[MUSIC PLAYING]

PROFESSOR: In the last lecture, we considered the frequency response of linear shift and variance systems. In this lecture, what I would like to do is extend some of those ideas to the notion of the Fourier transform, which I'll refer to as the discrete time Fourier transform.

So let me begin by reminding you of one of the key results from last time, namely, the notion of the frequency response of a linear shift invariant system. For a linear shift invariant system, we saw last time that one of the key properties was that for a complex exponential input, the output is a complex exponential sequence with the same complex frequency, but with a change in amplitude. And the change in amplitude, that is the function h of e to the j omega, we refer to as the frequency response of the system.

The expression that we had for the frequency response, in terms of the unit sample response, was that the frequency response h of e to the j omega, is given by the sum of h of n, the unit sample response of the system, multiplied by e to the minus j omega n. So this is a relationship that tells us, in terms of the unit sample response of the system, how to get the frequency response.

There were two properties of the frequency response that I stressed last time. One of them was the fact that the frequency response is a function of a continuous variable, omega. That is, complex exponential inputs can have the frequency variable omega, vary continuously. Omega is a continuous variable, whereas n is the discrete variable. So this was one important property.

The second important property is that the frequency response is a periodic function of omega. And the period is equal 2 pi. H of e to the j omega is equal to h of e to the j omega plus 2 pi k, where k is any integer.

Now one of the sort of heuristic explanations or justifications for why the frequency response is periodic, as I've tried to indicate in several lectures, is basically tied to the fact that complex exponentials or sinusoidal inputs, are very periodically with frequency over an interval 0 to 2 pi. That is, once we've looked at them for omega between 0 and 2 pi, if we go further than that in frequency, all we see are the same complex exponentials or the same sinusoids over again.

So that, in essence, is the reason why the frequency response is a periodic function of omega.

Well, one of the things we'd like to do is develop an inverse relationship, which permits us to get the unit sample response from the frequency response. One heuristic notion or style of developing such an inverse relationship, we can get by observing that since the frequency response is a function of a continuous variable omega, and in fact, it's a periodic function of omega, then it should have a Fourier series representation. That is, it's a continuous periodic function. A continuous periodic function has a Fourier series representation.

Well in fact, if we look at this expression, we see that indeed what this is is a Fourier series expansion of this periodic function in terms of a complex Fourier series. The Fourier series expressed in terms of complex exponentials.

Well, what are the Fourier series coefficients? The Fourier series coefficients are the values of the unit sample response, in other words, the values h of n. So in fact, this has the form of a Fourier series. That means that these are the Fourier coefficients. That means that we can get these Fourier coefficients in terms of h of e to the j omega, the way we always get Fourier series coefficients in terms of a periodic continuous function.

Well, the resulting expression then is that the Fourier series coefficients, or equivalently the unit sample response, is equal to 1 over 2 pi times the integral over a period, which I've taken as minus pi to plus pi, of h of e to the j omega, e to the omega n, d omega. This is just simply the expression for Fourier series coefficients in terms of the continuous periodic function that we're dealing with. Well that's, in a sense, a heuristic derivation.

We can in fact verify that this expression is valid simply by substituting in to the integral, the expression for h of e to the j omega. And if we do that, we end up with, for the expression for h of e to the j omega, we have the sum of h of k, e to the minus j omega k. Substituting in, the result is this integral. If we interchange the order of summation and integration, then the resulting expression is shown here. And we can observe that this integral--

First of all, let's consider it for n not equal to k. If n is not equal to k, this is just simply the integral over a period of 2 pi of a complex exponential whose period is 2 pi. The integral of the real part is the integral of a cosine over an integral number of periods. And the integral of the imaginary part is the integral of the sine over an integral number of periods. Obviously then, this integral, if n is not equal to k is going to be 0.

So for n not equal to k, the value of this integral is 0, whereas for n equal to k, this exponent is 0, the exponential is unity, unity integrated from minus pi to pi is equal to 2 pi. So for n equal to k, the value of the integral is equal to 2 pi. Well that means that the only term in this sum that is going to contribute to the answer is the term for k equal to n, because for all the others the integral is 0. So for k equal to n, that's the only non-zero term, for k equal to n, the value of this integral is 2 pi, which will cancel out this 1 over 2 pi and consequently we get h of n.

All right, so this is, in a sense, a formal justification for the fact that, indeed this is the inverse relationship to express h of n terms of the frequency response. Although, in fact it's more useful in terms of insight to think of it as the Fourier series coefficients of a function, a periodic function, of the continuous variable omega.

Now this in essence, is a Fourier domain representation for the unit sample response of a system. And in fact, we can think of this as a Fourier transform or a time domain, Fourier domain transform pair, relating the unit sample response of a system and the frequency response of the system. And we can go back and forth.

Now one of the obvious facts is that any sequence, if we wanted to, we could think of as the unit sample response of a system. Well, that suggests then that this transform pair isn't restricted to just sequences that we explicitly identify as the unit sample response of a system. It in fact permits the representation of an arbitrary, not quite arbitrary as we'll see in a while, but the representation of a more general class of sequences than just the ones that we explicitly identify as the unit sample response of a system.

That is, we can generalize this set of ideas to the Fourier transform, which is a frequency domain representation of an arbitrary, again not quite arbitrary, but for the time being we'll consider arbitrary sequence, x of n. Well, the Fourier transform then is defined as x of e to the j omega, which is the sum of x of n, e to the minus j omega n, the Fourier transform x of e to the j omega, of a sequence, x of n, is essentially the frequency response of a system whose unit sample response would be x of n.

So then obviously from what we just finished discussing, we have an inverse relationship that tells us what x of n is in terms of x of e to the j omega. And that then is this relationship.

Well this then provides a transform pair, or a frequency domain, time domain relationship between sequences and frequency domain functions. And it's useful, in fact, to interpret this expression somewhat heuristically again as basically corresponding to a decomposition of a sequence, x of n, in terms of complex exponentials with incremental amplitude.

This is basically an expression that says that x of n is a sum, except that it's sum sort of in the limit which corresponds to an integral, of a set of complex exponentials with amplitudes that are essentially given by the Fourier transform x of e to the j omega. Well in fact, you can see this a little more explicitly if we consider this integral as the limiting form of a sum, the limit, as delta omega goes to 0, of the sum of x of e to the j k delta omega, times delta omega over 2 pi, e to the j k delta omega n. This limit is by definition, what this integral is.

So an important point to keep in mind is that basically the Fourier transform corresponds to a decomposition of a sequence in terms of a linear combination of complex exponentials with incremental amplitudes. There are a number of reasons why that's an important point of view. One of the reasons is that it leads to a very important property of linear shift invariance systems, which I refer to as the convolution property, and which states that if I have the convolution of two sequences, x of n and y of n, then the Fourier transform of those is the product of their Fourier transforms. And this of course, shouldn't be an h, it should be a y. Or this shouldn't be a y, it should be an h.

The important property is that the convolution of two sequences has, as its Fourier transform, the product of the Fourier transforms. Well let's look at again, somewhat heuristically, an argument that at least justifies this, keeping in mind that we could go through this more formally, plugging sums into integrals and integrals into sums, but let's not do that. Let's look at this from a somewhat heuristic point of view.

First of all, we have that again, a property that we began the lecture with, that for a linear shift invariant system, a complex exponential input gives us at the output, the same complex exponential, that is the same frequency, and only one of them, multiplied by h of e to the j omega 0. That's a consequence of linearity and shift invariance. We know also, that because of the fact that the system is linear, if we have a linear combination of complex exponentials, then the output of this system is going to be the same linear combination of complex exponentials, with the amplitudes multiplied by h of e to the j omega k. In other words, each of these exponentials has as an output ase of k times h of e to the j omegus of k, e to the j omegus of k n. A sum of those at the input gives us a sum at the output because of the fact that the system is linear.

All right, so let's take an input, which is an arbitrary input, more or less arbitrary, an input x of n, which I can express in terms of its Fourier transform, in terms of this inverse Fourier transform relationship. All right, now m as I emphasized just a minute ago, is a decomposition of x of n as a linear combination of complex exponentials. This is a linear combination of complex exponentials at the input, so what's the output? Well, the output is then going to be a linear combination of complex exponentials with the amplitude of each one of the input exponentials, multiplied by the frequency response of the system at that frequency. So the result is that with this as the input, what we have to get at the output is this, the same linear combination, the only change being that the complex amplitudes of the input are multiplied by h of e to the j omega.

Well, this of course, is the expression for the output. We put in x of n, we know that we're getting out y of n. So obviously this then must be the Fourier transform of the output of the system. So what this says then is that the Fourier transform of the output y of e to the j omega, has to be this, it has to be x of e to the j omega times h of e to the j omega. And basically-- well, and we know also that y of n is going to be equal to x of n convolved with h of n. So basically, this justifies the convolution property, that is the convolution of two sequences has as its Fourier transform, the product of the Fourier transform of each of the individual sequences. Very important property. And although we can go through a formal derivation of this, in fact the basic reason for it is tied to the arguments that I've outlined here. And in terms of insight, I feel that it's more important to understand this way of looking at it than to understand a formal derivation.

Well, this is interesting, also important, and also should be familiar to you in terms of things that you're used to thinking about for continuous time systems. Obviously, in continuous time systems the same type of property holds. That is, that the Fourier transform changes convolution to multiplication. And in fact, it permits the description of a linear shift invariant system to be in terms of multiplication rather than in terms of convolution. And as we'll see in a number of lectures, that basically the basis for the notions of filtering and some other very important notions, and notions of modulation etc.

OK, key property, this is a key property of linear shift invariant systems and the Fourier transform. There are, of course, lots of other properties of linear shift invariant systems-- sorry, other properties, there are other properties, obviously, of linear shift invariant systems. There are other properties of Fourier transforms that are important, both in terms of interpreting Fourier transforms and in terms of computing Fourier transforms of a variety of sequences. A lot of these properties will be developed in the text and also in the study guide, so you'll have to do the work rather than me. But let me just indicate one class of properties that, again, should be very familiar to you if you relate your thinking back to continuous time Fourier transforms. And that is the class of symmetry properties for the special case in which the sequences that we're talking about are real sequences.

Well the basic symmetry property is that for x of n real the Fourier transform is a conjugate symmetric function, x of e to the j omega is equal to x conjugate of e to the minus j omega. And we can see that, essentially in a straightforward way, that is the derivation is effectively straightforward. Here's x of e to the j omega, the sum of x of n, e to the minus j omega n. Here is x of e to the minus j omega. Well, the only difference between that and that is that we replace omega by minus omega, so this becomes a plus sign. Of course, it's not this that we want from this expression, it's the conjugate of that. So we want to complex conjugate this.

Well, we will do the same on the left--hand side-- on the right-hand side. And so we would conjugate that and conjugate this, which replaces this plus sign with a minus sign. But we're talking about a real sequence, x of n, so x conjugate of n is just x of n. In other words, this is just the x of an all over again. And so we have that x conjugate of e to the minus j omega is the sum of x of n, e to the minus j omega n, which is just what we have up here. So obviously then, these two are equal. So for a real sequence, the Fourier transform is a conjugate symmetric function. This is what we'll call conjugate symmetric.

Well let's press that a little further. We have x of e to the j omega, which I can represent in terms of its real part and its imaginary part as obviously x real, x sub r of e to the j omega, plus j times x sub i of e to the j omega. And the conjugate symmetric counterpart is this with omega replaced by minus omega and the expression conjugated. So we have x sub r of e to the minus j omega, minus j times x sub i of e to the minus j omega. And we know that these two are equal. Well, if these two are equal, then these two are equal, and so are these. The real part of x of e to the j omega must be equal to the real part of x of e to the minus j omega.

In other words, the real part has to be the same if omega is replaced by minus omega. That means then that the real part of the Fourier transform is an even function of omega. Meaning that if we replace omega by minus omega then the real part doesn't change. The imaginary part, on the other hand, does. In particular, on the basis of what we're saying here and the equality of these, the imaginary part x sub i of e to the j omega, must be equal to minus, don't forget the minus sign, minus x sub i of e to the minus j omega. And that says then, that if we replace omega by minus omega, then the sign of the imaginary part changes. So the imaginary part, in fact, is an odd function of omega. The real part is an even function of omega.

Well, from this, or from this, we can also show that the magnitude of the Fourier transform is an even function of omega. And the angle of the Fourier transform is an odd function of omega. And those, of course, are identical to what we know is true for the continuous time Fourier transform. Remember however, again, that these are periodic functions, whereas in the continuous time case, they're not.

All right well, this is an introduction to the Fourier transform. One of the things that we've refrained from doing in all of these lectures is tying our development too closely to the notion of continuous time signals, and in particular, avoiding to some extent, the notion of interpreting discrete time signals as just simply sampled replicas of continuous time signals. And we'll continue to do that throughout this set of lectures. But in particular Fourier transform, I think that it's instructive to tie together, at least in terms of some insight into the relationship, the continuous time Fourier transform of obviously continuous time signal, and the discrete time Fourier transform for a sequence that's obtained by periodic sampling. That is, equally spaced sampling of the continuous time signal.

So what I'd like to do now is focus on that relationship, emphasizing again, that not all sequences arise by periodic sampling of continuous time signals. But for the cases in which they do, the relationship between the continuous time and discrete time Fourier transform is instructive. So let's take a look at the relationship between some continuous time and discrete time Fourier transforms when we obtain the discrete time signal by sampling a continuous time signal.

Well, we begin of course, with a continuous time, time function which I denote by x sub a of t, a sort of meaning analog. And the steps that we would go through to convert that to a sequence are first of all, to go through a sampler, which I've indicated here. And the output of the sampler is then a sampled continuous time version of this signal. This essentially is the input signal multiplied by an impulse train.

So I've indicated that here, here's the continuous time input. Here is the continuous time output of the sampler, which is an impulse train with the envelope of the impulse train being the continuous time function x sub a of t. All right, well this isn't the sequence. This is just simply an impulse train. To turn this into a sequence we need to go into a box, which I've labeled c slash d, meaning continuous time to discrete time converter. And the output then is a sequence x of n. The sequence values being samples of x sub a of t at the sampling instances n times capital T. In other words, it's converting the areas of these impulses into sequence values.

Now I've illustrated it here for one choice for the sampling interval capital T. Let's, down here, illustrate it for a sampling interval that's twice as long. Well, of course we have the same continuous time input. The sampled output, x sub a, tilde of t, has the same envelope, but the spacing of the impulses is twice what it is here, and that I've indicated. The envelope, of course, is the same. But at the output of the continuous time to discrete time converter, what do we have?

Well, we have the areas of these impulses lined up along this axis, again, at integer values of n. That is, the spacing of the lines, when we look at the sequence here, must be exactly the same as the spacing of the lines here, it's just that the values are different because we picked out samples at different instance. So in fact, the envelope here is indeed a compressed version of the envelope that we had here. Very important point.

The point being that no matter what the sampling rate is the sequence values, when we line them up as a sequence, are going to fall at integer values along this argument, n, always at intervals that correspond to 1. Whereas the output of the sampler had impulses occurring at a spacing of capital T, which in this case, was capital T 1, and in this case is capital T 2.

All right well, this is essentially the sampling process plus the conversion to a discrete time signal. And now let's take a look at what this means in terms of the Fourier transform of the discrete time signal as compared with the Fourier transform of the continuous time signal. Let's do this in two steps.

First of all, we have the sampler. Here is the input x sub a of t, continuous time function. Here is the output, x sub a tilde of t, a continuous time function. And the output of the sampler is the input to the sampler, multiplied by an impulse train. Or equivalently, it's an impulse train with the areas of the impulses given by the values of x sub a of t at the times that the impulses occur. Naught, by the way, is what I'm using as the notation to designate a unit impulse, unit continuous time impulse.

Now in the Fourier transformed domain then, the continuous time Fourier transform, is the convolution of the continuous time Fourier transform of x sub a of t, convolved with the Fourier transform of the impulse train, which is an impulse train in the frequency domain. That's a result that you should be familiar with for continuous time signals. Or equivalently, it's given by 1 over t, times the sum of x sub a of j omega, plus j 2 pi r over capital T. Basically, what that means is that the Fourier transform of this continuous time signal is equal to the Fourier transform of this continuous time signal, but repeated over and over again in frequency at intervals of 2 pi over capital T. So this is sort of a standard sampling theorem kind of result in the continuous time case, and a result that you should be more or less familiar with. Let's look at this now from a different point of view. Again looking at x sub a tilde of j omega.

Well x sub a tilde of j omega is the integral of x sub a tilde of t, e to the minus j omega t d t. Substitute in for x sub a tilde of t, the relationship in terms of an impulse train, and interchange summation and integration, and I think you could verify in a very straightforward way that what you end up with is an expression for the Fourier transform at the output of the sampler as given by the sampling values times e to the minus j n capital omega T. Capital omega is a continuous frequency variable.

Now we want to look at the relationship between the Fourier transform, continuous time of x sub a of t and the discrete time Fourier transform of x of n. So we have the next step, which is to put this impulse train into the continuous to discrete time converter. And I remind you that we just developed two results. One result was that x sub a tilde of j omega was given by this expression. Also we developed that it was given by this expression.

Well what's the Fourier transform of the output of the continuous to discrete time converter? Well, it's just simply x of e to the j omega, our discrete time Fourier transform, which is equal to the sum on n of x of n. But x of n is x sub a of n capital T. So we have x sub a of n capital T, times e to the minus j omega n.

Well let's compare this with this. We see that they're exactly the same except that for omega, little omega here we have capital omega times capital T there. Consequently, we can say that the discrete time Fourier transform is equal to the Fourier transform of the impulse train, the continuous time Fourier transform of the impulse train, with omega times capital T equal to little omega.

And we had another expression for the Fourier transform of the impulse train, which we derived here. Consequently, the final result that we end up with is that the Fourier transform of the discrete time sequence is equal to 1 over t times the sum of x sub a of j omega, plus j 2 pi r over capital T, with omega replaced by little omega, divided by capital T. In other words, we had the expression that capital omega times capital T is equal to little omega. And this then tells us what the Fourier transform of the sequence is in terms of the Fourier transform of the output.

Well, this is just the equations. Let's take a look at what this looks like graphically. I've depicted here the continuous time Fourier transform of some time function, and I picked the Fourier transform that looks like a triangle. Well, first of all we derived the fact that the impulse train that results from sampling little x sub a of t has a Fourier transform, which is this, periodically repeated in frequency with a period in frequency equal to 2 pi divided by capital 1, where capital 1 is the sampling rate.

And on the basis of the expression that we derived relating the discrete time Fourier transform and the continuous time Fourier transform, the discrete time Fourier transform looks exactly like this, but with a re-normalization of the frequency axis, because capital omega times capital T is equal to little omega. Yes, capital omega times capital T is equal to little omega. So where capital omega is equal to 2 pi over capital T, little omega is equal to 2 pi. So this point, which was at pi over capital T ends up on the omega, little omega axis at pi. So this picture just simply gets scaled according to the relationship that capital omega times capital T is equal to little omega.

Well, this is, for one choice of the sampling period. Obviously, if capital T was large enough so that 2 pi over t 1 was small enough, then each of these individual replicas of the frequency response would interact. And we wouldn't have just the simple separation of the spectra as we have here. We'd have an interaction, which of course, is the interaction and the relationship between when that interaction occurs and capital T, is the basis for the well-known, and hopefully, theorem that you're familiar with namely, the sampling theorem.

Well, to illustrate that if we had a different sampling rate, say twice the sampling rate that we have over here, so that capital T is equal to 2 times capital T 1, then starting with the same continuous time spectrum, what we have now is the spectra, again periodically repeated, with a period again, which is 2 pi over capital T, which in this case is 2 pi over capital T 2, but now they interact and of course, we have to add these up. That was the expression that we just derived.

So in that case, we get some interaction or aliasing, as it's referred to. And due to this aliasing, the periodic, one period of this periodic spectrum no longer resembles the original spectrum. This is the Fourier transform for the sampled continuous time function. That is, this is the Fourier transform for the impulse train. And now if we renormalizing the frequency axis so that we express this in terms of the discrete time frequency variable, little omega, then we simply scale this picture so that we end up with 2 pi over capital T, corresponding to 2 pi in little omega. Pi over capital T corresponding to pi.

And we see in fact, as-- we better see, that is, it better turn out this way, that the spectrum in the discrete time case is a periodic function of omega. In other words, it's periodic. Furthermore, the period is given by 2 pi, whereas here the period was 2 pi divided by capital T.

All right, this perhaps takes a little digesting. And you'll have some chance to do that as we work some problems. We, meaning you, as you work some problems in the study guide and digest this a little while you're reading the text. But it's an important relationship and it's important to understand it.

Now, this is basically the Fourier transform, the relationship between the discrete time and continuous time Fourier transform. One of the difficulties with the Fourier transform, which I've avoided illustrating explicitly in this lecture, is that the Fourier transform doesn't exist for all sequences. In particular, it doesn't converge for all sequences. And this is a problem which we can get around by generalizing the notion of the Fourier transform to what we'll call the z transform. And the z transform, as you'll observe, is like, in the continuous time case, the Laplace transform. And this is what we'll go on to in the next lecture.

Thanks

[MUSIC PLAYING]