[SQUEAKING] [RUSTLING] [CLICKING]

DAVID PERREAULT: OK why don't we get started?

[MIXED CONVERSATIONS]

DAVID PERREAULT: Spring seems to be giving everybody energy here. I wanted to continue talking about EMIs and filter design, and most of what I'm talking about is in *Principles of Power Electronics* chapter 26, plus some additional handouts that are going around and are also posted on Canvas.

Just as a reminder, last time we talked a little bit about how one goes about measuring EMI. And we said, OK, if I have some converter-- it could be a buck converter is our canonical example here, and it's drawing some really pulsating current from its input that's going to look like sort of ix is some big pulsation at the switching frequency, right? And then I could say, well, maybe I have some input filter that maybe here I'm just going to say is some capacitor. At minimum, I always have a capacitor here.

And we said, well, if I just connected this up to some source impedance, Z sub s, or some voltage source V sub s with a source impedance Z sub s, I would probably run into trouble with repeatability, because depending upon what Z sub s looked like, I would get different voltages, Vx. And maybe I'll call this current i sub y on the other side of the filter. And so it's kind of problematic from a repeatability point of view.

So if I want to test to some standards I need some better way of basically controlling the impedance we're looking back into. And so to do that, when we're doing EMI testing alone, we come and then insert an additional stage in here, which I'll draw in orange. And this would look like something like this. I'll have some box that is part of my test setup, called a LISN-- Line Impedance Stabilization Network. And it will, at minimum, have some L LISN and C LISN.

And then I will connect this up to some impedance analyzer that is usually 50 ohms. So if I looked at the impedance looking back this way, I would-- high EMI frequencies, when this becomes an open circuit and this becomes a short circuit, I'll look back into the spectrum analyzer impedance, which is usually 50 ohms, and I'll make my tests that way. So we get rid of the variability in the source impedance for testing purposes only by doing that.

And I should say, this is actually a slight simplification. In most systems, you actually put a separate LISN on the hot lead, the power lead, and on the ground lead, and you measure the response of the LISN on both. And the usual method you do is you're looking at the frequency content of what goes into the spectrum analyzer here, typically for conducting EMI between 150 kilohertz to 30 megahertz. But that again depends upon the standard you're comparing it to. Some military standards go down to much lower frequencies. And all the details. What are the LISN values, what does the test setup look like, et cetera, that's all part of a standard that you adhere to test your converter.

But what we-- the basic idea is as follows. If I think about I have some ripple frequency. I'll call it ix tilde here, which might be some large signal. I put it back into my filter capacitor, or my filter, more broadly. Here is just Cf. And then that goes back into some impedance which, in the test stand only, ends up looking like our LISN. And I can measure either this current-- I'll call it iy, or the more accurately, in terms of the actual test specifications, the voltage across the LISN impedance in order to make my measurements. Any questions about that?

So we said, OK. We need to design a filter here. A capacitor is not going to cut it because for typical specifications, who knows what this is? Maybe this is tens of amps. Typically, the current I'm allowed to squirt back into here is on the orders of tens of microamps. It's really small.

So we usually need to design a higher-order filter than that. And so what we've thought about is building with a basic filter stage. So if here's my converter that's driving noise, and I'm going to call this ix tilde, and again, whatever this term by the converter, I will put this into some filter. And here maybe I would think of having some filter capacitor Cf and some filter inductance Lf as a basic building block for a filter. And the notion is, I can look back. And on the other side of this filter, depending upon whether I'm-- in my test stand, I'm going to have-- I'm going to look back into something that looks like Z LISN, which at high frequencies might be 50 ohms and at lower frequencies is something else.

But let me just say, this goes back into some source impedance which for the moment, for our present discussion, let me just pretend it goes back into a short. And I'm interested in this current. I'm interested in this current iy tilde, the amount that's going back into the source, in terms of ix tilde.

And we said, all right, if I wanted to think about the attenuation of this filter, and maybe I could think of this as being going back into Z LISN under some other conditions if I'm in the test stand. Maybe this filter block would give me an attenuation that would go the magnitude of iy over ix tilde might look something like this. At DC, everything gets passed back.

And then well above 1 over the square root of Lf Cf it will roll off as 40 dB per decade, or minus 40 dB per decade. Why? Because the capacitor impedance is falling by a factor of 10, every factor of 10 in frequency. And the inductance impedance is rising by a factor of every factor of 10 impedance, factor of 10 for every factor of 10 in frequency. And hence so I get a factor of 100 attenuation for every factor of 10 in frequency.

And that's great. And I'm intentionally ignoring what happens in the vicinity of cutoff, because keep in mind, the cutoff some low frequency, we're going to put our PWM frequency and ripple frequency way out here somewhere so we're getting lots of attenuation.

But we have to be a little bit careful because we said when we lay these filters out, first of all, there are internal parasitics. So if I go buy a filter capacitor, he's not perfect. So maybe he has some internal ESR, RESR. And what that would do is that would reduce, instead of falling off like this, it would start to fall off at minus 20 dB per decade, because above some frequency, this shunt path just looks like a resistance.

And then if I keep going up, and I have some parasitic inductance here, L parasitic, then I just have a current divider between Lf and Lp. And in fact, this might even flatten out. So I might be thinking I'm going to be way down here, but really my attenuation's up here, and I don't get nearly as good a performance as I hoped.

So this could be due to internal component parasitics, or it could be due to layout parasitics-- how I place everything on the circuit board. Parasitic inductance and resistance is one thing I worry about. Parasitic capacitance is another thing I worry about. So that's just a quick review of last class.

But I'd like to now move forward a little bit, And let's start thinking about-- let's start thinking about what happens in the vicinity of cutoff. So if my converter is in the hundreds of kilohertz, you know, maybe this is-- this cutoff is in the kilohertz range. It depends on the system. But at any rate, my cutoff of my filter is going to be well below the switching frequency.

So let's think about what's going on there. If I have some converter, and it's drawing some ripple current, and then I put in my filter-- and I'm going to call this i sub y tilde-- let me just forget about the LISN impedance for a moment. What am I going to be looking back into? That depends what I'm going to put my converter into. The LISN is only there for testing. This converter and this filter has to live in the real world where this source impedance might even be 0, for all I know

So what's going to happen in the vicinity of cutoff? Well, let's think about this. If I plotted iy over ix versus frequency, and this is all on a log scale, I should say, what do we say was going to happen? At low frequency, I'm going to get a transfer function of 1, and at high frequency I'm going to go minus 40 dB per decade.

The question is, what happens here at cutoff? Where is cutoff? 1 over the square root of Lf Cf. And what do I look back into at exactly the cutoff frequency? I have an inductor in parallel with the capacitor at its resonant frequency.

So while I get this roll-off at high frequency, and I get all the current paths passed back at DC, basically, at the resonance frequency, Lf and Cf basically look like an open circuit. Or what's going to happen is, if I start to draw any ripple current near the resonant frequency, I'm going to get huge circulating currents in the input. And I'm going to drive a big current iy back into my input.

So the current gain here, actually, does something like this. It goes off to infinity in this simplified case. It would do something like this, which is pretty bad. If you imagine that the converter somehow had some small amount that it was drawing at some low frequency near this cutoff, it would drive this input filter nuts.

Another thing I could look at is I could say, well, what is the output impedance of the filter look like? In other words, the filter seen by the converter. If I look back into this filter, and I say, what is this impedance Zx, which is equal to Vx tilde over ix tilde, where this is Vx tilde. What does that look like?

Well, I would argue in this case that at low frequencies, this is close to a short circuit, so it looks like omega Lf. And at high frequencies, this starts to look like a short circuit, so it becomes 1 over omega Cf. And you get the following behavior. I get omega Lf, I get 1 over omega Cf, and right at this cutoff frequency these two are in parallel and these are resonating against each other. One has a positive admittance. The other has a negative admittance, and they cancel. And this impedance would go off to infinity.

So you're looking back into an open circuit. That makes sense why this draws kind of high current. You're trying to draw a current from an open circuit, and so you're going to get a lot of voltage ripple here and a lot of resulting current ripple back into the input. Any questions about that?

So what would I do about that? What would be the kind of natural thing to do? Better put a damping leg on your filter. So while we're using this filter way out here to cut off ripple from the converter, I also have to think about what it's going to do here.

Now the most natural thing to think about doing, putting in a filter, is maybe just putting a resistor-- maybe I could put a resistor-- I'll call it a damping resistor-- Rd-- right here. What would Rd do? Well if Lf and Cf are at resonance, Rd is the only thing left, because these kind of cancel each other. And my output impedance would just go to Rd. So instead of peaking off to infinity, I would peak off to some value Rd, would be the maximum of my output impedance if I did this.

And likewise, if I looked at the peaking in my transfer function, it would be limited by the fact that I was damping this. And you know, maybe it does something like this. And I'm no longer unhappy. Any questions about that?

Anybody see a problem with that scheme? Yeah, you're going to burn a lot of energy in this. I mean, if I come back to my original converter, I need to make Rd, typically, if I want this to be nicely damped, on the order of square root of Lf over Cf or lower for good damping. And I'm basically putting a resistor across the input of my converter. That's not very happy. So what would I do in the real world? I'd say, well, I can't accept that being a very good idea.

But maybe what I could do is take this damping resistor and put it in series with a damping capacitor, Cd. OK. So I'll put Rd in series with Cd. OK, what does that do?

Well, if Cd is bigger than C sub f near the cutoff, 1 over the square root of Lf Cf, Cd looks like a short circuit, and it looks just like Rd is in shunt. And so I'll get this nice kind of damping.

But a DC, Cd is an open circuit, so I don't get any DC current through here. And so that helps me. OK, so this Rd Cd branch is sometimes called the damping leg, And it can really do a nice job of improving your filter design.

Now typically, Cd depends-- and we're going to talk about the details of this-- a typical number for Cd is usually on the order of 2 to 10 times Cf, or two or greater times Cf in order to get good enough damping. But we'll get into the details of that. But before I do that, any questions?

So let's think about what would happen if I put this damping branch in here. How do I pick Rd. Or another question I could ask is, what does the resonance of this filter look like as I vary Rd in this system?

So let's take a look at that. I'm going to replot it here. And I'm going to look at Zx, the filter output impedance. And we're going to come back to why I'm obsessed with what the filter output impedance looks like. And that's the impedance seen looking back, what the converter looks back and sees. And in doing my analysis here, I'm just assuming that the source impedance out here looks like a short circuit. I'm making that simplifying assumption.

But let's think about this. I said in the undamped case, what would I have seen? I would have seen something like this-- 1 over the square root of Lf Cf. This is omega Cf. Oops. I'm sorry. This is omega Lf. And then this is 1 over omega Cf. And this is versus frequency on a log log plot. And this is the magnitude of Zx, looking back there.

If I let Rd go to infinity, it's basically like Cd is not in the circuit. So as Rd goes to infinity, what I'm going to see is something that looks like this. This is the case where Rd goes to infinity. This is just going to do what the undamped system did.

So this is the case where Rd goes to infinity. So, clearly, I can't make Rd too big, or the-- even putting in the damping doesn't do me any good.

Well, what happens if I make-- so, clearly, I need to make Rd small. What happens if I make Rd super small? Well, if I make Rd go to 0, what happens? If Rd was 0, i would just have Cd in parallel with Cf. So if I made Rd go to 0, what I would see is something like this. Instead of having a cut off 1 over omega Lf Cf, I'm just putting Cd in parallel with Cf so I'd have some other place-- here it is-- 1 over the square root of Lf times Cf plus Cd.

And I'd have another cutoff. Actually, let me put that a little bit over to the right. Let me put it here just so you can see it more easily. So I'm going to put that cutoff right here. And what I would get is the same-- I would get a different high-frequency asymptote. Actually, let me put it right here. Be accurate.

OK, so this would be 1 over omega Cf plus Cd. And I would peak up, and I'd get something that went off to infinity like this. So this is the case where Rd goes to 0 for some value of Cd that I've picked. Does that make sense to everybody?

So this is a Goldilocks thing. If you make your damping resistor too big, you get bad damping. If you make your damping resistor too small, you get terrible damping. And you got to, for a given value of Cd, you've got to pick the right damping resistance.

What would my output impedance look like if I picked the right one, where you can show that the best case, through a nice bit of analysis, actually goes through these two points. So if you picked Rd optimum, what you would get is something that did this. And then it would eventually fall off like this.

So that's a pretty-- kind of crummy drawing. But this is sort of best Rd, the best Rd in terms of limiting how big the output impedance gets, which we could take as some kind of measure for how badly damped the filter is. Questions about that?

So what do I do if I don't like that value of output resistance that I get? Well, if I make Cd bigger, I can slide this down, and then I'll pick a different value, which will then bring this intersection point down. And we'll get better damping.

So I'm not going to derive this. This is actually nicely derived in one of the handouts that I've provided. But here's the final result. If I define R out as being equal to square root of Lf over Cf-- that's the characteristic impedance of my original Lf Cf branch-- and I define n as being equal to Cd over Cf-- so that's how much bigger is the damping capacitor compared to the filter capacitor-- what you can see is that Rd opt is equal to R0 times square root of 2 plus n times 4 plus 3n divided by 2n squared times 4 plus n.

And this gives a Zx max-- if I call this Zx max-- a value that's equal to R0 times square root of 2 plus n over n. So that means, the bigger I choose n, this goes down, and my peak output impedance of the filter goes down. All right? And I get essentially better damping. OK. Larry.

AUDIENCE: Your notes say 2 times 2 plus n.

DAVID PERREAULT: Whatever my notes say are correct. 2 times 2 plus n. Yes.

Now by the way, these are nice equations. There's a very clever analysis that shows all this. And certainly, this is well known, and it's in one of the handouts.

The nice thing about this is that you can say, well, how small do I or how well the damp do I want this thing, I will pick n big enough to give me that, and then that gives me my value of damping capacitor, and away I go, and I get enough, make my output impedance small or my system well damped.

Now the thing we don't like about doing this is you say, well, why don't I just make n really big, and then this will be really well damped. Well, the reason is, I've got to pay for that capacitor. If I'm choosing Cd as 10 times Cf, well, that's expensive, in terms of cost but also in terms of volume, because it has to store 10 times the energy of Cf, right?

On the other hand, Cd and Cf, if I come back to this original design, Cf is important at the ripple frequency We were talking about-- up here, we were talking about, well, I really care about the parasitics of this filter capacitor in order to get good, high-frequency attenuation.

So Cf's job, the filter capacitor's job, is really to help me attenuate the high ripple frequency. He wants to be a really good capacitor for high frequency.

On the other hand, Cd is really operative near the cutoff, which is orders of magnitude lower. So Cd doesn't have to be a great capacitor. In fact, if Cd had parasitic inductance, I probably don't care that much. In fact, if he has parasitic resistance, I still may not care that much, because I can use that parasitic resistance as part of my filter damping leg.

So this might be, for example, an electrolytic capacitor, whereas this might be film or ceramic capacitor. So the jobs of these two capacitors are different, and the construction of them. And you can get a lot of energy storage in electrolytic, for example, as compared to a ceramic or a film capacitor typically. Any questions about that?

All right. So what we're going to do is we're going to say, jeez, how well damped do I want this to be? First of all, design my filter Lf and Cf to attenuate the ripple frequency. Then I'll say, how well damped do I need this to be down at the cutoff? And I will say, OK, I will pick n to be big enough, the ratio of Cd over Cf to be big enough, and then out of that will-- out of this will drop what filter damping resistor I need. And then we kind of have a filter.

By the way, if you don't like remembering these equations, it's not too hard to simulate or just calculate for a few values, and then you'd zero in on it. It's just the Goldilocks thing. It's too hot, it's too cold, you find the one that's just right, and then you're there.

I should say, this isn't the only way to damp a filter. If I thought of a-- if I thought of a different damping mechanism, right, I could come back to the same basic filter structure. Here's Cf. OK. Here's Lf. And he's going to be coming back into the source impedance, which are just presented as a short circuit. And he's going to be driven by some converter ripple ix tilde. And I want to worry about what iy tilde here is, this current gain of the filter.

I could come back to this and say, you know what? Instead of using a shunt damping leg like that, in parallel with Cf, maybe I could go get a damping leg like this. I could have Ld and Rd here. To the extent that this is a short circuit, this now also gives me parallel damping.

In this case, I would make, if this is Lf, in this case, I would make Ld less than Lf, so that nearer the cutoff frequency, 1 over the square root of Lf Cf, Ld now looks like a short circuit, and it's Rd that's damping Lf and Cf.

When I get to high enough frequencies, Ld starts to look inductive. It looks like a higher impedance than Rd, and I get my roll-off.

Now this is also a very popular damping scheme. And you can do the same thing for this analysis. There's always going to be, for a given value of Ld relative to Lf, there's an optimum value of Rd to give you a good damping. The downside to this filter-damping strategy is that, whereas this additional shunt path only kind of improved your filtering well above cutoff, here, because Ld is small, you're kind of hurting your filter impedance This path's impedance gone down compared to this path's impedance. So this doesn't do as good a job slightly above cutoff because Ld is small.

So what will happen is, instead of getting-- if I looked at iy over ix, iy over ix, instead of cutting off if minus 40 dB per decade, what's going to happen is I'm going to cut off at minus 20 dB per decade because all I have is this damping resistor being important. And then, once Ld comes in, then I'll start falling off, and I will eventually roll off at minus 40 dB per decade.

But my roll-off is not going to be as good. So what do I have to do? I have to pick a correspondingly lower cutoff frequency to account for what Ld is doing to my filtering. So that's the downside of this inductive damping scheme.

The upside to the inductive damping scheme is that whereas Cd costs me a lot of energy storage because he's got to block the full input voltage, Ld is small and carries no DC current, so it can be a pretty small inductor. So they both have their upsides and downsides. They're both widely used. Just depends on the situation. Try them both out.

In one of the handouts that I provided, I have actually shown all kinds of different filters, some higher-order versions than this. But typically, an Lc filter as a building block is a pretty good structure. Any questions about that? Larry.

AUDIENCE: Do we have to consider any of this at the power stage? Because you also have notes saying--

DAVID PERREAULT: Oh, yes. Well, yes. Actually, let me answer that exact question. The answer is yes, you do.

Why was I talking-- I kind of focused in this picture. I focused on what the output impedance of the filter looks like. Well, let's think about this for a second. Suppose I have some power stage. Here's my power stage. This is my power converter. And he's got some you know R out here that he's controlling some V out on. And he's being powered from an input voltage that I'm going to call Vn.

And then I'm going to have some filter. And let me consider a simple filter. Here's Cf. Here's Lf. Let me put-- I'm just going to use a simple resistive damping leg, Rd, just for simplicity's sake.

And now this is going to get connected to my power source. So here is Vn. Or I'm going to call this Vs and Zs. But again, let me just consider the case where Zs is 0, for simplicity.

All right. What does this power converter do? Let's assume he's a perfect power converter. He's going to go and keep the output voltage at the reference no matter what the input voltage does. That's his job. Input voltage varies. I don't care. I'm going to keep 12 volts at the output no matter what.

So what would I say that he does? Well, he's a constant power load in that case. No matter what the input voltage is doing, I want to draw V reference squared over R0 if this is some fixed load at the output.

So if I thought of this current as being some i n, What I get is i n times Vn is equal to whatever my power output power is going to be, which is V out squared over R out.

Or I could write this as in terms of the-- if I reference, write it in terms of voltage, I could write this as V out-- I'm sorry, Vn is equal to P out over i n. And I might evaluate I might consider that as some operating point at which I'm sitting Vn and i n.

So if I were to plot that-- let's plot that. What we're going to get is Here's Vn, here's i n. This converter, his job at his input port is basically to maintain the system on this curve, where Vn is equal to P out over i n. And if he's going to sit at some operating point-- maybe that's here. Maybe I have some DC operating point Vn and capital I n that we're going to sit at.

So how does this behave for variations? Suppose my input voltage wiggles a little bit. What does my input current do? Well, it's exactly the kind of linearization we caught before. I can look at what this tangent line looks like. And I can write d Vn, D i n evaluated at Vn i n is equal to minus p 0 over i n squared. All I did was differentiate this, and evaluate it at the operating point.

And this is equal to what I'm going to call r sub l. This is r sub l. This slope is r sub L. Does that make sense to everybody?

Anybody notice anything about this resistance here, r sub l, this slope?

AUDIENCE: It's negative.

DAVID PERREAULT: It's negative. So this negative sign is the source of a lot of pain and embarrassment, because essentially, from a small signal perspective, the power converter looks like a negative resistance to the source. And negative resistances tend to make you unhappy because why? Because they can make things oscillate. All right?

So if I then thought of what this converter looks like, OK, I might have, looking this way, what I see is rl is less than 0. Looking this way, I see a negative resistor.

What do I see looking this way? I see Lf, Cf, and Rd. So if I were to redraw this, what I have is essentially Cf, Lf, and Rd in parallel. And so this, we've been calling-- this output impedance we've been calling Zx, right? And he's looking in the other direction. He's looking into rl less than 0. And this is my converter and his load, in the case where the converter is perfect. All right.

So what's going to be true? Well, basically Rd and rl in this example are in parallel. So the net resistance, R net, the effect of resistance damping the Lc tank, is equal to Rd rl over Rd plus rl. Our net is greater than 0. Our net is greater than 0 for the magnitude of Zx-- or I should say, for Rd is less than the magnitude of rl.

So if I'm looking in here, and that slope-- this slope is minus 1 ohms, I'd better make sure Rd is less than 1 ohm. Otherwise, basically this LCR stage is going to start to oscillate. And that's not going to be happy.

And this is true in this instance, but it's true in a lot of other instances, where you have to deal with these kind of self-regulating loads, is they just tend to destabilize your system. And in this case, the system we're destabilizing is the input filter. But in other systems, you can destabilize other things. So in a DC power system, you might worry about what it's doing to long inductive runs of cable or capacitive filters on things. You can start to make oscillations in constant power loads.

So generally, the reason we obsess over what's our output impedance of our filter is because I need to keep it small enough that when it looks into a negative input impedance-- effectively a power converter-- the whole system is going to be stable. And here I've shown you in the simplest case for Rd. But generally what we need, the more general statement of this, is that we need the maximum-- output impedance of my filter across all frequencies ought to be well below the magnitude of P out minus P out over I n squared.

So I've got to pick my damping. And that's going to drive me to go pick some minimum-- if I were to use this capacitive damping scheme, that's going to drive me to pick some minimum value of the damping capacitance so that I can minimize the maximum value of my output, my filter output impedance, across frequency. Does that make sense to everybody?

Now just to drive home in your head that this isn't some-- just some kind of mathematical funniness, Muncie has been very kind to set up an actual power converter with an actual filter. And this is actually a buck converter with an Lc input filter.

And then what she has done is made sure that we can basically put a switch on this damping leg. We can put the damping leg in or we can take the damping leg out.

And if we power this up, that's the input voltage to the converter. So it's actually on the-- it's exactly the voltage Vn here that we're seeing. This is the input voltage that we're looking at-- Vn.

If we have the filter in place, Vn basically sits at Vs. But if I take out that damping leg, that's what happens. The input voltage is still fixed, but this voltage is going crazy. And you say, well, I could live with that. I don't care. Well, it depends. Sometimes you can hear a lot of noise.

This is just with the converter at some value, some value of load. But imagine this thing we're trying to drive, draw current varying near that frequency, it would go really crazy.

So you really can't tolerate this kind of thing. And you would see this in the out-- converter response at the output and so forth. So you really do have to seriously account not only for the fact that you need to damp the LNC, but you need to account for the negative input impedance of the converter in order to get acceptable response. Any question about that?

Now what I've told you so far is necessary. You'd better design your filter for that. But it comes out to be a little bit trickier too, because we've talked about modeling and control of your converter. We've said, OK, I want some certain transfer functions. We use averaging to derive kind of local average models, and you get your transfer functions and all that.

When we did all that, we assumed that you were working from a perfect source. The picture we drew for our boost converter was just you as a voltage source. It's just zero-impedance source. We pretended there was no filter there.

In the real world, there is a filter there. And that filter can indeed impact your converter. If I come back to the baseline here, maybe I don't have my LISN here. That's fine. There's no LISN when I'm running in the real world.

And maybe my source impedance is low, but I still have some converter and filter. And who knows what this looks like? This looks like whatever thing you drew here. But I could replace this whole thing with the source voltage Vs and an output impedance Zx, the output impedance seen from my filter.

So this is really the model of my source and filter. It's the source voltage in series with the output impedance of my filter. That's a Thevenin model for my system.

And my assumption in designing the converter and controller was that Zs was 0, or Zx was 0. But Zx isn't 0 once I've designed my filter. How close do I need to make it towards 0?

Well, what I would like to do is have the case where all of my models that I used to derive my converter, or control of transfer function, all of that stuff is valid. It's not affected by the filter I designed later. And a sufficient-- not necessary, but a sufficient condition to doing that is two things. And I'm not going to prove this to you. But there's some very clever analysis. People have shown this, and in particular, Erickson and Maksimovic at University of Colorado Boulder, did some very nice analysis of this. And a sufficient conversion-- a sufficient condition is the following.

First, if I assume my converter operates perfectly like this, we would like the output impedance of my filter, Zx, the maximum value, to be much less than the magnitude of the negative input resistance if I have perfect control. And that would be perfectly the case if I had a perfect converter that had infinite bandwidth and controlling the output anyways. So I kind of need to meet that requirement.

But what they showed is, in additionally, a sufficient condition, although not necessary to getting a stable system, but sufficient to be that you can ignore the design of the filter or decouple the design of the filter from the design of your controller for your converter, is to also say, if I imagined, instead of having a perfect controller over the output, I just regulated my converter at fixed-duty ratio.

And I looked at the input impedance of the converter at fixed-duty ratio across frequency-- and that's going to be variable across frequency-- at every frequency, that input impedance under fixed-duty ratio ought to be smaller than the output impedance at the same frequency looking back into my output filter. If you put those two conditions together, you get a complete sufficient condition that you can ignore the effect of the input filter on the rest of your control design.

But a good starting place for all that is just to make sure that under perfect control, that your filter is well enough damped to deal with the negative resistance of a perfect controller. So that's the place I would start in thinking about any design.

So that's a little bit about filter design and filter damping. The last thing I should have mentioned was, well, what if this Lc filter isn't enough? You can cascade Lc filter stages and get more and more attenuation. And again, there's rules for how you want to cascade those filter impedances. And again, those are in the handouts.

Any last questions before we wrap up for the day? OK, next class we'll talk about one other additional item that's very important in filter design. Have a great day.