[SQUEAKING]

[RUSTLING]

[CLICKING]

PROFESSOR: Before diving into this, I'd like to remind you of what we talked about earlier in the course, which is switching loss. If I took-- for example, let me just take my friendly neighborhood buck converter here. so I have V IN. And I might have a switch. And maybe I have a diode. And I'll think of these as ideal.

And actually, perhaps let me just leave myself a little more room here. I'll have a diode. And maybe for the moment, I'll think of this as this ID, this almost constant current into the output that I'm controlling by PW-ing my converter.

And we can say, OK, well, I'm interested in the losses in the switch and the diode. Let's focus on the switch here. So I might have some switch voltage that I'll call the switch. And I might have some switch current that I'll call I switch.

And I might say, well, when my switch is on, I have I switch squared R on for the switch. And that's how I figure out my conduction loss on my switch. We also said you have to worry about switching losses in the device. And we might approximate switching loss this way.

We might say, OK, let me just focus in on the tiny part of the time cycle when I'm, say, turning this device on and off and turning this device on and off. I'm switching between the two devices. In an ideal world, that happens instantly, but in the real world, it doesn't. And it can take, depending upon my switching frequency, maybe a percent of the cycle or something like that. OK.

So let's just think. What happens when I turn this switch off? Suppose this switch is carrying ID. I would like to turn him off. I can worry about what does v switch do and I switch do. And if I just thought of this switch as an ideal current source ramping from carrying high current down to zero, maybe that would do something like this. So here's I switch. And at some point, it's going to fall from high current down to zero.

Now the question is, does the switch voltage rise before or after the switch current falls? And to the extent that we can think of this diode as being ideal, the diode cannot carry any current until the diode has zero volts on them. He's not going to carry any current till he has zero volts on him. That means that the switch voltage must go up to V IN before this diode can start picking up current.

So I might expect then that the switch voltage, if I'm going to have my switching start here, the switch voltage has to rise from 0 to V IN at which point the current in the switch can start falling. OK. So I'm going to have some switching transition.

And I'm approximating these transitions as linear. What do they actually look like? That depends upon the device, and the device driver, and a lot of other things. But approximating those linear is often not a bad idea. And I might say that this time period takes some, say, time t fall.

And in fact, if you look in datasheets, you can find expressions for t fall under some nominal drive conditions. It'll give you some approximate notion of how long that will take. But we can control it by the driver we choose and other things.

OK. But here's the point. During this time period, the switching transition when the voltage on the switch is rising and the current on the switch is falling, I have a big overlap in voltage and current on that switch. So right. So if I looked at the power in the switch, the power dissipated in the switch, yeah I had some i squared R going on here. It wasn't 0 but it was pretty darn small.

Here I may have, if this is 100 volts and 100 amps, maybe that's like 5 kilowatts peak dissipation. It's 5 kilowatts for a very short period of time that just goes-- actually, I haven't drawn that very well-- that just goes over this duration. But I get a big pulse of power in the switch.

And this is taking place over a very short time period, Maybe if I was switching a megahertz, it would be a microsecond period. And maybe this takes 10 nanoseconds. So maybe that's a percent of the cycle. But on the other hand, this is a pretty big instantaneous power. So I actually have to care about figuring out that power dissipation.

If I waited for the other transition-- this is when I went from the switch conducting. The switch turns off and the diode turns on. If I waited in the other direction, I said, what is the switch loss? What's going to happen here? The opposite. Before the switch can pick up any current or for the switch to pick up current, as long as the diode is carrying any current, the switch voltage has to stay high.

And so the switch current will come up before the switch voltage collapses. And if I looked at the other transition, I would see something like this. Later at the other end of the cycle, I would see something like the switch current rises and then, after the switch current rises, the switch voltage can fall. And I might call this time period t rise. And again, I will get a big power pulsation during that time period.

And then if I use this simplified model, what I'd get is a switching loss that was equal to V IN, because that's the peak voltage, times I d, that's the peak current, times tr plus tf over to times the switching frequency. This is the area under each of these times tr. And I'd get this total power dissipation. Any questions about that?

So I care about this because, you know, we're motivated to turn up our switching frequency until it hurts. If I have a low switching frequency, I have big components, I want my components to be small so I'm going to keep turning up my switching frequency until something tells me that I shouldn't.

And here you have it. You get frequency-dependent losses. And one of those losses is the switching frequency. This is proportional to how many times a second I switch. And we can control that to some extent by making t fall and t rise short but that also has EMI implications so we can't do that arbitrarily. So we're stuck with some amount of switching loss. And I'm going to pick my switching frequency in part to balance the switching loss against the conduction loss or get my total loss down. And this is a good way to proceed. You figure out some balance between the size of the device you're going to use this r on, how fast you can switch it, and get some total loss.

We also did, however, talk about how might I slow up these kind of transitions. We said sometimes those transitions can cause a lot of EMI. And one of the tricks that we showed earlier in the term was we said, OK, well, I could come in here. And if I put, essentially, a circuit that looked like it did something like this, I'll put a capacitor across the device, ideally. But what I will do is I will make it so that instead of putting it directly across the device, I'll do it so that in the forward direction, it looks directly across the device and, in the reverse direction, I have a resistor here.

If I did that, what that would do is when I turn the transistor off, I have another path for current to go through and I can charge the-- I can charge up this capacitance so voltage doesn't need to rise so quickly. So if I do this, if I put this kind of circuit in, I can delay the voltage rise in the device. It might do something like this. So I kind of push this out here. And I reduce the VI overlap. And I also, instead of a fast rise, I get a slow rise.

And that works. And it does reduce the loss of turn off. That's very nice. The penalty I pay is I put a bunch of energy in this capacitor. And when I go to turn the switch back on, that energy ultimately gets dumped into this resistor. And in practice, typically, I make the efficiency of the converter worse, not better. I've reduced the loss in the device, but I've more than offset that with loss in the resistor I get here. OK.

I could do, if I had a device that needed it, I could do the same kind of thing at turn on. I could come in here and I could say, OK, let me add an inductor here, and then I will add a resistor and a diode. So when I go to turn the device on, this will hold the voltage and keep the current from rising very quickly so instead of the current rising like this, the current rising might do this. And I will again reduce the overlap between the voltage and current by delaying the current rise.

But I store a bunch of energy in this inductor. When I go to turn the switch off, that current will circulate in here and dissipate in the resistor. So again, I have reduced the overlap at the switching transition, but I have lost energy in the process on balance. But I might want to do it for reasons of changing, say, the di/dt or the dv/dt's over here. So that's all a little bit of a review. Any questions about that?

STUDENT: Yeah. Are we expected to look into snubbers for the project?

PROFESSOR: That's a very good question. So, no. Right? Why? Because the devices that you're using, MOSFETs and Schottky diodes, tend not to be very susceptible to failure of having a safe operating area limitation or a dv/dt or a di/dt limit in their operation. If I were working with giant gate turn off thyristors, then that would be a different story. And I would have to put snubbers on them.

But for the kind of converter you're designing and the design project, pretty much you can get away without snubbers. In the real world, if you were to build one, you might build it and then say, oh, I'm getting a little bit too much EMI and I associate that with the slew rate of the voltage, for example. And then I might go in, throw in a small snubber to slow down the voltage transition.

So sometimes people will leave provision to put a snubber in if they need it in practice, but it's not something you would necessarily have to design upfront unless the types of devices you were using required it. And that would be typically very high power devices that are a bit more fragile for MOSFETs and Schottky diodes and stuff. Typically, you don't have to worry about it. Excellent question.

OK. So that brings me to the topic I wanted to tell you about today and I only really wanted to give you a flavor of today. But I want to give you a notion of how this works. This kind of poses a question. I said putting the inductor and the capacitor in parallel with the device here does reduce the overlap, but I burn a bunch of energy.

And likewise, putting the inductor in series reduces the overlap, but I burn a bunch of energy. Could I do that without burning a bunch of energy? Could I somehow get that energy back? And what would be the benefit of that? This loss, if I only thought about the overlaps, this loss would be greatly reduced. And that would either let me, A, get rid of that loss and get a very high efficiency converter, or, B, crank up the switching frequency and make my component smaller or some combination thereof.

As you make things smaller, it has to get more efficient in order to meet thermal requirements. But nonetheless, what that would do is let me either get to higher efficiencies or higher switching frequencies. OK. And people do this. They come up with circuits, either changes in the circuit design or changes in the control, that will let people get this so-called soft switching. This pink and yellow transition we would consider hard-switched.

A soft switch transition is one in which you reduce the voltage and current overlap at the switching transition and get much lower switching loss. Now I say that many, many very high performance converters do this. This is a very common thing at the high performance end of the spectrum. In fact, if you looked at the power converter that's in this little miniaturized laptop adapter, which is, A, very small, and, B, exceeds 95% end to end efficiency so they can make it small without it getting hot, that design uses soft switching techniques and turns up the frequency to make it small and uses part of that benefit to make it efficient.

So what I want to do is show you the kind of category of techniques that you might use there. I will say, though, that while this is very, very common in the high performance end of the spectrum, where you're trying for great miniaturization or really fast response, something like that, perhaps the majority of power converters designed in the world don't do this. They just get an optimum.

They live. They try to get good devices, live with the switching loss, make the design trade-offs. All the stuff I showed you so far in the term is the majority of things, that's the way they're designed. Nonetheless, when you're really trying to push the boundaries or, for whatever reason, you want to go to very high frequencies, suddenly you have to use these other kinds of techniques. OK. Any questions about that?

OK. So let's think about what we might do to build so-called soft switch converters. And how soft is soft? Well, there's all degrees of softness. But we'll talk about it. Now there's two flavors. There's one that I'll call zero voltage switching, ZVS, and one that I'll call zero current switching, ZCS. So you'll see these acronyms used everywhere. And essentially, I'll describe what they mean.

You also have any-- if I think about a given switch, like here I'll think about the active switch here, he has two transitions. He has a turn on transition and a turn off transition. I got to think about both of those because both of those give me some loss. So I might think about turn off. And I might think about turn on.

OK. So let's think about these four possibilities. OK. What is ZVS turn on? ZVS turn off, if I thought about, say, the turn off I had here, I had something where the current was falling like this at turn off. And under a hard switch transition-- so this is I switch.

Under a hard switch transition, we expect kind of V switch to rise before the current falls like this. And that would be a hard switch transition. What we do in a zero voltage turn off is we typically do something to delay the voltage rise on the switch. Maybe it does something like that. Well, that ought to look pretty familiar to you. Right? That's what the snubber did over here.

So I might get this kind of behavior by taking my switch that I'm going to try to turn off and just put a capacitor across them. And so if I put an external capacitor across the switch, when it turns off, it'll delay the rise of the switch and I'll reduce the overlap. So it's no different than what the snubber does, but now that I've put energy in this capacitor, we've got to figure out a way to get that energy back without dissipating it. So I can't just turn on the switch because I dumped that energy and lose it. OK.

What about zero voltage turn on? OK. What would zero voltage turn on look like? Well, we said, OK, what we expect normally at turn on is the current in the switch is going to rise at the turn on transition. Right. And usually, after the current rises-- here's the switch voltage-- the switch voltage falls. So that's what I would get and I'd have this overlap period.

Well, what if I could be somehow super clever in my circuit through magical means and just happen to work it so that instead of the voltage falling because I'm turning the switch on, I do something else in the circuit to make the voltage go to zero before I try to turn the switch on? Maybe I have a circuit where somehow through natural action of the circuit, the voltage goes to zero and then I can turn the switch on with no voltage across it and then I get no overlap? So zero voltage switching, either on or off, is arranging the switching operation such that I move the voltage around in a way that gives me little overlap between voltage and current at the transition. And we'll talk about how you might do that, but that's the idea. Any questions?

So what would a zero current turn on look like? Well, we said, OK, for turn on, usually, what happens is I have my switch current and he's rising because this is turn on so here's I switch. The current is going to go from zero and then I'm going to turn him on. It's going to be something. And normally, what happens is the voltage falls after the current rises. Right?

So what if I could-- what if I could do this? What am I doing? What if I could do this? What if I could say, all right, instead of having the switch current rise early, I did some magic game such that he didn't rise until later? I somehow delay the switch current rise until after the voltage has fallen. I would then have little overlap and that would be good.

Well, this is very similar to the notion of if I had my switch, you know, maybe I could go put an inductor in series with the switch. So even after I turn the switch on, the voltage doesn't appear across the switch. It could-- this is v switch. It could appear across this inductor. And the current in the inductor will be delayed in its rise. So in some sense, this snubber helps give me zero current turn on except that I'm dumping the energy. So I can play this same game except that then I got to figure out how to recover the energy in this inductor.

And you can see that these two zero voltage turn off and zero current turn on are kind of duals of one another. I'm just swapping voltage and current waveforms. And I'm swapping the role of the capacitor and zero voltage turn off and the inductor and zero current turn on. There are other ways to get ZVS turn off or ZCS turn on, but I'm showing you a common way to do it.

What's the last one? Well, we said I want to do zero current turn off. And we said what usually happens is that my switch voltage rises and my switch current can fall only after that happens. Well, what if I were really clever and said, OK, here's I switch? What if I were really clever and somehow I could make the current go to zero before the time I wanted to turn off? If the current were naturally zero when I was going to turn off, I can turn off the switch with no overlap in voltage and current. I'm done.

All right. I've got to come up with some magical circuit means to do this. We've kind of got ideas for how we might do these two. We've got to figure out how we might do these two guys. But that's the basic concept. I'm either going to manipulate the voltage to make the voltage and current overlap smaller or I'm going to manipulate the current to make the voltage and current overlap small or both for that matter.

But I'm trying to create this kind of magical opportunity to change my switch configurations at low loss. And that's not for free. If I just go to design a buck converter, I don't get that for free. But maybe I can do something more complicated that makes that work.

So does everybody, first of all, get the premise of soft switching? OK. Let me show you a few examples of different ways people do this. And I'm going to show you examples in DC to DC converters today. We'll talk about more ideas of doing soft switching in other kind of converters next class.

But here is the first example I'd like to show you. Now this circuit-- and I'm showing you an actual circuit because I wanted to show you the real circuit and the switching waveforms from it-- is a version of a buck converter. And you could do this same technique in a regular buck converter. It's just that this version of the circuit I'm showing you here is what we would call an inverted buck converter or a common positive buck converter. That is, the switch is ground referenced and the input and output share a common positive terminal instead of a common negative terminal, which they normally would. But you could see that this is still a buck converter. When the switch is on, I get V in minus V out across the inductor. And when the switch is off, I get minus V out across the inductor. It is a buck converter.

But how are we going to run this buck converter? Well, let me tell you how we're going to run this buck converter. And the one thing you can't see in this picture-- first of all, is there any question as to why this thing is a buck converter? OK. So here's the circuit. And it looks like this. So this is my buck-- this is my basic buck converter. If I drew this for you and put it on an exam and say what kind of converter that is, you'd tell me it's a buck converter. Missing something here. I'm missing my output voltage.

So here we go. Here's my output capacitor and my output resistor, for example. Here's V out and here's IL. OK. That's a buck converter. We're going to have to do some work to get soft switching in this converter. But let me tell you the one thing that's missing in the circuit diagram picture. And that is that there's a capacitance here, an effective capacitance here, across the device.

Where does that capacitance come from? Well, it comes from the fact that this switch has capacitance. And so that's in parallel with the device. The diode also has some capacitance. And that's connected between this node and a fixed potential. So from an AC perspective, it also looks like capacitance sitting here. So what we're going to do is we're going to use the net capacitance at this node to help give us zero voltage switching.

And I'll remind you that a capacitor across the device is just what we need to get zero voltage turn off. So we're going to do this trick because I've put a capacitor across the device. So if I see a design with a capacitor across the device, I think, yeah, I'm probably doing zero voltage switching. So how does this thing work? OK.

Well, I'm going to turn-- let's assume the switch is on. I'm in my normal state when the switch is on. The inductor current's going to rise. OK. We're going to let the inductor current rise. And we're actually going to have a large ripple ratio in this inductor current. And eventually, the inductor current is going to get up to some level where I tell the switch to turn off. OK.

Now what happens when this inductor current IL is big and I turn off? Well, if this capacitor's sitting across the switch, even though IL is big and coming here, this capacitor delays the rise in voltage on the switch. And I can turn off the switch under approximately zero voltage because of the capacitance of this node. All right. So that's exactly this picture here. OK.

So what happens, eventually, this inductor charges up this capacitor until the diode turns on. So I get a zero voltage transition. You can see that in this operating diagram here. The voltage across the switch goes up until it hits VN, which is exactly the condition in which the diode is on and the inductor is going to circulate its current into the output. Does that make sense to everybody?

Now what I have to do in this circuit is somehow ensure that this voltage across the switch be switch, which I guess I should have kept my color scheme consistent. V switch is going to get back down to 0. How does that work in this circuit? What magic do I do in the circuit? Well, in this circuit, what we're going to do is we're going to let the diode conduct. And because there's a negative voltage across the inductor, the inductor current is going to ramp down. Until the inductor current hits 0. OK once the inductor current hits zero, the switch is off. The diode is off. And all I essentially have is a resonance. If I think about V out as being constant, I have a resonance between the buck conductor and the capacitance at the switch.

So both devices are now off. By the way, this diode is turned off at zero current, naturally, because the current's ramped down to zero. OK. But what's going to happen then, the inductor and the capacitor are going to ring with each other and the voltage across the transistor is going to ring down. How far is it going to ring down? That makes for a great homework problem. But it's going to ring down to V in minus 2V out.

So if the output voltage V out was exactly half the input voltage, this node would ring exactly down to zero. If it was a little bit bigger, it wouldn't quite get to zero. If it was a little bit less, it would ring all the way into zero and the internal body diode of this MOSFET would turn on because don't forget there's a free diode right there.

So what that means is this circuit, as long as I let the inductor current get all the way down to zero and then the diode turns off, this voltage can ring all the way down to zero and I get the opportunity to turn on the switch when he already has zero voltage across them. The resonance between the buck inductor and the capacitor at that node give me a free zero voltage switching opportunity. The current in the device is zero. The voltages naturally run down to zero. And I can turn them on for free. OK. So I've got zero voltage turn off at the switch, zero voltage turn on at the switch. Any questions about that?

These are actually real waveforms. The green is the inductor current ringing up and down with this high ringing. The blue waveform is actually the drain voltage of the transistor. So you can see where it's on. This is the gate drive turning off. The voltage goes up. Then it eventually rings down so I can turn them on again for free. In this example, this was coming from 100 volts in to 35 volts out in this snapshot. And it was running at about 8 megahertz. So if you do your calculations for 100 volts input and running at 8 megahertz, you try to do that hard switch, you're going to be very unhappy.

So one of the things is, by soft switching this thing, we to get to very high frequencies while maintaining high efficiency. The efficiency of that converter at that point is about 97%. Yeah.

STUDENT: About how large is the parasitic capacitance across the transistor? And do you choose the transistor based on the parasitic capacitance?

PROFESSOR: That's an excellent question. In this design, you have to-- the switching frequency is going to be determined by the inductance value and by the net capacitance. If all of the net capacitance is provided by the switch in the diode, then you very much care. You can't throw a random switch in there. You need one with the right capacitance that will get you at the right frequency and everything else. If you added-- you could, and people often do add external capacitance, so that then that's the controlled capacitance and the device capacitance doesn't matter as much. But when you're doing this game, yes, you care much more about what the parasitics look like because you're actually using the parasitics here.

And I should say, by the way, I'm pretending that the device capacitance is a lossless capacitance. Were that it were so but you actually get some loss in that capacitance. So sometimes you would like to use a high quality external capacitor so that that ringdown is less lossy than it would otherwise be. But in this case, we pushed up the switching frequency so far that we could just do it with the device capacitance. OK.

Now. The fact that this inductor is resonating with the capacitor also means you've got to be very careful about how you size that inductor. You can't pick any inductance value. You've got to pick the right one. But it turns out to be an extremely small inductor because it has a lot of ripple in it. OK.

How do we control power in this converter? Well, basically, we can turn on the switch until the inductor current hits the current I want. It's like peak current control, if you will, but with high ripple. So this is a variable frequency converter. In fact, over its operating range, this converter varies between 5 and 10 megahertz.

So it's not fixed frequency. So if you're going to do this game in this design, you've had to design the whole converter and its operation, its control all around getting soft switching. So it's not like I can say I want my switching frequency to be this and just go design it with PWM. No. The whole thing is designed around the soft switch operation. Any questions about that?

Where might you do this? I'll show you one more example. I should say it has a limited zero voltage switching range. Technically, we need the output voltage to be greater than half the input voltage which is, you know, limits your buck conversion ratio that you can get with this converter. Maybe you can get down to a third of the input voltage before your switching losses become too big. It won't be quite soft switched but it'll be close.

So there are a lot of constraints around this. What's your switching frequency range? What's your allowed voltage conversion range? What's the power range? You've put yourself under a lot of constraints that may be a normal PWM converter doesn't have, but you get to go to a high frequency and make everything small and react fast.

This is another example. This is a different converter. This is a higher voltage. This is for a wide input voltage range, higher power. This is 1 to 4 megahertz. You can see that we're still getting very high efficiencies across a wide power range, so in this example 98%.

This is the converter it went in. It's actually a computer power supply. So it has a rectifier and then a pair of these buck converters. You can see them here. And then there's another converter to do isolation. And these are the twice line frequency energy caps. But the beauty of doing this, getting to this very high frequency, this is one kind of soft switch converter, the resonant transition buck converter. There's another kind of soft switch converter based on what's known as a dual active bridge. But this whole converter end to end, AC in to DC out, is like almost 96% efficient. It's 80 plus platinum rated efficiency. And it's way out there in the combination of what you can get in terms of efficiency and power density. In fact, if you look at these designs that are way out here, giving you high combinations of efficiency and power density, they tend to all be soft switch designs. If you don't care about power density, you can get high efficiency but make it big by using a low switching frequency. But if you really want to make something small and efficient, this is the way to go. OK.

That's the zero voltage switching example. And you can see, by the way, there's a duel between ZVS turn off and ZCS turn on and ZCS turn on and ZVS-- I'm sorry, ZCS turn off and ZVS turn on. OK. So you can do any combination of these techniques that you want. They're not all created equal, however. OK. Why?

All devices have some internal parasitic capacitance. So even if I turn this switch on with zero current coming from here, if he has voltage across them, I still dissipate energy in this switch. So if I want to think about the energy associated with the switch capacitance, I really need zero voltage switching. On the other hand, if I had a switch that was in a package with a huge inductance, turning this on and off the voltage will not help me if I have energy stored in this parasitic inductance. So depending upon whether I had a lot of capacitance or I had a lot of inductance, I might think about doing one or the other parasitically. To go to really high frequencies where many people in my group design converters, we tend to focus on zero voltage switching because that lets you push up the switching frequency in the face of the device capacitance as being very important. But that's a design decision. And in fact, that's why there are so many different techniques. I'm only showing you examples of these techniques. There's a lot of different ones because different devices, different applications, and operating range all might drive you towards a different circuit that would have advantage or disadvantage. And that's part of the art of designing a high performance converter is knowing which technique to pick in what application and what topology and what control mode to use.

Let's talk about this one. This is a zero current switch, quasi-resonant buck converter. This is part of a particular flavor of soft switch converters. And this paper reference here, this is actually a very famous paper from the late 1980s. And you will see commercial converters using the quasi-resonant technique, either ZVS quasi-resonant or ZCS quasi-resonant. It's not as popular as it used to be, but I thought I'd illustrate a zero current switch converter.

So if I look at this converter, what do I see? If I took away Cr and Lr and I just took them out, what kind of converter would this be? It'd be a straight up buck converter, right? What we are doing is we're putting a resonant inductor in series with the switch. So if you told me nothing else about this converter and I looked at that, I'd say, boy, they're probably using zero current switching because I got this inductor in series with the switch and I don't want to-- that will let me turn on the switch and hold the current low in it. That would be ZCS turn on, this trick. And then that probably means I need ZCS turn off, too, because if I turn off the switch when the current's in that inductor, I'm going to be unhappy.

OK. So how does this thing work? Suppose the diode is conducting. So this is acting like a buck converter. I'm going to commutate between the switch on and the diode on. But suppose the diode's conducting. This inductor here is just a giant inductor. So think of him as having a constant current. When I want to switch, I want to turn the switch on, fine, I just turn the switch on. He turns on under zero current because he's in series with this inductor. OK. So I get this kind of switching transition. He can turn on at low loss, at least ignoring his switch capacitance.

All right. Once the switch is on, now I have V in across the inductor and the diode's conducting so I basically have V in across this resonant inductor. And the current in the inductor is going to ramp up linearly. That's this picture here. The current in this inductor is going to ramp up linearly. Eventually, the current in the inductor is going to exceed the output current and the diode is going to turn off.

Once the diode turns off, basically, now Lr, though it's still at zero voltage, Lr is going to resonate with Cr here. And I'm going to have a current pulse I Lr that's going to do this. It's going to ring up. It was linearly charging and then it rings up. And then this voltage is going to get bigger than the input voltage and it's going to ring down. And eventually, the current is going to reverse and charge back this way.

Once the current is charging back this way, that means it's going through this diode. I can turn the FET off. When the FET's off, once I Lr goes to zero again, everything just stays off. And during that time period, this diode is going to be conducting again.

All right. So we're going to basically take a buck converter and put this Lr and Cr in it so that, basically, when I turn on the transistor, I get a resonant pulse of current and return to zero and then the diode is going to conduct again, essentially, because when this-- I should have been clear. This rings up, VCR peaks, and then VCR rings down to zero and the diode picks up the current. So what happens is I have a situation where this diode is carrying the output current except for pulses when I trip this device on. And depending upon how often I trip this device on, I get more and more pulses from the input.

So this is another technique. We're going to use resonant elements introduced into the circuit. It doesn't look anything like a PWM converter. I'm getting these resonant pulses from the input. And it turns out that the average voltage at this node, which determines the average output voltage, is precisely determined by how often I trip this switch on relative to the resonant frequency. So it's actually really a frequency controlled converter. OK. If I trip the switch on a lot, I get a higher output voltage. If I trip the switch on occasionally, I get a lower output voltage.

So why might I do this technique? Well, notice that the current in the switch is very smooth, very low EMI. That's kind of nice. The overlap switching loss is gone because of the zero current switching, although I still have the capacitive discharge loss in the switch. So it has some advantages. On the other hand, I can still get ringing between the device capacitance and the resonant inductance. I can still have problems with ringing. So this kind of technique is good to moderate frequencies but not super high frequencies. It might be good for a device type that likes to turn on and off with low current in it. OK.

Other limits-- because I have these pulsey kind of waveforms, I tend to have a lot of reactive power transfer in this circuit. It's not like I've got square wave voltages and currents and I'm kind of optimizing how power flows through the circuit. And I also tend to have increased device stress. So this diode, instead of being rated for the input voltage like a normal buck converter, might be rated for twice the input voltage.

So I've gotten some things. I've gotten some reduced switching loss, but I've had to entirely compromise everything about my converter to get it. Any questions about that?

Now I would characterize that as one flavor of soft switch converters where you design the whole converter around the switching regime to get you the nice soft switched operation. So how you control it, how you pick the components, all of that stuff is 100% determined by getting soft switching and keeping soft switching where you want it. So it's just like there's all kinds of classes of soft switch converter that you can talk about. But a big category is circuits that you design the circuit around getting the magic operating waveform somehow.

Another class is how can I augment a standard PWM converter with some extra circuitry that will let me run it just like a normal PWM converter, everything we've been doing all semester, but still soft switch the transistors, recover the energy in any snubbing, and still run it like a normal converter?

And I'm almost out of time, but I'll point to one example. And I'll just give you a notion of this. And in fact, this circuit or a close version is described in KPVS, the snubbers and soft switching chapter.

Here is the idea. If I ignored this Lr, D1, and S1, what would this circuit-- and Cr-- if I ignore these elements, what would this circuit be? Just be a standard boost converter. Right? And what we want to do is think about designing a boost converter.

So here's my standard boost converter. If I worked out what the switching losses of this thing were, I have, basically, the fact that he's going to turn on and turn off loss of this guy and then diode reverse recovery is also often a big deal.

OK. Well, what happens if I came along and said, look, I want to help so I'm going to go put a snubber capacitor across my device and give myself zero voltage turn off? This would still work just like a boost converter. I can turn them off. He captures the energy. And the only problem I have now is that when I turn the switch on, I'm going to dump that capacitor energy.

So maybe I can come up with a circuit that just when I want to go turn the switch back on will suck the energy off that capacitor and let me turn on the switch with zero voltage. And the way they do that-- very clever design-- is they put in another little baby boost circuit like this. And what it'll do is the orange circuit doesn't do anything. It's normally off. When I turn off the switch, the capacitor gives me zero voltage switching. But other than that, it looks just like a normal PWM converter.

When I get to the end of the cycle and I want to turn the switch back on, what I do is I quickly trip on the orange circuit and basically use this inductor to suck all the charge off of this capacitor and recover the energy on this capacitor, at which point I can turn on this converter with zero voltage switching-- turn on this transistor with zero voltage switching. And then the orange circuit, what I do is I turn this thing off and this inductor current goes to the output and gets delivered.

OK. So the orange circuit is, essentially, his entire job, he only runs once in a little for a very short time to basically recover the energy that's on the capacitor so I can get zero voltage turn off and zero voltage turn on. This is called a zero voltage transition converter. The main converter in white is soft switched. This converter is got zero current switching. But still, this guy's still turning on and off at zero voltage. He's got very high peak pulse power. But nonetheless, as long as I use this orange circuit at the opportune time, I can run the rest of the converter just like it's a normal PWM converter.

So I'm adding an auxiliary circuit comprising the green and orange circuit to let me operate the regular circuit just like a normal PWM converter. And that's a whole category of switch capacitor converters. How do I augment my normal PWM design so that it also gives me soft switching and I can get either a higher efficiency or higher frequency?

Any final questions for today? Yeah. Laurie.

STUDENT: What would be the losses on the orange circuit? Is that the same thing?

PROFESSOR: Yes, it is. I mean, he is taking the full peak current and he's operating into the full high output voltage. And so it's not negligible. But if you do it right, it's a lot less than your original high power circuit was. So he's got high peak power. You can optimize the orange circuit differently than the white circuit. And it will overall be better but it's not negligible loss. You still have to think about it.

OK. We will take this up next class.