

[SQUEAKING] [RUSTLING] [CLICKING]

DAVID PERREAULT: Which capacitor converters, which is covered in *Principles of Power Electronics* chapter 6. And we saw some very interesting phenomena in terms of the effective output resistance of a switch capacitor converter. And I'd like to take a few minutes just to talk about what are the origins of that behavior.

So let's just take a really simple case. Suppose I have a voltage supply that I'm going to call V_s . And I'm going to use it. And at $t = 0$ I'm going to close a switch through a resistor R to charge a capacitor, C . And at $t = 0^-$, we're going to say that has some initial voltage V_i on the capacitor.

So this capacitor I'm going to close the switch, and then I'm going to charge the capacitor up from V_i to its final value, which is going to end up at V_s . And if I plotted the behavior here, right, I could find, for example, if I plotted i_c , what would I expect to happen? Well it's 0 before $t = 0$, obviously. At $t = 0$, it's going to be $V_s - V_i$, which is V_i / R , right? So I'd have some initial value, $V_s - V_i$ over R .

And then it will exponentially decay down to 0 with some time constant τ is equal to RC . And so as we're charging up this capacitor, we dissipate some energy in the resistor, and we can calculate that.

What would happen if I changed the resistor value here? Suppose I used a much smaller resistor. So if I did that, what would happen to the initial current? Well, if I use a smaller resistor, the initial current would be higher. So here I'll pick-- this will be $V_s - V_i$ over R smaller.

And then it will also decay with the same RC time constant. So it'll go from 0. It'll jump up and it'll decay. Now the RC time constant will be shorter. So it'll do something like this.

So τ is again equal to RC , but R is smaller than-- the yellow R is smaller than the pink R . So it starts with a higher current and decays quicker. All right?

All right, well, let's calculate this stuff just analytically. So let's ask this. How much charge do I deliver in charging this capacitor? How much charge do I deliver from V_s to V_c ? Or what's the equivalent, the integral of i_c here, all right?

So we could calculate that as Δq is the charge I'm going to deliver to the capacitor, is the integral from 0 to infinity of $i_c dt$, right? And I can calculate that in a number of ways. I could write this $i_c dt$ as equal to dq , which is also equal to the integral of $C dv$, where this is-- where $C dv$ is going from V_i to V_s , which is my final-- it charges up eventually to V_s , right? So I could write this charge is just simply C times $V_s - V_i$. OK? That's the amount of charge that flows out of the source on to the capacitor to get it from its initial voltage up to its final voltage V_s . OK? Any questions about that?

So let's just ask, where's the energy going in this circuit? I said qualitatively, you know, we're going to dissipate some power in the resistor to charge up the capacitor. Let's calculate that. How much energy comes from the source. ΔE_{source} -- and I'm talking about the energy out coming out of the source to charge everything-- is simply going to be V_s times Δq , which is going to be $C V_s^2 - C V_i^2$. OK. That's how much energy comes out of the source to charge the capacitor.

How much energy goes into the capacitor? We can again use-- oh, here we can use energy methods very clearly. ΔE_{cap} is-- well, the energy that ends up on the cap is $\frac{1}{2} C V_s^2$. The energy we started on the cap before we close the switch was simply $\frac{1}{2} C V_i^2$ quantity squared. And that's how much energy, and that's how much change in energy is existing on the cap. Makes sense to everybody?

So that means that the lost energy, ΔE_{loss} , which, of course, some energy is coming out of the cap, some energy is coming out of the source, and some energy is ending up in the cap. The rest, obviously, ΔE_{loss} is in the resistor in this model-- there's nowhere else for it to go-- is going to be $\Delta E_{\text{source}} - \Delta E_{\text{cap}}$, which I could write as this minus this, which is going to be $C V_s^2 - \frac{1}{2} C V_s^2 - \frac{1}{2} C V_i^2$. And if I collapse this down a little bit and rearrange it, what I get is $\frac{1}{2} C (V_s - V_i)^2$. OK? And this is in the resistor? OK? Any questions about that? Yeah.

AUDIENCE: So in this model, like, you didn't assume that the resistor is in the circuit?

DAVID No, I did. Of course I did. Right? I mean, I still have conservation of energy, right? I didn't calculate an integral.

PERREAULT: Like, I didn't write the integral in terms of the resistor. But I still, I didn't-- it's not that I didn't assume it was there.

AUDIENCE: I guess I think you could have drawn the circuit and written down the same equations and arrived at the same result.

DAVID Absolutely right. Very good observation.

PERREAULT:

So what's going on here? Did I cheat? No, I did not cheat. I just-- I'm lazy and I don't like doing hard integrals. Or even easy integrals, for that matter.

What's going on here? It's exactly the picture I drew. OK. If I had some nominal value of resistor, 1 ohm, I'd get exactly this curve. And if I squared this-- this is a current. If I squared this, I could write the exponential decay, and then I could integrate it, and I'd get some loss. I would get exactly that loss. And if I change the value of resistance, I'd get this curve, and I'd get exactly the same loss.

So what's being noticed here is the-- none of these answers, the amount of energy that comes out of the source, the amount of energy, the one under the cap, and the amount of energy I lose, don't depend on the resistance at all. You get the same answer. The waveforms are different. It depends. The current waveform that I get depends upon the value of the resistor.

What happens if I keep making that resistor smaller and smaller and smaller? What I get is currents that have higher peaks and shorter durations, until, like, as R approaches 0, this more or less becomes a kind of impulse. But it's an impulse with an amount of charge Δq that $C V_s - V_i$.

And the amount of energy I lose in the resistor-- yeah, I have a smaller resistor. It's going to 0, but its current through, it's going to infinity, and it goes through it in just such a way that I lose the same amount of energy no matter what the resistance is.

So it's not that I'm cheating. It's just that this is what happens when you take a limit of a resistance, go to 0, but you're getting a limit as the current goes to infinity. And if I put that into the circuit and I took that integral, and I took the limit as R goes to 0, I'd get the same answer. So I'm just lazy, and I kind of skip over the integral. OK? Yeah. Go ahead.

AUDIENCE: In the case that there is a non-zero R, it makes sense that the energy that gets lost through the resistor [INAUDIBLE] R equals 0. Where is the energy lost?

DAVID
PERREAULT: The energy loss-- I mean, the actual distribution of the energy loss, now you're asking about idealized situations. Somewhere in this loop, I've got to dissipate energy. And so there has to be some-- it's sort of like saying if I connect a voltage source across a short circuit, what happens? Like, well, you've got to have some approximation. You've got to have some limiting resistance going to 0.

So in the real world, if this switch and resistor were an ideal MOSFET or something, where would the energy go? It would depend upon where the actual resistance was distributed in the circuit. And to the extent that the resistance is in the plates of the capacitor, it will get dissipated there.

But if you want to if you want to force ideality to where things are 0, then you can get squirrely results if you don't account for the fact that you have infinite currents. OK.

This kind of charge, or energy-- charge in energy methods are very convenient for working things, results out where you don't want to have to integrate things. And it turns out, in fact, it doesn't matter what this element is. Basically, I could replace this with a non-linear resistor. Like, so, if the MOSFET was nonlinear, you'd get the same answer because it doesn't matter.

The time waveform, the fact that I drew this as a perfect exponential, even if this resistor was a non-linear, so long as I ended up with V_c at the same voltage as V_s , I'm going to get the same dissipation. So in fact, even if this had some inductance, and it rang back and forth a little bit, and then settled down. I get the same total loss. It just doesn't matter.

AUDIENCE: This is just a result of the fact that this is in some sort of [INAUDIBLE].

DAVID
PERREAULT: Yeah, that's right. I mean, it satisfies conservation of energy, and it satisfies the fact that the change in the capacitor voltage determines the charge flow both here and here. Excellent observation.

What does this imply? Well, what is the efficiency of charging this capacitor? If I thought, OK, I brought this capacitor from some initial voltage up to V_s . What fraction of the energy I sourced got to the load, right?

So if I thought of the capacitor as the load. So I could think of the charging efficiency as being equal to the energy ΔE_{cap} -- that's the energy that went to the capacitor-- over ΔE_{source} . And what is that? Well, I could write that as equal to $\frac{1}{2} C V_s^2$ minus $\frac{1}{2} C V_i^2$ quantity squared divided by ΔE_{source} is $C V_s^2$ minus $C V_i^2$.

And I can divide out the capacitance here in this. OK. So I can write this as equal to-- if I factor out $1/2 V_s$ squared, or if I divide top and bottom by V_s squared, I can write this as $1/2$ times $1 - V_i/V_s$ quantity squared divided by $1 - V_i/V_s$, which I could write as $1/2$ times $1 + V_i/V_s$. That makes sense to everybody?

And in fact, actually, you can write it in a slightly different way. If I thought of ΔV as being equal to $V_s - V_i$ -- that's how much I'm charging up the cap-- this can also be written-- and I won't go through the math-- this can also be written as $1 - V_i/V_s$. Sorry. $1 - 1/2$ of ΔV over V_s .

So why am I dragging you through this? What I want to understand is how does this charging efficiency go? Well, if I charge the capacitor from 0, what would be the charging efficiency? Well, if V_i was 0, this would go to 0. This would go to this. The numerator would go to 1, the bottom denominator would go to 1. My charging efficiency would be 50%.

So if I charge the capacitor up from 0 to final value, I put $1/2 C V_s^2$ on the capacitor, and I dissipate half $C V_s^2$ in the resistor. So it's not a very happy energy transfer. I lose half my energy in charging up that capacitor.

On the other hand, suppose V_i was-- suppose the initial voltage on the capacitor was half of V_s . If I do this calculation, this is a half of $1 + 0.5$, or it becomes 0.75. So if I charge half way up from V_s to V_s , I have 75% energy efficiency in charging that.

And in the limit as V_i approaches V_s , the efficiency goes to 1. So if I want to have high efficiency, I only want to charge the capacitor a little. And the way you could think about that is, what's this happening? When I have the capacitor voltage is very small and the V_s is very big, I've got a big voltage across this resistor, I'm dumping a lot of energy in the resistor compared to the capacitor. Whereas, if V_c is very close to V_s , I only have a tiny voltage across the resistor. Most of the energy goes into the capacitor. So if I charge a capacitor only a little bit, I only lose a small fraction of the energy. Does that make sense to everybody?

I should have noted, by the way, that there is a way to make this circuit have greater efficiency. And what that would be is if I put an inductor in series here, and I slosh charge from V_s onto V_c , but let V_c finish and then reopen the switch when V_c is at some voltage greater than V_s -- so this limit on the efficiency that it doesn't matter how I charge it is only true if V_c ends up at V_s . If I make V_c somehow end up greater than V_s -- for example, through a resonant sloshing of energy-- then I can get higher efficiency. But in a switch capacitor converter, or in a converter where I'm going to let the final value of the capacitor go to V_s , the efficiency is fixed by the charge transfer in the initial value. Any questions about that?

So how does this relate to what we were talking about last time? Let's remind ourselves of the circuit that we saw last time. And it was a two to one switch capacitor step-down converter. When we say two to one, we're referring to voltage conversion, because that's just the way people do it.

So here's the circuit we're talking about. Here's V_s . I'm going to have a first switch 1. I'm going to have a capacitor where I'm going to have some voltage V_c on it. Another switch 1. And then this is going to go to some big capacitor. I'll call it C big. That's going to hold the voltage. And he's loaded with some current, I out.

And then this is one charging state, and then the other charging state looks like this. This is switch 2. And switch 2.

So what we do is, in, say, the switch 2 state-- I'm sorry. In the switch 1 state, we charge from V_s , charging this capacitor C up by the difference between V_s and V_{out} . And then in the 2 state, we discharge C into the output so that the output current is going to be always twice the average input current.

And we came up with a model. We derived a model for everything inside of this box. So I could take everything in here and come up with an equivalent average circuit model, a DC average model. It would look like this. I'd have V_{in} . And the rest of it I can model with a two to one transformer. And then a resistor, R_{eq} . And in the slow-switching limit, this is $1/4 f_c$, where this is the switching frequency.

And this is everything inside that box. And then I can model the output resistance. Output current, I_{out} . And this is V_{out} . That makes sense to everybody? Everybody recall this?

So notice that this is the slow switching limit, and the slow switching limit, R_{eq} on the log scale, versus F_{switch} on a log scale, goes as $1/4 F C$.

And why does it do that? Because for a given load current, the charge transfer on and off of this capacitance is exactly determined by the load current. I transfer a certain amount of charge per cycle through the cap with Δq coming out of V_s and ultimately $2 \Delta q$ going to the output.

And as I make that current bigger and bigger, that capacitor, if I make the current bigger and bigger, that capacitor has to charge up and down more, and my efficiency should droop. Or equivalently, as I make C bigger and bigger, for a given charge transfer, his voltage swing shrinks. Or if I turn up switching frequency, then the amount of time-- I have to transfer less charge in each switching cycle.

And the capacitor voltage doesn't move as much. And the resistance goes down, the charging loss goes down. Why? Because essentially, in each cycle, I charge and discharge that capacitor less.

So this reflection that the output resistance or the loss due to charging and discharging this resistance, this capacitor through the switch resistance, is inversely varied with frequency, or increases with load current, that the loss is exactly related to how much up and down I'm charging the capacitor. Any questions about that?

And then the only thing we said last time was, well, this isn't the slow switching limit where I have this kind of charging, right? The capacitor charges-- the capacitor current does this, and gets to its final charge state, right, assuming that C big is so big that it's staying constant.

If I turn up the switching frequency enough, I get on to a time scale that's short compared to this time constant. Maybe it's only doing this. The current never charging or discharging the capacitor voltage very much. And the capacitor voltage isn't moving around very much.

And in that case, I would get to a new limit. And we said the resistance in the fast switching limit-- so this is the SSL-- would look like this. And in the FSL, R_{eq} ends up being $4R_{switch}$ -- I'm sorry-- $2R_{switch}$.

And that makes sense, because in the fast-switching limit, I'm basically always charging the output through two switches. And so I ought to see two switch resistances in the fast-switching limit.

And if I said, well, what does it look like across the switching frequency for a given load across the switching frequency, it does something like this. It's one thing in the fast-switching-- slow-switching limit. It's another thing in the fast-switching limit. Any questions about that?

So if I'm going to draw more current, I'm going to get more droop. If I draw more current, there's more droop across this equivalent resistance. Or if I turn up frequency until I hit the fast-switching limit, I can make that resistance smaller and make the droop smaller.

Let's see a demonstration of this. So I'm going to pass this to Dave Otten, who's built us a really beautiful demo to illustrate this property.

DAVID OTTEN: OK. So right now, I'm going to-- the purple waveform which I'm about to turn off is my input supply, and the green waveform is the output of this converter. And right now I have no load on it, OK? So let me turn off the purple so you can see it.

Now the yellow waveform is the top of our flying capacitor. So you can see that it sometimes is up at the supply. Sometimes it's at the output. And then the blue waveform is the bottom of the flying capacitor. Sometimes it's ground and sometimes it's at the middle, OK?

So now I'm going to make sure we can see the green. And I'm going to turn on a small load. And you can see that the output voltage came down as a result. And if you look at the yellow waveform, you know, you can sort of see an RC kind of a time constant on it. And it looks like we could switch sooner.

If I increase the load a little bit-- so this is twice as much. You can see that that RC time constant might be faster. And you can also see that the output changed by quite a bit in terms of its output droop.

So now I'm going to start turning up the frequency. And we can watch the output increase. So there's the frequency coming up. Let me change the timescale a little bit so we can see it better now.

It's less clear, exactly, what the RC time constant is here. Maybe it's settling in the time that we give it, but maybe it hasn't quite made it yet. And you can see that the output voltage is rising. Oops, went the wrong way. So I can turn it up by quite a bit.

And this is sort of-- so I went from 5 kilohertz to 250 kilohertz. So-- whoops. Sorry, I changed the wrong knob. Here we go.

So at this high frequency, you can see that there really isn't much ripple on the output. We're doing a pretty good job of switching it fast enough so that the capacitor voltage doesn't change. So this is the most efficient point for this converter, particularly if you don't count the power that goes into the gate drives and the power that goes into switching the switches. That decreases the efficiency a little bit. But in terms of the power part of it, this is its best shot.

DAVID PERREAULT: So any questions about the fact that we can really get some benefit by turning up the switching frequency? We drop the resistance, which increases our converter efficiency, neglecting gate charge losses and that sort of thing.

And we would keep doing that. We'd usually like to keep doing that until we get up here somewhere. So we're at maximum efficiency. We wouldn't want to go out here, because there are other losses that I'm not counting. So there's always an optimum switching frequency.

And the nice thing is is the load decreases. We can often change the switching frequency to keep efficiency really high in these converters across a very wide power range. Any questions about that?

So let's consider a few other things about switch capacitor converters. One thing I'd like you to notice is, in fact, that if you have the right switch implementation, switch capacitor converters are bi-directional in their power flow capability, just like other converters. Remember, we could run a buck converter essentially backwards and turn it into a boost Converter That same trick works with switching power converters-- works with switch capacitor converters.

So let me just choose another converter. And what I'm going to do is I'm going to go take the output, put it over here, take the source, put it over here. I'll draw it backwards. And the only other thing is I'll change the voltage polarity here just so you can-- just for clarity in the circuit. And what would we get if I did that? It would look something like this.

So OK. Let me draw. I'll keep my source. And I'll keep my source on my load and I'll flip the circuit around.

So here we go. Here's V_s . OK. I'm going to have-- I'm first going to go through with a switch and a 1, and then up, switch 2. And then I'm going to have a capacitor. And I'm going to draw it as V_c , like this.

And then I'm going to have a switch that's a 1. And I've got to come to the output. So I'm going to have a C big. And I'm going to have a load current on this side.

And now-- so what am I going to get here? I've got to draw my other switch. So he's ending up on this side. So let me get just get this right. Switch 1, and switch 2 goes here.

Switch 1 comes from the source and comes here. Yes. So 2 and 2, like this. OK?

So what am I doing in this circuit? 1, 1, 2, 2. It's the same thing. All I did was basically flip the internal circuit around and kept the source and load in the same place.

So what are the two switch states here? If I have 1 and 1, what I'm doing is this. If switch 1s are on, what I got is this. And here's my C big. All right. And here's my I out.

And if I have my switch, so this is state 1. If I have state 2, essentially, I am doing this.

OK. So what happens? Like how might I think about the charge transfers here in states 1 and 2?

In this state, state 2, what I'm going to do, if I have a active load over here, I out, is I'm basically in state 2, I'm basically just charging this capacitor from the input. So this capacitor V_c is going to charge up to V_n or to V_s .

Then, in state 1, what I'm going to do is I'm going to discharge him like this. So I'm putting something that's charged to V_s , in series with V_s , and I'm discharging it into the output.

So the output ideally ought to charge up. If I had no load current, what will happen is the output will charge up to $2 V_s$ or $2 V_n$ -- let me call this V_n just for consistency.

So this becomes a 1 to 2 converter. All right? I can do the same kind of analysis on this converter that I did on the converter from last class. And if I did that, we would get this equivalent circuit model. I would get V_n . The equivalent circuit model for this whole thing would be like this. 1 to 2.

And then I would have an R_{eq} that would be equal to $1 / (F \text{ switch } C)$ in SSL. And then here's my load current, I_{out} .

It's the same kind of analysis. We do the same charge transfer thing. We get an equation. We see that this circuit model right here matches what goes on in the slow-switching limit. Does that make sense to everybody?

The neat thing about this is that actually the model hasn't changed. I'm using idealized switches. Let's just say I assume they're perfect switches that can block and carry bidirectionally. So I have a full switch implementation, full four-quadrant switching implementation.

I could take this circuit model and reflect this R_{eq} through the transformer. And what would I do when I go through the transformer? I have to scale impedances by the square of the turns ratio.

So another model I could develop for this, an equivalent to this model, is the following. I could put a resistor here that would be scaled by $1/4$. So I'd have $1 / (4 F \text{ switch } C)$. 1 to 2. And this would be my alternative circuit model. And now here I have I_{out} .

So here's my point. This is just another way to draw the circuit model, same circuit model. I've just reflected the resistance through the transformer.

And if you look at this, you're going to see it's exactly equal to the one up there. It's $1 / (4 F \text{ switch } C)$.

So the energy loss doesn't really depend upon which way we're transferring power through the converter. In other words, am I sticking current in here and delivering energy to here or vice versa? I'm going to get the same efficiency in either direction-- so long as my switch implementations and so forth can support that. Any questions about that?

So once you've analyzed the 2 to 1, you've also analyzed the 1 to 2. You just don't know it yet. We could have just, like, swapped the source and the load and been done, and then I would have probably slid the resistor through the transformer just to put it into standard form, and we would have been done.

Let's talk about-- let's talk about challenges of this converter. If I thought about building this converter, OK, let's think about where the physical components are big.

This capacitor here is my energy transfer capacitor. And we've already said, if I can make him a bigger energy transfer capacitor, the bigger C gets, the smaller the resistance gets. So it's good to have this capacitor be pretty big relative to the amount of power I'm pulling out of it. So I can be close to the fast-switching limit.

On the other hand, I've treated this capacitor here as if it was infinite. So if I do nothing else, usually this capacitor C big is big compared to C .

And his job here, if you notice, is that he's holding up the output. And one of the two states, in this state, there really is nothing else making V out constant except for C big. He has to be big enough so that V out's not going to droop while this guy is in his other state. Does that make sense to everybody?

So when you size the capacitors in this thing, if I built this circuit, usually C big would be substantially bigger than C , which is in some sense a shame, because I'm using C big only as a filter but not to transfer energy through the converter.

So the question is, is there something clever I could do about this to make better use of my capacitor energy storage? And the answer is, there is. And here's the trick.

Suppose I was to build this circuit. And let me just-- let me just draw-- I'll redraw this circuit. So here's V_n . And here's my circuit. I'm going to redraw this just up here.

And I get-- here's $2 V_c$. Here's 1. I'm sorry. This is switch 1. This is switch 1. I also have a switch 2. And I have a switch 2.

And he comes down to this ground input, right? And then, I have C big over here. And here's my output. And I'm going to deliver energy here.

What is this current? What does this current look like? Let me call this-- I'll call this i_{x1} .

What does i_{x1} look like? Well, let me think about, in the fast-switching limit, what i_{x1} looks like. Well, when I'm in the fast-switching limit, I'm dumping some charge off of V_c into the output. And in the fast-switching limit, this capacitor is not charging or discharging much. So the current flow is limited by these switch resistances. So in the fast-switching limit, i_{x1} looks like this. So this is state 1. And then in state 2, he's just 0 because in stage 2, the switch is open.

And so I dump a bunch of charge into C big in state 1, and then nothing in state 2. So C big's left to take this total charge and then deliver the average current to the output.

Well, what would happen if I came along and said, you know what? Let me build a second converter. I'm going to build the complement to the first converter. So I'm going to build something that looks like this. He'll be my orange converter.

So maybe what I'll do is I'll build a second converter exactly like this. And I'm going to make this turn on in state 2. I'm going to have another capacitor, V_c prime. OK. Here's state 2. And now I'll deliver his charge to the output. And I'll call this i_{x2} .

Now in state 1 I'm going to do the complement. So I'm going to do this. And here's-- I'm going to have this guy in state 1 and this guy in state 1. All right.

So all I've done is I've built a copy of this circuit, but I've switched the switches exactly oppositely. Does that make sense to everybody?

So what is i_x ? Maybe I should call them i_x prime here. What is i_x prime? I guess I call them i_{x1} . I guess I will call this i_{x2} . I have i_{x1} . i_{x2} is going to be the exact complement of i_{x1} because he happens in the opposite portion of the cycle.

And then if I said, well, what's the total current coming in? What is this current i_x ? Well, i_x is just now doing this. He's following i_{x1} . He's following i_{x2} . So he's almost a constant current into the output. And so this capacitor is only filtering the difference between this and the load current.

The load current I_L , or I_{out} , is just the average of this current. So I_{out} looks like this, all right?

And so this circuit really doesn't need C big anymore. He only needs a tiny capacitor here. So I could take this giant capacitor and divvy them up among this guy, if this is C prime and this guy is C . And basically, instead of having an energy transfer or capacitor and a giant filter capacitor, I can have two pretty big energy transfer capacitors and a tiny filter capacitor. And I can get the same behavior, the same performance, for much less capacitance or lower output resistance for the same total capacitance. Any questions about that?

This trick is very commonly used. This is called interleaving. So this is an interleaved 1 to 2 step-up switch capacitor converter. And it just has the benefit that you can get the same performance for much less total capacitance because I don't need to have some filtering-- any filtering going on.

And I should say, by the way, this is a 2-phase circuit, with states 1 and states 2. You can have more phases. There are other games you can play. But this is the basic one.

What's the penalty you pay? To do this, I more or less have twice as many energy transfer capacitors and twice as many switches. And in a world where the count of the components matters, that's a pain in the butt.

On the other hand, if I cared about the size of my energy transfer capacitors, this thing is wonderful. And this is a very common technique, is interleaving your switch capacitor converters to alleviate the need for big either input filter capacitors or output filter capacitors. Questions?

From a DC perspective, all we do is basically put two of those circuits in parallel, and I get half the output resistance essentially. But what I've gotten rid of is the filter capacitor that's not in that DC model.

Now I should say, I've kind of spent all my time talking about 2 to 1 or 1 to 2 switch capacitor converters. There are a lot of different kind of switch capacitor converters. You can expand out to almost any rational conversion ratio.

And keep in mind, as we said, the rational conversion ratio gives you perfect conversion in current-- you know, the this current is exactly twice this current, right? So you get a perfect current conversion ratio and an approximate voltage conversion ratio.

But you can almost get any ratio, by picking the switching states in the capacitors, you can get almost any ratio here that's rational-- like 2 to 1, 3 to 1, 4 to 1, 3 to 2, you name it, OK?

There are bounds, and there's different ways to expand on this basic idea. Let me just show you a couple for-- out of many. Here's one idea.

Suppose I did this. Suppose I had my input source V_n , and I did this. I will transfer energy to a capacitor here, and then to here, and then to here. And then I'll have my C big here. All right.

So when I close this and switch 1, 1, 1, 1, I'm going to charge these capacitors in series between the input and the output. So I'll get a charge flow that does this. And I'll charge up all my capacitors to the output.

And then I can add another set of switches that are like this. And then maybe these guys will discharge in parallel to the output. So I'll get charge flows that flow out of this capacitor like this out of this guy, like this, out of this guy like this. So I'll charge them in series and discharge them in parallel to the output.

So in the final state, these capacitors will charge to V out. And then so I would get, in the state 1, I would get V out plus V out plus V out plus V out. So four V out at the input.

So this is a 4 to 1 step-down series parallel switch capacitor converter. So we charge the capacitors in series. We discharge them in parallel. Or, if I turn this around and swap where the source and load were, it would be a 1 to 4 converter. That's one kind.

There are many others. And I'm kind out of time. So I'm not going to draw any of those. There are others listed in the book. And I show another one. I show a ladder converter in my notes.

What distinguishes these converters is how many switches and how many capacitors do you need to do a certain conversion function. Some are better than others. What are the voltage ratings of the switches and capacitors? How much energy storage do you need to do a given conversion function? And whether they're the same or not.

So in this circuit, for example, this switch has different ratings than the other switches for example. So you might choose this one if you really cared about all the capacitors having the same voltage, but you weren't so worried about the switches being different. You might choose a different one if you had different constraints.

So there's a bunch of different converters that are popular-- series, parallel like this, one ladder, Dixon. There's a whole variety of them, that you pick the one that matches the design regime you're in. If you're in an integrated circuit, maybe you care about one thing. If you're really high voltage, maybe you care about another thing. Any questions about that?

I should also say that while I'm showing you DC to DC converters, you can also make switch-capacitor inverters or switch-capacitor rectifiers. And in fact I haven't drawn the right circuit. I'm kind out of time. It's in the notes. But if you can take a circuit that has an AC-driven input with diodes, and charges the capacitors in one part of the cycle, and then discharges them in the other in just such a manner that you shuffle charge up and get a rectifier that multiplies the voltage-- The so-called famous Cockcroft-Walton voltage multiplier is one example of a switch capacitor rectifier. And those are very often used in high-voltage power supplies where you want to generate very high voltages. And that's why, in fact, Cockcroft-Walton got the Nobel Prize for splitting the atom with really huge voltages generated by just such a rectifier. So while I've spent my time talking about switch capacitor DC to conversion, you can do inversion and rectification. And they all follow the same fundamental principles.

So I'm kind out of time to talk about this, but I'd be happy to take any final questions.

OK. In the next class we will take up a new topic. Have a great day.