PROFESSOR: Over the last several lectures, we've dealt with the representation of linear time-invariant systems through convolution. And just to remind you of our basic strategy, essentially, the idea was to exploit the notion of linearity by decomposing the input into a sum of basic inputs and then using linearity to tell us that the output can be represented as the corresponding linear combination of the associated outputs.

So, if we have a linear system, either continuous-time or discrete-time, for example, with continuous time, if the input is decomposed as a linear combination of basic inputs, with each of these basic inputs generating an associated output, and if the system is linear, then the output of the system is the same linear combination of the associated outputs. And the same statement is identical both for continuous time and discrete time.

So the strategy is to decompose the input into these basic inputs. And the inputs were chosen also with some particular strategy in mind. In particular, for both continuous time or discrete time, in this representation, the basic inputs used in the decomposition are chosen, first of all, so that a broad class of signals could be represented in terms of these basic inputs, and second of all, so that the response to these basic inputs is, in some sense, easy to compute.

Now, in the representation which led us to convolution, the particular choice that we made in the discrete-time case for our basic inputs was a decomposition of the input in terms of delayed impulses. And the associated outputs that that generated were delayed versions of the impulse response. Decomposing the input into a linear combination of these, the output into the corresponding linear combination of these, then led to the convolution sum in the discrete time case. And in the continuous-time case, a similar kind of decomposition, in terms of impulses, and associated representation of the output, in terms of the impulse response, led to the convolution integral.

Now, in this lecture, and for a number of the succeeding lectures, we'll want to turn our attention to a very different set of basic building blocks. And in particular, the signals that we'll be using as the building blocks for our more general signals, rather than impulses, as we've dealt with before, will be, in general, complex exponentials. So, in a general sense, in the continuous-time case, we'll be thinking in terms of a decomposition of our signals as a linear combination of complex exponentials, continuous-time, or, in the discrete-time case, complex exponentials, where z_k is complex here in discrete time and s sub k is complex here in continuous time.

Now, the basic strategy, of course, requires that we choose a set of inputs, basic building blocks, which have two properties. One is that the system response be straightforward to compute, or in some sense, easy to compute. And second is that it be a fairly general set of building blocks so that we can build lots of signals out of them.

What we'll find with complex exponentials, either continuous-time or discrete-time, is that they very nicely have those two properties. In particular, the notion that the output of a linear time-invariant system is easy to compute is tied to what's referred to as the Eigenfunction function property of complex exponentials, which we'll focus on shortly in a little more detail. And second of all, the fact that we can, in fact, represent very broad classes of signals as linear combinations of these will be a topic and an issue that we'll develop in detail over, in fact, the next set of lectures, this lecture, and the next set of lectures.

Now, in doing this, although we could, in fact, begin with our attention focused on, in general, complex exponentials, what we'll choose to do is first focus on the case in which the exponent in the continuous-time case is purely imaginary, as I indicate here, and in the discrete-time case, where the magnitude of the complex number z_k is equal to 1. So what that corresponds to in the continuous-time case is a set of building blocks of the form e^(j omega_k t), and in the discrete-time case, a set of building blocks of the form e^(j Omega_k n).

What we'll see is a representation in these terms leads to what's referred to as Fourier analysis. And that's what will be dealing with over the next set of lectures. We'll then be exploiting this representation actually through most of the course. And then toward the end of the course, we'll return to generalizing the Fourier representation to a discussion Laplace transforms and Z-transforms. So for now, we want to restrict ourselves to complex exponentials of a particular form, and in fact, also initially to continuous-time signals and systems.

So let's begin with the continuous-time case and the complex exponentials that we want to deal with and focus, first of all, on what I refer to as the Eigenfunction property of this particular set of building blocks. We're talking about basic signals of the form e^(j omega_k t). And the statement is that for a linear time-invariant system, the response to one of these is of exactly the same form, just simply multiplied by a complex factor, that complex factor depending on what the frequency, omega_k, is.

Now more or less, the justification for this, or the proof, follows by simply looking at the response to a complex exponential, using the convolution integral. So if we put a complex exponentials into a linear time-invariant system with impulse response h(t), then we can express the response as I've indicated here. We can then recognize that this complex exponentials can be factored into two terms.

And so we can rewrite this complex exponential as this product. Second, recognize that this term can be taken outside the integral, over here, because of the fact that it depends only on t and not on Tau. And so what we're left with, when we track this through, is that, with a complex exponential input, we get an output which is the same complex exponential, namely this factor, times this integral. And this integral is what I refer to above as H(omega_k).

And so, in fact, we put in a complex exponential, we get out a complex exponentials of the same frequency, multiplied by a complex constant. And that is what's referred to as the Eigenfunction property, Eigenfunction meaning that an Eigenfunction of a system, or mathematical expression, is a function which, when you put it through the system, comes out looking exactly the same except for a change in amplitude, the change in amplitude being the Eigenvalue. So in fact, this function is the Eigenfunction. And this value is the Eigenvalue.

OK, now it's because of the Eigenfunction property that complex exponentials are particularly convenient as building blocks. Namely you put it through the system, they come out with the same form and simply scale. The other part to the question, related to the strategy that we've been pursuing, is to hope that these signals can be used as building blocks to represent a very broad class of signals through a linear combination. And in fact, that turns out to be the case with complex exponentials.

As we work our way through that, we'll first consider the case of periodic signals. And what that leads to is a representation of periodic signals through what's called the Fourier series. Following that, we'll turn our attention to non-periodic, or as I refer to it, aperiodic signals. And the representation that's developed in terms of linear combinations of complex exponentials is what's referred to as the Fourier transform. So the first thing we want to deal with are periodic signals and the Fourier series.

So what we're talking about then is the continuous-time Fourier series. And the Fourier series is a representation for periodic continuous-time signals. We have a signal, then, which is periodic. And we're choosing T_0 to denote the period. So it's T_0 that corresponds to the period of our periodic signal. omega_0 is 2 pi / T_0, as you recall from our discussion pf periodic signals and sinusoids before. And that's 2 pi f_0. So this is the fundamental frequency.

Now let's examine, first of all, complex exponentials, and recognize, first of all, that there is a complex exponential that has exactly the same period and fundamental frequency as our more general periodic signal, namely the complex exponential e^(j omega_0 t), where omega_0 is 2 pi / T_0, or equivalently, T_0 is 2 pi / omega_0. Now that's the complex exponential which has T_0 as the fundamental period.

But there are harmonically related complex exponentials that also have T_0 as a period, although in fact, their fundamental period is shorter. So we can also look at complex exponentials of the form e^(j k omega_0 t). These likewise are periodic with a period of T_0. Although, in fact, their fundamental period is T_0 / k, or equivalently, 2 pi divided by their fundamental frequency, k omega_0. So as k, an integer, varies, these correspond to harmonically related complex exponentials.

Now what the Fourier series says, and we'll justify this bit by bit as the discussion goes on, what the Fourier series says, and in fact, what Fourier said, which was essentially his brilliant insight, is that, if I have a very general periodic signal, I can represent it as a linear combination of these harmonically-related complex exponentials. So that representation is what I've indicated here. And this summation is what will be referred to as the Fourier series.

And as we proceed with the discussion, there are two issues that will develop. One is, assuming that our periodic signal can be represented this way, how do we determine the Fourier series coefficients, as they're referred to, a_k. That's one question. And the second question will be how broad a class of signals, in fact, can be represented this way. And that's another question that we'll deal with separately.

Now just focusing on this representation for a minute, this representation of the Fourier series, which I've repeated again here, is what's referred to as the complex exponential form of the Fourier series. And it's important to note, incidentally, that the summation involves frequencies, k omega_0, that are both positive and negative. In other words, this index k runs over limits that include both negative values and positive values.

Now that complex exponential form is one representation for the Fourier series. And in fact, it's the one that we will be principally relying on in this course.

There is another representation that perhaps you've come across previously and that in a variety of other contexts is typically used, which is called the trigonometric form for the Fourier series. Without really tracking through the algebra, essentially we can get to the trigonometric form from the complex exponential form by recognizing that if we express the complex coefficient in polar form or in rectangular form and expand the complex exponential term out in terms of cosine plus j sine, using just simply Euler's relation, then we will end up with a representation for the periodic signal, or a re-expression of the Fourier series expression that we had previously, either in the form that I indicate here, where now the periodic signal is expressed in terms of a summation of cosines with appropriate amplitude and phase. Or another equivalent trigonometric form involves rearranging this in terms of a combination of cosines and sines.

Now in this representation, the frequencies of the sinusoids vary only over positive frequencies. And typically one thinks of periodic signals as having positive frequencies associated with them. However, let's look back and the complex exponential form for the Fourier series at the top of the board. And in that representation, when we use this representation, we'll find it convenient to refer to both positive frequencies and negative frequencies.

So the representation that we will most typically be using is the complex exponential form. And in that form, what we'll find as we think of decomposing a periodic signal into its components at different frequencies, it will involve both positive frequencies and negative frequencies.

Okay, now we have the Fourier series representation, as I've indicated here. Again, so far I've sidestepped the issue as to whether this in fact represents all the signals that we'd like to represent. Let's first address the issue of how we determine these coefficients a_k, assuming that, in fact, this representation is valid. And again, I'll kind of move through the algebra fairly quickly. The algebraic steps are ones that you can pursue more leisurely just to kind of verify them and step through them.

But essentially, the algebra develops out of the recognition that if we integrate a complex exponential over one period, T_0-- and I mean by this notation that this is an integral over a period, where I don't particularly care where the period starts and where the period stops, in other words, exactly what period I picked-- that this integral is equal to T_0 when m is equal to 0. And it's equal to 0 if m is not equal to

0.

That follows simply from the fact that if we substitute in for using or Euler's relation, so that we have the integral of a cosine plus j times the sine, if m is not equal to 0, then both of these integrals over a period are 0. The integral of a of a periodic of a sinusoid, cosine or sine, over an integral number of periods is 0. Whereas, if m is equal to 0, this integral will be equal to T_0, the integral of the cosine. And the integral of the sine is equal to 0.

Okay, well, the next step in developing the expression for the coefficient a_k is to refer back to the Fourier series expression, which was that $x(t)$ is equal to the sum of $a_k e^{j k \omega_0 t}$. If we multiply both sides of that by $e^{-j n \omega_0 t}$, and integrate that over a period-- both sides of the equation integrated over a period, so these two equations are equal-- and then in essence, interchange the summation and the integration so that this part of the expression comes outside the sum, and then we combine these two complex exponentials together, where we come out is the expression that I've indicated here. And then essentially what happens at this point, algebraically, is that we use the result that we just developed to evaluate this integral.

So multiplying both sides of the Fourier series and then doing the integration leads us, after the appropriate manipulation, to the expression that I have up here. And this integral is equal to T_0 if k is equal to n, corresponding to 0 up here. And it's 0 otherwise, which is what we had demonstrated or argued previously.

And the upshot of all that, then, is that the right hand side of this expression disappears except for the term when k is equal to n. And so finally, we have what I indicate here, taking T_0 and moving it over to the other side of the equation, that then tells us how we determine the Fourier series coefficients a_n, or a_k.

So that, in effect, then is what we refer to as the analysis equation, the equation that begins with $x(t)$ and tells us how to get the Fourier series coefficients. What I'll refer to as the Fourier series synthesis equation is the equation that tells us how to build $x(t)$ out of these complex exponentials. So we have the synthesis equation, which is the one we started from. We have the analysis equation, which is the equation that we just developed.

So we in effect have gone through the issue of, assuming that a Fourier series representation is in fact valid, how we get the coefficients. We'll want to address somewhat the question of how broad a class of signals are we talking about. And what's in fact amazing, and was Fourier's amazing insight, was that it's a very broad class of signals.

But let's first look at just some examples in which we take a signal, assume that it has the Fourier series representation, and see what the Fourier series coefficients look like. So we'll begin with what I refer to as an antisymmetric periodic square wave-- periodic of course, because we're talking about periodic signals; square wave referring to its shape; and antisymmetric referring to the fact that it is an odd time function. In other words, it is antisymmetric about the origin.

Now the expression for the Fourier series coefficients tells us that we determine a_k by taking 1 / T0 times the integral over a period of x(t), e^(-j k omega_0 t) dt. The most convenient thing in this case is to choose a period, which let's say goes from -T_0 / 2 to +T_0 / 2. So here x(t) is -1. Here x(t) is +1. And so I've expressed the Fourier series coefficients as this integral, that's from -T_0 / 2 to 0. And then added to that is the positive part of the cycle. And so we have these two integrals.

Now, I don't want to track through the details of the algebra again. I guess I've decided that that's much more fun for you to do on your own. But the way it comes out when you go through it is the expression that I finally indicate after suggesting that there are few more steps to follow. And what develops is that those two integrals together, for k not equal to 0, come out to this expression. And that expression is not valid for k = 0.

For k equal to 0, we can go back to the basic expression for the Fourier series, which is 1 / T_0, the integral over a period, x(t) e^(-j k omega_0 t) dt. For k = 0, of course this term just simply becomes 1. And so the zeroth coefficient is 1 / T_0 times the integral of x(t) over a period.

Now, going back to the original function that we have, what we're saying then is that the zeroth coefficient is 1 / T_0 times the integral over one period, which is, in effect, the average value. And it's straightforward to verify for this case that average value is equal to 0.

Now let's look at these Fourier series coefficients on a bar graph. And I've indicated that here. The expression for the Fourier series coefficients we just developed. And it involves-- it's 0 for k = 0, it's a factor of this form for k =/= 0. Plotted on a bar graph, then we see values like this, 0 at k = 0 and then associated values.

And there are a number of things to focus on when you look at this. One is the fact that the Fourier series coefficients for this example are purely imaginary. A second is that the Fourier series coefficients for this example are an odd sequence. In other words, if you look at this sequence, what you see are these values for -k flipped over. So they're imaginary and odd.

And what that results in, when you look at the trigonometric form of the Fourier series, is that in fact, those conditions, if you put the terms all together, lead you to a trigonometric representation, which involves only sine terms-- in other words, no cosine terms. Let me just draw your attention to the fact that, since a_k's are imaginary, this j takes care of that fact so that these coefficients are in fact real.

So what this says is that for the antisymmetric square wave, in effect, the Fourier series is a sine series. The antisymmetric square wave is an odd function. Sinusoids are odd functions. And so this is all kind of reasonable, that we're building an odd function out of odd functions.

As an additional aside, which I won't exploit or refer to any further here, but just draw your attention to, is that another aspect of this periodic square wave, the particular one that we chose, is that it is what's referred to as an odd harmonic function. In other words, for even values of k, the Fourier series coefficients are 0. They're are only non-zero for odd values of k.

Now let's look at another example. Another example is the symmetric periodic square wave. And this is in fact example 4.5, worked out in more detail in the text. Then I won't bother to work this out in detail here, except to draw your attention to several points. Here is the symmetric periodic square wave. And what I mean by symmetric is that it's an even time function. Now just kind of extrapolating your intuition, what you should expect is that if it's only an even time function, it should be built up or buildable, if it's buildable at all, out of only even sinusoids. And in fact, that's the case.

So if we look at the Fourier series coefficients for this, is zeroth coefficient, again, is the average value, which in this case, is 1/2. Here I've plotted pi times the Fourier series coefficients. So the zeroth value is pi / 2. The coefficients are now an even sequence, in other words, symmetric about k = 0. And the consequence of that is that when you take these coefficients and put together the equivalent trigonometric form, the trigonometric form involves only cosines and no sine terms.

Now you'll see this in other examples, not that we'll do in the lecture, but examples in the text and in the video manual, if in fact the square wave was neither symmetric or antisymmetric, then the trigonometric form would involve both sines and cosines. And that is, of course, the more general case.

Furthermore, in the two examples I've shown here, in both cases, the signal is odd harmonic. In other words, for even values of k, the coefficients are equal to 0. Although I won't justify that here, that's a consequence of the fact that this symmetry is exactly about half a period. And if you made the on time of the square wave different in relation to the off time, then that property would also disappear.

Now what's kind of amazing, actually, is that if we take a square wave, like I have here or as I had in the antisymmetric case, the implication is that I can build that square wave by adding up enough sines or cosines. And it really seems kind of amazing because the square wave, after all, is a very discontinuous function. Sinusoids are very continuous. And it seems puzzling that in fact you can do that.

Well let's look in a little bit of detail how the sinusoidal terms add up to build a square wave. And to do that, let's first define what I refer to as a partial sum. So here we have the expression which is the synthesis equation, telling us how x(t) could be represented as complex exponentials if it can be. And let's consider just a finite number of terms in this sum. And so x_n(t), of course, as n goes to infinity, approaches the infinite sum that we're talking about.

And although we could do this more generally, let's not. Let's focus on the symmetric square wave case, where because of the symmetry of these coefficients, namely that a_k is equal to a_(-k), we can rewrite these terms as cosine terms. And so this partial sum can be expressed the way I'm expressing it here.

Well let's look at a few of these terms. On the graph, I have, first of all, x(t), which is our original square wave. The term that I indicate here is the factor of 1/2, which is this term. With n = 1, that would correspond to adding one cosine term to that. And so the sum of those two would be this, which looks a little closer to the square wave, but certainly not very close to it at all. And in fact, it's somewhat hard to imagine without seeing the terms build up how in fact, by adding more and more terms, we can generate something that is essentially flat, except at the discontinuities.

So let's look at this example. And what I'd like to show is this example, but now as we add many more terms to it. And let's see in fact how these individual terms add up to build up the square wave.

So this is the square wave that we want to build up through the Fourier series as a sum of sinusoids. And the term for k = 0 will be a constant which represents the DC value of this. And so in the partial sum, as we develop it, the first thing that we'll show is just the term for k = 0. Now for k = 1, we would add to that one sinusoidal term. And so the sum of the term for k = 1 and k = 0 is represented here.

Now when we go to k = 2, because of the fact that this is an odd harmonic function, in fact, the term for k = 2 will have zero amplitude and so this won't change. Here we show the Fourier series with k = 2 and there's no change. And then we will go to k = 3. And we will be adding, then, one additional sinusoidal term. Here is k = 3.

When we go to k = 4, again, there won't be any change. But there will be another term that's added at k = 5 here. Then k = 6, again, because it's odd harmonic, no change. And finally k = 7 is shown here.

And we can begin to see that this starts to look somewhat like the square wave. But now to really emphasize how this builds up, let's more rapidly add many more terms, and in fact increase the number of terms up to about 100, recognizing that the shape will only change on the inclusion of the odd-numbered terms, not the even-numbered terms, because it's an odd harmonic function.

So now we're increasing and we're building up toward k = 100, 100 terms. And notice that it is the higher-order terms that tend to build up the discontinuity corresponding to the notion that the discontinuity, or sharp edges in a signal, in fact, are represented through the higher frequencies in the Fourier series. And here we have a not-too-unreasonable approximation to the original square wave. There is the artifact of the ripples at the discontinuity. And in fact, that rippling behavior at the discontinuity is referred to the *Gibbs phenomenon.* And it's an inherent part of the Fourier series representation at discontinuities.

Now to emphasize this, let's decrease the number of terms back down. And we will carry this down to k = 1, again to emphasize how the sinusoids are building up the square wave. Here we are back at k = 1. And then finally, we will add back in the sinusoids that we took out. And let's build this back up to 100 terms, showing the approximation that we generated with 100 terms to the square wave.

Okay, so what you saw is that, in fact, we got awfully close to a square wave. And the other thing that was kind of interesting about it as it went along was the fact that, with the low frequencies, what we were tending to build was the general behavior. And as the higher frequencies came in, that tended contribute to the discontinuity. And in fact, something that will stand out more and more as we go through our discussion of Fourier series and Fourier transforms, is that general statement, that it's the low-frequency terms that represent the broad time behavior, and it's the high-frequency terms that are used to build up the sharp transitions in the time domain.

Now we need to get a little more precise about the question of how, in fact, the Fourier series, or when the Fourier series represents the functions that we're talking about and in what sense they represent them. And so if we look again at the synthesis equation, what we really want to ask is, if we add up enough of these terms, in what sense does this sum represent this time function?

Well, let's again use the notion of our partial sum. So we have the partial sum down here. And we can think of the difference between this partial sum and the original time function as the error. And I've defined the error here. And what we would like to know is does this error decrease as we add more and more terms? And in fact, in what sense, if the error does decrease, in what sense does it decrease?

Now in detail this is a fairly complicated and elaborate topic. I don't mean to make that sound frightening. It's mainly a statement that I don't want to explore in a lot of detail. But it relates to what it's referred to as the issue of convergence of the Fourier series. And the convergence of the Fourier series, the bottom line on it, the kind of end statement, can be made in several ways.

One statement related to the convergence of the Fourier series is the following. If I have a time function, which is what is referred to as *square integrable,* namely its integral, over a period, is finite. Then what you can show, kind of amazingly, is that the energy in that error, in other words, the energy and the difference between the original function and the partial sum, the energy in that, goes to 0 as n goes to infinity.

A somewhat tighter condition is a condition referred to as it Dirichlet conditions, which says that if the time function is absolutely integrable, not square integrable, but absolutely integrable-- and I've kind of hedged the issue by just simply referring to x(t) as being well behaved-- then the statement is that the error in fact goes to 0 as n increases, except at the discontinuities. And what well behaved means in that statement is that, as discussed in the book, there are a finite number of maxima and minima in any period and a finite number of finite discontinuities, which is, essentially, always the case.

So under square integrability what we have is the statement not that the partial sum goes to the right value at every point, but that the energy in the error goes to 0. Under the Dirichlet conditions, it says that, in fact, the signal goes to the right value at every time instant except at the discontinuities.

So going back to the square wave, the square wave satisfies either one of those conditions. And so what the consequence is is that, with the square wave, if we looked at the error, then in fact what we would find is that the energy in the error would go to zero as we add more and more terms in the partial sum. And in fact, since the square wave also satisfies the Dirichlet conditions, the actual value of the error, the difference between the partial sum and the true value, will actually go to 0. That difference will go to 0 except at the discontinuities. And that, in fact, is kind of evident as we watch the function build up by adding up these terms.

And so in fact, let's go back and see again the development of the partial sums in relation to the original time function. Let's observe, this time again, basically what we saw before, which is that it builds up to the right answer. And furthermore what we'll plot this time, also as a function of time, is the energy in the error. And what we'll see is that the energy in the error will be tending towards 0 as the number of terms increases.

So once again, we have the square wave. And we want to again show the buildup of the Fourier series, this time showing also how the energy in the error decreases as we add more and more terms. Well, once again, we'll begin with $k = 0$, corresponding to the constant term. And what's shown on the bottom trace is the energy in the error between those two. And we'll then add the term $k = 1$ to the DC term and we'll see that the energy will decrease when we do that.

Here we have then the sum of $k = 0$ and $k = 1$. Now with $k = 2$, the energy won't decrease any further because it's an odd harmonic function. That's what we've just added in. When we add in the term for $k = 3$, again, we'll see the energy in the error decrease as reflected in the bottom curve.

So there we are at $k = 3$. When we go to $k = 4$, there again is no change in the error. At k = 5, again the error decreases. k = 6, there will be no change again. And at k = 7, the energy decreases.

And now let's show how the error decreases by building up the number of terms much more rapidly. Already the error has gotten somewhat small on the scale in which we're showing it, so let's expand out the error scale, the vertical axis displaying the energy in the error, so that we could watch how the energy decreases as we add more and more terms.

So here we have the vertical scale expanded. And now what we'll do is increase the number of terms in the Fourier series and watch the energy in the error decreasing, always decreasing, of course, on the inclusion of the odd-numbered terms and not on the inclusion of the even-numbered terms because of the fact that it's an odd harmonic function.

Now the energy in the error asymptotically will approach 0, although point by point, the Fourier series will never be equal to the square wave. It will, at every instant of time, except at the discontinuities, where there will always be some ripple corresponding to what's referred to as the Gibbs phenomenon.

So what we've seen, then, is a quick look at the Fourier series representation of periodic signals. We more broadly want to have a more general representation of signals in terms of complex exponentials. And so our next step will be to move toward a representation of nonperiodic or aperiodic signals.

Now the details of this, I leave for the next lecture. The only thought that I want to introduce at this point is the basic strategy which is somewhat amazing and kind of interesting to reflect on in the interim. The basic strategy with an aperiodic signal is to think of representing this aperiodic signal as a linear combination of complex exponentials by the simple trick of periodically replicating this signal, generating a periodic signal using a Fourier series representation for that periodic signal, and then simply letting the period go to infinity. As the period goes to infinity, that periodic signal becomes the original aperiodic one that we had before. And the Fourier series representation then becomes what we'll refer to the Fourier transform.

So that's just a quick look at the basic idea and approach that we'll take. In the next lecture, we'll develop this a little more carefully and more fully, moving from the Fourier series, which we've used for periodic signals, to develop the Fourier transform, which will then be representation for aperiodic signals. Thank you.