

[SQUEAKING]

[RUSTLING]

[CLICKING]

DAVID PERREAULT: OK. Thank you. Let's talk about control a little bit more. What I wanted to do today was a little bit of an example just to illustrate the trade-offs and what you need to do to think about a control design for a general converter.

And the example I'm going to do is very, very garden variety. It's not very-- it's not very extreme in any dimension, So let's think about our boost converter where I have some input voltage u , an inductor, a switch, modulating with q of t , a diode, an output capacitor C , and a resistive load with the output voltage on it.

And the parameters I chose for this are very, very simple, not aggressive. L is 10 microhenries. C is 50 microfarads. The switching frequency is 100 kilohertz. Pretty low. The output reference that we're going to try to make this go to via ref is equal to 24 volts. And we said some designs, that output might move around. The reference that you want to program it to might move around. Here, I'm just going to pretend it's fixed that's about the simplest case.

And what are the things we said we had to deal with in terms of variations in a practical converter? Well, one of them is the input voltage. My input voltage might not always be what I wanted to. So we might say the u nominal, the nominal input voltage might be 9 volts. But we're going to design this for 8 volts less than u is less than 10 volts.

Now that's a variable range, because I want to illustrate to you the impact of a variable input voltage range, but that's a pretty narrow range. I mean, many converters, 2 to 1 is very typical. 4 to 1 sometimes. Depending upon the application, even more than that. So I don't want you to get the notion that this is some extreme range. This is actually a really narrow range.

The other thing we said we had to deal with was variations in the load resistor. And I'm going to use a range of 2 ohms-- that's the heaviest load-- less than R is less than 10 ohms. And I'm just going to use a pure resistive load for simplicity. So that's about-- this is the heavy load case. That's high power for a given output voltage. This is low-power case. This is about a 5-to-1 power range.

All right, that's a little bit more respectable, but again, a lot of converters have to deal with very wide power ranges. And sometimes we bring in additional tricks that I'm not going to discuss to deal with really the wide load ranges. But you'll get a notion of the impact of load resistance by the load range. OK

All right. So that's going to be our example. Let's step back and just remember what we did to model all this stuff. We started off with the switched equations of state. So I'm just going to copy these down just so we can see the progression. So this is the dI/dt is equal to $u/L q$ plus $u/L (1-q)$. And dV/dt is equal to $1/(RC) V$ plus $1/C IL$ minus $1/RC V$ times $1-q$. And I should have said this is q of t .

So in this original model, all we've done is said broken this down into two states and use my switching function to switch back and forth between the states. And if I knew initial values for i_L and v and I knew q of t , how it was modulating, basically, at least in continuous conduction mode, I could exactly simulate the actual behavior of that converter.

So if I did something like-- if I simulated that set of equations and I put a change in the load at some time and I put in a new R , I could calculate the transient there. But we said, I don't really care about all this ripple. I can't do anything about that in my controller, so we want to get rid of that. And the thing we said we were going to do was apply this notion of state space averaging.

So if we apply an averaging to this-- this is $\bar{x}$ is equal to $\frac{1}{T} \int_{t-T}^t x(\tau) d\tau$. I can do that to the circuit or I can do that to these equations. Either way, I can generate an average set of equations. So I get an average model. And in this case, it was $\frac{d\bar{i}_L}{dt} = \frac{1}{L} \bar{u} - \frac{1}{L} \bar{v}$ and $\frac{d\bar{v}}{dt} = -\frac{1}{RC} \bar{v} + \frac{1}{C} \bar{i}_L$.

And so this is our state space averaged model. I could also generate a state space average circuit which would have the same dynamics. And that's helpful, Why? Because this thing here doesn't have any switching in it at all. d is how I'm varying duty ratio over time, not the switching function. It wouldn't have any switching ripple.

So in fact, if I simulated this model-- first of all, I apply the averaging, which kills the switching ripple, and I simulate the average model, I will get what we see up there, the solid lines that capture the low-frequency dynamics and kill the ripple. And that's pretty nice. That helps me. If I thought about my original goal for building a control system, it was along the following lines.

Maybe I have some reference voltage V_{REF} . And what I could think about from a control modeling point of view is comparing that to the output voltage-- and here, I'll think about the average local average output voltage. I'll put that into some compensator, which maybe I'll call $G_C(s)$. So this is the error. Maybe I think about the local average error.

And then I'm going to get a duty ratio and then I'm going to put that into the model for the converter. And the converter also has an input u , the input voltage. And out comes $\bar{i}_L$ and $\bar{v}$, which is the thing I want to control, and I feed that back. So I could take this set of equations and stick it in here, and suddenly, I would have a model for my control system.

The only fly in that ointment is the fact that this average model is nonlinear. And the nonlinearity really comes in here and here. Now that doesn't mean I can't build a controller with a nonlinear model. People do that all the time. But a lot of the tools that you have for control design are based around the notion of a linear system. That's the easiest thing to do. It's the most flexible. So people generally would like to have a linear model for their system.

AUDIENCE: So it's the duty cycle that introduces the nonlinearity?

DAVID PERREAULT: It is the fact that the response of the system has the product of the complement of the duty cycle and the voltage and the complement of the duty cycle and the current. So it's the fact that when I modulate duty cycle, I'm doing something that modulates how the inductor current changes. And that modulation is also proportional to the output voltage.

So it's the fact that-- you can see that in here, is basically it comes time to the duty ratio and the output voltage effect how this is coming as a product. We don't like that. We want a way around it. So what do we do? We linearize.

And we said, what we can do is we can say, suppose this converter is operating at some point, and all I want to know is if I move it a little bit away from that point, how will the response wiggle? I perturb the input just a tiny bit, the perturbation in the states and the output, maybe that looks close to linear.

So we went through our linearization process, which we did last time. And here's the result we got. Now we're going to talk about-- instead of about the average models, we're going to talk about perturbations, things like $\tilde{I}_L$, which is equal to the difference between the actual local average current in the inductor and the DC current in the inductor. So it's the deviation away from the steady state.

And when we did that, we said, OK, here's the kind of equations I'm going to get. I'm going to get $d\tilde{I}_L/dt$ is equal to $1/L \tilde{u} - d'/L \tilde{V} + V/L d$. And $d\tilde{V}/dt$ is equal to $-1/RC \tilde{V} + D'/C \tilde{I}_L - I_L/C d$.

What did I do to get this? Essentially we took the nonlinear differential equations and we did a Taylor series expansion. We threw out the DC terms because we're only looking at variations and we threw out the products of small variations. So it's only valid in a limited range, but surprisingly, especially for switching power converters where the ripples are triangular, it turns out, you can get pretty far away from the operating point and it's still pretty good.

So this model is pretty good. What's important about this model from our perspective is that I have-- this is a pair of linear constant coefficient differential equations. And constant coefficient in the sense that these are the coefficients to the variables. So now my variables are all small signal. These are all just constants. All right. And with this, that's a linear system. I can take a Laplace transform of it, get nice stuff, I can make transfer functions, everything else.

And the only caveat is that notice, that these constants are each-- like if my-- if I decide my DC output voltage changes or, for example, I operate from a different input voltage, I'm going to get a different steady state duty ratio D , and then, hence, a different $1 - D$, which is D' here.

So these constants do depend upon the load current, the input voltage, all those things. So I got to think that the solution I get is only for one point, and I've got to go find the numerical values at different points. But nonetheless, here, I have this model, which maybe is even better to put in here. The difference is, here, I might have $\tilde{d}$, I might have $\tilde{V}$, this is $\tilde{V}$ error, and so forth. And $\tilde{I}_L$. So that's a better model to use-- or a simpler model to use from a control design point of view.

And then we said, well, how do we mostly use that kind of thing? Well, what we can do is we can say, take the Laplace transform of this thing. This d/dt 's becomes s times $\tilde{I}_L$ and this becomes s times $\tilde{V}$. I can then eliminate $\tilde{I}_L$ from the equation, and I can actually come up with a transfer function.

And if we do that, we're going to get-- if u is equal to u -- in other words, is not changing, what I get is $H(s)$ is equal to $\tilde{V}(s)/\tilde{d}(s)$, which is equal to $-sV/RC D' + VD'/LC$ divided by $s^2 + 1/RC s + D'^2/LC$.

So this is now a transfer function. And I can just stick H of s in here and get my transfer function from d to-- $\tilde{d}$ to $\tilde{V}$, and now I'm in a position to do control design. Before I move on, are there any questions about any of these steps? These are the steps we covered in the last couple of lectures.

OK. So I've taken my switched equations-- or I've derived my switch equations. And I've averaged them. Then I've linearized them. And now I have a transfer function. I can start to think about my converter control. Let's think about, what is this transfer function look like?

If I plotted the poles of this circuit-- so here's the complex plane. Here's the real value of s and the imaginary value of s . What I would see is I will have one right half-plane 0 because I'll have a positive value of 0, which generates a 0 in the numerator. So somewhere over here, I'm going to have a 0.

And then if I plug values in here for typical parameters like I'm going to see, what we're going to see is we're going to have a pair of lightly-damped poles in the left half-plane. So this means the response is naturally oscillatory. That shouldn't surprise you because when we did a little load step there, we saw a second-order oscillatory response and settling down.

OK. And again, where these poles are, we'll notice that there's a D prime squared here. So where the actual poles are depends upon the operating point. Depends upon what steady state duty ratio I'm operating at. But it looks something like this.

Now you might ask, gee, you did all this math. Is this model that's popping out, is that just random? And the reason you might want to wonder that is just because this system has a right half-plane 0.

Now those who designed control systems hate half-plane 0's. They are very annoying. Why are they annoying? Because usually if I take some system. And I put some gain-- forget about what this control is. Say it's just a proportional gain. And I turn up and I close the loop and I turn up the gain, what tends to happen is the poles move.

The idea is I had some open loop poles, and when I put a feedback loop around them, the dynamics are going to change. And what happens is that the closed loop poles, when I close that, are somewhere else, but it tends to be that as I turn up the gain, the open loop poles tend to move towards the open loop 0.

And what that means is, if I took this thing and I say I put it in a proportional feedback loop, the poles would tend to move towards the 0. Now if the poles move towards the 0 and the 0's in the right half-plane, that means the thing tends to go unstable. So something with a right half-plane 0 is very, very easy to make. Unstable generally. So we don't like that.

But you could say, well, OK, you're mathematically finagling a bunch of stuff. Can't you just do a version of a model that's not going to have a right half-plane 0? Like, why do we have to be-- why does life have to be difficult? And the answer is the mathematical model you get actually very nicely reflects the physical dynamics of the system. It's not accidental that there's a right half-plane 0 in the equations. It's intrinsic to the behavior of the system. OK

What does a right half-plane 0 tend to do? If I take a system with a right half-plane 0 and I step the input-- so suppose-- here's my input, And I step my input up a little bit. What I expect to do-- eventually what I want is that the response, the output should go up, too.

A system with a right half-plane 0 tends to do this. I want my output to go up from-- this is the output. I want it to go up from some level to some higher level. But what a system with a right half-plane 0 will do is as soon as I step the input, the output will actually go the wrong way, and eventually we'll settle down and go the direction I want.

So right half-plane 0's are also considered nasty because you want it to go in one direction, but before it goes in that direction, it goes in some other direction. And this system's like that. But I would argue that this is like that because of the way the system behaves.

Suppose I wanted to make my output voltage bigger. Forget about closing the loop. Suppose I just want to-- I'm running at some duty ratio. Suppose my input voltage is 5 volts. I'm running at 50% duty ratio, and so my output voltage goes to 10 volts-- is at 10 volts. But suppose I suddenly want to make my output voltage higher. What do I do? What do I do to my duty ratio? I increase it.

So suppose at some point I said, no, 10 volts isn't good enough. I want to go to 20 volts. I will step my duty ratio up. So I was running at 50%. Whatever's in my inductor current, half the time is going to the output, and that's what's supplying the resistor. Suddenly, I increased the duty ratio.

Now suddenly, instead of going to the output half the time, the inductor currents going to ground half the time. So suddenly, there's less current feeding to the output. And so what's the output voltage going to do? It's going to start to fall because you're not feeding as much current to the output.

Eventually, because I've turned up the duty ratio, I'm applying a positive voltage across the inductor, the inductor current is going to build up. And eventually, the inductor current will catch up and it'll feed the right current to the output. But what you're going to get is exactly a behavior like this. The output is going to droop and then it's going to come up.

So the right half-plane 0 in my control model, this guy here, in fact, precisely reflects the fundamental physics of how energy is going through the converter. So you're so we're going to have to live with it. It's not a fiction of the math, it's a fundamental physical thing. Any questions?

OK. So now let's think about building a controller. I want to come up here and I want to build some compensator GC of s. And the idea here is that I'm going to have some error between my reference and my output that generates an error voltage. I amplify that such that if that error voltage is positive, d gets bigger, and my error then gets smaller and I get close to what I want. So this thing, at least at DC, ought to have a high gain.

So let's think about this. What do my plant poles look like here? My plant poles, I said, were one half-plane 0 and two left half-plane lightly-damped poles. All right. And there they're complex conjugates. So what would happen if I made GC of s just a proportional controller? I'm going to make this a constant gain. So I'm going to get a duty ratio that's proportional to the error. And what I would want is to turn that gain up so that a tiny error gives me a big duty ratio and I get a small steady state error.

What I told you is as I change values of GC of s-- if I use one-- a first value of GC of s, if I use a first value of K_p , maybe my closed-loop poles would move here. And if I kept a higher value of K_p , they'd move here. And if you do this, you can form what's a root locus and say as a function, as I turn up K_p , where do my closed-loop poles go? And what this thing would do is it would do something like this. And then the poles would split. And so these are for different values of K_p .

So that's a problem because if I turn up-- I want a high gain, because I want the error to be small, but if I make gain too big, once that pole moves to the right half-plane, this whole thing is going to oscillate. So this is the nasty thing about the right half-plane, is that these-- the closed-loop values will be moved from where the open loop poles are somewhere into the right half-plane.

So you wouldn't want to do a proportional compensator on this just because you're going to-- either you're going to-- either it's going to have very bad stability or what's going to happen is if you your gain to be a really low value. That means that you're going to have to end up with living with a pretty big error between the output voltage you're getting and your reference voltage. It's going to be very droopy output.

So what could I do? I could pick some other compensator. Now I'm not going to try to design a good compensator for you here, that's not the topic of this class. But what's the next simplest compensator maybe I could think of that would fix that problem?

And I would argue you could use $G_C(s) = K_I / s$. This is an integral controller. $1/s$ is an integrator. So what I could do is instead of making my K_P proportional to the error voltage, I could make it proportional to the inner integral of the error voltage. So K_P and K_I are parts of PID controllers. And the beauty of an integral controller is eventually it will integrate up the error and leave me with 0 steady state error.

So if I did that, I still have my plant poles-- I have two lightly-damped plant poles and a right half-plane 0. And if I did this, well, I'm also going to have another in my forward loop. I also have now the dynamics associated with the compensator. So here's his open loop pole. And when I start dialing those around, what it would look like is this.

This pole would move this way, trying to cross through infinity to come back to the 0. And then this-- these poles would do something like-- they'd first swing inwards, and then they'd come outwards like this, and they do something like this. And if you take a control class, they'll tell you exactly how to calculate all these geometries and what they do.

So this is-- if I add a certain value of K_I , maybe my poles would sit here. If I turned up K_I really big, maybe my poles would be here. And I wouldn't want that because now it would be oscillating. So I can still make this thing unstable by using too much integral gain. But even with a small integral gain, I can still get to-- 0 error here gives me a duty ratio, which gets my output, so my output voltage will track my reference even with a small integral gain. Any questions about that? Yeah, Lowry?

AUDIENCE: How do you know which line you're on-- [INAUDIBLE] line you're on?

DAVID How do I know which value on the line?

PERREAULT:

AUDIENCE: I guess-- oh, I see. Each of them represents a very different pole.

DAVID Yeah. No, each of these-- each of these points-- these two points represent one gain of K_I , these two points

PERREAULT: represent some bigger value of K_I , and so forth. So I can pick a value of K_I that would keep my poles over here.

All right. So-- I wouldn't do this in the real world, but maybe I could say, well, if I'm really careful, maybe I can get away with an integral controller on this thing and it'll do the right thing. So now let's start thinking about doing this design. Maybe I'll use an integral controller because it's the easiest thing I can think of. Let's start talking about my open-loop poles and 0's.

Firstly, we got to remember that this transfer function varies with operating condition. So if I'm at a different value of load resistance, I'm going to have a different 0 location and I'm going to have a different-- this doesn't show very clearly that the pole locations will differ, except that it's in this term, which actually relates to the damping.

So if I calculate this transfer function's poles and 0's for different values of R at the nominal input resistance, this is what you're going to see. The poles in the right half 0-- right half-plane, these are the-- I'm sorry, the 0's in the right half-plane, the location of the 0 moves with different values of the load resistance. If I look at the locations of the poles, they move with the load resistance.

And what will happen is as R gets bigger, this thing gets less and less damped. The open-loop poles are moving towards the $j\omega$ axis, which is right here. Why does it do that?

Well, if you think about back to the original converter, what's happening? Essentially, L and C are oscillating. They slosh energy back and forth dynamically. Not just directly because they're not directly connected, but modulated through the duty ratio. They're still trading energy back and forth, and that's the source of oscillation. That's why it's second-order and can be underdamped.

The resistor damps that. So it's some funny, modulated version of the LC oscillation with a parallel resistor. So basically, it's like a parallel resonant tank, and the resistor is what's damping that thing. If you make the R big, the damping is small, and it oscillates more. So it shouldn't be surprising that the location of the open loop poles gets closer to the $j\omega$ axis and less damped as R gets bigger. So I might really have a trouble if my R went to 100, but I'm only making it go to 10.

Noticing that also, things vary with input voltage because if I think about input voltage, D prime here is changing, So that's also changing the pole locations. In fact, very directly. And in fact, you can see that here. So this is a plot of what happens when I vary both R and input voltage, and hence, duty ratio. And all I did was I clunked the pole locations into MATLAB and said, plot it for me. So any questions about that?

So any compensator I'm going to have to build, the typical thing you do-- you could build a variable compensator that did different things for different operating points, but the most typical thing to do in a DC converter is to say, I know my plant characteristics vary. Let me design one compensator that'll work for any location on that. That's the simplest thing you can do. And so let's see what would happen if we tried to do that.

What I'm going to do is let me put this system in a closed loop. I'm going to use a compensator-- I'm going to use this compensator. I'm going to use an integral controller. And I'm going to use my plant dynamics there. And then for different values of K_I, I'm going to take the closed loop, I'm going to find the closed-loop pole locations, then I'm just going to plot them. And I'm just going to do that physically in MATLAB. Here we go.

So the open-loop pole locations are the ends here. And this is for a load resistance of 2 ohms. This is one end of my load range. This is the heavy load end of my load range. We can see that for K_I , because I have this integral gain of 5, I'm going to have a pole here that's due to the compensator, and then the poles that's due to the plant are going to be here. And the right half-plane 0 is out here.

So that system ought to at least be stable. And I plotted this for a bunch of different values, increasing values of K_I . So this root locus that I was showing you, I've calculated for a bunch of different values of K_I . So you can see what will happen. If I make K_I too big, I will get instability.

Now that's for one end of my load range, but keep in mind, the starting place of my poles-- my open-loop pole's dependent upon my load resistance. If I started with the other end of my load range-- 10 ohms, light load, this is what it would look like. For K_I equals 5, I have a dominant pole here. This is the closed-loop pole location. And this is-- I have a pair of super lightly-damped poles here.

And I could then say, all right, let me calculate the closed-loop pole locations. That's how-- once I have the feedback loop, if I kick the output, that's how it'll oscillate. With my control system in place. And I could calculate that for every possible value of R and input voltage that I might deal with. And if you do that, this is what you get.

So you could argue that no matter what here, it's at least stable. All the poles are in the left half-plane. This is the $j\omega$ axis on the right here. And I'm going to have a dominant first-order pole, and then I'm going to have some lightly-damped higher-order poles. Any questions about that? So I could at least say I built a stable system that won't blow up-- hopefully-- for any of the loads or input voltages I'm going to put on this thing.

OK. How stable is it? And this is where me and your average control textbook break paths in terms of what we're happy with. A typical measure for those who do control are things either like the phase margin you have or the gain margin you have on your loop. So what you do is you look at the loop gain and you plot that out and you can look at the open-loop loop gain with a controller.

This is what it looks like. So this would tell you at 10 ohms, which is the most lightly damped point in that, this says we have a gain margin of 17 dB and a phase margin of almost 90 degrees. So any control guy will tell you that's a super stable system. You're not you're not going to be unstable, no worries. And they're right. The output is not going to blow up. Whether it's acceptable or not is another question, and that's what you've got to think about.

And so let's look at this. Let's look at some load step responses. So what we do is we have the closed-loop control, and then we pick the load, and we see how does the voltage and current respond? In this case, we're stepping up the-- we're stepping down the resistance to see a load step from light load to heavy load, and then from heavy load to light load. So here's one from one input end to the other.

And we can see that here's the voltage response, and here's the current response here. And it pretty much settles out. I think this is a simulation thing out here. So you could say, that's not too badly damped. I don't mind that so much. It's all right. It's not great, but it's stable. But that's going from light load to heavy load.

What happens-- oh, and I should say, that's the actual switch circuit model. What is the average model say will happen? There it is. So the average model-- the thing I want to point out is even though that's a pretty big load step, the average model does a really good job of capturing what goes on. So you can really get behind averaging as a way to get a simple look at what your system is doing. And in fact, if you calculated what the oscillation dynamics looks like, the small signal model still looks pretty darn good.

All right. Well, what about if I go in the other direction? I'm going to take-- I have heavy load, a lot of load on the converter. Now I'm going to step the load light. Well, here we go, boom. The first thing you'll notice, by the way, is when the load went away, what happens to the output voltage? It rockets up. Why does it rocket up?

Because there was all this current in the inductor and the load resistance-- the load went away. And so all that energy that was in the inductor ends up in the output capacitor. And that output voltage swings way up. And there ain't a lot you can do about that. The energy in the inductor has to go somewhere, and it goes into the output voltage.

So in the real world, if I didn't want to see a ginormous spike on my output voltage-- like a 10-volt spike in my output voltage, I better throw a bigger output capacitor on there, or use a smaller inductor.

So when you're thinking in your design projects about how to size LNC and you want some cut-off, using a 1-Henry inductor and a picofarad load filter capacitor is probably a bad idea because you've got to eat the energy in that inductor, in a load transient. So word to the wise, you got to you got to size those very wisely and make sure you have enough energy storage in your capacitor that he's not going to be too unhappy to absorb the energy from the inductor in a transient.

But, more to the point, what else goes on here? Here's what happens to the energy in the inductor. The inductor current-- this is the inductor current here, it goes bam! It goes down to 0. Basically, the output voltage gets big. The converter says, oh my god, turn-- what can I do? The most I can do is turn my duty ratio to 0. And I turn my duty ratio to 0, and all that energy in the inductor ends up in the capacitor. And then it starts to switch again.

But you notice here, things are clipped at 0. In this time period, this converter is actually in discontinuous conduction mode. So actually, that average model isn't quite right anymore.

Now if we simulate this, we actually do pretty well, but the simulation-- actually, we stick it in a diode to keep the - we clamped it so-- in the simulation so that it wouldn't have the inductor current go negative. Why? Because this has a switch in a diode in it. Now if you had two switches, this average model's perfect, it works all the time. It is still clipped. In our system, we can never have duty ratios greater than 1 or less than 0, so there is some clipping that goes on.

But, the other thing I wanted to point out-- and this is related to what we said earlier. I showed you we have the gain margin of, what, 17 db and the phase margin was 90 degrees, that's perfectly stable. Anybody would say that's very stable. But look at the oscillation in the output voltage here. Like, to my mind-- I'm sorry, this is the oscillation in the current, but there's also an oscillation in the voltage. That's very ringy. I wouldn't like that at all in a real converter, especially if I'm feeding some sensitive load, and when I set the load, the output voltage is going erargh.

And what you can-- and where does this oscillation come from? This oscillation comes from these lightly-damped poles. The dominant pole here makes it really stable. It's not going to-- the output's not going to walk off and go to infinity or something, so that's why the model says it's stable. Nonetheless, these lightly-damped poles still cause oscillation in the output.

Now I chose on purpose an extreme example that makes it look pretty bad. What you can see sometimes is that you'll have something that's very, very stable, but there'll be some really lightly-damped pole, maybe not even a little bit down, but way down. And what you often see when you do that, if you leave this lightly-damped pole hanging there, is that your output voltage will look fine on a macro scale, but then you zoom in on it, and you find it's wobbling by 10 or 20 or 30 millivolts on top of your 24-volt output, and then you have to ask the question, is that tolerable?

And maybe I don't care about 30 millivolts of ripple at 20 kilohertz because there's some lightly-damped pole pair. I could live with that. But there's other applications where they tell you, oh, by the way, you've got to have your output voltage rock solid, and some lightly-damped pole pair that's going to get excited by noise and just wobble around is probably not acceptable.

So in addition to making your control loops stable, I would argue you've got to think about how precise does it have to track the output voltage? How much will it wiggle when something happens to it? And is that acceptable?

And in power supplies, very often your job is to make the world clean, solid voltage no matter what. And in those cases, you've got to think very carefully about how you want to deal with the closed-loop poles. Maybe you want to make sure that no poles are super lightly damped even if they're pretty far down and not contributing that much to stability, maybe they'll make your output response unpleasant. So, that was me talking a lot. Are there any questions about those examples?

AUDIENCE: What devices would require a more stable output voltage?

DAVID PERREAULT: If you were building a power supply to feed a sensitive instrumentation amplifier, for example, because if you have some instrumentation amplifier that's picking up some tiny signal, I don't know, for an EKG or whatever it is, that amplifier has some power supply rejection ratio, meaning the output of the amplifier will give you the response to the input, but there's a tiny little response to the power supply variation. And that little response may be bad.

So if your power supply is wobbling around, a little bit of that gets into the output, and so you don't like that. If what I'm feeding is a resistor, maybe you don't care. But so it depends on the application, but a lot of the time, if it's a power supply for instrumentation amplifier or for an RF transmitter or for a receiver or something like that, you've got to keep it very clean. Yeah, Jack?

AUDIENCE: Is there a phase margin for the transfer function from D to V.

DAVID PERREAULT: Yeah. That's true, yes.

AUDIENCE: That's why you showed us a load step that that phase margin was different.

DAVID That's right. But on the other hand, from a small signal point of view, all the pole locations are-- once I have a linear system, the pole locations are the same for any input to any-- for any input. That's a good question, but the phase margin--

AUDIENCE: I guess [INAUDIBLE].

DAVID I'd have to think-- I'd have to think about whether that's true. But it's also true that you would still see that-- if I did a tiny little load step, you'd still see that kind of thing. And if I did a tiny reference step and came back to what I'm treating as the input, you'd still see that.

AUDIENCE: Yeah. I guess that's [INAUDIBLE] phase margin that's [INAUDIBLE] is pretty [INAUDIBLE].

DAVID What's that?

PERREAULT:

AUDIENCE: [INAUDIBLE] 90 degrees is pretty [INAUDIBLE].

DAVID Right, but the only thing that's telling you is what the dominant pole is doing. It is very damped. It will-- the majority of the transition will reflect that dominant pole. So the reason this looks damped is because it's dominated by this guy down here. Let me put on-- right? But that doesn't mean there's not some small oscillations.

So the small oscillations aren't going to make your circuit blow up. And that's why they say, oh, it's great, it's very, very stable. And it's true. The problem is, if these wiggling around, even if it's very, very small, might be unacceptable as a response at the output, that's the point I wanted to make to make. Does that answer your question? OK.

So this is just to give you an idea of the way people-- first of all, the process people use to get a model, and then the way they think about what you might use that model for, understanding if some of you have had control system design and some of you hadn't. You notice the question I dodged is, well, if this is a bad compensator and this is a bad compensator, what's a good compensator?

And the answer to that is, there's a lot of things you can do. One of them-- if I just wanted to use something like an integral controller or a more fancy version that's not just an integral controller, but something that was a good duty ratio control scheme, one thing I could do is change my plant so that it had friendlier dynamics.

I already said, if you don't want the voltage overshoot, I'd use a bigger capacitor. I can also do things to the actual converter to make it friendlier for my dynamics. And what we saw as one example was that-- well, what's really happening is this resistor is damping the LC oscillation modulated through the switching. But the problem is, when R gets big, it's not well-damped.

Well, what happens if I walk up to this thing and I say, you know what? I will throw a damping leg on the output of this. I'll put another resistor in here, R_D . And here, I have see C big. When we talk about filter design, you'll also see these see these kind of damping legs. And you've seen it in snubbers, too, right?

The idea is, at frequencies near the oscillation point with L and C, if C looks like a short circuit, then RD is in parallel with R and damps this. And then if R gets big or small, it always has RD in parallel with it. So if this was 2 ohms and I was happy with the response I saw for 2-ohm damping, then maybe this damping leg added on to my converter would satisfy me. And I also have some more big capacitance here to help me absorb larger voltage transients.

So one thing I can do if I don't like my dynamics, and you might not, either because the overshoots are too big or because it's not as well-damped as you want, as you can do things like add damping legs or other things to the plant-- to the actual converter or change the component values so that it is better damped, or that it has friendlier response. Nothing wrong with that. Do it all the time. The upside is it works.

The downside is, now I've got to go back out and buy an extra capacitor and a damping resistor-- it's more expensive. There's some things you can't change. I mean, some things you must change in the plant. If I want to reduce that amount, that voltage is going to swing up when the load steps down, I better go get more capacitance. There's no way around it because there's nowhere else for the energy to go unless I'm going to somehow dump it. So sometimes you have to change things. But you don't tend to add hardware if you can do it in control.

And the way I might think about doing something like that, I'd certainly use a bigger capacitor. The other thing I would probably do in this kind of converter is a different control scheme. You notice that in this control scheme, we reduced everything down to a transfer function from duty ratio to voltage and I just ignored what the inductor current was doing.

Well, it turns out that control guy will tell you, well, you should be doing full-state feedback. In other words, don't just feedback V, but feedback I, and use that knowledge to improve the control so that you basically don't make your poles worse, but you make them better.

And it turns out that just by doing better control, by looking at the inductor current and the voltage, I can make a better controller that will lead to a much nicer-- more nicely damped response. And we will talk about that-- that technique is called current mode control or current program control, and that will be the topic of next class. Have a great day.