DENNIS FREEMAN: Hello, and welcome. So before starting today, there's one of those kind of required announcements for events coming up. So even though it seems like we just did exam 1, next week is exam 2. So next Wednesday evening, 7:30 to 9:30. Same as before other than it's Walker. So don't go to Building 26. Go to Walker. OK.

Other than that, the rules are pretty much the same. No recitations on the day of the exam. Coverage will be up until the end of this week. That includes Lecture 12, Recitation 12. It includes Homework 7, but Homework 7 will be handled the way Homework 4 was. So Homework 7 won't be collected. It won't be graded. There will be solutions posted.

We'll post previous exams. You'll be allowed to use two pages of notes, presumably the page you used last time and one more. Although, we are not going to check whether it was the page you used last time. This is supposed to be a convenience.

No calculators. No electronic devices. Just like last time-- I hope you all found it to be this way-- the exam was designed to be done in one hour. And you have two hours, so there's not supposed to be any time pressure.

Review sessions during the normal open office hours. And conflicts, please let me know so that I can arrange to get somebody to proctor. And so that I can get a room. So for those two reasons, try to tell me about conflicts before Friday.

I guess the other important thing is remember that the theory of the exams in this class is that they ramp up. The first exam only counted 10%, with the idea that that's to let you get acclimated to things. You're sort of walking in the shallow end of the pool is the idea. This one will count 15%. So it's coming up. The next one will count 20. And the final one will count 40. So this one counts slightly more than the last one, but you're still kind of in the not-so-deep end of the pool.

Questions? Comments? Questions on the exams? Good.

So I want to start talking about a new topic today, which is really an elaboration of what we talked about with frequency response. But I want to start by telling you the big picture, the sort of 30,000-foot view.

We've been focusing in this class on multiple representations for thinking about how linear systems work. Differential equations, block diagrams, impulse responses, frequency response-- all kinds of different ways of thinking about it, with the idea that if you know all of them, then you can use the one that's the easiest to solve your particular problem.

In almost all of the cases, we've tried to find representations that are small in conceptual complexity. So for example, differential equations. In principle, differential equations could be complicated. In fact, we're only interested in linear time invariant differential equations, which means that they always have a simple form. They always look something like y plus some coefficient times y dot plus some other coefficient times y dot. y double dot plus a whole bunch of things turns into say b0 x plus b1 x dot plus b2 x dot-- double dot plus a whole bunch of things.

Point being that although the differential equations could, in fact, be arbitrarily complicated, we only focus on the ones that have a simple representation in terms of a small number of constants. There they are-- a1, a2, a3, b0, b1, b2. Small number of constants. So that's the way of managing the complexity. There's only a small number of constants you need to worry about.

Similarly, when we think about things like poles and zeros, the whole point of a pole-zero diagram is that you reduce. So say we're doing s-plane. So say we're doing some continuous kind of a problem. We reduced the whole system to thinking about a small number of singularities.

Again, we're trying to think about the system in terms of a small number of numbers. We like this one because the numbers are in some sense more intuitive or help us to solve problems better than perhaps these numbers did. Both of these representations are characterized by a small number of numbers. But for certain purposes, one set of numbers is easier to work with than another set of numbers.

Last week, we started to think about frequency responses. Frequency responses were really good for thinking about systems, like audio. I did some examples with audio. Or just other kinds of systems. It's very good to think about this kind of a system as a frequency response because there's a certain frequency at which a very small input-- you can barely see my hand moving. Or, if I weren't shaking, it wouldn't move so much. And in fact, the mass moves a lot. So that's some unique property of a particular frequency that's interesting for us to know about.

The problem with frequency responses is that when we start to think about a frequency response, we're now thinking about, say, the magnitude of a system function as a function of omega, which might be horribly complicated, and an angle which might be horribly complicated. So we've kind of lost the modularity. We've kind of lost the ability to think about the whole frequency response by a handful of numbers. That's what Bode plots are. Bode plots are a way of thinking about the entire frequency response in terms of a handful of numbers. That's why we like it.

We like frequency response because it's a very good way to think about certain systems. Audio because we like to think about bass differently from the way we think about treble. Mass-spring dashpots because there are certain frequencies at which the system goes crazy. Airplanes because you don't want it to do that, right? So there are certain reasons why frequency responses are good. And Bode plots are a way of getting back to this idea that we can think about a whole function of frequency with just a handful of numbers. That's what we're doing. OK.

So just to get things going, remember where we're coming from. We think about frequency responses by thinking about eigenfunctions and eigenvalues. We think about how if your system is linear and time invariant, last week we showed that complex exponentials are eigenfunctions for all such systems.

You put in any complex exponential and what comes out is a weighted version of the same complex exponential. Property of linear time invariant systems Furthermore, the eigenvalues are really easy to find from vector diagrams.

The eigenvalue is the value of the system function evaluated at the frequency of the complex exponential. Really easy. And you can calculate that from a vector diagram, which capitalizes on the small number of numbers representation of the pole-zero plot. So it's all easy. We only do easy things.

And then, we're really interested in these sinusoidal responses. So we think about a sinusoid by using Euler's formula as the sum of two complex exponentials. And that lets us see that all we need to do is evaluate the system function at j omega 0, j times the frequency of interest, the magnitude, and phase. And that lets us compute, how much is the amplification at that frequency and how much is the time delay associated with that frequency?

And so that motivated us to look at, how does frequency response map to pole-zero? There's a simple mapping, the vector diagram. Think about for example, if you have an isolated zero, a single zero in your system. A Single zero here at minus 2. The magnitude can be determined by the magnitude of the-- by the length of the angle that goes from the 0 to the frequency of interest. Frequencies are on the j omega axis.

Angle can be found-- this angle can be computed from the angle that this vector makes with the x-axis. And that always works. Here it's illustrated for a single zero.

If you think about the magnitude and angle as a function of frequency, you map out features that look like so. As the frequency goes from 0 to very big, the arrow goes from short to very long. That's manifested by this. And the angle changes systematically from being near 0 to being up around pi over 2. OK.

And the same sort of thing works for a pole, except now a pole is in the bottom. So when it used to get big, now it gets small. The angle, which used to go positive, now goes negative. But it's the same idea.

And you can compose more elaborate functions by thinking about the individual arrows that correspond to factors. That's using the factor theorem to reduce the system function, which is always a polynomial in s, into factors. OK. So that was what we talked about last time

what I want to talk about today now is how to reduce the frequency response, which looks like a lot of numbers-- one for every frequency-- to thinking about it in terms of a small number of numbers. So the idea is that this magnitude response is simple if you look at it the right way. In particular, it has a simple value at low frequencies. If we think about the smallest frequencies you can have, that's omega equals 0. And the smallest frequencies, the magnitude response asymptotes becomes arbitrarily close to the line that you would get by substituting omega equals 0 here.

If omega were 0, then h of j omega would be minus z1. The magnitude of minus z1 is z1. So at very low frequencies, the response-- so when you get to frequencies near 0, the response asymptotes to that line. And if you get to very high frequencies, then the fact that you're adding a small constant to a number that's already very big, if omega were large, then that little constant, z1, wouldn't matter. So the magnitude becomes very close to the magnitude of omega shown by the green dotted line. So you can think about the blue line being some kind of an interpolation between those two dash lines, the low-frequency behavior to the high-frequency behavior.

And the relation looks even simpler if you plot it on a different kind of axis. If you make it log-log, then you can think about frequencies going-- so the log of frequency. Well, there's some-- if the frequency were 1, whatever that means. Or in fact here, it's important to note that I'm plotting here versus a normalized version of frequency. I'm normalizing by z1. What That has the effect that when omega is the same as z1, I get the log of 1, which is?

AUDIENCE: 0.

DENNIS FREEMAN: 0, right? That's a way of shifting the axis to make the interesting stuff happen at 0. That's all. So if you scale frequency by the singularity-- in this case, a 0-- then the critical frequency becomes the frequency zero. We all like 0, right?

So what you can see then if you plot it this way-- and I also scaled the magnitude the same way. So if you realize that there's a simple asymptote, you can see very clearly on this log-log plot that you get a good approximation to the function, which is showed in blue. That one's calculated exactly. And you get a good approximation by the low- and high-frequency asymptotes.

So instead of thinking about the whole function of frequency, what we do is we think about just the asymptotes. That's the simplification. We don't think about this complicated function. We only think about what happens at low frequencies and high frequencies.

Of course, this was too simple, right? This is what happens with a single zero. The point is that a similar thing happens with a single pole or when you combine them. Let's look at the case of a pole.

Here, the pole idea is a little bit more complicated. If you go to very low frequencies, think about s being j omega. The frequency response lives on the j omega axis. Always look at s equals j omega if you're interested in the frequency response.

If you think about s equals j omega, then if you go to small omega, this term goes away. And we're left with, again, a constant at low frequency showed by the red line.

If we look at high frequencies, now the curve's a little more funky. The curve at high frequencies-- the real answer is showed by this sort of bell-shaped curve. If you plot for different values of omega where s is j omega, plot out the magnitude, you get this sort of funky, bell-shaped curve.

The point is that-- and the high-frequency asymptote. If you just say, what if j omega were so big that the p1 wouldn't matter? Then, you would get 9 over omega being the asymptote. So that's this hyperbolic sort of thing here. OK. Well, that's ugly.

But if you go to log-log, it's very simple. That's why we like log-log.

If you think about the function y' equals 1 over omega and plot that function on log-log axes, reciprocal turns into a straight line. We like straight lines. They're easy. So what happens then-- again, I've normalized things. I've normalized frequency by the pole. That makes the critical frequency be omega equals pole. The critical frequency comes out 1. The log of 1 is 0. So the interesting behavior happens at log equals 0.

Below that frequency, the frequency response magnitude is well approximated by the constant whose log divided by some funny number that was put there for normalization is 0. And at high frequencies, it falls off with a slope of minus 1.

So again, for the case of the pole, we get this simple behavior if we focus on the asymptotes. OK. So now, that's kind of the theory behind everything. Make sure you're all up to speed on the theory.

So compare log-log plots of the frequency response magnitudes of the following system functions. H1, 1 over s plus 1. Where's the pole?

AUDIENCE: [INAUDIBLE]

DENNIS FREEMAN: Minus 1. And s2, 1 over s plus 10. Pole is? Minus 10.

Compare the magnitude. Compare their magnitude functions when plotted on log-log and answer the following question, which of these best describes the transformation between H1 and H2? You should shift horizontally. You should shift and scale horizontally. You should shift horizontally and vertically. You should shift and scale horizontally and vertically. Or, it's something completely different.

OK, turn to your neighbor. Say hi. Right now, say hi. And now, figure out what answer best classifies that transformation.

[SIDE CONVERSATIONS]

DENNIS FREEMAN: OK. So which transformation do you like the best? Everybody raise your hand with the number of fingers between 1 and 5. Let me make sure I know the answer. I think I know the answer. OK.

Greater audience participation. Blame it on your neighbor. My stupid partner thought it was-- OK. It's about 90% correct.

If I plot the magnitude response for H1 on log-log axes, what will it look like? Sketch it in the air. Sketch it in the air is like that sort of thing. Go ahead. Go ahead. Sketch. OK. Maybe that didn't work.

OK. What I want to do is plot the magnitude of H1 of j omega log versus the log of omega. But I only want to think about the asymptotes. Where should I draw the low-frequency asymptote for that function? What's the low-frequency asymptote for the log magnitude of H1?

What's the low-frequency asymptote-- forget the log part. What's the low-frequency asymptote for H1? 0. Yes.

AUDIENCE: [INAUDIBLE]

DENNIS FREEMAN: I'm sorry. I can't hear.

AUDIENCE: Why is it not symmetric? So if you look at the amplitude of the system function, j omega, it's symmetric. But the Bode plots are not symmetric.

DENNIS FREEMAN: The Bode plots are unsymmetric. Why are the Bode plots unsymmetric and the frequency response was symmetric? Somebody--

AUDIENCE: I get why, if you have a negative frequency.

DENNIS FREEMAN: Negative frequency, exactly. Why do we have negative frequencies?

AUDIENCE: [INAUDIBLE]

DENNIS FREEMAN: Because time runs backwards. Negative frequency corresponds to running time backwards, right?

The argument, it gets bigger when time gets smaller. OK. That's not the reason we do it, right?

Why do we use negative frequencies?

AUDIENCE: [INAUDIBLE]

DENNIS FREEMAN: Why do we consider negative frequencies?

AUDIENCE: [INAUDIBLE]

DENNIS FREEMAN: To make it real. To make what real?

AUDIENCE: A system.

DENNIS FREEMAN: Why do we think about negative frequencies? Come on. Yeah.

AUDIENCE: [INAUDIBLE]

DENNIS FREEMAN: We want sines and cosines. It's easy to find complex exponentials. We like Euler, right? We like Euler. We like that cosine of omega t can be written as 1/2 e to the j omega t plus 1/2 e to the minus j omega t. That's what we like. Whoops. I shouldn't put the parentheses there. That's what we like.

We invent the negative frequencies. They are completely an invention of our own. When we do the Bode plots, we throw them away.

We know that we can always figure them out because it was an invention of ours, so we don't need to write them down. If we ever needed them, we could figure them out. So we don't bother to write them down on the Bode plot.

What's the low-frequency asymptote of H1? If omega 0 goes to 0, what's the value of H1? 1. s equals j omega, omega equals 0. 1.

What's the log of 1?

AUDIENCE: 0.

DENNIS FREEMAN: 0. So the low-frequency asymptote for H1 is 0. What's the high-frequency asymptote of H1? If you put a high frequency-

AUDIENCE: [INAUDIBLE]

DENNIS FREEMAN: We're thinking about H1 of j omega. H1 of j omega is 1 over j omega plus 1. What happens when you make omega high? What happens to the magnitude?

AUDIENCE: [INAUDIBLE]

DENNIS FREEMAN: 1 over omega for the amplitude, right? You could say it approaches 0. That's kind of right, but we're interested in how it approaches 0. On a log plot, 0 is very far away. So we never think about being there. We think about getting there. So we get there with a slope of minus 1. So it's like this.

What's the crossover point? Omega equals? What's the omega at the crossover point?

AUDIENCE: 0.

DENNIS FREEMAN: 0.

AUDIENCE: We're plotting for negative [INAUDIBLE]. And I would think that since we're taking the absolute value [INAUDIBLE] and not having it just be 0. Because you're talking about as omega goes to 0, but not necessarily as omega goes to negative.

DENNIS FREEMAN: We're thinking about the limit as we go towards 0, because we're plotting on the log of omega. 0, omega equals 0, is way over there. Way, way, way off the blackboard. So we don't think about that. We think about how it got there. And it got there by being constant.

The trend-- if you start at the critical frequency in the middle, the value becomes closer and closer to 1 as you go to the left. That's why we say the asymptote is at log equals 0. Log of 1 is 0. The answer goes to 1. The log of 1 is 0.

Does everybody know what I'm talking about? I'm getting a lot of blank stares.

Then, if you try to go to a high frequency and think about the magnitude, you can go to a high enough frequency where the magnitude of this doesn't matter that that's there. The magnitude never matters that there's a j there. So the magnitude goes like 1 over omega.

If you plot the function, 1 over omega, a reciprocal function on log coordinates, it turns into slope of minus 1. So that gives us this region where it has a slope of minus 1.

What's the frequency at which they cross?

AUDIENCE: [INAUDIBLE]. In log scale, it's only [INAUDIBLE].

DENNIS FREEMAN: On a log scale, it's omega equals 0. On a linear scale, it's omega equals 1. So the crossover point where 1 over omega-- the high-frequency asymptote is 1 ever omega. The low-frequency asymptote's 1. They are equal when 1 over omega equals 1. Omega equals 1, right?

Log 0. So they cross at a point where the log is 0.

Now, how about H2? What if I plot H2 on top of this? What's the low-frequency limit of H2?

AUDIENCE: [INAUDIBLE]

DENNIS FREEMAN: 1/10. What's the log of 1/10? Minus 1. So now, H2, if I do log of H2, it's going to-- so the low-frequency asymptote is 1/10. It's minus 1 on a log plot.

What's the high-frequency asymptote for H2?

AUDIENCE: [INAUDIBLE]

DENNIS FREEMAN: 1/s. 1 over omega. It's the same.

If you go to a high enough frequency, the fact that there is divide by 10-- the fact that you've added 10 or added 1 doesn't matter. If omega gets sufficiently high, the thing you added to it has no relevance. You get the same high-frequency asymptote. So to transform the asymptotic view of the magnitude function from H1 to H2, what should you do to the curves?

[INTERPOSING VOICES]

DENNIS FREEMAN: Shift and shift. No scaling. That's why we like Bode. That's why we like log-log. Every pole looks the same. Every pole is straight across and down. Every zero looks the same, straight across and up.

There's a critical frequency. That frequency is the frequency equal to the position of the pole or zero. That's all of the rules. OK. That's why we like this. Does that make sense?

So the answer here was that we want to be able to shift horizontally and vertically. We don't need to scale. The shape of the curve is invariant to the position of the pole and zero. That's what we like. It's an invariant. We like invariant things. OK.

Then, if you wanted to construct a more complicated system, it's also easy. If you want to construct a more complicated system, you can use the factor theorem to convert the system function, which for a system that's built out of integrators, adders, and gains. If you build a system out of that, the system function will always be a ratio of polynomials in s. We proved that a while back. You can always factor such things.

We just found out a rule for how you think about the individual factors. And so all you need to worry about is the rule for combination.

If you multiply a bunch of zeros and divide by the product of a bunch of poles, that's pretty easy. If you take the magnitude-- the magnitude of a product is the product of the magnitudes-- that's also pretty easy. Everybody's with me?

And if you take the log, it's even easier. The log of a product is the sum of the logs. OK. Sum's easier to multiply. That's another reason we like this log-log thing.

So all you need to do is think about each singularity. Each one of them can be represented by this and down or this and up. And compose them by adding. So it's really easy.

So say I had this system function, a 0 at 0, a pole at minus 1, and a pole at minus 10. Think about the Bode representation. Bode's just a word that means think about the asymptotes. Think about the Bode representation for the top.

Well, that grows linearly with omega. A function that grows linearly with omega has a log that grows with a slope of 1. So that's this. So that's the zero. That's the contribution to the magnitude function of the zero.

The contribution to the magnitude function of the pole at minus 1 is this. The pole at minus 10 is this. All pole's look the same. So all you need to do is add them up. So you add the first two, you get the-- so backing up.

0. Pole. Add them. In this region, the sum of a constant and a straight line sloping up is a straight line sloping up. In this region, the sum of a straight line sloping up with this sloping down is flat. So that's why you get that.

Then, we add in this contribution. Add a constant to this. It just shifts it up and down. Who cares? That's easy. The important thing is that it breaks down. So the result of summing that is that it breaks down. Easy.

So instead of thinking about a frequency response as something that's horribly complicated, we think about it as having parts that came from each pole and zero. So instead of thinking about it as a collection of arbitrary numbers at different frequencies, we think about it as a little part that comes from the first pole, another part that comes from the second pole, another part that comes from the first zero-- blah, blah, blah. It's a way of thinking about it, reducing the complexity, the conceptual complexity, of the frequency response. So the angles are the same.

If we think about the low-frequency and high-frequency behavior of the angle starting with a 0. For low frequencies, the angle is 0. So that's the dot. For high frequencies, as frequency goes higher and higher and higher, the angle stands straighter and straighter up. The asymptotic value is? Pi over 2. So we have two asymptotes, 0 and pi over 2.

If we plot that on log-log axes, we get something slightly more complicated. The blue line shows the calculated value. The red line shows a way of thinking about that from a straight line approximation point of view. Very nice construction is put a straight line starting one decade-- a factor of 10, minus 1 on a log scale. Draw a line from minus 1 to plus 1, a straight line. And you actually get a very good approximation to the angle function over the whole range. Notice that there are two critical frequencies associated with the phase. There was one critical frequency associated with magnitude.

What was the crossing point of the low- and high-frequency asymptotes? In phase, we think about that same thing, but then we bump it up and down 1. So there are two critical frequencies. The same critical frequency we use with the magnitude plus or minus 1.

The same sort of thing happens with a pole, but now it's upside down. Other than that, it's identical.

Same thing. Having calculated the phase for a single pole or zero, it's easy to think about how you would generalize to multiple poles and zeros. The angle of a product is the sum of the angles. Don't need to take the log this way.

This time, when we're doing the phase, we don't need to take the log. There's a way of thinking about phase. Because it's in the exponent, e to the j angle by Euler's equation. There's a way of thinking about the logs already in the angle function, right? The angle was in the exponent, e to the j omega. So since the angle is already up in the e, it's kind of like an-- it's already a logarithmic function. And that shows up here.

The angle of a product is the sum of the angles. So same sort of thing happens if we had a complicated function. Just think about the angle that results for each of the poles and zeros and add them.

The angle associated with the zero at 0 is pi over 2. j is the same as e to the j pi over 2. e to the j pi over 2. The angle's pi over 2. So the angle associated with this is always pi over 2.

The angle associated with the pole at minus 1. Poles in the left-half plane cause the angle to start out at 0 and go negative. So we start out at 0 and go negative.

We find the critical frequency labeled in this axis by 0. s equals 1 is the critical frequency. So we go up and down 1 unit and we draw the straight-line approximation.

s equals 10 is the same thing, except now it's shifted to a higher frequency. Factor of 10, units shift on a log plot.

And now all we do is sum them, add the first to the second, add the third. That's our angle approximation. Again, the idea is to take something that's conceptually hard-- what's the value of the angle function as a function of frequency-- and turn it into something simple. A few straight line segments associated with every pole and zero. So this is just a summary.

Because we can represent a system that's composed of integrators, summers, and gains by a linear differential equation with constant coefficients, it follows that the system function is a quotient of polynomials in s. Because of the fundamental theorem in algebra, there's n roots. Because of factor theorem, you can break it up into factors. Because of all that, we can think about them one at a time and glue them together.

Gluing them together in the case of the magnitude works best if you use the log because then the product turns into a sum. In the case of the angle, it's for free because the angle is, in some sense, already a logarithmic function. OK.

OK, we'll see if you got it. So here is a complicated frequency response. This is the straight-line approximation to the frequency response of a system. Which system?

So which of the system functions-- 1, 2, 3, 4, or none of the above-- is represented by the Bode straight-line plot showed above? Raise your hands, which one is better-- 1, 2, 3, 4, or 5? Ah, 100%. That's exactly the right answer. Wonderful.

The idea is that Bode is easy. That's why we do it. It's easy. So tell me a rule. How do I think about this one? What would happen here? What's this saying? What would be the Bode plot of this one? Sketch it in the air. Exactly.

You start at 0, right? OK. So start at 0. Then what? What do you run into first? You start at 0. What do you run into first when you're doing this guy?

You run into this factor, or that factor, or that factor first?

AUDIENCE: [INAUDIBLE]

DENNIS FREEMAN: Yeah, the left one. So think about the order. By how close were they to the origin? Think about them in order from the origin. So you hit the first one first, the one at 10, then the one at 100. This would start flat, break down at 1, break down again at 10, break down again at 100. That's not the right shape.

This one starts with a 0. That's a break up. The 0 happens at frequency 1. So that's log frequency 0. That's right. Then, you break down, break down. OK. So break up, break up, break down. Except that it would actually break down first. So the way to think about this one is break down, break up, break up ordered from going away from the origin. OK.

So the point is it's very easy to take a system function and immediately draw the frequency response. That lets you take a pole-zero representation, which can be very concise in terms of the number of numbers you need to know, and quickly map out the frequency response, which can be very intuitive for thinking about how the system responds. OK.

So one more issue. We don't really like log plots when it gets down to talking about the frequency at which the log was 7.129. We just don't like that because nobody measures frequency in the frequency whose log is x.

So what we usually do, instead of plotting versus the log of omega, we plot versus omega on a log scale. OK. Very reasonable. We just put tick marks and we just label them exponentially. So the tick mark associated with omega is labeled 1, not 0. That's frequency 1.

The tick mark associated with frequency 10 is labeled 10. What that means, though, is that the growth inside that interval is on a log scale. That's why we write log scale here to keep reminding ourselves that the ticks between 1 and 10-- 1, 2, 3, 4, 5-- are not uniformly spaced. They're log spaced.

The distance between 1 and 2 is bigger than the distance between 9 and 10. That's the way logs work. OK. So we think about frequency on a log scale rather than log frequency.

Similarly, we think about amplitude on a dB scale rather than log amplitude. dB is for Alexander Graham Bell. It's decibel. A bell is a factor of 10. It's a little unintuitive, a factor of 10 is labeled 20, because Bell was really thinking energy. And you have to square energy. You have to square voltage to get energy. Bell was thinking a factor of 10 in energy. We like to think amplitude. Therefore, a factor of 10 for us would be 20 decibels. 20. OK, so that's a little weird. So we will think about this axis in decibels and we'll think about that one in decades. And so what that means is that the slopes are no longer minus 1 and 1. The slopes are 20 decibels per decade, or minus 20 decibels per decade.

Now, this is all completely meaningless. It's a slope of 1. And you could equally have labeled this axis in dB and this axis on a log scale, and then you would have the slope of 1 decade per 20 dB. And that, too, would be 1. We just don't do it that way. OK.

What do we do? What we do do is frequency on a log scale. So therefore, the unit of frequency is the decade. You have a frequency and a decade higher and a decade higher and a decade higher. The unit of frequency is decade or octave. Octave is a factor of 2. Octave higher, factor of 2. Octave higher. Octave makes more sense if you're a musician, right?

The distance between two C's is an octave. So that's a dB scale. And that results in a little bit of funny math.

So if you convert linear measures of amplitude to decibel measures, a factor of 1 in amplitude is a factor of 0 dB. A factor of 10 is 20 dB. We already talked about that. There are some convenient middle grounds. 2 is 6 dB. That's a little weird. So we'll go around saying it's 6 dB and we mean 2. But every engineer in the world will call it 6 dB.

Half of that is the square root of 2, which is 3 dB. It's just convenient, at least it is after you spend 30 years doing that.

The other point from the slide is that the asymptotic responses are really quite good. The magnitude for a single pole deviates from the straight-line approximation by only 3 dB. I can hear sounds that range 120 dB. That's the sense in which 3 is small. The kinds of signals we work with every day have ranges that are big compared to 3. So we think of 3 as a small thing.

The phases are within 6 degrees. And for a lot of applications, 6 degrees is a small number. Again, we're thinking 6 out of 180. So it's a small fraction of a cycle. OK, this is a good thing-- especially the fact that it's a trick question-- for you to practice for the exam. So let me skip it in the interest of time, because there's one more important thing.

Don't look at the answer. Use this as practice. There's one more important thing in trying to reduce poles and zeros to a frequency response. And that is that when you use the fundamental theorem of algebra, even though the polynomial can have real-valued coefficients, that does not mean the roots are real. Right

The roots to a polynomial with real value coefficients can have complex parts. So the remaining thing I want to talk about is, what do you do with these poles that have complex parts?

The imaginary part has to do with oscillations. And so we want to think about, what's a Bode plot looks like? What's the asymptotes look like when you have a system like this? This was a system that has a mass spring, and dashpot. You should have done this by homework by now. The differential equation is second order. It has real-valued coefficients, but the holes are complex.

What happens with complex roots? Well, if the com-- if the polynomial had real-valued coefficients, complex roots come in pairs. That's the only way to take a complex number and end up with a product that's all real. You take the complex number. It has to be paired with its complex conjugate. So that when you pair them, the result has real coefficients.

So we only need to worry about the case when the poles or zeros come in complex pairs. And for that purpose, it's convenient to think about-- you might think that if you had mass-spring dashpot system, you might be expecting something like s squared m plus sb plus k, something like that. Mass-spring and dashpot. Three parameters.

Well, you don't really need three. Three is nice because it has an association with how stiff is this thing and how massive is that thing, et cetera. But in terms of thinking about poles and zeros, you don't need to think about all three.

First off, you could divide the top and bottom by k and the 1/k then is just a gain factor. Gains are easy. They don't affect shape. OK. We don't care about that one. So we went from 3 to 2.

Now, there's another simplification because all of these are going to be oscillatory. Pole pairs work in an oscillatory fashion. This has a natural frequency that if I didn't shake it, it has a preferred frequency.

If I divide by that preferred frequency, omega 0, I can get rid of another parameter. So don't think about s. Think about s over omega 0.

Now, frequencies are normalized to 1. OK. So by dividing by omega 0, I turn every frequency-dependent system into something whose best frequency is near omega equals 1.

Then finally, for my third parameter, if I write my third parameter as 1 over q, something very magical and wonderful happens. The roots fall on a circle.

So what I want to show here then is as I change q, I don't need to think about omega. I can plot this on an s over omega 0 plane and it works for all omega. I do need to worry about q. So if I change q, here q is 1/2. So there's q of 1/2. 1/4. 1/8. Excuse me, I'm doing it backwards.

I want 1/q to equal-- I want 1 over 2q to equal 1/2. I need q equals 1. OK, that's better. That's q equals 1. q equals 2 in order to half the distance to the origin-- 4, 8, 16. Notice that as I change-- now watch this side. So q equals 1. So q equals 1, q equals 1. Low-frequency magnitude is flat. High-frequency is slope of 2. Sloping down with minus 2.

As I change q from 1 to 2, 4, 8, 16, the peak gets bigger. In fact, if you do a little bit of math, you can show that the peak gets bigger with q. The peak value, if you measure how big is the peak compared to where is the crossover, that distance is a factor of q.

Similarly-- and you can reason about that with vector diagrams. And we'll do homework problems to practice that.

Similarly, it got peakier. It got sharper.

If you do the vector story to try to figure out why it got peakier, the width turns out to be 1/q. So that's kind of impulse-y. The height got bigger with q. The width got bigger with 1/q. The product always 1. So that's a way of thinking about why it got peaky that way.

And finally, if you think about the angle, the angle changes. As you make q bigger and bigger, the angle changes very quickly. And it turns out that the angle changes abruptly over the same bandwidth. Bandwidth is how many frequencies are there between the low-frequency part and the high-frequency part. The phase change over the bandwidth is always pi over 2. OK.

So that's the whole story then. Think about isolating the poles on the real axes. They're just this. Isolated zeros, they're just this. Pairs can be more complicated because they can be peaky. OK. Thanks.