The Quickest Path To Power Integrity

by Steve Sandler, Picotest, Phoenix, Ariz.

Two things became apparent to me during our recent series of 2017 Hands-On Power Integrity Workshops. First, the majority of attendees weren't power supply designers, or even analog engineers. They were high-speed circuit designers and printed circuit board designers. In other words, they were power *consumers.*

The second surprise was that there was much less interest in *why* the power integrity (PI) performance was what it was than how to correct it. For those engineers who want to know what steps they can take to prevent or quickly mitigate PI-related problems in their system designs, I promise you that this article will provide the quickest and easiest path to power integrity.

What Are The Goals?

While there are many aspects to achieving power integrity, most engineers are struggling with the impedance aspects of the power distribution network (PDN) as impedance controls the voltage range of the power to the high-speed chips. As the high-speed devices switch, they consume large pulses of current with fast edges. The role of the voltage regulator module (VRM) (or any on-board voltage regulator) and the circuit board decoupling is to manage the voltage transients that result from these high-speed current demands.

This might lead you to believe that power integrity is about achieving the lowest possible impedance at the high-speed devices, but that's far from true. The first goal is to achieve a resonant-free impedance, implying that the ideal impedance is flat or resistive in nature when viewed in the frequency domain.[1]

It can easily be shown that the general nature of the VRM control loop results in an inductive output[2] and the addition of output capacitors serves to *absorb* the excess inductance. However, the output capacitors themselves also include some series inductance (ESL). Next, the circuit board planes are transmission lines, which are simultaneously inductive and capacitive. We can illustrate this easily using a 50-Ω coaxial cable as shown in Fig. 1.

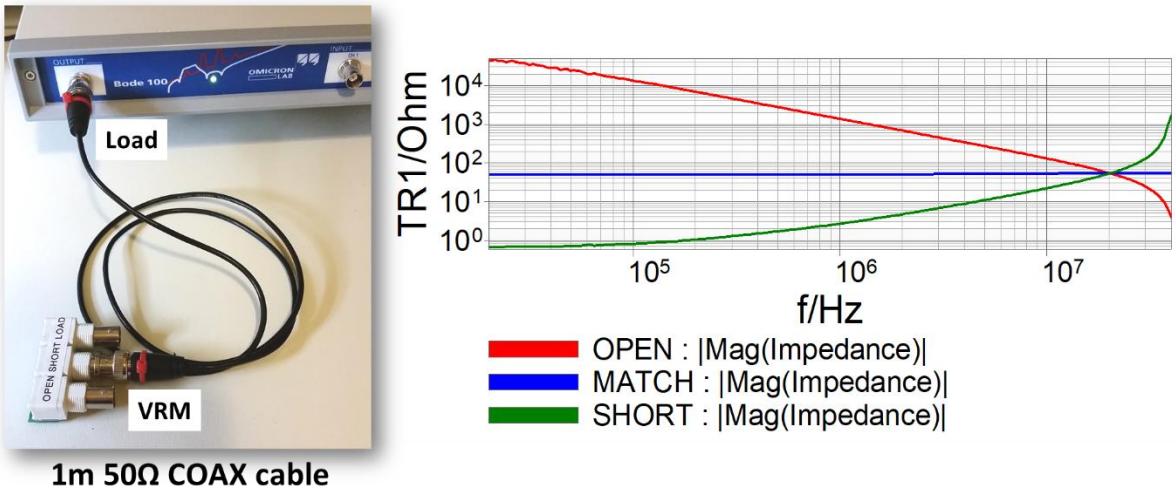

Fig. 1. A 1-meter coaxial cable is used to illustrate the inductive and capacitive nature of a transmission line.

In this measurement, a vector network analyzer (VNA) is connected to one end of the cable. We'll designate this end of the cable as the load end and the other end of the cable is where the VRM is connected. Here, we use a set of VNA calibration terminations to represent the VRM as shown in Fig. 1. If the VRM impedance matches the cable impedance the result is a perfectly resonant-free 50 Ω as shown by the blue trace. The capacitance of the cable exactly absorbs the inductance of the cable.

If the VRM impedance is less than 50 Ω the cable appears to be inductive, evident from the impedance magnitude increasing with frequency. This is demonstrated by the measurement of the shorted termination (green trace). In a similar way, if the VRM impedance is greater than 50 Ω the cable appears capacitive, evident

from the impedance magnitude decreasing with frequency, as happens when the open circuit termination is used (red trace).

The required VRM impedance will almost always be less than the characteristic impedance of the circuit board planes. As was the case with the shorted termination in our simple example above, this results in an inductive plane. Decoupling capacitors are then used to *absorb* the extra inductance, providing a matched impedance between the VRM and the printed circuit board.

The transmission line impedance can be defined by the inductance and capacitance of the line

$$Z_o = \sqrt{\frac{L}{C}} \quad (1)$$

And matching this means that we set the source and load resistance equal to Z_o

$$Z_o = \sqrt{\frac{L}{C}} = R \quad (2)$$

And since the VRM and the planes both tend to be inductive, we need to determine the capacitance required to absorb the inductance. Solving this equation for C leads results in:

$$C = \frac{L}{R^2} \quad (3)$$

And the ESR of the capacitor should be equal to R and the ESL of the capacitor should be much lower than L.

Equation 3 reveals that the required capacitance increases inversely with resistance *squared*. So, minimum capacitance is achieved by setting the resistance *as high as is tolerable* and not to as low as possible. The capacitance also varies proportionately with inductance, so secondarily, minimum inductance is desirable.

The target impedance establishes the value of R. Usually the target impedance is specified by the vendor of the high-speed device being powered, or its value may be estimated by dividing the part's supply voltage tolerance by the dynamic load step.

$$Z_{target} \cong \frac{V_{tolerance}}{3 \cdot I_{dymanic}}$$

Whatever method is used to obtain the target impedance value, for a given application I am afraid it *is what it is*. Minimizing the inductance involves the specification of the switching frequency of the VRM and the dielectric thickness of the circuit board planes. It even impacts the choice of connector, if one is used.

Dealing With It As It Is

One recurring theme I hear from designers is that they are putting out fires, not designing a new circuit from the ground up. In this case, one option is to accept the excess inductance as it is. Forget about reducing the inductance, just accept that *it is what it is* and deal with it. But, you might ask, how do I know what the inductance is?

Looking at the example in Fig. 2, which shows a measurement of a VRM's output impedance versus frequency (red trace), the inductance can be determined in a few different ways, depending on which is the simplest.

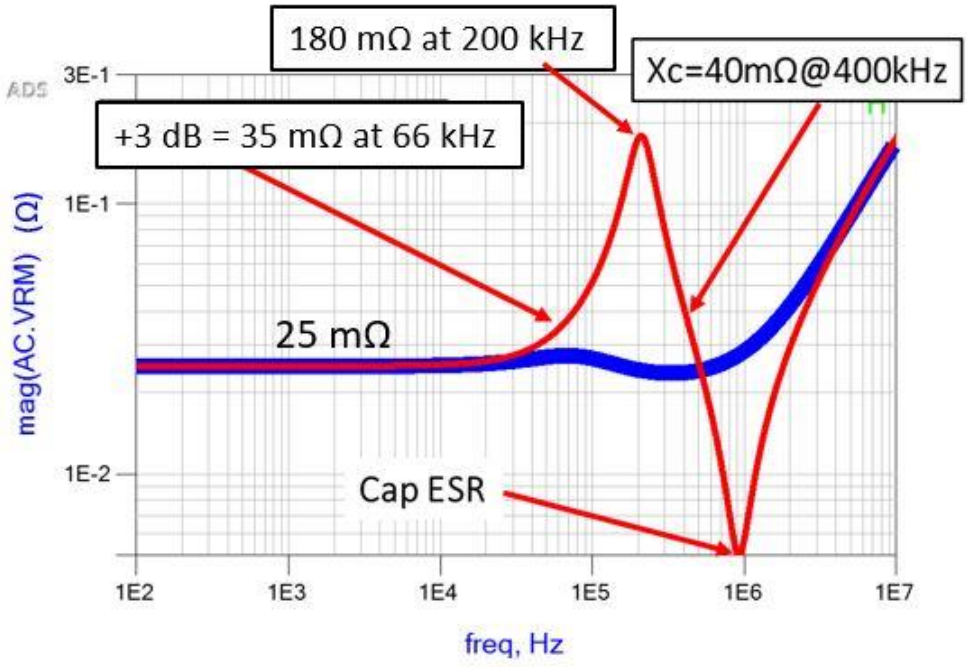

Fig. 1. The resonant peak on this measurement of VRM output impedance (first peak on red trace) provides insight into the excess inductance. A variety of methods can be employed to extract the excess inductance of this resonant circuit.

One way is to determine the from the 3-dB point. Since the flat impedance region below the resonant peak is about 25 mΩ, the 3-dB point is 25*1.414 or about 35 mΩ and that is the frequency where the inductive reactance is equal to the resistance. The inductance can then be calculated as

$$L = \frac{R_{flat}}{2\pi \cdot f_{3dB}} = \frac{25m\Omega}{2\pi \cdot 68kHz} = 58.5nH \quad (4)$$

Another way is from the resonant or peak frequency and the capacitance. Estimating the capacitance from the capacitive side of the resonance using the impedance point 40 mΩ at 400 kHz the inductance can be estimated as

$$C = \frac{1}{2\pi \cdot Z \cdot f} = \frac{1}{2\pi \cdot 40m\Omega \cdot 400kHz} = 10uF \quad (5)$$

$$L = \frac{1}{(2\pi \cdot f_{res})^2 \cdot C} = \frac{12}{(2\pi \cdot 200kHz)^2 \cdot 10uF} = 63nH \quad (6)$$

And yet a third method determines the inductance from the Q of the resonant peak. Some VNAs, such as the OMICRON Lab Bode 100, can directly display the Q of a peak or you can calculate it from the 3-dB points on either side of the resonance.

In this example the peak is approximately 180 mΩ at 200 kHz so the 3-dB points are at 0.707*180 mΩ or 127 mΩ. Measuring the bandwidth at 127 mΩ is approximately 80 kHz.

$$L = \frac{Z_{pk} \cdot f_{bw}}{2\pi \cdot f_{pk}^2} = \frac{180m\Omega \cdot 80kHz}{2\pi \cdot 200kHz^2} = 57nH \quad (7)$$

Regardless of which method was used to determine the inductance, the correct capacitor for absorbing the excess inductance of the VRM output can be calculated as

$$C_{fix} = \frac{L}{R_{target}^2} = \frac{58nH}{25m\Omega^2} = 92uF \quad (8)$$

The closest next value is 100 µF and the capacitor should have an ESR of 25 mΩ. The impedance that results from inserting the 100-µF capacitor with 25-mΩ ESR on the output of the VRM is shown in the blue trace of Fig. 2. Note the impedance is now flat with no resonances.

While we didn't optimize the performance, we were able to eliminate the resonant peak. If a low-ESR capacitor was used this would lead to another resonance at higher frequency or it would lead to much larger decoupling capacitors, so it is important to maintain this correct ESR value.

Or Achieving A More Optimized Solution

Another option is to optimize the design. Achieving power integrity involves balancing equation 3 to provide a flat, or resistive, impedance close to the target impedance. If the resistance is too low than a lot more capacitance will be required and if the resistance is too high the dynamic transient voltage requirements won't be met.

Since the resistance in equation 3 is a second-order term, this is the most sensitive parameter and therefore the most important to get right. If the resistance is too low, you'll want to raise it. This is sometimes controlled by a load line circuit and other times by the feedback amplifier gain. Correcting it in these locations will not degrade the efficiency, while adding series resistance outside of the control loop will. For the sake of efficiency, it's best to deal with this in the control circuitry.[3]

With the resistance set as high as tolerable the next step is to minimize the inductance. The first resonance is often due to the VRM control loop. This is easily verified by measuring the impedance at the VRM output with the input power turned off and with the input power turned on as shown in Fig. 3. Using any of the methods above we could determine the power supply excess inductance to be about 150 nH and the power supply capacitance can be determined to be about 800 µF from the point at 10 mΩ at 20 kHz with input power turned off.

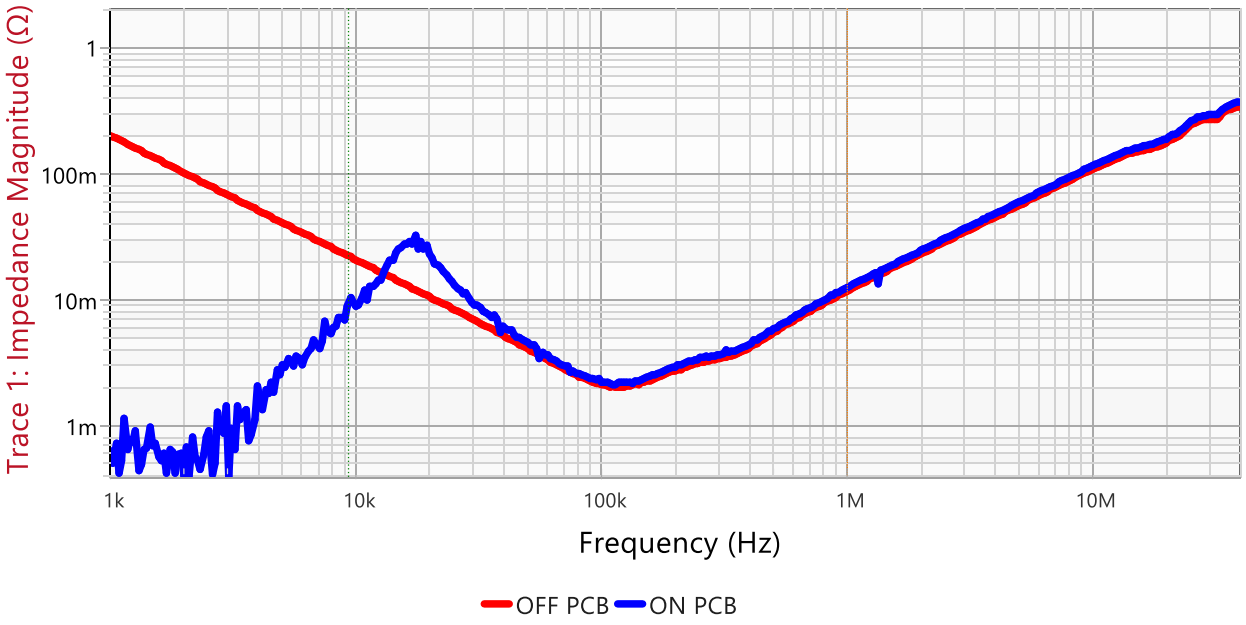

Fig. 2. PCB trace impedance with power off (red trace) and power on (blue trace). This design isn't flat and far exceeds the 6-mΩ target impedance goal. Any of the methods described above can be used to determine the VRM excess inductance value to be about 150 nH.

While we could resolve this resonance with much more capacitance, it isn't practical to add that much capacitance on the PCB as it simply wouldn't fit in most cases. So another option is to reduce the excess power supply inductance. Without knowing why that excess inductance is 150 nH, we do know that this is 18% of the 820-nH output filter inductance. Therefore, reducing the output inductor from 820 nH to 220 nH will reduce the excess inductance to about 40 nH. This is 18% of the 220-nH output filter inductor. If we do that, then the switching frequency should be increased about 50% to keep the ripple current reasonable while not degrading efficiency more than necessary.

The target impedance for this application was 6 mΩ and so using this value and the reduced excess inductance of 40 nH the correct capacitance (required to absorb the excess inductance) can be calculated to be approximately 1100 µF.

$$C_{fix} = \frac{L}{R_{target}{}^2} = \frac{40nH}{6m\Omega^2} = 1111uF \quad (9)$$

However, the maximum capacitance that will reasonably fit in the example design is two 470-µF tantalum capacitors. The output capacitor ESR should be equal to the target resistance and choosing capacitors that are close to 12 mΩ each will satisfy this requirement.

The result, shown in Fig. 3 isn't perfect, since the capacitance is a little bit low. Also, this regulator does not include a load line circuit and so the only way to increase the output resistance would be to add a 5-mΩ series resistor. The added series resistor would dissipate approximately 500 mW at full load and would improve the impedance flatness at the expense of degraded efficiency. Since component pads were not available on the PCB for the added resistor, we skipped it, but still ended up with a much flatter response and very close to the 6-mΩ target value.

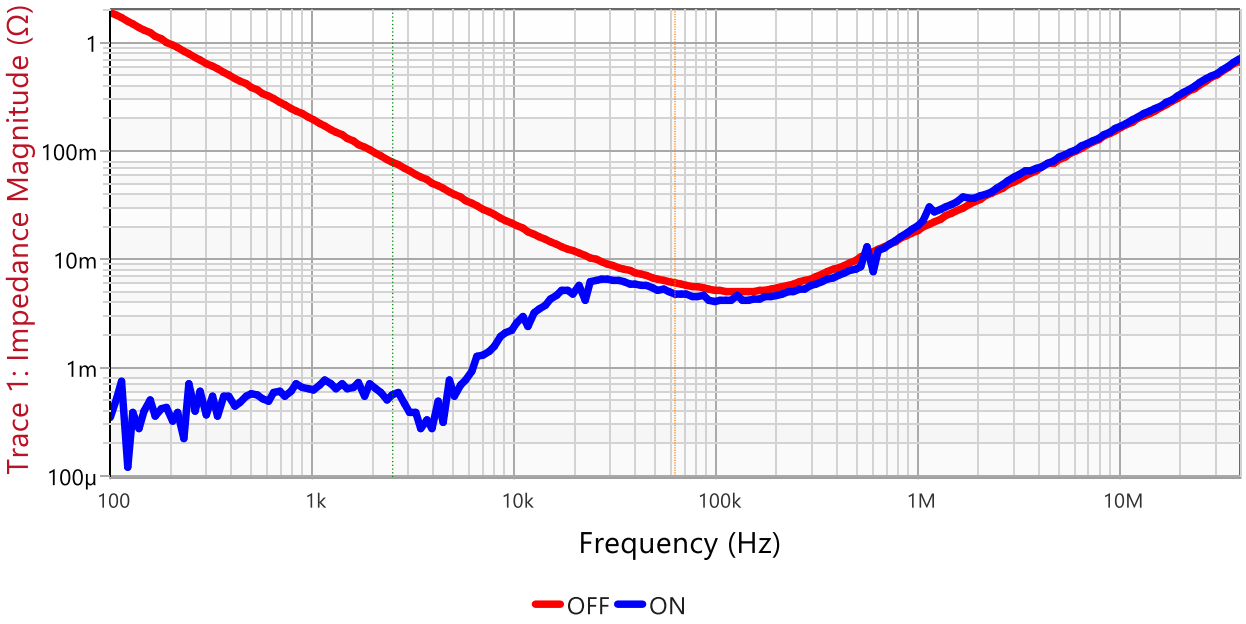

Fig. 3. Reducing the power supply's output filter inductance, increasing the switching frequency and changing the output capacitors got this design back to being reasonably flat and close to the target impedance goal of 6 mΩ.

There are certainly many other techniques that could be used to optimize a new design from the ground up. However, the methods shown here provide a quick method of assessment and two correction paths that will hopefully get you on your way again without much delay.

References

1. "Comparison of Power Distribution Network Design Methods: Bypass Capacitor Selection Based on Time Domain and Frequency Domain Performances" by Istvan Novak, DesignCon 2006.

2. "The inductive nature of voltage control loops" by S.M. Sandler, EDN, Feb. 5, 2015,

3. "PDN Basics For Power Designers" a three part short-video series by Steve Sandler, originally published in How2Power Today, September 2014; Part 1: What's A PDN?, Part 2: Keep Impedance Flat and Part 3: Impedance Matching Is Critical.

4. "Power Integrity for 32 Gb/s SERDES Transceivers" by H. Barnes, J. Carrel and S.M. Sandler, DesignCon 2018.

About The Author

Steven Sandler is the managing director of Picotest, a company specializing in precision test and measurement equipment. Sandler is also the founder and chief engineer of AEi Systems, where he leads development of high-fidelity simulation models for all types of simulators as well as the design and analysis of both power and RF systems.

Sandler has over 30 years of experience in engineering and is a recognized author, educator and entrepreneur in the areas of power, RF and instrumentation. His latest book, "Power Integrity: Measuring, Optimizing and Troubleshooting Power Related Parameters in Electronics Systems," was recently published by McGraw-Hill Education.

Do you have any comments on this column? If so, share them with the editor.