[SQUEAKING] [RUSTLING] [CLICKING]

DAVID PERREAULT: We started talking about modeling of control systems. And really, what we want to do is have some model for the dynamics of a DC/DC converter. And the example converter that we've been using is a boost converter. So here's our example boost converter. I might have some input voltage that I'll call u. I have some switch that I'm modulating with q of t, diode, and then some output voltage. I'll call the output voltage V.

And in this case, maybe I'll say, this is one state variable. V is one state variable, which is the capacitor voltage. And the other state variable in the system is the inductor current. And state variables are those variables that carry information about their history. You can't instantly change a capacitor voltage without an infinite current. You can't instantly change an inductor current without an infinite voltage.

So if you know what the inductor current or capacitor voltage are at one time, you know what it's going to do because that's what carries history around with it in any system. So it's the energy storage elements. And we said we figured out, for this boost converter, if I define q of t is the modulating function, it switches between 0 and 1 for q of t.

And then if I have q prime of t is equal to 1 minus q of t, if I'm modulating this at a constant duty ratio, q of t goes on and off, and I let this thing settle into periodic steady state, what we get is NPSS. We get V, the output voltage, the DC output voltage is going to be the DC input voltage divided by 1 minus D.

And I'm using capital letters here to indicate these are all periodic steady state values. But we said, if I were, for example, to suddenly change the load resistance, then even if I was operating at fixed duty ratio, eventually I would have to go through some transient because, to supply a lower resistance, I need more current, more power. I can have more current from the input. Something has to change. It has to go through some transient.

What we have here on the scope-- or what we have here is actually a boost converter run at fixed duty ratio. And what I'm going to do is I'm going to set this thing to single sequence. And what I have is-- on the yellow channel, I have the inductor current, just directly coupled. And on the purple channel, what I have is the output voltage, but AC coupled.

So I'm just going to look at the deviations of the ripple in the AC voltage. And all I'm going to do is press this button. And that's going to add a load onto the converter. So it's going to make the load current higher or the load resistance smaller. And here we go.

Oh, It didn't capture. Let's see. I don't know why that is. Oh, it would help if I actually coupled the output voltage. It seems to have-- oh, set this, set the minimum. I don't know why. All right, let me try that again. There we go.

So what we can see here is that the current goes from some value and settles out to some higher value. The voltage goes through some oscillation, and then settles out to some later voltage. And the truth of the matter is we're not actually seeing the full transient here. Let me do this, try this one more time.

Ah, there we see it a little bit better, except that actually what I triggered on was the step down in current. In any event, what we see is some kind of second order response. Superimposed on top of it is a ton of switching. This is going through thousands and thousands of switching cycles.

And what we would like to do is some model that captures this low frequency variation, but isn't worrying about the instantaneous switching of what's going on in steady state. So I'm going to hide this, put this up. So last time, we said, well, how can we capture the low frequency content in i, l, and v without worrying about the high frequency content?

We defined an operator that we called the local average operator. And this is x-bar of t is equal to 1 over t, the integral from t to minus capital T up to t of x of tau d tau. So basically, if I apply this local average to any waveform, a voltage, or a current or something like that, I'm going to look back over one capital period T, one switching cycle, and take the average over that switching cycle, so that if I have a bunch of up and down within a switching cycle, that gets filtered out.

But if over time that voltage or current moves up and down, it captures the slow frequency up and down while suppressing the switching ripple and its harmonics. And we said this operator has a bunch of properties that are interesting. One is it's linear. Two is it's time-invariant. That means this thing is an LTI operator. That means we can write a transfer function for it. It's just a linear filter. And we looked at that transfer function last time. And it, indeed, passes low frequency components and suppresses the switching frequency.

A third property that we discussed is the so-called differentiation property, which says that if I took dx some signal dx of t dt and then took its local average, that ends up being equal to d/dt of x-bar of t. So taking the local average of a derivative is the same thing as the derivative of the local average.

The last property is we said we can approximate x of t y of t local average of the product as being approximately equal to x-bar of t times y-bar of t if one of x and y have, first of all, small ripple, meaning it's not changing full scale within a cycle, and, two, slow variation, meaning it only changes gradually from cycle to cycle.

And actually, you can relax this former case in some instances. But we saw that it doesn't always work. But if one of x and y was, say, an inductor current or a capacitor voltage, where we do expect there to be small ripple on top of a big DC, we can do that.

Last time, we used this And we applied this to a circuit. So we said, we could take the boost converter circuit, and play some games, and come up with a new circuit in which all of the circuit variables were averaged. And so it gets rid of all the switching, but it carries the low frequency behavior of the circuit. And that's one way to do things. And people do it, the so-called direct circuit averaging, all the time.

What I'd like to do today is to show you another way we can use this local average operator. And it's essentially the same idea. It's just implemented in a different way. Some people like one method of thinking about it. Some people like the other method of thinking about it. Frankly, when I'm doing a new converter, I usually do the method I'm going to show you today, but I like them equally. It's just the question of what's more convenient to do.

And the technique I'm going to show you today is called state space averaging. And here's the idea. Suppose I wanted to know the-- suppose I wanted to know the dynamics of this circuit about how it behaved. Well, I could say, write a set of differential equations for that circuit that would describe its behavior in time. And keep in mind that this circuit really has two subcircuits. When q of t is equal to 1, the switch is on. So I could think about this circuit as being this.

So the switch is on. This is a linear circuit that models that behavior. If q of t is equal to 0, I can then say the circuit really looks like this.

So I could then write a differential equation, a linear differential equation, because these are all linear elements plus an input, when it's in this state. And I could write a different differential equation when it's in this state. And if I just use the right set of differential equations in each substate, I can get the behavior.

So let me just write this down. And I'm going to show that we can use this switching function, q of t, where q of t varies between 1 and 0-- I'm sorry, this is 1. We can use this switching function to build up a differential equation that captures our system.

So here's what I would argue we could do. We can say, what does the diL/dt look like? diL/dt would simply be the voltage across the inductor. In this state, It would just be equal to 1 over L times u. And that's true when q of t is equal to 1. So I can multiply it by q of t.

When q of t is equal to 0 and 1 minus q of t is equal to 1, diL/dt is equal to u minus v. So I could say, this is plus u minus-- this is plus 1 over L u minus v times 1 minus q of t. So basically, when q of t is equal to 1, this term, diL/dt, is 1 over L u. So it matches the top circuit. And this term is gone because 1 minus q of t is 0. And this term just disappears.

When q of t is equal to 0, this term disappears. And this factor becomes 1. And I just have 1 over L u minus v. So whether I'm in this switch state or this switch state, this differential equation captures the dynamics of that inductor current. Does that make sense to everybody?

I can do the same thing for the voltage. I can get dv/dt is equal to 1 over C times-- this would be minus v over r times that, minus 1 over RC times q of t plus 1 over C i sub L minus v over R 1 minus q of t.

So in fact, if I came up and I told you the starting state of voltage and current, inductor voltage, capacitor voltage, and inductor current, and I told you the q of t over time, I could basically numerically integrate this pair of differential equations and predict the trajectory, including all of the switching and everything else. Does that make sense to everybody? Any questions about that?

So this is actually a simulation model of this converter in continuous conduction mode, including dynamics. This is a pair of differential equations, switched differential equations that capture it. Well, I could take my averaging operator and apply it to the differential equations as well. I could apply it to each side.

All I would be doing is walking up and saying, OK, let me local average each side. So I'll put a local average on the left side of this. And then the local average of these sums of terms is just the sum of the local averages. So I can do this. And I can do this. And by the way, the scaling constants can come out of the local averages. So that's OK. And I can do this. And I can do this. And then I can do this, local average of all this.

So all I did was I walked up and I applied that operator to each side of the equation. Well, now, I can rewrite this. And I can apply some of the properties of the local average. First of all, the local average of the derivative is just the derivative of the local averages. So I can write this as di L-bar/dt Then this would be 1 over L times the local average of u.

And let me rearrange this, actually. I have u times q plus u times 1 minus q. So this would just be 1 over L times u-bar. Then I have minus 1 over L or plus-- minus 1 over L the local average of v times 1 minus q. Does that make sense, everybody? So all I did was rearrange this. I put up the q and 1 minus q terms so that the u term doesn't have q in it anymore. And then it's just minus 1 over Lv times 1 minus q. And I took the local average of that.

I could then do the same thing with the bottom equation. I can get dv-bar/dt, because I'm using the differentiation property, is equal to-- the v over RC term again goes up. So this is minus 1 over RC times v-bar plus 1 over c times the local average of iL times 1 minus q. Everybody down with that?

Now, let me define q-prime is 1 minus q. Let me define d is equal to the local average of q of t. If I'm modulating along in a cycle, in any given cycle, the local average of q of t is the duty ratio in that cycle, neglecting small shifts in time. But this is everything slowly varying with respect to a cycle.

That would mean that d-prime, which is the complement of the duty ratio, is equal to 1 minus q local average. Or it's equal to q-prime local average. So if I apply that, and then what I'm going to do is I'm going to say, OK, I can make this approximately equal to 1 over Lu-bar.

Then I can make this approximately equal. If I know that xy-bar is approximately equal to x-bar y-bar if one of them has slow ripple and small variation-- well, iL does and v does because this is a capacitor voltage and inductor current. So this would be minus 1 over L v-bar d-prime.

And this is approximately minus 1 over RC v-bar plus 1 over c iL-bar d-prime. So this becomes a set of partial differential equations. Maybe I should just rewrite this. This is the diL-bar/dt. And this is equal to dv-bar/dt.

Now, I have a set of state space averaged equations. And notice that there's no switching left in here. There's no q of t bopping up and down and turning switches on and off. This is just in terms of whatever my duty ratio is going to be tells me how this evolves. So if I knew how the commanded duty ratio varied, I would know how the local averages of i and l, iL, and v would vary. Any questions about that?

So how would I use this thing? I could come up and say, all right, I had my converter here. What I really wanted to do is say, OK, I have some reference voltage, v-ref. And I compared that to-- in this case, I'm going to say it's a local average version of the output voltage, v-bar. And I'm going to generate some v-error bar. OK I'm going to run that into some compensator.

That could be gc of s. Think of this as, for example, a lead lag controller, or a PID controller, or something like that, something that has high gain so that it's going to amplify the error. And out of this comes some signal that I'm going to call d, the duty ratio I want.

Now, that might be some voltage representing the duty ratio or current representing the duty ratio on an actual implementation. But from a theoretical perspective, the output of this compensator is just the duty ratio I want. I need to stick this into the plant. That's what a control person would call it. We'd call it the converter.

And of course, it has another input, which is, in this case, u, or u-bar. And out of this thing, out of this model would come iL-bar and v-bar. And so what I needed was a dynamic model for this thing. Well, guess what? Here is a differential equation model for the behavior of the converter. And I could just stick this inside of here. And this is saying, if I put in d, what happens to iL and v? And this model is my dynamics. Any questions about that? No switching involved.

Now, I should have noted, last time, we drew a circuit model. We didn't put a pair state space equation in here. We actually had a circuit model. What was that circuit model from last time? It really looked like this. It had u-bar iL-bar. It had a voltage source here, which was d-prime v-bar and a dependent current source, which was d-prime iL-bar. And then it had v-bar here.

If you write the differential equation for this circuit, this nonlinear circuit that we got by circuit averaging, it will give you exactly this set of differential equations. diL/dt is u. L diL/dt is u-bar minus d-prime v-bar. That's exactly this equation.

So I can describe it with circuit averaging. I can describe it with state space averaging. This just gives me the equation for it so I can go back and forth between the two as I want. They both work. There are other circuit models that represent the same differential equations, of course. You can do things different ways. But the underlying dynamics will always be the same, if you apply averaging. Any questions about that?

Now, what is the headache of this approach? Or it's not a headache. It's nothing to do with the approach. I shouldn't say-- this is a fundamental-- this is a fundamental property of the system, is that this is a nonlinear circuit and this is a nonlinear differential equations. Why?

Because I have a product between my state variable, my output voltage, and my control input, the duty ratio. And likewise, I have a-- or 1 minus the duty ratio. And here, I have a product between my other state variable and 1 minus the duty ratio. And that makes this system nonlinear.

Why do I care? Why is that a pain in the butt? Well, what would I like to do? If I go back to standard control theory, what I'd really like is perhaps some kind of transfer function that says, if I change the duty ratio, what does the output voltage do? Because if I'm going to write everything in terms of a transfer function, I can apply linear control theory and life is good.

How do I usually get the transfer function? I take the Laplace transform of this stuff, which converts this differential equation into an algebraic equation. diL/dt becomes just S iL, where S is just the Laplace variable. But this is kind of a pain in the ass in this equation. Why?

Taking the Laplace of a product. Anybody recall what you have to do when you take the product with a Laplace? You end up with a convolution. And that's not very pretty. So I can't just write a transfer function for this. I can't boil this thing down into a transfer function because this Laplace of the product is a problem.

So the question is, what can I do? And the answer is, it's nonlinear and it has nonlinear dynamics, meaning it truly is not linear. That's the way it is. But we could find out what it does for small variations. In fact, in my original example, I just showed you the output voltage. And I did this little bump, but it was around 0. That's because I was AC coupling it. I was showing you the deviation in the output voltage from its steady state value.

And what I can go do is I can say, OK, if I believe my converter is going to be operating in periodic steady state somewhere, and then I put in a small perturbation away from that steady state-- in other words, maybe I step my reference voltage just a tiny bit-- it should respond. And the dynamics are nonlinear. But for small variations, they're pretty close to linear.

So what I might do is I might go and say, all right, suppose that all of my variables in my system, x of t, was equal to some DC value, sometimes called the bias or the steady state, plus some AC variation, small variation.

And then I could express, well, the small variations are simply the total variable minus-- it's the deviation of the total variable from the steady state operating point. And what I can go do is I can express all my variables that way. I can say, OK, u is equal to some DC value plus any variation I put on the input voltage. v is equal to some DC voltage plus some small variation, where this is the deviation away from-- this is the AC-coupled version, essentially. iL is equal to some DC iL plus some iL tilde, some small variation, et cetera.

And what can we say mathematically? Suppose I had some nonlinear differential equation, dx/dt, where x is some variable or some vector of variables-- and this is a vector of equations-- is equal to some function of the state variables, the inputs, and could be a function of time, too.

So this is some arbitrarily nonlinear differential equation. It could be a vector differential equation, like this has two state variables. So x could be a vector of state variables, in general. And I could have a vector of inputs, r. And suppose I had this and I had some steady state operating point. I have f of x and capital R and t is equal to 0. This is an operating point. This is a nonlinear diffy q, or system of differential equations.

Why is this an operating point? Because what I'm saying is, if I put f of capital x and capital r in here, I get 0, which means dx/dt is 0, which means it's not going anywhere. It sits there. So this equation is precisely defining some operating point where it's just going to sit, some periodic steady state value where it's going to sit.

And what we can do is we can linearize this guy. What do I mean by that? Let me just do the simplest 1D version of this idea. Suppose I have one dimension, x, and I have f of x. And this has some function, some nonlinear-- this is not a straight line-- function.

Well, this is f of x. Notice that there's some place where f of x goes through 0. And I could call that x. And so the variation, dx/dt, is exactly defined by this function near here. But maybe near this operating point, I could come and say, you know what? Let me just approximate him with a tangent line.

And I will approximate f of x right near this operating point with a straight line. And then I can look at, if I have some variable x, the distance here to some other point. This distance I might call x-tilde. The deviation between x and its operating point is just the variation. So I can worry about this differential equation, variations about this point.

And how am I going to approximate that? I could approximate that as dx/dt, which I could write as capital d capital X/dt plus dx-tilde/dt dt. So x is capital-- this is x. This is capital X plus x-tilde. This is just 0 because this is the derivative of a constant.

It's going to be equal to-- well, dx/dt was f of x R and t. I'm just going to substitute x plus x-tilde et cetera in here. What I'm going to get is f of x R and t plus df/dx evaluated at x and R for-- evaluated times x minus capital X, which is actually just times x-tilde, plus df/dR evaluated at x and R times R-tilde.

So all I'm doing is instead of using f of x, I'm doing a Fourier-- I'm doing a Taylor expansion of this thing and looking at just the derivatives at that point. What's f of x R and t? This is just 0 because I just told you this thing was 0. And so I just get df/dx and df/dR times x-tilde and R-tilde. Does that make sense to everybody?

And all I'm really doing is I'm linearizing-- I'm approximating this nonlinear curve. If I have my deviations are small enough, I can't really tell the difference between the curvy f of x and the straight line orange approximation. That make sense to everybody? To do this, I do require that f is differentiable near here. But our systems are, so life is good. Any questions about that idea generally?

So that's messy to think about. But actually, it's quite simple in the real world. In the real world, what am I going to do? In the real world, all I'm going to do is I'm going to come up to this equation and substitute in these expansions and factor it. Life is easy.

So let me do that. So what I have is I have diL-bar/dt. And I said, iL-- oh, I should have said these are the local averages because I'm working with a local averages. diL-bar/dt is equal to-- is equal to d/dt of iL plus diL-tilde/dt. And we said this was 0 because iL is a constant.

And that's going to be equal to what? That's going to be equal to-- all I've got to do is substitute the stuff in there-- 1 over L u plus u-tilde minus 1 over L capital V plus v-tilde times-- it's 1 minus d-- I'm sorry, it's d-prime, which is 1 minus d. That's 1 minus d minus d-tilde. And that's the top equation.

And I get dv-bar/dt, which is equal to d/dt of capital V, which we said was 0, plus dV-tilde/dt, which is equal to minus 1 over RC v plus v-tilde plus 1 over c times iL plus iL-tilde times 1 minus d minus d-tilde.

I apologize for all the algebra. But guess what? You can't do this without it. So that's life. All right, so what can I do here? Let me just refactor this. And here's how I'm going to factor it. I'm going to put all of the terms that are just operating point capital letters together. Then I'm going to have mixed ones. And then I'm going to have all terms that are products of tildes.

So let me just rewrite this for you. So what I get is diL-tilde/dt is equal to 1 over L u minus 1 over L times d-prime v, so 1 over L times u minus d-prime v. Then I'm going to get mixed terms. I'm going to get plus 1 over L u-tilde minus 1 over L v d-tilde. Oops, it's going to be plus 1 over L v d-tilde minus 1 over L d-prime, which is 1 minus capital D, v-tilde. And then I'm going to get plus 1 over L v-tilde d-tilde.

Last equation, and then life will get easier. I won't have to write any equations. It's dv-tilde/dt is going to be equal to minus 1 over RC minus-- I'm sorry, it's going to be-- it's going to be 1 over C times minus v over R plus iL times 1 minus d-- that's the all-caps-- term minus 1 over RC v-tilde plus 1 over C d-prime over C iL-tilde plus or minus iL over C d-tilde, and then minus 1 over C iL-tilde d-tilde.

I apologize for all that math. But let me get to why I'm telling you this. Let's look at this factor. And I wrote u here. I should have been very careful to write capital U. These are all DC periodic steady state values. This is the DC input voltage. This is the DC output voltage at an operating point. Or should I say, in periodic steady state local averages?

What is 1 minus d times v? Well, if I came back over to my boost converter here, 1 minus d times v is just u. So this whole thing is 0. Likewise, this is the output current, capital V over R, capital IL times-- this should have been d-prime is simply the current going into the output from the inductor. This is also 0, if you worked it out. So basically, all capital terms in this thing, if you're dealing with the periodic steady state operating points of a converter, are going to drop out and go away. If they don't, you did your math wrong.

What about these guys over here? This is the small variation in the output voltage times the small variation in the duty ratio. I'm going to argue that this is a small thing and this is a small thing. This is small squared. I'm going to pretend he's approximately 0. Likewise, this is small times small. I'm going to say he's approximately 0.

Basically, these terms are precisely the curvature away in this nonlinear function from the straight line. And I'm saying, if I get close enough, I can ignore that curvature. If I make my variations away from that point small enough, I can't tell the difference, which leaves me with an equation in terms of diL-tilde/dt, which is in terms of two-- basically, it's in terms of this.

I'm going to have constants. This is 1 over L. This is 1 over L times capital V. So capital V is my periodic steady state value. That's just a constant, as far as the variations are concerned. This is minus 1 over L. This is another constant. d-prime is my steady state duty ratio. And d-tilde is my variation away from that. Or d-prime is 1 minus my steady state duty ratio. And I can do the same thing with these other terms. These, as far as this equation is concerned, are all constants. Does that make sense to everybody?

So I would argue then that this beast, this beastie-- if I take away all the stuff-- I throw away all the stuff, all I have is small signal variables times constants added up. So this is also a system of differential equations. But as far as a given periodic steady state operating condition, it's a linear constant coefficient differential equations. This is the friendly kind of differential equations, the ones we love.

Well, why do I like that? Because I could take that set of differential equations and come back here. And instead of worrying about v-error, I could worry about v-error variation. I could worry about the variation in duty cycle, the variation in my inductor current away from-- this variations away from the steady state operating point, the variation in voltage away from my steady state operating point. And this would be feeding back v-tilde.

And suddenly, I stick these differential equations in here. And these are linear. And suddenly, my life became easy because now I can apply all my linear control theory to it. Questions about that?

So how might I do this in practice? A lot of things I can do. And I should say, by the way, there's nothing that prevents me from doing the original nonlinear stuff and doing a nonlinear-- having a nonlinear plan and considering Lyapunov controllers or anything I want to consider. It's just complicated and perhaps unnecessarily so.

By doing this linearization trick around an operating point, I get a set of linear constant coefficient differential equations that then give me nice behavior. The only constraint, the only fly in this ointment is the differential equations, the coefficients in it depend on the operating point.

So I've got to say, geez, I'm operating at d-prime is 0.5. My steady state is at 50%. It's for this boost conversion ratio. And then I'll get the dynamics for that boost conversion ratio, for small deviations away from that. But if I suddenly change duty ratio to 75% and d-prime became 0.25, the equations are different and the dynamics are different,

So the fact that this is nonlinear, it means, at different operating points, the dynamics vary. Well, you're stuck with that. It's a nonlinear system. That's what it does. But if I'm interested in-- jeez, I kicked it a little bit, I changed my duty ratio a little, or I stepped my load, when I get close to that second operating point, I can predict what the dynamics are very exactly.

And by the way, even though this is nonlinear, it turns out that it's pretty good. The dynamics are quite nicely captured. So how might I deal with this? Well, I could do this. What would I do? I could turn this whole thing into to a transfer function, for example.

Like over here, what I would really like is some transfer function, perhaps H of s, where H of s goes from d-tilde to v-tilde, for example. It's a transfer function from variations in duty ratio to variations in output voltage if I kept my input voltage constant, for example.

Well, how would I do that? I'd just take the Laplace transform of this stuff, all my i sub L-tildes would become S i sub L-tildes. This would become S v-tilde. And then I could look at this thing, and I could go in here, and I could multiply this by S. I could then substitute my S iL-tilde in here and iL-tiles would go away. And I'd get everything that relates d-tilde and v-tilde. I'll just write down that solution for you because I'm running out of time. But if you do that, here's what you get. First of all, any questions about the process I just mentioned?

So what I would get is something like this. I would get H of s is equal to v-tilde over d-tilde, which is equal to minus s times iL over C plus v d-prime over LC divided by s squared plus 1 over RC s plus d-prime squared over LC.

So this is a second order differential-- this is a second order transfer function. It has two poles, S squared in the denominator. So this characteristic equation is second order. Well, we expect that. It's got to be second order. Why? Because what are my energy storage elements in my converter?

My energy storage elements in my converter are one inductor and one capacitor. So it is second order. Great, I expect that. These have two left half plane zeros. I can tell because all the coefficients are positive on the characteristic equation. It has one right half plane zero. We'll come back to talking about that.

So that's great. I get a behavior I'm expecting. Notice that the pole locations move around depending upon duty ratio. So if I'm sitting here and I'm coming from 5 to 10 volts DC, d-prime would be 0.5. And I would get one set of pole locations. If I was coming from 5 to 20 volts, I'd have a different value of d-prime. And I'd get a different set of pole locations.

So the poles do move around at different operating points. But that's life. I can deal with that. How accurate is this? Well, I plugged in-- I won't bother writing it on the board. I plugged in values for an example. And it's precisely the example I showed you. It's precisely the example I showed you-- precisely the example I showed you last class, in fact, this transient.

This is we get a step, we step-- in this case, we step the load. So I calculate everything for the new, the final value of what the inductor current would be. And I just plug into that equation. And what do you get for that equation? You get a 0 at 35 kiloradians per second and poles at minus 4,000 plus or minus J 16, 279 radians per second. That gives an oscillation period that ought to be about 0.4 milliseconds.

Well, if you look at this, that's pretty good. That's about what we got, 1.6 to 2. So even though this is a pretty large step-- it's a factor of 2 in load-- I can predict the transient response, the average transient response, with this nice little linearized model. That means, if I put this thing in a feedback controller, I ought to be able to build a compensator that will control the closed loop dynamics and give me the response I want.

So that is state space averaging in its application, and then linearization. So I'm going to average my circuit. And I get a nonlinear average model. Then I'm going to linearize the differential equations. And that gives me a nice linear constant coefficient model that I can model with a transfer function, at least for a given operating point. Any questions before we wrap up? Yeah?

STUDENT: When you said [INAUDIBLE], were you showing us some simulation or--

DAVID PERREAULT: Oh, what's this? Where did I get that number?

STUDENT: Yeah.

DAVID PERREAULT: What I did was-- well, it's not on the graph. What I did was I plugged into this. And I got the values of S. And the imaginary component of the pole locations is the damped natural frequency. And that should correspond to the oscillation period of a transient response.

So if I took an RLC circuit and I put a step on it, I would see a transient with the pole locations indicated. And all I did was said, what's the pole locations? And how does that match that transient? Yeah?

STUDENT: When did you choose to linearize throughout the starting or ending operating points?

DAVID PERREAULT: Well, OK, so I've got to linearize around the ending operating point because I'm asking how does it vary about where I'm going. And it's going to settle down to that place. Asking what the dynamics are about where I'm leaving from doesn't tell me very much because I'm not going to be there. I'm going away from there. Any last questions? So we'll take up control of power converters again next class.