[SQUEAKING]

[RUSTLING]

[CLICKING]

DAVID PERREAULT: OK, good afternoon. So last time, we started talking about power electronics with three-phase elements. We spent our time talking about just diode rectifiers. And that's actually quite important because when you're building really high-power systems, very often that's the kind of thing you have at your front end between your AC to convert to DC.

At the same time, a lot of applications require high-frequency switching power conversion. So the kind of things we were talking about single-phase inverters, the question is, what do you do with a three-phase inverter? And one possibility is just to, in some cases, replicate a single-phase system.

So you could imagine if I had, say, a three-phase machine with open-ended windings, maybe I would have something where I had a phase A and a phase B and a phase C. And each of them have a pair of terminals. And I say, OK, I can go build a three-phase inverter basically by building three single-phase inverters.

So maybe I come in from my DC bus, VDC. And I could come in and say, OK, let me just build a single-phase inverter, which would be, say, four power MOSFETs. And I'll show the diodes on the power MOSFETs just for clarity.

And I could come out, and here's one terminal, and here's the other. And I build this to drive phase A. And then I build another copy. This is phase A. I build another copy, drive phase B. And I bring it in from the same terminals. And then I build another copy, and I drive phase C. And I could again build in same DC bus. And away I go.

So I have 12 devices. I just have single-phase inverters, and I run it, And in fact, sometimes people do this. And in fact, actually, at MIT right now, we're building a megawatt system that basically has this structure. It's sets of single-phase inverters that drive a three-phase system, where the machine is a custom machine and has so-called open-ended windings like this. So it works perfectly well.

However, the vast majority of, say, three-phase motors are not just separate windings connected like that. Typically, if, for example, I have a three-phase motor, what I will have is either something connected in wye, where this is one motor winding, a second motor winding, and a third motor winding, where I'll have A, B, and C.

I only have three terminals. Or maybe I have-- this is a wye connection. Maybe I have a delta connection where the machine is wound like this. And again, I have A, B, and C. In this case, I might have a neutral. In this case, I have no neutral.

So in this case, I can't just arbitrarily apply voltages from three independent single-phase inverters. Why? Because if I accidentally connected phase A to high with one inverter and phase A to low with another inverter, I'd out my DC bus and everything would blow up. So if I'm going to drive something that's connected as a three-phase load in wye or delta, I need a different plan.

OK, so what do people do? It's very much the structure that we saw with the three-phase bridge rectifier. Instead of having diodes, however, I will have active switches. So maybe I would come up and build something like this. I'll have my DC bus come in. So I'll have two terminals, VDC.

I will have a first bridge leg like this. And let me make sure and show the diodes again. And this can generate VA. A second one can generate VB. And a third one can generate VC.

Maybe what I'll do is I will consider the midpoint of this bus for reasons we'll talk about. And I'm going to call this node V sub R. OK, that's my reference voltage. It's just the center point of the bus. Where I ground this thing, we can discuss that. But let me just pretend that this is exactly in the middle. It's VDC over 2 from here. And we'll see why we're going to talk about that.

And that's my inverter. And maybe I would connect it up to some three-phase load. So maybe it would look like this. I would have my three phases. Here's one. Here's two. And here's three. And maybe I would connect this point to what I call V sub N, or the neutral point. That's this neutral point if it was a wye-connected load. Or I could just have line to lines if there's no neutral.

So what's the advantage here? Well, you'll notice that this structure only requires six stitches instead of my 12. So it's simpler. It also turns out that as compared to that solution-- let's suppose I could connect my motor either way I wanted. I could connect it in wye or delta, or I could connect it as three independent windings.

Because this circuit, like this phase leg-- whatever current goes out, it returns through one of the other phase legs-- I get an advantage in terms of the utilization of the devices. So I will either require less semiconductor area to get a certain conduction loss in this structure, or I will have lower loss for the same semiconductor area-- lower conduction loss for the same semiconductor area as I will compared to the top solution. So this solution in most three-phase applications is vastly preferred.

Where's the exception to that? It turns out that if I wanted to generate a certain ripple in an inductive load, like a motor inductive back emf kind of load, it turns out that you end up having to switch this bridge at a higher frequency than you do the three sets of single phase. So in a system where it's a really high-frequency output and your switching losses are heavily dominated, the top solution actually ends up starting to be advantageous.

So in very specialized, high-speed machines, for example, you might go for the top solution of independent phase drive if you're allowed. But the vast majority of three-phase inverters or three-phase DC to AC converters, power could flow either direction. Where you have active control is with this three-phase bridge structure. Any questions about that so far?

So one thing I wanted to note about this is that I want to pay attention that these devices are-- especially if the load is going to be inductive, so I don't actually control instantaneously, say, what the direction of that load is-- I usually have devices with reverse diodes on them, or I use MOSFETs, which have it built in because that way when I turn one device off, I can always guarantee there's a path for the current.

So these are single-directional blocking, bi-directional carrying switches. So this is a kind of voltage source inverter. There are also three-phase current source inverters just as there are with single phase. But again, this is vastly the most common.

Just to illustrate this, I'll show you this guy here. This is actually an inverter out of an older model Prius. You'll see the two in here. That's the DC bus coming in. That's where the capacitors and everything else would hook up and go back to the battery. And then there's phase A, phase B, phase C outputs.

Each of these, what you'll see is a pile of devices in parallel in order to provide the required power. And when you look closely at this, you're going to see that there are essentially three terminal devices that go back to gate drivers. Those are the IGBTs in this older model. Today, they use silicon carbide FETs. And in antiparallel with them, because the IGBT can't carry reverse current, they have actual diodes. So in their design, the actual diode is separate from the actual switch.

And so I'll pass this around. You can look at it. And by the way, if your Prius happens to be missing an inverter, I swear I had nothing to do with it.

So what can I do with this inverter? Keep in mind-- go ahead, Jack.

AUDIENCE: I just have a question about the inverter? Why would they choose to do IGBTs with diodes instead of MOSFET?

DAVID PERREAULT: Ah, that's a very good question. Because if you think about a MOSFET, a MOSFET's on state looks like a resistance, whereas an IGBT's on state looks like a constant drop. And it turns out that at the time that was built, for a very high-voltage device, you could get lower forward drop for a given device area with an IGBT than you could with a FET.

And so maybe until 5 years ago or maybe 10, it was most common to use IGBTs especially at high voltage. More recently, with wide-band gap devices, people are starting to move over to sometimes IGBTs but mainly FET kind of structures because you can get away with that. And it turns out to be lower loss. So whichever switch gives you less loss is ultimately what you want to choose to use.

So let's think about what we can do with this. And keep in mind, we said one of the big advantages of having three phase is that we get natural cancellation of all the triple-N harmonics. And we're going to see that with this inverter structure.

But essentially, if I think about it, I have three bridge legs. And each one of them, either the top switch can be on or the bottom switch can be on. So it's like there's 2 to the 3, or 8, possible states for these inverter switches to be on, assuming one switch is always on in each bridge leg.

So eight possible states. If I come and show you what those states look like, they look like this. So if I think about the three inputs, I can have bottom, bottom, bottom; top, bottom, bottom; et cetera. There's eight possible states.

And these are what you get for AR, BR, and CR normalized to the DC bus voltage. And then this is what you get for VAB, VBC, and VCA at the line-to-line voltages. And this is what you get for the line-to-neutral voltages on this load.

And I want to emphasize that in this structure, I'm not proposing that we connect this reference voltage to the neutral. This neutral, if it's a wye-connected system, we're typically leaving this neutral floating. So it's not connected to anything. Or in a delta-connected system-- I drew this as if it were a wye-connected system. If it was a delta-connected system, there wouldn't even be a neutral.

So the neutral is not connected to this reference voltage. So how do I figure out what the voltage say between VA and the neutral is? So this would be-- I'd call this VAN. Well, if I imagined that my load, which might be a motor, acts like three identical impedances, then it's just a voltage divider.

So if I have one bottom switch on and two top switches on, VN goes to 2/3 of the way up. If two bottom switches are on and one top switch is on, it's 2/3 of the way to the bottom. If all the top switches are on, VN goes to the top. So the neutral just goes up and down instantaneously based on voltage division from these individual values.

Now, I'm not saying you could not tie the neutral back to the reference, the center of the bridge. In fact, that would be just this circuit with half bridges instead of full bridges. It's just not common to do because you lose the three-phase harmonic-- the three-phase triple-N cancellation if you do that. You lose a lot of the benefits if you don't let the neutral float.

So this just shows the voltages I can synthesize. Two of the eight states basically short out the load. Because if all the top switches are on, basically I'm tying all of these three nodes together. So they basically just short the three terminals together. If all the bottom switches are on, I do the same thing. And any other switch configuration, I get some other set of states on my inverter output. Any questions about that?

AUDIENCE: So where's the voltage divider between again?

DAVID PERREAULT: So for example, suppose I took-- let me draw it this way. Here's VDC. Suppose I turn the top switch on here. So that means I tie this guy to the top. If the top switch is on here, I tie this guy to the top. And if the last guy has the bottom switch on, I tie him to the bottom.

So if I have two to the top and one to the bottom, this node voltage, if these green boxes are all equal, impedances will be 2/3 of the way to the top by voltage division. Does that make sense?

OK, so that's the way you tend to think about it. So one way I could run this is very much like the inverse of the diode rectifier I talked about last time. If you remember last time, as the AC waveform happened and drove the diode bridge, the individual diodes turned on. If you remember, it was like 1, 2; 2, 3; 3, 4; 4, 5; 5, 6; 6, 1. So it just went around in this six-step pattern.

Well, I could actively control these devices to follow through those six patterns, ignoring 000 and 111. So those are the two I wouldn't use. And this is what it would look like if you did that. You would see that, say, each half bridge in this inverter would be on half the time.

So A might be on exactly like a square in phase with a sine wave. B would be on but shifted by 2 pi over 3. And C would be on with high shifted by another 2 pi over 3. So I'd have three square waves, VA, VB, and VC with respect to the reference.

What would happen at the load? Well, if I looked at VAB, VBC, and VCA, that's these next-- the next three here-- and I apologize; it's not the prettiest diagram-- are the line-to-line voltages. So all I'm doing is instead of looking at what this is with respect to the reference, I'm just looking at the difference between VA and VB.

And what you can see is that these three suddenly take on this six-step waveform, which looks just like what the diode rectifier was doing. It's just that now I'm imposing a DC voltage with that six-step pattern to say the machine or whatever it is I'm driving. Any questions about that?

Now, if I do that, because the line-to-line voltages-- these square waves, which is the line to reference voltages, have third harmonic. But because I'm taking the difference between two of them to get the line-to-line voltage shifted by a third of a cycle, all the triple-N components disappear. So the line-to-line voltages have no triple-N components.

So whatever my motor sees it's line to line, if I'm driving, say, a delta-connected machine, it will see no triple-N drive. And that's just free because I have this three-phase inverter structure. If everything's balanced, it can't generate three-phase line to line. Any questions about that?

What would it do as regards the line-to-neutral voltages? Well, if I did that, we said this thing is basically a voltage divider. But I could think about it basically being-- the neutral voltage as being of the average of VA, VB, and VC.

And what essentially happens is that each of VA, VB, VC has a fundamental, and then it has, for example, a third harmonic because it's a square wave. The fundamentals are 120 degrees out of phase, but that means that the third harmonics are exactly in phase. So the third harmonics of each of these square waves add at the neutral, and you get a neutral voltage that bops up and down at the third harmonic.

So this has-- the neutral contains the third harmonic and actually all the triple-N harmonics. And so as a result, just as the line to line didn't have any triple-N harmonics, neither did any of the line-to-neutral voltages. So if it's wye-connected machine, the line-to-neutral voltage, which is really what's driving current in this thing, no triple-N harmonics. Does that make sense to everybody?

So this kind of operation is what's known as six-step operation. At each AC cycle, I'm switching my devices as minimum as I can. And I just go through and step sequentially around in a cycle. And in fact, the Tesla egg demo that we brought, that's basically what it was doing. It had six switches, and it was just boom, boom, boom, six-stepping around to generate a rotating magnetic field which spun up the egg.

That's a fine way to operate. And in fact, it generates the largest fundamental component you can generate out of this inverter. Why? Because each half bridge is basically generating a square wave-- the largest fundamental component you can generate when you're not tying the neutral back to the reference point, I should say.

That's not the only way you could drive this. You could actually start thinking about doing things like harmonic elimination waveforms. Instead of generating square waves in between the top and bottom switching, you could start thinking about putting notches in those waveforms. Now, if you're going to put notches in the waveforms, you wouldn't try to eliminate the third harmonic, because the third harmonic is going to go away. You might try to eliminate the fifth and possibly the seventh harmonic.

So the same kind of games we talked about playing with a single-phase inverter, you can play with three-phase inverters. It's just that you're doing each phase and then you're shifting that by a third of a cycle to get to the next phase. Any questions about that?

So those are the kind of games you would play if you're at very high outputs where you don't want to switch your devices very much. But in a large part of the operating range for modern electric machines, you don't have to do that.

You can do a much nicer job of synthesizing a waveform that would look very close-- that would have a fundamental content that follows very closely, for example, a sine wave, if you were driving a machine that wanted sine wave voltages, or that would generate three-phase waveforms if you wanted to make a three-phase uninterruptible power supply that would synthesize a three-phase output that would look very much like sine wave outputs.

How would you do that? Well, let's think about that. What would I do? Well, what I could do is I could make each of the phase outputs, think about them independently, generate something that looked like a sine with respect to the reference, on average. Well, how would I do that? The same way we did things with DC-to-DC converters or DC-to-AC converters and single-phase DC-to-AC converters.

Maybe I would take and say, OK, suppose I wanted VAR to follow some sine wave. I could create some triangle wave that's a carrier wave that does something like this.

And I'm going to exaggerate the ratios between the switching frequency and the line frequency just for visualization. But here I have a triangle wave. Maybe I'll call this minus VT max and VT max. We might think of those as 1 volt, plus or minus 1 volt. And this is 0.

And this we might think of as T switch. T switch is equal to 1 over the switching frequency that I'm going to switch at. And then what I could do is I could create a signal. Let me call this--

I'll call this VAR ref. This is the reference voltage to which I want within a switching cycle VA to follow with respect to the reference. So what would I do? I'd then put this into a comparator if I called this V triangle.

So I could then just create a comparator and say, OK, into this terminal, I'm going to put VAR ref. Into this terminal, I'm going to put V triangle. And then out of this comes q1 of t, which is equal to-- which is equal to 1 minus q4 of t. And this is q1, so q1 of t goes here, and q4 goes here.

So whenever this switch is high-- whenever this switch is on, this switch is off and vice versa. Why would I do that? If you think about it, if Vref is all the way at the top, that means the top switch is always on. If VA is near the bottom, that means the bottom switch is always on. If the reference is right here in the middle, this triangle wave goes from a plus and minus, I would have it switching with 50% duty cycle between top and bottom.

And on average, with respect to VR, this would be 0. So I can linearly change the average voltage synthesized by VA with respect to VR over one switching period based on this triangle intercept PWM. Does that make sense to everybody?

Now, I should say I'm showing this with balanced triangles. You could do this with any other kind of waveform, like a sawtooth waveform or something like that, where at least there's a proportionality in time between the height and how long the comparator trips for. But balanced triangles tend to give you better switching harmonic content than nonbalanced triangle waves. Does that make sense to everybody?

So suppose I had a machine and I wanted to create some waveform that looked like something like VAR on average. So I'm going to draw this bar, meaning average over one switching cycle.

I would like this to be V sub m sine of omega line t, for example, where I'm assuming that omega line is much, much less than 2 pi f switch. That is, this is a very slowly varying sine wave. So what I'm saying is I have some triangle wave like this. Maybe I'll draw it bigger.

This is my PWM waveform, and I'm trying to generate a sine wave that maybe there's a hundred switching cycles or something like that per line cycle. So this is 2 pi over omega switch. And this is 2 pi over omega line versus time. Does that make sense to everybody?

And as long as-- and I should say, what I'm trying to do is I need my-- if this is V sub m sine of omega lt, I need V sub m scaled such that at least the-- I'm sorry, this is VA-- this is the reference voltage-- is less than VDC over 2.

So let me clarify how we're going to do this. And I apologize, that's probably not very clear. What I'm going to do is I'm going to make VAR ref equal to V sub m divided by VDC over 2 times V sub triangle. And then I'll make that sine of omega lt.

So this thing-- this thing is the thing-- this is VA ref that I'm going to compare to VT max. So this is-- I'm sorry, this is VT max. So if VT max was 1 volt, just for simplicity, I'm going to make V sub m over VDC over 2 the height of my reference waveform.

So I can create-- if I do this, then what I'm going to get is VAR as a result of running this thing into this comparator network because this is my triangle away from my controller. The peak of my voltage is switching between plus VDC over 2 and minus VDC over 2.

So what I'm going to get is VAR is going to be equal to Vm over VDC-- it's going to be equal to Vm sine of omega lt, as long as V sub m is less than VDC over 2. Is that clear? Or should I try to explain that better?

I can modulate this thing up and down such that what would be the biggest sine wave I could create without distortion with this scheme, or at least with what I've told you so far, would be if V sub m was actually equal to VDC over 2, which would mean the top of this waveform would exactly touch the top of the triangle wave. So I wouldn't have any period in which I'm exceeding duty ratio of 1.

The problem here is if my reference waveform exceeds the height of the triangle wave, then I'm going to get some time period where I stop switching, and I just have the top switch on constantly or the bottom switch on constantly. And then I get some distortion.

By keeping my reference always below the height of the triangle wave, I kind of keep the no distortion relative at least over a switching cycle, averaged over a switching cycle. Any questions about that?

So when we do this, and I say, OK, I'm going to do this for phase A, I do the same thing for phase B, except I shift it by a third of a line cycle, and phase C is shifted by a third of a line cycle, then at least as far as the local averages, I'm creating a voltage A with respect to reference that follows a sine wave, a voltage B that follows a sine wave, a voltage C that follows a sine wave.

And yes, each of these nodes is going up and down, but if this load is inductive, it will filter that, and then the currents will look very sinusoidal. Does that make sense to everybody? So I can really generate something that looks like a sinusoidal, at least on average, output, just the way I would do it with a single-phase inverter. So this is the same trick of triangle intercept of PWM that we talked about before. Any questions?

OK, so that's nice. This technique is very good when we are trying to synthesize a waveform with a lot of purity. The six-step waveform I showed you before had a ton of harmonic content, a ton of low-frequency harmonic content in it. This PWM waveform really only has content as long as you stay where your reference never exceeds the height of the triangle.

It really only has of harmonic content near the switching frequency and its harmonics. There's some difference with the line frequency and the carrier frequency stuff, but it's all pushed up to the switching frequency. So if I have a high switching frequency, I get a very nice purity waveform, at least averaged over a switching cycle.

But I said, what's the largest waveform I can synthesize that way, at least with what I've told you so far, is I can make this voltage V sub m up to VDC over 2. So this inverter can switch and it can synthesize-- I could think of-- I could think of this as synthesizing VAR, VBR, and VCR, which is going to also give me the same thing for VAN, VBN, and VCN, because there's no triple-N content here, of being V sub m sine omega line t so long as Vm is less than or equal to VDC over 2.

If I'm going to synthesize a sine wave like that, we often define something called m or the modulation index. Sometimes this is called the depth of modulation also, which is essentially V sub m divided by VDC over 2. This is how big a sine wave am I generating? How big is this sine wave relative to VDC over 2 or relative to the maximum I can do without distorting?

If at least with what I've told you so far, I make m greater than 1 or V sub m bigger than VDC over 2, then what will happen is I'll stop switching for part of the waveform, and I'll start distorting. Any questions about that?

Well, that's fine. And through this whole period, at least averaged over a switching cycle, if I'm doing this sinusoidal or sine triangle PWM, this neutral point averaged over switching cycle always has 0 average voltage, in other words, equal to the reference voltage, halfway between the bus voltage. So I'm not getting any third harmonic swinging around on the neutral.

Why do I think about this? Because if you think about an electric machine, and this is a typical application for this kind of thing, the back emf tends to be proportional to speed. And how much current I want to drive in the motor-- if I want more current, I try to drive more voltage, which will drive more current into the windings.

So very often, the amplitude of the voltage, if I thought about some machine, which is just sitting here connected like this-- and maybe I'm just trying to generate-- I'm trying to use my inverter to generate essentially three voltages, which are then going to drive current into these phases.

At low speeds, the back emf voltage of the motor is low, and I want low-- or I want low current. I want these voltage waveforms-- these are now at the line frequency. These are at V sub m sine of omega line t. I want these voltages to be small. But as speed goes up and the back emf internal to the machine goes up, I want a bigger V sub m, for example.

So very often I can-- basically, by changing the V sub m, by changing the amplitude of my reference, I change these voltages' amplitudes, and then I change how much current I'm driving into my load. So I often want to modulate that. And typically, as I'm going up in speed or I'm going up in power or both, I'm trying to make V sub m bigger. Does that make sense to everybody?

Now, the challenge is eventually I run out of headroom. Suppose the internal back emf of these motors is getting really big. So I need a pretty big voltage V sub m to drive current into my load. Once V sub m equals VDC over 2, suddenly I'm going to start distorting, right? I'm not going to be generating just sine waves. I'm going to have low frequency harmonic content in my output.

So the question is, is there a way you can extend the range over which you can generate, if you will, pure sine waves at the output of your motor and generate like no low frequency, line frequency harmonic content? And the answer is there's a trick for doing this. It can be done in a variety of different ways, but I just thought I'd mention it because in the motors world, this is a very common trick. And the idea is this.

Suppose if I thought-- if I thought of what the duty ratio-- if I thought of this switch q1-- if I thought of this as q1, the average value of q1 over a cycle as being the duty ratio of this switch, so modulation index of 1, the duty ratio of the top switch goes to 1, with what I've told you so far.

So I could talk about the duty ratio of this half bridge, just like it was a buck converter, where I'm talking about on state fraction of the top switch. The way this would look if I did my sine triangle PWM would be d1 of t would be equal to 1/2 plus m over 2 sine of omega line t.

So if I did my sine triangle PWM, if m was one, duty ratio goes to 1 when the sine wave is peak, and it goes to 0 when sine wave is at minimum. Does that make sense to everybody? And then if I was looking at d3, it would be this same three-- actually, it's not d3. I forget what the middle switch-- top switch-- it is d3. d3 would be equal to 1/2 plus m over 2 sine omega lt minus 2 pi over 3, for example. Does that make sense to everybody?

And d5, which would be the third bridge leg, would have 1/2 plus m over 2 sine omega lt plus 2 pi over 3. So this is what would result with this triangle intercept PWM done on the three phases.

And what we said was basically the average voltage is proportional-- the average voltage at that node with respect to ground is proportional to the duty ratio. With respect to the midpoint, there's this one half of the bus voltage in there.

The question would be I have a trouble that I'm limited to making m less than 1. Because when m exactly equals one, the peak of this drives the duty ratio of one just touches the top. If I try to make m 1.5, what would happen is for some part of the cycle, duty ratio basically saturates at 1, the top switches on, and I get a bunch of distortion. Is that clear to everybody?

So here's the trick people use. It's kind of a cool trick. There's different ways to implement this trick. So this isn't the way you have to do it. But suppose I did this. I could take and add m over 12 sine of 3 omega line t to each of these waveforms.

Why would I do that? Well, because when I think about sine, if I looked at this duty ratio, it would look like this. If I just took-- here, I would have a half, and I'm doing something like this. What I can do is I can add in a third harmonic component. What does a third harmonic component do? It does something like this.

So the sum of this piece, which is this, and this piece, which is this, has some distorted shape which does something like this. And what that means is I can turn up m just bigger than 1, and it turns out to be about 1.15, before basically the peak will hit VDC over 2 or before my duty ratio will saturate at 1.

And so that gives me an extra 15% of voltage I can synthesize. And the way I'm synthesizing it is I'm adding third harmonic distortion in. So I'm adding some third harmonic to VAR, to VBR, and VCR. But why doesn't that matter? Because the third harmonic, that third harmonic component that I drew in yellow, just appears at the neutral. And the line-to-line voltages don't have it.

So that's a technique that's sometimes called third-harmonic injection PWM. You can do it by adding third harmonic onto your-- instead of doing this, you add some third harmonic onto that thing, and then you get an ability to expand your duty ratio range or expand the fundamental voltage you can synthesize without distortion. Any questions about that? Yeah, Jack.

AUDIENCE: Can you explain again how it's-- like VDC is fixed. So how is this increasing-- overcoming back emf?

DAVID PERREAULT: Right. Well, let's assume the back emf is at the fundamental, right? So what I have to worry about is how much fundamental can I create here? The vacuum is at the fundamental. How much fundamental can I create here? And the problem with not doing this third-harmonic distortion trick is that if I try to turn up the duty ratio more, I get a lot of low frequency harmonics. I get third and fifth, other content in here, which is then going to drive low frequency content in the machine.

This trick, what it lets me do is I basically am going to create the fundamental, plus I'm going to create some third harmonic here, third harmonic here, and third harmonic here. But the third harmonic elements all cancel. They just show up at the neutral, and the neutral goes up and down by the third harmonic.

So I can keep this component pure because I'm not saturating my comparison anymore. I'm not making this signal go out of the range just for that little bit, for that extra roughly 15% of modulation index over 1. So I can go up to m of 1.15. Does that make sense?

Now, that gives me a little bit more head room to synthesize a little bit more voltage. And we're often constrained by how much voltage we can synthesize with our inverter. What do I do if I'm not happy about that? Well, I could just keep turning up m beyond 1. What happens if I turn up m beyond 1?

Well, if I turn up m beyond 1, and my reference waveform was going to swing up here like that, what I would do is I would get a bunch of PWM switching until the middle of the cycle or something where I just flat top. And so my PWM waveform for each phase might look like this. It's PWMing. And then it'll just sit at the top. And it'll start PWMing again.

And if I keep pushing it up and I keep making my reference waveform bigger and bigger and bigger, the limit of that is m goes to infinity. It's just that I'm just going to generate a square wave. He's going to be on for half the cycle. He's going to be on for half the cycle. And you're back to the six-step waveform that I talked about before.

And in that case, then I really do generate the maximum fundamental. The maximum fundamental is basically VAR is a square wave. It's VDC over 2 times 4 over pi is the fundamental of a square wave, of amplitude plus and minus VDC over 2. That's the biggest fundamental I can synthesize. And that sets an upper bound of how the fundamental voltage I can drive here.

The only problem I get with doing that is that waveform has third harmonic, which will cancel, but it also has 5th and 7th and 11th and 13th. And I will get harmonic currents flowing in these phase windings when those harmonic currents probably aren't helping me drive the machine. So I can synthesize something that gives me nice, effectively-- at least on a local average basis, average over switching cycle-- sine waves up until the amplitude equals VDC over 2.

If I use my third harmonic injection trick, I can get it up to VDC over 2 times 1.15. And then if I don't care about distortion, I'm going to let a bunch of low frequency harmonics exist. I can just drive up the modulation index even further and get up to 4 over pi VDC over 2. And people do that.

In fact, in the Prius, the Toyota guys eventually published a paper saying, yeah, we ran out of headroom and we wanted to run our motor faster for higher speed, so we just let it saturate and generate the harmonics. And the motor still worked, and we were OK. But you pay a price for that. So people do do this. But if you want your waveforms to be pure and not have low frequency harmonics and generate that noise in your load, whatever it is, with that low frequency content in your load, you've got to keep your modulation index down. Any questions? Yeah.

AUDIENCE: Is it possible to step up VDC with a [INAUDIBLE] before?

DAVID PERREAULT: Absolutely. And in fact, that's what people do. So why is this a concern? Because if you're in a Prius, you've got this bus voltage. And that bus voltage is your battery voltage. And that might vary over some range. And then you're trying to deal with this motor whose back emf varies over some range and you need to drive it. If you don't like that, then you put another stage in between your battery and your converter, and you can control the DC bus voltage to go up higher.

And then you're only limited by how high voltage these switches are. So you're always limited by that. But yes, people do actually control this bus voltage from the battery voltage, for example, in an electric vehicle to be something higher. And they might control it dynamically if they had to, depending upon the operating condition they want. All right. Well, have a great day. We will pick up a new topic next class.