

MITOCW | watch?v=3D51nqZ-97Q

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.

DENNIS Hello, welcome. As you expect, the big news is that there's a quiz next week. So hopefully, that's all
FREEMAN: square, everybody knows everything.

Questions, comments? Everybody's pretty comfortable with what's going on with the quiz? If you have a conflict, please contact me shortly, because I have to arrange rooms and proctors. No questions? Good.

So last time we looked at discrete time frequency responses. Lots of things were similar. A few things were different. The first difference we saw was that in discrete time, we like to think about complex geometric sequences rather than complex exponential sequences. And the effect of that is that the frequency response lives on the unit circle.

So if you try to think about a discrete time system, say, characterized by a z-transform and think about what would be its output when there is a sinusoidal input, just like CT, if you have a sinusoidal input, there's going to be a sinusoidal output at the same frequency. But the amplitude and phase can be shifted. But to compute the amplitude and the phase for a DT system, you look at the unit circle of the z-transform.

That's as distinct from the Laplace transform in CT, we looked at the $j\omega$ -axis. So now, everything we used to say about the $j\omega$ -axis, we say about the unit circle. Other than that, things are extremely similar.

And just to boost your confidence so you know exactly what's going on, here's a question. What if I gave you a system that had a single pole and a single zero, and say the pole and zero are related that way? How would you think about that system-- Low pass, high pass, all pass, band pass, band stop, high stop, low stop, all stop?

AUDIENCE: I say low pass.

AUDIENCE: Low pass.

AUDIENCE: Band pass.

DENNIS So some very early votes that are completely correct. You had the recitation, so then there is no

FREEMAN: excuse. So if I see any wrong answers-- and I see wrong answers.

[LAUGHTER]

Ah! So maybe I should have the recitation instructor-- well, no, no, OK, everybody votes. So high pass, low pass, band pass, all pass, band stop, none of the above, vote. The vast majority is 100% correct. The vast majority is correct, but it's not quite 100%.

So the answer is 4, all pass. Why? Can somebody make an argument why that should be all pass, other than you saw it yesterday in recitation?

[LAUGHTER]

I mean that's a good reason. Yeah.

AUDIENCE: It can no longer go through the axis below omega. The magnitude of the zero divided by the magnitude of the pole is pretty large. And then as you go to infinite omega, it goes to 1. So I think it's going to pass.

DENNIS So I'm a little bothered by infinite omega. Where should we be-- when we think about things as low
FREEMAN: pass, high pass, whatever, what are we thinking about when we say something is a low pass filter? So somebody says this. So what's that mean? That's correct.

So a large fraction of this class is learning to interpret sign language in the air. And so what's it mean if a system is low pass? What do I look for when I look at this diagram to see that something's low pass?

AUDIENCE: Look for lower frequencies.

DENNIS Where is low frequencies on this diagram?

FREEMAN:

AUDIENCE: If you cut in half--

DENNIS Cut it in half, top and bottom, take that side.

FREEMAN:

[LAUGHTER]

So low frequencies are there. So we're thinking about frequencies being e to the j omega, because what we're thinking of or thinking about, we'd like to think about cosine of omega n . But we want to

make that look like an eigenfunction. So we want to make that look like $e^{j\omega n}$ plus $e^{-j\omega n}$, that kind of thing-- then divide by $1/2$ probably. So the question is-- so this is all pass. And that's actually kind of tricky, if you think about how it would become that way. Yes.

AUDIENCE: [INAUDIBLE] at any point on the unit circle--

DENNIS Yes.

FREEMAN:

AUDIENCE: It stems from the [INAUDIBLE].

DENNIS That's precisely right. So somehow, if you took this distance and that distance, which is intended in

FREEMAN: this figure not to look the same, but if you took the ratio-- because this one is a zero, so it's in the top, and this one is a pole, so it's in the bottom-- and if you took that ratio, somehow if you took that ratio divided by that ratio and compared it to this one divided by that one, those numbers would be the same thing. That's what would have to happen.

You would compute the magnitude of the frequency response by taking the ratio of the magnitudes of the vectors that connect the zero to the point of interest and the pole to the point of interest-- the point of interest being on the unit circle. It's not at all obvious to me that that ratio is a constant, if you move along a circle. That's another one of those things that I think I need Greek to convince me.

AUDIENCE: [INAUDIBLE] because he figured it out.

DENNIS Oh, somebody knows.

FREEMAN:

AUDIENCE: Yeah, the guy [INAUDIBLE].

DENNIS Wonderful. So you should explain it to me. [LAUGHS]

FREEMAN:

AUDIENCE: He explained it to me.

DENNIS [LAUGHS] Putting you on the spot.

FREEMAN:

AUDIENCE: Like right now?

DENNIS Oh, that's OK. If you want to, that's fine.

FREEMAN:

AUDIENCE: [INAUDIBLE] So if you have any point on the unit circle?

DENNIS Yeah.

FREEMAN:

AUDIENCE: You draw like a line from that point to the origin.

DENNIS So if you have any point on the unit circle, say that one--

FREEMAN:

AUDIENCE: So you draw a line halfway-- could you do it on the other side, on that side [INAUDIBLE].

DENNIS So you like this side better, OK.

FREEMAN:

AUDIENCE: So you draw a line from that point to the origin. [INAUDIBLE] 1 over a and a on the x-axis.

DENNIS So 1 over a and a.

FREEMAN:

AUDIENCE: And draw set of lines from those to the point of [INAUDIBLE].

DENNIS Up here?

FREEMAN:

AUDIENCE: Yeah. So now like you have two triangles that have the same side and the same angle. And if you just use the law of cosines--

DENNIS You have two--

FREEMAN:

AUDIENCE: So you have a big triangle and a small triangle.

DENNIS So we have this triangle and we have this triangle. Those are the two you want?

FREEMAN:

AUDIENCE: And yeah. So like they have-- the line of the leftmost is like the same for both. And the angle at the center is the same for both.

DENNIS So this angle is equal to this angle?

FREEMAN:

AUDIENCE: No.

AUDIENCE: The central angle. The central angle with the--

AUDIENCE: You take the whole be triangle and--

DENNIS Take the big triangle--

FREEMAN:

AUDIENCE: And the origin.

AUDIENCE: And then the one you just--

AUDIENCE: The angle at the origin.

DENNIS And the angle at the origin.

FREEMAN:

AUDIENCE: Yes, that one.

AUDIENCE: So that angle is in both the big triangle and the little one.

DENNIS That angle is in the big triangle and the little one? Yes!

FREEMAN:

[LAUGHTER]

Got it!

AUDIENCE: And the size of the bottom are like $1/a$ and a for the two triangles.

DENNIS Yeah.

FREEMAN:

AUDIENCE: So you just use the law of cosines and--

DENNIS That's very good.

FREEMAN:

AUDIENCE: Isn't that pretty?

DENNIS What was that?

FREEMAN:

AUDIENCE: Isn't that pretty?

DENNIS That's amazing. So wonderful. Congratulations.

FREEMAN:

[APPLAUSE]

So you're also working on the triangle thing? Don't have that yet. OK. I'll keep checking in. So the way people who don't do quite so well in trig do it, so you can also just empty-headedly substitute z equals e to the j omega. And then if you want to make the top look a little bit more like the bottom, you could take out the e to the j omega part. And then you see that these look like complex conjugates.

So this is kind of the math way. And what you said was the trig way. I like the trig way much better, the [INAUDIBLE] way, yes. So complex conjugates, same magnitude done.

So the whole point was just to motivate the idea of thinking about an all pass. Why would you build an all pass? What's an all pass good for? It's 6.003 exercises, of course. So an all pass doesn't change the magnitudes. So what's it good for?

AUDIENCE: Changes the phase.

DENNIS Changes the phase. So if you were to think about what's the phase of the all pass, you get a function

FREEMAN: that depends differently. So you have to compute the angles of these vectors. So this one-- the angle at the zero starts out at π . The angle associated with the pole starts out at zero.

There's actually a minus sign in here, because I took $1 - z$ compared to $z - a$. So I put the z at the two places. So there's a net minus sign in the whole transfer function. So the π cancels out the π . So we start at zero at zero.

Everybody know what I just said? So the angle associated with the zero is π . But there's also a negative sign associated with the transfer function, because we would normally like to write $z - a$ divided by $z - a$ pole. And I didn't write it that way. The top one, I wrote backwards.

So there's a net minus sign. Because of the minus sign, the phase starts out at zero instead of π .

That's how that got to be zero.

So this starts out at pi-- the angle starts out at pi. It becomes less, and then it goes back to being pi. But the angle associated with the pole starts at zero and goes around the pi.

So the net from the pole is a pi shift, which accounts for why this goes from 0 to minus pi. You can all do that, right? So the idea is that the all pass would alter the phase. And that's one of the things I wanted to talk about.

What's the effect of changing the phase? And I'll do a demo now, thinking about how it affects audio. And at the end of the hour, thinking about how it affects video.

So to think about this, I took this fraction. And I used a equals 0.95. And I took a waveform that I generated with Python. I just added together three tones in the ratio 1, 3, 5 in frequencies, so it sounds like more harmonic. And when you add those together, you get a sound like this.

[TONES PLAYING]

What you hear is a very dominant tone-- a very dominant pitch, because I made the harmonics be integer multiples. So I made the overtones be integer multiples of the fundamental. So the fundamental was at 8,440. And I put in some third harmonic and some fifth harmonic. So you get something that sounds very constant pitch.

Then I took that waveform and put it through this filter. And that changes this looking wave form into that looking waveform. Those are real waveforms. And then it changes the sound, so that instead of sounding like this [TONES PLAYING], it sounds like [TONES PLAYING].

[LAUGHTER]

You can all hear the enormous difference, right? So even with headphones, I can't hear any difference at all. And that's kind of a property of the way you hear. You kind of hear-- I'll say some caveats about this. This is not exactly true.

So Ohm's law of hearing-- the same Ohm-- says that you hear sounds by the frequencies that they contain. Roughly speaking, you hear sounds according to the magnitude of the frequencies that you hear. Now, that's not quite true, but it's sort of true.

So let me do a more complicated example. So the first sound is a consonant vowel consonant.

[AUDIO PLAYBACK]

- Bat.

[END PLAYBACK]

DENNIS And then I put it through the same all pass. So I screw up the phases. So it goes from looking like this

FREEMAN: to looking like this and sounding like--

[AUDIO PLAYBACK]

- Bat.

[END PLAYBACK]

DENNIS Hard to hear the difference. Here's an even more complicated sound. This is Bob Donovan's thesis.

FREEMAN:

[AUDIO PLAYBACK]

-

[INAUDIBLE]

but you see

I have no

brain.

[END PLAYBACK]

DENNIS Put that through the all pass, and you get--

FREEMAN:

[AUDIO PLAYBACK]

-

[INAUDIBLE]

for seeking

[INAUDIBLE]

but you see

I have no

brain.

[END PLAYBACK]

DENNIS If I listen to that on headphones, I can actually hear a difference. But I study hearing, and I know what
FREEMAN: to listen for. It's still difficult to tell the difference. Now, that's because I kind of cheated.

To let you know how I cheated, let me do something more drastic to the phase. Let me play the waveform backwards.

[AUDIO PLAYBACK]

-

[INAUDIBLE]

[END PLAYBACK]

DENNIS Now, you could probably hear the difference.

FREEMAN:

[LAUGHTER]

So now, check yourself. How does the phase of the second signal compare to the phase of the first signal? So look at each other so you can blame each other. If you work on this in solitary, then I'll have to blame you.

[INTERPOSING VOICES]

DENNIS So it's a phase relationship between the x signal and the y signal. If I showed you the Fourier
FREEMAN: transform for the x signal and looked at the y signal, which is x of minus n, how would the two look--
the same or different? I sort of gave it away. Flipping in time does what to the magnitude of the
Fourier transform? Flipping time--

AUDIENCE: Change the magnitude?

DENNIS Doesn't change the magnitude at all. What does it do?

FREEMAN:

AUDIENCE: The phase.

DENNIS Flips the sign of the phase. That's exactly right. So if you're into math-y these sort of things, you can

FREEMAN: think of representing the two as Fourier series-- the a_k 's and the b_k 's.

And if you think about what happens if you have x of minus n , you can do a change of variable, make that look like x of n . Then it looks more like the Fourier series equation, except that now I lost my minus sign. So a_{-k} is the same as b_k .

The effect is to flip the sign of the phase. Why would that be so much more audible than the other things that I showed you? Flipping the sign of a phase, why would that be audible? What would it mean if I flipped the sign from 68 degrees to minus 68 degrees? Not a clue. Yes.

AUDIENCE: For a total time, it wouldn't change a thing. But for any sign terms, it would invert that over that.

DENNIS So it'd invert-- Let's see, if you went from plus 68 to minus 68-- so say I had cosine of ωn plus

FREEMAN: 68 degrees and cosine ωn minus 68 degrees. So I change the phase by flipping the sign of it. How would that change the relationship between the two signals?

What's phase? Intuitively, what's phase? Yeah, it's this way. It's a time shift. So the problem with this phase shift is that you're shifting all the different components by different amounts of phase. So you might have a little bit of 68 hertz and 72 hertz and 3 kilohertz and 512 hertz. And all of those are getting their phase flipped. So whatever it came out-- 13, 128, whatever-- it becomes minus that.

All of those phase shifts represent a time shift. All the components are being shifted in time. So you're shifting around things, all the different components, by some amount. And what your ear is actually very sensitive to is the low frequencies. When you shift the low frequencies by an amount of phase, that's a big time.

So if you think about the fundamental frequency that started when the utterance began and went to the end of the utterance, that was like two seconds. So whatever that phase is, it got inverted, which is a large fraction of two seconds probably. You can easily hear a time-shift of two seconds. Yeah.

AUDIENCE: Could you notice it, if it was just a pure tone?

DENNIS If it were just a pure tone, you would have no clue.

FREEMAN:

AUDIENCE: Or you can [INAUDIBLE] overtones very simple.

DENNIS So to make Ohm's law a little more precise, what your ear really does is it does band pass filters. If

FREEMAN: the components fall very close to each other, you can hear the phase shift. If the components are far

from each other, you can't.

Even the second overtone of a fundamental is far enough away that you can't hear the relative phase of it. The frequency components have to be very close together. So you can hear phase only in a very particular case when there's frequency components that are very close to each other and there's a relatively big shift in the phase of those two signals. Except for that, Ohm's law is true.

The point is just that the Fourier series gives us a way of thinking about a very interesting dimension of a signal-- one that separates things that are audible from things that are not audible. So when you're doing signal processing, you can pay less attention to the thing that's not audible. You can go ahead and distort that, and nobody will know. But you have to pay close attention to the things that's very audible, which is the magnitude. We'll see at the end of the hour that that situation is completely flipped for video.

But before we get there-- so the first thing to think about was the idea of circular frequency-- the fact that we're looking at frequency responses living on the unit circle, rather than on the $j\omega$ -axis. The second big difference with the DT is the frequencies are cyclic. So rather than having harmonics go the whole way up to infinity, as we did in CT-- In CT, we would have a Fourier series constructed of the fundamental, the second harmonic, the third harmonic, the fourth harmonic, the fifth harmonic, the whole way out to infinity, in principle at least.

In DT, they rap, so that you get a finite number of them. In fact, if you're taking a sequence of length N , here illustrated for 8, you get eight possible frequencies, half of which are negative. And you don't care about those anyway.

So the number of frequency components for a DT signal is finite-- for DT Fourier series. That has enormous implications. It means that when you do a Fourier series decomposition, there's a finite number of terms. That's enormous, because that means that, if you started out with a sequence of length N , say 1,024, and you take the Fourier series of it in order to think about the frequency content, you can replace the 1,024 numbers in time with 1,024 numbers in frequency.

There's no-- well, no is exaggeration. There is a factor of two kind of explosion of data, because I went from N real-valued signals, samples, to N complex-valued signals. That's not quite true either, because of the reasons about how information works. The-- how do I want to say it? The resolution that you need for the frequency domain representation is actually slightly smaller. To a first order, you're replacing $[? N ?]$ with $[? n ?]$

That's really good, because if you think about the alternative in CT, you would be replacing a signal with an infinite number. And that's just not reasonable, if you want to think about processing signals with digital electronics. There's no way that you could process an infinite number of signals in a finite amount of time. So the fact that you're replacing n samples in time with n samples in frequency really lends itself to doing signal processing using digital techniques.

Now, the problem is-- and just one way of thinking about the finite length is that you can think about it by matrix. That's just a mnemonic that helps me a lot, to think about how, if I have a four-length sequence, you go through some matrix. And what comes out is a four-length frequency response. You have some eight-length sequence, you go through some matrix. And then comes out some eight-length frequency representation.

And there's actually a lot of intuition that you can draw in making the parallels between matrix representations and the Fourier representation, because it is a transformation, a linear transformation, just like the matrix implies. There is one problem with that finite representation idea. And that is that it scales poorly.

If you think about having a two-length sequence, the two-length sequence has a two-length representation in frequency via the Fourier series. And that representation has how many coefficients? Four-- two and two. Now, if I do of four, I get four, how many coefficients? 4 by 4 is 16.

Now, some of these are 1's. Right So all of these are 1's. As you make the sequence bigger and bigger and bigger, however, the number of trivial entries gets smaller and smaller and smaller relative.

So as you go from a two-length sequence to a four-length sequence to an eight-length sequence, now we're up to 64 multiplies in order to convert a time-domain representation into a frequency-domain representation. Well, that's bad, because the kinds of things we want to do are not two and four and eight. The kinds of things we want to do are manipulations on signals of 10 to the sixth, 10 to the eighth, 10 to the 12th.

For my doctoral thesis, I did some fluid dynamic problems, where the calculations had 10 to the sixth unknowns. And I used frequency domain techniques, which I was able to implement with things that I'll talk about in a minute. If we had used this scaling by the square, it just wouldn't have worked. If you have 10 to the sixth unknowns and if you're using a method that converts from one representation to another that involves 10 to the 12th operations, that's just not a good idea. So that's bad.

So by the theory of lectures, what I'm about to tell you is--

AUDIENCE: [INAUDIBLE].

DENNIS There's is a great way to fix this. If there weren't, I wouldn't tell you. I wouldn't pose the problem. So
FREEMAN: the question is how to think about this scaling in order to not have the number of calculations that you need to do blow up with N^2 .

So in order to motivate this, I've used a funny notation, where the W represents the N -th primary root of 1. That just saves me from writing out $e^{j 2 \pi / N}$ 64 times, which when I did it, it didn't fit on one slide. I had to compress it too much, and you couldn't see it.

So just to save space, every time I say W_8^0 , all I mean is compute the eighth root of 1 and then raise it to the 0-th power. Compute the eighth root of 1 and then raise it to the seventh power. That's all that means.

So the idea is, is there a way that we can think about the structure of this matrix more simply? And in particular, is there a way to use divide and conquer? That's our best algorithm. It would be great if we could take the eight-point sequence and substitute in its place finding the frequency representations for two four-point sequences. And by the theory of lecture's, that's what I'm about to do.

So the idea is, what if I just simply took the even entries in this more complicated problem-- so I have eight entries, none of which are necessarily 0. So I have x_0 through x_7 . And I want to find the frequency representation c_0 through c_7 .

What if I just did the evens and the odds separately? That'd be great. In order to do eight in one fell swoop, I need 64 multiplications. In order to do each of these problems, I have 16 for each. That's only 32, I have a big win. But now, I'm left with, is there a way to relate these things-- these a_k 's and b_k 's back to the original c_k ?

So the idea then is, how would you think about the contribution of the evens, as opposed to the odds in the overall calculation? So think about the even-numbered entries in the signal, think about the transformation if you were to do a frequency representation for that via the Fourier series, so that you get a_0 through a_3 . The first thing you see is that these numbers have W^4 , the fourth root. I need W^8 , the eighth root.

But there's a simple identity. If I know the eighth root, I can just square it to get the fourth root. So I

can rewrite the fourth root guys in terms of equivalent eighth root guys. That was trivial. Nod your head, completely trivial.

And then I can say, well, OK here's my problem that I would like to solve. What if I was only thinking about the contribution of the evens? If I'm only worried about the contributions of the evens, then I don't need to worry about some of the columns, because they don't contribute to the evens. That's worth thinking about things in terms of the matrix helps us. We can visualize operations as operations on a matrix.

So now, if I'm only worried about the evens, I only need to worry about four of these columns. And if I squint at this, I can see that $W_8, 0, 0, 0$, that's the same row as up here. So this is the same as that. And this one, W_8 -- so those are all W_8 's, I don't even to say that-- $0, 2, 4, 6; 0, 2, 4, 6; 0, 4, 0, 4; 0, 4, 0, 4$. These numbers are just the eight coefficients.

Furthermore, it repeats. $0, 0, 0, 0$ is the same as $0, 0, 0, 0$. So that's $8, 0$, and that's $8, 0$. What I just found was that the way the evens effect the answer is precisely by taking their Fourier representation and repeating it twice in time.

Well, that's pretty cute. So therefore, the contribution of the four-length sequences to the eight-length sequences is showed this way. Stack up the frequency representations for the four, and that is their contribution to the eight.

The same sort of thing happens when I do the odds. Now, if I only look at the odds, I throw away a different set of four columns. And now, if I try to compare that to what I have up here for the odds, I don't quite get something as easy.

So if I look down this column, they should have all been 0. But if I look down this column, it's $0, 1, 2, 3, 4, 5, 6, 7$. If I look down this column-- $0, 2, 4; 0, 3, 6$ -- I'm not quite getting the right answers. But it turns out that, if I simply factor this number out, so that I make the first coefficient right, all of the other coefficients are right too.

So I got an easy rule now for how the Fourier coefficients for the odds contribute to the answer for the eight-length sequence. And now, all I need to do is glue those two together. So the way I can calculate the eight-length sequence is to calculate the a's and b's from four-length sequences and then paste them together this way. And I'll get the eight-length answer.

Now, pasting them together takes some multiplies. So it took me 16 to compute the a_0 to a_3 . It took me 16 multiplies to figure out b_0 to b_3 . So I'm up to 32. Then I have to do eight more to make the b's

look right. So if I just count multiplies, I'm now doing 40.

But that's a win, because that's compared to 64. And I get precisely the same answer. That algorithm-- that's an algorithm; that's a rule-based system for how you take one set of eight numbers and turn it into a different set of eight numbers-- that algorithm is called the FFT, the Fast Fourier Transform. That's mostly to confuse you, because it's not a transform, but it is fast.

So it's really a fast Fourier series. And in fact, the guys who figured it out knew that. And I don't know why we insist on trying to confuse you. That's perhaps our job. But the FFT, fast Fourier transform, is really the fast Fourier series.

And the idea is that it's a divide and conquer thing. So if you think about extending it-- So I motivated it by thinking about how to calculate the 8 by 8 from two 4 by 4's. But let's start from 0 and say, let's like capital M represent however multiplies it took to get an endpoint FFT.

How many multiplies does it take to do a 1-point FFT? If my signal is periodic in 1-- so what's the relationship between the series representation and the time representation if the length is capital N equals 1?

The relationship is identity. If my sequence was 17, 17, 17, 17, 17, 17, 17, 17, 17, 17, which is a sequence that has a period of 1, then the series representation is 17, 17, 17, 17-- not all that surprising. So it takes exactly 0 multiplies to calculate that one.

How many multiplies does it take to calculate a length 2? Well, you do two length 1 problems. And then because it's length 2, you have to do those two extra multiplies. Then take a length 2 and turn it into a 4, how many multiplies for 4? Well, you have to do two length 2, but now it's length 4, so there's 4 gluing multiplies.

To do 32, you have to do two length 16. And then there's 32 gluing. When you're all done, it takes eight N times log N. And the point is that N times log N can be a lot smaller than N squared.

So let's go back to the motivational example. The idea was, what if I had about 10 to the sixth unknowns? So that's about 2 to the 20th?

Yeah, that's right. So 2 to the 10th is 10 cubed roughly. 1,024 is 1,000 in CS terms. So I have about 2 to the 20th. So what would be the savings if I did the FFT rather than the direct form N squared?

AUDIENCE: That would be 2 to the 20th divided by [INAUDIBLE].

DENNIS So I need to take 2 to 20th times 2 to the 20th would be the N squared algorithm. And the other
FREEMAN: algorithm would be 2 to the 20th--

AUDIENCE: By 20.

DENNIS --divided by 20. It's enormous difference, that's the point. So just a little bit of historic context. That
FREEMAN: algorithm has been repeatedly invented in time. So we usually cite Cooley and Tukey. They wrote a classic paper in 1965.

The title of their paper was "An Algorithm for the Machine Calculation of Complex Fourier Series."
They knew it was series. That was 1965. That was when computers were starting to get big.

They actually went on and solved some very interesting problems, including neuroscience. They were the first people who solved Hodgkin-Huxley's model for a propagated action potential and showed that the equations that Hodgkin and Huxley had found for biology actually work. So that was 1965.

However, before them in 1942, Danielson and Lanczos had independently discovered it. In 1903, Runge had discovered it. And, in 1805, Gauss described it.

Now, there's a good reason why all these brilliant scientists-- so Gauss noted it and even used it, but he commented on how useless it was. Why would Gauss think it was useless? He was kind of smart.

AUDIENCE: Was Fourier alive then?

DENNIS So Gauss actually died before he published this. But it got published posthumously before Fourier
FREEMAN: started to work. So this came before Fourier.

AUDIENCE: [INAUDIBLE].

DENNIS So why did Gauss think this was useless? And that's 1805. And now, we come around to 1965, and
FREEMAN: Cooley and Tukey make this huge deal. And they're world-famous because of their rediscovery of something that Gauss knew in 1905.

AUDIENCE: I mean, you're not going to be doing anything that's like 128 by 128 by hand anyway.

DENNIS That's exactly right. So when Gauss was thinking about it, he was thinking about 16 by 16. And so,
FREEMAN: yeah, you save two, big deal.

He was much more actually impressed with logarithms, because he was doing the multiplies by hand.

And the log made it much easier to do multiplies, because you could substitute addition. So he actually commented on how useful logarithms are. And he didn't comment on how useful this is.

But now, we make a big deal in the other order. And it's entirely digital computation. So when we solve systems of 10 to the sixth or 10 to the 10th unknowns, this is an enormous deal. And logarithms are good too.

[LAUGHTER]

But this is an enormous deal. So one more piece of theory and then one more ending thing. The piece of theory is going from periodic signals to aperiodic. And not surprisingly, that works just the way it did in CT. In fact, there's nothing new here.

We can think about how would you make a Fourier transform compared to a Fourier series. With the CT, we developed the theory for series first, because it was easier. And then we derived the theory for aperiodic signals. By something called periodic extension, we took an aperiodic signal, like this CT signal, made it periodic, like that DT signal, and then took the limit as the period goes to infinity.

Had we stayed in time domain the whole time, that would be a trivial no-op. But the trick was that after we did the periodic extension, we then took the series of this and then passed the limit onto the series. So that's the trick.

So we think about generalizing the aperiodic into periodic. Then we take the series. Then we pass the limit. And we ask whether the series converges. And just like CT, the answer for DT will converge.

So you start with a periodic signal. You write out the DT Fourier series just as we've been doing. And you get an answer that for this square pulse looks very much like the answer we got before.

The tricky thing that I've done here is I've written this series on two axes, just like CT-- a k -axis, which is the harmonic number, and an ω -axis, which is the frequency. The frequency represents the position around the unit circle. So capital ω is the angle of the unit circle.

Then if I change the period, the effect of making the period bigger is to put the harmonics closer. But in a magical way, because the envelope-- the thing showed by the thin black line-- doesn't change as I take the limit, that's what enables me to think about there being a limit. As you make the period bigger and bigger, you get more and more lines in an envelope that doesn't change.

And then you can think about the envelope as being a function of the continuous variable ω ,

rather than the discrete variable k . And then you can think about, what would happen if you pass the limit? And the idea is that, if you do the synthesis equation, there's an a_k that depends on $1/n$, which you can write in terms of the fundamental frequency ω_0 . And then as you pass the limit, ω_0 is the spacing between these harmonics.

So if you multiply by the spacing, you get the area of that block. So the area associated with N equals 0 is this area. The area associated with N equals 1 is that one, N equals 2, N equals 3, et cetera. And so by doing the sum, the sum increasingly converges to the integral.

So you get precisely the same kinds of formulas we've been using before. The one thing that's a little bit weird, in the DT Fourier series, we had a sum in the synthesis equation as a sum in the analysis equation. Here, we get one of both.

The DT Fourier transform-- that's this thing-- is a function of the continuous variable-- uh-- discrete-time Fourier transform has a continuous frequency variable capital ω . Because it has a continuous-- even though it's a discrete-time Fourier transform-- discrete-time Fourier transform has a continuous frequency variable ω . So when we do the synthesis, we have to integrate. So one direction is a sum, the other direction it's an integral. Other than that, it's just the same.

So there's a striking parallel. All of these things are highly related. The CT Fourier transforms are highly related to the CT Fourier series, as is the CT transform to the CT series. And there's lots of connections. And we'll talk about that more next week-- how you can infer how one of these transform relationships would have worked from all of the others.

For the time being, the important thing is just that, when you do the DT Fourier transform, you can think about that just like the CT Fourier transform. The CT Fourier transform was evaluate the Laplace transform on the $j\omega$ -axis. Now, you evaluate the z -transform on the unit circle. The unit circle is no surprise. That's what we've been doing.

Because of the relationship between the DT Fourier transform and the z -transform, the DT Fourier transform inherits all the nice properties of the z -transform. We'll make a bigger deal of that in the upcoming two weeks. In fact, that's the next two weeks of this course, is to figure out the implications of all of those properties.

But for today, what I want to just close with is one more example of how you can think about these frequency representations helping you to imagine the kinds of signal processing that you do or would like to do. And so the last example now is one from image processing, instead of audio processing. I

started with audio processing. Now, I want to think about image processing.

Image has the property that I'm going to have to do something with 2D. And in this class, we primarily think about 1D. But you saw in the zebra problem that sometimes extending to 2D is not hard. And it certainly is the case that the Fourier method extending to 2D is really easy, or at least I'm going to say that.

So the way you can think about a Fourier series in two dimensions is illustrated by this picture. I took a picture that 512 by 512. That's the boat. And I made a new picture by taking this row, calculating the Fourier series for that row-- that had 512; so it's 512 by 512. So I took these 512 numbers, took the Fourier series, and wrote them here.

Then I took these 512 numbers and put the Fourier series here and repeat-- so fill the whole thing up. So I took 512, 512-length sequences. Then I took that string of numbers, took the Fourier series, and put it there. And then that one, take the Fourier series, put it there. Take that one, take the Fourier series, put it there, and build up an image that way.

So I, row by row, took all the Fourier series, pasted them together. Column by column, took all the Fourier series, pasted them together. When you're done with that, you get what we call a two dimensional Fourier series.

And what you see on the right then-- so here's my original picture 512 by 512. Here's the picture of the magnitude of the angle. You can clearly see the boat.

[LAUGHTER]

And it's much clearer here. So the phase looks kind of random. So to gain some insight into what's in the magnitude and what's in the phase, what I can do is generate a new picture by setting the angle everywhere to 0.

So I took the picture, the magnitude, and I did this to it. I got the magnitude and angle. I set the angles to 0. And that's the picture I got.

So I was able to munge the angle in the auditory thing, and you could barely tell. Can you tell that I've munged the angle? There's some distortion caused by the angles. There are some things that you can still see.

So if you go back to this one, there's these things here, and there's these things here, and there's--

And if you go to this one, there's some lines here, and some lines-- Maybe there's something there. But to a first order, it kind of messed things up.

So then I did the opposite-- just throw away the magnitude and substitute 1, put 1 everywhere there was a magnitude, but put the angle back. Now, you can almost see the ship. So throw away the magnitude is better than throw away the phase. That's kind of cute. Can somebody think of why throw away the magnitude might be better than throw away the angle when you're trying to look at a picture? Yeah.

AUDIENCE: Angle determines resolution. Magnitude determines [INAUDIBLE].

DENNIS That's exactly right. So if you muck around with the phase, you're changing the position. Even in

FREEMAN: audio, when you muck around with the phase, you're changing when the component comes out-- earlier or later. In the picture, you're changing when the brightness comes out-- earlier or later.

If you want to have a nice straight line, you better keep all the phases lined up. If you scramble them, they end up scrambled, and you don't see them anymore. And as you said, the magnitude is controlling the intensity. So my intensities are all screwed up, but you can kind of still see the lines. You can almost even read-- if you know Polish-- you can almost read what the title says.

So one more manipulation-- I took a different picture and substituted the magnitude of that different picture in place of the magnitude of the boat picture. So angle with the boat, magnitude of the mystery picture, and it came out pretty good. But by the theory of lectures, of course, you know that the way I did this was to get a picture of a different boat, so that it would be very similar. And so that's what the magnitude was.

[LAUGHTER]

So it doesn't really matter much. So I started with this magnitude and angle, substitute the mandrill's magnitude, and you get that. So it comes out pretty good. So the idea then is just that we've looked at Fourier representations as an alternative way of thinking about the content of a signal. And for many purposes, like audio and video, that alternative representation is easy and insightful.

No class tomorrow. It's Veterans Day. Party next Wednesday.