[SQUEAKING]

[RUSTLING]

[CLICKING]

DAVID PERREAULT: OK, so why don't we get started? So we've been talking about single phase DC-to-AC converters, and I actually haven't said which direction power is flowing. So the circuits that we're talking about can generally flow power in both directions. And so a general single-phase, full-bridge voltage source inverter looks something like this, And you would usually put a DC bus capacitor here to help with high-frequency ripple currents.

And you have four switches like this. And then you will put your output here and maybe there's going to be some load network that goes right here.

And in addition to that-- and we'll talk to this-- maybe sometimes you would have some kind of current sensor right here in order to sense what the output current is if you wanted to use that information. I just thought I'd show you this as an example. This is a 33-kilowatt single-phase inverter made out of-- using silicon carbide FETs.

You can see the very tight loop bus capacitors in here. These are ceramic capacitors. These are film bus capacitors. And by the way, this inverter was built by Dave Otten at MIT.

What you have here is a full bridge-- each using a bunch of parallel devices for each switch because there's a high-power design. And basically, you have one set of switches, and on the other side of the board, another set of switches that form one half bridge, and then the other pair on the other side of the board forms a second half bridge.

Air flows down the board this way to cool it. You can see things like gate driver chips along here that are distributed to run the parallel gate drives. There's also, on this side of the board, a current sensor. That's the current sensor we'll talk about.

So basically, this is precisely the structure we've been talking about. And as I said, this is about a 33-kilowatt block that runs off of a 720-volt bus for a motor drive application. So you can pass that around. Just take a look at it. Try not to drop it, however. I'll be in trouble if you guys drop it.

So that's the hardware structure. So far, we've been talking about different ways we can start to approximate a sine wave better and better. And as I said, in a lot of applications, a sine wave is what you're trying to synthesize because, for example, in that application, the motor is expecting-- has sine wave back-EMF, and you want sine wave currents going into the motor.

So you're trying to synthesize something that approaches a sine wave such that this voltage, Vx, has a sine wave voltage component plus a bunch of high frequency switching content. And we said there's different ways we could try to make that. work better in terms of getting close to a sine wave by switching.

One technique we talked about was harmonic elimination. And the idea is this. Suppose I would make my voltage, Vx, here have a bunch of pulses. So if I imagine that the low-frequency content was going to try to look like some sine wave and I'm switching between plus VDC, 0, and minus VDC, then maybe what I would do is I would have certain number of pulses, and I would control the angle of those pulses very carefully.

And if I do that and I want to make my waveforms half-wave symmetric so that I get no even harmonics-- so we would come along here, and we would do this. And this would be my pulsing waveform.

And what we would do is we would pick these angles very specifically in order to kill off harmonic content because we said we can figure out the harmonic content by multiplying the PWM waveform by the sine-- if it's an odd waveform at the correct frequency, integrating them when I get 0. That harmonic is not there.

And what that technique tends to do-- if I looked at the magnitude of Vx versus frequency, I will have some fundamental. And I want that because that's what I'm trying to approximate.

And then I'll have some harmonic content that depends on the angles, and I can get rid of the evens easily. But I might suppress one or more odd harmonics. And each pulse and a half cycle I can usually suppress, one harmonic.

What I do get, however, when I do a lot of switching up and down is when I go out in frequency, I will get an increase in the content, because there's a lot of up and down switching, a lot of high-frequency content in this V sub x waveform.

But if I'm running that through a filter-- if this. load is, say, a low-pass filter and it's doing this versus frequency-- so this is the magnitude of H of j. omega, It's very easy to filter off the high frequency stuff. It's hard to filter off the lower frequency stuff. So we basically trade low-frequency content for high-frequency content, and it makes a good trade. So that's harmonic elimination, and you can do that just with this structure.

The other trick we saw last class is harmonic cancellation. And the idea is-- suppose I could take multiple waveforms and add them up and phase shift them a little bit. I'll phase shift the fundamental by a little bit, but some harmonic of interest by 180 degrees so when I add them up, I kill off harmonics,

And we saw one example of this was something like this. If I took two waveforms that have plus 0 minus and I phase shift them, I can get something that looks like this, for example.

So it might go up to one level, another level, something like this, and so forth. And this one-- you can see, wow, that really does start to look like some kind of sine wave very closely.

So harmonic cancellation-- actually, I can view the six-step wave-- or the stepped waveform that we said-- if I have no third harmonic, the waveform looked like this. The Vx waveform looked like this. This is actually two square waves phase shifted just such that when I take their difference, I kill off the third harmonic.

So we're already doing that with this. But if I want to do it on a more sophisticated level, I need more sophisticated hardware that can let me synthesize this more accurate waveform.

If I can do that, the beauty of harmonic cancellation is that harmonic cancellation not only kills the low frequency harmonics that I want to cancel, but it tends to cancel multiples of that.

So we can get-- if I do a lot of harmonic cancellation, I can get a waveform that looks very much like a sine wave, whereas this illuminated waveform-- it's kind of ugly, and I'm just taking advantage of the fact that I can filter the high-frequency junk out. And there's nothing that says we can't combine those two. In fact, people do.

So the challenge with harmonic cancellation is now I need a means to synthesize-- maybe this is VDC. Maybe this is VDC over 2 and 0 and minus VDC over 2 and minus VDC. I've got to be able to synthesize more effective levels to make that work. Any questions about that?

So last class I showed you a transformer-coupled converter that could synthesize these waveforms. And perhaps the theme of today's class is-- there's a lot of ways you can do things, and people do do things because they have different advantages and different circumstances.

So let me show you some additional ways one might synthesize a circuit that can make these kind of fancier waveforms. Let's just consider-- we might call this, here, a bridge leg or a half-bridge leg. And this thing-- in this context, this circuit can synthesize a V01 that's either VDC or 0.

So this bridge leg can basically synthesize two levels. And because this leg can synthesize two levels, this is sometimes called a two level inverter. A circuit that could synthesize for one leg, 0, 1/2, and full, might be called a three-level converter. Kind of confusing terminology, but that's the way they do it.

How could I do that? I showed you one circuit last time. Let's talk about a bridge leg that might do this. And here's one way to do it. This is called a neutral-point clamped inverter or an NPC inverter. This general technique was proposed in about 1980 by a guy named Nabae in Japan.

And here's what it looks like. Suppose I took my bus and I split it into two by just, basically, putting a big tap point on my capacitor here. So here's C big. Here's C big. Let me assume that this is about DC over 2, or I will make sure it stays near VDC over 2.

Then maybe what I could do is do this. I will have S1, S2, S3, and S4. I'm going to have a stack of switches. And then I'm going to have another one that-- here, I'm going to call it S6. And this is S5. So I'm going to have six switches here.

I can also, in some cases, just replace these two switches with diodes. And in fact, in the original version, that's what most people did. So this one might be called an active neutral-point clamped inverter because it uses controllable switches here. You could replace with diodes.

But here's the idea. Let me think about this output voltage. I'll call this V01. If I turn the following switches on, S1, S2, and S5-- oops, S5-- I'm going to turn the top two switches on here-- V out 1 is just going to be VDC.

Suppose I turn on S2, S3, S5, S6. If these four devices are on, this node is then, through multiple paths, pinned to the center here. So then this should be-- V out should be approximately VDC over 2. It's going to be this center-point voltage.

If I then turn on S3, S4, S6 I'm going to get 0 out. So by picking different switch configurations. I can make this output voltage either be plus VDC, plus VDC over 2, or 0.

And if I took two of these things or I took another set of switches and tied it to the same center point, then I could generate this waveform here. Or I could generate this waveform plus PWM between levels and combine harmonic elimination and cancellation.

So this is a bridge leg for a so-called three-level inverter. Any questions about that? Yeah.

AUDIENCE: How do turn on S5?

DAVID PERREAULT: Ah, excellent question. So you don't have to. But here's the neat thing about this inverter. In fact, this is one of the key reasons why Nabae came up with what he was doing originally is the following. What is the off-state voltage of S1?

Well, if I have this thing tied to the middle point, S1 only sees VDC over 2. What about S3 and S4? If S1 and S2 are on and S5 is on, that means S3 sees VDC over 2 and S4 sees VDC over 2.

So by turning on the extra switches, S5 and S6, what I'm really doing-- in addition to providing a path to the middle, I'm ensuring that every single switch only is rated for half the switching voltage. And that's a big deal because-- like, Nabae liked to work on things like traction, where the voltages are really high.

And it turns out a lot of times the voltage you want to use-- you can't buy switches with high enough rating, You would like to make it bigger, but nobody makes the semiconductor switch you want. So if you can build your inverter with lower rated switches, you've got a very powerful advantage.

It also turns out that if you double the blocking voltage of a transistor, typically, the performance is a lot worse. They're a lot less than half as good, Like, the resistance way more than doubles. Usually, it goes as a quadratic, almost, or more. So you're actually, in fact better having two 600-volt switches than one 1,200-volts switch, at least from an on-state resistance and capacitance point of view.

So it's a more complicated circuit, but it reduces the stress on the components. And that's one of the reasons why I was doing it. And a reason why you would put diodes in here-- the original flavor of this device-- basically, just replace this with diodes. The diodes would turn on-- depending upon whether current was going this way, S5 would turn on. If the current was going this way, S6 would turn on. So you could use diodes here as well.

It doesn't do as perfect a job of protecting at least one of the switches if you do that, but it still tends to work just fine. great question. I should make a couple of notes. This structure really only works for DC to AC conversion. Why? Because what I have to do ultimately is-- what is the allowable average current into this node if I want to keep the voltage stable?

AUDIENCE: 0?

DAVID PERREAULT: It's got to be 0. Right. Why? Because I can't put a DC current into a capacitor and have its voltage stay stable. So what that means is if I have a motor drive that's going to draw 100, 500 hertz out here or something, I may have 500 hertz currents coming in here, but no DC current.

So I may have to control the switches a little bit just to make sure that this node stays balanced. And I can't deliver DC current out of this thing. I can only deliver AC current out of it because otherwise, this capacitor-- this node wouldn't say it half the bus voltage. But since I usually have big bus capacitors here anyways and I usually use this when I do want AC currents only, that works just fine. Any questions about that?

So that's one of a variety of ways to do it. Let me show you one more, keeping with the theme for today. And by the way, these are hardly the only ways you can do this. I'm just showing you two very popular ones. This is known as a flying capacitor multilevel inverter.

And here's the idea. Suppose I'll take my DC bus voltage here, VDC. Now I'm going to go out and get four switches. I'll call this S1-- actually, I'll call this S1, S2. I'll call this S2 prime and S1 prime.

And what I'm going to do is I'm going to put a big capacitor here, And I'm going to make VC approximately equal to VDC over 2. I'm going to control this circuit so this capacitor voltage stays at approximately half the bus voltage.

What do I get when I do that. Which switch is on? Let's see. S1 and S2. If these two switches are on, the output voltage would go to VDC. If I have S1 and S2 prime on, I get VDC minus VC, which is approximately equal to VDC over 2.

If I have. S1 prime and S2 on, that will connect here to here to the output. Then I would get VC, which is also approximately equal to VDC over 2. And lastly, if I have S1 prime and S2 prime, I'll connect the bottom two and V out 1 will be 0.

So here, I can synthesize plus VDC, approximately VDC over 2 and 0. Again, this one bridge leg can synthesize the positive half of this thing. I can make another three-level inverter with this. Any questions about that?

Now, I have two ways I can synthesize, roughly speaking, VDC over 2. Suppose I had a positive load current. If I did it the first way, I would have current flowing out of here into the load, and it would charge this capacitor. If I did it the second way, it would flow up from the bottom out through the capacitor and out and discharge the capacitor.

So these two mid-voltage states have opposite effects in whether they charge or discharge this capacitor for a given load current. So by choosing which of these I use at any given time, I can maintain that capacitor voltage near half VDC.

Again, notice that if this is kept near VDC over 2, each of these switches again only sees VDC over 2 blocking. Now, in practice, I can't quite derate them that much, but I have to derate them less than if I didn't have this kind of divide down.

So this FCML, as it's known-- this was, by the way, originally proposed in about the early '90s by a guy named Maynard. So this is sometimes called a Maynard inverter-- is really powerful. It gets really widely used.

And the other neat thing about this thing is that you can actually-- whereas in the other version I showed you, you cannot source DC current, because you'll eventually charge or discharge those bus caps, here, by flipping back and forth between these two states, I could source DC current out without charging or discharging this cap.

So this one's a little trickier to run, and it has this capacitor that flies up and down with the different states. But I can supply DC power with it. So people actually build DC-to-DC converters this way, too. They build three-level DC-to-DC converters. And I could show you the T-type-- there's a bunch of inverters that give you multi-level capabilities like this, and people tend to use all of them. And which one do you use? It depends what you're trying to optimize. People choose different things for different goals.

I should also tell you that just because I said we could do three levels, you can expand your imagination however many levels you like. And in fact, the neutral point clamped-in converter-- I'm sorry, the neutral point clamp converter is sometimes of a bigger family of converters called diode-clamped or active-clamped, multi-level inverters.

This one also expands. So for example-- I'll just show you an example. Suppose I-- instead of having one device, I had two devices here in the top and two devices in the bottom position, my next thing would-- I'd have three devices in the top and three devices in the bottom thing.

So I could do this. I could have one, two, three, four, five, six. Here's VDC. And I will have V out.

And what I'll do is I will have one flying capacitor here so that these two switches are never on at the same time. And I'll have another flying capacitor here. So I've got to be careful about whether these two switches are on at the same time, or certainly, these four switches.

And what I will do is I will control this one. So that this is approximately VDC over 3 and this is approximately 2VDC over 3. Same game. I can pick switch states to regulate both of those capacitor voltages.

Now, each switch only needs to see a third of the bus voltage. Now, I can't derate it down that far in the real world because I have all kinds of start-up and transient and other effects. But nonetheless, what that means is I have, first of all, lower voltage-rated switches overall and more levels I can synthesize.

So I can synthesize, roughly speaking. VDC-- 2/3 VDC, 1/3 VDC, and 0. So it's a four-level inverter. And you want to expand that out. People have built 10-level inverters.

Why would I want to build a 10-level inverter? Because if I'm building an inverter with this bridge leg, my output waveform, V01-- V01 might be able to look something like this. I can build this very beautiful sine wave kind of thing, Very little harmonic content. Needs very little filtering. So the filter I need is very tiny.

And by the way, if I'm building a grid interface inverter, I've got to build that filter inductor, so I can suddenly make that filter inductor really small. So there are motivations to go to high numbers of levels. What's the downside? Complexity. There's all kinds of complexity you got to throw at it. But you're trading complexity for performance.

And like I said, there's versions of this flying capacitor multi-level inverter. There's higher-order versions of the neutral point clamped converter. And people just decide what they're going to do based on what they're trying to achieve. Questions about any of that? Yeah, Jack.

AUDIENCE: [INAUDIBLE] the high-side FETs [INAUDIBLE]?

DAVID PERREAULT: Ah, excellent question. So I will cover details of how you do bootstrap drivers and that kind of thing. The higher the voltages, the trickier it gets. But the crudest thing you could imagine is you need two pieces. A, you need a little power supply that flies up and down. Sometimes that's built with an isolated power supply.

And in fact, on the board I'm passing around, there's a little black rectangular-- a couple of little black rectangular cubes. Those are the flying power supplies. Other times, you can build it with a capacitor and a diode and bootstrap it. The higher the bus voltage, the more difficult that becomes.

The other thing you need, in addition to a power supply to power the gate drive charging is a signal isolator. You've got to send the signal up and down. And up to a few 1,000 volts, that's not too hard. You can out and buy digital isolators. Or at lower voltages, a few 100 volts, you can buy ones that aren't even galvanically isolated and do the level shifting for you.

When you get to higher and higher voltages, it becomes harder and harder and harder, and then you're using these optical connections and everything else to do that. But there's a lot of practical issues with what you would choose to do there. And that might be another thing that would influence how many levels you want to use.

At the same time, switches tend to be more important-- more expensive than isolating signals If you're at high power. If you're at low power, you wouldn't want to do that too much because then the switches aren't expensive, but the isolators are. And that's why there's a plethora of techniques that people use.

I will say, by the way, that this kind of technique-- very popular in integrated silicon converters like CMOS kind of converters, because all of this stuff can be integrated on a CMOS die, and then you're using native devices with low voltage to synthesize a higher voltage. Great question. Any other questions? Yeah.

AUDIENCE: So when you were regulating the voltage, are we doing that by fixing the respective duty ratio? Or how are we getting that?

DAVID PERREAULT: Yeah, so what you might do in this example is you might say, OK, my controller is telling me I need to synthesize the middle level, And if I know which direction the current is. I would say, is this capacitor voltage a little low or a little high? And I'd pick my next state as to whether I want this to decrease or increase. So I basically use the degrees of freedom. I have to choose which way to switch to control that voltage.

In this case, you can build a little state machine that says-- given that there's two voltage that I need to regulate, which one's big, which one's small, and you can make a little state machine that will control both of them. Any other questions?

So if I can build multilevel inverters, I can do higher-order versions of harmonic cancellation. I can do harmonic elimination at the same time. Make really, really-- waveforms that'll have very nice purity. I don't need to filter much, and I get what I want. So that's pretty cool.

I'd like to talk about one other aspect. And that is-- I've focused on the situation where, boy, we want to-- to get what we want-- we don't want to switch too often. And that is especially true at high power levels because at high power levels with high voltages, the switching costs you a lot, and you really focus on how to minimize that.

But in some cases, I have really good MOSFETs at low voltage. I can switch a lot, I don't have to say, oh, I can only switch 10 times per AC cycle. I'm trying to synthesize 60 hertz. I can switch 100 times, more.

I think that inverter up there in the application that we're doing-- I forget how many switchings per half cycle we do. But it's-- I want to say I want to say it's something on the order of 50 or something. Actually, it's-- no, it's more than that. So it's pretty high. So you can switch a lot If you're at low power or lower voltage. Actually, that's not even particularly low voltage. Like I said, it's a 720-volt bus. And that one switches at 80khz to synthesize a couple of kilohertz output signal with some caveats that I'll explain shortly here.

But suppose I can switch a lot. How would I do that? Well, one of the things I told you was-- if you're trying to do this harmonic elimination, man, you got to time those angles really carefully, Otherwise, it doesn't work out too well. And even the same is true of harmonic cancellation because you're phase shifting stuff, So you've got to get your angles right in order to do very precise harmonic elimination.

But if you take the handcuffs off on how fast you can go, you can switch a lot of times. How do you want to decide how do you want to switch to synthesize a sinusoid, for example?

But I'll ask even a broader question. What if it isn't a sinusoid I want to synthesize? And in some cases I don't even know, ahead of time, what I want to synthesize. So, for example, if I want to build a switch mode audio power amplifier-- I want to play music out this thing-- I may not know ahead of time what the waveform I want to synthesize is because it's different every day. It depends what button the person pushes on their iPod or what you select on Spotify.

So I need some way to take an arbitrary waveform and synthesize an arbitrary AC output signal very quickly. And there's, again, different ways you can do that.

But I will show you one that's very common, and this is sometimes called sine triangle PWM. PWM is, again, pulse width modulation. So we're going to modulate the widths of individual pulses in order to control the output.

Sine triangle PWM or sine triangle intercept PWM speaks to the way I would do it if I was trying to synthesize a sine wave. But you don't have to synthesize a sine wave. We can synthesize whatever we want.

And here's the idea. Suppose I take some triangular carrier waveform here. And I'm going to make-- I'm going to draw this as triangles, like this, with some period, T, like this. And T is going to set my switching period and so forth. So 2T, et cetera.

And then suppose I have some reference waveform-- Vref-- that's doing this. So here's Vref. What I'm going to do is I'm going to have a comparator that is going to, say, take, at the positive input, Vref, and, at the negative input, the triangle wave, and output is going to come q of T.

If I was doing this on a half bridge, what would I be thinking? I would say this gets q of T and this gets q prime of T is equal to 1 minus q of T, So whenever he's on, he's off, and vice versa.

So what is the idea here? The idea here is the following. If I run through this comparator, what I'm going to get for my q of T is the following. Whenever the reference exceeds the triangle wave, I'm going to be on. So here, my switch is going to be on. Here's q of T. Is going to be on here. And maybe a tiny time here,

When I come to the next cycle and my reference is bigger. I synthesize a wider pulse. Here, the reference is bigger, and I've got an even wider pulse. My height shouldn't be growing here. The heights are all the same. Here, it's even wider pulse and so forth.

So what happens is the fraction of time-- so each of these is spaced-- T, 2T, 3T, 4T, et cetera. So the spacing of the switching is constant, I have a constant switching frequency.

The pulse width that I'm going to get in any capital T seconds is going to be proportional to the-- if I call this triangle height, H, d of T, in any cycle, is going to be equal to the magnitude of Vref over H. Why? Because for a triangle wave, if this was flat, it's just-- the percentage-- if I go from 1 down to 0, just linearly.

So the duty cycle in any given time, and hence the average value, V01 average over a switching period, is going to be precisely d, which is precisely the ratio of Vref to the top of the triangle. Any questions about that?

Now, if you recall, this is the exact same trick we talked about doing with DC-to-DC converters. We just had a sawtooth, and we compared it to the level of a signal.

Here, I used a triangle wave, like a balanced triangle wave. You could do this with a triangle wave. You could do this with a sawtooth, with a capital period T. You could do it with a reverse sawtooth, like this. All of those things work. But what happens is all of them will give you the same average on time-- or average high time, or the same average V01 over a switching cycle.

But what differs-- boy, that's a bad drawing. What differs among these is-- what differs among these is the harmonic content. So the high-frequency switching content-- how much is there and what does it look like-- depends upon whether I use this waveform or this waveform. At DC, the local average value doesn't vary.

And it turns out that the balanced triangle wave tends to give you less harmonics than anything else, or it gives you less high frequency content or better-shaped high frequency content. So mostly, for inverters, people use a balanced triangle wave like this. Mostly for DC-to-DC converters, where they don't care about the switching ripple, the detailed harmonic content of the switching ripple, they often use a sawtooth because it's a little bit easier to generate. But either of them are legit, and either of them do the same thing. Questions?

And I will say that, by the way, if I want to synthesize a sine wave voltage, this will let me do that. One way I could do this-- I'm showing this for one half bridge, but we don't have a half bridge, We tend to have a full bridge-- two half bridges that are working together.

And now, there are a lot of ways I can do my control, One thing I might do is say, look, this scheme-- if I have a triangle between 0 and some positive number, call it 0 and 1, I can take, then, a reference voltage and make my reference voltage between 0 and VDC. I can synthesize an average voltage between 0 and VDC with a half bridge.

Well, maybe what I'll do is I'll just do that right here with this half bridge. So based on the triangle comparison, I can synthesize this voltage to be anywhere between 0 and VDC. If I want my output to be positive, maybe I just turn on this switch. And I can synthesize anything between 0 to 0 and VDC on this side. So Vx is between 0 and VDC.

Or if I want my Vx to be negative, I'll turn on this switch. I'll make this signal q of T, and I'll switch this guy just the opposite. And I can synthesize my load voltage going negative, from the top. So in that case, I would have one half bridge that's doing pulse-width modulation and another half bridge that's doing what we call unfolding, deciding am I going to get a positive output or a negative output?

When would I like to do that? I would like to do that when I have one set of switches that's really good at high-frequency switching and another set of switches that's got really low conduction loss, but they don't switch very often because they only switch when the voltage reverses. Any questions about that kind of scheme?

Now, that seems good if I have two different kinds of switches. But on the other hand, a lot of time, if I'm just going to have four switches, like in that board I passed around, all the four switches are the same, That's maybe not a great way to do it. I'd like to balance my switching across all the switches, So maybe-- another way I could do it is shown here. So this is the scheme I just described to you. So we will make the-- we will put through the reference waveform that's-- if I'm synthesizing a sine wave, I will make one side modulate like this, and then I'll flip the switches on the other side to unfold it and get the signal I want.

Or I could say, OK, when my output voltage is positive, I'll modulate this side of the bridge and hold this switch on. When my output voltage is negative, I'll hold this switch on and modulate this side of the bridge.

And I can do it with this same signal. That way, the waveforms are balanced across the half bridges. One half bridge switches when it's positive. And the other one has one switch on. and the other switch switches when it's positive, and the other half bridge has one switch on. Does that make sense to everybody? Which one do I use? It depends on my design.

Now I'll show you the technique that actually gets used in the system that that board's going into, and that looks more like this. And this is kind of cool because this basically relates back to what we were talking about, about interleaving and harmonic cancellation. Only now I'm not thinking about harmonic cancellation of the fundamental waveform, but rather the switching contents.

Here's the trick. Suppose I have two comparators. And I'll have one triangle wave, and let me make the triangle wave bipolar. It goes from minus 1 to plus 1. And then I'm going to have a reference signal that goes like this.

And so what am I going to do? The top bridge I'm going to switch when the reference signal is bigger than the triangle wave, just like I told you. But now the triangle wave goes to minus 1, and the reference signal will go to minus 1. I've just shifted them both down.

The other bridge, I do the same thing, except that I put negative 1 on the reference signal. So basically, the other guy switches 180 degrees out of phase of the switching frequency. And I get an output switching waveform that's also pulse-width modulated. But if each of these clocks, or if this clock was at 100 kilohertz, the output ripple basically would be near 200 kilohertz.

So the left side of the-- the left, half bridge is switching at 100 kilohertz. The right half bridge is switching at 100 kilohertz. But the ripple that I get at the output, because the fundamentals are canceling of the ripple frequency, is, like, 200 kilohertz, or it's around 200 kilohertz because there's some very slight variation in the frequency content of the output because of the reference. But if I was putting out a DC signal, it would be 200 kilohertz. Any questions about that? And in fact, that's what we do in that design.

AUDIENCE: Where is the reference signal coming from? Is this whole idea just to amplify the reference?

DAVID PERREAULT: That's an excellent question. Where is the reference signal coming from. The reference signal is coming from whatever you want for your output. So in our case, we might say, for some operating condition, I want this many volts AC sinusoid at my output at certain frequency, and feed that in-- feed that through a DAC and put that into a sine triangle, or do the digital equivalent of that.

Or if I wanted to build an audio amplifier, I might have some D-to-A converter that's spitting out my audio signal. And this is the audio signal that I want to amplify. And the voltage across the output is going to be the amplified version of this reference signal. So if that's Mozart. Mozart comes out the other end at high power into your-- maybe this is a model for your speaker coil. That make sense?

In fact, I should say most audio amps today-- if you go by a car audio amp or something like that-- they tend to use switching amplifiers. They don't use linear amps anymore because they're just way, way more power efficient. Same reason why you'd go to a DC-to-DC converter instead of a linear regulator. People do audio amps with switching these days. Excellent question. Other questions? Yeah, Jack.

AUDIENCE: What is the input frequency limit for that signal?

DAVID PERREAULT: Ah, That's a good question. So there's a couple things that drive that. One of them is that the output frequency-- if I actually looked at the spectrum of the output, Vx-- it kind of contains mixing between the reference signal and the PWM carrier signal. So generally, you need the reference you want to synthesize-- typically, you want it to be much lower frequency content than the carrier. And how high a carrier-- that depends on how much distortion you want to allow in your-- say, your output voltage or current.

The other kind of thing you have to think about is, what does that look like compared to this load? So if this load was inductive resistive, the R over L constant-- you'd have to have your switching frequency high enough that the switching stuff's going to get filtered.

And you want the signal you want to get through that to be cut off enough that it really generates the current. Or if there's-- if you're interested in the amount getting across the resistor, you'd have to have an upper bound on that-- from that, too. Yeah.

AUDIENCE: So would this technique be used for a solar inverter? [INAUDIBLE]?

DAVID PERREAULT: Sure. That's a very good question. If what I was trying to do is do something grid-forming, meaning I want my inverter to create the grid. I want an AC output voltage. Yes, I would synthesize a sine wave that would be-- if I wanted to keep it simple, a single phase-- 120 volts. It's 170 peak. And I'd chop it at whatever frequency, and I'd filter it, and then I'd get, across my load, the voltage I wanted, hopefully.

If, however, I'm not a grid-forming inverter, I just want to stick current into the AC grid, then I don't control voltage. What I control? Current. I would like to be a current source into the grid, preferably looking like a sine wave.

Excellent question. How do I do that? One thing I might want to do is say, OK, I know that I want to put-- from my solar current-- from my solar panel. I have to generate so much power into the grid. That means I want a 1-amp sine wave. So I'll have some I reference that might be a 1-amp sine wave going into the grid.

What I'll do is I will compare that to my inverter. So I will come back to my inverter. I will measure this current, I'll call that I inverter. Here's I inverter.

And then that'll be-- this is the error current, I error. I'll put that in through some big gain compensator. Basically, I'll amplify the error, and then I will use that to create Vref. So the bigger the error, my current is too low, I get a bigger reference-- drives the current up, and you close the loop that way.

AUDIENCE: Where did you get Iref?

DAVID PERREAULT: Iref is coming from saying, how much current do I want to stick into the grid? And that's a question of power, and that comes from maximum power point tracking on your-- on your design.

And I should say, by the way, there's a lot of applications where we control current. Injecting into the grid from solar panels is one. Another thing is if I have a motor, torque is proportional to current, So in a lot of motor drive applications, I don't just control voltage. I say, how much sinusoidal current do I want to feed it? And so I've got to-- ultimately, I'm trying my command-- the thing that's I'm controlling to control torque, for example, is going to be current.

Now, this is one way to do things. And the beautiful thing about this is everything's more or less fixed switching frequency, I know exactly how many times I'm going to switch per cycle. And I can take that to the bank because I know maybe my switching loss and my loss of my inverter. Are there other ways you can do it? Sure. Here's another one. And I should say-- all these things I say, oh, you could do this, or you could do this. People do all of them. It just depends on your application.

Here's another idea. Maybe I would do this. Maybe I would have some current reference, Iref. And I would have I would create Iref plus delta I, and maybe Iref minus delta I A hysteresis band. I'll say, OK, if I'm here turn on the top switch so that I have higher voltage applied.

And the IL-- this is the I inverter-- will go up until I hit this top threshold where it's Iref-- it's delta I too big. And then I'll turn on the bottom switch, and the current will fall. And I'll just chop it like that.

And if I zoomed out, what I would get is some waveform that does this. And maybe I could synthesize a sine wave, or I could synthesize my music waveform, whatever I wanted.

So this is known as hysteretic current control. Its advantage is I guarantee that once I get into this, I'm always going to be within 2 delta I of my reference. So I kind of control the current very directly.

The downside is what rate does this thing switch at? Well, that's a bunch of analysis, but it varies. It depends upon the load and everything else, and the band. So you don't get to really say exactly what frequency you're going to run it. You can analyze it, but it's a variable frequency thing. Sometimes I really don't like that because it affects my EMI and my filter design and everything else. So depending upon whether you want to do this or that just depends on the application. People do both.

We are pretty much out of time. The last thing we'll show you-- and I'll just show you really quick. We talked about doing PWM. And I'll just show you-- let's see. I got to go back.

Here you go. This is precisely the digital equivalent of sine triangle PWM. And what you can see here is that there's a ton of switching going on in the blue waveform, That's the inverter output.

But in the red waveform, which shows the spectral content, at low frequencies, you have very much the sine wave you're synthesizing. The rest of the junk is kind of spit out to very high frequencies that are easy to filter, and I'm not even showing them. That's how high they are.

So if you really can pump up the volume in your switching frequency, you can make very pure waveforms. Even though that's extremely distorted, all the distortion and stuff you don't like is such high frequency you can filter it off and get rid of it. And we are over time. I apologize. We will take up a slightly different theme on this topic next class.