[SQUEAKING]

[RUSTLING]

[CLICKING]

DAVID PERREAULT: --control. And we said that we saw a few phenomena. So if I took, for example, a boost converter as my canonical example, I have V, the output voltage, as one of my state variables because the capacitor C and IL is the other state variable.

And what we really control here is q of t or the local average of q of t, which we call the duty ratio. And we said we developed linearized average models for this thing with the notion that if I had some reference voltage, and I compare it to the output voltage, and I went into a compensator GC of S, and I use that to generate duty ratio and then I had my converter, which might have some IL, it might have some V out, and I then compare that output to my reference, I get an error voltage.

And then I use that error voltage to generate the duty ratio I want. I stick that into my converter. And we generated some linearized average models for this thing. So I might think about the variations in all these things to linearize it. And we came up with a model for our converter, H of S, which basically converts between variations in duty ratio to variations in output voltage, and we design a converter.

And when we did that, we saw a couple of things. First, we saw that the dynamics were nonlinear. That's why we had to do this linearization thing. Now, are they always nonlinear in a converter? No. But many, many converters are. And in this case, why is that nonlinearity there? The nonlinearity is there because what we're seeing is our second order responses are some kind of oscillations between the inductor and the capacitor and damped by the load resistor.

And those oscillations are mediated through the switching or the modulation of the switches, which varies with time. And more to the point, the steady state value of that varies with DC input to output voltage and hence the dynamics vary. Because the modulation is mediated by the switching, we get nonlinear behavior. And what that means is that the dynamics, the small signal dynamics, physically vary with different values of, say, input voltage and output voltage.

So if I have a different value of input voltage and same value of output voltage, I should expect my dynamics to vary. So one, we saw nonlinearity which arises. The other thing we saw was a pole-zero map.

If I looked at H of S, what's going on in here, the one thing we saw was a right-half-plane zero in our transfer function. And we said that that right-half-plane zero comes from the underlying way in which the converter processes energy. It's not a function of the math. It's because when I change what I'm doing, the dynamics it goes through is reflective of what a right-half-plane zero does. So I'm kind of stuck with this right-half-plane zero.

And then the other thing we saw was that the pole locations, which are due to the L and C and how they interact with the resistance, so it's a second-order system, tend to be very lightly damped, like this, something like that. And the amount of damping depends upon the load resistance R. So I'm really relying on this load resistance to damp the oscillations, or that makes it hard to design a controller. Lightly damp poles and a right-half-plane zero mean I'm not a very happy camper in terms of designing a good controller.

So that's where we came to. We said, well, what can I do about this situation? Because this is a hard thing to deal with. You can deal with it. You can design a GC of S, a compensator that will deal with that, but it's not very pretty. One thing you can do is by choosing L and C appropriately, that helps you set dynamics the way you'd like them.

The other thing you can do is change the plant. So we said, oh, it's this resistor that's damping the poles. Maybe I could go add RD and CD such that CD is a short circuit near the corner frequency where the L and C are oscillating, so RD serves to damp it, and RD doesn't vary. So by picking additions to my plant, by changing my converter, I can help make things better damped and maybe get better pole locations here.

I'd have another pole location, but maybe it would be a low-order overdamped pole location. So that's one thing I can do. I don't like to do that because then I got to buy hardware. I got to buy a big capacitor and that kind of thing. The question is, what else can I do to make a better power converter controller? And that's what we're going to talk about today.

And the basic idea is this. I have a converter that has two state variables. Its state variables are IL and VC. And yes, it has as its input the input voltage. And so what I've been doing is taking-- between these two, I've been feeding back the output voltage.

What a control person would tell you is if you're not happy with your dynamics and you have a controllable system, feedback all of your states. Do full-state feedback. If you feedback all of your states, you can place your dynamics much better. And we're going to do a version of that called current mode control.

The general concept originated in the 1960s. The version I'm going to tell you about came about in the late 1970s. And there's lots of versions of this idea. I'm telling you about what's probably the most popular one in DC-to-DC converters if you go out to buy one or to design one. A lot of the control chips use this technique.

And here's the idea. For the converter, we said the input was-- well, ultimately the input is q of t. It's what's modulating the converter. Maybe I can say, OK, let me feed back the inductor current and compare that to a reference. And I'm going to call this reference IP. So this is a current reference.

And based on the comparison of IL with IP in a current controller, I'm going to generate the switching function. And then what I'm going to do is I'm going to feed back the voltage and compare it to a reference and put that into a compensator GC of S. And I'm going to use GC of S to generate the current reference.

So this is something called minor loop control. We're going to have a fast inner loop to control the current and then a slow outer loop to control the voltage. But we're going to feed back both current and voltage and thereby get better dynamics. That's the basic idea. Any questions about that?

AUDIENCE: What does fast and slow mean?

DAVID PERREAULT: Let me show you the scheme, and then you'll see what I mean by fast and slow. Remember, when we were talking about our models, for our average models up there, they're kind of good on time periods that are slow compared to half the switching frequency. So maybe you can go to 0.2 of the switching frequency. There's games you can play. But pretty much, you're looking at the average value. So you don't want your controller, GC of S in the top example, to be changing your command very quickly, certainly not within a cycle.

In this scheme, what we're going to do-- that will still be true of this outer loop control. The inner loop control is actually going to run on a cycle-by-cycle basis. And that has a lot of follow-on benefits. So what we're going to do is we're going to come up with a scheme where we turn on the switch every cycle. And then I'm going to turn off the switch based on what the inductor current is doing in a particular way. And I'm going to show you the particular way we're going to examine.

And so that's going to give us control over what the inductor current is doing. And then I'm going to use my outer loop to control that command to the inductor current. So the crudest way people like to think about this is I'm putting the inductor current inside a loop. So now I just get a commandable current. And then I sort of have a first-order system. Now, it's not quite that good. That's an oversimplification. I still have two state variables. It's still a second-order system, but it's not a bad way to think about it.

All right, so how does this scheme work? Let me show you the basis of the scheme. It looks something like this. What I'm going to do is I'm going to create-- I'm going to have a current reference. Here's my current reference.

Let me call this I peak. And this is-- within a switching cycle, this is more or less a constant defined by this compensator, but it's going to move up and down as my outer loop tells it to. Then I'm going to create a slight variation off that, and we'll explain why a little bit later. But suppose this is my switching period T. I'm going to create a slope with some slope. I'm going to call this minus MC.

This is some slope-- MC is a positive number-- which is a ramp that I subtract off of my command, and I'm going to do that every cycle. So that's T. Here's 2T and so forth. So I'm going to do this every cycle. Now, here's my inductor current. So I'm going to sense IL.

And what I'm going to do is at the beginning of each cycle, I'm going to turn on-- this is IL-- I'm going to turn on the inductor current. I'm going to turn on the switch, I'm sorry. And the inductor current is going to rise with some slope M1. So the beginning of the cycle, I come-- q of t turns high, inductor current ramps up. And what I'm going to do is I'm going to look at that inductor current, and I'm going to compare the inductor current to Ip minus this slope. I'm going to compare it to this pink signal.

And when I hit here, when the inductor current exceeds the I peak minus this compensating ramp slope, I'm going to turn off the switch and let the switch current fall like this. So this is, say, DT. The switch turns on at the beginning of the cycle. But it turns off-- instead of me setting a duty ratio explicitly, I'm going to implicitly determine the duty ratio based on this peak current.

And then, beginning of the next cycle, I do the same thing, I turn the switch on, it rises, it falls. And I get the same kind of general behavior. It's just that instead of saying turn the do switch off after some time, DT, I'm going to say turn it off implicitly when it hits some current that I command. Any questions about that?

Why do I like that? Well, one thing we're going to see is that we're going to get better control dynamics. And I should say, by the way, this one, this slope, is minus M2. And why am I talking about M1, M2, minus MC? That's because this same picture applies to any kind of converter.

For the boost, M1 is just equal to the input voltage over L, and M2 is equal to the output voltage minus the input voltage over L. But I could have drawn the same picture if it were a buck converter or the same picture if it were some other converter.

So why am I doing this? Well, a couple of reasons. One, we're going to see it gives us much nicer dynamics. Why? Because I'm basically explicitly controlling the inductor current. And then the other thing that's kind of nice is if I did duty ratio control, what is the inductor current doing? I don't know. If I put a small resistance on it, it's just going to keep cranking up the current until it does what it wants.

Here, I'm telling the switch to turn off if the current gets too big. So this gives me automatic current control. My currents can't get out of control. Now, in a boost converter, that's not quite true because if I put a short circuit on this, there's nothing I'm going to do because the switch won't control the inductor current. But as long as the output voltage is bigger than the input voltage, by having this peak current reference, I can always limit the amount of current the converter is processing and hence prevent it from overheating, for example.

All right, so essentially, this technique gives me cycle-by-cycle current control over the switch-- over the inductor current. Now, I have to sense the inductor current because I'm controlling IL here. Sometimes I do literally sense the inductor current. Sometimes, keep in mind, you're only looking at the current when the switch is on. So sometimes people will come and put a little resistor right here or use the switch resistance itself and sense the current in the switch. It gives you the same thing as far as this comparison is concerned. Any questions about that? Yeah.

AUDIENCE: How would you adjust the peak current level if your load changes?

DAVID PERREAULT: Ah, that's a good question. So what's going to happen? I'm going to do the same kind of deal. If the load voltage-- well, let me go back to the right controller. If the load voltage is too low, too low means below the reference, the error voltage gets bigger, this compensator, this outer loop compensator amplifies it. Maybe it integrates the error. It will make the peak command get bigger. If this peak command moves up, so will the switching time. And then the inductor will have more current, and eventually the inductor will feed the output, and it will bring the output voltage up.

So the ultimate output voltage is really controlled by this, looking at the error between the output voltage what I wanted, just like I did before. I just have an intermediate value where it controls the current at which the inductor is going to run to get there. And if you think about it, the inductor current times, the input voltage is the power I'm drawing. So it's in some indirect way controlling the amount of power that I'm going to draw from the input.

All right, so one thing it gives us is a cycle-by-cycle current limiting. The other thing it's going to give us is better dynamics. The question is, how do we analyze the dynamics of this thing? What did I do before? The thing I did before was we started-- we averaged the average circuit, either the circuit or the state space model, and came up with a duty-- and then we linearized it, and we came up with a conversion between the duty ratio and the local average inductor current and local average output voltage.

And we could try to do that for this scheme. It's a little bit complicated, and it turns out it's mathematically very difficult. You kind of walk into a wall of math if you try to play this game directly. But there are clever ways to get around it. I've already got a model for how duty ratio relates to the local averages of the output voltages and the local average of the inductor current.

If I can relate the peak current to the duty ratio, then I'm kind of in business. I can just kind of swap it out and create a new model. And that's what people do. It's much more simple, or it's more or less the only way to go about it that I know of that works out nicely.

And so here's the idea. This picture I drew is actually a periodic steady state picture. I used capital D here. I'm assuming it does something every cycle, DT. Let's just draw a slightly better picture for purposes of analysis. OK, so here we go.

Here's my waveform. Here's my peak current, I peak. And then I have my so-called compensating ramp, which falls with some slope minus Mc. And now I have my inductor current. Let's suppose the inductor current starts here. It rises up with slope M1, hits the intersection, that's DT. And then it falls with some slope M2 or minus M2.

But I'm actually treating this thing dynamically. I'm not necessarily saying this end point is the same as the starting point. So this is a dynamic picture of what happens over one cycle. And what I want to do is I want to relate somehow this duty ratio in this cycle to what the local average of the inductor current is.

And I'm not going to do that by local averaging. What I'm going to do is I'm going to make a geometric approximation. I'm going to say, OK, let me look over this window and say the average of this waveform, which is basically the average here, the average of this current is the local average.

So I'm saying IL of t is approximately equal to the average of IL over this window. Strictly speaking, this is equal to IL of capital T, if this is time T. And it's not quite the same thing anywhere else, but it turns out it's pretty-- it's not a bad geometric approximation.

Now, if I wanted to be really theoretically pure about this, I might actually-- instead of looking over this window, which is centered at 0.5T, I might look at a window centered at dT and look around here and it extend into the different switching cycles. But it turns out for practical purposes, you really more or less get the same-- for most cases, you more or less get the same thing. So the approximation I'm going to make is good enough given all the other approximations in the system.

So I'm just going to average it over this window. But do people see what the basic idea I'm going to use? I'm going to approximate IL bar as the average of this waveform. And then that's a function of the duty ratio in that cycle. And then I'm going to swap that into my local average duty ratio equation. Actually, I'm going to linearize it first, but that's neither here nor there. So let's do that.

What is my approximation to IL bar? Well, let's see. The peak of this waveform here is I peak minus Mc dT. That would be, if I wanted-- if I just use this value, that would be basically if the local-- if the waveform-- if the orange waveform looked like this green waveform, that would be exactly the local average in the cycle.

But what I really need to do is kind of subtract off the average of this area and the average of this area. And then I'll have the average underneath the inductor current curve. So if I do that, then I'm going to get, OK, I've got to subtract off the area of this triangle here. So that would be minus 1/2 the base, which is dT times the height, which is dT times M1.

And likewise-- and by the way, I better divide that by T because I'm taking an average. So this is 1 over 2T. And I better subtract off the area of this triangle divided by capital T. So this would be minus 1 over 2T, 1 minus d quantity squared T squared M2.

So if I rewrite that, IL bar is approximately equal to I peak minus Mc dT minus 1/2 d squared TM1 minus 1/2 1 minus d squared TM2. And this is the substitution I'm going to make. Or it's the basis of the substitution I'm going to make. Any questions about that?

So I've got a geometric approximation relating the duty ratio in the cycle, the peak current in the cycle, and the local average in the cycle. So let me linearize this. I'm going to linearize this because we're going to substitute it into our linearized equations, and I'm going to solve for d tilde.

And if I do that, the result you get is-- oh, and then the last thing I guess I should do is let me substitute in, in this case, u-- I'm sorry, M1 is equal to u over L. And M2 is equal to V minus u over L.

And if I make that substitution, I do it, what I get for my boost converter is d tilde is equal to 1 over McT I peak tilde minus IL tilde minus D squared minus D prime squared over 2LMC, u tilde minus d prime squared over 2LMC, V tilde.

So this is a linearized version of this guy for the boost converter. If you wanted to do a buck boost converter, you'd need different expressions for M1 and M2, and you'd do the same process, and you'd linearize it. You'd get a different expression. Any questions about that?

OK, so what's the point of doing this? The point of doing this is I can now substitute in for d tilde in my existing equations. And remember, d tilde was my control variable in all my existing state space averaged linearized equations. If I substituted for d tilde, I get rid of d tilde. And all this other stuff is already in the equations except for I peak tilde. Does that make sense to everybody?

So I'm going to now get a new set of linearized state space equations, except instead of detailed in the equations, I'm going to have a control variable Ip tilde. That's what I needed here. So basically, I can take this out and swap it for an H2 of S that's basically equal to V tilde over I peak tilde.

So I basically got my whole inner loop going. And now if I have Ip tilde and V tilde, I can now giant design my outer loop. What was the point of doing this substitution and getting my new H2 of S? If I look at H2 of S, I look at that transfer function, what am I going to find?

I'm still going to find the right-half-plane zero. I can't get rid of that. That's related to how the converter behaves. If I step I peak tilde, I'm still going to see-- I'm still going to see funny behavior in the output voltage before the output voltage response. And my right-half-plane zero doesn't go away.

The beautiful thing is, however, what I will tend to see is instead of two really lightly damped poles, I'll tend to see a dominant pole and then a much higher frequency pole. Why? Because essentially the way people like to think of it is this inner loop is making the current do what you want. And hence, the dominant pole is more or less governed by how the current from that inductor feeds into the output capacitor and charges up and down the output capacitor.

So it's not that I have a first-order system. There is a second pole here. It's just dominated by this low frequency pole. And that's a much easier thing to control. Any questions about that?

So right now we can go design GC of S. I still have a right-half-plane zero, but it's a much easier plan to compensate, and I can put my dynamics much more nicely than I could before. And I get free current limiting in the process.

And in fact, if you're designing most kinds of PWM-DC converters, most of the control chips on the market are going to use current mode control. This version that I'm showing you is called peak current mode control because you're controlling the peak of the inductor current. And that's the most common kind. There's also valley control. And there's a few different versions of this general idea, but this is the most popular one.

So that's one trick. Any questions up to now?

AUDIENCE: Is there an easy way to understand why [INAUDIBLE]?

DAVID PERREAULT: Excellent question. I'll cover that one next. Any other questions? OK, yeah, why did I even bother with this ramping MC signal? Why not just use I peak? And in some cases, you can. But here's the problem with doing that.

In some cases, if I do that, this is the situation I can find myself in. I'm going along, and suppose I just have I peak, and I get rid of the compensating ramp. So I'm at the beginning of my cycle. And here's my inductor current. I turn on my switch, I go up, and I hit the top. And I trip, so this is my duty ratio in this cycle. And I turn off. And then I turn off till the end of the cycle.

And here's the end of the cycle. Here's capital T. And now I turn on my inductor current on, and he ramps up. He hits here. Here's my duty ratio in this cycle. And then he turns off, and I hit the end of the cycle, 2T. And you can see that I don't settle down into periodic steady state operation at the switching period.

In this example, it repeats every other cycle. In some examples, it won't settle down at all. It'll just bungle around chaotically. I mean, literally chaotically. Now, is that a problem? Well, I mean, at the highest order, it's not. I'm still controlling the peak current. My output loop will still regulate and everything, but it is problematic from other perspectives.

One perspective is my ripple in this case is a heck of a lot bigger than if I were at some fixed intermediate duty ratio. So I get much bigger ripple. The control is jittery because it's bouncing around depending upon the operating condition. I have lower frequency content in my output. It's bigger ripple, and it's lower frequencies that can make a mess of things. So people don't like this behavior.

This is sometimes called subharmonic oscillation. Well, this behavior where it's rippling at a subharmonic is called a subharmonic oscillation. But as I mentioned, in some cases, it can actually be chaotic. Generally, this issue is called the ripple instability. This ripple doesn't settle down into periodic steady state behavior.

And what people figured out-- and the first paper I knew about this was out of the Caltech group in the late '70s, but other people thought of it too-- was to stick this compensating ramp in, which if you put enough compensating ramp in makes that behavior go away, and it'll settle down just like this case, to a kind of a fixed ripple. OK.

AUDIENCE: Did we cover how the MC is calculated?

DAVID PERREAULT: No, but I will in about 30 seconds. Excellent question. By the way, you might say, why did they think to do this? Well, A, they were clever. But B, think about it-- in most power converters, you have a ramp generator to set your clock frequency. That's how you get your oscillator. I'm going to turn on every cycle.

How do I know what that is? I generate a ramp that it hits the threshold, and that sets my clock period. So they had this ramp signal sticking around. So they just basically took a little section of that ramp signal and added or subtracted it from I peak, and they found that fixed their problem. So that's how they got to the idea, I think. Or that's why it's-- I shouldn't say that's how they got the idea. That's why it's a natural trick to play because you have the signal around at your disposal.

But the question is, how big do I have to make this MC to get this behavior instead of this behavior? And let's analyze that. So let me take-- let me analyze-- let me first start and say, OK, what would my periodic steady state behavior be? Here is I peak. Here is my compensating ramp. Here's period T.

So if my compensating ramp looks like this-- this is a slope minus MC-- maybe I could say my periodic steady state behavior looks like this. I start off somewhere. I ramp up with slope M1. Then I ramp down with slope minus M2.

And I end up exactly where I started at the beginning of the cycle. This is periodic steady state behavior. So this is what I'm going to call DT. This is the behavior-- this yellow waveform is the behavior I'm looking for. Let's suppose my inductor at the beginning of this cycle started off a little bit high. So suppose he started off here. He has some error, some deviation from periodic steady state that I'm going to call delta I sub n.

So at the beginning of the cycle, he starts off a little bit high. What would happen? Well, he would ramp up at the same slope because the output voltage is the same. So this would still be M1. He would hit the trip point here at something I'll call d star T. It's not the steady state value.

And then he would fall at minus M2, like this. And then he would end up with some error at the end of the cycle or the beginning of the next cycle that I'll call delta I n plus 1. So if he starts off here, he's going to end off here. And I have a different error from periodic steady state at the beginning of the next cycle. Does that make sense?

So let's see if we can quantify what that looks like. And again, this was done by this group, this Middlebrook group. And by the way, it's in a paper that we've posted on Canvas for you. And this behavior is also described in a book chapter that we've also posted on Canvas for you. Here's the idea.

Let's quantify what delta I sub n is. I could write delta I sub n in terms of this time difference, D minus d star T times, let's see-- I go D minus d star T. I go up slope M1 and then backwards, so up slope MC.

So this is going to be D1 times d1 times is M1 plus MC. So I could write this delta In because these waveforms are parallel in terms of M1 plus MC times d minus d star T. Does that make sense to everybody?

I could write this difference, delta In plus 1 also in terms of D minus d star t. Why? Because these two lines are parallel. And what I've really got is this is M1-- oops, I'm sorry-- M1 plus MC. I could write this--

I could write this in terms of D minus d star T times MC minus M2. They're both related to this time difference. It's just I'm kind of running up and down different slopes. Does that make sense to everybody?

So if I take the ratio of these two things, I could write this as delta In plus 1 over delta In. That's what I get in one cycle is just going to be this ratio, which I can write as the bottom over the top here, which is minus M2 minus MC. Actually, I'll write it this way. It's MC minus M2 divided by MC plus M1.

OK, that make sense to everybody? That means that if I started off with some error delta I0, I could write delta In as that ratio, which I could write as minus-- and I'm going to write this the way they often write it in papers-- minus M2 minus MC over M1 plus MC to the n-th power.

Let me make sure I've gotten that right. Yeah, M1 plus MC. So can anybody tell me what the condition on this factor is to make an initial error delta I0 go to 0? Yeah, this ratio-- what I require-- this suggests that I need the magnitude of minus M2 minus MC over M1 plus MC to be less than 1.

This is the-- I've got a discrete time model for what's going to have-- and it's going to decay if this term is less than 1. So what that usually says is that if I make MC bigger, the numerator gets smaller. And actually, the denominator gets bigger. So a bigger MC will make this decay more quickly. And in fact, if I made MC exactly equal to M2, it would be gone in one cycle.

Does that make sense to everybody? So what we'll do then, and we can look this up. It depends on the converter, what M1 and M2 are in the converter. But we will pick this compensating ramp slope to be big enough-- MC is a positive number to be big enough that my ripple will decay away in a reasonable number of cycles. Maybe I make this magnitude less than 0.5 or 0.7 or something like that, so it'll settle down.

Any questions about that? Now, I should be a little bit careful in what I'm saying. And by the way, what this means is if I had MC equal to 0, this would be minus M2 or M1 for the boost converter. What that means is this magnitude would be greater than 1 if it was d greater than 0.5. So if you took a boost converter and you didn't put any slope compensation-- if you're only running at 25% duty ratio, fine, you won't have any ripple instability because it'll settle down.

If you get near 50% duty ratio or above, it'll start doing this subharmonic oscillation or ripple instability stuff. This is sometimes called the compensating ramp or the stabilizing ramp for that reason.

Now, why don't I want to make MC really big? Well, keep in mind, if we came back to this original model, notice that MC appears in this substitution. My local average linearized dynamics of my converter are not independent of MC.

And when I told you, oh, by the way, I get nicely-- I get a really low frequency pole and then a high frequency pole both on the real axis, that's when MC is small. As I make MC bigger and bigger, what happens is these closed-- these open loop poles move. And if I make MC too big, I get undamped poles again.

So basically, the higher this value of MC, the closer the converter transfer functions and current mode control look like they do for duty ratio control or voltage mode control. So what do we do? We would usually do this kind of ripple instability analysis and pick an MC that was big enough to make this settle down and make me happy but not much bigger.

And then I'd use that value-- I'd plunk it into here in this average model and find out where my pole locations are and go on from there. I could also make-- obviously, I could make my compensating ramp slope depend upon M1 or M2 in some way. People don't usually go to that range. If you've got to have your converter operating in a really wide range, maybe you start to play those games. But generally, you're just picking some slope and setting it with a resistor or something. Questions?

Let me summarize. We didn't like the fact that we had these very lightly damped polls. So we did full state feedback, and we feed back current on this inner loop. And then we use the outer loop to set the current reference and do the outer loop. And that's the most common way people do things.

I'll show you a demonstration of doing this that Mansi was kind enough to set up for us. And all I'm doing is this is again our demo boost converter here. And it's actually using current mode control. And the notion is what we have it set up is that we can turn on and off this compensating ramp. So here we go. Let me power up the converter.

And that's pretty nice. I got more or less periodic steady state operation. There's actually a tiny bit of jitter in the duty ratio here, but it's basically switching once per cycle and nicely settled down. What happens if I turn off the compensating ramp? Every cycle is doing something different, and I can single sequence this so you can see one shot.

What you can see is sometimes it has a long duty ratio. Sometimes it has a short duty ratio. It kind of pips around like this. It's not unstable in the sense that my converter is going to blow up. But what it does mean is that the ripple is all over the place. And I should say that this is a function-- if I change the voltage, this tends to-- if I have low-- if I have a low boosting ratio, it goes away, even without a compensating ramp. But the higher I go, because of that ripple instability, I get some really kind of funky behavior.

If I throw in my compensating ramp-- I've just done something bad to my converter. There, I get nice steady state operation.

So that is current mode control for you. I wanted to tell you about it because in the context of DC-to-DC converters, it's probably the most widely used control technique. And to analyze it, you do everything we talked about so far, and you take one more step further to get a current mode model. Any final questions before we wrap up? So that was our very brief sequence on modeling and control of power converters. We will take up a new topic next class. Have a great day.