[SQUEAKING]

[RUSTLING]

[CLICKING]

DAVID PERREAULT: OK, why don't we get started? My computer decided to tank right now, but we can get started without that. And then we'll see what it says.

I'd like to continue our discussion of soft-switched power converters today. And if you'll recall, when we think of soft switching, we're thinking of these switching transitions when we're commutating between devices. And we said, if I'm going to consider soft switching, I can consider 0 voltage switching, ZVS, and 0 current switching, ZCS. And some switching transitions, you can get a combination of both of these, depending upon the design.

But we also have to think about switch turn-off and switch turn-on. So we said, in a turn-off transition, what we're doing is this the switch current, i switch, is going from whatever its maximum value is down to 0 in some finite time.

And if I did a hard-switch converter, just a straight-up buck converter or something, what I would see that, usually, what would happen is the voltage on the switch would rise, and then the current on the switch would fall.

To do a soft-switch transition, we want to somehow delay the voltage rise, which we can do by putting a capacitor across the device as one means of doing this, in which case we'd see a voltage rise which did something like this. Maybe it would do this, like this. And we would delay the rise.

The comparison to that is if I did a zero-voltage switch turn-on. In a switch turn-on, what we expect to have is that the current in the device falls. I'm sorry, the current in the device rises.

So if pink is the switch voltage and green is the switch current, what I get at switch turn-on is, usually, the current in the switch rising. So here's i switch. And usually, what happens is if I'm doing a hard-switch transition, the switch current comes up, and then the switch voltage falls some time after that.

In a zero-voltage turn-on, what I do is try to make some kind of circuit magic such that I will already make the voltage on the switch 0 somehow before I turn the switch on and the current rises. So I'm going to have some kind of circuit magic that lets me make the voltage 0. Or another way to look at it is-- and I only turn on the switch when the voltage is 0.

Now, that's ZVS. And ZVS just turns out to be best for really high-frequency designs. Because I also take care of issues with discharging the switch capacitance at the switching turn-off transition.

But I can also do zero-current switching. So zero-current turn-on is the dual of zero-voltage turn-off. So what we have here is that the switch voltage is falling. So this is v switch. And normally, what we would have is the switch current rising before the switch voltage falls.

But maybe I can delay that somehow and really reduce this overlap, like this, and get a reduced switch overlap between turn-on and turn-off. And I can usually do this delay.

One way to do it is to put an inductor in series with the device. So when I turn it on, the inductor takes up the voltage and, hence, we can have low voltage while the switch is still turning on.

And lastly, we can have a zero-current switch turn-off, which is the dual here, which is what I'm going to have is, when I'm turning the switch off, I'm going to have a rise in the switch voltage, like this. And normally, the current and the switch can only fall after the switch voltage rises.

So here's i switch. And what we're going to do is somehow play some kind of magic game where we're going to make sure that the current and the switch is 0 before I ever try to turn it off.

So I can use any combination of these kind of switching transitions in order to get soft switching. And we saw a few examples in DC to DC converters where we did that last time. Let me pause there and just ask if there's any questions about the general concept.

OK. So what do I get out of soft switching? Soft switching, it's implied that if I'm snubbing the switch transition, I'm doing it in a pseudo lossless fashion, meaning if I put a capacitor across the device for the turn-off loss, I'm going to recover that energy.

So one thing I can get out of it is I reduce the switching loss. So one thing I can get out of this is the switching-- I'm sorry-- P switching goes down. And I can either use that to make the switching frequency higher because the switching loss is proportional to frequency.

So if I've turned it down, I can turn the switching frequency up and then make my components smaller. So one benefit of this is I typically can push up my switching frequency and push down my converter size.

The other thing I can do is push up my efficiency. So if I have lower switching loss, I just take that out as efficiency or some combination thereof. And that's one major reason to do soft switching.

A second reason to do soft switching is a little bit more subtle. We spent a bunch of time talking about EMI. And one of the things is for a given switching transition, if I had a very rapid, say, dv/dt here, if I'm going to soft switch, usually, what I'm doing is I'm arranging so that the voltage slope is lower.

So I can reduce EMI in a soft switch converter in a lot of cases. So another reason people will often use soft switching is because they want to get nice, smooth waveforms and, hence, reduced EMI in the converter.

So if we were to think about-- let me think about, before we talk more about soft switching, let me think about this in a little bit more detail. If I were to come back-- and give me one moment here-- let's think about the soft switching impacts on EMI.

So if I thought about my buck converter, for example-- so here's my buck converter. And usually, I'm going to put some input filter here. But in the buck converter-- in the input filter, maybe I consider the input capacitor, too-- let me just think of this as IDC. If I were to look at this current, this current that I'm drawing through the input switch-- I'll call that ix-- is, more or less, a square wave.

It has some finite slope on it, but I'm looking at some kind of square wave slope With time, so if I were to look at this. And so the frequency content is what drives my filter. And then that's what sets what shows up, say, at my input.

Likewise, if I looked at-- be helpful if I put a diode here instead of a capacitor. If I looked at the switch voltage-- let me call this vy-- vy has the same sort of shape. I have high for some time DT and then low for the rest of the cycle, for T. And I've got this very sharp-edged waveform right here.

Now, that doesn't appear specifically out of the converter. But remember, maybe I might have some parasitic capacitance coming back here and either into my filter somewhere or towards my input somewhere. So basically, this very sharp voltage injects currents back. And that can appear as noise, either within my circuit or at the input or output somewhere.

So one of the things that influences how much EMI your converter generates is the di/dt here or the dv/dt here. Does that make sense to everybody?

So let's think about what's the impact of square wave switching. Because what are we going to do? Maybe if I use zero-voltage switching here, instead of having a really steep waveform vy, maybe what I would have is a more slew rate limited waveform here. And so what's the impact of this slew rate on dv/dt or the slew rate on di/dt if I had lower di/dt switching transitions on the fundamental waveforms?

And here is what we can say about a square wave waveform. And it doesn't matter what the shape of the square wave is. But let's just say this was a current i of t.

And here's the waveform I'm going to consider. I'm going to say, OK, it's going to have some finite slew rate up to a maximum current I. And it's going to have some slew rate down. And at the end of this, here is going to be some period T. And let me call this period to the middle of the rise some time period DT.

And then I'm also going to worry about-- I'm going to call this t rise. I'm going to call this t rise. Let me assume the rising and falling slopes are the same. They may not be in the real circuit, but, for purposes of analysis, we can consider that way. So I'm exaggerating this rectangle wave into a trapezoid wave just for illustrative purposes. Does that make sense to everybody?

Here is what the magnitude of the n-th Fourier coefficient looks like. You can derive this. I actually took this version from Ott's *Electromagnetic Compatibility* book, but it's also in KPBS.

It looks like this-- 2ID, or it's 2ID. I'm sorry, it's 2I times the duty ratio times sine of n pi D over pi n D times sine of n pi tr over T divided by pi n pi tr over T. So This is the n-th coefficient. And the whole waveform has the sum of these individual Fourier coefficients.

So why am I raising this? Each of these is a sinc function. This is a sinc function in terms of the duty ratio. And this is a sinc function-- this is a sinc function in terms of the percentage rise time.

And this is for the n-th harmonic, so it's at the n-th frequency, So this is multiples of the switching frequency, essentially, where these appear. And then this sinc function is related to the slew rate normalized to the period, the percentage rise time, if you will.

Why do I raise that? Because think about what a sinc function looks like, A sinc function looks like this. It falls off. It starts at a certain height 1. And it falls off like this, like this. At really low frequencies, a sinc function asymptotes to 1.

Above x equals 1-- if this is sinc of x, above x equals 1, you get an asymptote that does this, that goes as 1 over x. So what I'm looking at isn't the actual value of the sinc function at any point but what does the envelope of the peaks of the sinc function do. That make sense to everybody?

So at low frequencies, it's DC. It's constant. And at high frequencies, it falls off as 1 over x, whatever x is. And we have the product of two state functions for our waveform.

So what does that mean? Well, let's just start by saying, suppose I have a perfect square wave. Suppose I have a design where, yeah, these walls are so steep, it's, more or less, instantaneous. What would the frequency distribution of that look like? And what would the waveform look like?

Well, here we go. And I should say that I'm taking these very nice slides from a presentation done by John McCloskey, who's the chief electromagnetic compatibility engineer for NASA. And he put together a really nice presentation on this.

So this just shows, suppose I have a square wave, or a rectangle wave, of duty ratio D. And here's the individual harmonic terms. And as I add up more and more harmonics, you get a square wave. That make sense to everybody?

And because I've ignored, I've made tr over T to 0, this whole factor just becomes 1. So I get harmonics that fall off as a sinc function in frequency.

So what does that look like? That looks exactly like this. So here's exactly the envelope I'm talking about. This is in a linear scale. So at low frequency, it has constant, and then it falls off as 1 over x. If I looked at this on a log-log plot, what I get is constant up to x equals 1 and then falls off as 1 over x, which is minus 20 db per decade.

So why am I bringing this up? What I'm trying to say is, Where is the energy content in my square wave waveform? It's sitting at a bunch of harmonics. And I'm not really worried about the individual harmonics. I'm just worried about how the envelope of those harmonic amplitudes fall off.

So an individual harmonic might be really low, but they will be no bigger than the envelope dictated by the sinc function. Some of them will get 0. Some of them will be maximum. But if I'm just looking at how the energy goes down overall, it goes as 1 over frequency for higher frequencies.

So if I looked at my square wave, in this case, this is what happens. Essentially, I have the fundamental. And then everything tends to fall off with an envelope that goes as 1 over the frequency. Any questions about that?

So now, what happens if I then say, all right, I don't switch my waveforms and exchange there. I don't have really fast gate drives and no capacitance at my switching node. And I, instead, allow something that limits the slew rate and voltage or current, depending upon whether I'm talking about a voltage or current.

This would be exactly the impact of, instead of having really hard and fast switching transitions, maybe I put a capacitor across my device that limited the dv/dt to a reduced value. Maybe this would instead turn into some kind of trapezoid waveform. And what I get is this second sinc function factor, which is related to the percentage of time it spends rising and falling. Does that make sense to everybody?

So here's what happens. From the original square wave, I get something that starts off at DC and then falls at 1 over f. And then, at higher frequencies, I get this additional attenuation of another 20 db per decade. And that's because of the roll-off of the second product of the sinc function. Make sense to everybody? Yeah, Loren?

AUDIENCE: So why would they not have [INAUDIBLE] frequencies?

DAVID PERREAULT: Because, remember, the sinc starts falling, sinc of x starts falling, when the asymptote is above x. So it's going to start rolling off when whatever the argument of the sinc function is becomes greater than 1. So what that means is when tr over T is really small, it doesn't attenuate. And then, when I get to higher frequencies, it does attenuate.

So what that should mean, then, is that, compared to having pure, sharp square waveforms, if I can do something that slows the dv/dt or the di/dt down, I should get reduced high-frequency content. And that should make it easier to filter.

Let's see this in action. And this is what I love about this presentation that he put together. What the author of the original presentation did was build a setup, get a signal generator, measure the voltage waveform, and then measure the spectral content so you can see exactly what happens with the actual waveform.

And here's the applied waveform. He's just doing exactly what we said. And he's varying the parameters. And here we go. He's got a 100-microsecond waveform that's, like, 10 kilohertz. So this is low frequency, easy to measure, easy to generate and measure.

And what you can see is exactly what he said. It's a 50% duty ratio waveform. It's not 0 rise and fall time, but it's pretty small. And what you can see is exactly that it falls above-- the fundamental falls off as 1 over f, basically, exactly what we would expect to see.

And here, you can see, OK, the even harmonics are low, and the odd harmonics are high because this is, ideally, a 50% waveform. And there's no even harmonics in it, except that it's not precisely 50%. So at high frequencies, there is some content.

What happens if you go from 40 nanoseconds to 360 nanoseconds for rise and fall time? So that's 0.36%. When you get to a high enough frequency, boy, some of the high-frequency content starts to go away.

What if I were to slow it down more? What if I were to make this, instead of 0.36% rise time, 3.6% rise time? Wow. I'm really killing off the high frequencies. If I make it 10% rise time, I'm really killing off high frequencies.

Now, maybe in a reality, maybe I wouldn't let my transitions take that long, or maybe I would. If you looked at the transitions in some of those waveforms we showed you in the last class, they were pretty slow compared to the switching period because you're using it to get the ZVS.

So what you really get, one of the effects you get, is as you as you slow up the slew rates, you start to kill off a bunch of high-frequency content. Now, your filter is sized, typically, for the lowest frequencies. But when you put parasitics in there, you remember that we saw the circuit parasitics make the filter start to drop out its capabilities at high frequency because the parasitics.

So you might argue that the best way to get rid of the high-frequency noise is, instead of generating and filtering it, just don't generate it. Or the stuff that's most likely to couple through, say, a parasitic capacitance is the highest frequencies. So if you just never generate them, it's a lot easier to filter them.

So you might not design your whole converter around this thought, but it really does help you with reduced filter size. And we've seen that, with soft-switch converters, you can often really make your filtering job easier. So questions.

AUDIENCE: I figured it out.

DAVID PERREAULT: Figured it out, OK. Other questions? So I wanted to argue that one of the things we can get out of soft switching is reduced switching loss, which lets us either get efficiency or higher frequencies and, hence, possibly higher power densities, smaller component size.

And for a given switching frequency, we can reduce the EMI. And even when you're pushing frequency, that can help you because you might have a high switching frequency converter that doesn't have increased really high-frequency content.

All right. Well, last time, I introduced some power converter examples that were soft switch. And I showed you DC to DC converter examples. And I was a little bit rushed. But we said that-- and I should say that there's all kinds of ways that people try to break down different classes of soft-switching converters. There's a whole variety of techniques. And you can mix these techniques any different ways.

What's a good way? Well, it depends upon, first of all, your device characteristics. I make it sound like all devices are the same, but it turns out, they're not. So what soft-switching technique might really be suitable does relate back to the details of the type of semiconductor device you're using.

There are some devices that just don't like to turn on under 0 voltage. So you probably wouldn't want to use zero-voltage switching for that kind of device. So you have to match the soft-switching technique to the device type you're using.

And you have to match the particular circuit or technique to the circuit application. Because it depends upon the voltage range you want to operate over or the power range you want to operate over.

You might change which circuit implementation of soft switching is a good idea versus a bad idea. And that's why there's so many different techniques people use, or versions of the general idea people use, is because there's not one size fits all.

I showed you a few examples. And maybe the broadest split I could make in the ideas of how soft switching is implemented is there's some designs in which you will design the whole circuit to achieve some particular kind of soft-switch operation. And that usually means that the way the circuit's component values are chosen, the whole control scheme, everything is geared around meeting the requirements of soft switching.

Why? Because maybe I have to make the circuit do something magic, and I only want to turn the switch on when the voltage is already 0 across the switch or something. So that prevents me from just doing standard fixed-frequency PWM. And I just work really hard to make that happen. So that's one very broad category of circuits.

Another category of circuits is, jeez, I don't want to redesign and conceive this whole funny operation of this circuit in order to achieve soft switching. What I'd like to do is, more or less, design a PWM converter like we've been talking about all semester and somehow add in some auxiliary circuit that's going to help me get soft switching.

Both of those kind of strategies are used. And I'm going to show you one of each today. And since last time, we talked about DC to DC converters, today, I'm going to talk about DC to AC converters, or inverters. And you can use the same kind of ideas on rectifiers, naturally.

So let's talk about a couple of examples of soft switching in DC to DC converters. Here is one soft-switch inverter. So I want to take a DC input voltage. And here, for simplicity, I'm illustrating something with a DC bus split in two. You don't have to do it this way. It's just easy for educational purposes to illustrate this this way.

So this is my DC input, Vdc. My output is the voltage across this capacitor here, Vcf. So I want to generate a low-frequency AC voltage across that capacitor. And that AC voltage is going to be applied to the load.

So maybe you could imagine this is some low-frequency AC output that I want to synthesize to deliver power at low-frequency AC, single phase. And I'm willing to switch this inverter at a high frequency in order to do that.

And this first design is called a resonant pole inverter. Well, that's one of its names. It's very often used. It's one of these techniques where the whole circuit is designed and controlled to achieve soft switching and to achieve high-frequency operation.

And what we can see is, here's my switch. And I'm showing like an IGBT, or bipolar, with a diode across it, but this could be a MOSFET. And each of the two MOSFETs in this bridge leg-- this is clearly a bridge leg, or half bridge-- has a capacitor across it.

So what does that suggest to you about what kind of soft switching I'm going to be doing? I got capacitance across my devices. It's going to be zero-voltage switching because when I put a capacitor across my device, I can delay the rise and get zero-voltage turn-off. And I'd better turn the switch on when I have 0 voltage. Otherwise, that energy in this capacitor will be dumped into the switch.

So this is going to be a zero-voltage switch circuit. This capacitance might be an external capacitor. It might entirely comprise the parasitic capacitance of the device-- just depends on the design.

So what's the idea here? We can assume that over any switching cycle, the voltage Vcf across this capacitor is constant. It might vary anywhere between plus Vdc over 2 and minus Vdc over 2. But within a switching cycle, I can think of it as quasi constant.

So what are we going to do? Suppose I want to deliver an average, a local average, current over a switching cycle iL into the load filter combination. Well, what I'm going to do is I'm going to, for example, turn on the top switch S1. Suppose the top switch S1 is on.

So long as Vcf is smaller than Vdc over 2, I've got a positive voltage across the inductor, and the inductor current is going to ramp up. So that's going to be some positive-rising current here. When I get to some peak-commanded current that I desire-- this is a current-controlled converter-- I will turn off S1 under zero-voltage switching because this capacitor is snubbing the switching transition.

What will happen then? When these two switches are off, the current will actually split between these two resonant capacitors, and the bus voltage will resonate down until the bottom diode turns on.

So now I've got the bottom diode on. And I'm still delivering current to the output. But now, the voltage across the inductor is negative, and the current is going to slow down in the inductor. But notice something. As long as this diode's on, what is the voltage across the bottom switch? 0.

So, in fact, by turning off the top device and having this inductor ring with these capacitors and turn the bottom diode on, I am, essentially-- boom-- creating a case where the bottom switch voltage is 0. That means, at any point after this node voltage rings down from Vdc to 0, I can turn on this switch.

So I have zero-voltage turn-off, zero-voltage turn-on. I will turn on the bottom switch. And what will eventually happen is, because the current's slowing down, the inductor current will reverse. And now I'll be building up negative current in the inductor.

What I'm going to do is I'm going to let this negative current build up to a minimum particular value. And then I'm going to-- I'll call that ip minus, for example. And eventually, I can turn off this bottom switch. And the precharged current in this inductor is going to resonate with these two capacitors again.

And again, I turn off the bottom switch with zero-voltage switching because of the capacitors. This inductor will ring with these two capacitors until the top diode turns on. And once the top diode turns on, I'm getting zero-voltage turn-on on the top device again.

And now I'm back. When I turn the switch on, I can get back to here. So I'm alternating between top device on and bottom device on but with dead times in between that let the bus voltage ring down and ring up. Any questions about that?

Now, what's the magic trick of this? In order to get zero-voltage switching-- let's consider this case. I'm delivering negative current in the inductor. The inductor current's gone negative. In order for the voltage in this state to ring all the way up such that the top diode turns on, there's a negative minimum current I need in this inductor.

That negative minimum current is, actually, a function of the voltage. So if I just didn't have enough current, this voltage would ring up, but it wouldn't ring up all the way. And then it would die, and I wouldn't get there. So I have to have a certain positive current to make sure I can ring down and a certain negative current to make sure I can ring up.

But if I ignore the short durations of the switching transitions, the average current delivered into the capacitor over a switching cycle is really the average of the peak positive current and the peak negative current. So this is very much a hysteretically controlled inverter-- hysteretic current control.

But what I'm doing is I'm making the positive current I deliver always positive enough that I can bring the voltage down and the negative current always negative enough-- and notice that the current's reversing here-- so that I can always ring the bus voltage the other way. And I'm delivering the average of the two in a switching cycle.

So I can work out the math that says how negative does the negative current have to be, how positive does the positive current have to be, create a current-based switching cycle that will always give me soft switching and will let me deliver either positive or negative average current into the capacitor independent of the capacitor voltage. And hence, I can deliver a low-frequency AC output component by modulating the positive and negative peak current over a cycle.

The exact values that you need to do that make for wonderful homework problems, not that hard to work out. But this technique is taking advantage of the fact that you have a high ripple current in the inductor. So the inductor current has to go in both directions in order to, generally, give you your output.

So you have to turn off the top switch when the inductor current's positive and turn off the bottom switch when the inductor current's negative. I have to have a high ripple to make this thing work. But if I'm willing to get that, I can get soft switching in my devices and deliver the average current I want.

Any questions about that? This, again, has variable switching frequency. It has all kinds of challenges. It has high ripple. But it can be very compact and efficient.

I should say one of the ways people think about this-- and we're going to come back to this because we're going to look at classes of converter that use this kind of zero-voltage soft-switch technique a lot-- usually, what you have is when you're going to turn off your device, you usually have some inductor current that carries a current in the right direction to give you the switch direction you want.

So usually, when you look at this kind of circuit, you say, OK, when the device is turning off, it's loaded by some impedance that, typically, looks inductive or, basically, is carrying a current in a direction that will help charge those capacitances in the right direction to give you soft switching is the way to think about it. Questions?

OK. This circuit works under variable switching frequency and variable duty ratio based on hysteretic current control. So everything about it is governed by the desire to get soft switching. So if I told you, you know what? We want to do harmonic illumination PWM with this inverter leg. Can you do that?

Not really, right? You don't really have the means to do it. You can't just generate switching angles at your command because everything's governed by what the ripple current in that resonant inductor is doing.

So this is nice that it lets you generate a low-frequency component plus high frequency into an output and get filtered and get what you want. But it's very constrained in how you have to operate it. It's in that first category of converters I talked about.

What if I, instead, said, you know what? I've got some motor drive. And I, basically, know how I want to control it in terms of switching between my top switch and my bottom main switch. But I don't like this notion that I don't get any control. I just want to switch this inverter because I already have a control system in mind for what this inverter's going to drive.

And maybe in the thing you want to drive-- this has an LC filter at the output-- maybe I just want to run this inverter into a motor. So I don't really want to have this extra stage of high-frequency stuff. What could I do?

Well, maybe, I would imagine this. And this is often done. Suppose I said, all right, here's my bus voltage. Maybe I'll, again, just for simplicity, let me split my bus voltage in two. And I could do this with just splitting capacitors.

Maybe I would just have a pair of devices, a first device and a second device. And maybe this thing is just going to drive some load current iL of t. And I don't know what this load current is, but maybe this is some motor or something.

And what my goal is is to, basically, generate this voltage vx according to some command. I either want to make VDC or 0. I want to make vx go to VDC or 0. I want to use standard motor drive techniques. I don't want to deal with all this funny behavior.

And yet maybe I would like to reduce the slew rate of vx. Why? Because I do have EMI going out towards my motor. So maybe I want to reduce that dv/dt. Or maybe I want a PWM at pretty high frequency with high efficiency, use a high-power system. So how could I do that?

And here's one thing you could do. Maybe we can look at this thing and add some auxiliary circuitry that lets me switch these two devices between these two devices any time I want but will give me a soft-switch transition.

Here's how we could do it. Maybe I will go-- the first thing I'll do is I'll add some capacitance across each device. And that would include the intrinsic capacitor in the device. So if I'm going to do that, is this going to be a ZVS converter or a ZCS converter? ZVS because I know that I'm going to be playing with ZVS switching transitions.

What is my auxiliary circuit going to look like? It's going to look like this. I'm going to add a resonant inductor, Lr, and then maybe a pair of back-to-back switches. But let me just think of this-- well, I'll draw it as a pair of back-to-back switches, just to be clear. If I had a bidirectional switch, I could do this with a bidirectional switch, but maybe I'll do it like this.

Actually, I could have put them back to back the other way. It might have been better. But here's my back-to-back switch. So this thing here is just a controllable bidirectional switch. Now this bidirectional switch has an inductor in series with it. So do you think I'm going to switch this switch with zero-voltage switching or zero-current switching? Zero current.

So this circuit, the pink circuit, is an auxiliary circuit, which is going to give me zero-voltage switching on my main devices and, itself, is going to operate with zero-current switching. And what we're going to do is this pink circuit is only going to operate transiently at the times when I want to switch between the bottom and top main device.

This circuit is called an auxiliary resonant commutated pole inverter. It was invented in the early '90s. It was actually an extension of a notion called the McMurray snubber. This was developed at GE.

And McMurray, before this particular invention, was a very famous power electronics researcher and designer at GE. And a guy named Rik De Doncker came up with this circuit. And he actually originally demonstrated it in high-power motor drives.

And here's how it works. For any transition, regardless of this load current, if I want to go between this device and this device, there is some way I can operate the auxiliary circuit to let me transition between the two devices with zero-voltage switching.

So here is the notion. This is actually from De Doncker's original PESC '91 tutorial on the ARCP converter. So he uses funny device symbols. I think these may have been MCCs at the time. But he's using these device symbols. And here's the auxiliary circuit that's going to be zero-current switched. Here's the main devices that are going to be zero-voltage switched.

And what you'll see is that this is-- when you trip the auxiliary circuit, you can go from the bottom switch being on to the top switch being on. And hence, you can get very nice soft transitions, which helps reduce the EMI of the circuit and reduces the switching loss.

Now, there's all kinds of switching transitions to the circuit. You have to think about what's the load current and everything else to figure out what sequence of steps you go through to get between the bottom switch on and the top switch on, or vice versa. I'll just show you one transition as an example. And this is, again, from his original tutorial.

Suppose the bottom diode-- and now he's showing it with GTOs-- suppose the bottom diode's conducting to the output. So this is one case. And I get to the time of my cycle. I say, you know what? The bottom diode was on, but now I want the top switch to be on. What do I do?

Well, what he says is this. He says, OK, let me trip on this auxiliary circuit. I'm going to close this switch. I'm going to close this switch with zero-current switching. And what's going to happen is, because the bottom diode's on, current's going to ramp up in this resonant inductor.

And then what he's going to do is he's going to turn the bottom switch on at 0 voltage because he's going to turn this device on when he's able. And since this diode's conducting, he can turn this guy on at 0 voltage. He charges up current in this resonant inductor. And then he turns off the bottom device.

And then it's very much like the other circuit. This resonant inductor rings with the two capacitors until the top diode conducts. Then, he can turn on the top device. And then he just waits until this current ramps down and goes to 0 and this auxiliary switch turns off at 0 current. And then the top device is conducting the output current.

So I've gotten from the bottom diode conducting to the top switch conducting by triggering the auxiliary circuit. And if you think of every other possible switching transmission and direction of the load current, you can find a sequence that will let you use the auxiliary circuit to make that transition for you.

And so this is a complicated circuit, relatively speaking, especially when he did this in the late '80s, early '90s. This is controls and sensing, and everything was really complicated. On the other hand, at high powers, you don't really mind.

Compared to the previous inverter I showed you, the control for that, it has a lot of stuff, but it's pretty simple. You just get a sense current and make your switching decisions based on commanded currents. This is much more complicated.

But, on the other hand, it lets you PWM any way you want. They were running motors with this, and they were running all kinds of special motor controls. So this is a really good circuit for this kind of application.

And the things you get are reduced switching loss in the main devices for a small extra auxiliary circuit, which is also soft switched. And you get-- this is another set of computations going the other way-- you get, because of the snubber capacitors, you get reduced dv/dt on the switching transmission and, hence, reduced EMI driving your motor.

So that is a lightning introduction to soft switching and some of its benefits and just an introduction to some of the ways you might go about it. Yeah?

AUDIENCE: How do you so abruptly turn off the current in the Lr?

DAVID PERREAULT: Well, you don't. What you do is suppose you're going from bottom to top. If I turn this switch on and leave this switch off-- this diode conducts-- I will charge current positively in Lr. And eventually, this node's going to ring from the bottom to the top.

And then Lr has a negative voltage across it. And so Lr will reduce down to 0. And when it hits 0, this switch just naturally turns off because this diode will block. So what you can do is, by how you control this switch, this current will just resonate back to 0 and stay at 0. Does that make sense?

Kind of. Let's talk a little bit after class. But yeah, the idea is to let voltages or currents ring up or down and then be clamped. In this case, it's clamped by this switch going naturally open. In the other case, it's ring up is clamped by the body diode of the FET. And you go along that route.

So that's a lightning introduction to soft switching. Next class, we'll start talking about a whole class of converters called resonant converters, which are a subset of the world of soft-switch converters that are very, very widely used in high-performance power converters. In fact, the converter inside my laptop adapter here is based on a certain kind of high-frequency resonant converter.

So have a great day, and we'll take this up next class. And yeah, Kaito is wandering around somewhere. I will find him. Have a great day.