[SQUEAKING]

[RUSTLING]

[CLICKING]

DAVID PERREAULT: So I wanted to keep talking about DC-to-DC converters today and keep talking about the impact of ripple and operating modes of DC-to-DC converters.

So we said-- we did some analysis last time, and what we found is that under reasonable approximations, the inductor current waveforms in a PWM DC-to-AC converter tend to look like some DC value plus some triangular ripple. And the same thing for the capacitor voltage waveforms.

And we said that we could represent this kind of waveform-- sort of a triangle wave with an offset-- as having some DC value plus triangular ripple, where we could find the-- we could define the average value and then the peak value. And because of the properties of a triangle wave, the peak is right in-- the average is right in the middle, between the two peaks. So you swing delta x peak to peak over to above the average and delta x peak to peak over to below the average.

And so we can define a ripple ratio for such a waveform, a notion of a percentage ripple, which we define in this class as half of the peak-to-peak swing, or the mean-to-peak swing normalized to the average value. And one of the nice things about that is we can then-- if I knew the average value, from our calculations, I can actually find the peak value as being the average times 1 plus the ripple ratio.

If we apply that-- we also talked about the device stresses in a direct converter, such as a buck or a boost converter, and an indirect converter, such as the inverting buck-boost converter.

And we had written it as-- for a direct converter, the peak switch voltage, the peak diode voltage, which is also the peak capacitor voltage, is this maximum of V1 and V2 times 1 plus the ripple ratio. This is the average DC value we see, in terms of the DC value. Plus, the ripple on the capacitor voltage increases the peak voltage, as we see, a little bit.

And likewise, we can talk about the switch current and the diode current peak values, which are related to the maximum of the average input and output currents. So for a buck converter, the output current is the bigger one, and the input voltage is the bigger one. But we see peak stresses and voltage and current that are related to the maximum of the input and output DC waveforms, and the ripple.

Then we looked, also, at this indirect converter. And we found the difference is instead of seeing the maximum of the input and output, we see the sum of the input and output currents and voltages. So this tells us, no matter what, for a given conversion ratio of V1V2, I1I2, we're always going to have higher stresses in an indirect converter compared to a direct converter.

And that makes sense because of how we're processing energy through the converter. The indirect converter-- you're basically storing all of the energy in the middle of the converter. And in the direct converter, you're some of it gets to flow right through the converter, and you can take advantage of that. So this is just a quick review of the conclusions we came to last time. Any questions about this so far?

The other thing we did was then we took that and we said, OK, well, I know that my inductor current, for example, has some triangular ripple on it, So this IL looks like some average value and some ripple,

And we could calculate-- if I wanted to limit my peak-to-peak inductor current ripple below some value, because that's my input ripple current, or my peak-to-peak output voltage ripple below a certain value, I can come up with minimum capacitor and inductor values to do that that are related to the conversion ratio. And I expressed it here in terms of the allowed ripple ratio of the designs.

So now we have a means to start sizing inductors and capacitors, We can't practically use nearly infinite ones because they get too big and expensive. If you have too small a one, you get too much ripple, and bad things happen. And so the question is, how do you pick those values?

And one of the things that you use to pick those values is some of notion of an allowable ripple, which we could express in terms of peak-to-peak ripple, or we could express it in terms of ripple ratios. Typically, in this example, you're not usually allowed very much ripple ratio on the capacitor because if you want a DC output voltage, you can't have 10% ripple in your output voltage. That would usually be bad. So maybe that's a percent or something. It depends on your specifications.

Inductor current ripple-- maybe a bit more is acceptable. Maybe you could have 25% ripple in that current, especially if you add on another capacitor to help you filter it back here.

So there's the notion of sizing. And we'll talk about-- this is one way you can start to pick your inductor and capacitor values based on ripple. There are other factors, too. We said size, So you can figure the biggest value you can get before it gets too big is probably a good thing.

But also, in terms of transient performance-- the dynamic performance of your converter-- we'll talk about that aspect a little later. But that can also come into to how you pick the values. And the reason I mentioned that is because in your design project, you're going to have to pick values, and you're going to have to think about these different aspects.

The other thing we said was, OK, if I have some minimum values based on ripple, what does that say about the energy storage requirement of my inductor and my capacitor?

And the reason we talk about energy storage requirement is because we ought to-- it's reasonable to expect that the size of your component is probably monotonic with its energy storage capability. It's going to get bigger if you have to store more energy, usually. There's a complex trade-off to design actual-- especially magnetic components. But nonetheless, it's reasonable to expect that.

And what we can see from this expression is that if I looked at the energy stored in the capacitor, it gets bigger if I have smaller ripple ratio. This factor here tends to get big if I make ripple ratio very small. It gets bigger proportional to the power of my converter needs to handle. That's kind of reasonable. If I want a kilowatt converter, it's going to be physically-- have more energy storage than a 100-watt converter.

It depends on the conversion ratio. And in this case, for a boost converter, a bigger conversion ratio is a bigger duty cycle. So the more voltage conversion I'm doing, the more energy storage I'm going to need to do it. And the last thing, which weighs heavily on a lot of designers' minds, is-- we expressed it last time in terms of the switching period. Here, I'm expressing it in terms of the switching frequency. The energy storage gets smaller if I go to a higher switching frequency.

So usually, I'd like to switch as fast as I can, commensurate with other constraints, because my energy storage got smaller. My components will then hopefully get smaller. Life is good. So there's been, always, a lot of effort to go as high as the frequency you can until we run into some other limits, and we'll talk about those limits-- loss being one of them.

And I will say that everything we just said about the capacitor we could say in the inductor. And I should note that in this set of lecture notes, I think I wrote 1 minus D here, and this should be D, which is correct in the last set of lecture notes. But nonetheless, the same ideas pertain. Smaller ripple means larger energy storage requirement. Larger conversion ratio means larger energy storage requirement. And higher frequency means lower energy storage requirements. Any questions about that?

OK. I'd like to keep talking about this notion of conversion and the ripple you see in your converters because we haven't actually gotten to all possible cases yet. And let's just go think about this boost converter. And what we'll do is let's think about the inductor current, and I'll think about the inductor voltage V sub L here.

So if I'm switching my converter, I'm going to have some switching signal which is driving-- which is telling when this gate is turning on and off. So I'll have some switching function. I'm just doing this. So it's high for some time, DT, and then it's low for some time, and then it's high for some time after a period.

And then I can say, OK, well what is my inductor voltage do during that time? Well, my inductor voltage is positive at V1 when I turn the switch on. So it looks like this at value V1. And then when I turn the switch off, I have V1 minus V2 across it, and I get some negative voltage.

And what we know is that in periodic steady state, this area must equal this area. And this ends up being V1 minus V2. And if I were to plot the current in the inductor, what would I see over the same time period? I have some DC value. In the inductor current, which is I1. The DC input current here is the DC current coming from the source. And that's my average value.

And what we'll have is in the first part of the cycle, I store energy in the inductor, and I get this positive triangular ripple. And in the second part of the cycle I get negative slope, and I have my periodic steady state waveform like this. And this is delta IL, peak to peak. So that's what our waveforms look like.

So let's think about this in terms of ripple ratio. OK, so what do I know? Delta IL, peak to peak is equal-- delta I is V over L delta T. So V is V1 divided by L times DT.

So this is an expression of how much I rise. I'm applying V1 for time DT, so the delta I is V1 over L DT.

What is the input current? I1. OK, well, what's I1 going to be? Well, the DC output current is going to be V2 divided by whatever this resistor is. So if I wanted I2, I would have V2 divided by R. But this is I2. How does I2 relate to I1? I2 is going to be I1 over 1 minus D. because I1 is bigger than I2.

That's just-- this factor is I2. And then this gives me-- 1 over 1 minus D gives me the conversion to I2. Any question about that?

So then if I were to write the ripple ratio, the ripple ratio in the inductor in this example is simply equal to delta IL peak to peak over 2 divided by IL. And this is also equal to IL. And so this is going to give me V1DT over L. So this is going to be 2L-- I sub L is V2 R times 1 minus D.

And I can rewrite this in terms of-- I can get rid of V1 and V2 here. Why? Because V2 is going to be V1 over 1 minus D.

So I could rewrite this as-- V1 over 1 minus D is V2. Or V2 times 1 minus D is v one. So I could rewrite this as D times 1 minus D quantity squared TR over 2L.

And what does this tell me? Well, this says how-- essentially, I'm asking, what's this ripple compared to the average value of my current?

And we can see that the ripple ratio goes up as what? It goes up when either R goes up or L goes down. Those are two things that will make the ripple ratio bigger.

And that makes sense. Making L smaller is going to make the ripple ratio bigger because I have a higher ripple on the inductor. Any questions about that? Yeah.

AUDIENCE: The DC values of the currents are being-- ignoring losses and that.

DAVID PERREAULT: The DC values of the current ignoring loss?

AUDIENCE: Loss.

DAVID PERREAULT: Yeah, there's no loss. This is a perfect converter. Everything's lossless. You throw in loss, and things happen, Right now, I'm treating everything as being perfect. So let's start to think about the impact of changing resistance and inductance. So let me replot this. So let me plot IL again.

And we said, OK, maybe I have some average value, and I have some ripple that does this, and so forth. And this is at some resistance value.

What would happen if I increase my resistance value? Well, up here, I see if I increase my resistance value, I increase the ripple ratio. But how does it change the IL waveforms? This is IL. But if I used a bigger resistance value, what might happen is it might look like this. I will have some lower current to supply the same voltage at the output. If I doubled my resistance, I need half the current to get to the same output voltage. If my output current halves, my input current halves. And that means my inductor current, on average, halves.

So OK-- I reduce the average inductor current, but V1 and V2 have stayed the same. And if V1 and V2 have stayed the same, the slope on the ripple has stayed the same. So my inductor current, at a higher resistance, might look like this. This is parallel. This is parallel, and so forth. And I get this.

So this is IL. This is IL at R small. This is IL at RL medium. What would happen if I kept increasing my resistance value? Well, all we saw is the waveform doesn't-- the waveform shape just doesn't change. The average just moves down. The ripple is the same, but the DC value is changing because I need more or less average current to support the output and more or less average power draw to support the output.

So eventually, what's going to happen is we're going to get something that-- we hit another mode of operation. We're going to get something where it does this. I'm going to turn my switch on, and then it's going to come dowr and then stop.

So here, this is supposed to be the same DT, and here is the same T. But I'm going to get this time period where the current comes down to 0. Now, why am I not drawing the current as going negative? I can't really get a negative current through the diode, So once the inductor current ramps down to zero, the diode is going to turn off, whereas the diode was conducting before.

So I could build a converter where this was an active switch and allowed bidirectional current, and people do that. But if I'm going to build it with a diode, I'm going to get this alternative state. So in this alternative state, what do I have? Here, I have the switch on. Here, I have the diode on. And in this time period, everything is off.

So we haven't considered this mode of operation, We always assumed it was either the switch or the diode that was conducting. Now I'm telling you, hey, if I make R big-- and how big is big? We'll get into that.

But if I make RL big enough, the average current goes down to the fact that my current ripple is going to make me hit 0 current, and I'm going to get this new funny state. Does that make sense to everybody? So this is what happens with variations, with resistance.

What happens with variations in inductance? Well, slightly different. Maybe what I have is this. I will have, again, some value, IL.

Actually, maybe make it a little smaller. I have some average value IL that I'm generating. And if I have some big inductance-- maybe my ripple is small, so it does this. So this is IL at L big. I have a big inductance in there. The ripple is relatively small.

If I got a smaller inductor-- I turned off my converter; I pull out one inductor; I put in a new inductor that's half the value. What would happen? Well, the ripple would increase, And I would get something that did this. The delta IL peak to peak is inversely proportional to L. So my ripple current would-- so it's a terrible drawing. My ripple current would increase.

If I turn the converter off and then I put in even a smaller inductor, what would happen? If I get a really small inductor-- I pull out the smallest inductor I can find in my kit, maybe I get a waveform that does this. It starts at 0, comes down, hits 0, and then stays at 0.

So I'm going to get, again, a case where I have the diode on. And then I have-- I'm sorry, I have the switch on. In this case-- and this is only pertaining to this bottom picture-- the diode's on. And then I get a time period where both are off. And this is a different mode of operation. Any questions about that?

This mode of operation is known as discontinuous conduction mode. It's often abbreviated, DCM.

Now, when we say, discontinuous, there's nothing discontinuous about the inductor current, obviously. The inductor current is always a continuous waveform. But what it's what's meant by that is it's not continuously flowing positive and negative. It has this third state, where both devices are off.

And keep in mind that that's all the possible cases because I cannot have a case where both the switch and the diode are on and the voltage is positive. That can't exist. So we exercise all the things that can happen. Any questions about that?

So when does this happen? Well, the ripple ratio-- if I were right at the edge of discontinuous conduction, you could imagine a situation where I pick some intermediate value of resistance, for example, in this case, where it does exactly this. Oops, I didn't draw that very well.

I'm sitting right at the edge, where, yes, my waveforms are continuous. My diode never turns off, but only barely so.

In that case, the peak-to-peak ripple is exactly twice the average current. Or this case is ripple ratio equal to 1. It's rippling up from 0 to twice the average and back down to 0. It's still a triangle wave. It's not this funny clipped triangle wave. It's still a pure triangle wave. And that's exactly the ripple ratio equal to 1 case. Does that make sense to everybody?

So what do I get from that? Well, we can come back and just calculate what happens. I have my expression for ripple ratio is equal to 1, so I can write that again here. I get R sub L is equal to D times 1 minus D quantity squared TR over 2L. and that's equal to 1 at the edge of DCM. That's where I transition between the two modes.

So from this, we can say that DCM happens for R is greater than 2L over D times 1 minus D quantity squared T or for L is less than D times 1 minus D squared T R over 2.

This value of L is sometimes called the critical inductance in the literature. There's nothing very critical about it. But all it's saying is if I go out and pick a big inductor, bigger than this, for the specified operating conditions, I'm not going to run into discontinuous conduction.

But understand, in a lot of cases, for some operating mode, you're going to have to run into discontinuous conduction. Why? Because think about it this way. I have to run for different load resistances, Mostly, in a power converter, you build the converter, and then people hook loads up to it, and you don't know what load they hook up to it. There's some maximum power rating, but maybe they want to put very little load on it sometimes, Maybe sometimes your microprocessor is running. Maybe sometimes it's not.

And so you might have to deal with a very wide range of load currents, up to some maximum. Well, if you have to deal with a very high load resistance or a very low load or a light load, that means your average is going to come down and down and down. And the only way to keep it into CCM would be to use smaller and smaller and smaller ripple so you never hit that mode. And the only way to do that is to get a bigger and bigger and bigger inductor.

But guess what? You've got to run at high load, too. And you run into this case where suddenly, if you've done that by a much bigger inductor, suddenly you've gotten to the point where you've got a lot of inductor energy storage. If your ripple ratio is small at light load, it's really tiny at high load, and this factor will kill you. You just have way too big a component.

So as a practical matter, if you're going to run over a really wide power range, you might have to deal with DCM. So we often design our converters for continuous conduction mode at full power, but they might go into DCM if you're at very light load power. Any questions about that?

So what can we say about operating in this new switch state? Well, let's think about that. What do our waveforms look like? Let's recreate our switching waveforms. We said we were going to have some duty ratio for our switch. And we were always focusing on the duty ratio of the switch. And that's useful in CCM because the diode does exactly the opposite of what the switch does. But now, if we're going to talk about DCM, we can't talk about that.

So here's q of T. And q of T is doing something. So here's T DT. Here's q of t for the switch. But I might have-- then think, OK, maybe what I have is some q sub D. This is the switching period for the diode, q sub D of T. The purple waveform tells me when the diode is on because now I have a time where the switch is on, when the diode is on, and when neither are on.

And maybe this happens for some time-- this time period might be D plus D2 times T, so the diode runs with some duty ratio of D2. And to be a DCM, D2 is less than 1 minus D. Make sense to everybody?

So-- OK, we have this additional mode. What does the voltage on the inductor look like? We had the voltage on the inductor here. What does he look like?

Well, in the first part of the cycle, he's at V1. Here's the voltage on the inductor. He started at V1 in the first part of the cycle. Then when the switch turns off and the diode turns on, he's at V1 minus V2. And then he's on for a while, until this second time period, D plus D2, ends.

What happens if both the switch and the diode are off? Well, if both the switch and the diode are off, there's no current in the inductor, or there's no DI DT in the inductor. If there's no DI DT in the inductor, then there can be no voltage across the inductor. And ideally, this voltage across the inductor would have to go to 0.

So then I'm going to get a time period where the inductor voltage is 0 until the end of the cycle. So this is the ideal waveforms for a discontinuous conduction mode where I'm plotting V sub L of T.

And what do I have for periodic steady state requirement? The average voltage across the inductor still has to be 0. That means that this area now balances against this area. So the average-- the integral of this over the cycle still has to be 0. Does that make sense to everybody?

And if I just plotted the current to complete everything, in this case, the current starts at 0 and ramps up. And the second part of the cycle, it ramps down, and then it sits at 0 for the rest of the cycle until T. And so this is I sub L.

So this is what we're going to have in DCM. What is the conversion ratio in DCM? And I say, well, why do I have to ask that question? Because previously, how did I find the conversion ratio? The way we got that voltage conversion ratio is-- what we did was we said that V1 DT is actually equal-- has got to be equal to V2 minus V1 times 1 minus DT. The average voltage across the inductor is 0. But that was based on this waveform. Our waveforms are now different, so now we've got to recalculate it.

But we can find the conversion ratio the same way we do it. We know that the average voltage across the inductor has to be 0.

So let's just calculate that. In PSS, what do I get? I get the average voltage across the inductor is equal to 0, and that's got to be equal to D times V1 plus D2 times V1 minus V2, and that has to equal 0.

So what I get is if I rearrange this equation-- what I should get is D plus D2 times V1 is equal to D2V2. Or V2 over V1 is equal to D plus D2 over D2V1, or D2 plus D over D2, which is also equal to 1 plus D over D2.

So that's not exactly an answer in terms of it, because it depends on the duty ratio of the switch. But then I got to go figure out what the duty ratio of that diode is going to look like.

But what do I know? I know that D2 is less than 1 minus D. Whatever it is, it's less than 1 minus D because this is the duration D2, or D2T. Right. And that's less than 1 minus DT.

So if I took the CCM-- what is it-- this is DCM. What is it in CCM? In CCM, V2 over V1 is equal to 1 over 1 minus D, which I could rewrite as 1 minus D plus D over 1 minus D.

All I did was add a minus D and a plus D in the numerator. And then I could rewrite this as 1 plus D over 1 minus D. So here, we can say that the conversion ratio in CCM can be rewritten. It's a less convenient form, but it's 1 plus D over 1 minus D. So it's clear for any value of D, the conversion ratio is going to be greater than 1.

Here-- because D2 is less than 1 minus D, V2 over V1 is bigger in discontinuous conduction mode than it would be in continuous conduction mode.

So what's going to happen is for a given duty ratio, whatever I get for my V2, it's going to be bigger than I would get in continuous conduction mode. The output voltage is going to get pumped up a bit. So I can tell that. That's pretty straightforward.

It's a little bit messy. But we can go figure this out because I have the means to calculate what D2 is because I know exactly what these ripple waveforms look like.

If I write the exact expression-- if I basically eliminate D2 from this expression, here's what you get. You get V2 over V1 is equal to a half plus 1/2 the square root of 1 plus 2D squared RT over 2L-- or RT over L, for DCM boost.

So I'm just writing what this is for a boost converter. It would be different for different kind of converters. This is just an example.

But what's the difference? Well, we said, look, if I'm in CCM, the voltage conversion ratio is just a function of duty ratio. It doesn't matter the inductance. It doesn't matter the load resistance. It's just going to-- very straightforward relationship. Suddenly, when I get into DCM, it depends on everything. It depends on my load resistor. It depends on my switching period. It depends on my inductance. It's a mess. It doesn't mean you can't do it. It's just not a simple relationship anymore.

So suddenly, it's a lot more complicated than it was in continuous conduction mode. And it also affects the converter's dynamics, as we'll see. Any questions about that? But that said, we can't avoid it. Yeah, go ahead.

AUDIENCE: How do we know that we stay in periodic steady state when we go through these different modes of operation?

DAVID PERREAULT: That's a good question. It's a question of, will it reach periodic steady state? That's the same analysis you go do, You can go do this analysis and find out whether it'll end up in periodic steady state. It generally will settle-- I mean, because you can solve for these things, the boost converter will end up in periodic steady state unless-- suppose I disconnected this load resistor in this case. That's equivalent to making R equal to infinity.

If I come over here and I make R equal to infinity, V2 over V1 goes to infinity. That means it's just going to wander off to infinity. The output will wander off to infinity.

Why would that happen? Because look, this is IL in DCM, So every cycle, I take energy out of the input, put it in the inductor, turn off the switch, and that energy goes to the output. And I build up some energy on the capacitor. And the feedback mechanism that gets you to steady state is the resistor drains some of that energy. But if you don't ever drain the energy, this just wanders off to infinity. I actually had a converter explode on me once because of that. The electrolytic capacitor literally blew up in my face.

AUDIENCE: Were you wearing glasses?

DAVID PERREAULT: What's that?

AUDIENCE: Were you wearing glasses?

DAVID PERREAULT: You're damn right I was-- I can still see.

[LAUGHTER]

I always wear glasses when I'm working on a higher powered thing these days. And I should on any converter because literally-- I was looking down to it. And I bumped the connection to the load with my elbow, and the capacitor just literally exploded in my face. They often don't vent so dramatically. They're not supposed to, but in this case, it exploded. So, yes, wearing glasses is a good idea in these kind of things.

I actually did that this morning on our demo converter, too, but that's not what happened. But what happened was the MOSFET started breaking down. In other words, it would just start dumping energy into the MOSFET, which isn't good for the MOSFET. So fortunately, I didn't kill the demo because I was being stupid by taking off the load.

But the point is you don't necessarily have to get to periodic steady state in all converters. And this is an example where you can get to really high voltage. And in fact, if you look at little photo flash chargers for those little cameras that you use at weddings, the actual film cameras, yielding cameras, they often use a boost or a flyback in DCM to charge the output voltage up from a volt and a half up to a few 100 volts. Same trick because you can get a very large conversion ratios this way.

So your DCM is harder to model. How would I think about modeling it? And I'll let Monsi start to set up and roll down screens and stuff.

One way I could think about this is the following. Maybe I would think about-- if I looked at my converter here, this is my inductor current waveform. What does my diode current waveform look like-- I sub D?

I sub D is 0 in the first part of the cycle. Then it conducts in the second part of the cycle, until D plus D2, and then it's 0 for the rest of the cycle. So this is DT. This is D plus D2, T, and so forth.

So maybe, I could think about modeling my converter from a control perspective as just whatever this current ID is into the output, So maybe I would model it by saying, OK, let me think of my converter as this current ID, which looks like this-- DT D plus D2 T and T, and just repeats-- into my output capacitor and resistor.

So maybe I would start to start to think of my converter as some sort of funny current source. And maybe I can relate this current source back to the component values and just start thinking about my converter as a current source because that's really what I've turned it into. I'm going to be squirting some current into that output capacitor every cycle.

Well, what does that look like? What I know here is that this current, I peak-- I peak is equal to V1DT over L. And this time period-- let me call this time period to discharge delta T, which is just D2 T.

What do I know? I know that V is equal to L DI DT. So delta T is equal to D2 T is equal to L I peak divided by L I peak divided by V, which is V2 minus V1.

And I can put in I peak here. So that's just going to be V1 DT V1DT over L divided by V2 minus V1. That make sense to everybody?

And so let me ask, what is the average-- what's important over the switching cycle is what's the average current that I'm sticking into the output? I should ask, each cycle, how much on average am I squirting into the output?

Well, I sub D-- the average value of I sub D is basically the area there, which would be 1/2 the base of that, which is this delta T, V1D2 over V2 minus V1. 1/2 the base times the height, which is IP, which is V1DT over L.

This is the area under that triangle-- 1/2 base times height-- times 1 over T. If I rewrite that, then, what I'm going to get is 1/2 V1 T V1 squared T over 2 V2 minus V1 L times D squared.

Let me just see-- double check to make sure I did my math right here. It's something like this, anyways. V1 squared T over 2L. V2 minus V1 D squared. OK, yeah.

So really what I have is, now-- I can think of all the converter leading up to the output cap as just a current source that's a function of the duty ratio I control. So I've turned this thing into a current source that squirts current into the output. And I think about it that way and control it that way. And if I know what my load resistor is and the voltage I want, I know how much averaged current I want to stick into it. So that's how I might model a DCM converter.

It's more complicated, in some sense, than my continuous conduction mode converter because it's a function of all this junk. My switching period and my and my inductor value and all that stuff matters, where it didn't before. But I have to live with that because it's this more complicated operating rule. Yeah.

AUDIENCE: So does that mean when you go from CCM to DCM, you're probably going to change the duty cycle to maintain a target voltage?

DAVID PERREAULT: That's right. You're going to use a lower duty cycle than you would otherwise use in the boost converter. Yes. That's right, because the conversion ratio goes up. And so your controller-- and we'll talk about this later-- your controller usually has to manage that boundary. You have to have a controller that works for either case. Let's just take a look at this. So Monsi has set up an actual boost converter that's doing this. And let's take a look at the waveforms.

In yellow, you can see-- you can essentially see the switch voltage waveform, which is essentially this waveform moved up where this is 0. So in the first part of the cycle, you get 0 in the yellow wave-- from the second part of the cycle, it's at V2. And then you get this time period where you see this voltage ring. What the heck is that?

Well, I said we had this time period where both the switch and the diode were open, and so that this voltage here ought to be exactly at V1. But what really happens in the real world is that-- guess what. There's some switch capacitance and some diode capacitance. And so what happens is this voltage, instead of going to 0, rings with those capacitance-- the inductor ring with the parasitic capacitances.

People don't like that. That's unpleasant, It causes noise and stuff. But you can't do anything about it. If you're going to have both devices on. In some cases, actually, people actually even put in another switch to short it out. But generally, you don't-- generally, you just let it ring, and that's life.

But you can now see that you have the rising inductor current, the falling inductor current at 0. It's exactly doing what this yellow waveform is doing here. Any questions about that? Yeah.

AUDIENCE: Can you mention the source that caused that ring again?

DAVID PERREAULT: It's basically the inductor. resonating with parasitic capacitance at this node. Yeah. It's an unfortunate thing. People don't like it. It causes noise. It's a little bit annoying. That's life.

So I said-- sometimes you have to live with DCM operation, especially if you operate over really wide load ranges or you want a finite value for your inductor sizes. Sometimes you actually want DCM operation. People often design their converters for DCM operation.

Why is that? Well, suppose I thought of-- I had a reference current. And in this case, since I'm talking about a boost converter, let me think-- suppose I had a reference current for current to draw to deliver the power I wanted from the output, from the input.

If I have a small ripple converter and I say, OK, ideally speaking, to support the load, I had this low current reference and then I stepped it up for this higher current reference. So this is I ref for the reference current-- for the inductor waveform,

If I have a small ripple converter and it's rippling like this and then I step the reference-- I want much more current to deliver more power, what has to happen? Well, the inductor current is going to slew up, and eventually, it can get there, and I can deliver this again.

But you notice that there's this time period where the inductor current takes a lot of time to change. If the ripple is small, the inductor value is big, it takes a lot of time to put the voltage across the inductor to slew the current in it. And that means for all this time period, you're not delivering the right output power, or you're at least not drawing the right input power. So this converter-- its ability to change its operating point is limited by the small ripple.

Suppose I had a DCM converter. Well, maybe I have something like this, and I-- again, I step my ripple. But maybe in a DCM converter, my inductor current looks like this. And the next cycle, I want more? Fine. Boom. I can deliver that almost instantly.

In other words, the high ripple means that I can change the average current in the inductor much faster, maybe even within a cycle.

So there's-- on the one hand, we often want small ripple because we want nice waveforms and not a lot of ripple and we got to filter it. On the other hand, the ability to change the energy storage and the energy processing of the converter quickly might take you towards using DCMs. But for a very high transient response converter, having high ripple and possibly DCM operation can be a good thing.

Any last questions before we wrap up? OK. Next class, we're going to start talking about, how do I actually do things like design our inductors? Have a great day.