MITOCW | 15. Modulation/demodulation

PROFESSOR: There's quite a bit I want to go through here. so we're going to talk today about modulation, which you've already gotten some notion of, and that's basically the task of matching a transmitted signal to the physical medium. And then we'll talk about demodulation, as well. Whoops.

So just to remind you of how we got into this story, we started off talking about bits that we had to get across to a receiver from a source. And we've been spending quite some time now focusing on this piece of the system, which is taking the bits, converting them to actual physical samples of a voltage, for instance, and then trying to get them over some physical medium. And then, at the other end, converting back to bits, all right?

So this is really a key piece of the system. If you can't get it over the physical medium, then you don't have anything. So we've been spending quite some time on that. We've talked about models for signals and models for systems, LTI models for systems, all in the time domain. And then we came to the frequency domain, which we said will make things a lot simpler, and actually is the way that people think about transmission on physical media, typically.

OK, so the actual math. Well, we've seen-- this is just review. We've seen that, most recently, that you can actually represent any signal as a weighted combination of exponentials, so this was the transform domain representation. And the weights here were given by the discrete-time Fourier transform. So you give me a signal, I can find for you the discrete-time Fourier transform, which then tells me how to assemble complex exponentials to get the signal of interest, all right? And so this Fourier representation of the signal is really the frequency domain thinking.

And then we saw how to apply that actually to a system. We have an LTI system, therefore characterized by a frequency response. We first introduced the frequency response as a way of thinking about what happens to cosines. Put a cosine in, and you get a cosine out that's scaled by the magnitude of the frequency response and with the phase shifted by the angle of the frequency response. And then we went from cosines, or in parallel we talked about exponential inputs, so inputs of this type.

And now we have more generally a signal that's represented as a weighted combination of exponentials of that type. Out comes the same weighted combination of exponentials, except each one is scaled by the frequency response as appropriate. And then comparing that with what we expect as a spectral representation for the output. We get this key relationship, which is relating the input and the output of a system, an LTI system, that's governed by a frequency response.

So now we're starting to think in terms of the spectral content of the input, all right, which is the frequency domain description of the signal that goes into the system. Then the frequency response of the system shaping that spectral content to give you the spectral content of what comes out, all right? So this is the language and the picture that we have, and it's all as simple as multiplication once you've figured out what the spectral content is of the signal of interest, once you have the frequency response of the system. So we've got to know how to do those pieces.

And then we talked most specifically about a physical medium that's close to what you're doing in the lab, which is the medium of, well, an acoustic channel driven by a loudspeaker, and at the other end are a microphone to pick up the signal. And I showed you these typical characteristics of loud speakers, the kinds that you'll find listed everywhere, three different speakers.

I mentioned last time that, when you look at frequency specs for speakers, people will typically only show you the magnitude specification because, for audio applications, the phase distortions are a little less important. They tend to not be picked up by the ear. All of these-- let's see they have passbands from-- this pointer's a little weak here, but-- passbands from around 100 Hertz to, let's say, 10 kilohertz.

So in that region, they pass signals more or less uniformly in at least the magnitude characteristic, and then near the edges they taper off. And some speakers will have bigger passbands, and will taper off closer to DC. Other speakers actually will not pass frequencies till you get till about, oh, what is that, 120 Hertz or so on this characteristic, but you've got to get way up before you get anything through that speaker. But nominally, we can think of speakers, since they're aimed at audio applications-- what? The ear here is something on the order of, let's say, 100 Hertz to 10 kilohertz.

OK, but the phase characteristic is important, too. It's not important, maybe, when you're talking about sending audio on a speaker, but in Audiocom in the lab, you're actually sending pulses across it. You're communicating something other than audio. You're actually trying to get a signal whose particular shape matters. It's not how you hear it, but what it looks like before you sample it, OK? So in settings like that, the phase characteristic is important, as well.

Now, you haven't explicitly probed the frequency characteristic of the speaker you're using. You could do that, but instead you've been looking at things like step responses in the time domain and constructing eye diagrams, but you could look in the frequency domain and characterize your particular channel for your laptop sitting in a particular place. You could look to see what the magnitude and phase are like.

OK, so I want to go through this exercise of looking at the spectral content of a signal you want to get across this audio channel, and then looking at how the audio channel shapes it, and then what you pick up at the other end. So just to give you a feel for how one thinks through this.

So the input in the typical application you have for Audiocom you have-- let's see, if you wanted to signal just a 1 and then all 0's. You would have 256 samples at height 1, and then everything from then on 0, OK? So what I'd like to do is think through how this pulse gets across the medium, but thinking it through in the frequency domain, all right?

So the first thing we have to figure out is, what's the spectral content of this pulse? By the way, if we understand it for one pulse, then we can understand it-- then we know it for a sequence of pulses because if we're modeling the system as time-invariant, once we figure out what one pulse does, we can figure out what a later pulse will do. It's just the same response delayed in time, OK? So the key to it is understanding what happens with one pulse.

So the spectral content of the signal is what we're interested in. And my question is, do you have any guesses as to what the spectral content might be, just roughly, qualitatively? Where do you think the energy of the signal is concentrated? What frequency ranges? Any thoughts? I'll need a hand up and a loud voice so I can figure out what's-- at least one? Yeah?

AUDIENCE: Low frequency.

PROFESSOR: Low frequency is a good idea because, for most of this signal, you've got essentially nothing happening, right? It's just flat. So you expect high spectral content at DC. But there is this sharp transition, so you might expect high frequencies associated with that. So do you think it might be low frequencies and then high frequencies, not much in between, or any thoughts?

OK, well, let's look at what it actually is. Let's work it out. So we're talking about a signal that's at height 1. For-- let's do the general case. So let's say it's for n samples, so from 0 to n minus 1 its height 1, and then it's 0 outside of that, OK? So suppose this is x of n. How do we determine the spectral content?

Well, we've got to write the-- we've got to compute the DTFT, right? So what's the DTFT? It's a summation x m e to the minus j omega m over all m. But in this case, it simplifies, right? Because there are only a few non-zero values of the signal. So this is going to be-- let's see. It's going to be x0 e to the minus j omega 0 plus x1 e the minus j omega 1, and so on. That's going to be 1 plus e to the minus j omega plus e to the minus j2 omega plus e to the minus j N minus 1 omega. So that's the DTFT.

But till you work with that and get it in a form that you can make sense of, you still don't have a feel for where the frequency content is, right? You've got to-- the best way to get at that is to think of what the magnitude of this will be. And even then, it's not obvious how to think about the magnitude of a sum of complex numbers like this, so you've got to play with it a little more.

OK, well, this is a geometric series, right? Each term is obtained from the previous one by multiplying by to the minus jam omega. And so if you've got the sum of a finite number of geometric series of this type, what do we have? We have that as the sum, right? You agree? So this was the factor by which we multiply each term. Sorry. And we've got N such terms, so you're summing N terms of a geometric series.

Well, we might be getting closer here to extracting a magnitude, but you really want to do a little bit more massaging here. Let's see. If I make this e to the j omega N over 2, then here's e to the j omega N over 2, minus e to the minus j omega N over 2. And this is a trick we've done-- we've played a few times before.

Right? I've just rearranged things. Let's see. How have I helped myself here? Have I helped myself at all? So what is-- what does that simplify to? Well, the factor in front I can write as some phase term e to the minus j omega N minus 1 over 2. And what's this? Anybody? Numerator? Does the numerator remind you of anything? Sine? Sine omega N over 2?

And the denominator, sine omega over 2, right? So now it starts to look a little bit more manageable. If I wanted to get the magnitude of this, well, the magnitude of this is going to be the magnitude of this piece times the magnitude of that piece. What's the magnitude of the first term here? Just 1, right? It's e to the j something, so its magnitude is 1.

So here's the magnitude of the DTFT, so that's the spectral characteristic, and that's something that we can plot.

AUDIENCE: Question.

PROFESSOR: OK. Sorry. Question, hi.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Sorry. Say that again?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Did I make a mistake somewhere here, or? Oh, this thing? This term here? I was trying to combine numerator and denominator here.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Which part? Sorry. I have to stand where you are to see if I made a mistake because it's hard to see close up.

AUDIENCE: Over the-- like, when we go from the 1 minus e to the minus 2 again to the other one, why are we dividing by two?

PROFESSOR: Oh, what are we-- oh. Why are we dividing by 2? Because when you multiply it out, that's what it takes. I'm trying to group things to get something interpretable. I don't know what this is, but I know what the numerator over sine looks like, so I'm trying to make this a little bit more equally distributed, right? So if I pull out that factor, what's left? Looks like part of a sine. OK?

So I'm just-- so it's this time this gives me the numerator here. And this times this gives me the denominator here. So it's just rearranging terms. We've used this trick before. Any trick that works twice is a method, OK? So we really have a method here. It's not just a trick. If it works three times, you can make a religion of it.

OK, so that's the derivation we have here. What's the height of this at the origin? Let's just focus on that term.

OK, so this is the magnitude we're talking about. What's the height at the origin, at omega equals 0? Well, now you can use L'Hopital's rule, right? Because omega is something that varies continuously. So for small values of the argument, you're really looking at something of height N. And then, when is the first time that this goes to 0?

Well, for small values of frequency, the numerator is not changing sign, and this first goes to 0 when you get to omega equals 2 pi over N, to private capital N. So actually, instead of saying all that, I should just draw you a picture. There's a picture of one particular case. So this is a case where, actually, the pulse didn't start at 0. It was symmetrically located around 0, OK?

It was a pulse of length 11 symmetrically located around 0. And because it was symmetrically located, this phase factor went away, and all you're left with is the sine omega N over 2 divided by sine omega over 2. So you're looking at the actual DTFT of a pulse of that type, OK? So this started at minus 5 and went to plus 5, and was 11 samples long and was 0 everywhere outside of that.

So that's what this function looks like, the sine omega N over 2 divided by sine omega over 2. Does it remind you of a function you've seen before? Sinc? A sinc function? It's very close to a sinc. The sinc, though, had just frequency in the denominator. It didn't have sine of something in the denominator.

And the reason this appears is, remember that transforms and frequency responses have to be periodic with period 2 pi. So it certainly wouldn't be possible for the transform of a signal to be a sinc because there's no periodicity in the sinc. But when you work it out carefully, you find that it's something close to a sinc, but one that has exactly the right periodicity, so this thing will repeat periodically with period 2 pi, exactly the way it's supposed to.

So it's sort of sinc-like. For small values of omega, the sine is essentially just omega over 2, and it is essentially a sinc. But when you get to larger values of omega, this thing starts to play a role.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Yeah? Sorry?

AUDIENCE: Does the magnitude mean anything?

PROFESSOR: I haven't-- I'm not plotting the magnitude now. I was plotting the actual DTFT for this case. So in this symmetric case, the actual DTFT is the sine N omega over 2 over sine omega over 2. So I was plotting the actual DTFT. And the magnitude of the DTFT I get just by taking absolute value, right?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Well, I did this for a case of a pulse starting at time 0, OK? So this factor came purely from where I located this on the time axis. So different positions of this pulse on the time axis will modify this factor but won't touch this factor, OK? Shifts in time correspond to multiplication by e to the minus j omega something, right? You've seen that in recitation.

So whenever you shift the pulse in time, what it does to the transform is it leaves this part intact and affects this factor, but that doesn't change the magnitude of the transform, OK? So in the case of a symmetrical pulse, symmetrical about the origin, actually that phase factor goes away, and the actual transform is just that piece without the e to the anything. Sorry. I jumped over a few steps in describing that.

Just to go back a second. This is not a sinc function, but it's actually referred to as a periodic sinc. It's got a fancier name, also. It crops up all over the place. Height N at the origin and the first zero crossing at 2 pi over cap N. So as you make the pulse wider in time, you make it narrower in frequency, right?

As N becomes larger, you make this wider in time. The main lobe of this frequency distribution gets more concentrated, and frequency gets closer to being a DC signal. Makes sense, right? The longer that this stays constant, the more the signal looks like just DC and the more the frequency is concentrated at the origin.

But what you can see here is there's actually a full spread of frequencies. It's not that there's just low frequency for the flat parts and high frequency for the vertical edge and nothing in between. There's actually a full spread of frequency components that it takes to make up that step.

OK. If you had a pulse that wasn't centered-- this is just to show you. Here is a pulse-- actually, this is not centered. It's only 10 long, but the magnitude here you're only seeing half the frequency scale, so 0 to pi essentially, except this is in terms of f, which is omega over 2 pi. You get the same kind of magnitude characteristic, but now because you've shifted it off-center, you've got a linear-phase characteristic, and what you're seeing here is a linear-phase characteristic, except every time you have a flip in sign, you jump the phase by 180 degrees, right?

When you change the sign of something from a plus to a minus, that's like adding 180 degrees to the phase or subtracting 180 degrees to the phase. So you can spend time on all of this and make sense of it. But the basic idea is that you get the sinc-like distribution and frequency.

OK, so let's get back to the particular pulse that's of interest to us, which is that pulse. So it's the same kind of thing, except N is 256. And what I've plotted for you here is the magnitude of the DTFT. It has the sinc-like shape. I haven't actually plotted it as a continuous function of omega. Instead, I've used the FFT. You remember? We talked about the Fast Fourier Transform.

So what the fast Fourier transform is going to do is, if the actual magnitude DTFT was some continuous thing like this, the fast Fourier transform is going to give me samples of it, as many as I want. But the more samples I ask for, the more work I have to do, of course, OK? I asked for 48,000 samples of the DTFT so that I could get a nice big spread here.

If your samples came from sampling at 48 kilohertz, for instance, then the rightmost end that corresponds to pi in terms of actual frequency would correspond to the sampling frequency divided by 2, so that's 24 kilohertz sitting there. So I actually have 24,000 points for 24,000 Hertz, so I've got one point at every Hertz position, but I could pick anything else.

The other thing I wanted to mention was that the reason I could do this is because I'm using the FFT. Because I told you if I just did sort of simple-minded implementation of the formula, I would take order p squared computations, where p is the length of the signal that I'm looking at. If I use the FFT, I go from p squared to-- I go down to p log to the base 2 of p, all right?

Well, the number of points I have here is 48,000, so going from p to log 2 p is going from 48,000 to 16, which is a factor of 3,000. So the difference is I'm sitting at the terminal and I hit the Return key to get the FFT, and maybe at 0.1 seconds later, I get the answer, versus if I didn't use the FFT I'd wait five minutes, and then I wouldn't be trying to put together these figures for you. OK, so the FFT really makes a real practical difference, and it really revolutionized how numerical computations were done.

OK, so here you now see the full spectral distribution if you're willing to let your eye interpolate between these samples that I've got. You see the full spectral distribution of that rectangular pulse, OK? So a short pulse in time. It's got certainly high DC content, but it's got tremendous frequency distribution, all the way out to high frequencies. In fact, all the way to the end. You're still seeing frequency content. It's visible to the eye. So all the way out to 24,000 Hertz, and you could keep going all right?

This is not a sinc. It's a sinc-like function because if you extended it, it would go back up again. It's got that period, 2 pi, exactly as it should. But for all intents and purposes, for small values of frequency, the sine omega over 2 here-- this is essentially omega over 2, and this is a sinc-like function, OK?

Now, this is different from what you saw before. Before you had a constant segment and frequency. At least, I think it's different from what you've seen before, unless you've done examples in recitation. But before what we had was, for instance, trying to get a frequency response that was like this, right? We ended up with a unit sample response that was a sinc function in time. This is going the other way. This is a rectangular function in time, giving rise to a sinc-like distribution in frequency.

All right. Let's zoom in a little bit just to see what we have here. And this is exactly what we expect to be seeing, the sinc-like distribution. The height there is 256, as it should be. That's the N. And then this should be at 2 pi over capital N. If you translate that to what it means in actual frequency, this first null there is at 187.5 Hertz. That's 48 kilohertz divided by N, so that's--

OK, so that's that number. So that gives you some idea of how this thing is spread in frequency. OK, so if this pulse is applied directly to the loudspeaker-- well, here, the loudspeaker passband goes from about 100 Hertz to, let's say, 10,000 Hertz. You see there's huge amounts of the energy that's not in the passband of the loudspeaker. All right? So this is not going to do very well if you just directly apply that pulse to the loudspeaker. You have to match the frequency content of the input to the frequency response of the system you're trying to get this over. OK.

Now, just as an experiment before we get back to sending this across a loudspeaker, let's take a look at what happens if we send that rectangular pulse over a lowpass channel. So let's save the bandpass channel, which is a little more involved, for later, and just look at what happens if we send this pulse over a lowpass channel.

So I'm thinking about a channel whose frequency characteristic-- this is h of omega-- it passes low frequencies, and then it truncates higher frequencies. OK? And what I'm going to do is send in an x of N, which is this pulse in time. 256 at height 1, and then everything 0.

All right, so if I'm thinking of it in the frequency domain, then I take the spectral content of the signal, which we've just worked out, and multiplied by the frequency response characteristic, which just basically selects out the frequency content that's in the passband and rejects everything else. And then that gives me the spectral content of the output, and I can translate that back to what happens in the time domain.

So here is-- well, actually, let's take a zoomed-in version. I've taken a lowpass filter, where this cutoff corresponds to actually a cutoff at 400 Hertz if you're thinking in terms of the underlying waveform. So what we've done is take the rectangular pulse, put it through a lowpass filter-- and I'm assuming an ideal lowpass filter-- passes everything in this frequency band and nothing outside of it, OK?

So here again, we see we're selecting out part of the spectral structure of the input, but there's huge amounts of the energy of the input that are being left out of the output of the filter. Now, in the time domain, look at what this corresponds to. It's an approximation to this pulse, but it's one in which all the high-frequency content has disappeared because we've only let the low-frequency pieces through.

So what you have is a very rounded kind of pulse. It spreads out well over the 256 mark, so here's the 0 to 256, but this thing actually spills over into adjacent bit slots. We've taken out the high-frequency components, so it can't make any sharp turns anymore. It's got this lower-frequency wiggling. As you can imagine, if I made this even smaller, my wiggling would become even more leisurely here, and I would spill further into the adjacent bit slots, OK? So this is what lowpass filtering will do to that rectangular pulse. Any questions on this piece? OK.

Now, we've actually seen examples of this type. I flashed these up last time. It's the same idea, except it was not just a single pulse. It was a succession of pulses like this. So you could do the same thing. Have a succession of pulses like this. Take its DTFT to assess the spectral content. I'm showing you not actually the DTFT but something proportional to it here.

This is-- so ignore the labels there. Think of this as essentially the DTFT. To within the scale factor it is. So you can see that I have spectral content all the way out to the edges. But now send it through a lowpass channel which zeros out all the spectral content outside of some central region, and what you have coming out the other end is something that can't take the sharp turns anymore, so it's much more rounded.

And as you narrow it down still further, you get even more rounding, and you get a spilling over. You see this sharply confined rectangular pulse now spills over into the adjacent slots. And you can go still further on that. The top plot here is the same as the bottom plot in the previous one. But this is just showing the sequence.

So as you narrow it down further and further and further, what comes out the other end gets much less distinctive in its features, OK? It can only-- this only has very low-frequency content, so it can't do any sharp turns. And this is an eye diagram corresponding to-- in each of these received signals-- just to show you how detection gets difficult if you've got a lowpass channel and you've got this pulse that's not well-defined.

Now, how might you actually-- how might you get a better-defined pulse for a given channel? So if I gave you this channel-- we sent in this pulse of length 256 samples, and we got something that we didn't like because it spilled over into other slots. The reason it spilled into other slots is we were cutting out too much of its spectral content. What could you do to this pulse to get more of its energy across that bandpass channel? Yeah? Sorry?

AUDIENCE: Set a longer pulse.

PROFESSOR: Set a longer pulse. How does that work? Now, you see, if you make N longer, you make the pulse longer, you shrink this correspondingly in the frequency domain, right? We said that on these sinc-type characteristics, the height was N. This first null was 2 pi over N. Make the pulse longer, you pull the main lobe in tighter, and more of this is going to go through, OK? And you can see that clearly. That's something you can explore in your experiment.

So if you wanted to get more clearly defined output for a given lowpass channel, you might want to increase the length of your pulse. Of course, that's going to slow down your signaling rate, so there's a trade-off involved, right?

Now, we've actually-- this is just to sort of step back and point out that what we're seeing here are some properties that are inherent to Fourier transforms, OK? So it's typically the case that if you've got a signal that's wide in time, it's narrow in frequency. And if you make it wider in time, it gets narrower in frequency.

In fact, as I mentioned up there, the uncertainty principle in physics really comes from this result. It says that the spread in time times the spread in frequency has some lower bound, OK? So this is some number that's strictly positive. So you can't make a signal arbitrarily concentrated in time and concentrated in frequency. If you make one small, the other one will have to grow correspondingly. So actually, the uncertainty principle in physics is precisely a theorem in Fourier transforms if you study it.

Here is another such complementarity or duality. The smoother you make a signal in time, the more sharp it is in frequency. Let's see. We saw that here, for instance, right? We had a signal in time, namely the unit sample response of the ideal filter. This was-- this didn't have any sharp edges to it. It was a sinc function.

You got something that's smooth in time, it ends up having sharp edges in frequency, OK? And the more smooth you make it in time, the sharper it gets in frequency. And this we've already seen. This is the kind of trade-off we're talking about. So these are characteristics to be on the lookout for. In fact, I just did a little experiment here.

What if we decided not to try and send a rectangular pulse over, but we smoothed out the pulse a little bit to get rid of that sharp edge? So that's another way that you can try and get a pulse over a lowpass channel like this. So you see, what I'm trying to do is not have the sharp discontinuities, the 0 to 1 and the 1 to 0. I want a more rounded behavior in the time domain so I get a sharper concentration in the frequency domain, and you can actually see how that works.

So what I've actually done here is, instead of a signal that-- well, if you'll permit me, I don't want to draw these stem plots, the lollipop figures, because they get painful to draw, but just assume that this is 256 such things. That's the rectangular pulse we were trying to send before.

What I've done now is instead say, well, let me around the edges, so I'm going to have a half cycle of a cosine for that edge, OK? So that's what-- these are the samples I'm going to use at this edge, and I'm going to have a half-cycle of a cosine at this edge. This is actually something that's used quite a bit in applications.

All right, so what have I done? I've removed the sharp edge and gotten a much-rounded transition. In fact, if you're thinking of continuous functions, the original had a discontinuity, whereas here I've got to take two derivatives before I encounter a discontinuity because the function itself and its slope at these ends are well-matched, so I've got to differentiate twice before I get a discontinuity. So actually, it has quite some smoothness to it.

Smooth in time means more tightly concentrated in frequency. So all I'm doing on the next slide is showing you, so your eye can compare, what the spectral content is of the original pulse and of this rounded pulse. And for the rounded pulse, I actually just flipped it over so that you can compare more easily. So this is the negative of the DTFT magnitude of the shaped pulse.

And you can sort of see right away here that the frequency content has essentially settled out. It's almost all contained in this smaller region, whereas the rectangular pulse had a frequency content that went way off to high frequencies, right? Went off to 24 kilohertz, for instance, in our example. So the frequency content of the shaped pulse is much more tightly contained, and you have a much better chance of getting that across a lowpass channel.

So this is just a little bit of shaping that can make a big difference in terms of adapting the signal you're trying to send to the channel. So if I look in the time domain, sending these two pulses over the same lowpass channel, you can do a visual comparison of what comes out.

So here's the original 256 rectangular pulse coming out the other end. Here is my shaped pulse coming out the other end. And you can see the shaped pulse is much more tightly confined in the bit slot that I assigned to it, so it's much more tightly confined around the 256-sample width, OK? So this is another thing that is done, and it's done by thinking in the frequency domain. People designed these pulses thinking in the frequency domain. They're not doing convolution.

OK, that was all lowpass, but we want to look at bandpass, so just a couple of quick examples there. So we're back to the speakers. And we're taking our rectangular pulse and applying it to the speaker. So here's what the spectrum looks like after ideal bandpass filtering. So I'm not actually filtering with the speaker characteristic. I'm assuming an ideal bandpass that extends from 100 Hertz to 10,000 Hertz and zeros out everything outside that range.

So I send in the spectral characteristic of my rectangular pulse. I shape it with the bandpass. So what comes out is something that has this spectral characteristic. And you can see that the frequency content is sharply limited 10 kilohertz above 0, and then there is actually a central region that's entirely missing. So remember, this had to go originally up to 256, but because we've chopped out the center portion, we're only going out to 150-something out there.

So actually, if I zoom in, you can see that a lot more closely. So this is a zoomed-in version of what comes out from a loudspeaker, from a bandpass filter, if I send in a rectangular pulse. The very low frequencies are entirely missing, and we saw in the previous characteristic that the very high frequencies are missing, as well.

So what's the shape of the pulse that you get out? Not very good. Because the low frequencies are missing, this thing tends to sag in the middle. It can't hold up DC. And it can't make the very sharp transitions, either, so there is a more leisurely transition. But again, this can't stay at that level. It actually asymptotes.

So the actual pulse occupies-- before it settles, occupies way over the 256 bits that I've allotted to it, OK? So taking that rectangular pulse and directly putting it on the speaker is going to give you something not pretty at the other end and something that you cannot signal with. So the question is what to do about that, OK?

And the answer to that is this thing that we call modulation. We've already seen it in different forms. Let's think about it now in the frequency domain. OK, so here's what we're going to do. Want the big stick of chalk. OK.

We're going to shape the spectral characteristics of the signal. We started off with some time-domain signal links, x N, corresponding DTFT x omega. It wasn't well-matched to our channel characteristics, and so what we're going to do is multiply by some carrier frequency. And this is simple amplitude modulation. We're just referring to it as modulation here. We get some signal out here.

And the question is, what is the frequency characteristic of that signal? OK, we've already seen what the frequency characteristic of the input is. What's the frequency characteristic of that signal?

So just to give you an example, we had our x of omega looking something like the sinc shape. Remember height N there? And the question is, what's the spectrum of this? And there's a computation up there that I don't want to go through, but it shows you that the answer is actually quite simple.

So if I call this-- let me call this t of N because it's the signal we're going to transmit. Here's what the spectrum of the transmitted signal looks like. There is omega. There's pi. Here's omega c. OK, so the prescription is simple. Take the spectral characteristic and replicate it at omega c and replicate it at minus omega c, and scale by 1/2.

So what you're going to get is this characteristic here, this characteristic here, and the height will be N over 2. It's that simple. So you can go through the math. When you're done with the math, what it says is that the spectral characteristic of this modulated signal is the spectral characteristic of the envelope of the baseband signal, or the envelope, but translated to the position of the carrier.

So you can begin to see how this is going to help us shift a signal to get it across a bandpass channel. We started off with something that was not well-suited to the loudspeaker we now have a way to shift its energy to get it right in the passband of the speaker by picking the carrier frequency appropriately. All right?

I think I'm going to skip some of this. But let's look at what this does. It's really this picture, but I just want to show you how it works with actual waveforms. So here is the rectangular pulse times the cosine. I've picked-- what did I pick? 1 kilohertz as my carrier? Yeah. 1,000 Hertz was the carrier.

1,000 Hertz sits comfortably in the passband of a speaker, so it's a reasonable choice. By the way, in typical AM, the carrier frequencies are much higher than-- or the ratio of the carrier frequency to the rate of variation of the envelope is much higher than what we have in these examples. So the audio channel is actually very challenging.

OK, so this is my modulated signal. The question is, what does its spectrum look like? So I run it through the FFT, and indeed I get the replication here. Let's zoom in a little bit so we can see it a little more closely.

So what I have is those two sinc-like spectral characteristics, but translated to sit centered at 1,000 Hertz and minus 1,000 Hertz. Remember that this is 0 out here, OK? And the height, well, it's now it's 128, which is half the 256 that I had before.

So in terms of positioning this within where the loudspeaker will transmit it, we were taking the lower cutoff of the loudspeaker as being around 100 Hertz. This is where 100 Hertz sits. So you can see that a huge amount of the energy of the pulse is getting through.

It's at the wrong frequencies. We'll have to deal with getting it back. But at least getting the energy across is working here, OK? And the upper cutoff of the speaker is way off over here. So this is 10,000, but I'm showing you a zoomed-in version, so the upper frequency is way off.

In fact, that also brings up the idea that you could actually do this trick multiple times. You could actually pick another carrier frequency somewhat higher than this with some other modulating signal on it and tuck that in there, as well. So you could simultaneously transmit messages on multiple carriers through that same speaker, and you'll be exploring that, as well. OK.

So what does the-- what's the time-domain signal that corresponds to this looks like? So when you get it, you impress this modulated signal on the loudspeaker, on this bandpass filter. What's the output of the bandpass filter? You can see it's almost exactly what you put in, OK? There's a little bit of distortion at the different places, but it's basically exactly the pulse that you put in. So almost all the energy has gone through, and you don't have significant distortion because you've squarely placed the energy in the passband of the filter.

OK, now how do we recover? How do we get back the original baseband signal? Well, it turns out it's very easy. Let me actually do it in pictures, and then we'll look at the math. This is what's coming in to our receiver, and now we've got to process this to get back a signal that has this spectrum.

We've learned a trick, which is modulation. Multiply it by a cosine, and you'll take the spectrum and replicate it at omega c and at minus omega c. I'm going to use the same trick again. I'm going to take the received signal multiplied by a cosine of the same carrier frequency, and what's that going to do?

Well, my scale is getting bigger here each time. So here's my omega c. Here's my minus omega c. What's coming in is this signal. I'm going to multiply it by cosine omega c, so what does that do in the spectrum? It takes this spectrum, replicates it at omega c. So what does that do?

Well, it puts a piece here, and it puts a piece here, at 2 omega c, right? Because I've taken this spectrum-- you've got to imagine the change of scale so I can draw all of this. I've taken this spectrum and I've placed it centered on omega c. And this is now-- this is the N over 2 here.

Oh, but now it's going to be N over 4, right? Because I divide by 2. And then I take the same spectrum and I center it on minus omega c. OK, so I've got the N over 4 piece here, but I'm going to have the other-- oh, sorry. I drew it in the wrong place.

I'm going to center it at omega c. So this is going to end up at minus 2 omega c. And the replication here. So there's going to be a second one sitting here at the origin. So the net effect at the origin is that I get the original spectrum, but scaled by a half, and then I get vestiges of this, if you like, centered at twice the carrier frequency, OK?

And that's actually what the algebra shows. The algebra is very simple. You're receiving this. Multiply it again by a cosine. So take the received signal and multiply it again by a cosine. Well, if you substitute for what the received signal is, you get x n times cosine squared. Cosine squared, if you use a standard identity, splits into these two terms.

Let's see. Do I have that right? Yeah. So here's the 0.5 x n sitting here at the origin, and here's another term, which is x n times the cosine. So this is like a modulated signal, but modulated by twice the carrier frequency. So what does this translate to?

Well, it's actually 0.5 times x n times cosine 2 omega c. What does that do in the spectral domain? It's going to take half of that and place it at minus 2 omega c and plus 2 omega c, so it's completely consistent with this picture. So what is it that we have to do now to extract the signal of interest?

Just a lowpass filtering here, OK? So if you can select out this piece with a lowpass filter, you've recovered the signal of interest. You can adjust the scale factor, too, so you can have a lowpass filter with a gain of 2 and you've recovered your original signal. All right, we'll build on Monday. And that'll be the last lecture on this material. Relative to the calendar, we're just sliding forward, so we'll wrap up this stuff on Monday and then continue with packets.