[SQUEAKING]

[RUSTLING]

[CLICKING]

DAVID PERREAULT: OK. Why don't we get started. We've been talking about DC to AC power conversion or inverters. And we've been talking about the full bridge-- what we would call the full bridge inverter or, more generally, you could call it the full bridge converter. So we have four switches.

And we can use this. And depending upon how we switch them, we can use them to generate an AC output voltage. So we've been drawing a generic element that we're going to have as a filter.

But imagine you constructed the structure like this. Suppose I did this. Suppose I had an inductor, a resistor, another inductor. Then I could consider that I'm synthesizing on this side of the circuit if I have q of 1 of t and its complement, so I have a duty ratio of the top switch on this side of the circuit, I could generate some output v x1. And that might be filtered through this inductor. Might generate me v o1.

And likewise on this circuit, if I have q2 and q prime 2, I might generate if this is d1, this side I'm running with duty ratio d2 of t. And so I'm going to have an output voltage v x2 and v o3 And if I thought of then this load network here, maybe if I thought of v x, I would have v x is equal to v x1 minus v x2.

And I could then think of v out 1-- I'm sorry. I could think of my output voltage v out across the resistor as being v out as being equal to v o1 minus v o2. So another way to think about the full bridge is as a pair of back to back buck converters. Here's one synchronous buck converter switch pair and his filter and another synchronous bucks switch pair and his filter. And really what I'm doing is I'm taking the difference between the two outputs to get the AC output. Right.

So, for example, if I made the average value of v x1, which I get to control anywhere between zero and v DC by how I switch these switches over a PWM cycle, I could make that equal to v times 1/2 plus cosine of omega AC t and maybe I can control v x2 to be equal to some voltage v times 1/2 cosine of 1/2-- sorry, minus cosine 1/2.

This should be 1/2 plus 1/2 cosine. This should be 1 minus 1/2 cosine of omega AC t, which would then give me v x. If this is v x1 and this is v x2, v x is just the difference of those two and the DC part cancels out and I just get the difference in these AC parts, which gives me v cosine of omega AC t.

So by basically synthesizing something with a DC component from the left buck converter and something with the same DC component, but a complementary AC component from the right hand buck converter, I generate an AC output voltage. So you could really think of a full bridge inverter as being two synchronous buck converters back to back. Any questions about that?

Now the average voltage that we get is just due to PWM. So I can control this average voltage and this average voltage. The details of how much ripple I get and what its high frequency AC components are depend upon how we do the PWM. And we talked last time about different ways we could do that. But it's the same exact concept.

All right. Why do I mention that one? Just because it's another way to look at what a full bridge inverter does. Another, just as an aside, which we're not going to dive into, is just because I used buck cells here doesn't mean that's the only converter I could use. So I could take two synchronous flyback converters, for example, and use each of those as a half cell and take the difference of their outputs and get a flyback-based inverter or a flyback inverter. And people do that.

So people use different topologies to do this. This full bridge structure is, generally, the most preferred because buck converters are really low stress. They're really simple to control. They have a whole bunch of benefits. And it's about as simple as you can get to do the job. But if I had to put an AC transformer in the middle, maybe I'd think about something like a flyback inverter to give me isolation, and large step up, and everything else. Any questions about that?

The second point I'd like to make is I've talked about power going from AC to DC-- I'm sorry-- DC to AC. And what the power flow looks like depends upon this current iL. What I'm really limited in this circuit by is what range of v x1 and v x2 I can synthesize. There's nothing wrong with using this circuit to flow power AC to DC so long as I have a source that could support it.

So one thing I could think about is coming along and building the same basic structure. So here's v DC, And I'll put a pair of switches here. And I'll put a pair of switches here. And there's nothing that says I can't combine the filter elements. So here's iL again. I'll just combine those two inductors as one element. And maybe here I have v AC. Right. There's absolutely nothing that says that I can't now use this circuit to run power from AC to DC. It's just a question of what the phase of iL is compared to v AC.

And in fact, if I think about it then, that's very much like from AC to DC, this acting as a pair of back to back boost converters. It's taking energy from here, boosting it up to a higher DC output voltage. So just like we talked on early in the lecture, I can either think of this as a buck inverter or a boost rectifier. Any questions about that?

And in fact, people often do this. Where might you do this? Well, imagine for a second that this whole thing-- this box-- maybe this is the back EMF and winding of a generator. So I can then extract energy from the generator so there's a mechanical input in. It generates a back EMF voltage and I take power out and throw it out to the DC side or maybe this is a motor drive.

And sometimes I'm driving power into the back EMF of the motor to generate torque. And sometimes I'm regeneratively braking and absorbing power from the AC side back to the DC side. So I can either generate positive torque with respect to speed and deliver mechanical power or absorb mechanical power and slow down. Any questions about that?

So just because I often call this an inverter, we might really think of it as a converter, meaning it can be an inverter or a rectifier, depending upon how I'm controlling it. And this is very common. The other places it could be is the AC grid. So if I had a solar panel stack here and I wanted to shove energy out into the AC grid, perhaps I use exactly this structure. OK. So there's another place you might do it.

And that's just many applications of an inverter. Yeah.

STUDENT: When you're generating breaking, do you have to use the switch the way the switches are switching [INAUDIBLE]?

DAVID PERREAULT: Well, we talked last time about how I might control current. So I could think to measure the current here. And I could either generate the PWM to generate the average voltage vx. If the current is too low, I turn up vx to make the current high. And I put a closed loop so I can control the current or I could use a hysteretic current controller is another way we talked about.

So really, it's really a question of whether a power is flowing AC to DC is just me deciding, well, relative to this voltage, what do I want the phase of that current to be. It could be exactly in phase and then I'm just purely delivering DC to AC. It could be iL could have the exact opposite sign to v AC and then I'm just purely delivering energy AC-DC or something in between. So yes, it's exactly how I control the switches determines the power flow direction in this generalized picture, understanding that there's limits on what the load can do.

Now if I had a current source inverter, here I'm talking about deciding power flow direction based on the sign of this current where the voltages I apply are the same sign. If I had a current source inverter, I'd do the reverse. The current would have the same direction. I'd do it at a different voltage polarity. Any other questions? Yeah.

STUDENT: I have a question on the one on the top. You have the average voltages, but they have the cosine omega t term. So is that some special AC average that gives you the fundamentals here?

DAVID PERREAULT: Yeah. That's a very good question. What I meant to convey here-- and I wasn't very clear about it-- is imagine I have a really high switching frequency, 100 kilohertz, but omega AC is maybe some lower frequency, 60 hertz. So from the perspective of the switching frequency, maybe omega AC is really low and so think about as DC. So instead of being, you know, average over infinite time, it's average over just some time this long compared to the PWM cycle, hundreds or thousands of PWM cycles.

So I slowly vary my duty cycle, d of t, not within a switching cycle but slowly compared to the switching period of the inverter and on the time scale of the AC output I want to synthesize. So if I want to synthesize music, I'd better have a switching frequency that's much higher than 20 kilohertz bandwidth you can hear kind of thing or, in my case, 8 or 10 kilohertz, but that's life. Any other questions? Yeah.

STUDENT: Can the top one be used to excite the excitation of the [INAUDIBLE] for rectification?

DAVID PERREAULT: For rectification. Yeah. I don't see why not. Yeah. As long as you have a DC voltage that's greater than whatever the AC output in amplitude is. Right. This thing, because it's based on a buck converter, I can't synthesize an AC waveform to drive current that's bigger than the thing I'm trying to drive.

STUDENT: If you were to have the stator in the middle and then use that to excite the rotor field winding.

DAVID PERREAULT: Oh. Could I? Yes. You can play games like that, but maybe that's better. Yeah. That would require more explanation than I have time to give, unfortunately. But yes, you can. So that's-- whether I think about this as an inverter, or rectifier, or just a converter is up to you.

But I would like to note that there are some times when I only really want to go one direction. In fact, there's times when I only want to take power from the wall and put it out to the DC side. If I'm building a computer power supply, at least in ones today, there's no real reason for me to do anything except absorb power from the AC side and put it to the DC side. Right. And if that's the only thing I want to do, I can do the conversion. I could do that with this circuit. I could connect this to the AC grid.

I could connect this thing to the AC grid, and control the switches, and dump energy into the DC side, which then I'd use to power my computer, for example. Right. But if that's all I want to do, this is kind of a more complicated circuit than you need to do it. So for things like computer power supplies, which are kind of ubiquitous these days, or something like my laptop adapter there, that's a pretty expensive circuit to be using because it has way more capability than we need.

So mostly if we're going to deliver power from AC or DC, we use simpler circuits. One way to do this and the oldie-fashioned way to do is to use what's known as a phase controlled rectifier. And there's a chapter in the book on it. At high powers, it's still done. But I'm going to skip over that as an approach because it's less relevant to high frequency converters and also because, if you think about things like computer power supplies, not so much in the sub-75 watt range but higher power like your desktop, there are very strict requirements in how it looks to the AC grid.

So we talked about uncontrolled rectifiers at the beginning of the term. And they tend to have ugly AC waveforms. And that means they have low power factor. So a goal of many switch mode rectifiers is to take energy from AC to DC only, but to do it at a very good power factor. OK. So let's think about how we might do that. Maybe I might draw this generalized box as something like this. I have some AC voltage, v AC. Right. That's the line voltage typically. Maybe I'll call it v line.

It goes into a box. OK. That box is a rectifier. And maybe it just feeds-- I'll put a capacitor here. And we'll talk more about this capacitor later. But maybe it feeds some load v DC. And I would like, generally, this line current i line to have very nice waveforms.

OK. What would nice be defined as? Well, you know, here is v line. This is the line frequency. This is 60hz or 50hz, depending upon where you are. The best I can do in terms of power factor is to make my whole converter look like a resistor. So my goal might be to have a current waveform that looks just like this.

Here's i line. And maybe I would make i line something like i line of t is just equal to some constant g times v line of t. Now, clearly, the amount of power I'm going to get is going to vary with g. But for any value of g, I'll get some amount of average power out of the line. And this is about as good as I can do in terms of power factor. OK.

So done that way, this is called a unity power factor rectifier because the line current, ideally, just looks like-- it makes the converter look like a resistor. It's a variable resistor depending upon this constant. So as I vary the g to vary the amount of power I draw, I vary the equivalent resistance. But nonetheless, at any given load, it looks like a resistor. OK. Does that make sense to everybody?

So that's called a unity power factor rectifier or a UPF rectifier. More generally, in modern terms, people talk about a PFC or power factor corrected rectifier. And what's meant by that is I've shown you the idealized case where this draws a perfect sinusoid.

You may have converters that don't draw perfect sinusoids but are close enough that the power factor is very, very high. And so that would be called a PFC converter. And if you look at many applications like computer power supplies or for different kinds of tools, they all have slightly different requirements put together. But generally, the power factor has to be pretty high. They limit the distortion in the current or they limit how non-sinusoidal the line current looks.

So what I'm going to do is I'm going to tell you now-- I'll give you as in one example of this, I'll show you what's at the front end or some variant of what's at the front end of most computer power supplies for single phase. And this is the idea. OK. Whereas we had all these active switches, it turns out I don't need all those active switches. And the active switches, and their controllers, and all that stuff are expensive.

Maybe I could do something like this. Maybe I could come in. And here's my v AC. Here's my AC line voltage. I'll run that into a diode rectifier. OK. So if I plotted v line of t here-- and again, I'm plotting things on the line frequency time scale so I'm going to have a switching circuit that runs much faster than that. So this time period t line might be 1 over 60 hertz.

So I have v line here. Maybe what I will get then is a voltage vx, depending upon how I make the rest of the circuit look like, vx that looks like this. So this is vx of t. I will make vx of t. When v line is positive, these two diodes are on and v line equals-- vx equals v line. And when v line is negative, these two diodes are on and vx equals minus v line. So now I get something that does this. OK. That's a pretty bad half sine wave. So right. So here's v line. Here's v x. Does that make sense to everybody?

Now. The beauty of this is I've made it unipolar. And so then I can build the rest of my circuit. And maybe one way I'd build the rest of my circuit is just to build a boost converter. So here we go. I only have one active switch. q of t runs this switch. And he's now referenced to this node. And then I can put a diode and a capacitor and get v DC out. And what's my goal?

My goal might be to control my input current i x. And maybe I'll make this current i x. And i x maybe I'll control to look like this, on average. There's going to be some switching component to it. But maybe I'll make i x, if this is a voltage v s-- peak voltage v s, maybe I'll make his peak g times v s. All right. So it's going to have a sinusoidal reference current. And I'm going to track it.

OK. Well, if this is the current i x, it has to ultimately come from the diode bridge. So when these two diodes are on, which happens when v line is positive, it goes through one way and then i x goes through the other when the other two diodes are positive. So then if I plotted i line, what would i line look like? i line would look like this.

So this is i x. And i line would look like this. It would look like i x in the positive half cycle minus x in the negative half cycle. And this would be i line. So suddenly, what happens is I only have to control this simple DC to DC boost converter to generate a time varying current over the line cycle. And the diode bridge naturally unfolds that. Give me a sinusoidal input current. Any questions about that?

This is kind of cool because I've thrown out four active switches and a bunch of control for basically one active switch. And yeah, I got some diodes, but so what. Those diodes don't require any control. They're cheap. They're simple. Life is good. OK.

Now this kind of boost rectifier-- this would be called a boost rectifier-- is very much in the spirit of my full bridge. I just don't have to worry about all the phase shift. Now, that said, it can only rectify. Power is only going to flow this way because of these four diodes. That's the price I paid. I can't phase shift my current with respect to the line or do any other fancy things. But, generally, I don't want to. I don't want to get energy out of the line and put it into a power supply for my computer.

The other thing about this, of course, is this particular circuit is a boost converter. Like if I look at it from here, that just looks like a boost DC to DC converter. That means that this DC output voltage must be bigger than the peak line voltage. OK. Now what that means practically is most places in the US the peak line voltage is 170 but it can be plus 10%, minus 15%, around that in practice or sometimes even worse.

But then I also have to make it run in some places in Japan where the line voltage might go down to like 90 VACRMS. And I got to make it run in Europe where it can get much higher. So most of the time, what people will do is they'll make this like 400 volts for a single phase line, which means it can take anywhere in the world's single phase AC voltage.

The other thing I should point out is notice that if this is hot and this is neutral, which is near ground, notice that this node is not referenced to this node. So don't make the mistake of saying, OK, I'm going to connect my ground clip here because grounds are neutral and bad things are going to happen. Usually, bad things means your circuit explodes. And I've been there.

Let me just illustrate this for you. This is a simulation. So here's my sine wave voltage source, four diodes. Here's my boost converter. OK. Let me run it. This is just a simulation we can run.

OK. Let's take a look at some of the waveforms. Let me look at the-- actually, let me take a look at-- let me take a look at the AC input voltage. Right. Looks like that's the AC input voltage.

Let me look at the AC input current. And actually, I looked at them with the wrong phase. Let me just change the phase here. Minus 1 times that current. So this is very much generating a very sinusoidal looking input. If I looked at this current, for example, the current in the inductor, he is the rectified version of that.

So it's exactly what I said. But of course, keep in mind, we're looking at a long time scale. This is 60 hertz. If I zoomed in close, what I would see is a lot of PWM switching. So I have a switching frequency. Actually, I don't remember what it is in this example, but it's probably like 500 kilohertz or something like that. And I'm using it to create a 60 hertz rectification and drawing energy from the line at essentially unity power factor. Any questions about that? Yeah.

STUDENT: So I may have missed this, but why isn't there a way you could introduce [INAUDIBLE]?

DAVID PERREAULT: Ah, that's a very good question because what I'm doing is if I had some passive kind of thing going on, I might. But what I'm really doing is I'm-- I can do this control in a bunch of ways, but I'm sensing this current. And then I'm chopping this-- I'm chopping the switch in order to regulate the current in the inductor i x.

So this would be like what i x is doing according to following a reference which is exactly in phase with the input voltage. OK. Any phase shift that would go on I can deal with. And keep in mind this inductor at the line frequency is really small so this doesn't have much impedance at the line frequency. It has impedance at the switching frequency. So I need to switch way faster than the line frequency.

STUDENT: How is it possible to control both i x and DC [INAUDIBLE]?

DAVID PERREAULT: So I'm not controlling-- oh, v DC?

STUDENT: Yeah. Like if the output power is set by the load and you're trying to maintain the DC at 400 volts while also having a [INAUDIBLE] on it. How do you control that with the one switch?

DAVID PERREAULT: Excellent question. Let me put that on hold. I'm going to answer that in about two minutes. But it's a perfect question. How do we do that? Let's think about that. I mean that's a very good question. In fact, that's the main thing I want to convey in the next section of the lecture.

How would I control that? Well, here's the scheme that we'll generally use. The magic thing I have is this constant g. So this constant g is essentially what sets the power. So what does the power-- what does the power look like here? p of t is equal to i line times v line, which is ultimately equal to g times v line square root of t. Right.

So if this is-- this could be then g times v s squared sine squared of omega t, which I could also write as g vs squared over 2 times 1 minus cosine of 2 omega t. So if I were to draw that out for you, what the power would look like is this. So this is p of t. It would have an average value that's gvs squared over 2. Right. And then it would pulsate up and down like this. That's all that equation says.

Why does it pulsate up and down? Because when v and i are equal to zero, when the line voltage is zero, you're not drawing any power from the line. When the line voltage is really big and the current's really big, you're drawing a lot of power from the line. And on average, you get gvs squared over 2. OK. So the whole magic of the control is controlling what g is because I get to pick g because I'm picking how big I'm tracking my current to be.

So here's how you'd go do it. You might say, OK, let me come in and let me take v line of t. OK. And I'm going to multiply v line of t, the instantaneous value, or, actually, maybe I just measure v x as the magnitude of v line of t. And I multiply it by some constant g. And this gives me i ref, which then I feed to my current controller for my boost converter. So an input to my control is what I chose g of t to be.

Now what I'll do is I will measure v out. I'll have an output reference v ref. You know, maybe that's 400v. I will measure v out. So I'll go to the output of the boost converter. Or I called it v DC so let me call it v DC. I'll take the difference between these two and I'll have an error voltage. And then I'm going to run that into a compensator h of s. OK. h of s. And this is a high gain.

And that's going to give me g of t. So if v DC falls below vref, v error gets bigger. I amplify that and I increase g of t. So I'm going to pick g of t big enough that I get enough power that the output voltage matches whatever I want, in this case, as I said, 400 volts. Does that make sense? So this is a typical kind of controller for a PFC control.

Let me explain a couple of things here. Keep in mind what's the bandwidth of this thing? Well, g is supposed to be a constant over a cycle so that the AC current that I draw isn't wildly varying over half a line cycle. So the bandwidth of this controller tends to be really slow, like less than-- it shouldn't vary its output voltage very much over half a line cycle because if it starts doing that, then the current no longer looks sinusoidal. And it can only respond to slow variations. But as we'll see, what that mainly implies is because the bandwidth of this loop is low, this capacitor size here needs to be pretty big to hold things up. All right. And in fact, we're going to talk about that in a minute. That size of this capacitor right here tends to be very dominant things in most PFC rectifiers.

Just to illustrate that, I've got one with me here. This part on your very left is basically the line filter, because we're switching and I got to keep the PWM junk out of the line. Has a fan. Then you'll see if you look carefully, you'll see these four diodes here. One, two, three, four. Those are the four diodes.

The inductor is a little bit hard to see because it's buried in here. Then the main thing that you see is this block right here, which is all these capacitors. And that's this guy right here. So we're going to come back and talk to that. And then the rest of this thing is this top board is essentially this first stage. Then we got to talk about what we do after that.

It's a little bit hard to see the inductor. But does that answer your question about control? So we need a very low bandwidth controller but we modulate that. We ultimately modulate one switch to control the current that then controls the output. Multiple loops. Yeah.

STUDENT: I don't understand how [INAUDIBLE] remaining half cycle.

DAVID PERREAULT: Because I'm telling it to you? In some sense, right? Suppose this compensator is acting really slowly so that g is, roughly speaking, a constant for some time period. Then all I'm doing is I'm measuring the line voltage. I'm making a multiple of it and creating a current reference. And then I have another inner loop, maybe a hysteretic current control, that switch-- if the inductor current is too far below i ref, I turn the switch on until it ramps up.

And once it gets a little bit too big, I turn it off and then the diode turns on and the current ramps down. And I'll get this kind of choppy current that roughly follows i ref. That make sense? That's all happening at some high frequency, the switching frequency. The component I care about for power factor purposes is the local average of that current because that's what's being fed back. And there's some additional filtering here to get rid of the ripple. Any other questions?

So that's the basic idea of a PFC rectifier control. Now note one thing. I mentioned this before. This node is nothing like this node. These aren't at the same potential. So it's a little bit hard to use because if I'm plugging into the wall and one of one of the sides of this thing that I'm plugging into is neutral, that's very close to the ground. It's ultimately a DC connected to ground.

Yeah. I better not connect this node to ground because bad things will happen. But on the other hand, my computer-- first of all, it doesn't need 400 volts. It needs, ultimately, one volt. Right. But intermediate in your desktop machine, it probably steps it down to 12 volts and then there's another converter that goes from 12 volts to 1 volt to run the processor. So ultimately, we need, first of all, isolation, because this node isn't galvanically isolated from here. And we do need our microprocessor really referenced to ground because you don't want to touch the case of your computer and get shocked.

So you need some isolation. We also need a large voltage conversion. So usually, after this PFC stage, the thing I drew as a resistor here is some other DC to DC stage. So what does that DC to DC stage look like? It could look like any kind of isolated DC to DC converter we're talking about. It could be just maybe I split the bus in two and I have a pair of switches and a transformer, and then I run the other side of this transformer into a diode rectifier.

And that comes out to my 12 volts, for example. And that powers my next power supply. Or maybe I have multiple outputs from next stage. You have plus and minus 12 and plus 5, whatever I need. But I will usually have some subsequent stage. And in fact, the other part of that converter, the other board I'm passing around, is the subsequent stages to provide from the high voltage rectifier down to lower voltage loads. And 400 volts is kind of standard, A, because it matches across what all the different universal input voltages across around the world for line inputs and, as a consequence, you can get really good energy storage capacitors and a bunch of other really good cheap elements to build your converter there. And so then you have two stages.

You have a first stage that does that, then you have a transformation and isolation stage, which gives you some low voltage, typically 12 volts, but it can be other voltages, and servers, and things like that. And then you have usually another converter that gets like this more or less a buck converter or some variant that gets you very high bandwidth control over your microprocessor supply. So hence you can see we're going through lots of power conversion stages here. Questions?

STUDENT: What do you refer to by bandwidth in this case?

DAVID PERREAULT: So what I mean by bandwidth is this gain stage-- maybe that's an integrator or maybe it's a PI controller-- oughtn't vary this signal on a time scale that approaches half of a line cycle. So maybe if half of a line cycle is 120 hertz, maybe this thing varies only like 10 hertz or something like that or it can be very slow. OK. That means you need a lot of energy storage to hold things up in the intermediate time but that's necessary.

OK. That brings me to my next point, which is the following. When we talked about-- even when we talked about inverters, I kind of brushed over-- I just said, oh, I have a perfect DC source and he's going to provide whatever I want. But I didn't worry how much power pulsations were coming from this DC source. Right. In reality, I almost always have some capacitor here.

In this application, I said I'm going to have-- this isn't really necessarily exactly going to be a resistor. This is going to be another power supply stage. But this next power supply stage, really, it wants a DC voltage. It wants to draw a DC current. So really, this capacitor is going to take the difference between instantaneous-- whatever power this thing is going to provide instantaneously and whatever the next stage instantaneously wants. OK.

So if we thought of-- let's see. Where's my original drawing? If we thought of this box over here as being an AC to DC rectifier, I just plotted for you what the power versus time that is coming out of-- that's the power versus time that's crossing this node. This p of t is basically what's coming out of this device into this network. And what you can see is it's actually pulsating at twice the line frequency.

So this basically every half cycle-- maybe I'll move this up. Every half cycle, both the line current and the line voltage go back to zero. And when the line current and line voltage are zero, there's no power coming from the line. In fact, it doesn't matter what current I want to draw from the line. If there's no line voltage, I can't get any current. I can't get any power out of it. So that's why I have this square law power versus time. OK.

Now if I imagined that the orange is what's being sourced-- so that's what power is coming to here. If I imagined what I want going to my load is DC power, what does that look like? Well, if I have a DC voltage and a DC current drawn by my next stage, my next converter stage-- it's not really a resistor-- it basically wants to look like this. That's what I'm controlling is I'm controlling the average power going to the load right here.

So what about the difference between the instantaneous power and the average power? Where is that being stored? Conservation of energy is flowing in one place and flowing out the other. If they're not equal, it's got to be stored somewhere. And precisely where it's stored is in that capacitor. OK. Now the most common place on a PFC converter is to store it across this high voltage DC output because you can buy very good capacitors there.

But what I'm trying to tell you is the energy going into that capacitor, basically, in the first part of the line cycle, I said this is kind of t over 2. This is t over 4. In the first part of the line cycle-- and if I went negative, it would do this. So this is, in this time period, basically, the line isn't sourcing enough energy. It's sourcing less than the average.

So during this time period, energy is coming out of the capacitor. And then in this time period, I'm drawing more energy out of the line than my load's using. And I'm recharging the capacitor. So at twice the line frequency, I'm charging and discharging that capacitor. So that capacitor has to be physically big enough or it has to be energy storage-wise big enough to store that much energy. And that sets a bound on the energy storage requirement. Now the same thing is true in the inverter.

Suppose this was the AC grid and this was a solar panel. I don't want my solar panel pulsating. I don't want my solar panel-- he doesn't want it delivering 100 watts in one time frame and 0 watts in the next for an average of 50. He wants to source DC current and DC voltage. So that means-- but if I'm going to deliver energy into the grid at unity power factor, which is what I want to do, that means this power into the AC grid is pulsating. It's just the reverse of that waveform.

So I still need some big capacitor to buffer the energy. So whether I'm going AC to DC or DC to AC, if I have pure DC power flow on one side and unity power factor AC on the other, there is no way around the fact that I got to store some energy somewhere. And you'll find people that say, I got this magic technique that's going to do it, except that conservation of energy rather forbids that.

And I have feel like I got this magic way around it. No, you don't because, you know, conservation of energy holds. So somewhere in your circuit-- it doesn't have to be across the input or the output, but if you're going to have AC voltage and AC current on the other side that's at unity power factor and a DC voltage on the other and a DC current-- so this is i AC and this is i DC-- somewhere in your circuit you either need an inductor or a capacitor. You need some third place that's going to store the instantaneous power between v DC i DC, and the pulsating AC power v AC i AC. Does that make sense to everybody?

Now I'm running out of time here, but it's very clear if I basically just integrate, if I take the integral of this right, that's an amount of energy. And that's how much energy that capacitor has to store. If I work through the integral, what I get is e store is exactly equal to p average, the DC power, divided by the AC line frequency. So it's worse at 50 hertz than 60 hertz by a little bit because I have to store it for longer.

OK. So when we said I want to increase the switching frequency of my converter to make the components smaller, That's all well and good, but if I'm delivering energy to the line at unity power factor, I need to store a certain amount of energy inside my converter for unity power factor anyways that's related to the line frequency. And that's going to set a minimum size on the capacitors, or inductors, or whatever I use. Most people use capacitors because they have better energy density, at least at small scales, like the single phase line scales.

So any scheme you're going to come up with, you're going to have to do this. Now it turns out that this stored energy is the minimum that I need if I have unity power factor. If you allow me to cheat a little bit, maybe I don't have perfect unity power factor, you can effectively reduce that a bit. And essentially, you're using the line as your storage mechanism, which is allowed over a limited range. Real computer power supplies don't have to have perfect power factor. They just have to have good enough power factor. And so you can reduce this maybe by a factor of 2 or something if you're willing to let your line waveforms distort a bit. But generally, it's on this scale.

And the details of how much you can do depend very much on what power level, what your application space is. There's rules regarding how much distortion you can have in line connected equipment depending upon, like I said, application and power level. But what that means is you've got to get clever in how you store this. And by the way, this is the amount of energy-- this is the swing in energy. So if I'm going to store that with very small ripple on that capacitor, that's 1/2 cv max squared by minus 1/2 cv min squared.

If my ripple on my capacitor voltage becomes smaller, the energy storage rating of that capacitor is even bigger than this number by a lot. So there are pretty-- there are games you can play, but this is a minimum energy storage you're going to need to store. And generally, if you want small ripple on the capacitor, you even need much more than that. And the lecture notes and the text go into the details of those calculations.

So I'm out of time to talk about this. Are there any final questions before we wrap up and move on to a new topic? OK, great. So next class, we'll have a little fun. And then when we come back after spring break, we're going to start to talk about what happens when I go up in power and we start to use three phase systems. So have a great-- if I don't see you, have a great break. And hope to see you tomorrow.