QuietPower columns, June 2021

Ask The Experts: PDN Filters

Istvan Novak, Samtec

In recent years I have been getting a lot of questions about PDN filters from my course participants and from friends, colleagues and even from strangers. Long gone are the days when the essence of power distribution design recommendation was “place a 0.1uF bypass capacitor next to each power pin.” Power distribution networks used to primarily contain wires, traces, planes, parallel bypass capacitors, but very few had systematically designed filters. Just for clarifying the nomenclature, *Figure 1* defines what I call a parallel PDN and a PDN filter. For the purposes of this article, the differentiation between a parallel PDN and a PDN filter is the intentionally included series resistive and/or inductive element in the PDN filter.

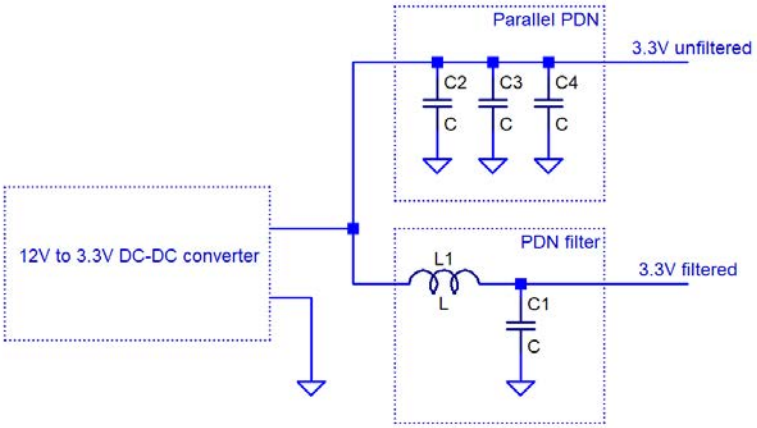

***Figure 1**: Definition of PDN filters.*

The figure shows a very simple implementation of a PDN filter. In actual circuits we may have much more complicated circuits: we may have higher order filters with cascaded sections, filter sections connected in branching topologies and we may also have filters that have to filter in ‘reverse direction’, preventing the noise from spilling out from the load connected to the output of the filter. Or, maybe we need filtering in both ways.

In filter circuits similar to what we show in *Figure 1*, the usual question I get is: how to describe the filter? Should we use Transfer impedance (Z_{21}) or Scattering transfer parameter (S_{21})? As we will point out, in a lot of applications none of these two would serve us well, we will need something else: a voltage transfer ratio.

To systematically design a PDN filter, first we have to establish the requirements. In PDN design work, the biggest pain for a board designer is that input requirements that would guide the designs are very rare, many times almost non-existent. Establishing PDN requirements is often left to the PDN and board designer. There is one situation though when the board designer should be in the position to establish the filter requirements properly: it is when we need to attenuate the output ripple of a switching regulator for a sensitive pin of a chip that has a specified maximum allowed noise. To quote some simple numbers, the typical peak-to-peak regulator ripple may be around 10mVpp, which is usually too much for a sensitive analog circuit, such as a reference clock or clock buffer circuit, PLL, SerDes supply, or sensitive analog circuit, more likely requiring around 1mVpp maximum noise. In such cases, relying on the allowed ripple voltage limit of our sensitive circuit and having an estimate of the regulator's switching frequency and output ripple, we can determine an attenuation requirement, at least at the switching frequency of the regulator.

As an illustration, lets take a filter circuit that is supposed to attenuate ten times the switching ripple of a DC-DC converter running at 500 kHz for an oscillator circuit that takes just a few milliamperes of current. In addition to the series inductive element and the output capacitor, the circuit in *Figure 2* also includes components that model the source side of the filter. R1 and L1 represent the impedance of the DC source and C2 and C3 with their parasitics model the board capacitors. The L5-R5 circuit is the model for a small inductor. As opposed to a ferrite bead, where we would anticipate a large increase of series resistance and a substantial drop of inductance at high frequencies, this simple inductor model has frequency independent series resistance and inductance. The filter output is a single capacitor, modeled by C4-R4-L4. Note that this model intentionally has a very high ESR for C4 that we can get either from a small electrolytic capacitor (though it will likely have a much higher ESL then) or we can use a regular ceramic capacitor with low ESR and add a 0.91-ohm series resistor. We also knowingly ignore the printed circuit board details and therefore we limit the AC simulation to below 100 MHz. The filter behavior at higher frequencies will be influenced and probably dominated by component placement, the PCB layout and stackup; something that is beyond the scope of this article.

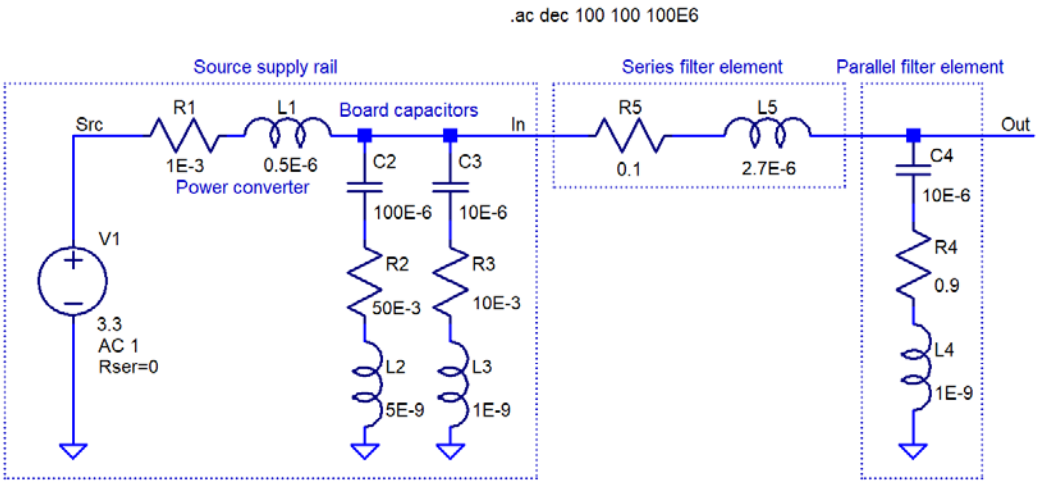

Figure 2: *Schematics of a simple one-stage PDN filter.*

We simulate the filter circuit in the frequency domain in the 100 Hz – 100 MHz frequency range with a 1V ideal AC source with zero source resistance behind the R1-L1 elements modeling the DC source. *Figure 3* shows the simulated voltages at various nodes. The V(src) source voltage shows a flat line at 0 dB, since this is the constant source voltage we enforce. V(In) and V(Out) are the input and output voltages of the filter. And finally, V(out)/V(in) is the voltage transfer ratio from input to output. Note that while V(in) and V(out) both have some minor peaking around 20kHz, the light-blue line of V(out)/V(in) shows no peaking, but it also has just 20 dB/decade slope, whereas the absolute output voltage drops steeper.

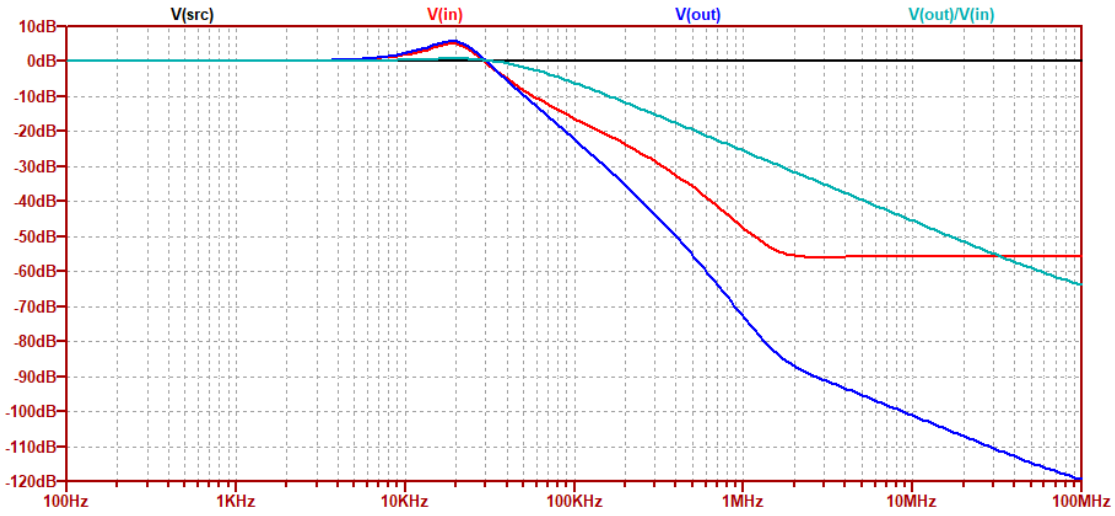

Figure 3: *AC simulation results of the filter from Figure 2.*

The various impedances can be simulated by attaching a 1A AC current source to the pieces, as shown in *Figure 4*. The AC simulation uses linearized models and therefore the 1A AC test current that may otherwise damage some of the real components, can be used here with no concern.

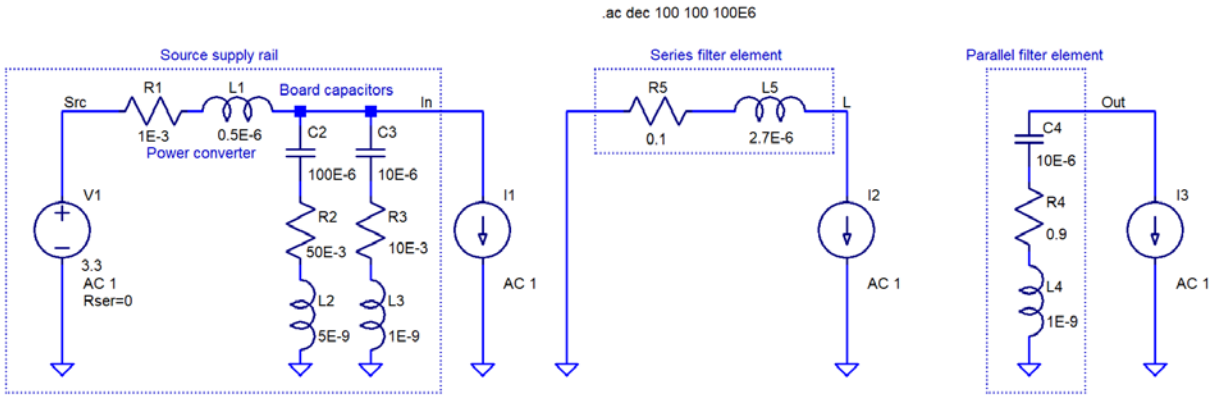

Figure 4*: Simulating the impedance of blocks in the PDN filter.*

In *Figure 5* we see that the minor peaking is related to the impedance of the main supply rail that serves as a source and provides input to the filter. The impedance of the series inductor starts out at the 0.1-ohm DC resistance and at about 5 kHz it goes inductive. Due to the large series resistance, the impedance of the output capacitor of the filter becomes resistive at about the same frequency where the source impedance has the peaking.

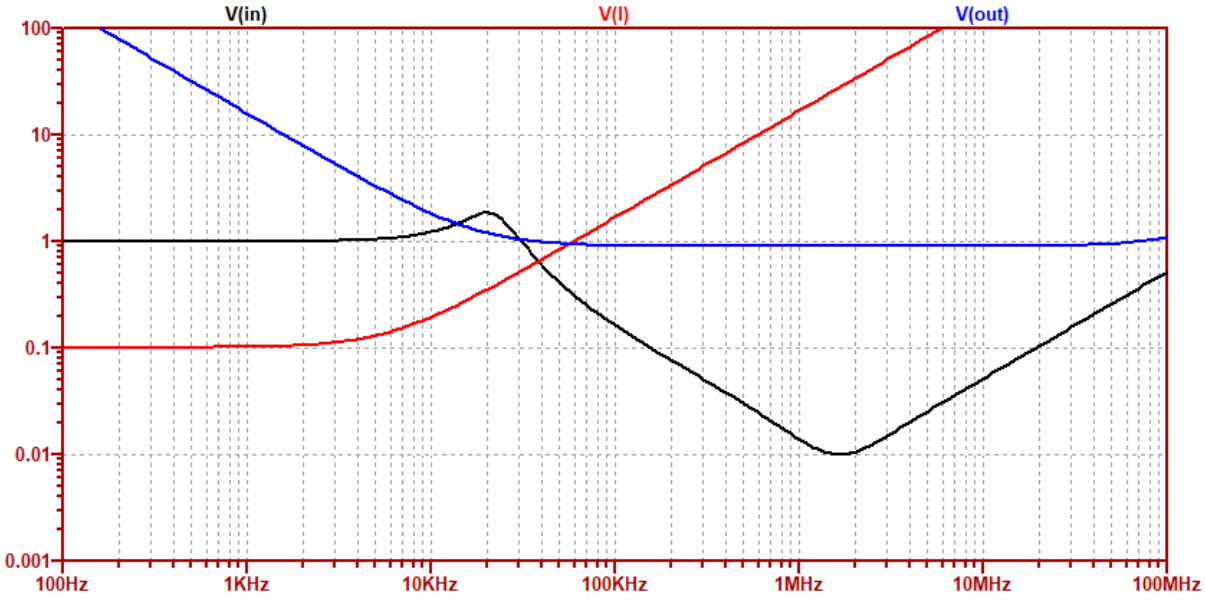

Figure 5: *Impedance magnitude of the source rail, the inductor and the output capacitor of the filter. Because we used a 1A AC source, the voltages at these nodes directly represent impedance.*

We can also simulate the input and output impedance of the full filter with its actual terminations. As shown in *Figure 6*, we can simulate the output impedance of the filter with the source-side impedance connected.

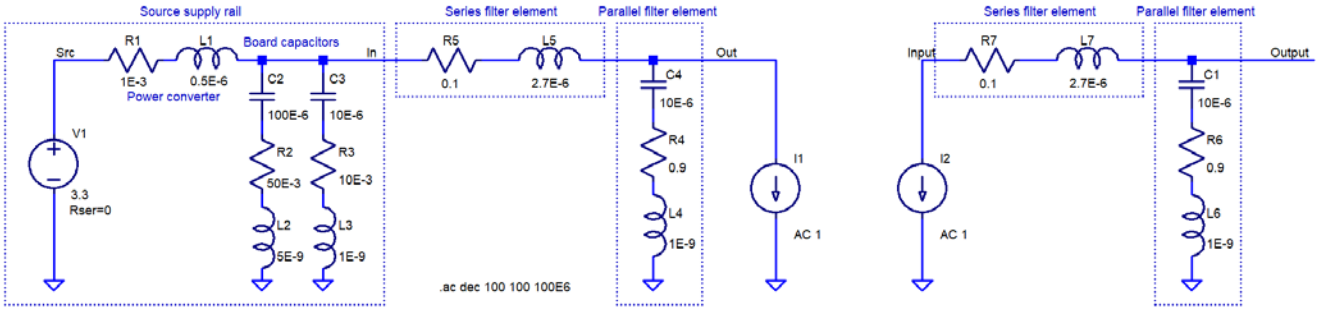

Figure 6: *Simulating the full input and output impedance of the filter.*

In the reverse direction, the input impedance is simulated with the assumption that the load has very high impedance and therefore we can leave the output of the filter open. On the input side of the filter, we drive the series inductive element without a shunt capacitor included. This is an acceptable approximation for unidirectional filters when the impedance of the source rail is much lower than the input impedance of our filter. When we need to work with multi-stage and/or bi-directional filters, where the driving source impedance is not much lower than the input impedance of our filter, we need to identify and include both at the input and at the output the capacitors that belong to our filter circuit.

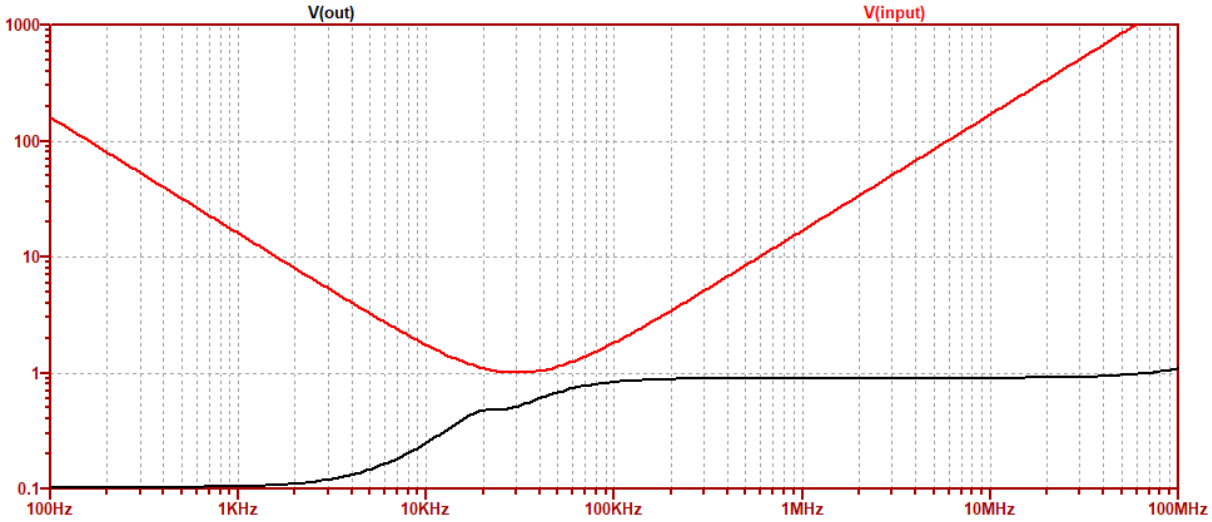

***Figure 7**: Input and output impedance of the full filter.*

The input and output impedance curves of the full filter are shown in *Figure 7*. The input impedance curve compares to the blue line in *Figure 5* with the difference that now we have a 2.7 uH inductance in series. This creates an impedance profile as if the capacitor on the output had a 2.7 uH ESL, which effectively removes the flat impedance plateau that the 0.91-ohm ESR created. The output impedance of the filter above 100 kHz matches the 0.91-ohm ESR of the output capacitor. At lower frequencies the output impedance approaches the 0.1-ohm DC resistance of the inductor.

We can also simulate the network parameters that we may otherwise consider using to describe the filter: Z21 and S21. To do this, we need to modify the simulation circuit by specifying the ports, the termination impedance and add one line that links the input and output: '.net I(Rout) V2'. The SPICE deck is shown in *Figure 8*, the result is in *Figure 9*. The three curves in the plot clearly illustrate the difference among them and show us why in our chosen case the voltage transfer function is a better metric. First, we can notice that the S21 and Z21 curves run in parallel. The vertical difference between them is 25x, which comes from the parallel equivalent of the two 50-ohm terminations we need to calculate S21. We also notice that both curves start at very low values at 100 Hz: this may seem to be good news, but in reality, it is misleading. In this test we assume that the input side of the filter is driven by a high-current low-impedance

supply rail. We get a low S21 or Z21 value simply because the low source-side impedance sunts out our test signal. In reality, whatever noise appears across the source rail will be forced across the input of the filter, which represents a much higher impedance and as such it has very little influence on the input noise.

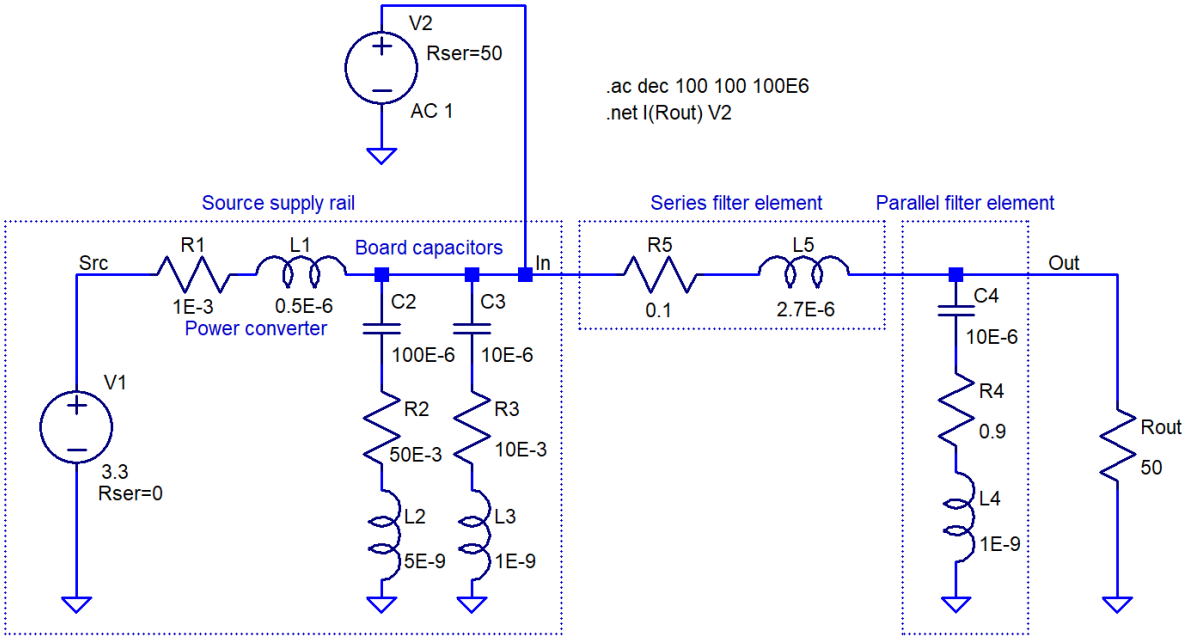

Figure 8: Simulation deck for network parameters.

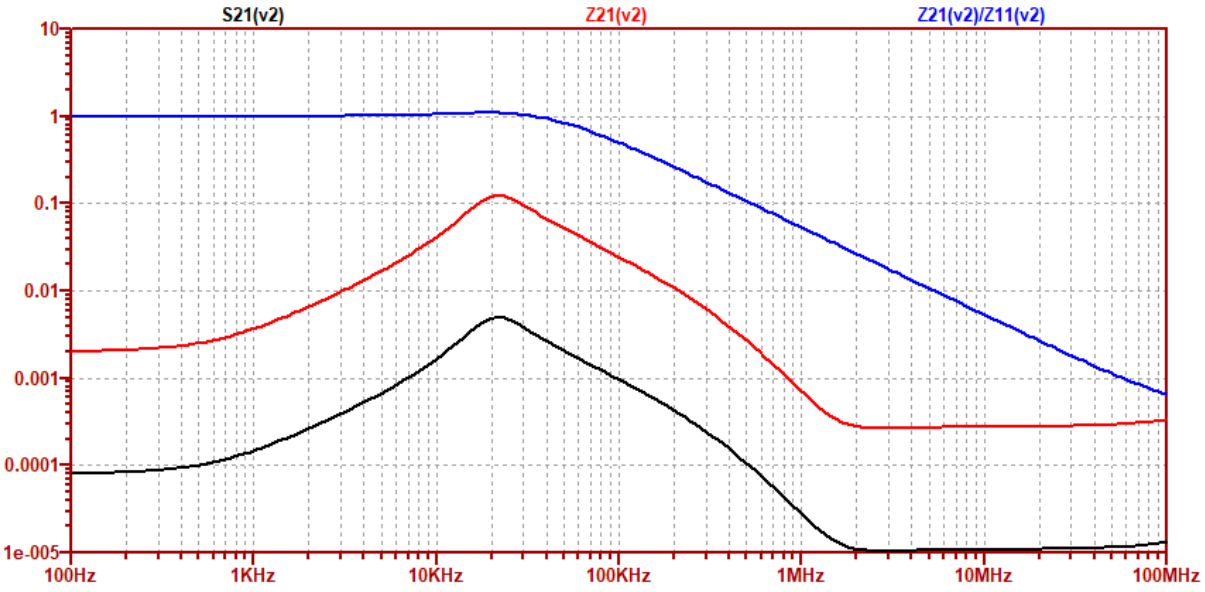

***Figure 9**: Z21, S21 and voltage transfer ratio.*

When we calculate the noise transfer with the source rail's impedance in place, the finite source impedance associated with S21 and the infinite source impedance associated with the current source for Z21 will produce this large attenuation at low frequencies where the low impedance of the DC source dominates the impedance of the source rail. The other extreme case, when we leave out the source impedance altogether, would produce equally unrealistic results. We show this in *Figures 10* and *11*.

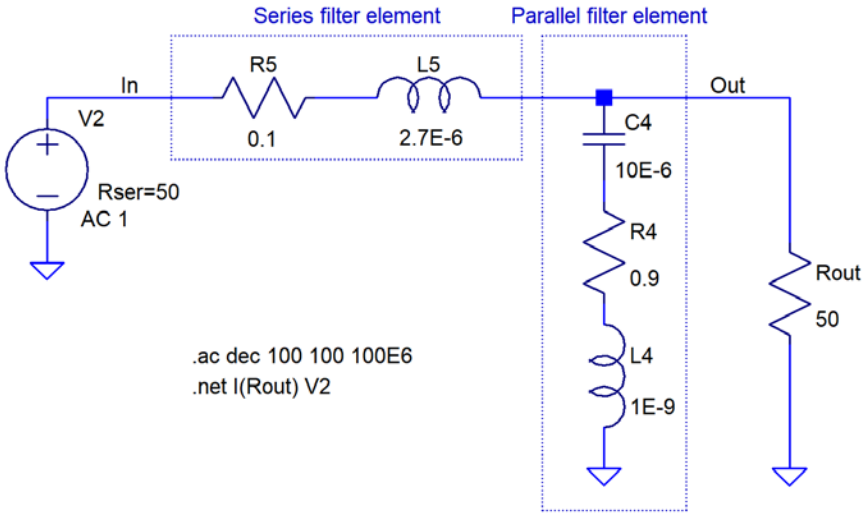

***Figure 10**: Simulation deck for filter transfer function without including the impedance of the input-side supply rail.*

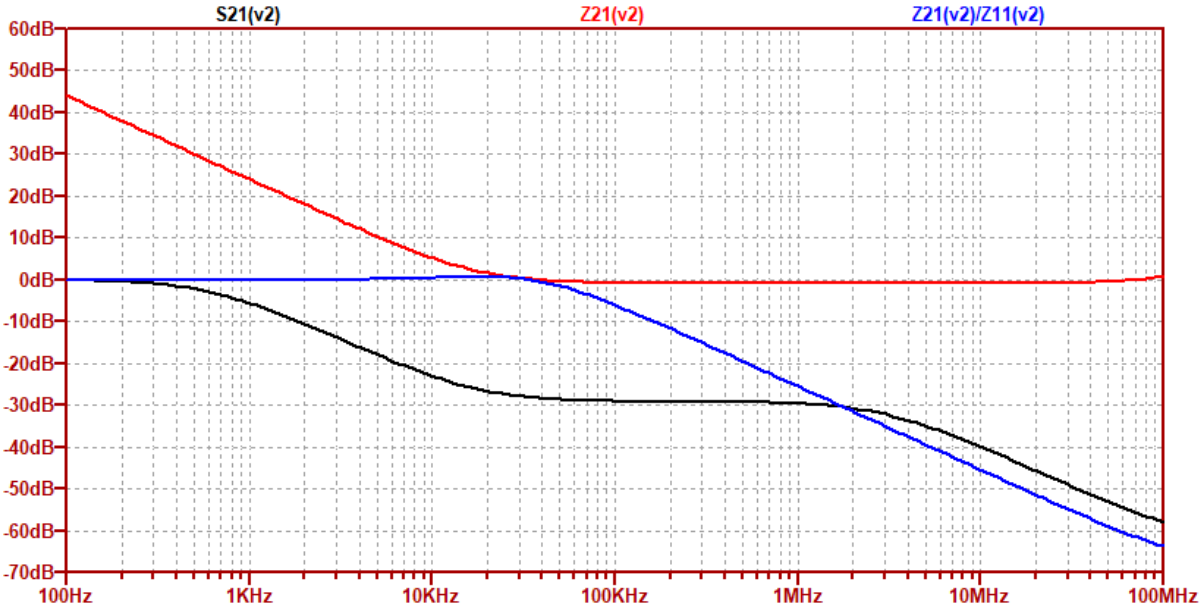

***Figure 11**: Transfer parameters for the circuit of Figure 10.*

S21 starts out at 0 dB at 100 Hz but then it drops unrealistically due to the 50-ohm source impedance, which would be very unexpected from a high-current rail impedance. Z21, on the other hand, starts with a large gain and then it settles at the ESR of C4. *Figures 9* and *11* illustrate that both S21 and Z21 hugely depend on the impedance of the source rail, whereas the voltage transfer ratio, by its definition, is not.

We need to remember that details matter: the illustration and conclusions are valid for the stated category of cases when we connect a filter to a main rail, which represents a source-side impedance much lower than the load impedance of the filter output. If you are interested in further details, want to learn about filter measurements and curious about correlations between measurement and simulation, check out the further reading listed below.

Further reading:

[1] Frequency-Domain Characterization of Power Distribution Networks, Artech House, 2007, Section 5.4

[2] Using Ferrites and Inductors in Power Distribution Networks (PDN), Samtec gEEk spEEk, December 3, 2020. View it at https://vimeo.com/487288222

[3] Do You Really Need that Ferrite Bead in the PDN? Available at http://www.electrical-integrity.com/Quietpower_files/QuietPower-55.pdf