

[SQUEAKING]

[RUSTLING]

[CLICKING]

DAVID
PERREAULT: OK, why don't we get going? What I wanted to do today was continue our discussion of resonant converters and specifically talk more about resonant DC-to-DC converters. And last time, just as we were running out of time, we talked about rectifiers, right.

So I might have a rectifier for a resonant DC-to-DC converter, and I might like to model that rectifier as a resistor because everything I've said about resonant inverters, right, we're talking about, say, an LCR tank, so the question is, What's the resistance we have?

So for example, suppose I just have a classical full-bridge rectifier. OK. And let's assume that I'm going to derive this from a perfect sinusoid, $I \sin \omega t$. OK. Where am I getting this $I \sin \omega t$?

This is going to be coming from my resonant converter, and if I have the fundamental harmonic approximation, for example, in a series resonant tank, I can think of the tank current as being perfectly sinusoidal. So here's the sinusoidal current, right, and maybe I can say, OK, I have v_{rect} here as the input voltage to the rectifier.

Now, I'm going to consider a capacitively loaded full-bridge rectifier, OK, where we have C big here. OK, so the output voltage v_d is approximately equal to its DC component, OK, and maybe I have a load resistor R_{dc} over here. OK. So then the question is, What do the waveforms in this thing look like?

And I could plot, for example, this current. I'll call this i_x . i_x , right, if this is a sine wave, when the sine wave is positive, these two diodes are on, and i_x looks like $I \sin \omega t$. If I is negative, these two diodes are on, and i_x looks like $-I \sin \omega t$.

So i_x looks just like the rectified sine wave, right, where this is 2π over ω in terms of this being t . OK. And the peak value here is just I , right. And if I said, What is the average value of i_x ? a simple analysis will show that the average value of i_x here, OK, will be equal to 2 over π times I . It's a pretty straightforward analysis.

OK. Why do I care about the average value of i_x ? Because clearly, in periodic steady state, the average current, it goes to the load resistor, and the ripple current goes to the capacitor, is the idealization. OK, and we know we don't get any DC current in the capacitor in periodic steady state.

So if I wanted to analyze this thing, what do I get for output power? The average output power is simply equal to - the average output power is simply equal to the average of v_d times i_x , which, if this capacitor is really big and v_d is approximately constant, this is approximately equal to capital V_d times the average of i_x , which we could say is equal to 2 over π V_D times I . And this is the power that's going to the load. This is the DC power.

And of course, if these are just ideal diodes, then all that power came from the AC side, right. Looking at the AC side, what do I see on the AC side? What I see is a voltage v_{rect} that looks like this. OK. It's a square wave. And this is, again, 2π over ω , and this is π over ω .

All right. This is v_d , and this is minus v_d , OK. And so the power on the AC side, right, would just be the integral of $i \sin \omega t$ times v_{rect} . OK. But we know that we only transfer power through voltages and currents at the same frequency, right.

So if I make the fundamental harmonic approximation, then what I can say is, look, the only component of this voltage v_{rect} -- and, boy, he doesn't look exactly square, but let's pretend he's square, OK-- the only component of v_{rect} that generates power is the fundamental, right, because the current's fundamental. So it's only going to be the fundamental current times the fundamental voltage that are delivering power, or real power, and all other power components have to be reactive, right.

So I might, then, just think about what $v_{\text{rect } 1}$ looks like. $v_{\text{rect } 1}$, if I did the analysis, the fundamental of a square wave is $4/\pi$ times the peak value of the square wave, right, so I might think of about a waveform that was the fundamental component, where this peak value is $4/\pi v_{dc}$ -- or $4/\pi v_d \sin \omega t$ is equal to $v_{\text{rect } 1}$ of t . Does that make sense to everybody? I'm sorry, this peak value, this peak value is $4/\pi v_d$, and this whole waveform is $4/\pi v_d \sin \omega t$.

OK. So what I get-- with ideal diodes, at any rate-- is a square wave that's in phase with the drive current. Remember, the drive current looks something like this. This is $i \sin \omega t$, right. So here's $i \sin \omega t$ -- $i \sin \omega t$.

So what I get is a voltage whose fundamental component is in phase with the current, right. So for power transfer purposes, I might as well think of this thing as a resistor, right. Essentially, it's a describing function approach. It's the amplitude of the voltage divided by the amplitude of the current.

So I might think of an equivalent rectifier impedance, like looking into the rectifier from the AC side. R_{ac} , OK, is going to be equal to $v_{\text{rect } 1}$ over the current, which is going to give me $4/\pi v_d$ divided by i . OK. This is the equivalent AC resistance that relates the fundamental drive current with the fundamental of the voltage. Any questions about that?

Now, it turns out I can write this a couple of other ways which are useful to think about. One of these is as follows. Because I have this relationship between the current and the power, I could write this as being equal to $8/\pi^2 v_d^2$ divided by the power.

OK. So why would I think about it that way? Well, let's pretend I don't know what the resistance looks like. But I know I'm delivering a certain amount of power into the load, and I know I have a certain DC output voltage. Then I could express the equivalent AC resistance, owing to this rectifier, based on $8/\pi^2 v_d^2$ over power.

OK. Or equivalently, I know the relationship between the average current and v_d is related to-- the average of i_x times R_{dc} has to be v_d , right. And if I apply that relationship, I can also write this as equal to being equal to $8/\pi^2 R_{dc}$. OK. It's just a question of whether-- if I had some other load besides a resistor, I couldn't write it this way, but I could always write it this way. OK. Any questions about where I'm getting that notion of an equivalent resistance from?

So I can take a rectifier and model it as an equivalent resistance for purposes of power transfer, assuming that it's being driven by a perfectly sinusoidal current. OK. And that's going to become really handy. Why is that going to be handy? That's going to be handy because we would like to model the behavior of a DC-to-DC converter, OK, and what I've been talking about up to now is series resonant inverters.

Right, so let's think about what a series resonant inverter-- I'm sorry, a series resonant DC-to-DC converter might look like. It could look like this. Maybe I have an input voltage v_{dc} , and then I will use a half-bridge inverter.

Right, so this is just a simple half bridge, OK, and I'm going to call this v_{inv} . That's the inverter voltage, v_{inv} . What does v_{inv} look like? v_{inv} looks like this.

OK. Here's v_{inv} . It's modulating between-- someday, I'm going to get to this 50% duty ratio here-- he's modulating between v and 0, and this is 2π over ω_s , whatever my switching frequency is. And this is π over ω_s in terms of time. OK.

All right. Then I'm going to connect him up to some tank, right, and in this case, I'm going to use the series resonant tank. Right, so I'm going to have a capacitor C , an inductor L . OK. And I could kind think of that as running into whatever my net load is. OK. So let me just call this voltage over here $v_{sub r}$.

OK. And then maybe I go and I say, OK, I need a transformer here in the middle of my AC converter because now I have AC waveforms. Let me put in an n_1 to n_2 transformer. OK. And I'm going to use this thing, and let me call this voltage v_{rect} . And I'm going to run this into my rectifier. So here's my rectifier.

All right, so what am I doing? I'm going DC to AC, running through my frequency-selective resonant tank. All right. Maybe I have an additional voltage transformation in here. Here, I'm doing it with a transformer.

You can do it with other kinds of transforming elements, too. I'm running into a rectifier and getting my output voltage-- so I come from a DC input to a DC output. OK. Any questions about the structure of the system?

So because I'm going to make the fundamental harmonic approximation, I'm going to assume that my current here, $i_{sub L}$, is approximately sinusoidal, that this LC tank is high enough Q when loaded by this rectifier system that I can ignore all the harmonic currents in i_L . OK. If I do that, then the only thing that matters for power transfer, just as the only thing that mattered on the rectifier was the fundamental of the rectifier voltage, the only thing that matters on the inverter side is the fundamental of the inverter voltage. Right, so maybe I could just go focus on, for example, $v_{inverter 1}$. It's fundamental.

All right. Well, what does the inverter fundamental voltage look like? It's a square wave between 0 and v_{dc} , but its DC components are going to get blocked anyway. So it's between plus v_{dc} over 2 and minus v_{dc} over 2, times 4 over π for its fundamental, and I'm going to get an $v_{inverter 1}$ that's basically equal to an amplitude of 2 over π v_{dc} .

And this is $v_{inverter 1}$, OK, and this is the only thing that matters in terms of transferring power. OK. So maybe, then, I could think of this system as being $v_{inverter 1}$ right, the sinusoidal fundamental component of the square-wave voltage on synthesizing, running into a tank, first C , then L , right, and then I'm going through a transformer, OK, and back into a rectifier.

But how did I model the rectifier? I modeled the rectifier as this R_{ac} . Right, I just did this. What's my value of R_{ac} ? It's right here.

OK, so I get some value of R_{ac} . This is N_1 to N_2 , right. Maybe I could say, What's the value of $R_{sub\ r}$ here? What's the resistance I see in here? Can anybody tell me what that is?

If I had-- if this were, honest to god, a resistor here, right, then the ideal transformer scales impedances by the square of the turns ratio, right. So that means I ought to scale this up by N_1 over N_2 squared. So this ought to be $R_{sub\ r}$ ought to be N_1 over N_2 quantity squared, R_{ac} , right.

So I could think, looking this way, I could imagine I'm just seeing an AC resistance scaled by the transformation ratio. And all of that, so long as $i_{sub\ L}$ is sinusoidal, all of this preserves power transfer in the system because we're only transferring power through the fundamental. Any questions about that? Yeah?

AUDIENCE: So is the voltage gain all coming from transformer?

DAVID PERREAULT: That is an excellent question. No, it's not, but that's precisely-- you're very astute, because that's precisely how we want to think about the system, is in terms of the voltage gain. We get a lot-- we may get a lot of the voltage gain from the transformer.

That's why we put-- that's one reason why we put the transformer there. One reason is I just want this voltage and this voltage to be isolated, so I get isolation. But I can also-- if I have a high-voltage input and a low-voltage output, maybe I do a big step-down here.

But I also get some, some, voltage transformation from the tank, right. So maybe I think about this thing in the following way. I have some $v_{inverter\ 1}$, right, then I have my LC network. And let me just preserve the order of my components, even though it doesn't matter.

If I'm going to operate above resonance, this series resonant tank looks net inductive. Right, If I'm going to go above this natural resonant frequency, ω_0 1 over the square root of LC, this looks going to be-- whatever the cancellation between C and L is, this is going to give me some of L effective, which is some equivalent impedance. And so this is going to give me some net impedance, which then voltage-divides with $R_{sub\ r}$ to give me a voltage I'm calling $v_{sub\ r}$, or $v_{sub\ r\ 1}$ in this model.

So this is $v_{sub\ r\ 1}$, right. So I basically have a voltage divider between my inverter-generated voltage, the voltage generated by a half bridge, divided down by this with $R_{sub\ r}$, right. And that's what we said, was one of the nice things about a resonant tank is the value of this effective impedance varies very quickly with frequency.

That's why the admittance looking in here is varying very quickly with frequency. And hence, I get a frequency-dependent voltage divider between the effective resistance posed by my load and the resonant tank components themselves. OK. Excellent question. Any other questions?

So how might I think about that? I might think about this gain. I might think of about my resonant tank gain, my resonant tank voltage gain, just owing to the operation of the resonant tank dividing with the equivalent load.

I'll call that some gain $G_{sub\ t}$, which I'm defining as being v_{r1} , the magnitude of v_{r1} , divided by the magnitude of $v_{inv\ 1}$. OK. And this resonant tank gain is going to be frequency dependent, OK, because as I go up in frequency, if I were dead on resonance, what would the resonant tank gain be here?

If I were to run this to where L and C exactly cancel in impedance, right, this would be a short circuit at the fundamental, and I'd have a gain of 1 between $v_{\text{inverter 1}}$ and v_{r1} . If I go above resonance or, for that matter, below resonance, this becomes a nonzero net reactance, and I get voltage division and I get a lower gain. OK. So I'm going to have-- this tank gain is a function of frequency.

OK. The other thing I need to think about is, What is the effective quality factor of this tank? Right, and we said, when we're talking about the quality factor, what we're really meaning is, How sharply tuned is the behavior change? And I've got to be a little bit careful about that, but we could say that the Q_0 , the effective quality factor of the tank, I might write as the square root of L divided by C divided by $R_{\text{sub } r}$. Does that make sense, everybody?

So how does this play out in practice? First of all, let's recognize that this tank gain is a function of frequency, right, and it's also a function-- and the tank quality factor, the effective tank quality factor is, because it's a function of R_r and R_r is a function-- is determined by R_{ac} , it's also a function of power as well as voltage.

OK. So how might I model this whole DC-to-DC conversion system? OK, let's break it up, and now I'm going beyond like-- sort of a resonant inverter's very simple. I just have to think about my gain from my inverter to my fundamental and what happens and how I control power. But now I have sort of DC on the input side, DC on the output side, but I'm mediating it through purely sinusoidal waveforms.

OK. So maybe here's how I might think about my DC-to-DC converter. Maybe I would think about this thing as, you know, I come from my DC input V_{DC} , and maybe I have an inverter gain-- and I'm going to call this the gain of the inverter, OK-- which basically gives me the AC fundamental divided by the DC voltage. So this is sort of $v_{\text{inverter 1}}$ divided by v_{DC} .

OK. I've already told you what that is. It's basically I come in from v_{DC} in DC, and I get $2/\pi$ v_{DC} AC. So in my example here, this is equal to $2/\pi$, right.

If I use a full-bridge inverter, maybe it would be 4, and square wave, it would be $4/\pi$. If I use some other inverter that I haven't told you about, it could be some other value, again. But this is just to say, OK, going through the switch portion of my inverter, there's some gain between the DC input and the fundamental AC output that I'm getting from my inverter. That make sense to everybody?

The next thing I have is the tank gain. It's basically this model here, right. This is $G_{\text{sub tank}}$. So maybe I will put that in yellow, right. So now I have $G_{\text{sub tank}}$, OK, and $G_{\text{sub tank}}$ is the magnitude v_{r1} divided by the magnitude of $v_{\text{inverter 1}}$.

OK. And this thing is a function of the tank. It's a function of, like, is it a series tank or a parallel tank, and the values of the tank. It's a function of frequency. It's a function of power. OK, but importantly, it's a function of the switching frequency, right. So basically, I get to control this tank gain by controlling my switching frequency.

All right. The output of this thing, we said was v_{r1} , the fundamental of the voltage here. OK. And this should be-- what does this voltage look like here? It's actually a square wave.

This is a square wave. On the other side of the transformer, I get a square wave. But it has a fundamental component that we looked at up there, so we can kind of treat it as fundamental.

Then I have my transformer gain. I put my transformer in green, right. So let me call this $G_{\text{sub } x}$. That's for the transformer.

This is magnitude of $v_{\text{rect } 1}$ over v_r . And this is simply equal to, in my example, N_2 over N_1 . Does that make sense to everybody?

And what's on the output side of this guy? It's simply equal to $v_{\text{rect } 1}$. And then the last thing I have to think about is, What is the effect of this inverter?

And what I'm interested in is, for this kind of model, the conversion between the fundamental of the rectifier voltage to the DC output voltage. So this gain's sort of an AC-to-DC gain, but I'm going to use it from the fundamental to the output. OK. What would that be?

Well, I've already I've already shown you that, right. The fundamental at the input of the rectifier is 4 over π v_d , and the output of the rectifier, the DC output, is just v_d . So the gain must then be π over 4 .

Right, so here's the rectifier. What did I call that? I guess I called it $G_{\text{sub } r}$, the gain of the rectifier. And this is, in my-- this is magnitude of v_d divided by $v_{\text{rect } 1}$, right. And this is equal to π over 4 because it's a full-bridge rectifier.

If I had to use a so-called voltage-doubler rectifier, which would basically be a half bridge with a capacitor, it would be π over 2 . If I used a voltage-multiplying rectifier, like a Cockcroft-Walton rectifier, it could be a different gain. OK. So I can pick, Do I use a half bridge or a full bridge or some other inverter? Do I use-- what kind of rectifier do I use? What kind of tank do I use? all of which helps me shape this gain.

And the output of this thing is then v_d , right. So what I get is, in my output, I get v_d over v_{dc} , right. This is the voltage gain of my converter, right. I come in from v_{dc} , and I get v_d at the output. And that's simply just going to be $G_{\text{sub } i}$ times $G_{\text{sub } \text{tank}}$ times $G_{\text{sub } x}$ times $G_{\text{sub } r}$.

OK. And what I'm going to do in this whole game is I'm trying to build a power converter that's going to regulate v_d for a given input voltage. The thing I'm going to control is the tank gain. Right, everything else is fixed.

I choose it to get the gain, some range I want, but how do I deal with variations in power and voltage, or especially very interesting volts? I vary the tank gain. Does that makes sense to everybody? Questions?

So let's think about this. This is the same picture, right. I've got my half-bridge inverter, which gives me my inverter gain here, and that generates a-- I consider the fundamental AC voltage here. I go through the tank gain-- which is a function of frequency-- the transformer gain, the rectifier gain, and then I get my DC output again.

OK. What is-- the important thing in all of this, the thing I'm controlling, is the tank gain. And my model for my tank gain is essentially right here, right. If I thought about this as being an equivalent resistor-- actually, it's right here, right. This models the tank gain right here.

Well, this circuit, I could analyze. If I tell you $R_{\text{sub } r}$, then I know the operating, I know Q_0 , and I can calculate the tank gain versus frequency, right. What does that look like? It looks like this.

All right. So what I've plotted is, for Q_0 , which would be R_r is equal to the square root of L over C , q_0 equals 1. It looks like this. All right. So if I had an equivalent resistance that looked like that, then if I'm operating right at resonance, the tank gain's 1.

Why? Because this thing looks like a short circuit, and I'm just applying the fundamental to the load. If I go above resonance, right, this looks like a net reactance, and I get a voltage divider and it droops down. And if my operating tank Q is low, it droops down kind of slowly.

OK. If, on the other hand, my effective R_r were small, right, my tank quality factor would be high, square root of L over C over R would get big, and maybe it would look like tank quality factor of 5, in which case the change in gain with frequency is very steep. Does that make sense to everybody? So what do we have to do?

Well, suppose I know my input voltage, my DC input voltage, and I know the output voltage that I'm regulating the output to. OK. That sets a voltage conversion ratio, right. And so I know this conversion ratio, and I have to pick G_t to match that conversion ratio if I'm going to-- if I'm going to operate there.

That means, essentially, suppose I needed to operate at a net-- total conversion ratio, right, this is 2 over π times π over 4 times whatever my transformer gain is. Factor all those out. Let's just pretend those are all 1 for the moment, even though they're not 1 here.

Suppose I needed a-- or, to get to the correct ratio here, suppose that I needed a tank gain of a half, right. That means I'd have to operate somewhere here. OK. So I'd have to be sitting on this line of exactly the tank gain's a half.

And which curve am I sitting on? Well, it depends what power I'm operating at, right, or, equivalently, which load resistance I'm operating at, because for a smaller load resistance, the AC resistance gets smaller, the tank gain gets higher-- the tank Q gets higher. But between those two, for a given operating power, that's going to determine one of these tank gain curves.

And I'm going to go sit here, and maybe if my tank loaded quality factor is 2, that means I need to operate at 1.5 times the resonant frequency to get the conversion I wanted, right. Now, if my output voltage stayed the same but my input voltage went down, maybe I would only need a tank gain-- instead of tank gain of a half, I need a tank gain of $2/3$, I'd jump up here, OK, and in here, and I'd have to operate at a different frequency to sit at the same power level.

OK. So basically, when I think about a resonant converter, it's like a resonant inverter in the sense that I'm doing voltage division with the tank or I'm controlling power by controlling the net impedance of this thing to feed power through to my load resistor. But at the same time, I have to think that because I have a rectifier, it's not a fixed resistor I'm running into, but it's a resistor that's a function of, equivalently, DC load resistance or operating power and output voltage.

OK. So maybe what I could do is, by knowing the range of powers that I have to operate at and the range of voltage conversion ratios I could operate at, I can kind of put a bounding box for a given tank gain on what frequency range I need to run over. OK. And there's an example in *Principles of Power Electronics* that works over just this such a range.

So that's the basic notion of how we think about resonant DC-to-DC converters, OK, or a typical resonant DC-to-DC converter, where we're going to basically pick all the designs to help us get the voltage gain we want. I might pick my inverter and rectifier topologies to help me get part of the gain, and the rest of the gain, I'm going to get out of picking the right tank.

And I'm going to pick the tank components such that the quality factors, the operating quality factors I'm going to get are going to put me in sort of an allowable range of frequency, and then I'm going to sort of operate there to get the conversion I want. Questions about that? Yeah?

AUDIENCE: We're playing with the LC reactants to do this sort of voltage division, but is the power factor of these kind of systems bad because of that? Or--

DAVID PERREAULT: So that's an excellent question. Yes, so if I thought about the power factor seen by this inverter, right, first of all, he's putting out a square wave, and then I'm only delivering power via the sine wave. So that's sort of, in some sense, a distortion factor, right. But also, yeah, the current's phase-shifted from the voltage, right. So if I compared this to just doing some kind of PWM isolated converter, like we've talked about before, where we're putting more or less square waves through our converter, we end up with higher peak voltage and current stresses in a resonant converter than I do in a PWM converter.

OK. So you are paying, definitely paying something. And I think you'll concede, this is a heck of a lot more complex to think about than, say, a forward converter, right. It's-- and the models, and we haven't even talked about control modeling for these things. It's really tricky.

OK. So you give up a lot to do a resonant DC-to-DC converter, right. So why would I do that? The reason I might do it is because, we said before, if my load looks inductive, I can soft-switch these-- I can soft-switch these transistors, right.

So I can get zero-voltage switching of my inverter and, hence, run him at a really high frequency. Likewise, I'm putting sinusoidal currents in here, so these diodes kind of switch nicely, too, or if I'm using-- depending on if I'm using diodes or FETs or whatever I'm using there. OK. So the thing I can get about it is the ability to push up to pretty darn high frequencies and keep good efficiency and, hence, get down to small sizes.

So a lot of high-performance designs, despite all the complexity of doing this, will do that. And in fact, if you look at this, this is a rectifier. Right, this comes from line voltage and eventually gives me the voltage from my laptop.

This has a resonant converter in it. It's tiny for 65 watts, and it's north of 96% efficient. And the reason it can be not hot and very tiny is because they went to the great lengths of designing the right kind of very high-frequency resonant converter.

OK. So would you always do this? No. Right, the design time and everything else is going to cost you a lot. But what you will find is a lot of very high-performance systems, they're going to go through the effort of figuring all this stuff out to get you there. Questions about that? Yeah?

AUDIENCE: What would be-- for this laptop charger that you have, what would be the load variation on it?

DAVID PERREAULT: That's a very good question. The load variation's essentially infinite, right. So that brings me to another point.

I've talked about essentially one mode of power control in this thing, right. I've talked about essentially varying the switching frequency to vary where you're sitting on these curves, right. And that is a good method of power control, but it'll only get you so far, right.

Or another way to think about it is, if my current goes to zero and my power goes to zero, my effective AC resistance is going through the roof, right. So what will happen eventually is maybe you enter a different mode of operation where you can't keep soft-switching. So people will do all kinds of different other things to control power.

They might burst the converter. They run the converter at full power or some nominal power for some time and then shut it off, burst it on and off, and use the output capacitor here to hold the voltage so that when it's running, it's running at high power and then it runs at zero power for a while. That's called on-off control or burst control.

If I had a different kind of inverter in the input, maybe I can do phase shift or some degree of kind of PWM control, OK, or what's known as out-phasing control if you're doing it with sinusoids. So there's a lot of other control techniques you can bring in to deal with the really light-load cases.

The other piece of it is, is there's other kinds of tanks. This tank really doesn't go down to zero power because when the load resistance goes to infinity, the tank goes to zero, and you'd get really wide frequency ranges. That wouldn't be very good, right.

On the other hand, if I used a parallel resonant tank, I could do something different, and it can go down to zero load just happily. It has circulating current. Maybe the efficiency decreases, but it doesn't mind. Why? Because if I look at a parallel tank, in a parallel tank, Q goes up as r gets bigger. And so my frequency range goes down, so I can go to zero load with that.

Or maybe I use an LCC tank, which has some of the properties of both, or an LLC tank, which has two inductors and one capacitor. That's a very popular one for very high-performance converters, OK. So you can basically trade simplicity and still use frequency control but then go to burst control or outphasing control or other control mechanisms.

And people do that. And is it worth the complexity? It depends on the application.

In some cases, straight up flyback converter is the right thing to do. Right, you don't have to do this to do a DC-to-DC converter, but in some high-performance applications, this is very popular. Other questions?

I should say, by the way, why else-- looking at this, why else might I choose a parallel resonant tank, besides the fact that it goes down to zero load happily? Voltage gain, right? Here, a series resonant tank, the best I can do is make this thing look like a short circuit and get a gain of 1.

With a parallel resonant tank, especially if I'm going to operate up here, I can get a lot of voltage gain out of that. And in fact, if you want to build a really high-gain DC-to-DC converter-- and a design that was done in my group some time ago went from 200 volts to 40,000 volts-- one thing you might do, you might pick a voltage-multiplying rectifier, a rectifier with a huge voltage gain. You might pair that with a parallel resonant tank that would help give you voltage gain from the tank, right.

So all these design decisions-- I might choose, instead of a half-bridge inverter, maybe I'll choose a full-bridge inverter, because instead of getting a gain of 2 over π , I get a gain of 4 over π . Right, so there's, all kinds of games that can be played to help you get to some large conversion ratio that you might want to get to, which might be hard with a standard PWM converter.

OK. I would like to say-- and we're getting close to the end here-- that not all resonant power converters are quite as-- I don't want to say "simple as this," because this is not simple, but that are quite as clean as this. In this version of a resonant converter, we have this intermediate current waveform that was purely sinusoidal, right, and that helped us think about the analysis, just thinking about sinusoids throughout the converter. You don't always get that. Sometimes you get waveforms that are quasi-sinusoidal, and sometimes they'll call that a quasi-resonant converter, OK, depending upon the type.

But I wanted to show you an example, a very a very useful home example. And what this thing is, this is your standard home induction cooker. Right, if you want to go out and spend-- I don't know, what did you spend for these?

DAVID OTTEN: \$50.

DAVID PERREAULT: \$50 for one of these things. These were so cheap, we went out and bought multiple of them and took them all apart. Actually, I have the following slides that Dave prepared. So I don't know if, Dave, you want to present the topology and everything that's going on and then show it, show it off, that'd be great.

DAVID OTTEN: No, no, go ahead.

DAVID PERREAULT: All right, I'll describe it. OK, I'll describe the circuit. Here's the circuit, OK, and it turns out a lot of the-- a lot of the systems are designed like this.

And here you can see the coil. Now, this coil both serves as the inductor-- both serves as sort of an inductor to store energy, but it also couples in to the pan to deliver the fields from this inductor, hit your pan, and make a effective resistive load on the tank.

OK. And then here you can see it's an interesting circuit. OK, here's the topology. OK. We can pass these around, but please, I need them back. But let's see what's inside this toy.

DAVID OTTEN: Note the pieces of ferrite on the back of the coil, to deflect the field into your pan and not into the circuit.

DAVID PERREAULT: That's right. Yeah, so the field's supposed to go into the pan side, and the ferrite chunks there help you out. And Dave has kind of taken this thing apart and instrumented it.

OK. Here you see the top. This is where the pan goes. The bottom, you can see these ferrite pieces to force the fields to the other side.

Here's the circuit board. Here's what it looks like. OK. It has a sort of input filter and a line-frequency rectifier followed by a low-frequency filter, so that at least at one point in the line cycle, you could imagine this being a DC voltage. OK. Essentially, this voltage goes almost as the rectified sine voltage.

OK. And here's the circuit. This is the inductor slash transformer, right, so that that's the coil right here. There's a capacitor across the coil, OK, so this is going to give us our resonance swing. OK. But this circuit only has one switch.

OK. So what's essentially going to happen is-- so think about this voltage as being a constant input voltage. We're going to turn on the switch and apply a constant voltage across the inductor and ramp up current in the inductor. Then we're going to turn off the switch.

And this is connected to a voltage source, and this thing is going to go through essentially a parallel resonance. It's going to-- the voltage is going to ring up, and it's going to ring back down. And when it rings back down, we can turn on the switch.

OK. And so in this AC ring, we're coupling energy via induction into the output, which is essentially sort of the resistance. Everybody get the basic idea? So this thing, it doesn't quite have sinusoidal waveforms at all, but it is using a resonance.

And in fact, we do get zero-voltage switching with a single switch here. There are many inverters that actually only need a single switch, one of the most famous of which is the Class-E inverter. This is a variant on that concept.

To control power, we're going to vary how long we hold up the switch and ramp up current in the inductor, but at a given power level, that's going to be fixed. Now, the beautiful thing about this is, is the amount current ramps up in this is proportional to this voltage, which varies over the line cycle. So almost for free, this thing operates at super high power factor because the switching frequency of this thing is way above the line frequency.

And so we're going to draw power that is proportional to the square of the line voltage, which lets us get very high power factor. OK. So basically, with one active switch and some resonant components, plus a diode rectifier and some filters, we get basically an induction heater, which is why they can do it for \$50. I will now turn it over to Dave, and he can sort of explain what's going on.

DAVID OTTEN: OK. So I've got-- so let me tell you-- let me tell you what we're working with here. So I want to adjust this so that it triggers. So the pink waveform is the input voltage after the rectifier. So it doesn't go all the way down to zero, but you can see it has some humps in it.

OK. The blue waveform is the voltage across the switch. So you can see it's zero for a period of time. That's when the switch is on. When the switch turns off, then you see that the voltage resonates up and then down again.

Now, the real thing that we're interested in is the current in the coil. That's the yellow waveform. So you can see that when the blue waveform is flat, the switch is on, the current ramps up. At some point, the switch turns off, the voltage builds up, the current goes down.

OK. And so that's the current waveform that's actually being used for the heat. Now, if I change the power-- so that was 1,200 watts. I'm going to go to 700 watts, and I'm going to adjust the thing so that it triggers. OK. You can see that the current is lower.

OK. And now while it's running, I'm going to try to turn up the voltage-- turn up the power. And you can see that the control system is going to increase the time that the switch is on. That allows the current to build up more. It turns out that it makes the frequency go down, but I think that the time is what's more important.

OK. Now, there are several ways to look at this circuit. OK, right now, I'm sort of showing you one high-frequency signal. Now I'm going to change the time scale so that we look at one full cycle of 60 cycles. So let me adjust the triggering.

Maybe one more. OK. And then let me go down in power because this one triggers more reasonably. So now you can see the feature that I was talking about, where during one period of 60 cycles-- this is really 120 hertz, so it's twice that-- you can see that the input voltage, which is the pink, goes basically to zero and up again. Then the current is the yellow waveform, so that has an envelope of a sine wave. OK, and then the voltage matches that.

OK. I should point out this is the parallel resonance circuit, and so the peak voltage is 700, 800 volts. OK. The peak current when we get up to full power, 1,800 watts, is about 70 amps, so it's serious. OK, let me just turn up the power, and you should see the current is building up.

So this is 1,400 watts. This is 1,600. And whether we get 1,800 actually has something to do with the line voltage, OK. So I'll just turn it down.

Now if I-- we were talking about different modes of control before I programmed the thing for 500 watts. OK, so it runs, and then all of a sudden, it stops. OK, so this circuit doesn't really want to work at really low power, and so it pulses on and off.

Normally, if you're just trying to heat water, that doesn't really make a difference, and so they can get away with this type of operation. Notice how the soft start, it slowly, slowly builds itself up. OK. OK.

DAVID PERREAULT: Yeah, OK, so just to show that this thing is real-- let's see how fast we can cook with 1,200 watts. And there we go. Wow, that's pretty quick. Anybody hungry? It's just after lunch.

DAVID OTTEN: All right.

DAVID PERREAULT: All right, there we go, one cooked egg. We'll wrap it up for the day. But there you can see how even resonant principles can be applied.

We've got zero-voltage switching pushing 1,800 watts with a single active switch. So that's sort of on the clever end of the trickery you can do with resonant converters, but as I mentioned, they're in a lot of kinds of applications. Have a great day.

[APPLAUSE]