

MITOCW | 14. Spectral representation of signals

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high-quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

PROFESSOR 1: OK, let's launch right into it. Jacob and Uri were inspired by the echo channel to try out a simulation of what I'd put on the board last time or on my slides. So what you're going to hear is Jacob's message going into the echo channel.

Remember, that was something with a unit sample response of the type $\delta[n]$ plus. And actually, I think in their case, now it's $0.999\delta[n - 1]$, so it is something like this. Sorry? Oh, $n - 4,000$, OK.

And so you'll hear the original message, I think. You'll hear the message going through the echo channel. And then you'll hear the message cleaned up with the receiver filter. That's just the inverse filter. And then you'll hear what noise does to it and two flavors of noise, I think.

ECHO VOICE: This is a test of the 602 deconvolution system. If this is a real deconvolution, you are instructed to go as quickly as possible to the frequency domain to make an assessment of the effectiveness of this approach. This is a test of the 602 deconvolution system. If this is a real deconvolution, you are instructed to go as quickly as possible to the frequency domain to make an assessment of the effectiveness of this approach.

[BELL DINGING]

PROFESSOR 2: I'm just going to increase the delay, so we can hear the echo more clearly.

ECHO VOICE: This is a test of the 602 deconvolution system. If this is a real deconvolution, you are instructed to go as quickly as possible to the frequency domain to make an assessment of the effectiveness of this approach.

PROFESSOR 2: So you can all hear the echo in that. And the next one will be it cleaned up--

ECHO VOICE: This is a test of the 602 deconvolution system. If this is a real deconvolution, you are instructed to go as quickly as possible to the frequency domain to make an assessment of the effectiveness of this approach.

PROFESSOR 2: So this was just one in deconvolution, assuming the channel had no noise. The next one is what happens if there's even a small amount of noise in the channel-- so little that you couldn't hear it in the echoed signal.

ECHO VOICE: This is a test of the 602 deconvolution system. If this is a real deconvolution, you are instructed to go as quickly as possible to the frequency domain to make an assessment of the effectiveness of this approach.

PROFESSOR 2: So you can hear the noise building up because of the deconvolver. If you actually had noise at a particular frequency, you end up with an even--

ECHO VOICE: This is a test of the 602 deconvolution system. If this is a real deconvolution, you are instructed to go as quickly as possible to the frequency domain to make an assessment of the effectiveness of this approach.

PROFESSOR 2: So one of the difficulties with deconvolution is it can be perfect if you have no noise. But even small amounts of noise can really mess up the deconvolution. It magnifies the small noise sources.

PROFESSOR 1: OK, great. Thank you. Thank you. And that was a few lines of MATLAB, right? Same thing can be done in Python. All right, we continue. So we're going to talk today about spectral content of signals after having spent some time on the frequency domain-- the frequency response of LTI systems. Let's get this up here.

So we've talked about frequency response. You've seen the definition. If I give you the unit output response of a system over here, you know how to compute the frequency response. And then we went through how you can go the other way with the DTFT to compute the frequency response and the inverse DTFT to compute the time domain signal from the frequency response.

And what we said last time was that what you're doing with a unit sample response, you could actually do with any signal. So you can take any signal $x[n]$, compute from it the DTFT, the Discrete Time Fourier Transform, with the same formula. What you get is this object that can be used then to reconstruct the signal-- again, the same formula, all right? So there's just a change of perspective. There's nothing different here.

The key observation now, though, is that the formula on the left, the inverse DTFT, is actually allowing us to represent $x[n]$, the time domain signal, as a weighted combination of exponentials of this type. And the reason that's important is that we know how to deal with signals of that type very easily. We already know that if you have $e^{j\omega n}$ going into a system with frequency response-- $h(\omega)$, so an LTI system with that frequency response.

Let me make this a specific frequency ω_0 . What comes out is the frequency response evaluated at ω_0 , multiplying what went in. All right? Nothing more complicated than that.

So what now if you had an $x[n]$ going in that was a weighted combination of terms of this type? And I'm going to take a weighted combination that's actually a continuum. It's not just a finite number of terms. I'm going to actually take this particular weighted combination.

So over some interval of length 2π , I take a weighted combination like this. So this is going over all frequencies in our interval. Let's take minus π to π . I might as well write in minus π to π -- keep it explicit. You can think of this as being approximated by a sum in the usual fashion. I'm not going to write this out. But we know how to approximate integrals by sums.

What the sum will have is, for instance, a typical term of the type $e^{j\omega_0 n}$. And the weight that multiplies it will be $x(\omega_0) d\omega$, right? So we think of $x(\omega_0) d\omega$ as being the amount of the exponential at frequency ω_0 that's in $x[n]$. So here is a representation of how the signal is made up.

So then what would you say is the output of the system? If that's the input, and this is an LTI system, what's the output going to be? Any ideas? Somebody? I thought I saw a hand. No ideas?

What if instead of this, I had $a_1 e^{j\omega_1 n}$ plus $a_2 e^{j\omega_2 n}$ going in? Suppose that had gone in? What would be coming out? Folks, we're just two weeks from a quiz, here. Yeah?

STUDENT: [INAUDIBLE]

PROFESSOR 1: So can you tell me explicitly what I would get in this case? If this was $x[n]$, then $y[n]$ would be?

STUDENT: [INAUDIBLE]

PROFESSOR 1: $a_1 h$.

STUDENT: [INAUDIBLE]

PROFESSOR 1: Is it just ω , or?

STUDENT: It's ω_1 .

PROFESSOR 1: ω_1 , right? It's the frequency response evaluated at the frequency that you're interested in, and then $e^{j\omega_1 n}$, and then the response to the other term. This is what superposition is about. So that wasn't so hard.

What if instead, $x[n]$ is given by a continuum of such exponentials? Integrals are essentially linear combinations, but taken to the limit, where it's not just a finite number. So I'm not asking for a proof of anything. I'm asking for your conjecture as to what the answer might be.

This is how math is done, by the way. You conjecture what the result might be based on gut instinct, based on well-educated intuition. And then you go back and construct a proof and hide all your tracks. But engineers like to work with intuition, and often will stick with that. Yeah?

STUDENT: [INAUDIBLE]

PROFESSOR 1: OK, so what would-- can you give me the explicit expression? What's your guess? This is going in. It's a weighted combination of exponentials with weights that are given by $x(\omega) d\omega$. So--

STUDENT: So you would get $x(\omega) d\omega$ [INAUDIBLE].

PROFESSOR 1: Times what? I didn't hear the last piece.

STUDENT: Times [INAUDIBLE].

PROFESSOR 1: Where's the $e^{j\omega}$ coming? What's the j ? I'm missing-- oh, $e^{j\omega}$. All right. So start again. Oh, you don't want to?

STUDENT: No.

PROFESSOR 1: OK. You're on the right track. You got us started. Anybody else? I could show it to you on the next slide, but that will take all the fun out of it, right?

STUDENT: [INAUDIBLE]

PROFESSOR 1: OK, it's probably that I didn't hear. Can you tell it to me?

STUDENT: $x(\omega) d\omega$ [INAUDIBLE].

PROFESSOR 1: $x(\omega)$ times $h(\omega)$ or $e^{j\omega n}$? What did you say?

STUDENT: [INAUDIBLE]

PROFESSOR 1: I'm willing to--

STUDENT: [INAUDIBLE] omega n times [INAUDIBLE].

PROFESSOR 1: That's still the part that went in, right? Now it's got to get mapped by a frequency response.

STUDENT: Yeah, that's what I said.

PROFESSOR 1: OK, so maybe you'd said that, and I didn't hear it. But now we've got to assemble this over all possible frequencies. Well, we have the $1/(2\pi)$. Maybe this was said already, and I just couldn't hear it. All we're doing is we're saying here's a weighted combination of exponentials going in.

The weights are given by the x omegas or x omega d omega, if you want to think of it that way. So what comes out is the combination of responses to each of those. These are the things that go in. For each frequency, you multiply by the corresponding value of the frequency response and do that for all the frequencies of the input. So this is just applying linearity and superposition.

So if I compare that with this expression, which just tells me how the time domain signal y_n relates to its DTFT, right? This is just writing the same thing for y that I wrote for x over there. If you compare the two, what you discover about the DTFT of the output? Where's the weighted combination? It's whatever multiplies e to the j omega n in this expression, right?

So it's just going to be h omega x omega. So we've done a complete analysis of the input-output response of the system for essentially an arbitrary input with just a simple multiplication. So look what we've done. We've taken the input that we were given, computed the DTFT of it, which gives us the spectral content, and I'll spend some time giving you intuition for that.

That's the spectral content of the input signal. What we've discovered is that the spectral content of the output is the spectral content of the input scale by the frequency response. And once you have the spectral content of the output, you can reconstruct the time domain signal.

So the big difference here is there's no convolution. You're just doing a multiplication. Instead of doing y of n equals h convolved with x_n , we're just doing a multiplication. So once again, you see the convolution of the time domain maps to multiplication in the frequency domain.

All right, so we'll build up to the story again. Let's get some intuition for spectral content. And let's take a particular example. So suppose I have an x of n that's a one-sided exponential. You've probably done things like this in recitation already.

So this is a signal that starts at time 0 and then starts at the value 1 and halves it each time. So it's a discrete time exponential. And so what's the DTFT? Well, you're going to sum from m equals minus infinity to infinity. But actually, this only exists for positive-- for non-negative time. So you're going to get 0.5 to the n . Or let's keep it at m .

That's just the definition. Isn't it the definition? I'll just use the definition. So you can now sum an infinite series here. And what do you get? You get $1/(1 - 0.5e^{-j\omega})$. This is just summing a geometric series because each term here follows from the previous one by multiplying by the factor $0.5e^{-j\omega}$.

So it's a geometric series with that ratio. And so this is what the sum works out to be. So if you wanted to figure out the spectral content, you first compute the DTFT. And then the most helpful way to get a feel for the signal is to look at the magnitude of the DTFT. And that's what's actually plotted in this case.

This is taken from somewhere that used slightly different notation, but we'll talk through it. What happened to the top of my slide, here? OK, it doesn't matter. What I've plotted here is the magnitude of x and the phase of x . To get that, you'll actually have to convert this to magnitude in angle form. I'm not doing that for you. I'm assuming you've had practice or will get practice in recitation.

But the result of that is a magnitude that looks like this and a phase that looks like this. The horizontal scale, just to make the point that this is a periodic object, just like frequency response, it actually goes from minus 4π to plus 4π . But the interval of interest is really just minus π to plus π . All right?

So it's just that central portion that's of interest-- similarly here, minus π to plus π . And then it replicates periodically outside of that. So we didn't really need to show it to you outside of that. This is just to make the point.

Another thing to observe is that the magnitude is an even function of frequency, and the phase is an odd function of frequency. So these are elementary checks that you should make. If you get an answer that doesn't satisfy those properties, you've gone wrong somewhere along the way.

If you look at the top plot in this set, the top plot is exactly a signal of that type, except I've chosen a different number. Instead of 0.5 to the n times u of n , it's something else. But here's a geometric-- sorry, a discrete time exponential or geometric series. Now I've just plotted the DTFT from minus π to π to show you what the spectral content of the signal is. And I'm just plotting the magnitude.

Ignore these labels. These are the same figures you saw earlier for little h and big H . I'm using them again, except now I'm thinking of this as a signal, and this is its DTFT. This is a signal, x sub n , and this is its spectral content. The relationship is exactly what we had with frequency response.

So where is the spectral content concentrated? Is it at low frequencies or high frequencies or intermediate frequencies for this first example? Concentrated around 0 , so you'd say it's concentrated at low frequencies. There is content on all frequencies, though, so this doesn't dip down to 0 anywhere else. You have to assemble a combination of sinusoids at all frequencies to construct this signal here.

Here is a signal. I've had to change the horizontal scale because this is a signal that evolves much more slowly. Again, I can ask what's the spectral content of it? I get something which has a peak near 0 frequency. So it's only got low frequencies in it and has very little high frequency content.

You can also start to develop ideas for how fast-- how high a frequency you ought to expect to see here. So for instance, what's the fastest wiggle that you see in this signal? About how long does it take to-- if you were thinking of underlying sinusoids, what's the fastest wiggling you're seeing over here? Well, to my eye, this kind of rises and curves within about 18 or 20 samples, right?

So this might be a half period of an underlying sinusoid. So if I thought of the period as being-- these are just rough calculations. But it helps you understand what we mean by spectral content and helps you as a way of checking answers. But let's see. If I said 18 was approximately a half period of an underlying sinusoid, and I don't see anything faster than that, period is 2π over ω_0 . So I'm saying that's approximately 18.

So the frequency that I expect to see, the fastest frequency there, is 2π over 18 or π over 9. Is that roughly consistent with what we're seeing there? Does the spectral content drop off? Well, here is π over 4. Here's π over 8. Somewhere on π over 9, we've run out of underlying components. It's because the frequency content is limited to that range that the associated signal doesn't wiggle any faster than this.

Here's another example-- a signal that actually-- well, in this case, it seems to have some fairly regular periodicity to it, and then it damps out. By the way, in all these cases, I'm assuming-- actually, in all these cases, I'm assuming the signal's identically 0 outside-- outside of what I've shown you here. If you take the DTFT of this, you find that the spectral content is what-- low frequency, mid frequency, high frequency?

STUDENT: Mid?

PROFESSOR 1: Mid frequency right? Because this is low frequency. Here is 0. This is high frequency. This is just a reflection on the left side. So at some intermediate frequency, there's a peak in the spectral content. And again, you can go through the rough calculation I just made.

So we see some oscillation here. Let's see. Let's estimate the period. That's 1, 2, 3, 4, 5. Let's say it's about a period of 5 for those oscillations. So 2π over ω_0 is approximately 5. So ω_0 is approximately 2 over 5π . So we expect to see a spectral peak somewhere around 0.4π . Here's 0.25π . Here's 0.5π . We're about right. So make these sorts of checks. Yeah, question?

STUDENT: So for that first calculation with the [INAUDIBLE].

PROFESSOR 1: Oh, sorry. Yeah, I should have done twice this, right? What I did was estimate-- thanks for catching that. I estimated at about 18 is the half period. And so the period I should have had-- 36 here. Good. This is all ballpark, of course, but no reason to make it worse than it has to be. Vibrating, right? And that's not exactly where it sits down there, but it's in the right region.

Now here's the part that we've already seen on the board. Once you know the spectral content of the input, you can assemble-- or you can think of the time domain signal as being made up of those components. Correspondingly, that's what the output is.

And all we're doing is invoking the fact that this is an LTI system for which we know the frequency response gives us the output for exponential inputs. And then we compare that with what we expect to be seeing for the DTFT of y , and we make this conclusion. So this is exactly what we had earlier.

One thing to keep in mind, the DTFT, the frequency response, the DTFT-- these are all complex functions of ω in general. Each of them will have a magnitude and an angle. So make sure you understand why the magnitude of y is the product of the magnitudes of h and x and why the angle of y is the sum of the individual angles, all right? It's basically the fact that for a complex number c , you can write it as the magnitude times e to the j angle. So that's really what's being invoked there.

So really, what the story is about-- we've done a lot of math along the way. But this is really the story. And I've only exposed a little part of it for you because I've only dealt with DT signals. But the same thing holds for continuous time signals. A huge class of such signals can be written as linear combinations of sinusoids.

And when I say "linear combination," it could be a combination of a discrete set, finite or infinite. Or it could be a continuous combination of exponentials or sinusoids under an integral sign, but the idea is the same. If you've done 1803, you've seen this kind of thing happening, at least for periodic signals. And then the other piece of what we rely on is that LTI systems are very easy to understand in terms of their action on sinusoids. So once you put these two pieces together, you've got a very powerful way to analyze LTI systems.

So just to go back to the kind of example I had last time, in, which you'll be-- or you're already dealing with in the lab-- we're talking about an audio channel, for instance. The frequency response, in this case, is the magnitude. Some characteristic here-- this is a bit of a cartoon. But let me show you more typical experimental plots of frequency response.

This is frequency response magnitude. In most of these plots, people don't show you the phase. Part of the reason is, or maybe the major reason is, that for audio, the ear is not all that sensitive to phase. If we were doing the analogous thing for video, then you'd be very concerned about phase. But in audio characteristics, people will typically only show you the magnitude because the phase distortions aren't picked up quite that readily.

So here is three speakers. If you look on the site there, you'll see many more tested. This is the frequency range. Now, I should make some comments about that. We're talking about doing minus π to π . We've been talking about filters with frequency responses that we show from minus π to π .

So for instance, if I had a band pass filter, it would be something with-- in the ideal case, something like this, right? This is because we're writing things in terms of big ω for a discrete time filter. These are actually written-- the scale here is hertz.

And they're talking about the action on an underlying continuous time signal. So you actually need a way to go from an underlying continuous time signal that sampled at a particular sampling rate-- let's say f_s samples per second-- to a corresponding ω for the underlying discrete time sequence.

So the question is, how does f_s map to ω ? And I had a slide last time. I haven't gone through it in detail. Maybe we'll have you work through it on a homework problem. But this is actually the mapping.

If you have a sequence that comes from an underlying continuous time sinusoid by sampling at f_s samples per second, and you're doing all your calculations in the discrete time domain, if you want to think about what that means for the underlying continuous time domain, you want to map π to f_s over 2. It's not the ω that maps to that. It's the π .

So for instance, in the lab, I think you're using 48 kilohertz, for instance, at least for some part of it, as a sampling rate. You get a discrete time sequence out of that. You do various DTFT-type computations-- spectral content or frequency response. Then you plot them on this kind of a scale.

If you want to think about what the underlying continuous line frequency is, well, that's 24 kilohertz in this case, and minus 24 kilohertz, it's at this end. So when you're trying to visualize what this characteristic is telling you about what you're seeing with a discrete time sequence, that's really the mapping.

The other thing about these characteristics is that people only plot the positive frequency part. So they ignore the negative frequency because of the symmetries there. In applications, when they give you a frequency response, they will typically just give you the positive frequency part of that. All right, so this is what we're seeing here.

This is the characteristic of the LTI system that you're going to be sending signals into. And then you've got to characterize the signals that you're going to send through it-- voice or music or whatever. And this is a figure I showed you last time. But basically, you're looking at the spectral content of the signal of interest and seeing how it matches up with the channel that you have.

And if you compare with-- well, let's actually look at the previous case and get a few landmarks, here. So let's take the Sony speaker, for instance, down here. OK, so it's got a fairly flat frequency response for a range of frequencies. But you've got to get fairly high up before you get there.

For frequencies lower than about 100 hertz, this is not doing a very good job of propagating the sound. The frequency response is measured by having a microphone at a fixed distance from the speaker in an anechoic chamber. And you can see that this one is actually perhaps the poorest of the-- it is the poorest of the speakers in that it does it very poorly with the low-frequency sound.

So for this particular one, you would hope that the spectral content of what you're trying to send through the channel lives somewhere in maybe 300 to-- 300 hertz to 10 kilohertz-- if you want to get it across the channel-- from the speaker with high fidelity. But if you've got low-frequency signal that you're trying to send, and you use the speaker, well, you're going to be out of luck. It's not going to propagate it very well. So thinking in terms of frequency response and spectral content is really key to making sense of a lot of this.

All right, let's get a little more practice with this. And I just want to show you that once you've learned how to deal with frequency response, there's not new stuff that you're going to do to deal with spectral content. It's just a change of perspective.

So let's see. If I asked you for a signal that had its frequency content uniformly distributed in some finite range, can someone tell me what that signal is going to look like? I'm asking you for a signal whose spectral content is uniformly in some range minus ω_c to plus ω_c . Have you met such a signal before? Anyone?

STUDENT: [INAUDIBLE]

PROFESSOR 1: Sorry?

STUDENT: [INAUDIBLE]

PROFESSOR 1: It was the unit-- we've seen the same kind of thing with the unit sample response to the low pass filter, right? So if this was a frequency response, then the associated unit sample response would be the signal we're talking about. So remember what that was called? We called it a sinc function-- sinc function in time.

So if the spectral content is this in frequency, then the signal that you're talking about is going to be a sinc function. Now you can actually work that out. You don't have to take my word for it. You want flat spectral content in the range minus ω_c to plus ω_c . So the signal that you're going to get as a result you can extract from this computation. And this is exactly the same function that we saw last time.

So the DTFT does what your eye may not do very well. If I had just given you the signal and asked you to take a guess as to what the spectral content of that is, you're not very likely to have ended up deducing that the spectral content is flat in some range and 0 outside of that. So the DTFT is valuable in actually doing this analysis for you.

So there is a signal that has flat spectral content. More examples of a similar type-- and again, we first encountered these in the context of unit sample responses and frequency responses. But now, I just want to change perspective and think in terms of time domain signals and their associated DTFTs.

So if we look at the top one there, this is the case we just saw, except I've truncated the sinc function. And so what I get is not the perfectly uniform distribution of frequencies in some interval. There's a little bit of a wiggle to it. But this is essentially the sinc function and its spectral content.

Here's another signal whose spectral content is at high frequencies and essentially 0 in the low-frequency range. What does it look like in the time domain? Well, you can actually work it out. And here's what you see-- that this has actually more wiggle to it than the sinc does. Alternate samples seem to take opposite signs, at least of the dominant ones over here, reflecting the high-frequency content of that signal.

Here's something that's intermediate. This also has the oscillation in sign, but it's not necessarily in alternate samples. It's a little bit more leisurely. Here's something that has low frequency and high frequency, but not intermediate frequency. So you see a component that's rapid wiggling, but you also see this lower-frequency content in there. So this is what the DTFT does for us.

Now there's an issue of how you compute these, because if you look at the formula for the DTFT, you could certainly do analytical things with that expression. And that's the case that-- we've treated cases of that type, where you write down an analytical formula for the DTFT. And then you do things with that, like plotting. But if I gave you some numerical sequence here, there's certain simplifications.

For one thing, you really aren't going to expect to compute this at a continuum of values of ω . You're not really going to expect to construct the values from $-\pi$ to π at every real number ω in that interval, right? That would take you a long time. So what you're likely to be doing is asking for what the DTFT is at some grid of points.

So you'll form a little grid. And it's on that grid of points that you want the DTFT. That's the only practical thing you can do. You're not going to compute it at all ω outside of toy examples like that. So if you had a numerical sequence collected in the lab, for instance, this is what you'd be aiming to do.

What's the other thing that's likely to be the case if you've got a numerical sequence collected in the lab? Any thoughts here? It's unlikely that my summation is going to go from minus infinity to infinity because I'd be waiting a long time to collect that signal, right?

So in practice, what we're dealing with are signals of finite duration, typically assumed to be 0 outside of that interval, though you might have reasons in some context for assuming otherwise. We're always going to take finite length signals. So the summation will be over a finite interval. And we're going to want to compute the DTFT on a finite group of points. And that makes for some simplifications. So let's see here.

You've probably heard people talk of FFT, or the fast Fourier transform. The fast Fourier transform is not a new kind of transform, so the name is a little bit misleading. It's a good way of computing samples of a DTFT. So you don't have to learn a new transform. We're still talking about this object, the DTFT. The FFT is an efficient way of computing the DTFT on a grid of points, given a signal of finite duration.

Now, I've got a lot on this slide, and I hope all of it is right. But let me talk you through the basic idea, here. So we're going to compute the DTFT on a finite grid of points. So that's the omega k's that I've shown you over there.

We've only got a finite duration signal. Let me say that it exists only from 0 to p minus 1. So the signal is 0 outside of that interval. And therefore, all the other terms drop out of this. So there's nothing new in this formula. This is just acknowledging that I only want to compute the DTFT at a grid point, and I only have a finite duration signal.

Now the interesting thing is that if your signal is 0 outside of this interval, well, that means your signal is completely specified if x_n is known to be non-zero only on the interval, let's say, 0 to p minus 1. So that's p values. Then you would hope that just having p samples of the DTFT will allow you to go the other way.

We know for sure that if I gave you the entire DTFT, if I gave you the entire DTFT, that you could go the other way, because we have this expression. If I gave you the entire DTFT, you would just plug it into here. And you'd get the time domain signal.

What's interesting, though, as it turns out-- and you might expect this. Since your signal takes non-zero values only at p points, you only need p samples of the DTFT to get an exact reconstruction. And here is the formula. And the derivation is not hard. I've omitted it here. It's the same kind of idea that we used in the full case. But you can actually-- using the values of the DTFT at these grid points, you can reconstruct the signal $x_{sub\ n}$.

So with these simplifications, you actually have a simple pair that gets you through the numerics. If you followed these formulas exactly as they're written, you'd end up doing work on the order of p squared, because you see each of these summations involves taking p products. And then you've got to sum them. But you've got to do it at p different frequencies.

And the same thing on the other side-- you've got to do p products, but you've got to do it p times. So it's order p squared computation. The fast Fourier transform is actually a clever way of using the symmetries associated with these exponentials to group the computations and make it much faster.

And you can actually reduce it to order $p \log p$. So it's a huge simplification. I've got some illustrative numbers down there. So the FFT actually is a major reason for advances in numerical computations, including signal processing of various kinds-- the fact that you can get this reduction from p squared to $p \log p$.

All right, I don't think I need to say much about the grid of points. But let's move on to thinking about spectral content of signals going through channels. All right, so we'll get closer to signals of the type that we're interested in, which are these signals in this case for on-off keying that are signaling 1's and 0's and that we're trying to get across a channel-- for instance, the audio channel. So this might be a typical finite length sequence. This particular case, we had chosen 7 samples per bit. That's why the shortest interval you see has 7 non-zero bits, there.

Here's the spectral content of the signal. And actually, I've taken this figure from an earlier version of the course, where we talked about the discrete time Fourier series, not the discrete time Fourier transform. The discrete time Fourier series turns out to be something very similar to the formulas I showed you for the FFT.

So the discrete time Fourier series, apart from a scale factor, is essentially a story built around this relationship. And that's developed in some detail in Section 13.2, but we're actually bypassing about this term to try and keep the story simpler. So when you see these plots, you'll see on a, at magnitude a of k , that's a symbol associated with the discrete time Fourier series.

These are actually Fourier coefficients associated with the periodic replication of this signal outside. So it's a discrete time version of the Fourier series you may have seen in 1803. All you have to do when you see a plot like this is think of it as a scaled version of the DTFT samples.

So we're just talking about samples of the DTFT taken at a grid of points. The scale factor may be off by a factor of p -- the length of the signal. But the shape is entirely told to you here. So think of this as samples of a DTFT. We're going from minus π to plus π .

What's down at the bottom here is imagining that you've sent this signal over a channel that could absorb the entire spectral content of the signal. So suppose you had a channel whose bandwidth-- suppose I had a low pass channel whose bandwidth has got some cutoff. This is low pass.

Suppose this bandwidth could absorb the entire spectral content of the signal. In other words, what I mean is that all these DTFT numbers that are significant actually fit in under this. So suppose your channel was such that it didn't attenuate the DTFT coefficients.

So the spectral content is unmodified when it gets through the channel. And so if you resynthesize the signal using this formula at the receiving end, you'd get back the same thing again, because the channel's not induced any distortion. Let me go past this and actually show you what happens when you start to distort the-- what goes across.

So here, what we're doing in this succession of experiments is sending that same signal through a channel with successively smaller bandwidth. So in the first case, everything goes through. There's no distortion. You get the same thing back again.

In this case, the channel actually has a cutoff that ends up zeroing out all spectral content outside of some frequency range. So it's a low pass channel whose bandwidth is not enough to take the spectral content of what you're feeding across. So what do you expect to happen?

Well, the higher-frequency components of the signal have been zeroed out. So what should happen? You expect the signal to be more rounded because it can't make these sharp transitions. It takes high-frequency content to make sharp transitions.

So what happens is you trim the spectral content by sending it through a channel that's not wide enough to contain all of the signal is that you get a more rounded signal at the other end. So you sent this. This is what you're receiving. This is the distortion that the channel has imposed on your signal. And you can imagine if you tried to find a place to sample this, you might run into some trouble.

If you go even more extreme, here is an even narrower channel. What comes out is even more rounded than what we had there because you've taken away more high-frequency components. The signal just can't wiggle that fast, so it takes its leisurely time going through its paces here. And you can imagine that you can be thrown off when you try and take samples. This is actually even more evident on the eye diagram, here.

So these are eye diagrams-- again, the same kind of thing. As you successively transmit fewer and fewer-- let's say less and less of the high-frequency content of the signal, what gets picked up-- what gets received is a more rounded version of what was sent in. And the corresponding eye diagrams that you construct-- well, at a certain point, I guess somewhere around here, you'd be a little nervous about trying to find a place to threshold and decide on what signal you have.

So this is not a noise issue. This is a distortion issue. It's distortion induced by the channel. And it can all be understood in terms of what the channel is doing to the spectral content of the input. I think we'll continue next time to get more insight into this and start on the topic of modulation.