QuietPower **columns**

Do You Really Need that Ferrite Bead in the PDN?

Istvan Novak, Samtec, May 2020

People who do power distribution design know that it is not easy. I would argue that power distribution design is even more challenging today than signal integrity. Not downplaying the technical challenges of high-end signal-integrity designs, my point is that signal integrity today can be considered a rather mature discipline, where system designers, circuit designers and PCB designers alike can rely on a number of standards and other forms of help and support. Power distribution design, on the other hand is much less mature as a discipline and the industry has not figured out yet what is the best way to help designers and users. This has multiple reasons, not just the fact that power integrity started to evolve into its own accepted discipline a couple of decades after signal integrity. For now, let's stay with the observation that many times users have to rely on application notes from chip vendors to figure out how to design the PDN for the active device. Within this still vast area of application notes, we focus on just one question that greatly divides even the experts: is it OK, is it necessary, or is it harmful to use ferrite beads in PDN?

There are two distinctly different camps: people in the first camp will categorically denounce the use of ferrite beads in PDN, while people in the second camp will insist that the PDN they came up with need the ferrite beads.

I start with a generic observation: I admit it may be a little bit of a generalization, but anyway, here it goes. It appears to me that people who suggest using ferrite beads tend to be writing application notes for their company's chips, while people who vehemently oppose it are the ones who try to help users of those chips to come up with a safe design.

As many times in life, the truth is somewhere in between and this creates a third camp of people, mainly the practicing system and circuit designers, who may be torn apart between the opposing opinions, suggestions and recommendations. To add insult to injury, designers working on fast-paced projects may often lack the time, resources and information to come up with the correct answer themselves.

Before we start the illustrations, we need to nail down a few things about the nomenclature. Over the years and decades, as the complexity of our electrical systems have grown, people have been using different names for the same thing. Just one example: the switching regulators of any kind, whether it converted AC to DC, or DC to DC, or DC to AC, decades ago were all called switching-mode power supplies. Later the

regulators converting DC to DC were called DC-DC converters, and in recent years we started to refer to those circuits and devices as VRM (Voltage Regulator Module). For the purposes of this article, and to differentiate between the fundamentally two kinds of ways how we can try to keep the power noise within limits, I will call the passive part of the power distribution network that has only parallel bypass capacitors *Parallel PDN*, and will call a circuit that has an intentional series element (inductor, or ferrite bead or a resistor or a combination of these) for the purposes of blocking the propagation of noise, a *PDN filter*. Note that on printed circuit boards the plane shapes, patches, traces, vias, which are part of the current-carrying path, will also have series resistance and inductance, and most designs consider them as side effects that have to be minimized. On the other hand, these parasitic series elements can also be part of intentionally created filters around them [1].

These two different kinds of circuits can be combined to create various power distribution trees. *Figure 1* shows a small part of the block schematics of a larger system PDN that we may have today, illustrating these two kinds of passive circuits. Here we assume that the PDN filter is placed on the output of a voltage regulator (for instance, to further suppress the switching ripple), but we can also use PDN filters on the input side of a noisy DC-DC converter to block noise from spilling out to the upstream distribution network [2].

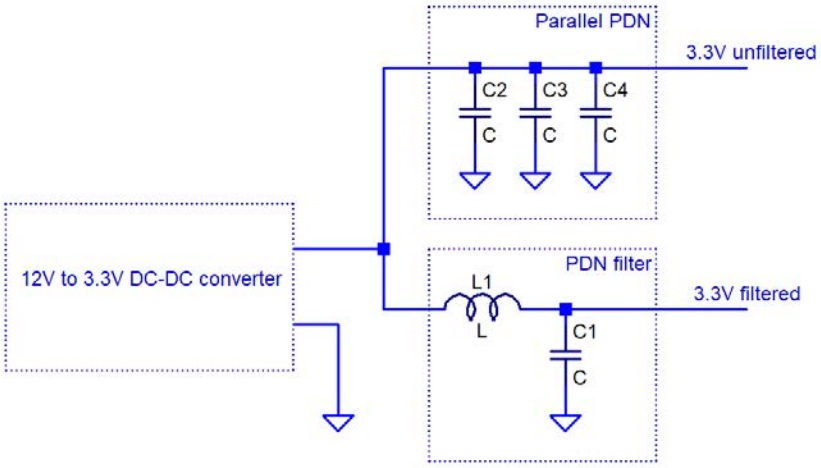

Figure 1: *Block schematics of a system PDN illustrating the definitions of Parallel PDN and PDN filter.*

This block schematics is highly simplified: three capacitors are shown in the Parallel PDN path, but it can be a mix of any number of same-valued and/or different-valued capacitors. Similarly, the PDN filter can be more complex, having an entire Parallel PDN on its output, composed of multiple capacitors and also the series path can be more complex, for instance having series and parallel resistors around the inductive component.

To illustrate the situation and to help explain what should go into the design decision considerations, take a look at the PDN filter circuit suggested for an encapsulated oscillator circuit (*Figure 2).*

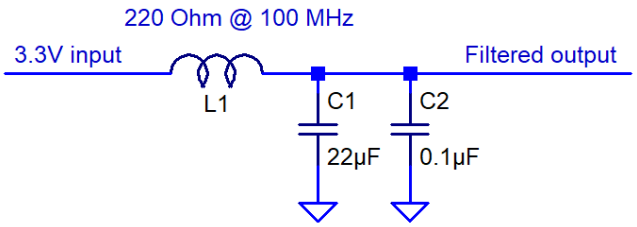

***Figure 2**: Suggested schematics of a PDN filter with series ferrite bead.*

The intended use of this filter was feeding an oscillator circuit with a few milliampere current consumption. If we choose to use filter components by Murata, we can select for instance a BLM21PG221SN1 0805-size ferrite bead for L1, a GRM219R60J226ME47 22uF 0805-size X5R 6.3V capacitor for C1 and a GRM155R61H104ME14 0.1uF 0402-size X5R 50V for C2.

These components have detailed simulation models [3], including temperature and DC bias dependence, so that we can use a circuit simulator to predict what happens. Or, we can purchase the components, build the filter and measure it. And actually, if we have the time, we should do both. But before we do either simulation or measurement, we need to decide what to look at. For filters connecting a high-power noisy rail to a low-current sensitive consumer, the proper metric is the input-to-output voltage ratio [4]. In simulations we can achieve this by using a voltage source for excitation and simulate the output voltage. In measurements, since having a source with zero internal impedance is not feasible, we need to use an instrument with one source and two inputs to measure complex voltage ratios. These instruments are commonly called Frequency Response Analyzers and they allow us to measure the ratio of input and output voltages of the filter.

When we build and measure and simulate this filter's input-to-output voltage transfer function, we get the plots shown in *Figure 3*.

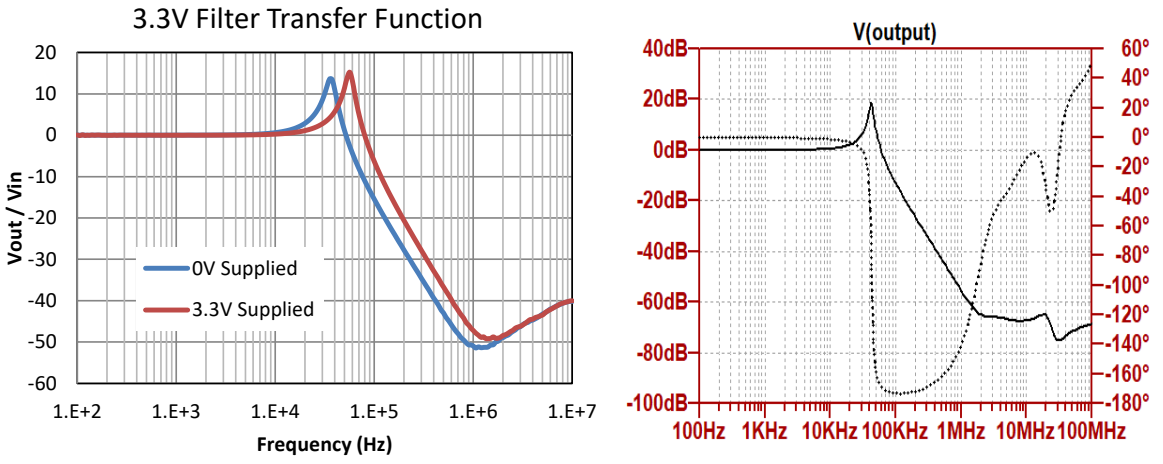

***Figure 3:** Measured (on the left) and simulated (on the right) voltage transfer function of the filter circuit in Figure 2.*

There are two notable features on these plots. First, on the left plot we see two lines, one for each condition of DC bias. The blue line shows what happens without input power, when the input voltage is zero, the red line shows the transfer function with 3.3V DC applied to the filter. The DC voltage across the capacitors and the load current consumed by the oscillator may change the characteristics of the filter components. In general, we should not worry about the blue line, because when there is no input DC voltage, the circuit does not work. Except, the difference between the blue and red lines is a reminder that unless we take the biasing conditions of the DC voltage and DC current into account, we may get the wrong answer. For instance, with this filter at 200 kHz, we get only 20 dB attenuation when the circuit is powered up. Without the DC bias we would see an attenuation of 30 dB. The simulated response with 3.3V DC bias is shown on the right.

The second notable feature, equally troubling, is the peaking of the transfer function at 57 kHz. The peak value is 14 dB, which translates to a five-times voltage amplification of the noise. So instead of making the noise voltage smaller, this filter makes it five times bigger at 57 kHz. We may argue that the purpose of this filter is to suppress the output ripple of a DC-DC converter and in these days the typical switching frequency is above 200 kHz. This is true, but wide-band noise of the converter and the noise generated by the loads connected to the unfiltered rail still could produce significant energy at lower frequencies, including 57 kHz. The peaking of the filter behaves like a high-Q band-pass filter and if we have a noisy unfiltered rail, we can watch on an oscilloscope the output of the filter producing a sine-wave looking ripple.

Also, when we look at the impedance presented by the filter to the load (*Figure 4*), it has a significant peak at the same frequency where the transfer function had its peak.

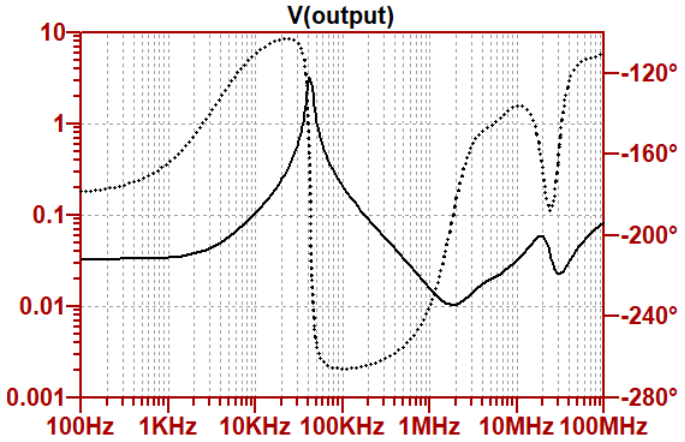

Figure 4: *Impedance of the filter looking back into its output.*

Whether this impedance peak poses a risk and potentially might create a problem, heavily depends on the nature of the load. If the load is 'quiet' in the frequency range of the peak, the peak in the filter's output impedance creates no problem. For instance, if the load is an

oscillator with no digital logic (say a synthesizer) in the same package, this impedance peak would create no issue, but the peaking of the transfer function is still a risk of getting noise on the oscillator supply pin, which will create a higher-than-expected jitter on the output.

If we conclude that the peaking in the transfer function and/or in the output impedance is a risk we want to eliminate, the solution is to add sufficient damping to the circuit. We can do it in multiple ways. We can simply add a larger capacitor with higher ESR on the output, like an electrolytic or tantalum capacitor, or we can add a discrete resistor in series to the 22uF ceramic capacitor. We can also try the add a parallel resistor across the inductive element. Finally, for filters that need to handle only very little DC current, we can add a series resistor.

If we want to totally eliminate the peaking in this particular filter without modifying the series path, we need to increase the total capacitance, so adding a separate lossy capacitor would be the simple way to go. Though in LTSPICE ESR and ESL of a capacitor can be added as attributes of the capacitor symbol, for sake of clarity, those elements are shown in the following figures as separate components. *Figure 5* shows the SPICE deck and result for the transfer function when we add a 220uF tantalum capacitor with 0.1 Ohm ESR. On the left, the schematics in LTSPICE has a voltage source connected to the input. With flat unity source voltage, the output gives us directly the transfer function. The plot on the right shows the magnitude of the transfer function with a solid line and the phase with a dashed line.

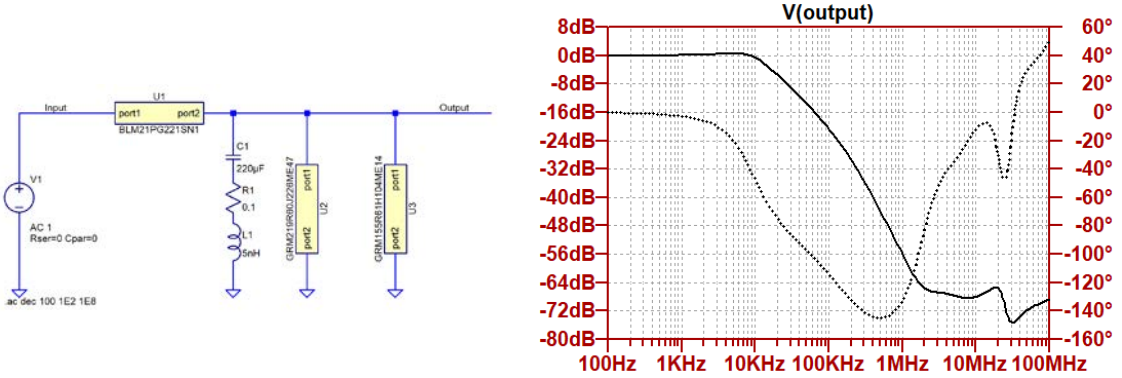

Figure 5: *Voltage transfer function with a lossy capacitor added.*

Figure 6 shows the SPICE deck and result for the output impedance. To simulate the impedance shown by the filter at its output, we short the filter input and connect a one ampere AC current source across the output. With 1A current, the voltage will be equal to the output impedance. In general, using a 1A swept-frequency AC current source is a convenient way to simulate impedance for any linear network.

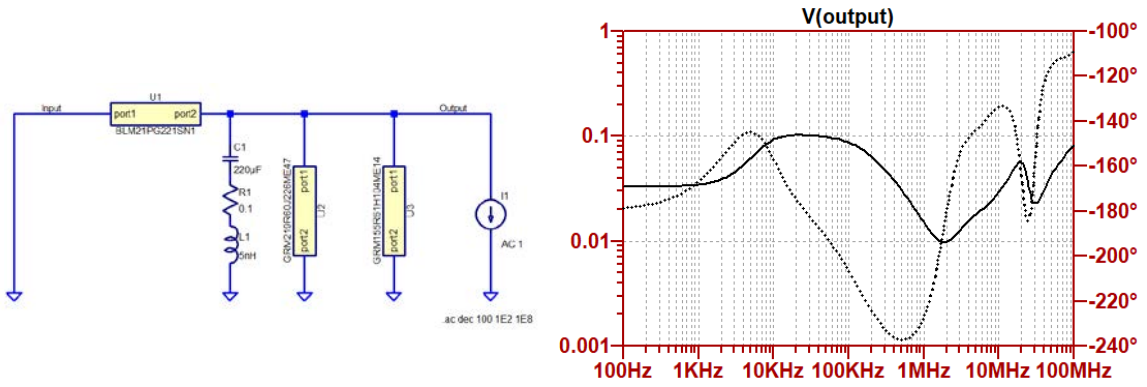

Figure 6: *Output impedance with a lossy capacitor added.*

Once we see the performance of the filter suggested in the application note and we understand the potential risk generated by the peaks in the transfer function and output impedance, we understand what is behind the claims of people who say that ferrite beads should never be used in PDN filters and we should not follow the advice of application notes. We can ask why those peaky filters are suggested in the first place? Were those circuits not tested on a real board? Most likely they were, but we need to see the motivation and the constraints behind those efforts: people creating the application notes for their chips have no way of knowing the many different possible applications their users will come up with. The testing is done in an evaluation board on the bench. This immediately removes most of the external noise and leaves only the noise that is created by the circuit being tested. If the self-generated noise is acceptably low, the circuit will properly function, though we could argue that in a low-noise environment we would not need such a filter in the first place.

We can also see that if we design the filter properly, it should work fine and it gives such isolation benefits that would be much harder to achieve without a series inductive element. In this short article we have not looked at the other benefit of using ferrite beads as opposed to regular inductors: ferrite beads have an AC series loss that increases with frequency sharper than what we get from regular power inductors and as such, it further helps blocking the propagation of high-frequency noise. If you are interested in the full process how to design PDN filters with no peaking, you can take a look at [4].

References:

[1] Alex Waizman, et. Al, “Novel Isolation Scheme for Mitigating PDN Coupling,” DesignCon 2018

[2] Charles Zhang, “Analysis and Design of Input Filter for DC-DC Circuit,” Texas Instruments Application Report, SNVA801–November 2017

[3] Murata SimSurfing tool, https://ds.murata.co.jp/simsurfing/index.html?lcid=en-us

[4] How to Design A Good PDN Filter, http://www.electrical-integrity.com/Paper_download_files/DC19_Tutorial_SLIDES_HowToDesignGoodPDNFilter.pdf, DesignCon 2019