

MITOCW | 12. Filters and composition

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

PROFESSOR: So we're going to continue talking about LTI systems and how to work with them. We're thinking of an LTI model for a channel, something with input x_n , output y_n . And we've seen that the output can be obtained from the input through this convolution operation, right? So for instance, this was one way to write it. And the shorthand notation here was h star x , evaluate to the time n . The m here, in general, goes from minus infinity to plus infinity.

In general, you put a unit sample function of the input, the response can extend from minus infinity to infinity if you've got a non-causal system. If you've got a causal system and you put a unit sample in, then the response starts from 0 and goes on to the future. But it's often useful to be able to represent and analyze non-causal systems. I mean, if you have all the data stored in your computer, then you can look forward and back from the time that you're at. And therefore, you can do things that can be looked at as non-causal or analyzed as non-causal.

Now, there's one issue that we've kind of swept under the rug, which is you can write nice looking formulas, but do they mean anything? So we went through a plausible derivation, but if you end up with a summation from minus infinity to infinity, then you've got to ask yourself when does this make sense? Because you know that adding an infinite number of things can cause problems. So you need conditions for all of this to make sense. For instance, if h was 1 for all time and x was 1 for all time, there's no way this is going to make sense. Because it's going to blow up at every value of n , right? So you clearly need conditions.

So what we're looking for is conditions for convolution to be well behaved. And I'm going to give you one important condition. Well actually, let me give you two. So here's one. Suppose we have a causal system and input that starts from 0 at time 0. When I say starts out, I mean that for all prior times the input is 0. So we've got an input that is 0 for all negative times. And then, at n equals 0 we start to get some action, OK?

So let's see, what happens to this infinite sum in this particular case? If I have a causal system, that means that h of n is 0 for n less than 0, right? That's what you should think of right away when you're told a system is causal. The unit sample response can only extend from 0 onwards. So if the unit sample response is 0 for a negative time, and the input starts at time 0, what simplifications can you make to that convolution representation? Somebody? Yeah.

AUDIENCE: [INAUDIBLE]

PROFESSOR: 0 to infinity does it? OK, so the answer was that instead of starting at minus infinity, you only need to start the summation at m equals 0. Because the h of m is going to be 0 for all negative values of m . So you can start that summation at 0. And then, where does the summation-- where can the summation end?

AUDIENCE: n ?

PROFESSOR: n ? Yeah, because this x here is 0 for negative values of the argument. Negative values of the argument are values of m greater than n . So this only needs to extend over a finite region. Well, if you have a finite sum, then you're happy. Nothing's going to blow up. So this is certainly one case where everything works fine.

There's another case which I'll describe which is where your input is bounded and what we mean by that is that your x_n has an absolute value that's less than or equal to some maximum value-- some finite maximum value for all time. So no matter what n you pick, you're always within this interval, OK? So your input is, here's your n , here's your plus m , here's your minus m . And your input is constrained for all time to just lie between these limits. So that's a bounded input.

And here's the other part of the condition. Absolute value of h_m summed over all m is finite. We say that h is absolutely summable-- absolutely summable. So what does that do for us? Well, it actually allows us to bound the outputs. So now, it turns out y of n , that's the absolute value of the output. Well, that's less than or equal to the absolute value of this convolution expression, which is in turn less than-- well, it's equal to that, right? It's equal to that, but it's less than or equal to this.

Absolute value of a sum is less than or equal to the sum of the absolute values. And that's less than or equal to capital M -- this is the max value that we allowed up there-- times summation h_m . So this whole thing is finite. So basically, you can bound the output. You can guarantee that the output is bounded provided these two conditions are satisfied. So this is actually a very important pair of conditions. It's one that we encounter all the time and in practice.

And because of this result, we actually refer to an LTI system that satisfies this absolute solvability condition. We say that the LTI system is bounded input, bounded output stable-- bounded input, bounded output stable. And I've run out of space from my stable there. So if somebody tells you they have a bounded input, bounded output stable system, if we're talking about an LTI system, what they mean is that that condition is satisfied. So that's an important condition.

OK, I have all that on a slide. But if I race it past you on a slide, then it's hard to track. But this is something that we want to do. So for instance, if I had an LTI system whose unit sample response was the following-- let's say it's 0 for all times up to and including time 0, and then it takes the value 1 over n , from then on, OK? So let's see, one way to write this is 1 over n u of n minus 1. If I tell you that the h of n is this, that automatically takes care of zeroing out everything from 0 backwards, and then putting in 1 over n from then on.

So is that BIBO stable system-- bounded input, bounded output stable? Yes? How many think yes? You'd like to think it's stable, because a unit sample response is decaying. But actually, it doesn't satisfy the absolute solvability condition. The sum of 1 over n from 1 to infinity is-- it actually blows up. It doesn't converge. If it falls off any faster than that, then you're in good shape. But this is actually bad.

If you had something that was 1 over n squared-- OK, so this is not BIBO stable. But if you had 1 over n squared u of n minus 1, it is BIBO stable. If you had, let's say, $1/3$ to the n u of n minus-- well, we'll say u of n , that's BIBO stable. So this falls off as 1 over n squared. That's fast enough for the sum to converge. This falls off exponentially faster. It's as a geometric series. It's a discrete time exponential. So that's fast enough. So that's also BIBO stable.

All right, so this is time domain. We know how to analyze any LTI system with this. You tell me what the unit sample response is and I can tell you what the output is for any given input. But this would be a nightmare if we had to do design with this, because convolution is not-- it's simple enough to implement for a particular case, but it's not a simple operation to think in terms of.

The reason is that the output at any one time is obtained by scrambling all the inputs for all time, combining them in this weighted linear fashion. And then, if you move to the next time step, you're again scrambling all the inputs, but with the weights shifted a little bit, so you've got to start from scratch again. So it's very hard to know what you can say in general using the time domain.

It's not clear at all how you'll do design, for instance, of filters to filter out noise and so on. So this is actually-- it's important. It's a full characterisation of LTI systems. But if we had to stop there, it's a fair bet that we wouldn't be anywhere near where we are for engineered systems, certainly not in digital communication.

So the key thing is actually to start thinking in terms of frequencies. So we're going to look-- we're going to spend the next several lectures looking at the frequency domain. And so what is the frequency domain? Well, we're going to be focusing, essentially, on sinusoidal inputs and inputs that are related to them. But let me actually just start back with something simpler.

Here's my question. And I hope the answer isn't already on the slides. Maybe it is. Is it true that if my input was periodic-- so for instance, if I had an input that, let's say, did this, it ramped up over maybe four time steps, and then started again, ramping up over 4 times steps, and continued that indefinitely. So here's my x of n . It's an x of n that has some basic period that then repeats periodically. So it satisfies that condition x_n equals x_{n+P} for some number P that's the period. What is P in this case? 4? OK. So every 4, this repeats.

Is it true that if I had a periodic input to an LTI system that I'll get a periodic output? Any reason I should expect that? Yeah?

AUDIENCE: Depends on the system.

PROFESSOR: Could you speak up a little?

AUDIENCE: It depends on the system.

PROFESSOR: It depends on the system. I'm telling you it's LTI but nothing more. So it depends on which particular LTI system? That would be my intuition, too. Yeah?

AUDIENCE: Just like in the piece that it can average over a long enough time sample, to where it would be constant. So I'm assuming it's constant.

PROFESSOR: Well, but that's an average. That's one number. What I'm asking you is could this output as a function of time also be periodic? Can you guarantee it? Yeah?

AUDIENCE: [INAUDIBLE]

PROFESSOR: I'm not exactly understanding your prescription. So are you telling me how to prove that y_n equals y_{n+P} plus something for all n ?

AUDIENCE: If we choose enough-- if we take enough x 's can we make y constant?

PROFESSOR: According to this-- well, I raised the convolution expression. But the convolution expression requires us to take x 's from minus infinity to plus infinity. In general, the output at time n depends on all the inputs. So it's not clear how you can block things off. You had another idea?

AUDIENCE: Well, because it's time invariant, [INAUDIBLE]

PROFESSOR: OK, good. This is a time invariant system, right? So if I shifted the input by capital P , the output should also get shifted by capital P . But shifting the input by capital P gives me the same input back again. So the response that I get for that shifted input must be the old response that I had, which tells me that the output also has the property that if I shifted by capital P , I'd get the same thing again. So this is guaranteed to be period capital P .

Now, actually there's a little twist to that. Because we usually think of the period as being the smallest interval for which you can repeat. It's conceivable that this output has a smaller interval, and therefore that the period is some integer fraction of this capital P . But this is the general idea. So you see that just knowing that the system is LTI, you can already tell a lot about what the response will look like.

One of our favorite periodic inputs-- well, if I asked you to tell me what your favorite periodic input is, what might it be? Any-- sorry?

AUDIENCE: Constant input?

PROFESSOR: Constant is good. That's periodic. A little less trivial than constant?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Sinusoid, right? So here's the nice thing about sinusoids. It turns out, for an LTI system, if I put in a sinusoid, not only is the output periodic with the same period. It's also a sinusoid. So that's an even greater restriction here. You see, in this particular case, I have an input that's periodic. I'm guaranteed the output is periodic with the same period. But the actual shape of the waveform can be all messed up relative to this. It may have no obvious visual relationship to this.

But if you have a sinusoidal input, then it turns out that more is true. So it turns out, if you put a sinusoid in, what you get out is a sinusoid of the same frequency. What might change is the amplitude of the sinusoid and the phase angle on the sinusoid. But it'll be the same frequency sinusoid that comes out. So that's a fairly dramatic restriction.

And that's actually key to frequency domain methods. What it means is we can focus on what an LTI system does one frequency at a time. I'll look to see how it behaves when I excite it with a particular value of this big ω here-- that's the frequency-- look at the response. I know the response will be exactly at that frequency. So all I have to capture is how much did this input get scaled by-- in other words, how much did the amplitude change by-- and how much did the phase get changed by? I just need to know the magnitude and phase transformation of the cosine at each frequency.

If I know that, I know everything about the system. And it decouples my design. It allows me to think frequency by frequency when I design with an LTI system. So this is actually a great simplification.

One other remark I've made here, by the way, it's certainly the case in continuous time, if I gave you x of t equals cosine, let's say, $\omega_0 t$ plus θ , this is always periodic. This is periodic with period 2π over little ω_0 -- a little ω_0 . Because any time I increase t by an integer multiple of this, I'm going to get an integer multiple of 2π added into the argument of the cosine. So I'll be back where I started, right? So this is always periodic.

But with a discrete time sequence, you actually have to be a little more careful. We can think of some particular discrete time sequence. We refer to the ω and the continuous time case as the angular frequency in radians per second. Here, we're thinking of this as angular frequency in radians per sample, for instance, because the thing it multiplies is n . But basically, the units of big ω are angle.

Well it turns out that this may not be strictly periodic in the sense that shifting it by an integer will get you exactly the same waveform. And that's all related to whether this frequency-- whether 2π over ω -- the thing that you would like to compute as a period-- is rational or not. It turns out, if 2π over big ω is rational, then the period is the numerator of that rational. But otherwise, it's not rational.

But you can think of it as being samples taken from some periodic quantity. So there's an underlying time varying-- sorry, there's an underlying continuous time periodic waveform. And you take samples of it. And depending on the frequency of the underlying sinusoid, the sequence of samples may be exactly periodic, or may be close to it in the sense that they're samples taken from a periodic signal.

In either case, actually, we tend to not fuss about that. We'll talk about a cosine like this as a sinusoid of frequency ω radians per sample, and we'll talk about the period as being 2π over ω even when it's not strictly periodic. So 2π over ω is our notion of period.

So a couple of examples here, you can-- I'll leave you to look through those on the slides. But you can easily construct examples where 2π over ω is rational, and draw a picture, and convince yourself that it's actually periodic. So in this case, for instance, if big ω -- big ω is whatever multiplies n , so it's going to be 3π over 4 . 3π over 4 is big ω . So now, I need to look at 2π over ω . And so that's equal to 8 over 3 . The numerator of that is what the period is. So if you were to actually sketch this out, you would find that every 8 samples, it repeats.

Whereas, if you look at this example, the thing that multiplies in is 3 over 4 . So now, the period 2π over ω is 2π over 3 over 4 . So now we're talking about 8π over 3 . That's not rational. You're not going to get a periodic sequence. But we still refer to 8π over 3 as the period of this discrete time signal, OK? Just because it comes from sampling an underlying continuous time periodic signal. All right, that's a detail. I just don't want you to trip up on that later on.

So here's the basic statement-- what I said earlier. If you have an LTI system-- by the way, I like to represent this with an h dot, so that slipped past there. I think I've talked about this before. Let me not spend time on it-- notation, notation. So if the input is a sinusoid of some frequency, and amplitude, and phase, the output is guaranteed to be a sinusoid of the same frequency, potentially different amplitude and different phase.

So what I want to do is establish that for you, starting with our time domain characterization, which is LTI convolution. It turns out, actually, sinusoids are not the only things that have this property. In fact, it might be good for me to show you another example, too. So let me give you another example of a waveform that you can put at the input-- or signal that you can put in the input to an LTI system, and it comes out the same shape despite all the convolution.

So here's the example. Suppose x_n is what I think of as a discrete time exponential. So this is r to the n for some real number r . So maybe-- maybe r is $1/2$. So this is-- so this is x_n . It's a discrete time exponential. You're used to thinking of that as a geometric series. But when you're talking about signals and systems, you like to think of that as a discrete time exponential. It does have an exponential fall off.

What if that goes in to this system? Well, the output is given by this expression always, for an LTI system. So let's just plug in what x is, that summation over all m , h_{n-m} , r to the n minus m . I'm just substituting in this expression for x of n . The n piece of this doesn't depend on the summation index. So I can actually simplify this further. And here's what I have.

So look what's happened. I sent in a discrete time exponential. And out comes the same discrete time exponential, but scaled by some number. This is just a number, right? It's an infinite sum. It's just a number. It works out to be a number. You'll have to look for conditions under which it's guaranteed to exist. And certainly, if the system is BIBO stable, and-- well, if the system is causal, and the exponential is decaying, then this is guaranteed to exist.

OK, so here's an example of another kind of signal-- related to the sinusoid as we'll see-- that has the property that you put it through the LTI system. Despite all this convolution stuff that's going on, I don't even have to know what that h is. I can tell you right away that what comes out is the same exponential but scaled.

So now, we're going to try and establish the property for sinusoids. And the way to do that-- the efficient way to do that is actually to work with exponentials again. Except they're not exponentials of the type that I have there. They're complex exponentials.

Now, when you learn complex numbers in high school, maybe you thought you wouldn't have to deal with them again. And then you came to calculus and you had complex numbers, and you thought maybe after that, you don't have to deal with them again. So I'm here to tell you that you're always going to have to deal with complex numbers. So you have to get comfortable with them.

So we're talking about, essentially, points in the plane-- point with a real part and an imaginary part. So here's the complex number c . Here's the real part. Here's the imaginary part. I never thought I'd need to define j , but actually, it turns out that if not everyone in the room is an electrical engineer, then maybe you're used to thinking of i as being the square root of minus 1. Electrical engineers like j , because they reserve i for currents, right? So j is the square root of minus 1 in all electrical engineering.

The key identity that you need is Euler's identity. Let me actually write it on the board and leave it there, because we'll be coming back to it multiple times. Well, actually I can have it on this board, because it's really the same. It's really this picture. e to the j theta is some complex number. Its real part is cosine theta, and its imaginary part is sine theta. That's all that Euler's identity is saying, right? Here's a complex number. Its real part is cosine theta, its imaginary part of sine theta.

So what's the magnitude of e to the j theta? Every complex number has a magnitude. e to the j theta has magnitude what? 1. Because it's cosine squared plus sine squared, right? And the angle of e to the j theta-- the angle of a complex number is just the angle from the real axis up to this vector. So the angle is?

AUDIENCE: [INAUDIBLE]

PROFESSOR: I'm hearing more complicated answers than I expected. Theta, right? You're right. It's arctan sine theta over cosine theta, which is theta, right? All right, so this is Euler's identity. And this is really critical. And if you can remember this, then you can remember all sorts of other identities that might trip you up from time to time.

So for instance, let's see, if I have e to the j theta 1 times e to the j theta 2, does that simplify to this? That's OK, right? That's just combining the exponents. Well, expand this out using Euler's identity. Expand this out using Euler's identity. Expand this out using Euler's identity, and you discover, for instance, that cosine theta 1 plus theta 2 equals cosine theta 1, cosine theta 2, minus sine theta 1 sine theta 2.

OK, you know all these identities. But if you're ever pressed to derive them, the place to start is Euler's identity. If you're pressed to derive Euler's identity, then maybe go back to Taylor series. But if you want to carry one thing in your head, carry Euler's identity. All the rest follows.

All right. So here is-- well, actually, let me make that point a little later. These are easy. e to the $j0$ is just-- this unit length vector line along the real axis is just the number 1. e to the $j\pi$, it's the unit vector lying along the negative real axis. So that's the number minus 1, right?

So here's how we use complex exponentials to prove the result that I claimed earlier, namely that sinusoids in gives you sinusoids out. A sinusoid of a given frequency in gives you a sinusoid of the same frequency out. What we're going to do is actually combine sines and cosines into one calculation. This may look a little funny, because now, we're suddenly putting a complex input into this LTI system.

But if you think about the math that we did for LTI, it didn't really care whether we were feeding in real numbers or complex numbers. So we could have a signal that has a real part and a complex part at each time. Convolution would work exactly the same way. We arrived at convolution through linearity and time invariance arguments, and all those work for complex signals. So we could actually be putting in a signal of this type and seeing what comes out.

Well, this is very close to the calculation we did earlier with a real exponential. The only difference now is it's a complex exponential. But here's the computation. We start, as always, with the time domain representation. That's convolution. Substitute in for x of n minus m . x of n is this signal here. I should have labeled it as x of n . So just stick x of n minus m here. And then, you discover that there's a part of this that doesn't depend on the summation index. So you pull that out. So you're left with this summation inside and this piece out.

So look what's happened. You've put in a complex exponential. And out comes the same complex exponential scaled by something. What is this complex exponential? Well, its real part is a cosine signal, and the imaginary part is a sine signal, all right?

So this object is something that we'll encounter and use a lot. It's referred to as the frequency response of the system. You've heard the term used, undoubtedly, in other settings. But here's the definition. It's a function only of big ω . Because once you've summed over m -- little m -- everything else has gone away. So only depends on big ω . And here's what it is. It's the summation over all values of m , h of m , e to the minus j ω m . That's the frequency response.

So if you give me the unit sample response of a system, I can find for you what the frequency response is. So let's actually do an example. Let's see. Let's take one of your-- let's take an averaging filter that you've looked at. To make it easy on you, I'll sketch it as a function of m . It doesn't matter what that's called. So here's $1/3$, $1/3$ $1/3$. We've come to recognize this as B unit-sample response of a 3 point averaging filter. It's a causal 3 point averaging.

So what's the frequency response of this system? Well, it's h of big ω . And then I just follow this prescription. So the only values of h that are non-0 are for m equals 0, 1, and 2. So it's going to be $1/3$, 1 plus e to the minus j ω , plus e to the minus $j2$ ω . And I'm done, right?

Now, to actually get a feel for this what you want to do is write it in different forms. For instance, a very useful way to write this is in terms of the magnitude and the angle. So here's a way to represent any complex number. This is a complex number, right? I can write it as magnitude times e to the j angle. So that will actually turn out to be a much more efficient way to think about frequency responses. This is the magnitude of the frequency response. And this is the phase.

So how does that bring us back to sines and cosines? Let me actually go a little bit out of order here, and here's the basic statement. From the result that we have up there-- and I'll show you on the next slide how to derive it-- what you can show is that if you put a cosine into the system, what comes out is the same cosine, except its amplitude is scaled by the magnitude of the frequency response, and the phase angle is increased by the angle of the frequency response. So this is really-- if you know this as a function of frequency-- if you know the magnitude and phase angle of the frequency response as a function of frequency, you can describe the response of the system to any cosine input.

What if the cosine was a little more complicated than the one there? Suppose I had-- suppose I had cosine ω n plus, let's say, π over 4 going in. What comes out? What do you think comes out? Anyone?

I showed you a particular case here. How much new work would you have to do if it was actually a slightly shifted cosine? Take a guess. Yeah?

AUDIENCE: Just replace n in the output with n plus π over 4.

PROFESSOR: Yeah, you just change this by adding in an extra π over 4. So what comes out-- that's what you said, right? Is that what you said? OK. Sometimes I'm guessing because I don't hear that well. So here's plus the angle, and then plus the additional π over 4. And this actually will follow just from time and variance of the system. So you can actually-- so this is actually pretty general. If you're going to remember one thing about frequency response in terms of what the operational significance is, for instance in a laboratory experiment, this is the result to remember. This is why frequency response is important.

And the proof, very easy. Once you have the basic result with exponential inputs, the proof is easy. Because a cosine can be written as the sum of these two exponentials. How do I get that? Just from Euler's identity, right? Use Euler's identity for each of these-- for e to the j , e to the minus j , when you add Euler's for this and Euler's for this, the sine terms cancel out. So you get 2 cosine in the numerator and you divide by 2. So if you're stuck for a derivation of a result like this, go back to Euler's.

So cosine can be written this way. So when I feed a cosine $\cos(\omega_0 n)$ of the input, what I'm actually feeding in is a linear combination of two exponentials. But I know how to write the response to an exponential. If this exponential-- this sum of exponential goes in, what comes out is the corresponding sum of responses. So it will be this exponential times the $1/2$ there, scaled by the frequency response. That's what frequency response does to an exponential. And this exponential comes out the same, but scaled by the frequency response-- evaluated at the frequency of that exponential.

So maybe you've-- you're having trouble visualizing the result that I'm invoking. But it's the one that we just proved. If you have $e^{j\omega_0 n}$ going in, let's say, at some particular frequency ω_0 to a system with frequency response $H(\omega)$, what comes out is the same exponential that went in, scaled by the frequency response evaluated at the frequency that we're talking about-- ω_0 .

So when I label the system with the frequency response, this is a general ω . But the value of it that I'm interested in is the value of-- at the frequency of the inputs. So if the input is $e^{j\omega_0 n}$, what comes out is that same $e^{j\omega_0 n}$, but scaled by the frequency response at that frequency. So that's what I'm invoking here. Invoking it twice, well, this is just the real part of this quantity, because it's a complex number and its complex conjugate, and so on. So you put it all together, and you actually very directly have this result.

So we're using complex inputs as just a trick for getting the results that we'd really like to get for real inputs. If you didn't want to do that, you could actually just put in the cosine $\cos(\omega_0 n)$, and crank it all the way through, and you would get it. So you could say y_n equals summation over all m , $h(m)$. Then we have $x(n-m)$ here, right? So it's going to be $\cos(\omega_0(n-m))$ going in. And now, use appropriate algebraic identities and you'll get the same result.

So we didn't have to use complex exponential inputs to get the result. It's just a convenient way of getting it. Again, if you're going to carry one result in your head in the complex case, this is the one to carry. It is very simple. It says complex exponential in, you get the same exponential out.

So let's play a little more with this particular filter since we see that magnitude and angle are so important. We're talking-- we're back to this 3 point averaging, right? Here's the frequency response. And I'd like to get the magnitude out. And you could certainly use Euler's identity on each of these pieces, group all the real parts, all the imaginary parts, and so on. It ends up actually being a bit of a mess to try and write down cleanly.

So let me show you a trick that works for this kind of thing. And it'll give you some practice in thinking about these complex exponentials. Do you agree that I've just rewritten the same thing as I had above? OK, does this simplify? Can you write it as something real? We've got a complex quantity and its complex conjugate there, so we should be able to collapse them into real, right? The way you recognize a complex conjugate is that the j has gone to a minus j .

So what is this whole thing? Somebody?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Where did that come from? Can I have a hand, just--? Oh, yeah.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Yeah, so let's say $1 + 2 \cos \omega$, right? So that's simplified nicely. OK, so are we in a position to say what the magnitude of h is? I've got an h that is represented as the product of this. Oh, by the way, sorry, I-- this should just be a 3 here, not a $1/3$. Everyone shook their head in agreement when I wrote that down, but-- what's the magnitude of the product of two complex numbers? Is it the product of the individual ones?

The magnitude of h is going to be, let's see $1/3$ magnitude $e^{-j\omega}$ times magnitude $1 + 2 \cos \omega$, right? The magnitude of a product of complex numbers is just the product of the individual magnitudes. What's the magnitude of this? 1. So we're actually done. We actually have a very simple expression for the magnitude of the frequency response of this moving average filter. That's all it is. Oh, sorry, I need the absolute value.

So let's see, I think I have that sketched out in one of these. I actually have three moving average filters drawn out here. The one that I've just worked out is this case. This is the 3 point moving average filter-- a height of $1/3$ for each of these at 0, 1, and 2, and everything else is 0. Here is the frequency response magnitude.

The notation that's used on the figure is slightly different. So some people, including me in the next course I teach, will write the frequency response as this instead of h of ω . But that's really unnecessarily fussy. It's important when you're talking about z transforms at the same time that you're talking about Fourier transforms. But for us, it's not important. So you'll see slightly different notation, probably in the notes as well. But just think of that as just h of ω .

So here's the filter we were talking about. Here, supposedly, is the frequency response magnitude. What we should be seeing is the magnitude of this quantity. And let me see if you believe it.

So what we have is-- so what I have is, at ω equals 0, I've got something that starts at 1. And then, when I get out to $-\pi$, this has come down to the value $1/3$. And so this is what I have for-- this is the quantity within the absolute sine, before I take an absolute value. And when I take the absolute value, this flips over. And that's really what you're looking at. You're looking at frequency response magnitude, which is this. And you've got to figure out the phase accordingly. And that I'll leave you to do in recitation.

But I want to ask you one last thing. Why do I not bother to plot the frequency response beyond $-\pi$ to π ?

AUDIENCE: [INAUDIBLE]

PROFESSOR: So the reason is that the frequency response tells us what the response is to inputs of this type, right? The frequency response says if this goes in, how much it gets scaled by when it comes out. Well, if I increase ω by an integer multiple of 2π , I get the same exponential back again. So there's no new information outside of $-\pi$ to π .

Another way to think of it is we're really talking about complex numbers and their angles. Once you've made a full circle from $-\pi$ to π , there's no new space to cover. So you'll see frequency response is only plotted from $-\pi$ to π . All the interesting action is there. All right, we'll develop more intuition for this in recitation and in the next couple of lectures.