[SQUEAKING]

[RUSTLING]

[CLICKING]

DAVID PERREAULT: OK, why don't we get started? So what I wanted to do in the next few lectures is talk about filtering. And this is ripple filtering or EMI, Electromagnetic Interference filtering. And why am I going to spend so much time talking about filters? And the reason is because in most power supplies, and in fact in a lot of power electronics circuits generally, the filters take up a large fraction of the overall volume and cost.

So if you look at a typical offline power supply, 25% of the volume might be devoted to just filters because we're switching, we're transferring energy. And the ripple we generate, we really need to suppress. And we want to suppress it maybe to meet some application desire, but we might want to suppress it also because we're not allowed to sell something that you plug into the wall that has too much ripple.

So if we think about it, there's kind of two kinds or two flavors of specifications we often have to meet. One of them is in the time domain. If I'm developing a power supply, somebody might tell me, oh, look, I'm only going to allow you to have 20 millivolts, either peak or RMS, of output ripple because we want such and such a level of noise in our output. And that's one kind of specification you might need to meet.

More typically, however, and almost universally, if it's something like you're talking about an input filter, you're going to plug it into the wall, is in the frequency domain. That is, how much energy can you inject in a given frequency bandwidth?

So a typical thing was you might detect this with a 9-kilohertz resolution bandwidth for where you're picking up energy and look in the span, say, between for conducted emissions, look in the span between, like, 150 kilohertz and 30 megahertz is most commercial standards for conducted EMI. Above 30 megahertz, you typically use antennas and stuff to pick it up. And we're not going to talk about radiated electromagnetic interference. But in fact, if you don't pass conducted interference, you won't pass radiated either because basically when you conduct energy into your input leads, your output leads, that's an antenna, and it broadcasts off noise.

So it's important to meet these things if you're going to sell commercial products that plug into a wall or into a vehicle. Into the wall, you've got to deal with the FCC. If you're going to plug it into the vehicle, you've got to deal with the vehicle manufacturer. And you're going to have to pass these specifications in order to sell something. And it's quite important. Let me just give you an example.

Once upon a time, I consulted for-- I'd call it a garage startup, but really, they were in the guy's basement, so it was a basement startup. And what they were developing-- this was early days. They were developing a solar inverter to sell. And so they were developing the power stage to convert energy into the grid and getting the efficiency up and doing all these things. But they hadn't gotten around to designing the EMI filter yet.

So the guy's neighbor came up to him and said, are you having trouble? Around every day around 3 o'clock, I can't watch television because, I don't know, everything just fuzzes out on me. And what the guy in the startup didn't tell him was 3 o'clock was when he started testing every day. He was dumping too much ripple into the grid and the next house over suddenly couldn't watch television. So these are real problems. It's not just you got to pass a spec. It's you will interfere with things.

If you think about it in a car, if you're plugging something into the USB outlet or into the cigarette lighter or something, any ripple you conduct is running around on the wires in the car. And it's right next to a radio. So your radio is a very sensitive detector of very, very small quantities of noise. And you will mess with that if you do not have very, very good filtering.

So what we're going to find is that you've got to do filtering, but it's not just, oh, I need to make my ripple a little bit small. You've got to make it really small in radio frequency bands in order to not interfere with things. So we're going to talk about conducted emissions. And what I'm going to mostly focus-- most of what I'm going to say about input filters-- like, if I'm plugging my power converter in the wall, it's between the power stage and my power converter in the wall. So that's the input side filter. Most of it applies to output side filters too. But we'll focus on input filters.

So let's just think, to get a sense of the challenges or issues, I want to start by talking about, how do you measure these things and what drives that? So let's just think about a very simple buck converter. So suppose I go off and build a buck converter. So here's my buck converter. And I'm not going to worry about the output filter. Let's just assume he's connected to a resistor.

But here, I have the input filter. And maybe I can consider the capacitor for the buck converter the first part of the input filter. And here I have I'm going to draw some current ix. I'm going to put some input filter here. But for the moment, let's just assume my capacitor is the input filter I've got. So this thing right here is right now my filter. And now I'm going to connect this up to some source.

And maybe this source is the battery in my car as vs. And then there's some source impedance, zs. And that may be-- at DC, that might be very small. But at AC, maybe it's not. So there's some source impedance that I'm driving in. And here's the challenge. What would be a model for this as far as conducting noise into the input?

And maybe I would consider my noise to be whatever ripple here is-- I'll call this iy. So if I'm driving ripple current back into this input, then I'm generating noise in the input. And maybe that's what's interfering with my radio. So how would I model this? Maybe the easiest thing for me to do is say let me pretend this is a constant current inductor. I'm not going to worry about the inductor ripple. And a simplified model for this would be something like this.

I'd have some current ix, which is some square wave. And then maybe I would only care about ix tilde, the AC component of that. So maybe I'll just worry about the AC component of this current. It's coming into the filter capacitor and into the source impedance.

And maybe I will worry about iy. Or maybe I will worry about the voltage here, vx. I could specify either thing to worry about the filter. Well, let's just think about the effect of the source impedance on my measurements.

If zs goes to 0, if I had a really good source, then what would happen? If zs goes to 0, iy goes to ix. If there's no source impedance, this capacitor does nothing, and all the ripple goes back into my input. Of course, if zs goes to 0, then v sub x tilde goes to 0.

So if I had no source impedance, I'd get-- all of my input would see the ripple current from the converter, but I wouldn't have any voltage ripple at the input of my converter. On the other hand, if z sub s at ripple frequencies went to infinity, what would happen in that case? Well, then I'd have a really high impedance here. The ripple current that I'd measure would go to 0 because this is just a very high source impedance, and I wouldn't-- all the current would go through the capacitor. And v sub x would go to-- it would go to i sub x divided by j omega c.

So my point is this. Depending upon what I care about, whether I'm measuring voltage here or current here, the source impedance affects both of those measurements. And that's a problem because here's the way EMI works. So if I want to go-- if I want to go sell a power supply, a commercial power supply that I'm going to plug into the wall, I have to go test it and make it meet standards.

And then I don't do anything. I just take that result and I put it in a desk drawer. And as long as nobody complains, great. But if somebody comes back and says, you're emitting too much EMI, then I've got to produce this test result, and we've got to come to some agreement. Or they've got to be-- they've got to be able to ascertain, did my power supply pass the test or not?

And what I'm telling you here is that if you try to just plug a power supply into a source, the result you get depends very, very heavily on this source impedance. So if I'm designing a power converter for a car, that source impedance depends upon what model car, what week it is, what else they have plugged into the bus. It depends on all kinds of stuff that I don't control.

So if I'm going to sell a power supply to plug into the BMW or into some other brand car, they have no way of coming to an agreement with me as to whether my power supply passes because on a different model or with different stuff plugged in, you might get a completely different result. Does that make sense to everybody?

So what do you do about this? And this is what underpins EMI test procedures. And I will say that for any procedure you want to test, like SAE has a set of specifications that they recommend to test procedures. If you want to qualify for something into the wall, FCC Title 17 has a whole set of procedures that you follow.

And these procedures are designed to deal with this issue, that you've got to frame a test that if I test it this week and I test it next week and I test it the week after that in some other place, I'm going to get more or less the same result. So it's not that it necessarily reflects the real environment, but it reflects a standardized environment that I can rely on.

So here's how they do it. What we would really do is the following. Or actually, I'm going to give you a slightly simplified version of what you would really do. Suppose here's my load. And I'm talking about an input filter because I'm going to plug this thing into the wall. So here's my converter. And it's generating some ripple current ix. And here's whatever my filter is. And here I'm just going to draw my filter as a capacitor.

What I'm going to do is I'm not going to plug this into my power source, which could be my vehicle battery. But instead, I'm going to go get another box and connect this up to a box that looks like this. It has a series inductance, L LISN, and a shunt capacitance, C LISN. Sometimes it has another capacitor here to ground. And then this is going to have some output. And then I'm going to connect this into the power source.

So here we go. Here's zs. And here is vs. So this is whatever my usual source is. If this is my vehicle battery, that's what I'm going to plug it into. And this element here comes in a box. It's a standardized element. And then I'm going to go plug this output of this thing into a spectrum analyzer.

And this is going to be-- the input of the spectrum analyzer has a 50-ohm input. So basically, my spectrum analyzer looks like a 50-ohm input. And the idea of this box is called a Line Impedance Stabilization Network, or LISN-- Line Impedance Stabilization Network.

And its job is to present a consistent measurement point for EMI. So what does this thing do? It's its own kind of filter, but the idea is all the DC current I want to provide to my converter from my source will just pass through this inductor. This will have some cutoff frequency, and L LISN and C LISN in this capacitor, depending upon the LISN, will be specified.

And below some frequency, all the current is just going to come from the input, and it will draw power just as normal. When I get to really, really high frequencies, tens of megahertz or whatever, the idea is that this path will look like an open circuit. C LISN will look like a short circuit. And I will be-- when I look back into here, if I look back this way, at radio frequencies, what I'm going to see is 50 ohms.

And so whatever noise I conduct from my switching is going to be looking back into 50 ohms. And if we plug this test setup into here, where the zs is one thing, and somebody else plugs into it and zs is another thing, you're still going to get pretty close to the same results. So this box, this LISN, is only for testing. It's not something that goes with your converter. It's just part of the test stand to make measurements. And I give you examples of values. What are the values of those things? It depends.

Are you building something for a car where they have one set of standards? Are you building it for FCC, where you plug into a wall, they have another set of standards? But all that is documented, and you test according to that standard. Any questions about this so far?

AUDIENCE: So that capacitor over there, on the left side of the list of [INAUDIBLE].

DAVID PERREAULT: It depends on the LISN. So it depends on the standard. Some will just have an inductor and a capacitor. Others will have another capacitor on the other side. That's a function of-- what the LISN looks like exactly inside depends on the standard. Other questions?

So presenting a non-LISN impedance lets me make repeatable measurements on my converter. And if I measure it this week, I get the same result as next week. Usually the way you would do this-- I should say, by the way, there's a lot of-- if you look up to get very, very repeatable measurements, they have all kinds of standards. Your converter has to be one meter or some specified height over a ground plane. It's like they make all the conditions exactly identical, which aren't real-world conditions, but they're just designed to make the test repeatable. And you follow that.

So very often, if you're trying to qualify a power supply, you'll ultimately go to a test house that that's what they do. What you do in the lab is you can set up an approximation to that usually. I mean, you can set up a real test stand if you're a company that's making lots of parts. But at home, you can go buy a LISN and do an approximation to that. And it'll be pretty good. It's good enough for designing a power supply if you put a bunch of margin in your answers.

And the way this thing-- the way the specifications are read is, for conducting EMI, you'll look at this response, spectrum analyzer configured a certain way with certain kinds quasi peak or peak detectors. There's a whole bunch of rules around it. And you're generally measuring the energy in different frequency bands and seeing, do they go above an allowed limit? And those allowed limits are really small.

So if I thought I had-- this was tens of amps, here I might be required to get tens of microamps. So it's really small currents that you've got to attenuate down to. So when I show you, oh, we've got a capacitor here, the model for this would then be basically a capacitor in parallel with 50 ohms, and I'm trying to measure that voltage ripple on the 50 ohms, you're going to need more than a capacitor to make this work. So very often, you're going to need a multi-stage filter.

I will say, one other thing, by the way, I've drawn this as if you have a LISN that you stick in between your power source and your converter, and you do, but typically, most specifications require that you put a LISN on this lead and a separate LISN on this lead. So there's actually two LISNs involved in a DC system test. We'll come back to why you would do that later. But the basic concept is not changed. It's the same idea.

Now, let's just think of what we would need. I said, OK, a capacitor is not really going to cut it for this measurement. So let me think about my converter as some current ix or ix tilde. This is the ripple from my converter.

If a converter-- if a capacitor is not going to be enough, what would be my basic building block? My basic building block might look something like this. It might be some filter capacitor Cf. And the first stage could be just my input capacitor, my converter, which I would always have and some inductor Lf. Now, this is not the LISN inductor. This is my filter inductor. So this is my filter.

And then eventually I look back into the LISN impedance. And this would be z LISN looking here. And at ripple frequencies, maybe that looks like 50 ohms. So a basic building block is an Lc filter. What is the transfer function-- if I said this was ix tilde and this was iy tilde, what does the transfer function of this thing look like?

The magnitude of iy tilde over ix tilde versus omega. And this is going to be a log scale, and this is going to be a log scale. At DC, all the current goes back. So at DC, I'm going to get this at really low frequencies. Then I'm going to have some region that's 1 over the square root of LfCf, which is going to be the cutoff of this filter.

And at really high frequencies, this is going to fall off as 1 over omega squared, or minus 40 db per decade. So the idea here is I can go get a bunch of attenuation. I'm going to put this cutoff of this filter well below my 150 kilohertz that I care about attenuating and live off of this attenuation out here.

Now, why is that going as 1 over omega squared? Just as a reminder, the impedance of the capacitor, the magnitude of zc, goes as 1 over omega C. So this is getting to be a smaller impedance as frequency goes up. And z sub L is going as omega L. So basically, you get an omega squared factor because he's getting-- it's a current divider, and he's getting smaller inversely proportional to omega, and he's getting bigger proportional to omega. So you get a kind of an omega squared attenuation.

So the way you might want to think about this is your filter's working because of the mismatch between the shunt path and the series path. And if I don't like this, I can build higher-order filters and so forth. And we'll talk about that. I haven't said what happens in the vicinity of cutoff. That's an important thing. And we're going to have to think about that. But let's for the moment not worry about-- not worry about what's going on here. We'll come back to that point.

So any questions up to now? So I want to build this filter. And as I said, I need a lot of attenuation. If I'm going from 100 amps to 100 microamps, that's six orders of attenuation. So I'm going to probably need a pretty low cutoff filter here. I'm going to need a bunch of attenuation.

And one of the things that people find about designing EMI filters, your average power electronics guy hates it. Why do they hate it? They hate it because you can very get-- you can very often get very kind of seemingly anomalous results. Like, you're doing stuff, and you're just getting noise, and you don't want this noise. And it becomes a problem. It's hard to get the good result that a simple picture like this would suggest that you can get.

So let me tell you-- just give you a heads up about some of the things you got to pay attention to when you're doing this. And the first of these things is component parasitics. So if I were to look at this filter, and I'd say, jeez, what does this look like?

Well, I said my capacitor-- if I looked at my capacitor, and I said, what is the magnitude of z sub C? I said that the magnitude of z sub C should fall off as 1 over omega C, right? So this is the magnitude of z sub C. And if I had an ideal capacitor, it would do that forever.

But what about a real capacitor? Well, a real capacitor, even though I think of it as an energy storage element, the plates do have some resistance. There's actually some resistance of the wiring going into that thing. So as a practical matter, a real capacitor is going to have some ESR, series resistance.

So what's that going to do? Well, there's going to be some frequency above which this ESR has a bigger impedance than the capacitor, the ideal portion of the capacitor. So what I would expect is beyond some frequency, there's going to be some place where our R ESR is dominant. So instead of falling off forever as 1 over omega, it flattens out eventually.

And then if I keep going up in frequency, I'm going to get to some point where I have some equivalent series inductance, L ESL. Why does that exist? Because there is space between the capacitors, and if I have current through it, there's some magnetic field in there. And there's just a little bit of inductance. It doesn't have to be very much, but we're talking about going to very high frequencies.

So at some point, this ESL will even get to have a high-- because this is growing with omega, it's going to go up. So my overall contour of my impedance doesn't fall forever, but rather, it sort of does this kind of falls, flattens out, and goes up.

And maybe this-- I've shown you a high ESR capacitor. Depending upon whether the ESR is greater than or less than the square root of L over C ESL-- L ESL over C, this will be very sharply tuned or very broad. But either way, you're going to get this kind of behavior. Why am I harping on this? Because if I thought of this as my real element, now I plug this thing in place of my capacitor. What would I expect my attenuation to look like in that case?

Well, I said it was based on the mismatch between this impedance going down and this impedance going up. So eventually, instead of turning into an LC filter, it turns into an LR filter. So instead of falling off as 1 over omega squared, it's going to fall off as 1 over omega. And when I get to really high frequencies, ESL is going to dominate, and then I'm just going to have a-- I'm just going to have a current divider between L ESL and L filter.

And in fact, what that's going to mean is my attenuation is just going to flatten out. So if I design my filter thinking, oh, jeez, that's how I get my six orders of magnitude, and I come back and I'm really dealing with something where at that frequency, this capacitor doesn't look like a capacitor, it looks like a little inductor, my EMI might be way higher than I thought it was going to be. My filtering is just going to get clobbered by the parasitic inductance of my capacitor.

Now, I will say that, by the way, you know this-- and by the way, where does this kind of phenomena happen? Well, it depends on the size of your capacitor, the physical size of your capacitor. It depends on whether it's a good capacitor or not, et cetera. But very often, like 10 megahertz kind of frequencies is where these things are happening, as an example, or even 1 megahertz-- depends on the size of the component. And that's right in the band that you're caring about and trying to attenuate.

So these bad parasitics of the components-- and usually the capacitor is the first thing to give out on you if you design the inductors right-- will cause you to get a lot less attenuation out of this LC filter than you thought. Any questions about that?

So that's kind of a problem. What would I do about it? What I might do about it is not just have one capacitor but have a bunch of capacitors, some of which might be good at-- they're not a very high C value, but they don't have as much inductance. And so they look better at high frequencies than bigger capacitors that have better performance at low frequencies but not as good as high frequencies.

So I can certainly do it by being careful about how I pick this component or picking multiple components to cover different frequency ranges. And if I have multiple components, by the way, I can get resonances between them. But that's another story. What else can I do? I might say, well, OK, if I'm limited by ESL, I will use a bigger Lf, but you're still only getting a divider that's a proportion between these two, the ratio of those two.

If I want really more attenuation, what I can do is do a higher-stage filter. So here's ix tilde. Maybe I build a first stage-- I'll call this Cf1 and Lf1. And I'm leaving out damping and other things I'm going to talk about in the future. And then maybe I have Cf2 and Lf2 and so forth.

And then maybe I have my-- I'm looking back into R LISN back here. So maybe I can get 40 db per decade out of this and 40 db per decade out of this. And maybe I do better. And it may be better to do that even in a lot of cases because even putting aside the parasitics, I can move the cutoff frequency down 1 over the square root of LC. By using bigger and bigger components, I can move that cutoff filter down.

Here, maybe I use smaller components, but instead of getting a 40 db per decade roll-off, I could theoretically get 80 out of this. So even if the net cutoff's higher, that fast attenuation will help me. So it's often better to go for smaller multi-stage filters than one big fat single filter stage trying to get the attenuation I want.

And by the way, when you do this, for reasons I'm not really going to go into the details of this because we don't have time-- when you design these kind of things, it might be your inclination to say, oh, let me use the same LfCf cutoff here and here and put two filters with the same cutoff. It turns out to be better to stagger these things. So usually, you'd make this one a low cutoff and then this one a higher cutoff to deal with the higher frequency and so forth. It turns out to be a better kind of design strategy for these kind of filters.

So you can cascade filter sections until you get the attenuation you need. And if you look in a typical offline power supply that you'd plug into the wall, they usually have a couple filter stages like this. Any questions about that?

I should warn you too-- I said, oh, people hate these-- they hate dealing with the EMI because you go to measure stuff and then what you measure doesn't seem to make sense. And it all does make sense, but very small phenomena can mess you up. So think about this. This filter and it's cutoff regime, putting aside these parasitics I was talking about, could give you 80 db per decade. You can get really like, wham, lots of attenuation. And that's good because you need a lot of attenuation.

But what's the downside? Well, one thing is the component parasitics themselves. And I focused on the capacitor. The inductors do have-- if you put a capacitance, inductors do have a parasitic parallel capacitance. You also have to worry about that one. But putting aside that parasitic, you also have layout parasitics, meaning if I do a crummy job and I don't stick the components into the board all the way or I put two components next to each other, they can have capacitance between them.

So I can accidentally insert inductance, or I can accidentally insert capacitance. And one of the things that can happen is imagine this is my power supply input leads, and over here is where it's going to go off to-- I'm showing the LISN here, but this is where it goes off to the power supply input. What if I put those leads near each other? Well, then maybe what I would end up with is some small parasitic capacitance from this node past my filter to this node.

And so what happens? This filter is doing some great job, massively killing the noise. But this, because this node has some switching ripple at it, it's being injected through this. Maybe it's a picofarad of capacitance. It doesn't even have to be very much if you have a high attenuation filter getting over to here and completely bypassing your filter. So you're getting bad results, you're measuring stuff over here, and you make this filter better and nothing happens. You make this filter better and nothing happens. Why? Because it's not going through your filter. It's going around your filter.

And that's the kind of thing that drives people nuts. So you have to think very carefully about your layout. People really try to have a quiet side and a noisy side so that they stay physically far apart so you don't get capacitive bridging. You can also have magnetic bridging. If I have this-- if this loop with a lot of switching ripple has mutual inductance to this loop, even if there's no capacitance, I can get magnetically coupled right across through my layout and completely ruin my filter performance.

We had some-- there's a power supply I worked on with a friend for commercial application. And we designed it, and it didn't quite make the-- couldn't pass the spec. We weren't passing the spec. We went and we moved on the board a trace about 3 or 4 millimeters. Suddenly we were passing it fine. It was exactly that phenomenon.

There was a noisy port that was too close to the input. And we were capacitively coupling the energy into the input. And all we had to do was move the trace. And that trace, that parasitic isn't showing up in your schematic. It's not showing up anywhere, but it's there in the real world. And you have to go find it. We moved that one plane, and we were fine. So you've got to-- you've got to spend a lot of time thinking about this.

What they do in commercial EMI filters is they often will have a filter stage like this and put this inside a shielded box and then have some feedthroughs going to the next filter stage and so forth, and they can build commercial EMI filters that way. And that's great. The sad news about being a power electronics designer is usually they won't let you do that. Nobody's going to let you have shielded boxes and so forth in your converter. So you've just got to be very clever about how you lay things out and keep things from coupling magnetically and keep things from coupling capacitively. And that's often an iterative process to get it to the point where you've realized what's going to happen with all the parasitics.

So that's a way to, first of all, warn you what's to come if you go do this and to not get frustrated. Just look for where nonidealities, either layout parasitics or component parasitics, are going to come bite you, and then fix those. You've also got to worry about the measurements being right too.

On one project I did with a colleague, we took our prototype to a test house, one of these commercial facilities, tested it, qualified the power supply. Yes, you can plug this thing into the wall. Went on our merry way. Company calls us up. And this was a power supply that went into a water purity meter that then went on to other semiconductor equipment. And get a call from the people we designed the power supply for.

Yeah, our customers came back to us and said that their equipment was failing EMI tests, and they tracked their failure of the EMI test back to this one box that we make. And of course, we had designed the EMI filter for it. What's going on? We've got to fix this problem because their customer's unhappy because they're not passing FCC. But of course, we're puzzled because we think we've passed the test.

Come to track back, what happened is we took it to a commercial test facility, and they had been doing work on their test setup. And their test setup had gotten messed up and wasn't grounded properly. So basically, the results they were showing from their test setup said it was passing, and it wasn't. We took it to a different facility. It was failing. And that's when we went back and changed the design, improved it, and suddenly it passed everything.

So there's a lot of layout challenges. And there's a lot of measurement challenges. And you just got to be willing to dig through them. And eventually, you'll get something. But when you're dealing with power supplies, you don't try to design something that's going to pass by a few db the specifications, like by a small margin. You want to give yourself-- you want to you want to really pass it. Because the reason is, even with all that craziness to try to keep everything identical, you're going to have variations in your design in one test setup or another or component variations in the manufacturing that aren't quite the same.

So you need to give yourself a bunch of margin around the specifications. And if we had had more margin in the first design, they probably wouldn't have even come back, even though the test results were-- it wasn't our fault that the test system was wrong. Actually, we got nicely paid to debug the EMI facility's problems for them, but that was another story.

So just to illustrate these phenomena-- and so what are we going to talk about? We're going to talk about a bit more about filter design and a bunch of other phenomena that come up. But I just wanted to illustrate this whole question about, first of all, component parasitics, which are very real even if you do your best layout. But also, if you do a bad job, if you accidentally put a capacitance here or you accidentally-- because you don't place this capacitor very well, you have extra lead length going to it and you insert some inductance, you'll make everything worse.

So you want to be really careful about your circuit layout. And we have a demo of this. And what this is, essentially, is a set of four filters here. These are just LC filters. And I'm going to let Rafa demonstrate. And what we have is a network analyzer. What a network analyzer does is it drives at one port, like our signal ix, and then it measures into another port, which would be like our signal at the LISN. And let's see the behavior of the same filter-- it has the same physical component, just placed slightly differently-- on the performance of the filter. And with that, I'll hand it over to Rafa and let her explain.

GUEST SPEAKER: Do I have to use the microphone?

DAVID PERREAULT: I think so.

GUEST SPEAKER: OK. All right, so, like Dave mentioned, we have four different filters with four different layout strategies. In some of them, we are being careful about both the parasitic inductance and capacitance. In some of them, only one.

DAVID PERREAULT: So yeah, I'll let Rafa-- I'll let Rafa do the testing while I do the explaining. So we have four filters. The one she's going to start with, it's an LC filter just like this. And all that happened is we were sloppy about, both, how we inserted the capacitor. So there's a little extra inductance in here. And the inductor, we place traces nearby. So there's a bit of extra capacitance across it. So that's the first trace we're going to see. And let's see what this does.

GUEST SPEAKER: So this is the filter where we don't-- where we are not careful about the capacitance and the inductance.

DAVID PERREAULT: So this is bad-- this is parasitic capacitance and parasitic inductance. And so this is the attenuation of the filter. And we can see that the filter attenuates down-- and by the way, we're going from 1 kilohertz to 20 megahertz here. You can see even before we get very far out, the filter ceases to attenuating and then it gets bad. And these were the good filter components, but we're just getting bad results.

The next one we're going to try-- what we did was we said, OK, we're going to have OK layout of the capacitor, so no parasitic inductance or as little as we can get from the component itself, but we didn't fix the bad layout across the inductor. And so here we go. And what you can see is, yeah, this is a little bit better because of what we got, but it's still terrible.

The next one, we have a bad layout with parasitic inductance in series with the capacitor, but we fixed the problem with the inductor. So we're fixing one at a time. And I should say, by the way, that this test, notice we don't have a ground plane set up. We don't have all this other stuff. So this test sub wouldn't qualify for getting repeatable results in all kinds of systems. But here we go. We have again-- I didn't fix both problems, and hence I still have bad performance.

So this is the frustrating thing about these things, that, yeah, you fix one thing but not the other, and you get the same kind of bad result. Now, let's do the final case. And this final case is we've taken the same components, but we've tried to do a good job of laying out the filter. So this is our attempt to-- we're still going to have component parasitics, the internal resistance and inductance of the components. But we're going to try to do a slightly better job about how we configure them on the board.

So here we go. What do we get now? Well, OK, there we go. Boom. Now, you say, well, is that very good? Did we change it very much? But notice that this log magnitude scale is 20 db per division. That means the noise of the blue trace is about a factor of 100 below the noise at the white trace. So that's a really big difference.

And for the others, it's maybe a factor of 50 or something like that. So it's a substantial amount below. And this is just one example of just one filter stage, whether you've done a good job or a bad job with the parasitic elements. So when you go out to do this stuff-- this is the whole thing I want you to take away from this-- really think about all the second order effects, the parasitics of the components, and how you use them.

And we try really hard to do things like keep the noisy input-- or the noisy converter away from the quiet input we're going to plug into the wall. Some days you just don't get away from it because they tell you, no, you can't change where our connector is. And just because the connector is right next to the noisy power supply, go figure it out. And you have to live with it.

But if you can, try to keep your noisy stuff far away from your quiet stuff. Try to reduce the parasitics by being very careful about how you lay out your components-- low inductance insertion, low capacitance across the components, and that will go a long way towards getting around weird, anomalous behavior and getting to be as good as you can.

And then even when you do that and you're getting some worse performance, getting worse around 20 megahertz, fine. Now we'll cascade another filter stage to handle the higher frequencies. So that's about it for today. Are there any questions before we wrap up this first introduction to EMI measurements? OK, we will see you next week.