

[SQUEAKING]

[RUSTLING]

[CLICKING]

DAVID PERREAULT: So you've gone through, and you've designed your converter, and you've worked out the magnetics and the controls and the circuit topology, all that stuff. And now it's time to put it on a board or put it in a package system and get the performance you want. And what you'll find is that depending upon how you do this, the details of how you put your circuit together, you can get great performance or terrible performance.

And so what I'd like to do is spend a little bit of time today just talking about practical implementation issues. And this is something that I'm-- what I'm hoping is you'll keep this in the back of your mind until the day you're actually building one of these things. And maybe the first one you build might cause you some unhappiness. And then you'll remember this lecture and look at it and say, what can I do to fix that? Or you can remember it all now, and then you'll never have this problem. And I certainly found over time that the more attention you pay to this kind of thing I'm going to talk about right now in your first design, like, when you first build the system, the less trouble you'll have with it getting it across the finish line.

OK, so let's talk about our friendly neighborhood buck converter. So suppose I have a buck converter. And let me assume I'm coming from some source impedance, that we don't have to worry about it. But I have a capacitor here to take the high-frequency input ripple current and hold the voltage steady. And I'll have a switch, and I'll have a diode, and my load inductor-- I mean, my output inductor and my load.

As far as the switching of this circuit is concerned, when the diode is on, the path through the circuit's like this. When I turn the diode off and the switch turns on, essentially, I'm commutating to this branch. So I have a commutation between this branch and this branch of the circuit, and I'm assuming all the high-frequency current's going through the input capacitor.

Likewise, when the switch is on, I'm carrying it this way. And when I turn the switch off, I essentially, from an AC perspective, I have a loop going this way. So essentially what dominates the high-frequency behavior in terms of high currents is this switching loop here. You can imagine that there's a large AC current running around this switching loop.

This branch has almost constant current. If you've done a good job of filtering, this input will have an almost constant current. But this is where the big AC current is.

So the first thing I think about when I'm putting together my circuit is how do I minimize that loop in terms of its inductance? And so one of the things we're going to talk about today is minimizing loop inductance. And you've got to-- the first thing you want to do when you look at any converter is, which loops do I care about? And this is the key loop.

So the first thing I'm looking at when I'm laying out my power stage is, what's my diode? What's my capacitor? I'm sorry, what's my diode? What's my switch? Here's the capacitor. How am I going to close this loop to minimize that?

And we've talked about how electromagnetic interference from that loop can couple into other systems. We'll talk about that in more detail.

We should also recognize that where is the high switching node in this circuit? This node has a big square wave on it with respect to ground. So this node flying up and down can inject charge into other places in the circuit and cause noise.

So making this node, this flying node, bigger than I need it to be is kind of a bad thing, because that can capacitively inject noise into other places in the system. So not only are going to think about how to make this tight. I also want to think about what's the size of this flying node. And if I imagine that I-- maybe I could place my-- maybe I could place my inductor over to the right and have a long lead going between the switching diode and the inductor. Or I could place, if this is my switching in diode node-- or I could put the inductor here and then have a long lead going to the output capacitor, these two things seem the same because the lead's either in series with the inductor on one side or another. But this node is flying up and down, and this node is about fixed. So, in voltage.

So to the extent that this is forming a capacitor plate to other stuff, and this is forming a capacitor plate to other stuff, this is much worse than this one. So you've got to think about both this loop and this flying node in a buck converter. In a boost converter it would be a different looping node, or in some other converter would be different. But that's the first thing you want to think about. And it can really have a big impact on performance.

And just to illustrate that, Rafa has actually set up a buck converter. And I'm going to let her show you four cases. OK. Why don't you tell me-- and since I have the mic, which case you're going to do first?

RAFA ISLAM: We can start with the worst-case one.

DAVID OK.

PERREAULT:

RAFA ISLAM: Have a big switch gap. And it'll be--

DAVID Yeah, so suppose I did a bad thing in 2-dimensions. I made this loop terrible, and I had excess area on this
PERREAULT: switching node. So I'm just doing a bad job of laying out my circuit. What would that waveform look like? OK, so we can see the first one.

And what you see is a whole pile of ringing in this circuit. So what she's looking at is you're looking at the voltage at the switching node, right?

RAFA ISLAM: That's right.

DAVID Yeah, so it's this voltage we're looking at, this switching node voltage. And what you see is there's all kinds of
PERREAULT: excess ringing there. Why? Because there's this inductance in this loop which is ringing with things, and then it's settling down. And that's pretty bad because that can inject noise elsewhere in the circuit. It can even interfere with the circuit itself. So you'll see a lot of people have problems, that they build their circuit, and maybe it works at super-low voltage, the power. And they turn it up, and everything starts mis-operating and doing funny things. It's because you're having interference with your own circuit. So not only can you interfere with the world, you can interfere with yourself.

Now that's the result of two bad things. Which one are we going to look at next?

RAFA ISLAM: We're going to move the capacitor closer to the switch.

DAVID We're going to have a bad switching loop.

PERREAULT:

RAFA ISLAM: Yes, smaller, smaller [INAUDIBLE].

DAVID Yeah. Right. So we're going to reduce the capacitance here. We're going to improve one aspect of that.

PERREAULT:

RAFA ISLAM: Sorry. [INAUDIBLE]. OK, that should be good.

DAVID So that's much more pleasant in terms of the amount of switching. You can still-- you can see it's still noisy.

PERREAULT: Stuff's still going on, but not nearly as bad as when I was really sloppy about both the loop inductance and the node capacitance.

Now we'll do the case where we have--

RAFA ISLAM: Small gap.

DAVID Yeah, the capacitor is bad, but I still have a big switching loop.

PERREAULT:

And again-- yeah, you got some ringing. You might not like that. You might start thinking about putting snubbers in there. But it's not as bad as that first case.

And then the last case is when perhaps better attention was paid to everything. And this is-- I don't want to say the ideal case, but this is when attention was paid to the layout of the circuit. So if you look at the two ends of that spectrum, you can see there's a pretty darn big difference between the actual performance of the converter.

Now this is in a low-voltage kind of low-power thing that it's not that bad. But when you're operating at hundreds of volts and hundreds of amps, these things get really, really bad. So you really want to think about layout. So thanks very much.

RAFA ISLAM: Sure.

DAVID So suppose I want to think about reducing the inductance of my loop? What does that mean in practice? Well, it means that I certainly want to place these three components close together such that if I think about the current flowing around in this loop, it forms a very tight area. So that's one way to think about it. If I imagine the current, AC current, circulating in this loop, how much flux is it throwing out in the world? And that relates to the inductance.

Sometimes, of course, I have to put components at a distance. There's some interconnection between them that I have to live with. Why don't we think about what is the parasitic of a pair of wires or a pair of traces interconnecting something?

So let's think about this example. Suppose I have two traces. Here's the first trace. And I'm going to carry current down it. And then I'm going to return the current somewhere else. So let me return it in a trace that's right underneath it, like this.

And I will close off the end. I'm going to drive current in one and return it on the other. And let's say it's got a trace of a width W , of length L , and a spacing that I'm going to call δ . What would be the inductance of that loop, where I'm going to carry current down one side and back the other?

And if I look at this end on, maybe I'm going to put current in one end, or maybe current will go in here. It goes down the conductor. Comes back the other side, coming out here.

And from that, what I'm going to get is some magnetic fields. And they're going to be very concentrated in between the traces. I'm assuming these traces are quite close together. And it's going to look something like this. Here's the magnetic fields. Here's the H fields. Looks like something like this, if I'm being schematic about it. Does that make sense to everybody?

So the question is, if these are the H fields, how might I analyze this thing? And let me do an analysis, a very simple analysis, just to say what are the trade-offs in the interconnect?

So we can say, OK, the integral of $H \cdot dL$ is equal to the integral of $J \cdot dA$, which is just the net current. I'm ignoring displacement current in this.

So what I'm saying is that if I integrate the H field around any loop, I'm going to get the net current enclosed. So suppose this total current here is I . All right. Well, that's fine.

The simplest thing for me to think about might be the fact that the fields are really concentrated in between the traces and they're kind of very diffuse out here. So maybe I can make the approximation-- and this is purely approximation that all the field that I'm interested in is in here. So then if I took this integral around this loop, I might say the integral of $H \cdot dL$ is just $H_{\text{sub } x}$ -- let me call this the x direction. It's the H inside the inside between here times-- this is the width of the winding-- is equal to I . Or equivalently, I could approximate this H field as $H_{\text{sub } x}$ is equal to I over W . All right? Has the right units-- amperes per meter.

How much energy am I storing in this thing? Well, again, I'm only going to pay attention to the-- really the energy I'm storing here. But that ought to be most of the energy. I'm going to ignore this sort of-- the fringing fields out here, even though I know they exist.

So that means I get the integral-- the magnetic energy storage is equal to the integral of $B \cdot H \, dV$, which is going to give me $\frac{1}{2} \mu_0 H^2$ times the volume, which in this case is going to be equal to $\frac{1}{2} \mu_0 I^2$ over W^2 squared, times the volume, which is W times L times δ . So I'm finding all the energy stored in between the plates of this conductor in return. Any questions about that?

So this is also equal to $\frac{1}{2} L I^2$. That's the energy I'm storing in the magnetic field. And if I work that out, this cancels. Then I get, if I just pattern match here, what I get is L is approximately equal to, with these simple simplifications, $\mu_0 L$ over W and δ .

This is the equivalent of saying capacitance is $\epsilon_0 A$ over d . I'm ignoring fringing between two plates. This is the magnetic equivalent of that. So the inductance of this trace to carry a current up and back is going to be basically $\mu_0 l$ over $w \delta$.

So if I want a low-inductance loop, what do I want? Firstly, obviously, I mean, I'm assuming there's no magnetic material in here. Clearly I don't want magnetic material in there if I want it to be low inductance. I want it to be short. Put everything as close together in terms of the length I'm carrying the current over as possible.

If I have to carry a current in a trace, I would like the trace to be wide. The wider the trace, the more the current and field spreads out, and the lower net energy stored. Because the width, the width helps me reduce the field magnitude. And even though I'm carrying it over a wider space, I have more volume, I get a lower inductance.

And then the other thing I can see is here I want δ to be small. So if I'm going to carry current out and return it, I want the spacing between those traces to be really low.

So the thing I'm thinking about when I'm thinking about what am I going to make this interconnect loop do, it's not only that I want-- I can think about, I can imagine this original thing I drew and say, yeah, I want this loop to be tight. But this might actually be in multiple dimensions that I have to think about. And if I have to carry current over some distance, carrying a wider trace and putting those traces closer together helps me make that inductance low. Any questions about that?

So if I'm-- an example of this would be not only for the switching loop, but the gate drive. I got to put the gate drive, and I want to send current out and bring the current back. When I'm thinking about that, if you want the energy storage to be small-- at least the magnetic energy storage to be small-- use a wide trace. Put the return very close.

Now, what would be the equivalent of that for this loop? What happens, the parasitic capacitance-- suppose I have some trace. This is my switching node, and this is, say, ground or something, or some other node. And I'm asking, what's the capacitance between these two pieces of metal that are very close to each other?

That, I think is more familiar to people. People recall that the capacitance is equal to $\epsilon_0 A$ over d , which in the way I've been writing this with these two plate sizes, would be ϵ_0 times-- the area would be width times the length divided by δ . Does that make sense?

So what does this say? In this case, if I want to make capacitance small, I also want to make the length of this thing short. But notice that whereas inductance goes as-- let's see, where do I have it here? L over δ . So let me write C . C is equal to $\epsilon_0 L$ times W over δ , right? If I thought about my inductance per unit length, is μ_0 over $W \delta$. Capacitance-- what did I do here?

AUDIENCE: I think [INAUDIBLE] top [INAUDIBLE] numbers.

DAVID PERREAULT: It should be. Yes. That's completely wrong. That's what happens when I skip boards. The inductance needs to be μ_0 . Yeah, all right. That's what I did. δ over W . If I want the inductance [INAUDIBLE] low, make W big and δ small. This should be W over δ . It's the induct-- I'm sorry-- the capacitance [INAUDIBLE].

Inductance per unit length is then $\mu_0 \Delta / W$. So L' , because if I have to do it over some length, then I can say what this is. And I want Δ to be small and W to be big. If I calculated the capacitance, the capacitance per unit length is $\epsilon_0 W / \Delta$.

So if I'm thinking about some trace in its return, or two pieces of metal close to each other, whatever I do to reduce the inductance increases the capacitance and vice versa. So in some cases, if I do the best thing I can do, I'm going to be trading capacitance for inductance, and I got to decide which I want, which I don't like less.

Now, that doesn't mean you can't-- this trade off-- and I should say, by the way, $1 / L' C'$ ends up being the speed of light squared, OK? So it's a physical constant. It's $\mu_0 \epsilon_0$.

But that doesn't mean I can't do anything bad here that doesn't help me. If I thought about, I'm going to carry some current over some distance. I'm going to carry current here and it's going to return underneath here, if I put some piece of metal out here that doesn't really carry any current, he can add capacitance without reducing inductance.

So I can certainly do things to make both inductance and capacitance worse, or make one of them worse without helping the other. But when I get it down to the best interconnect strategy, at some point, if I put my trace and return close, I'm making the inductance low and the capacitance high and vice versa. And you just got to think about which piece of trace is it and what is it doing for you. Any question about that?

Let's talk a little bit more about this switching node here. One thing it does, which we saw in the demo, is it'll cause a bunch of ringing because you have excess capacitance at this node that you don't like. But there are other capacitances that node, like the device capacitance and so forth. But what really often hurts you is the capacitance of that node not to ground, but to other stuff in your circuit.

So let's just imagine this node is driven by the switching. So there's a volt-- there's a square-wave voltage here. So if I thought about modeling this, I might say, OK, here's my big square-wave voltage. This is going up and down like this. And this piece of metal here forms one plate of a capacitor. So maybe I think about that as being one plate. Here it is. Here's one plate of a capacitor. All right.

What's the other plate of my capacitor? Well, it depends what's around it. But imagine nearby, I have in my circuit, for example, suppose I have some control circuitry down here. It's ground referenced. I have drivers and sensors and comparators. They're going to do all the computation for me.

And maybe within this there's some very sensitive node. Maybe that's the input of a comparator, or right where my error amplifier is, my op amp. And I can have some very high impedance nodes.

So maybe I have a plate there. That plate is just the metal of that node. And it has some resistance too. So maybe my comparator has some resistance to ground. So this white could be inside this box. Maybe there's a comparator here, and he has a positive input. And there's some metal that I've got sitting right here. That's the other plate of this capacitor. Does that make sense?

So if this voltage is here, what is this voltage? The induced voltage? Well, a simple RC divider, what does that tell me? It tells me I ought to be doing this. I ought to be seeing a voltage that spikes up and rings down, and then spikes down and rings up, and vice versa. So this is the voltage I'm going to see here. Does that make sense to everybody?

And maybe the energy in these spikes isn't very big. But if that's the input to a comparator, maybe I'm very unhappy, because if that comparator is determining when the switch turns on and off, and I get a spike on that signal because of this capacitive injection, I could be really unhappy. Any questions about that?

So what can I do about that? Well, one thing is we said the area, the capacitance is related to the area. I assume the two plate areas were the same here. This is some small area that's this node. This is some big area that's associated with carrying currents for my switching node.

One thing I want to do is make that as small as possible. Also, keep the distance between this thing and this thing bigger so that the capacitance goes down. So I can put them far apart, as far as I can practically do it, and reduce the areas. And that will help me minimize that capacitance that's causing noise. Does that make sense to everybody?

What else could I do about that? Well, another thing I could do about that is start to get smart about where those components are and what I put in between them. So one thing I could do is put in what's known as a Faraday shield, or lay out my circuit so I effectively get a Faraday shield.

What's a Faraday shield? Let me just say, this is a model for what's going on. Suppose I did this instead. I said, OK, I have my switching node. It's switching up and down. And I still have this-- maybe this is still a-- I minimized it, but I've still got a certain plate area there. And I've still-- I need some area for my-- I need some area for my switching node, and he's still going to be a high-impedance node.

The other thing I can do, by the way, is if I make this resistor smaller. That starts to become less nasty. But there's only so much I can do about that.

Well, what else could I do? Maybe I could come in and say, OK, let me put some piece of metal that sits between the noisy node and the thing-- the sensitive node. And maybe I'll just take that piece of metal and I'll ground it.

So what would be the circuit model for this? The circuit model for this would be the following. I have my-- I mean, this is sort of a circuit model already. But what I have is essentially one capacitor that's tied to ground. I have the capacitance between pink and green, and I'm grounding green.

And then I have the capacitance between green and yellow. So here's my other capacitance, that goes to my sensitive node. And here's now the voltage I'm inducing on the sensitive node.

What voltage would I be inducing due to this square wave in this case? Am I going to see a voltage over here in this circuit? No, right? I mean, he's got a grounded capacitor over here.

So, essentially, this piece of metal acts as an electrostatic shield. If it was all the way around, you'd call it a Faraday cage. But it prevents this noisy signal from injecting onto this.

Well, how would I do that? One way to do that is maybe I put this switching node on one side of my circuit board. I have a ground plane in between them, and I put my control circuitry on the other side of the circuit board that's sort of shielded or hidden from the high switching node. That makes sense to everybody?

Now you can't always do that. What else could you do? If you have really sensitive stuff-- so suppose I had a radio receiver right next to my switching node circuit. Is that a good thing? No, but some days you don't get a choice, right?

You can actually buy, for a printed circuit board mount, little shielding cans. So basically it's like it'll be-- it'll run around on the circuit board and provide a side edge, and then put a little cap on it. And you can put your whole-- your radio receiver or whatever, very sensitive noise, noise-sensitive element inside of there.

So actually, if you take apart a Wi-Fi router or something like that, you'll often see these little kind of metal boxes at the top will pop off, and some of the RF stuff will be in there. That's both to keep the RF noise in and to keep the external stuff out. But you could do that with your control circuitry too, if you had to.

But mostly what do you do? You try to be smart by not putting it close, or maybe putting it on the other side of the circuit board, and getting yourself an electrostatic shield in there. Questions about that?

I should give you a little warning. When I did this approximation, and hopefully get the right actual equation here, to get the inductance and the capacitance, in both cases, I was neglecting fringing fields. So when does this inductance equation really work well? It works well when you have to really kind of wide conductors put close together. So maybe I send my gate drive signal out and I return it right underneath. That becomes a pretty good approximation to the loop inductance you're going to get. Or in a capacitance, if I have two pieces of metal that are of the same size, and they're really close together, that's a pretty good approximation.

When I start to pull them apart, what happens? Well, in the capacitance case, for example, this approximation is really just capturing the fields that are directly between these plates. That's how you get to this equation. Once you start to have significant fringing fields, the distance between the plates is long compared to the areas or the lengths on the other dimensions of the plate, then these fringing fields start to become important.

So you can't necessarily think of, oh, jeez, my capacitance is going to go exactly as the inverse of the distance between whatever things I'm going to have. That's true when they're really close together, but less true as you start to pull them far apart.

So a classic analysis you might look at is like this. Suppose I took a sphere of radius r_0 . And I'm going to put a little charge on that sphere, plus Q . Total net charge on the sphere.

And let me make the other plate of my capacitor a sphere that's around it at some big R . That itself forms a capacitance, where I have electric fields that are-- I'm going to have mirror charges distributed on my outer plate, on my outer sphere, and I'm going to have electric field pointing outward here like this. Does that make sense to everybody, the picture I'm drawing?

I've got an inner sphere and an outer sphere. And I want to calculate the capacitance between them. And why am I doing spheres? Because I'm lazy and the math is easy. So think of whatever you got as being a little sphere.

But here's what I'm going to do. Let me make the outer sphere go off to infinity. So I can say, what is the capacitance of this little sphere of metal to the rest of the universe? Well, what is that? Well, we can calculate that.

If I take the integral of $\epsilon_0 \mathbf{E} \cdot d\mathbf{A}$ over-- the integral of $\epsilon_0 \mathbf{E} \cdot d\mathbf{A}$ is equal to the net charge I'm enclosing. So if I then figured this out, suppose I took a sphere around my inner sphere, and I said, what's the radial electric field? What is the $E_{\text{sub } r}$, the radial electric field, going to be? What is this electric field going to be at some distance r from the center? That's going to be equal to $E_{\text{sub } r}$ is just going to be equal to Q over ϵ_0 times the area, which is $4\pi r^2$.

What's the voltage that's going to be between the outer sphere and the inner sphere? Well, the voltage is going to be equal to minus the integral of $E_{\text{sub } r} dr$. Or I could write that, just to be easy, as infinity to r_0 -- the integral from infinity out to the surface of my little piece of metal that's got charge on it, of Q over ϵ_0 times $4\pi r^2$, which is just going to be equal to Q divided by ϵ_0 over $4\pi r$.

And that means that Q , if I write this, Q is just going to be equal to V times ϵ_0 times $4\pi r_0$. It's the radius of this little sphere, r_0 . But we know that Q is equal to CV . So this must be the capacitance.

So I'm kind of taking the other case, where all of my field is fringing field out to the universe. And in this case, the capacitance to the universe, from some little sphere of metal that I've got sitting here, is exactly the dielectric, which presumably, if it's free space, is ϵ_0 times 4π times the radius of the sphere. So it's related to the square root of the area of the sphere.

If I increase the area of this thing, the capacitance to the universe is not 0, even though the universe is infinitely far away. It's some finite amount related to the area or the radius of the thing I've got. So that's like-- this is basically a long way of saying, when I make things spaced far apart, fringing becomes very, very important, and I can get non-0 capacitances, or significant capacitances. And in that light, the only way you can really make that capacitance smaller is to basically make r not smaller, or basically make the area of the thing that you're worried about spraying field out smaller, so its capacitance to other stuff will be smaller. Questions about that?

So this is just a case where Physics 101 leads you to some conclusions about how you should lay your circuits out, which comes back to make your loops very tight. And I could make this-- I could have done the same analysis for finding inductances of loops and so forth. You make your loop tight. You spread the area-- you spread the current out to get the inductance low, and make the loop very tight. Don't use excess loop. Don't use excess trace area, especially at this loop and this node.

If you're going to go analyze a juke converter or some other kind of converter, which loop you care about is different. Which nodes you care about are different. But they're usually the ones with the square-wave currents in them and the square-wave voltages on them, or the high dv/dt or di/dt elements on them. And small changes can make a big difference. Questions about any of that?

Let me just talk about a couple of other things that I've glossed over, or just not-- it's not that I've glossed over. I've just not worried about them. And these, again, are practical implementation issues which are important, but you've got to worry about them after you worry about the rest of your system.

Let's think about driving the gates of our converters, gates of our transistors. So if I think about my MOSFET-- so suppose, I don't know, suppose I make a boost converter, for example. So I've been talking about buck converters. Now I'll talk about boost converters.

I'm going to have some circuit here, and I'm going to drive him. And usually what's going to be here is there's going to be some circuit which is making switching decisions. It's powered from some low-voltage supply. And it's got some input that's telling me when to switch this transistor.

What is the input of that transistor look like? Well, you know, this thing in him, the last thing between me and my switch, is some of gate driver, which is going to be powered from some voltage. And he's got some signal coming in. And he's going to try to turn this transistor on and off.

So if I thought about my transistor, here I'm going to drive this with some square-wave voltage. And I need to remember that this transistor is going to have some gate-source capacitance, he's going to have some drain-source capacitance, but he's also, importantly, going to have some gate to drain capacitance.

And this gate to drain capacitance actually has a significant influence, especially in the case when, suppose I'm driving this between 0 and 10 volts or something, or 0 and 5 volts if it's a lower, a logic-level kind of device. All right? But maybe this drain node is switching by hundreds of volts.

So this capacitance is the so-called Miller feedback capacitance-- this guy right here. Can have a big impact on the switching characteristics here. Because what happens? If I look at the charge I'm sourcing in here, and the voltage V_{gs} I'm getting. And this voltage, if this is a N-channel MOSFET, it's this voltage that-- this V_{gs} voltage is what's controlling how he switches on. If I looked at V_{gs} versus $Q_{sub\ g}$, the gate charge I'm putting in, what would happen?

When this thing is off, Before he turns on, I'm going to start charging the gate. And this is the capacitance my gate driver is driving. $dQ = C\ dV$. All right?

But what happens at some point is when V_{gs} gets bigger than the threshold voltage, I'm going to start turning on this FET. And what does that happen? This voltage, which was up at some high voltage-- this was the output voltage-- maybe that's 200 volts-- starts to swing down. And when it starts to swing down, I need to feed charge in from the gate to discharge this capacitor. And what that means is I'm going to be trying-- I'm going to be pumping charge in here. But the gate voltage isn't going to be going up appreciably. He's just going to sit flat. And all the charge I'm pumping in here is basically going to drive the drain down.

So this capacitor basically suddenly makes the equivalent capacitance I'm driving at my gate very large. Eventually, once I've pumped enough in, this node comes down to close to ground, this voltage ceases changing, and then this goes up like this.

So this happens somewhere near V_{th} , the threshold voltage of the FET. Does that make sense to everybody?

Why am I pointing this out? Well, two reasons. Firstly, when I think about the amount of charge I need to put in, maybe this is my threshold voltage. I don't know, whatever that is. Maybe this is 10 volts and this is 50 nanocoulombs.

So what I care about is I need to source this total charge to get to 10 volts. And if I only paid attention to what this capacitance was, I wouldn't really calculate the total charge I needed.

And how much energy-- how much do I lose from this gate-drive power supply? Well, we talked about this when we talked about switched capacitor converters. If I'm supplying, here's some value $Q_{\text{sub G total}}$, that supply is always coming from this gate-drive power-supply voltage-- 10 volts or something. So the amount of energy that comes in to charge up this gate is going to be equal to $Q_{\text{G}} \times V_{\text{gs max}}$, or V_{supply} , where this is my gate-drive power supply, V_{supply} .

If I wanted the power dissipation, that's how much I lose. Some of that energy goes to charging the gate. Some of it goes-- it gets dissipated in the resistors here. We talked about that when we talked about switch capacitor supplies.

But then in the end, what happens is I dump all that energy anyways when I turn the switch off. And so each cycle I'm going to lose $Q_{\text{G}} \times V_{\text{supply}}$, where my power is going to be $Q_{\text{G}} \times V_{\text{supply}} \times \text{switching frequency}$. Does that make sense to everybody?

So P_{gate} might be equal to $Q_{\text{sub G}} \times V_{\text{supply}} \times F_{\text{switch}}$ -- the number of times I'm turning on and off the switch per second. Notice that it has nothing to do with the-- necessarily with the gate drive internal resistances for the same reasons we talked about when we talked about switch capacitor converters. Any questions about that?

So let's think about-- so that's one thing I want you to pay attention about is you got to really look at these threshold curves to get your total loss. There's another thing you might want to think about when you think about the driver itself, and that's the following. What would this driver, this guy, look like inside?

Inside this driver, you might model as one MOSFET and another MOSFET. And then he's connected to the gate that he's going to drive. So I'm going to either turn on the top. So this is my gate driver. I'm going to turn on-- this is V_{supply} . So I'm going to turn on the top switch, and he's going to charge the gate, and then I'm going to turn on the bottom switch and he's going to discharge the gate and turn it off. Does that make sense to everybody?

Let's think about-- and by the way, sometimes-- let me just think about this. I'm going to simplify this as just being a capacitance, a simple linear capacitance. But what we know is it's anything but a simple linear capacitance. Maybe I'm thinking about this equivalent capacitance now for simplicity, this chordal capacitance.

Let's think about charging this up and discharging it. If I turn on, one of my tendencies might be to turn this on as fast as possible. But maybe-- let me just imagine I have some resistance here. Sometimes I have resistance in the driver, sometimes I have resistance in the gate, and sometimes I actually add external resistor.

Why would I add external resistor? We talked about EMI before. When I turn this transistor on, I am discharging this capacitance. So I have some capacitance here, and I'm discharging it into the switch. Capacitance on my diode. I'm discharging it, or I'm charging it into the switch. That causes a huge pulse of current in this loop which could cause a lot of noise if it rings a lot.

And yet, the amount I lose, at least in the capacitive portion, doesn't depend on how fast I turn it on, right? So at some level, maybe I'm motivated to turn on my switch somewhat slowly so that I don't discharge this capacitance really, really fast and cause a lot of noise.

So a lot of the time, we actually put additional resistance in here to slow up how fast I turn on the device. And some gate drivers will tell you what that is. Sometimes you put in a discrete resistor to do that.

What about turning off the device? When I turn off the device, well, he was on, so this already has 0 voltage. When I turn them off, this device isn't going to turn on instantly. It's not going to turn on until the inductor charges up this capacitance again. So in some sense, this capacitance, it hurt me when I turned it on, but now that I'm turning him off, he's actually acting like a little bit like a snubber and slowing down the rate of rise of this thing.

But the longer I take to turn him off, the more overlap I get. If the longer his current takes to fall, the more voltage I'm building up across this transistor and dissipating them. So it doesn't cost me anything to turn them off fast, whereas it cost me something to turn them on fast, in terms of noise. So very often we would like to turn on the transistor somewhat slowly, and turn them off really fast.

So you don't want-- this resistor, if I load it all onto this resistor, I don't want this resistor to be the same for turning on and turning off. And what you'll see in a lot of drivers is that what they'll do is, for turn on, this resistance, or equivalently the positive sourcing capability will be a bigger resistance. And the resistance for turning off will be a much smaller resistance. So you can turn the transistor off fast, but turn them on only slightly slow. And that will help reduce your noise.

If you don't like what that looks like, you can add your own resistor that will slow down the turn on. Maybe you don't want to slow down the turn off as much. So for the turn off, maybe what you do is you put a diode and a smaller resistor, or no resistor here. So he charges slowly through the yellow resistor and he discharges quickly through the green diode.

So you can control rate of turn on and turn off to control that switching and the noise injection elsewhere in your circuit. And you think about that when you buy your driver, and how you drive-- you connect your driver to your switch. Questions about that?

So I'm out of time to talk about these practical implementation issues. This is like, you could spend a whole course on this. But at least I hope you get a bit of a sense of how you might do some of these things in laying out your circuit. And so at least I've attempted this term to take you from what's the basic operation of a switching circuit, how do you design all the components, how do you control it, how do you deal with the filters, how do you lay it out?

And I'm afraid that's all we have time for. I wish you the best in finishing your projects, and a really nice summer. So have a great day.

[APPLAUSE]