[MUSIC PLAYING]

ALAN OPPENHEIM: Hi. In this lecture, there are two sets of ideas that I'd like to discuss, both of which are related to our topic of the last several lectures, namely, the Z-transform. The first topic that we will focus on is what I'll referred to as the geometric interpretation of the frequency response. Now, you recall that in lecture 5 when I introduced the Z-transform, we observed the relationship between the Z-transform and the Fourier transform. The notion behind the geometric interpretation of the frequency response is to relate, in terms of a geometric picture, the rough characteristics of the frequency response of a system and the pole-zero pattern for the system function for the system.

The second topic that we'll focus on and eventually tie back into this geometric interpretation is the general issue of some of the properties of the Z-transform. And we'll illustrate some of these properties with some examples. But first of all, let's turn our attention to the subject of the geometric interpretation of the frequency response. And let me begin by reminding you of one of the important properties of the Z-transform for a linear shift invariant system, namely the fact that if we have a linear shift invariant system with a unit sample response, H of n, input x of n, and output y of n, then of course, y of n is the convolution of H of n with x of n. And the Z-transform of the output is the product of the Z-transform of the input and H of Z, the Z-transform of the unit sample response, which is what we've been referring to as the system function.

And in first developing the Z-transform, we tie together the notion of the Z-transform and the Fourier transform, in particular, the important interpretation that the Fourier transform is the Z-transform evaluated on the unit circle. So in terms of the system function H of z, if we evaluate that on the unit circle for z equal to e to the j omega, the result is the frequency response of the system, H of e to the j omega. Simply a statement that the Z-transform evaluated on the unit circle is the frequency response of the system.

Well, in terms of a simple example, we drag out our usual simple example. The system function is 1 over 1 minus az to the minus 1 or z divided by z minus a. And in terms of the z plane representation, or pole-zero representation in the z-plane, we then have a representation of this rational function in terms of one pole-- the pole at z equals a-- and one zero, that is the zero at z equals 0.

Well, to get the frequency response associated with that system function, we want to look at the Z-transform evaluated on the unit circle. We can do that geometrically by interpreting this complex number as a vector in the z plane and this complex number as a vector in the z plane, in which case the magnitude of H of z will be the ratio of the magnitudes of those vectors. And the angle of H of z will be the angle of this vector minus the angle of this vector.

In other words, suppose that we want to look at H of z, at some value of z. And let's say, for the moment, not on the unit circle, or it could be on the unit circle. Let's just pick a general value z equals z1. We have a vector corresponding to this complex number, which is the vector, or complex number, z. And that's a vector going from the origin out to this point. We have one vector, which is the vector z or z1 because we're evaluating this at z equals z1. So this is a vector z1. And the vector z corresponding to the complex number z minus a is the vector z or z1 minus the vector a. The vector z1 is this vector. The vector minus a is this vector. And the sum of those two, the vector that corresponds to the complex number z minus a, is the vector whose tail is at this pole and whose head is at the value of z, at which we want to evaluate H of z.

And so we end up then with an interpretation of the value of H of z in polar form with its magnitude equal to the magnitude of this vector divided by the magnitude or length of the pole vector. And the angle of this complex number, H of z at z equals z1 is the angle of the 0 vector minus the angle of the pole vector. That's a general statement. And z1 can be any place in the z plane.

So we have then the statement as I've just made it, which we would now like to apply to the interpretation or the generation of the frequency response of a system. So we have, then, our example that is one pole at equals a and 0 at z equals 0. We have the 0 vector and the pole vector. And as we generate the frequency response for z1 equal to e to the j omega, we have the point on the unit circle at which we're evaluating the Z-transform changing. If we think of angular distance around the unit circle as omega, then as omega varies and we consider the relative behavior of these two vectors, what we essentially trace out is the frequency response of the system.

So let's look at this and for this example see if we can get, for this example, a rough idea of what the frequency response should look like. Well, first of all we observe. And this is an important observation, since it's a point that comes up frequently, is that as we travel around the unit circle, the vector from the 0 to the unit circle changes in angle, of course. But it doesn't change in length. The 0 vector for this example, or in fact any vector from the origin out to the unit circle, as we vary omega and, therefore, trace around the unit circle, this vector doesn't change in length.

So if we're interested in, say, the magnitude of the frequency response, any poles or zeros at the origin, of course, have no effect on the magnitude. The only effect that they have is on the phase of the frequency response that is on the angle of H of z. And in fact, as we sweep around the unit circle, what happens to this angle? Well, that angle is just simply equal to the value omega that we're looking at that is the frequency value omega.

So then, in fact, poles or zeros at the origin as, of course, we could see from a strictly algebraic argument. But from a geometric argument, it should be clear that poles or zeros at the origin introduce, in the frequency response, a linear phase term. And they have no effect on the magnitude since the length of a vector from the origin to the unit circle is obviously 1. As a consequence of that, by the way, often when talking about the frequency response or equivalently when looking at pole-zero patterns, it's common to minimize or ignore the presence of zeros or poles at the origin since they just correspond to a linear phase term.

Now, let's look at what the behavior of the pole vector is. The pole vector as we start, say, at omega equals 0, we have a vector going from the pole to this point on the unit circle. And as we sweep around in frequency until we get to omega equals pi halfway around the unit circle, the length of this vector is monotonically increasing. Well, what does that mean about the magnitude of the frequency response? It means that the magnitude of the frequency response, which should be the magnitude of that vector-- the magnitude of the frequency response is monotonically decreasing as we sweep from omega equals zero around to omega equals pi.

At omega equal to pi, we have exactly the reverse situation. The length of the vector from the pole-- and let me draw that vector-- from the pole to omega equals pi as we sweep around back to omega equals 0, the length of that vector decreases. So the magnitude of the frequency response increases.

So for this particular example, by just observing what happens to the length of this pole vector as we trace around the unit circle, we can see that the frequency response starts at some value. That's obvious. And as it sweeps around to pi, it monotonically decreases. And when we come back from pi around to 0, the frequency response increases again. And I don't know if I've drawn it so that it looks clearly this way. But in fact, this piece is just this piece reflected over.

There's another important point that we can observe geometrically. It's a point about the frequency response that, of course, we've emphasized several times in several ways. This just offers an opportunity to re-emphasise it in yet another way.

As the frequency variable goes from 0 around to 2pi, we trace out a magnitude and then, of course, also a phase for the frequency response. When we get back to 0 or 2pi, if we go from 2pi around to 4pi, what should we see for the frequency response? Well, we should see exactly the same thing that we saw before because what we're doing literally, actually, is going around in circles.

We started here, went around there to get the 2pi, omega equals 2pi. For omega from 2pi to 4pi, we go around again. For 4pi to 6pi, we go around again, et cetera. So just geometrically by looking at this geometric interpretation of the frequency response, it gives us an opportunity to emphasize once again that the frequency response is a periodic function of frequency.

Now, this is for one specific example. The more general statement, then, is that the magnitude of the frequency response is equal to the product of the length of the 0 vectors-- the vectors from the zeros to the unit circle-- divided by the product of the length of the pole vectors-- the vectors from the poles to the unit circle. And the angle of the frequency response is the sum of the angles of the 0 vectors minus the sum of the angles of the pole vectors. Incidentally, it usually is the case-- or it is for me, anyway-- that often, it's possible to get a rough picture of what the magnitude of the frequency response is like by looking at this geometric picture. It's usually somewhat more difficult except, perhaps, for linear phase terms, et cetera, to get a very clear picture of what the phase looks like.

Well let's just look at one more example, which is a kind of example that we haven't discussed explicitly up to this point. And this also provides us with an opportunity to introduce this idea. Let's consider the case of a Z-transform which has a pole-zero pattern consisting of a pair of complex conjugate poles in the z-plane. These are complex conjugate poles with a radius equal to some value, say r, and an angular spacing equal to omega 0. So this pole is at omega 0. This pole is at an angle of minus omega 0. And using this notion of interpreting the frequency response geometrically, let's just sketch out roughly what we would expect the frequency response to look like.

Well, let's see. If we start at omega equals zero, then we have a vector from this pole and a vector from this pole. The resulting value of the frequency response is 1 over the product of those two vectors. As we move around the unit circle from omega equals 0, let's say, to some point that's closer to omega 0, then this vector gets changed to that one. And this vector gets changed to this one.

Well, we can see what's happening to the vector lengths. This vector is getting shorter as we approach omega equal to omega 0. This vector is getting longer. But you can imagine-- and I think it should be relatively clear from the picture-- that this vector is getting shorter faster than this one is getting longer. It seems relatively clear from a geometric picture that the product of the lengths of these two vectors is smaller when omega is closer to omega 0 than, let's say, when omega is equal to 0.

So consequently, we would expect the magnitude of the frequency response to start at some value. As we get in the vicinity of omega equal to omega 0, we would expect that frequency response to peak. As we pass that pole, the frequency response will now begin to decrease. And clearly, when we are at omega equal to pi, the product of the lengths of these two vectors is considerably longer than the products of the lengths of the vectors at omega equals 0.

So on the basis of that argument, we can see that roughly what we expect the frequency response to do is begin at some value, increase as we approach the pole at omega equals omega 0, decrease as we pass the pole heading toward pi. And then from pi around to pi, or equivalently from 0 around to minus pi, we would expect the reciprocal behavior. That is, the frequency response would increase, peak in the vicinity of this pole in the lower half of the z-plane, and then return to the same value that it started from. And of course, it'll be periodic as we run around the unit circle.

Well, this is a resonant characteristic, reminiscent in the analog case of what we would expect from a complex conjugate pole pair close to the j omega axis that's in the s-plane. And in fact, of course, in the s-plane, we have exactly the same kinds of geometrical arguments to allow us to roughly sketch out the frequency response. The only real important difference is that in the discrete time case, it's the unit circle, which is the locus in the z-plane that we're looking at.

In the continuous time case, it's the vertical, or j omega axis in the s-plane that we're looking at.

Well, this shouldn't be particularly evident from what we've done here. But by inference or by carrying your intuition from the continuous time to the discrete time, can you guess at what you would expect the unit sample response of a system with this pole-zero pattern to look like? Well, it would look like the discrete time counterpart of what happens in the continuous time case. That is, it would look like a damped sinusoidal sequence with the damping influenced by the distance of these poles from the unit circle. The closer the pole is to the unit circle, the sharper this resonant peak will be and the less damping on the sinusoid.

One additional point to remind you of-- I made reference to this in an earlier lecture. But notice that in talking about the Z-transform poles and zeros, et cetera here, I didn't say what the region of convergence was. Or at least, I didn't say explicitly what the region of convergence was. But did I say implicitly what it is?

Well, sure I did because what I said or what I've been assuming is that the systems that we've been talking about can be described in terms of a frequency response. In other words, the unit sample response has a Fourier transform. Well for that to be the case, the region of convergence has to include the unit circle. And then, we can use all those other rules of regions of convergences to allow us to figure out from there how far on both sides of the unit circle the region of convergence extends.

All right, well this is, then, a geometrical interpretation of the frequency response. It's often useful as a rough guide in getting a general picture of what the frequency response might look like. Although, for complicated cases, it's often difficult to get precise details about the frequency response, which of course, we could also get algebraically.

Well now, I'd like to, at least temporarily, change gears or topics and talk about the issue of properties of the Z-transform in talking about the properties and working some examples to illustrate the use of the properties of the Z-transform. In fact, I'll have occasion to make reference back to this geometric interpretation of the frequency response to help with at least one of the examples. So now I'd like to turn our attention, then, to the question or the topic of the properties of the Z-transform.

Well, first of all, why do we want properties of the Z-transform? The Z-transform has properties. Why do we want these properties? One of the reasons-- it's a very practical reason-- is that the properties of the Z-transform help us in calculating Z-transforms and inverse Z-transforms. And they also, obviously, provide a certain amount of intuition and insight with regard to generally dealing with Z-transforms and their inverses.

Well, there are a lot of properties. In fact, there are trivial properties. There are very complicated properties. There are properties that we more commonly tend to carry around. And I've listed a few that shouldn't be considered to be exhaustive but generally tend to be the properties that turn out to be the handiest. That is, these are the properties, at a minimum, that you should carry around in your back pocket.

Well, let's see. We're talking about a sequence, x of n with a Z-transform, x of z. One of the properties that we've taken advantage of already throughout the entire discussion of the Z-transform is the fact that the Z-transform maps convolution to multiplication. If I have the convolution of two sequences, then the resulting Z-transform is the product of the Z-transforms.

The second property, which is often very useful, is referred to as the shifting property, which says that if I shift x of n by an amount n0-- and you should think, by the way, of if n0 is positive, does that mean shifting to the right or shifting to the left? Well, I'll let you think about that. It's something that you should nail down. If I shift the sequence x of n by replacing the argument by n plus n0, then the resulting Z-transform is z to the n0 times x of z. That is, shifting corresponds to multiplying by z to the n0.

Another useful property has to do with taking a sequence and turning it around in n. That is, replacing n by minus n. The resulting effect on the Z-transform is to replace z by 1 over z. Another useful property is the result of multiplying a sequence x of n by an exponential, a to the n. a might be complex. Or it might be real. And the result there is that the z transform is x of 8 to the minus 1 times z.

Another one which is useful is multiplication of a sequence by n which results in a z transform, which is minus z times the derivative of x of z. And then I've indicated-- tried to be somewhat explicit-- that that's not the end of the list. There are lots of other properties-- a number of others that are presented in the text, others besides that that aren't presented in the text, ones that you can dream up yourself, ones that your friends know that you don't, et cetera.

Now, we could, of course, go through the proof of all these. The proofs of properties tend to all be in somewhat of a similar vein, as a matter of fact. And the style, once you see what the trick is or roughly how you go about proving properties, then you can just prove properties and prove properties. And we won't do that in this lecture, with the exception of illustrating the style of proving properties with a couple of examples.

And the two examples that I've picked, somewhat arbitrarily as a matter of fact, is the Shifting Property, that is Property 2, and the result of multiplication by an exponential, which is Property 4. But this is only to illustrate the style of proving properties. And you can guess where, actually, you get a chance to see the proof of some of the others.

All right, let's take a look at Property 2, that is the Shifting Property. Well, to prove that replacing n by n plus n0 results in z to the n0 times x of z, we want to consider a sequence, x1 of n equal to x of n plus n0 so that its Z-transform, x1 of z, is the sum of x of n plus n0 times z to the minus n. Well, a simple idea here is a substitution of variables. Let's replace n plus n0 by a new variable, m, or n is equal to m minus n0, in which case, we can rewrite this expression as x1 of z is the sum on m of x of m, because this is now m. n is replaced by m minus n0. So we have z to the n0 times z to the minus m.

Well, the z to the n0 can come outside the sum. The limits on the sum, incidentally, are m equals minus infinity to plus infinity because if n0 is finite, as n runs from minus infinity to plus infinity, so does m. That comes outside the sum. And what's left, then, is the sum of x of m z to the minus m, which is just x of z. So consequently, then, what we end up with is that x of z is z to the n0 times x of z. x1 of z is z to the n0 times x of z, which is, of course, the way I advertised.

Well, there are lots of times when, in fact, the Shifting Property comes into play. One thing that we can use it for immediately is to tie together a couple of things that have been floating, more or less, in the background. I have alluded several times to the fact that systems whose system function is rational correspond to systems that are characterized by linear constant coefficient difference equations. That is, linear constant coefficient difference equations, those are the systems that end up with system functions that are rational functions. And in fact, we can see that in a straightforward way by just simply applying the Shifting Property.

Well, let's consider a linear constant coefficient difference equation of the general form is the sum from k equals 0 to n. a sub k y of n minus k is equal to a sum of b sub k times x of n minus k. If we take the Z-transform of both sides of this equation and use the fact-- incidentally, this is another property that I've more or less been using, although I've never stated explicitly-- that the Z-transform of a sum is the sum of the Z-transforms. Taking, then, the Z-transform of this equation, y of n minus k using the Shifting Property will give us z to the minus k times y of z. And x of n minus k will give us z to the minus k times x of z.

Consequently, the sum of a sub k z to the minus k y of z is equal to the sum of b sub k z to the minus k x of z. If we solve that equation for y of z over x of z, which is the system function, then we end up with the sum of b sub k z to the minus k divided by the sum of a sub k times z to the minus k.

And there are a couple of important observations. One is we ended up with a rational function, which is what we expected, or what I said we were going to get. And a second is that the coefficients in the numerator polynomial are exactly the same as the coefficients on the right hand side of the difference equation. And the coefficients in the denominator polynomial are exactly the same as the coefficients on the left hand side of the difference equation. This by the way, is exactly consistent, or analogous, with what happens when we apply the Laplace transform to linear constant coefficient differential equation. Exactly the same thing happens.

It should be clear then, incidentally, that if I give you a system function that's a rational function of z, that you could construct, in a straightforward way, the difference equation that characterizes that system because you can pick the coefficients off from the numerator. Those are on the right hand side. And you can pick the coefficients off from the denominator. Those are on the left hand side of the difference equation.

The Shifting Property is a property that arises very often and, in fact, is a very, very useful property. Now, the second property that I want to outline the proof for is the property related to multiplication of the sequence by an exponential a to the n. I am forming a new sequence x1 of n, which is a to the n times x of n. And to derive the relationship between x1 of z and x of n, again, we can look at the Z-transform, x1 of z, which is the sum of a to the n, x of n, times z to the minus n.

Well, it's straightforward to rewrite what's inside the sum as x of n times a to the minus 1 times z to the minus n. x1 of z is the sum of x of n times a to the minus 1z-- all that raised to the minus n. Well, that looks just like the Z-transform of x of n but with z replaced by a to the minus 1 times z. So this says, consequently, that the Z-transform of x1 of n is equal to the Z-transform of x of n but with z replaced by a to the minus 1 times z. So we stick in here a to the minus 1 times z. And that then relates the Z-transform of the original sequence and the Z-transform of the sequence multiplied by a decaying exponential.

Well, it's interesting to look at what this property implies in terms of the movement of the poles and zeros in the z-plane. That is, a useful notion or a useful fact to have, again, stored away in your hip pocket is the effect on the poles and zeros of a system function or a Z-transform of multiplying the sequence by an exponential-- maybe a complex exponential, maybe a real exponential. Well, let's focus on a pole or a 0. And the result that we get by considering just a simple pole or 0 will, of course, generalize to all the poles and zeros.

So let's consider x of z to have a factor either in the numerator and denominator of the form z minus z0. Well, x1 of z will then have a factor derived from that one, but for which z is replaced by 8 to the minus 1 times z. And then we have the minus z0. Or if we pull the a to the minus outside, we have a to the minus 1 times z minus a times z0.

So there are two effects here. One effect-- if we think of x of z as a product of zeros divided by a product of poles-- one effect is that there is a constant that, perhaps, collects out in front. But we never see that constant anyway in the pole-0 pattern. The more important point is that whereas here we had, say, 0 at z equals z0, that 0 is now shifted to a times z0. So the 0 at z0 gets replaced by a 0 at eight times z0. A pole at z0 gets replaced by a pole at a times z0.

Well, more specifically, then, here's our polar 0 which gets converted to that polar 0 multiplied by a. And if we write this or think of it in polar form, as r0 e to the j theta 0, then the result is a times r0 times e to the j theta 0. Well obviously then, if this number a is a real number, then the only effect on the location of the pole is to change its radial value and not change its angle. More generally, if a is complex, then to write the resulting polar zero in polar form, we would replace this by its magnitude and add to the phase-- the phase that corresponds to that number a.

Well consequently, first of all, if a is real and positive, actually, if a is real and positive, then if we had a pole of x of n say here, here being any place, then if a is real, what happens to that pole is that it moves either in or out. But it moves radially as we vary the value of a. So that's for a real.

What's the movement of the pole if a is pure imaginary? Well, if a is pure imaginary, then the magnitude of a is-- I'm sorry. Not if a is pure imaginary, but if the magnitude of a is equal to 1 and it has only a phase component, in that case there's no effect on the radial value. There's only an effect on the angle. And in that case, the movement of the pole is such that the radial value stays the same. But the angle of the pole changes.

So in general, of course, if a is complex, the poles can have a little movement that way and also a little movement that way. In some cases, if a is pure real, the movement will just be this way. And if the magnitude of a is equal to 1, the movement of the pole will just be that way.

All right, so we have first of all, a list of properties. But in particular, there are two that we've spent a little time on. Finally, let's look at, actually, one example or one and a half examples to see how some of the properties might be useful in obtaining the frequency response, or the Z-transform of a system.

And the sequence that I want to focus on is a sequence that will play an important role throughout digital signal processing and in particular, in some lectures coming up, which is a sequence that I'll refer to as a boxcar sequence. It's the sequence which is equal to unity for n between 0 and capital N minus 1. And it's equal to 0 otherwise. It's basically a rectangular sequence. 0 for n negative, 0 for n greater than or equal to capital N and unity otherwise. It's the counterpart of the rectangular time function.

Well, there are a couple of ways of getting at Z-transform. Since we had some properties, let's use one of them, in particular, the Shifting Property. We can think of a boxcar sequence as the sum of two sequences. One is a unit step. And the second is the negative of a unit step starting at n equals capital N to subtract off these other values. We can express x of n, the boxcar sequence, as a unit step minus a unit step delayed by capital N.

Well, if you don't see this graphical picture exactly, you can just see quickly that this is true since for n greater than or equal to capital N, both of these arguments are non-negative. So the value of both of these units steps is unity. And they subtract off to zero.

All right, then to get the Z-transform, we can add the Z-transform or this piece and this piece. The Z-transform of a unit step, well that's our old friend a to the n times a unit step, except in this case, a equals 1. So the Z-transform of this piece is 1 over 1 minus z to the minus 1.

This one is this one shifted. So we can apply our Shifting Property to multiply this by z to the minus capital N, since that's the amount of our shift, so that this piece, then, has a Z-transform z to the minus n over 1 minus z to the minus 1. Or if we add these two together, we have 1 minus z to the minus capital N divided by 1 minus z to the minus 1. Or we can rewrite that, just to focus on something a little more clearly, multiplying top and bottom by z to the n minus 1.

We can rewrite that in this form so that we have in the numerator z to the capital N minus 1, in the denominator, this factor times z to the minus 1.

Well, let's look at the pole-zero pattern. First of all, at z equals 1, we have a pole. So there's that pole. Second of all, at z equals 0, we have n minus1 poles. Well, let's stick in the n minus 1 poles. And let me just draw that with an asterisk and an indication that that's n minus 1 poles. That's from that term.

And the zeros, well, where are the zeros? They're the roots of the numerator, which are at the N roots of unity. Where are they? They're distributed. The N roots of unity are distributed around the unit circle equally spaced in angle starting at z equals 1. So there's a zero at z equals 1. But there's also a pole it equals 1. So in fact at z equals 1, there's neither 0 nor a pole because the two cancel out.

If we look at the other zeros, then let's take a specific case, that is n equals 8. We'd expect to see eight zeros, except for the 1 at z equals 1 that got canceled out. So there are seven left-- one there, there, there, there, here, here, and here. Then, there was the one at the origin. Let me just indicate that and the fact that it got canceled out by a pole. So in fact, that one isn't there. So the pole-zero pattern for the boxcar sequence, then, is n minus 1 poles at the origin plus zeros equally spaced in angle, but with the 1 at z equals 1 missing.

Now, what does this mean in terms of the frequency response? Well, we can very quickly generate the frequency response, or a rough idea of the frequency response, geometrically by referring back to the set of ideas that we introduced at the beginning of the lecture and ask what the behavior of the pole-zero vectors are as we go around the unit circle. The pole vectors, first of all, introduce only a linear phase term and have no effect on the magnitude. We had agreed on that before. And so for the magnitude, we only need focus on the zeros.

Well, one thing is obvious and that is that as we go around the unit circle, the frequency response is obviously 0 when we hit each one of these zeros. That corresponds to pi over 4, pi over 2, pi over 2 plus pi over 4-- whatever that is, pi, and the next increment of pi over 4, et cetera. Furthermore, you can see that at least, it's not implausible. Actually, we really could argue this somewhat precisely, that at omega equals 0, that's the place where we're the farthest away from all the zeros. As we start moving around the unit circle, if we're in-between two of these zeros, maybe we're a little farther away from one of them, but we're closer to another one. And in terms of the product of the length of the zero vectors, that will tend to stay smaller than the product of the lengths of the zero vectors at z equals 1.

Consequently, the frequency response starts at some value and of course, goes down to 0 at pi over 4. That's because of this 0. Then, it comes back up again, but not quite as far, and then goes back down to 0 at pi over 2. Then it goes up again and comes down again. It goes up not quite as far. And that's not particularly obvious geometrically. And then, the same thing again going to 0 at pi. And then of course as we come back around the unit circle, we see the same thing, the same type of behavior, repeated again. So roughly, we can get a geometric picture of the frequency response for a boxcar sequence by looking at the location of the poles and zeros in the z-plane and the behavior of the vectors as we travel around the unit circle.

Well finally, let's just look at the Fourier transform of the boxcar somewhat more formally because in fact, it's an important sequence. And it's important to have a fairly complete precise statement of the Fourier transform and a complete picture of what it looks like. Well, let's see.

We had the Z-transform as 1 minus z to the minus n over 1 minus z to the minus 1. Substituting in z equals z to the j omega, then we have x of e to the j omega as this. We can factor out a factor e to the minus j omega capital N over 2, leaving e to the j omega capital N over 2 minus e to the minus j omega capital N over 2, and a denominator factor e to the minus j omega over 2 times e to the j omega-- over 2 minus e to the minus j omega over 2.

This piece we recognize as 2j sine omega capital N over to 2. And this piece we recognize as 2-j times sine omega over 2. And consequently, putting these two terms together and inserting this substitution, the Fourier transform is e to the minus j omega capital N minus 1 over 2-- that's a linear phase term, by the way-- times sine omega capital N over 2 divided by sine omega over 2.

And this function, sine omega capital N over 2 divided by sine omega over 2, is the discrete time counterpart of what we usually find in the continuous time case as sine x over x. That is, this is a sine nx over sine x kind of function. And it plays exactly the same role in the discrete time case that the sine x over x function plays in the continuous time case.

And that's not unreasonable, actually, because this arose by looking at the Fourier transform of a rectangular sequence, whereas the sine x over x function arises by looking at the Fourier transform of a continuous time rectangle. So the sine nx over sine x function, which is what we have here, is an extremely important function-- as important as sine x over x in the continuous time case. And consequently, let me just show you this function plotted out a little more precisely than I would be able to do at the board. In particular, let me show you a viewgraph which illustrates the function sine omega capital N over 2 divided by sine omega over 2. And I've sketched this now over more than just a 0 to 2 pi interval to, again, stress the fact that this is periodic-- as all Fourier transforms are-- periodic with a period of 2 pi.

The important characteristics of it-- some important characteristics. There are a lot. But some important characteristics of it are that it has an envelope. It's basically a sinusoidal function with an envelope that is the reciprocal of a sinusoid. The period of this sinusoid is from 0 to 2pi, whereas this one wiggles faster, depending on the value of capital N. I've sketched it here for n equals 15.

But in fact, it has a lot of the character of a sine x over x function. That is, it has a big central lobe, decays down and wiggles and gets smaller as it's decaying. But then, of course, the fact that it has to be periodic is what distinguishes it in the discrete time case from the sine x over x function in the continuous time case.

OK. Well, this concludes our discussion of the Z-transform. We've now talked about two transforms. We've talked about the Fourier transform and the Z-transform spread out over about five lectures. And in the next lecture, the next set of two or three lectures, we'll be talking about yet another transform, which is a transform that's really somewhat special and linked very closely to the notion of discrete time signals and discrete time signal processing. That transform is the Discrete Fourier transform. And besides being a mathematical tool, as the Fourier transform and the Z-transform have been, Discrete Fourier transform has some important computational realizations and computational implications that will be one of the important things that will want to capitalize on in applying digital signal processing to real problems. Thank you.

[MUSIC PLAYING]