in tandem, and we did the imaging condition, magnification, and all those beautiful things with it. And then we also looked at the thin lens-- I'm sorry-- at the thick lens, where the thickness is defined-- as you recall, it is defined relative to the paraxial approximation. If the distance between the two curved surfaces exceeds the cosine theta, or whatever it is that the paraxial approximation limits for the negligible distance, then we call the lens thick. And we saw that there's some differences. For example, the focal length is a little bit different in the formula than it this for the case of the thin lens.

But the primary reason why we did it is because I wanted to define this concept of principal planes. So the principal planes-- I'll move slowly so the camera can follow me. So the principle planes are basically the locations where the rays bend at once. So in a thick optical element, whether it is a thick lens or it is a composite of several lenses in tandem, as you can imagine, the rays will bend several times. They may bend up, downs, straight, and so on and so forth. But the idea is that if you take a parallel ray from infinity, after however many times it bends-- in this case, it bends twice, once here and once here-- eventually it comes out at the final propagation angle.

So the principal plane is the plane that you obtain if you extend the incoming ray, you extend the outgoing ray, and you find where they meet. That is the principal plane. In fact, in general, it is a principal surface. If you really do the geometry right, it is not a plane. It turns out to be a curved surface, but within a paraxial approximation, it is a plane. So that's why we define it with sometimes as a principal plane or as a principal surface. I kind of try to stick to the term principal plane in the notes.

And then, relative to the principal plane, as we did last time, you can define the various quantities. The effective focal length is simply the distance from that principal plane to the focal point. So, basically, as far as the rays that are coming from the left are concerned, the principal plane is acting like a thin lens, because the-- is acting like a thin lens with a focal length equal to the effective focal length of the composite, because this is the ray coming from infinity. And when it comes out, it is as if I had put a thin lens in this particular location.

The opposite is true for images at infinity. So if you take a ray from this point here, it actually goes to the first principal plane, and then takes off to infinity. Now, of course, again, this is what I get when I take the incoming ray and the outgoing ray, and then find where they meet. In general, this ray also suffers multiple refractions. But if I take the first one and the very last one, then I can have a thin lens at the first principal plane.

So this is the idea, and we can see the reason why they are convenient when we do the imaging condition for the composite optical element. So we'll not do this again. Se Baek went through it in great detail last time. But you can see that, basically, if we know the principal planes, then we don't have to worry about the rest of the physical system any more, because we can discover things like imaging condition, magnification, and so on and so forth simply by applying the rules that we learned in the case of the thin lens.

For example, in this case, in order to find the image, I take a ray that is horizontal. Therefore, it would have come from infinity. I drag it all the way to the second principal plane. Then I know it has to go through the focus. And then I extend it further. When I take another one that goes through the first focus, then I know that at the first principal plane, this ray will have to become horizontal, and again take off towards infinity. So, therefore, this ray goes like this. And then where they meet is actually the location of the image.

So this is the idea. And then all the rest of the formulas about the imaging distances, the magnifications, lateral, angular, so on and so forth, they all fall out from the similar triangles in this diagram, just in the same way we did them for the thin lens, except now we have these sort of additional space between the two principal planes. So we have to be careful when we deal with a ray that is coming horizontal at the input. It has to bend at the second principal plane. And when we think the ray that goes through the first focal point and comes out horizontal, that one has to go through the first principal plane.

And the reason I'm emphasizing this is because-- this is the generic case, where the first principal is to the left, the second is to the right. But there is optical systems where they may actually be swapped. You may have the second principal plane to the left of the first principal plane. There's no rule that forbids that. It can happen. So that's fine. We still have to apply the same rule, but we have to be careful where we bend the ray.

So one of the homeworks-- you will see the homeworks that we posted today and are due, actually, next week-- you will see that this happens in one of the optical systems that we give you there. The second principal plane is actually to the left. So I'll let you work that out by yourselves, but keep that in mind. Always identify which principal plane you're using for any given imaging condition.

Any questions about all of this? Can someone say at least good morning, so I know if our sound has been restored?

AUDIENCE: Hello, hello, can you hear us?

GEORGE BARBASTATHIS: OK, either nobody wants to say good morning or there's no sound.

AUDIENCE: No. No sound, I guess.

GEORGE BARBASTATHIS: OK. I'll keep going, and, hopefully, some time at some point we'll have sound. Or at some time you'll wake up enough to be willing to say good morning. So, today, we'll deal with a topic called stops and apertures. And, basically, we'll spend the rest of the hour learning the meaning of all of these terms. So there's quite a few of them-- aperture, pupil, entrance, exit pupil. And then we field stop, window, entrance and exit, and so on. So we'll spend some time learning what these things are and why they're useful in an optical system.

So starting with the first one, the aperture stop, the aperture stop has to do with the fact that any optical instrument, in practice, has a finite size. Therefore-- OK, let me back up. So the instrument has a finite size. But a point source, as we learned, it actually-- at least an ideal source-- it emits kind of isotropically, all the way 4 pi steradians around the location of the point source. So what this means, then, is that the optical instrument can only capture a fraction of the light that this-- that any given point sources is emitting. So the aperture actually helps us quantify how much is the light that an optical instrument can admit given the location of a point source.

So it is defined, actually, with respect to a point source that is located at the optical axis, that is, at the center where all of the optics are rotationally symmetric about. And there's two concepts here. One is the physical stop. So the physical stop is-- you can think of it as the rim of the optics, the physical limit beyond which the rays that are coming from the source, they miss the optical elements in the system.

Those of you who have cameras, you're also familiar with physical stops that, in relatively complicated optical systems-- for example the lens that goes with the camera-- inside there's a little hole that is very often adjustable depending on the light level and so on. This hole, nowadays automatically, in the old days with photographers having manual control, you would actually open and close that opening to admit more or less light. This is called the aperture stop.

And we also are familiar with something else, the pupil of our eyes. If you look at someone's eyes, there's a little hole. And if you look in a bright environment, this hole is relatively small. If you move to a dark environment, our eyes have a mechanism that kind of measures the ambient light and opens up or closes down the pupil. So this is the aperture stop for the optical system of our eye.

So the aperture stop, then, is the physical element that limits the acceptance of the light rays. So, in this case, the aperture stop, you can think of this sort of black, thick line here, you can think of it as a physical, for example a metal plate with a hole in the center. So that hole would be the actual aperture stop in this particular instrument.

The numerical aperture is a mathematical quantity. It is the angle subtended from the point source to the rim of the physical aperture stop. And I was a little bit sloppy when I called the angle. It is not the actual definition. It's the sine of that angle multiplied by the index of refraction.

Now, we've already seen Snell's law several times. You can guess immediately the reason people define the numerical aperture this way is so that this quantity is conserved when you go through flat refractive interfaces. So it's very convenient because, well, as we'll see later, the numerical aperture is a very important quantity in characterizing an optical system. So, therefore, by sticking the index of refraction in this definition here, we'll make sure that the numerical aperture is preserved as rays go through various elements in the system, at least elements without power, without optical power, that is, as they go through flat interfaces.

So that's the definition. Now, of course, in many optical systems, the index where the object lives is air. So, therefore, n is 1. And also, in this class at least, we deal with the paraxial approximation, where the sine of the angle is approximately equal to the angle itself. So, therefore, indeed, the numerical aperture is approximately equal to the angle itself.

However, as you have seen from the homework, if you already started it, in many optical systems the object lives in a space of refractive index different than 1. For example, in immersion microscopy, we put a droplet of oil. We drop a droplet of oil on top of the sample. Then we stick the objective so that there is oil in this space here between the sample and the objective. Well, in that case, we really should include the index of refraction in the definition of the optic-- of the numerical aperture.

So, of course, you cannot speak yet, but I do want to ask a question myself now. And if you can answer-- well, those of you who can answer easily. Those of you in Boston, if you want to answer, just please wave and we'll figure out how to do it. But the question I wanted to ask, what do you think the value of the numerical aperture is? Is it smaller than 1, bigger than 1? What is it?

Do this if you think that the numerical aperture is bigger than 1. Do this if you think that the numerical aperture is always less than 1. What's the verdict? Always less than 1. Anybody who thinks the opposite, that the numerical aperture can be bigger than 1? It's really strange. I feel like we're transported to a different century now [INAUDIBLE]. Anyway.

So, actually, in principle-- and, in fact, people do it-- I can construct an optical instrument with a relatively large acceptance angle. For example, theta may be as high as 75 degrees, 80 degrees. It is not uncommon in high-end microscope objective lenses to have a very high acceptance angle. Of course, this would violate the paraxial approximation, but that is no problem. The fact that the instrument is not paraxial actually does not make it an invalid instrument, so it's fine. All it means is that our simple theory here does not apply, and we have to be more sophisticated in the analysis, but we can have a large acceptance angle.

For example-- so, in that case, the sine of this angle may be close to 1, maybe 0.9 or 0.95. And, also, in addition, remember that we can use immersion. So we can put the object in a liquid of index of refraction-- for example, if it is water, that would be 1.3. If it is one of the immersion oils that are used in microscopy, what is the index, typically? 1.4, 1 point--

AUDIENCE: [INAUDIBLE]

GEORGE BARBASTATHIS: Oh. Just the same as glass?

AUDIENCE: [INAUDIBLE]

GEORGE BARBASTATHIS: OK. So even higher than for immersion instruments, 1.54-- 0.14. I'm sorry. 1.514. And so, in that case, if I do 1.514 times 0.9, of course it will be bigger than 1, right? I don't know what it is, but it would be definitely bigger than 1. So, actually, I can have a numerical aperture that is bigger than 1. That is kind of a useful point to remember.

Related to the-- so the numerical aperture is used very commonly in microscopy when we characterize objective lenses. Typically, we characterize them by their numerical aperture. In photography, especially, people prefer to use an alternate quantity known as the f number. So the f number is kind of the inverse of the numerical aperture times a factor of 1/2, which always confuses the hell out of me. But, anyway, the f number, it is written like this. And people replace the hash symbol with the actual number. For example, if you look at the lens of a camera, it may something like f/8.

This means that the f number of this camera lens is 8. So the rule of thumb is where does your money go? When you pay money for a very good quality microscope objective, one of those that cost $10,000 or so, typically you pay-- in a microscope objective, you pay for a large numerical aperture. When you walk into a store and you buy a high-end, expensive camera lens, one of those Nikon or a Zeiss, really expensive multi-element lenses that look about as big as my computer here, what you pay there is for a really low f number. So, generally, large numerical aperture is good. Or, equivalently, low f number is good, provided you can afford it, of course.

So, at the moment, the way we defined it so far, the numerical aperture, it seems to be related only to the energy acceptance of the optical system, that is, how much-- how many watts of optical power the optical system admits. So this is of course important, but you might wonder why am I making such a big deal, and why people pay so much money for high NA. Well, again, the energy acceptance, as you can imagine, is incredibly important because it will determine how many useful photons from the object you admit into your system for imaging purposes. But, as we will learn later, the numerical aperture also defines the resolution of an optical system.

And, right now, you probably know from everyday life what resolution means. People talk about the resolution in your camera, in your television, and so on and so forth. We will do a precise definition of resolution later, in about a month from now. We'll define it very precisely, and we'll see why and how it is related to this numerical aperture quantity. But, again, the rule is that the largest the numerical aperture, or, equivalently, the smaller the f number, the better the resolution of an optical system. So we will see this come up again and again.

So that's the basic definition that we use. Now, related to this concept of aperture stop and numeric aperture, there is two sort of derivative definitions, and these are called the pupils, the entrance and exit pupils. So in order to define those, I kind of need a multi-element optical system, so I sort of concocted one over there. I just made it up. This is nothing in particular. I just stuck some lenses in tandem, And i put the aperture stop somewhere in the middle.

So think of this again as a physical element, a hole that I put in there. So the reason we need these additional definitions of the pupils is that if the aperture stop is varied inside an optical instrument, as it happens, for example, in your camera lens-- if you look at your photographic camera that you have at home, you will see that the aperture is varied inside a relatively large lens barrel if you have a sophisticated camera. The aperture is a physical element. It is a hole, right?

So if you look at your lens from the one side that looks towards the object, or from the opposite side that looks towards the CCD, you will see the aperture, but you will not see the aperture itself, you will see the image of that aperture stop through the elements that precede or succeed it. So because of that, people have defined this concept of the entrance and exit pupil. So the entrance pupil is what you get if you take this aperture stop and now you treat it as an object in an imaging system, and you image it through the preceding optical elements.

Now, here, I've committed a mortal violation. I imaged from right to left. We said last time that every time we'll write down an optical system, the light is assumed to go onwards from left to right. So, here, this is an exception. When we deal with aperture stops, we actually image from right to left, because we really need to look at this as the object, and image it through the optical elements that precede it. So if we do that, then we get the entrance pupil.

So, in this case, it would look something like this. So this is now not a physical component any more. It is the image of the physical component that was inside the system, and we imaged through the elements that preceded it in that optical system. Similarly, we can define the exit pupil if we image the physical aperture stop through the elements that succeed it. So the two definitions are kind of conjugate.

Now, with respect to this entrance, exit pupils and aperture stops-- this is a mouthful-- you can define two significant rays that come up in the design of optical systems all the time, and we'll see why. The first one is called the chief ray, and the chief ray is defined as the ray that goes through the center of the aperture stop, the physical aperture stop.

So that's one. The other is called the marginal ray, and it is the ray that goes through the rim of the physical aperture stop, that is, the ray that just clears the aperture. If you take the next ray, that would propagate at a slightly higher angle with respect to this one. This ray would actually hit the physical block that is located at the aperture plane and would miss. So this last ray just before the aperture stop blocks the rest of the light, this is called the marginal ray.

So you can already see the importance of this, because the marginal ray kind of defines again the angle of acceptance of the optical system. Except, this time, notice that I've done this for an off-axis object point. Recall that a numerical aperture, the way I defined it before, it was still the angle of acceptance but for an on-axis object point. Here, the marginal and the chief rays, I can define them for any point I like, any point I like that is in the object space.

And this step, let me sort of-- OK, I already decided that these rays help us define the angle of acceptance. Now, before I show the next slide, and since these two rays, these two important rays, where do you think they might go with respect to the entrance and exit pupils? By definition, they go through the center and edge of the aperture stop, respectively, the chief and marginal ray. So, therefore, with respect to the entrance and exit pupils, where do you think these rays would pass?

So the answer is that they should pass through the center and an edge, respectively, as well, because the entrance and exit pupils are images of the aperture stop. So, for example, if the chief ray goes through the center of the aperture stop, then it should also go through the center of the entrance and exit pupils because these are images of the aperture stop through the preceding and succeeding elements. And, similarly, the marginal ray hits the rim of the aperture stop. Therefore, it should also hit the rim of the entrance and exit pupils, respectively.

Now, for the marginal ray, there could be magnifications involved. So, for example, the distance from the optical axis to the rim, it may be different in the aperture stop plane than it is, for example, in the entrance pupil plane. And that is because, in this case, the optical system has magnified the aperture stop. So the entrance pupil in this case is bigger than the physical stop itself. So, therefore, the marginal ray, it hits the optical axis at a different location.

And, also, as you can see, there's an inversion. It hits at the bottom at the aperture stop, and then it hits at the top of the entrance pupil. This is fine, but what I-- because of the way optical systems behave. They can produce inversions and magnifications, and as we have seen. But what I want to emphasize is that they go through the rim because of that property that the aperture is imaged at the entrance and exit pupils, respectively.

And, of course, the two rays will also meet again at the corresponding image point at the image plane. Because the two rays emanated from the object, if the imaging system is good, that is, if it does form an image properly and we don't have some kind of a weird defocus or some other kind of inconvenience, then the two rays should better meet again at the image point, because the entire spherical ray bundle that emanated here, that diverged here, it would also converge by the time it arrives here. So, therefore, these two rays are just members of this spherical ray bundle, so they will meet again at the object point.

So any questions? I know this is frustrating. Even if you have a question, you cannot speak your mind. But any questions? Can you just wave at least so I know to save time later to answer the question?

AUDIENCE: So how do you define the distance between the-- I guess he can't hear me. Ah, right.

GEORGE BARBASTATHIS: Let me move on with a slightly different concept, and then we'll come back to this topic again. So the next topic is a different kind of stop called the field stop. So remember the definition of chief ray. We said that the chief ray is a ray that goes through the center of the pupil. I mean the entrance pupil, the aperture stop, and so on. So here is a chief ray from a tall object, and here is a chief ray from a short object.

| of the chief rays, they will go through the center of the aperture stop by definition. And, of course, they will also go through the center of the entrance pupil. The entrance pupil used to be here, if you remember. I'm not drawing it any more because I don't want to clutter this diagram, but it used to be here. So the two rays, they meet at this point because they're the images of the aperture.

However, this ray will go off axis at other points in the optical system, and there will be some point in the optical system where one of these rays may get cut. This element in the optical system that limits the chief rays, that element is called the field stop. And the physical significance of the field stop is actually very important and even more intuitive than the aperture stop. The field stop is what defines the field of view.

And now the field of view has a formal definition. It is this angle here that I drew. But even in colloquial language, when we say the field of view, we know exactly what we mean. In an image, the field of view is basically the size that can be visible. And outside the field of view, the image is not formed.

So an example is if I look at you-- well, actually, this is great because I look at you through a camera, and the way the camera is pointed right now, I can only see the central row of desks out in Cambridge but I cannot see those of you who are seated on the left or on the right. So you might be picking your nose, for example, and I would not be able to see it. And that's because you are outside my field of view.

And thank you for the control, for turning the camera. Now what is happening, my field of view is shifting. So the field stop in this case is actually the dye, the electronic dye that captures the image electronically on your side, then transmits it to us here in Singapore. This is the physical field stop that limits how many of you I can see at any given instance.

So the field stop is second to the numerical aperture and the aperture stop, is another very important quantity. It tells us how big is the size of an object that we can see. So if you look, for example, at a telescope-- if you look at the night sky through a telescope, the field of view is basically how many stars can fit, what portion of the sky you can fit within your image when you look through the telescope. When you look through a microscope, similarly, the size of the specimen that you can observe in the microscope, that would be the field of view.

Now, intuitively, we would like to define the field of view in physical size, in meters, or millimeters, or microns, or whatever the case might be. For reasons that will become apparent a little bit later, we prefer to define the field of view as an angle actually, the angle subtended towards the object. And you can guess-- I don't even have to make you wait. You can guess why we define angle instead of physical size. That is because, very often, optical systems have variable magnification.

So with the same optical system might move a lens, and then the object would move, but the field of view might not change. So if I define it as a size and then move it here, then the size would be bigger, whereas the angle generally is preserved. So that's why we define the field of view as an angle in radians or degrees instead of physical distances.

AUDIENCE: Hello, hello, hello.

GEORGE BARBASTATHIS: So that is the meaning of the field stop. And, similarly, through the entrance and exit pupils that we saw for the aperture stop, we can define equivalent quantities for the field stop, but we'll call the windows. So if we image the field stop through the preceding elements, we get what is called the entrance window. If we image it through the succeeding elements, we'll get what is called the exit window. And, generally, in an optical system, these windows, they should coincide with your object and image planes respectively. So the entrance window should coincide with the object plane. The exit window should coincide with the image plane. If that's not the case, there's something wrong with the optical system. It can happen. It is not impossible, but there's something wrong with the optical system.

So now that we did that, let's put all these definitions together. There's a lot of confusion in the diagram, but it shows everything. It shows the chief and marginal rays. It shows the two physical stops, the aperture stop and the field stop. And then it shows their images with respect to the preceding and succeeding elements, that are called the entrance and exit pupil and windows, respectively. I realize this is a bunch of terms that I threw at once. So treat this as an information session, then go back and make sure you read them again from the notes and from the book so you make sure they sort of solidify in your mind.

But this diagram, actually, I spent a lot of time making this diagram because I actually drew it by hand. Instead of using an optical software, I did it by hand, so it took me a lot of time to make it accurate. But, to the best that I can-- I tried, it's pretty accurate, except for this lens here. This lens doesn't look very good. It's a negative lens. It should have sort of diverged this ray bundle even more. If I had done this properly, this would then become monstrous, so I didn't do it. But, anyway.

But you can see how it works. Here's the chief ray. Let's trace it through the system. So here's the chief ray. It goes through the center of the entrance pupil. Then it goes through the edge of the field stops. That means that I chose a chief ray that propagates at an angle equal to what? Equal to the field of view, right? Because this is the last chief ray admitted through a system. If I put a chief ray at a bigger angle, that is a taller object, it would miss the field stop.

Then the same chief ray goes on. It goes through a center of the aperture stop again, then again through the center of the exit pupil, and finally hits the edge of the exit window, which is also the tip of the image. That's the chief ray. Now let's go through the marginal ray. The marginal ray will hit the rim of the entrance pupil, then will hit the rim of the field stop again. Why? I'll give you 30 seconds to think about it. Why does the marginal ray hit the rim of the field stop again? My computer has actually a timer, so I can count the 30 seconds very precisely.

AUDIENCE: [INAUDIBLE]

GEORGE BARBASTATHIS: Oh, there is audio, so you can even speak your mind now. Thanks, by the way, for-- thanks, control, for fixing it. Anybody? Why does the marginal ray also go through the rim, the edge of the field stop?

AUDIENCE: Because the entrance window is an image of the field stop? Or it's the other way around. So the marginal ray is-- I mean, it's the edge of window, so it's going to be on the edge of the field stop also.

GEORGE BARBASTATHIS: That's exactly right. These two planes are images of each other. So since the marginal ray meets the chief ray here, then it should also meet the chief ray here again. Another way to put it is that this is actually an intermediate image plane for this particular optical system. So then the-- in other words, you would see an image of the object itself here if you were to stick your eye.

Then the imaginary ray will go again. Now it goes again through the edge of the aperture stop again. And that is, of course, because it went through the edge of the pupil, and since the pupil and the aperture stop are images of each other, then it has to go through the edge again. Now it will become repetitive. It will hit again the rim of the exit pupil for the same reason. And then, finally, it would hit the edge of the exit window because the exit window is an image of the field stop.

So these multiple imaging requirements, they sort of make this all makes sense, right? But this is how they work. The marginal ray, as the name suggests, hits the rims of everything-- the aperture stops, the pupils, the windows, and all that. The chief ray, it goes through the center of the pupils, entrance and exit, and the aperture stop, and hits the edges, the rims of the field stop and the associated entrance and exit windows.

And, finally, now that we have everything together, here is again the definition of the field of view. It is angle between the last two chief rays admitted by the system, last two meaning that it is the last two that the field stop would admit. And the numerical aperture is the angle between the two marginal rays that depart from the center, from the center of the window, that is, from the optical axis itself.

We do not define the numerical aperture from off-axis points. The reason is that, in poorly designed optical systems-- first of all, let me say, in an ideal optical system, the numerical aperture should be equal no matter where you are. Within the field of view, the numerical aperture should be the same. In practice, that is very difficult to do, and if you do a particularly poor job as a designer, the numerical aperture can decrease drastically when you go off axis here.

In fact, in the next example that I will show, this will happen. The numerical aperture becomes drastically smaller as you go off axis. That is called vignetting. It's one of those French words that have crept in optics and is difficult to pronounce. You will see it written in a second. I will just write it down here, vignetting. It is very annoying, because if you have vignetting in an optical system, then the edges of the image, they appear darker. Remember, the numerical aperture is the amount of light that the optical system admits. Well, if I admit more light in the center, I have a larger numerical aperture in the center.

A numerical aperture decreases as I go off axis, it means that the edges of the image will be darker. It will be dimmer. And that is, of course, very undesirable. And, generally, when we design microscopes, telescopes, and such, we try to make sure that this will not happen. But, anyway, because this could happen, that's why the numerical aperture is defined at the center, on axis, so that it has an unambiguous definition.

So having said all that, let's look now at an example that will actually reveal this vignetting phenomenon, in addition to other things. So this example is relatively simple. It is just a single thin lens. And then, in front of the thin lens, we've put an aperture stop. You can think of it-- actually, this is a very simplified model of the eye, of the human eye. We have a lens in front of it. In the human eye, we call it the pupil, right? Of course, in the human eye, in actuality you also have an additional lens in front of the pupil called the cornea. But, anyway, that's why this is perhaps not an ideal model. But, anyway, that's the example that I chose to do.

And, in addition, it has a field stop which is located at the image plane. Now, generally, if you see a stop at the image plane, this is a good candidate to be the field stop, right. Because, clearly, if I image anything outside this opening-- this is what aperture means, by the way. Literally, aperture means opening. So if I image anything outside this opening, then it will hit the block, and therefore it will not be image. So, therefore, this is very clearly the field stop.

Physically, the field stop is, for example, as I mentioned earlier, it is the size of the chip that registers the digital image, right? It's a very simple instance of a field stop. In old-fashioned cameras that use film, then the field stop would be, in that case, the film itself. Wherever the film is chemically active, this is where you record the image. Outside, it's just lost. Whatever is imaged outside that area is lost.

So, in this case, it is pretty clear who is the aperture and who is the field stop. So you can see very clearly that if you take an object point on the axis, then the physical aperture here will block the angle of acceptance from this particular point. And, therefore, this is the aperture stop. And, also, because there is no elements to the left, this is also the entrance pupil. If there is no element to image it through, the aperture stop will also be, simultaneously, the entrance pupil.

What this means is that if you look at the imaging instrument, you will actually see the pupil itself before you ever see anything else. So that's why the two things are actually coincident in this case. What about the exit pupil? Well, the exit pupil is what I will see on this side if I image this element through the succeeding elements. So the succeeding element is this lens here.

So what will I see if I image this through the lens? Well, this element now-- forget about the actual object of the system. When we question-- when we ask what are the entrance and exit pupils, this is now our object. We're trying to see where this would be imaged through the optical system. And we notice is that this is between the focal point and the lens. Here's the focal point. Here's the lens. This is in between. Therefore, this will form a virtual image. So the exit pupil would be somewhere out here. It would be actually to the left of the lens.

That is not surprising. All it means is that if I go here and then look at the optical instrument this way, what I will see, I will see a virtual image of the pupil, very similar to the virtual images that you saw last time in the demo in the class when Pepe brought the magnifier objective. So that's the exit pupil over here.

And what about the marginal and chief rays? Well, here is the-- here they are. Here's the marginal ray. Here's the chief ray. The chief ray, by definition, has to go out through the center of the aperture stop. So that is indeed. The marginal ray hits the rim. And what about the field of view? Well, the field of view will be defined by the pupil-- I'm sorry. It would be defined by the field stop.

Now, notice that the field stop is also the exit window because there is no optical elements to its right. So, therefore, that is the exit window. The entrance window is at the object plane, since this is the object plane and this is the image plane that satisfy an imaging condition. That is, if you image the field stop through the preceding optical elements, you'll arrive here. So that's the entrance window.

And, of course, the field of view is defined by the two chief rays that just clear the field stop, so here they are. And the angle between the two is the field of view. By the way, there is one ambiguity in the literature. Some people define the field of view as the full angle. Some people define it as the half angle. It's a matter of taste, how you want to define it. But just make sure you tell people what you really mean, if you really mean the full thing or not.

And here they're all together, all these quantities. And-- yes? Oh, let me ask-- yes, never mind. And, in this case, I actually worked out the math. I will not go through it here. I will let you do it by yourselves. If I put some symbols, for example if I call the size of the aperture stop, I call it a. The size of the field stop, I call it w prime. Then I have distance z between the aperture stop and the lens, distance s prime between the lens and the field stop.

And then this calculation shows you where-- how you find the other elements of the system. For example, when it comes to the entrance pupil, what should you to do? Well, for the entrance pupil, we have to consider this optical system. If that's our aperture stop, then I have distance z. Then I have the lens. And then-- well, I'm sorry. The entrance pupil is the-- I should have said the exit pupil.

So where is the exit pupil? Well I have to image this aperture stop through the system. So how do I image it? Well, let's say that the exit pupil is at some distance z prime to the right of the lens. Well, how do we find z prime? We apply the imaging condition, 1 over z plus 1 over z prime equals 1 over f. We can solve that. 1 over z prime equals 1 over f minus 1 over z. This is also known as z minus f over f times z. Or z prime equals f times z over z minus f.

Now, I said that the aperture stop is between the focal point and the lens. That means that z is less than f in this case. Therefore, this quantity-- since z is less than f, this quantity turned out to be negative. If it is negative, it means that my assumption before that the exit pupil is to the right was incorrect. My assumption would have been correct if z prime had turned out to be a positive quantity. Since this is turning out to be negative here, it means that the actual exit pupil is to the right. I'm sorry. It is to the left, opposite from my assumption. So, therefore, the exit pupil is over here, as I already said, or as I already said before.

And, basically, this result, these are all the quantities here. You can see the marginal ray from the field edge, the marginal ray from on axis, the chief ray, and so on and so forth. And the last comment here, I did not really prove it mathematically. You can do it if you like. It's a bit of a more involved calculation.

But even visually here, you can see that the angle between marginal rays on axis is, even in my handwritten diagram here, it is visibly larger than the angle between the two marginal rays off axis. This means that this particular system here is subject to vignetting that I mentioned earlier. That's the term here, vignetting, which means it is not a very well-designed optical system.

In sophisticated optical systems, there's rules that we follow in order to avoid vignetting. When I describe the microscope next week-- I mean, not next week. Next Wednesday, in 2 days from today, I will show you how microscopes are designed in order to avoid vignetting. And I see that I've already gone 2 minutes over time. So let me conclude, unless there's any questions that someone will ask, now that we have sound.