DesignCon 2018

How Spatial Variation of Voltage Regulator Output Impedance Depends on Sense Point Location

Alejandro ‘Alex’ Miranda, Oracle Corporation,
alex.miranda@oracle.com

Joseph (Abe) Hartman, Oracle Corporation,
abe.hartman@oracle.com

Kavitha Narayandass, Oracle Corporation
kavitha.narayandass@oracle.com

Ethan Koether, Oracle Corporation
ethan.koether@oracle.com

Gustavo Blando, Oracle Corporation,
gustavo.blando@oracle.com

Istvan Novak, Oracle Corporation
istvan.novak@oracle.com

Abstract

Power converters are part of almost all electronic devices, and yet, there are still some unknowns about them. Measuring the impedance of the power delivery network is a common practice in designs based on the target-impedance approach. But what happens when the impedance we observe changes depending on the location we measure? Or with the converter sense point location? Also, one of the key parameters we can approximate from the impedance is the stability of the converter, but can we still rely on the impedance to back calculate the stability of the converter if it is location dependent?

Author(s) Biography

Alejandro 'Alex' Miranda is a Senior Hardware Engineer at Oracle focused in power integrity for SPARC Low/Mid range servers. Prior to joining Oracle, Alex was a Hardware Engineer at Intel working on the design of DC/DC converters and PDN simulation for datacenter products. Alex earned his BS in Electrical Engineering from Instituto Tecnologico de Celaya, Mexico

Joseph 'Abe' Hartman is a Principal Hardware Engineer focusing on system signal and power integrity at Oracle. Abe has worked as a signal integrity engineer at Amphenol TCS, Juniper Networks, and Enterasys. Abe also worked at General Motors. Abe holds a MS in Electrical Engineering from the University of Massachusetts-Lowell, a MS in Engineering Science from Rensselaer Polytechnic Institute, a BS in Mechanical Engineering and a BS in Electrical Engineering from Kettering University, Flint, MI.

Kavitha Narayandass is a Hardware Engineer at Oracle. Her work focuses on Power Integrity for SPARC servers. Prior to joining Oracle, she was a Hardware Design Engineer, working on Circuit Design and Signal Integrity simulation at Wipro Technologies, India. She holds a Bachelor's in Electrical and Electronics Engineering from Bharathiar University, India and is pursuing her Master's degree in Electrical and Computer Engineering at Georgia Institute of Technology, GA

Ethan Koether is a Hardware Engineer at Oracle Corporation. He is currently focusing on system power-distribution network design, measurement, and analysis. He received his master's degree in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology.

Gustavo Blando is a Senior Principle Engineer leading the SI/PI team at Oracle Corporation. In addition to his leadership roles, he's charged with the development of new SI/PI methodologies, high speed characterization, tools and modeling in general. Gustavo has twenty plus years of experience in Signal Integrity and high speed circuits.

Istvan Novak is a Senior Principle Engineer at Oracle. Besides signal integrity design of high-speed serial and parallel buses, he has been engaged in the design and characterization of power-distribution networks and packages for mid-range servers. He creates simulation models, and develops measurement techniques for power distribution. Istvan has twenty plus years of experience with high-speed digital, RF, and analog circuit and system design. He is a Fellow of IEEE for his contributions to signal-integrity and RF measurement and simulation methodologies.

I. Introduction

The power distribution network for many of our electronic designs today can be described with the simplified equivalent circuit of *Figure 1*. The *DC source* could be an analog or switching DC-DC regulator, converting the incoming battery or main power supply voltage to a regulated DC voltage matching the need of the *Die*. In a typical scenario the *DC source* provides not only voltage translation, but also regulation, keeping the output voltage constant, by compensating for varying input voltage and load current. This is achieved by a feedback loop, which monitors the voltage at a specified location. In *Figure 1* this is represented by a single sense line, which can be tied to the power plane under the *Package* or can be connected to the sense line if the *Package* has a dedicated sense connection going all the way to the *Die*. In the figure we show only high-side sense with the two potential connection locations, but voltage regulators are also available providing differential high-side and low-side sense, in which case the regulator can also compensate for voltage drops on the power return (ground).

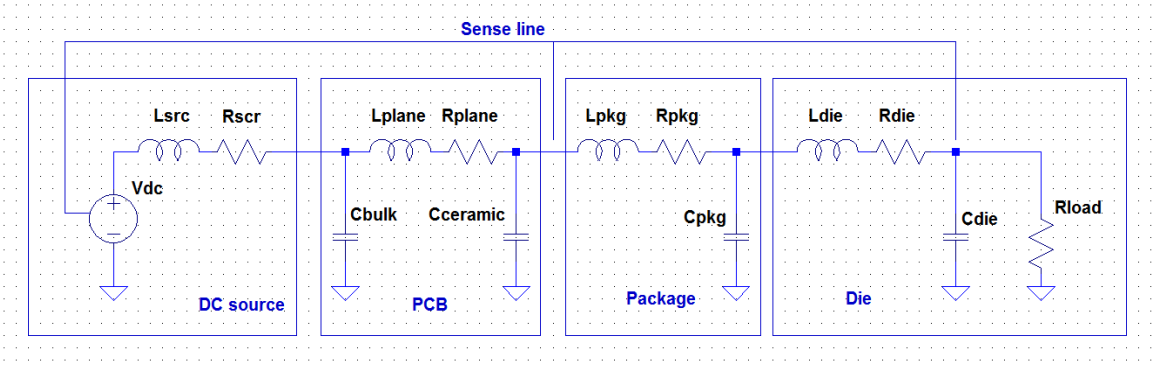

***Figure 1**: Source-to-silicon simplified ladder model of a point-of-load PDN.*

System and board designers have to design and validate this PDN by selecting the components, building it, testing and characterizing it. The design has to be based on specifications, among others the maximum allowed voltage fluctuation in response to changing load currents.

We usually need the V_{AC} AC transient voltage across the load to be much smaller than the nominal DC rail voltage V_{nom} DC load voltage that we want to maintain across the die. The typical ratio is one-and-half to two orders of magnitude: V_{AC} is a few percent of V_{nom}. With the point-of-load scheme we assume that noise at the silicon is primarily due to the fluctuations in current demand of the silicon: 'crosstalk' noise initiated by other loads in the system is neglected. Similarly, the noise generated by the power source is assumed to be either negligible or just a small portion of the worst-case noise on the supply rail. There are also series resistances along the path and the DC average current of the load will produce a DC drop between the DC source output and the load, or for cases where the DC source has a remote-sense connection, between the remote sense point and the load. The DC source also has a finite tolerance and drift on its set output voltage. This is illustrated in *Figure 2*. The DC set-point accuracy and any uncompensated DC voltage drop is represented by the green bar on the voltage graph.

The transient noise as a result of load-current fluctuations plus any DC-DC converter output ripple (in case the DC source is a switching regulator) are represented by the orange bar. Altogether the DC voltage tolerance and DC drop together with the worst-case AC noise have to stay within the V_{max} and V_{min} voltage limits established for the silicon. For sake of simplicity here we assumed no added margin, which can be easily added to this calculation if needed. Without restricting generality, symmetric voltage allocation was assumed, where the DC tolerance and AC tolerance bars are evenly placed above and below the nominal voltage. Dependent on the actual circuit details, in some of the real systems both the DC and AC tolerance bars may be arranged in an asymmetric way around the nominal voltage. From the noise budget sketch of *Figure 2*, we can calculate the maximum allowed voltage fluctuation ΔV as a result of transient currents.

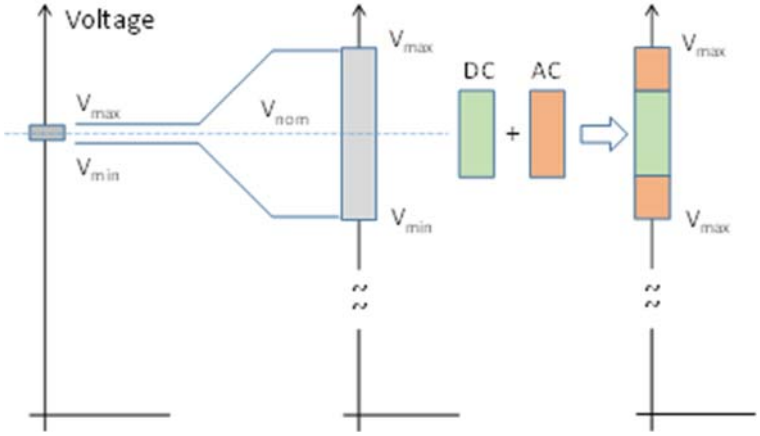

***Figure 2**: Voltage diagram of a supply rail, illustrating some of the major design parameters.*

Though the fundamental specification items are in the time domain, many designers choose the frequency domain for the design process [1], based on the *Target Impedance* concept. The Target Impedance is the ratio of the allowed voltage fluctuation ΔV and the maximum current fluctuation ΔI.

$$Z_{target} = \frac{\Delta V}{\Delta I}$$

This impedance target calculation works well to predict the worst-case noise as long as we create a relatively flat and frequency independent PDN impedance at the load. With a flat (resistive) PDN impedance with Z_{target} value, we guarantee that for any arbitrary sequence of ΔI current steps the resulting transient noise will not be bigger than ΔV [2]. Strictly speaking this also assumes that the entire PDN is linear and time invariant. In the chain of components of *Figure 1*, the two ends may be nonlinear [5]: the *DC source* and the *Die*. The discrete capacitors and PCB elements (including the package), however, are sufficiently linear in most practical cases. Because the capacitors and the PCB cover multiple decades of frequency, the frequency-domain *Target Impedance* approach is a

usually good design approximation. For testing and validation, on the other hand, the dilemma is real: do we want to do it in the frequency or the time domain?

There are many instances when we cannot, or for some reason do not want to, create a flat PDN impedance. As it was pointed out and explained in [3] and [4], for non-flat impedances, not exceeding the target impedance, we pay a penalty of increased worst-case noise, where the penalty ratio depends on how much we distort the flat impedance. When the impedance profile is not flat, the worst-case transient noise increases. We illustrate this with the hypothetical PDN circuit of *Figure 3*. It is a resistive PDN with a frequency-independent 0.1-Ohm value, except a single capacitor, whose Series Resonance Frequency (SRF) is 1 MHz. The impedance profile is shown in the lower left, the step response is shown in the lower right. The parameter is the impedance minimum driven by the capacitor's SRF.

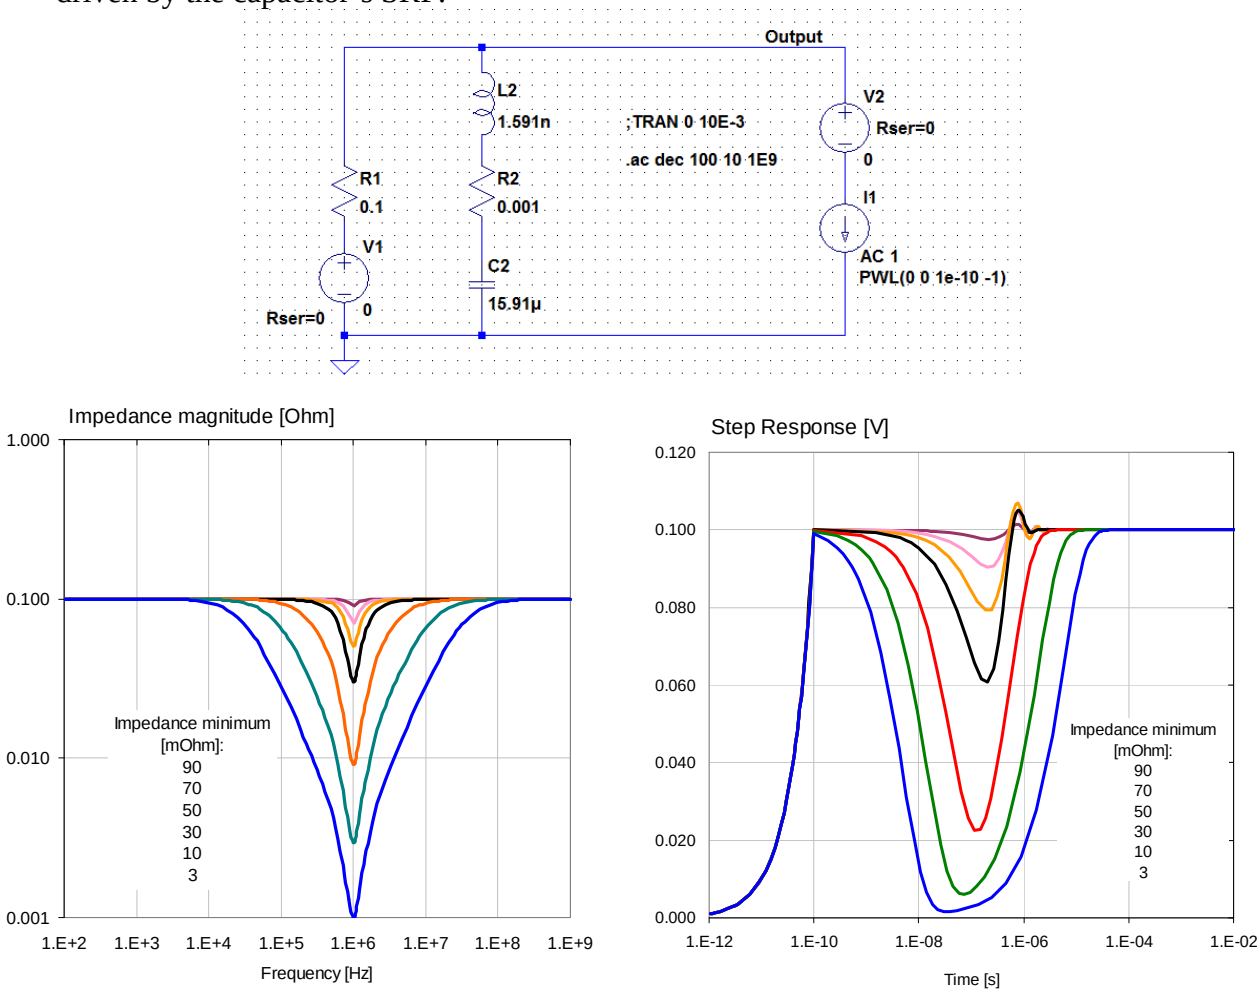

***Figure 3**: A simple PDN schematics (top), impedance and step response (lower left and right).*

We can take the step responses with different impedance minimum values, and calculate the worst-case transient noise as a function of the relative depth of the impedance minimum; the calculated increase of worst-case noise with respect to its flat-impedance baseline, is shown in *Figure 4*.

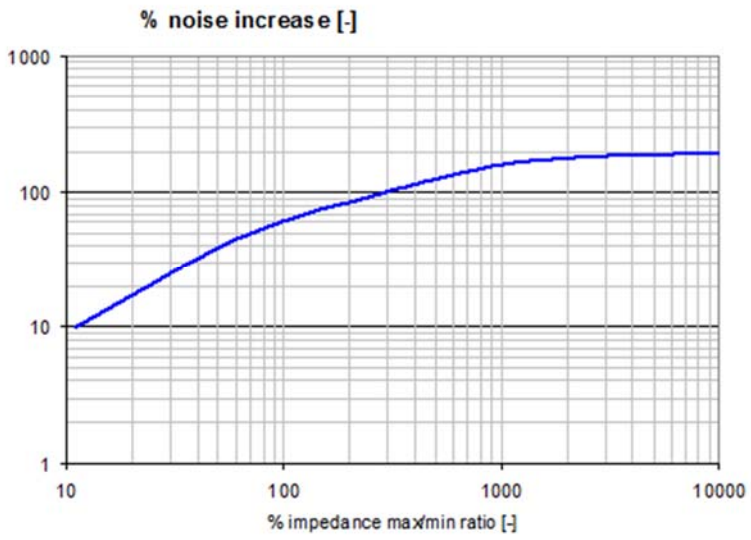

Figure 4: *Relative noise increase as a function of relative max/min ratio of impedance profile on a flat impedance with a single second-order notch.*

The above illustration shows that though designers are usually concerned about impedance peaks, but dips are equally bad.

The increase of worst-case noise is additive for each disturbance (dip or peak) of the impedance profile, as long as the signatures are sufficiently separated in frequency. An illustration is shown in *Figure 5*. The schematics (top) now has two capacitors, each driving the impedance down at their particular SRF.

The impact of the dip in the frequency domain is the same in the time domain regardless of where on the frequency axis the dip occurs. This includes zero frequency as well, which means that the lack of voltage positioning in the DC source (low output resistance), creates an additional 'dip' in the frequency domain, at DC. As we keep adding dips in the frequency domain, the worst-case transient noise keeps going up.

These illustrations are intentionally simplified and exaggerated, just to show the effect. Luckily in real designs having a large number of deep dips in the impedance profiles would be hard to produce even on purpose, which puts a practical limit on how bad the worst-case noise can be compared to what would be generated by a flat impedance. In the following sections we will explore practical situations on two evaluation boards and a system board. The evaluation boards are used because of their simple geometry, making it easier to follow certain constraints and to understand the results.

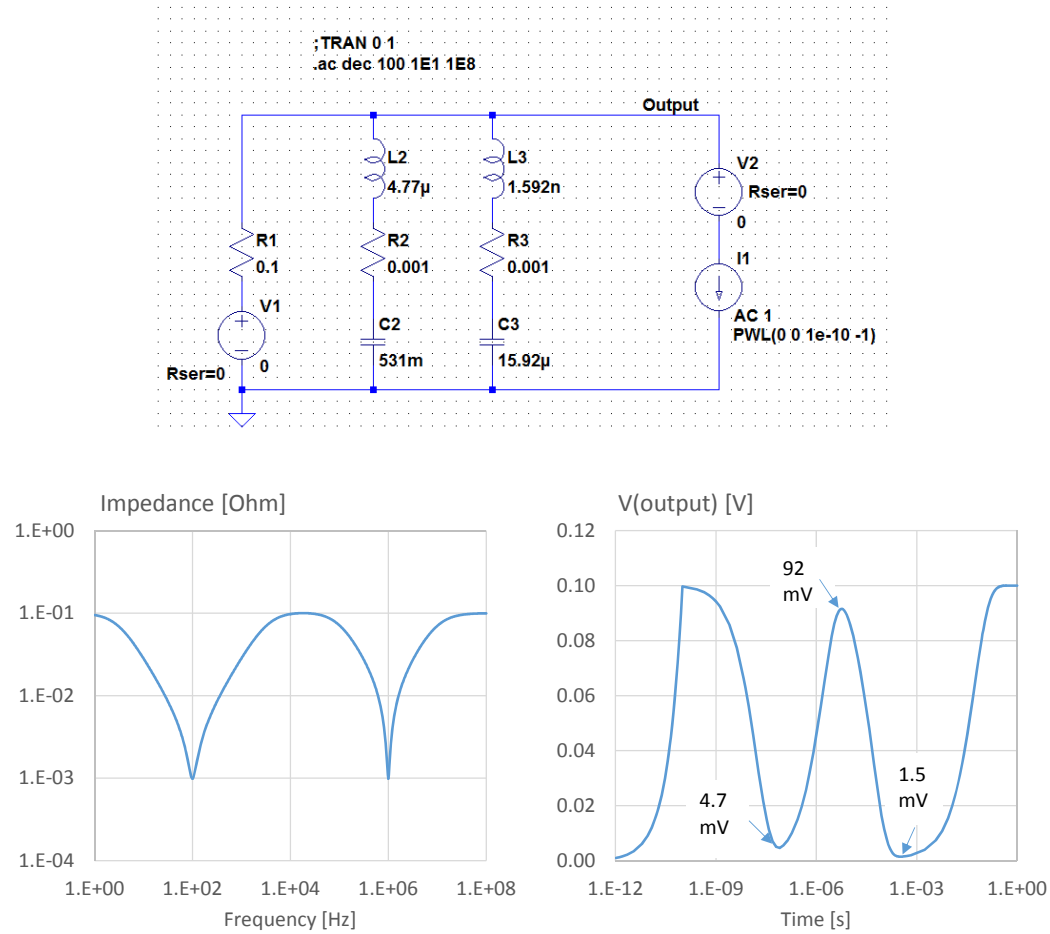

Figure 5: A simple PDN schematics (top), impedance and step response (lower left and right with two notches in the impedance profile.

II. Time-to-frequency domain correlation

To show the correlation between time and frequency domain and its limits, we use a small DC-DC converter evaluation module [10]. First we use the evaluation board without changing anything on it and measure its output response in the time and frequency domain with two different connection schemes. *Figure 6* shows the preferred connections and one of the possible wrong connections. The correct connection will attach both the injection source and the response sense lines where the response is picked up, at the same location, free of any voltage drop and filtering that is not part of the intended power path to the load [11]. This particular evaluation board has small metal turrets for connecting the input power and the load. At low frequencies, where we want to test only the response of the DC-DC converter, the preferred connection is to attach the source and response sense cables to those metal pins. If we want to make sure that the resistance and inductance of metal turrets will not be part of the measured data, one cable can connect on the top side of the board, the other connection on the bottom side of the

board. These rules uniformly apply for both frequency and time domain instruments. If we make one or both of the connections for instance across a capacitor (as shown by the red arrows in *Figure 6*), the pad and connection resistance and inductance of the capacitor hook-up will create a frequency dependent transfer function for the injected and/or the sensed signal. *Figure 7* shows a simplified equivalent circuit and a typical frequency response [12].

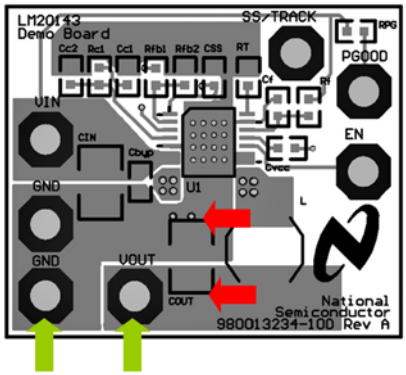

***Figure 6**: Top view of the evaluation board. The preferred connections to the instruments are shown by the green arrows. One of the potential wrong connections are shown by the red arrows.*

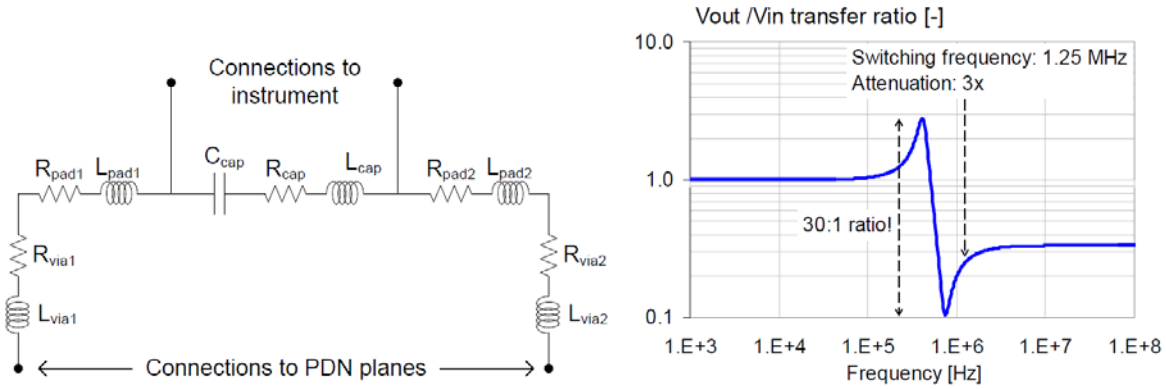

***Figure 7**: Equivalent circuit of a connection across a capacitor (left) and its resulting transfer function (on the right).*

Figure 8 shows the frequency-domain response measured with a VNA with 0dBm source power and 1.5A external DC load and the time-domain response with 0.5A 100ns current step from 1.5A to 2.0A. In each response we also show the particular response back-calculated from the other domain. The data was taken with no change to the evaluation board. The input voltage was 4.5V. The left graph shows multiple traces: step response from rising current stimulus measured in the time domain (solid line), inverted and shifted step response from falling current stimulus measured in the time domain (solid line) and step response back calculated from the measured impedance profile (dashed line). Note that all responses line up well, in fact the step responses from the rising and

falling current stimulus run on top of each other and only one color is seen. This suggests that with this level of excitation the converter is sufficiently linear so that we can go back-and-forth between time and frequency domains with high confidence. The bottom plot in figure 8 shows the same step response, except plotted on logarithmic horizontal scale. This reveals that the bandwidth of the excitation pulse extends to the rising slope of the impedance profile above the series resonance frequency, creating a small blip.

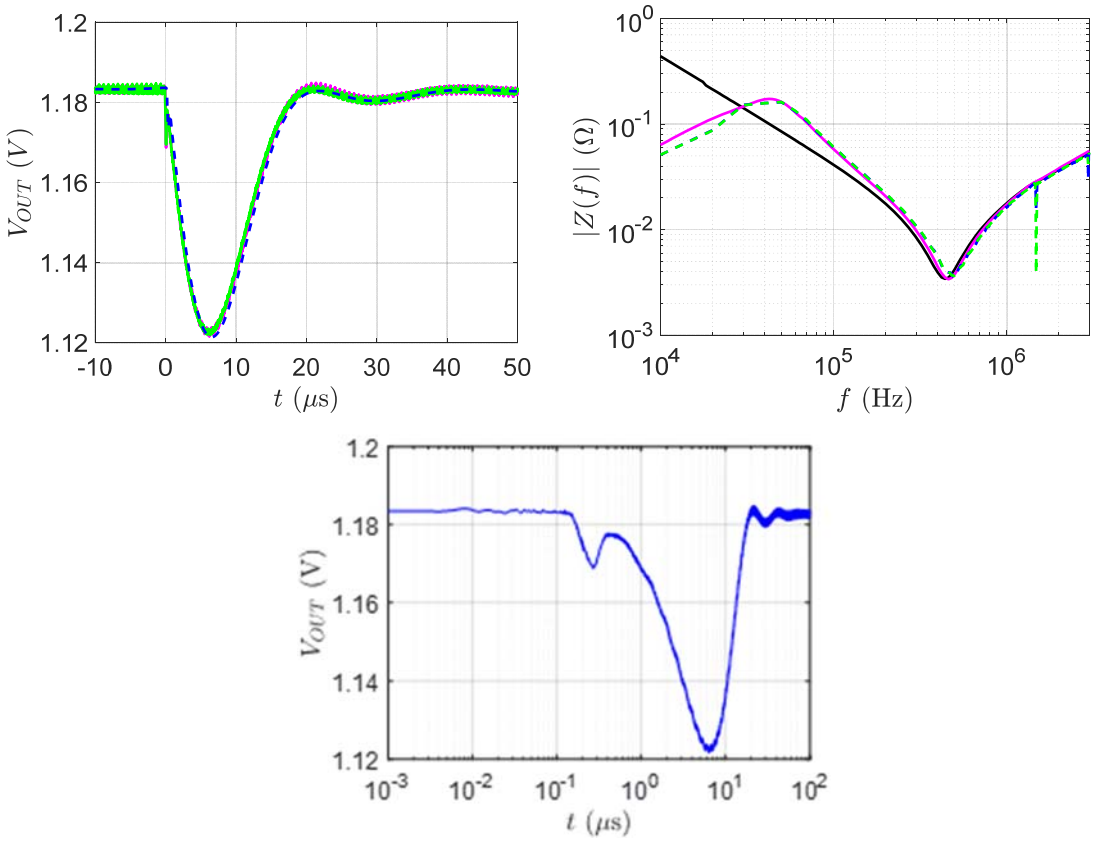

Figure 8: *Time-domain response (top left) and frequency-domain response (top right) of the board with 'good' connections from Figure 6. In each plot the solid line is the response directly measured; the dashed line is the equivalent response back-calculated from the other domain. Rising (load-release) responses are mirrored and overlayed. Bottom plot: data from top left, plotted on logarithmic horizontal scale.*

Figure 9 shows the worst case transient noise for 0.5A current steps, based on the Reverse Pulse Technique. On the left of *Figure 9.a*, the step response is shown with the peaks and valleys identified with their time stamp and value in the table on the right. In *Figure 9.b* the measured worst-case noise is shown in the time domain following the worst-case pattern excitation. The worst-case noise calculated simply from the peaks and valleys and the DC shift from *Figure 9.a* is 130 mVpp. The worst-case noise in *Figure 9.b* is 127 mVpp.

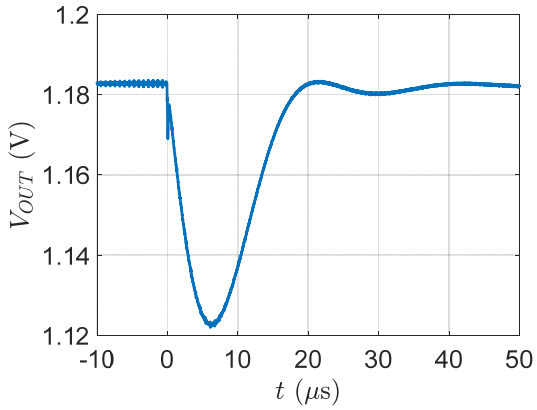

Time (us)	Vout (V)
0	1.184
6.143	1.123
21.37	1.184
29.82	1.179
42.99	1.184
50.04	1.182

Figure 9.a: *Step response for 0.5A step current (left) and the table of peak and valley voltages and time stamps (on the right) for the case shown in Figure 8.*

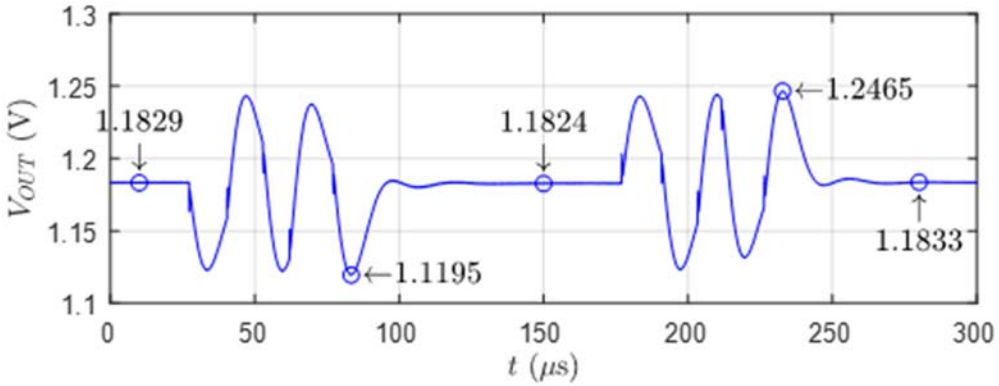

Figure 9.b: *Measured worst-case noise based on the excitation pattern derived from the table in Figure 9.a.*

The effect of nonlinearity is illustrated in *Figure 10*. Using the same setup and same DUT as for *Figures 8* and *9*, now the current step stimulus is increased to 1A. Note that this converter has a maximum current rating of 3A and there is only a single ceramic capacitor on the output. The current stimulus stepped from 1.0A to 2.0A. The figure shows the step responses for the rising and falling current steps overlaid and normalized to 0.5A. For sake of reference, the 0.5A step response from *Figure 8* is also included (purple trace).

Note that while the load-attack response (rising current step, blue trace) is effectively the same, the load release response (green trace) is significantly different. To make the visual comparison easier, not only the vertical values were scaled to the same 0.5A response, but the load-release response was also inverted and shifted. This scenario illustrates a case when due to nonlinearity the time and frequency domain data are not interchangeable.

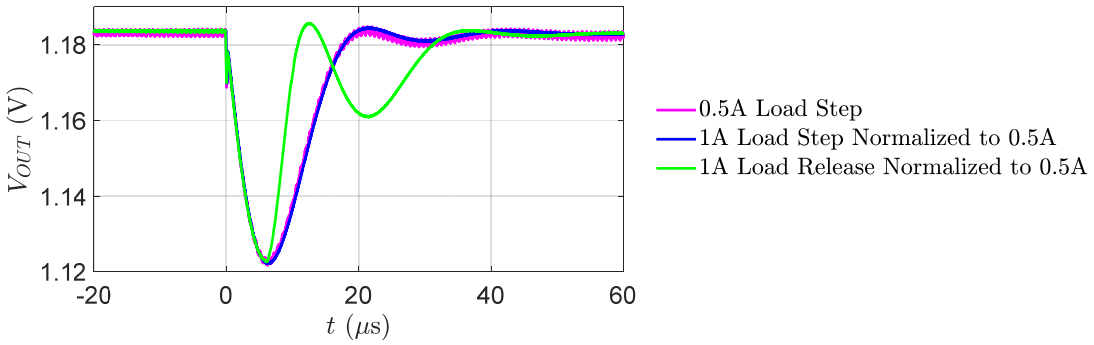

Figure 10: *Rising and falling step responses with 1.0A to 2.0A current steps.*

III. Transient response and impedance distortions

There are multiple ways how the impedance profile and transient response can become distorted. Some effects are real, some are just measurement artifacts. In this section we look at some of the major options.

III.1 Wrong connections

First we show how making the wrong choice about the connection to the DUT can create deep impedance notches. We use the DUT from *Figure 6* and use the wrong connections of the cable picking up the response, on top of a capacitor, marked by the red arrows. We can often see this suggestion in application notes of power converters, recommending to pick up the output signal across a capacitor body. While this makes the signal less noisy by attenuating any high-frequency ringing, this also creates a strong frequency dependence near the series resonance frequency of the capacitor (see *Figure 7*). For this illustration the source connection, both for the VNA and for the transient injector, was left at the base of the turrets where they were positioned for *Figure 8*. *Figure 11* shows the result in the frequency and time domains with 0dBm VNA source power and 0.5A current steps. The unrealistically deep notch could be accidentally confused with the low ESR of the output capacitor, which is clearly not the case; the deep notch is just a measurement artifact. Interestingly though the measurement artifact equally impacts the passive impedance as well, when we measure the DUT without input power. In that case we expect to measure only the impedance of the capacitor on the output. This is something to keep in mind for component characterizations as well. The step response waveforms look very similar for the good and bad connections (compare the top right plots of *Figure 8* and *11*), however, the step responses must be different because transforming them to the frequency domain reveals the same difference at the resonance frequency. The reason why we do not see a more dramatic difference in the step response (and in the worst-case noise) is because the 0.2-Ohm peak at 50 kHz dominates the response. The effect of the notch depth is more pronounced when the impedance is more flat, as will be shown next.

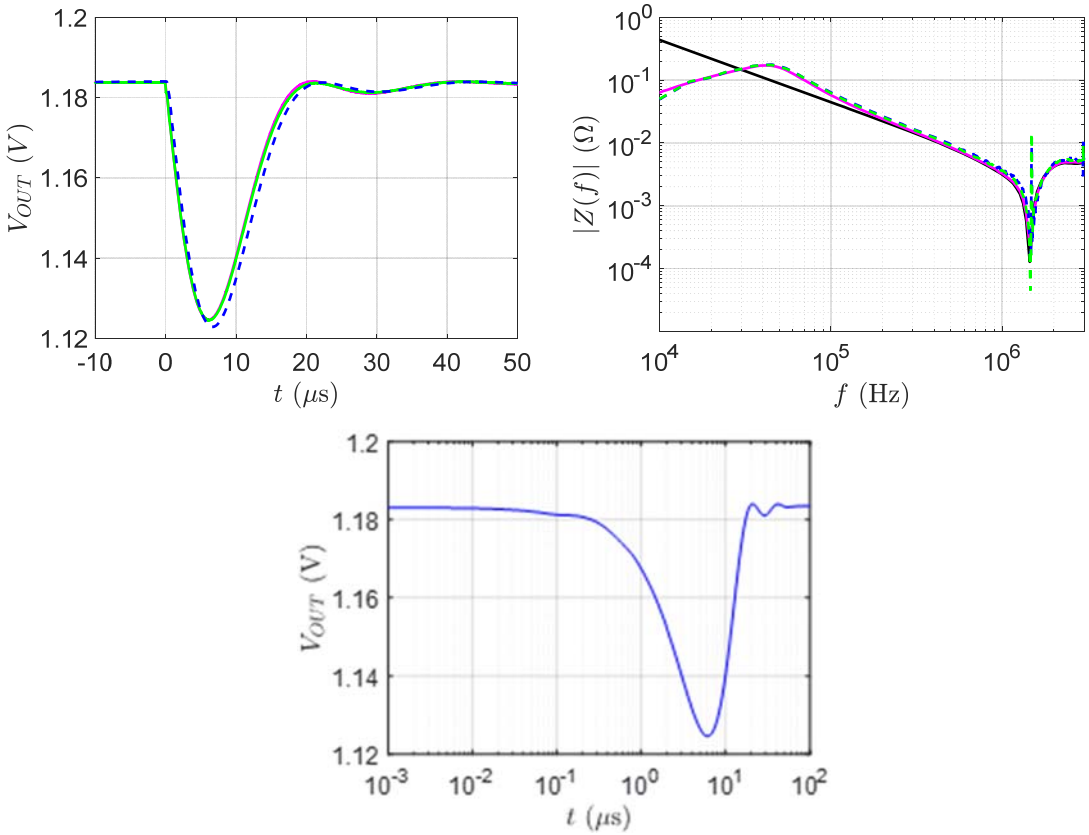

Figure 11: *Time-domain response (top left) and frequency-domain response (top right) of the board and connections from Figure 8. In each plot the solid line is the response directly measured; the dashed line is the equivalent response back-calculated from the other domain. Rising (load-release) responses are mirrored and overlayed. Bottom plot: same step response with logarithmic horizontal scale.*

With logarithmic horizontal scale the step response does not exhibit the initial dip that we see in *Figure 8*. However, by comparing the two step responses more closely, we notice that the main dip is also slightly different. With reference to *Figure 3*, we can note that we changed two things simultaneously: while the low-frequency portion below 1 MHz remained the same, the bad connection created not only a deeper dip, it also lowered the impedance above the series resonance frequency.

Our next DUT was another evaluation board with a single-phase DC-DC converter with 15A max DC current rating. Some early measurement data on this DUT was included in [8], reproduced here in *Figure 12*. The figure caption identifies the permutations. The evaluation board was just a couple of inch square in dimension, the output capacitors were all MLCC parts in a tight array on the top and bottom side of the board placed in between the converter output (location 1) and the board edge (location 2). This board was extensively re-measured in many different variations and connection options. In addition to the changes in the Gain-phase curve and output impedance due to the phase

shift created by the spatial plane resistance and output capacitors (as reported in [7]), another possible measurement connection error was also discovered. Often times in reworks, or in what-if measurements the sense-point connections are moved around to different locations by replacing PCB traces with twisted pairs.

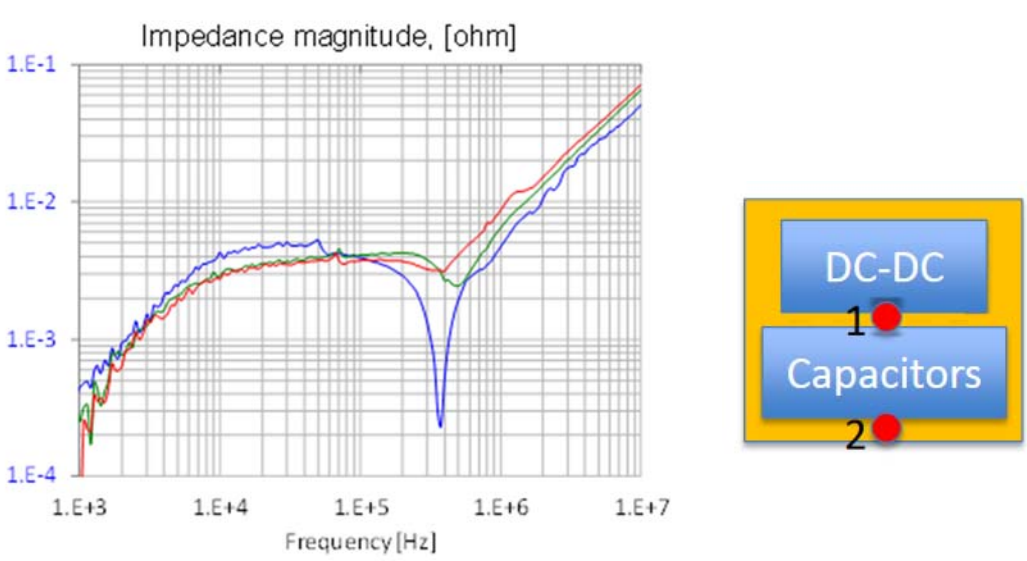

Figure 12: *Output impedance curves at three different combinations of sense-point and measurement locations (from [8]). Red curve: measure at point 1, sense connected to point 2. Green curve: measure at point 2, sense connected to point 1. Blue curve: measure at point 2, sense connected to point 2.*

From the large number of cases measured, here we show three scenarios (*Case A*, *Case B* and *Case C*), all three used connections along the edge of the evaluation board. The three connection geometries are shown by the sketch of *Figure 13*. The sketches show the edge of the board where the connections were made. The top side of the board had solid power plane, the back side was solid ground plane. The measurement cables were RG178 with approximately 200-mil open pigtail at the end for the soldered connections. The sense lines were carried to the edge of the board by a twisted pair. The two coaxial cables were left at the same connection points for the three variants, the only change was the location of sense wire connections. In *Case A* the twisted pair was solder to the planes in between the two coaxial cables. In *Case B* is was outside of the two coax cables, still close. In *Case C* the twisted pair was further removed by approximately one inch.

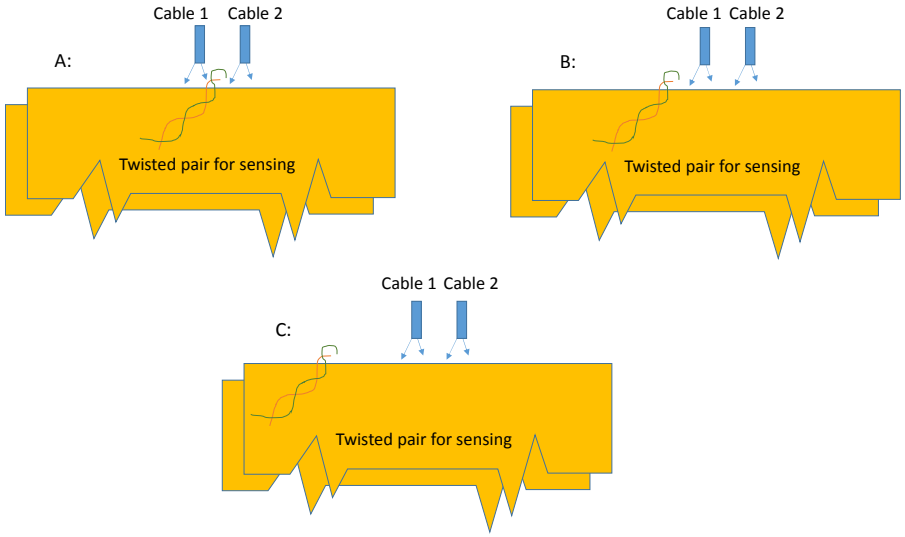

Figure 13: *Connection sketch on the board edge of the evaluation module.*

The impedance profile in *Case A* exhibited a sharp and deep notch as shown in *Figure 14.a*, pushing the impedance minimum to very low values.

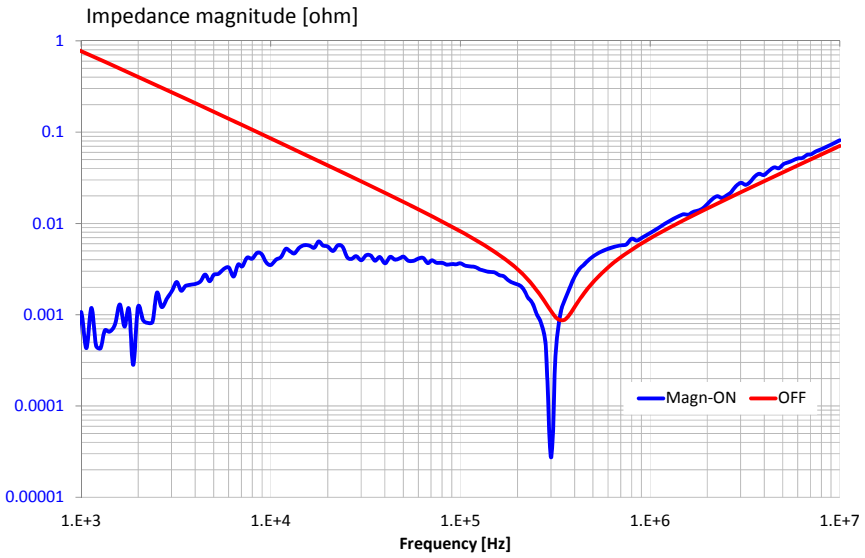

Figure 14.a*: Impedance profile with connection scheme A from Figure 13.*

The sharp notch was very sensitive to the cable orientation and loop size. By slightly rearranging the cables it became clear that the coupling between the sense-wire loop and injection-cable loop created the variation and there was minimal or no effect from the coupling to the response pickup cable loop. This is understandable, since the small loop at the end of the injection cable is terminated in virtual short, producing strong magnetic

field and magnetic coupling. The response pickup cable loop is terminated in 50 Ohms, the magnetic coupling from that loop is much smaller.

The Gain-phase plots were also measured, the phase margin was 72.5 degree, the data is shown in *Figure 14.b*.

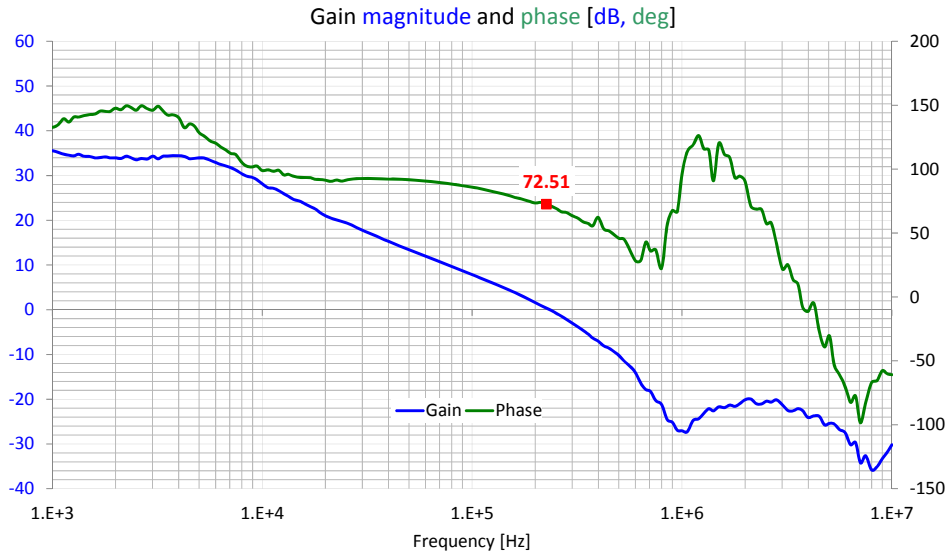

***Figure 14.b**: Gain-phase plot with connection scheme Case A from Figure 13.*

In *Case B* the impedance notch stays the same whether the converter feedback loop is ON or OFF. When the coupling was minimized by moving the sense wires further away from the injector cable loop, the impedance minimum was pushed up, creating a very shallow dip instead (see *Figure 16.a*). The measured Gain-phase plot indicates that the phase margin has not changed.

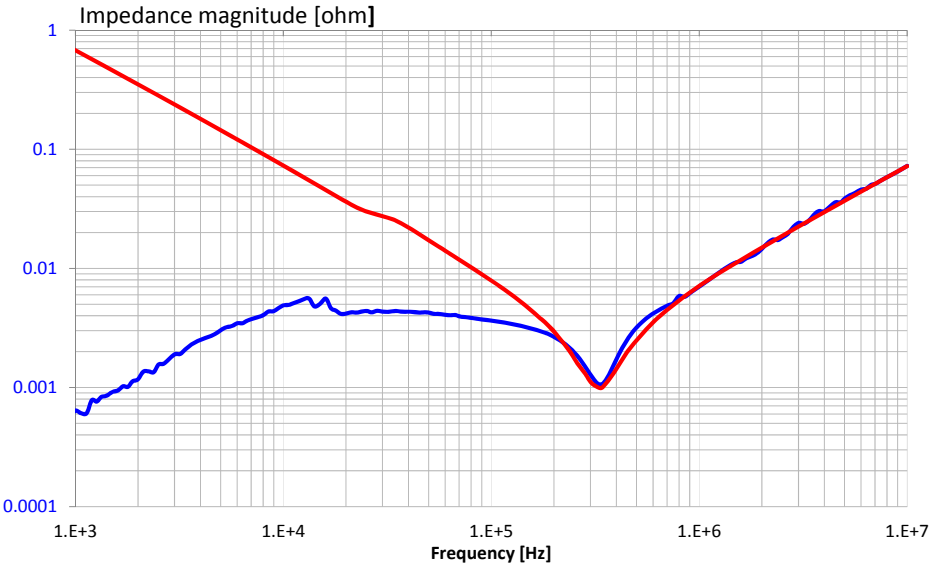

***Figure 15.a**: Impedance profile with connection scheme Case B from Figure 13.*

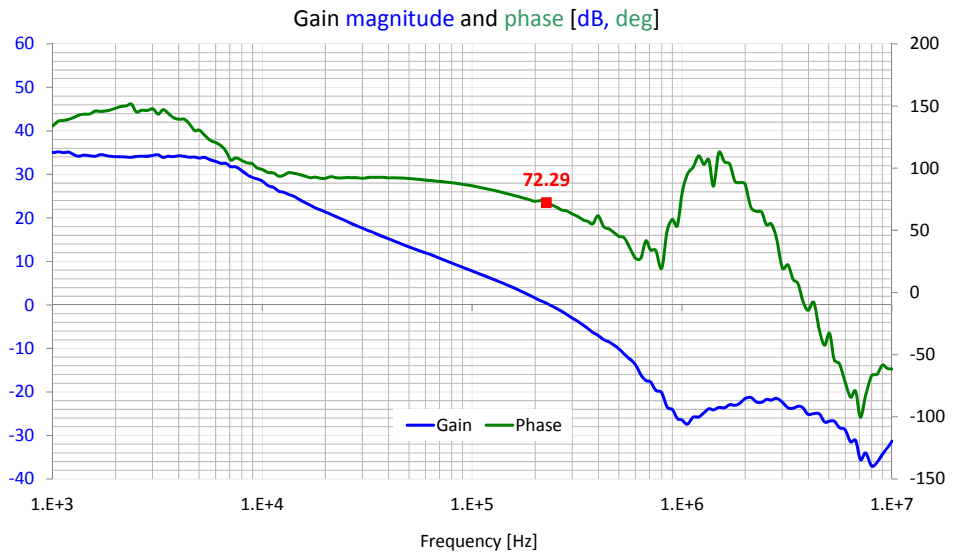

***Figure 15.b**: Gain-phase plot with connection scheme Case B from Figure 13.*

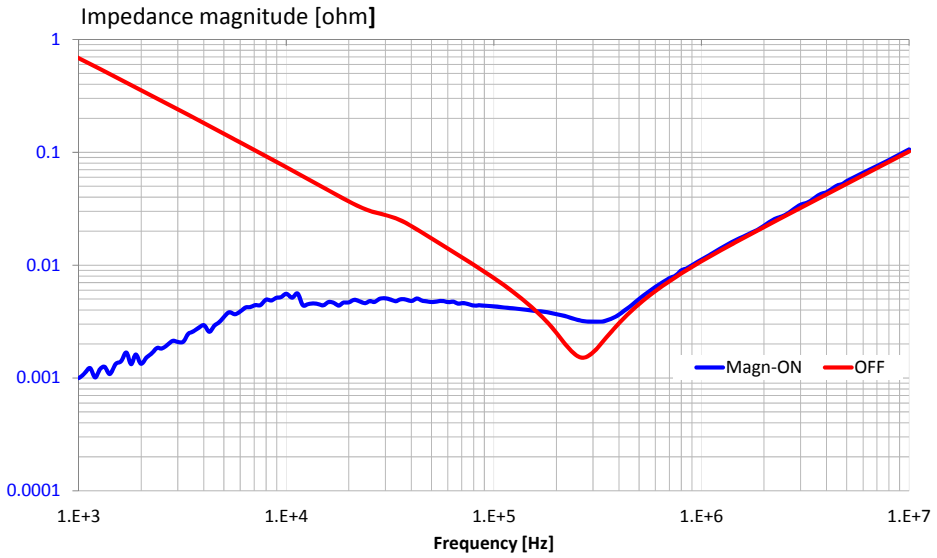

***Figure 16.a**: Impedance profile with connection scheme Case C from Figure 13.*

The measured Gain-phase plot shown in *Figure 16.b* tells us that the phase margin has not changed.

In all three cases the impedance measurement was performed at the same location and when the active control loop of the DC-DC converter was not working (OFF), the impedance reading was the same. It is important to note that while the measured Gain-phase plots are very similar and the phase margins are identical, any approximation of the phase margin based on the measured impedances will show different results. NISM [13] is not applicable in this case since there is no impedance peaking around the crossover

frequency. The Gain-phase curves back calculated from the ON and OFF impedances can be applied, but it will produce different results for the three cases.

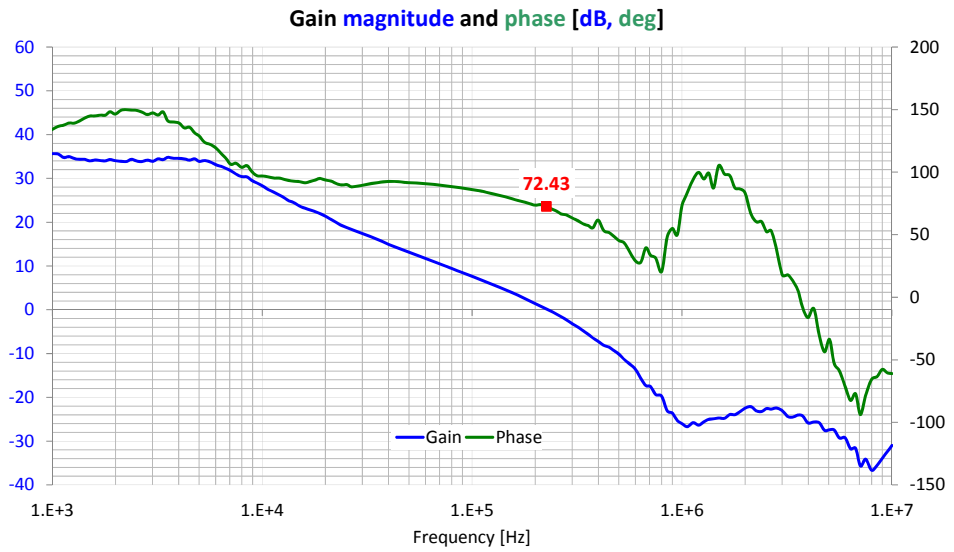

***Figure 16.b**: Gain-phase plot with connection scheme Case C from Figure 13.*

In the above illustrations the impedance profiles were relatively flat except the dip at the series resonance frequency of output capacitors. In such cases the variation of the depth of the impedance minimum will create a worst-case noise variation depicted in *Figure 4*, suggesting that the noise prediction can be 100s of percents in error.

The DC-DC converter in *Figures 12* through *16* also had only ceramic capacitors at the output, but the lower-loss power components and different active loop created a naturally low, flat impedance profile. When the converter does not inherently provides this, we can add extra capacitors across the output. We illustrate this option by taking the 3A-rated evaluation board from *Figure 6* and re-measured in the time and frequency domain with an extra 1500uF polymer tantalum capacitor across its output terminals. Here we show the results with 'bad' connections, connecting the response pickup cable across the output ceramic capacitor. *Figure 17* shows the impedance plot on the left and the step response on the right. The left graph shows multiple traces: step response from rising current stimulus measured in the time domain (solid line), inverted and shifted step response from falling current stimulus measured in the time domain (solid line) and step response back calculated from the measured impedance profile (dashed line). Note that all three responses line up well, in fact the step responses from the rising and falling current stimulus run on top of each other and only one color is seen. The VNA source power was 0dBm, the transient excitation current step was 0.5A, stepping from 1.5A to 2.0A. The added capacitance reduced the 55 kHz peak; it moved to 5 kHz with a peak magnitude of 50 mOhm instead of 0.2 Ohm. Note that though this impedance profile is not as flat that we showed in *Figures 12* through *16*, nevertheless the impact of notch-depth variation will create more uncertainty in the worst-case noise predictions.

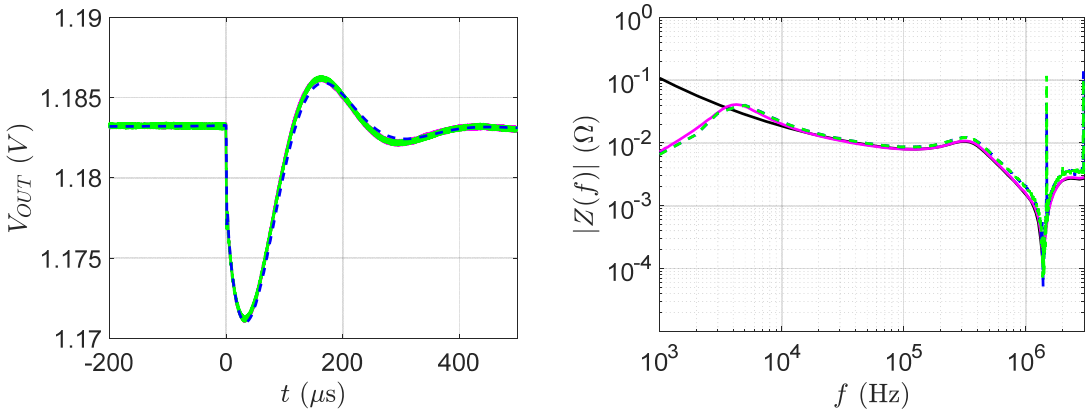

Figure 17: *Time-domain response (on the left) and frequency-domain response (on the right) of the board with 'bad' connections from Figure 6, with a 1500uF polymer tantalum added to its output. In each plot the solid line is the response directly measured; the dashed line is the equivalent response back-calculated from the other domain.*

Figure 18 shows the step response with the peaks and valleys identified in the table on the right together with the worst-case noise waveform.

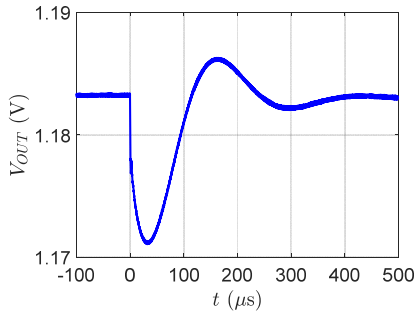

Time (us)	Vout (V)
0	1.183
33.63	1.171
164.4	1.185
303.7	1.183
433.0	1.182
565.0	1.1823

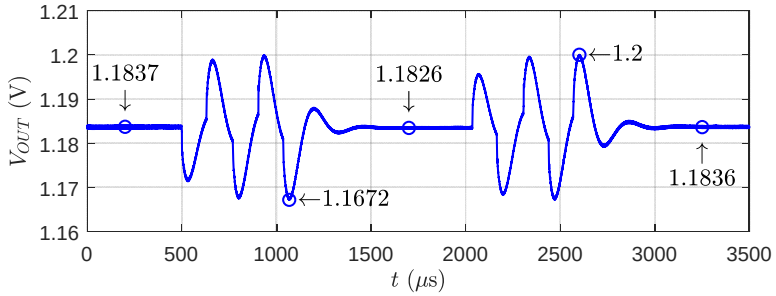

Figure 18: *Step response for 0.5A step current (top left) and the table of peak and valley voltages and time stamps (top right) for the case shown in Figure 17. Bottom plot: measured worst-case noise based on the excitation pattern derived from the table.*

III.2 Impedance profile variations by design choice

In the frequency range we have shown in the illustrations above (up to 10 MHz), we may tend to think about the power distribution network as a lumped network showing no spatial frequency-dependent effects. This is based on the assumption that in the PCB the

planes connect all components solidly, with negligible phase shift due to time-of-flight delay. It is true that our typical PCB geometries are much smaller than the wavelength. Also, the phase shift due to the time of flight delay (say 0.5 ns for a few inches of distance) is negligible compared to the period of our highest frequency that we look at here (100 ns for 10 MHz). However, the resistance of the planes in the series path together with low-impedance bypass capacitors in the parallel path can create appreciable filtering and phase shift even at low frequencies. The horizontal plane resistance will also increase under large components with multiple pins/vias perforating the planes. Some of these effects are illustrated in this section by using one supply rail on a memory card. The selected power rail has a DC-DC converter in between two ASICs with a plane shape connecting them. There were three different kind of capacitors on this rail: 470uF 9mOhm polymer tantalum capacitors near the source, 47uF MLCCs also near the source and a large number of 4.7uF 0402 MLCCs in the pinfields of the two ASICs, bridging adjacent power/ground pins on the back side of the board. An approximate sketch of the power rail is shown in *Figure 19*. Note that the ASICs are bigger than shown by the blue rectangles; those highlight only the pin array connected to this particular rail.

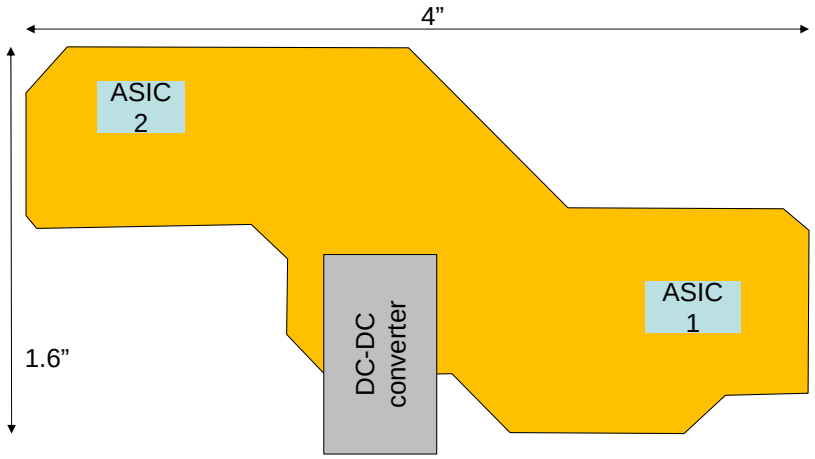

Figure 19: *Sketch of a power rail with one DC-DC converter and two ASICs. The blue rectangle shapes indicate the pin array of the ASIC connected to the plane shape.*

The DUT board had the ASICs unpopulated, giving the possibility to measure in the pinfield either on the top pads or on the bottom pads. This power rail had high and low side sense lines for the DC-DC converter with a choice through jumper resistors to connect the lines at different locations: Case 1) ASIC-1 pinfield, Case 2) ASIC-2 pinfield, Case 3) averaging the ASIC-1 and ASIC-2 pinfield sense voltages and Case 4) at the center of the plane shape near the DC-DC converter. Within the ASIC pinfields a number of connection permutations were tested; three of them are marked on the close-up sketch of *Figure 20*. Option 1 used two adjacent pad pairs on the top and it approximates self impedance in that area. Option 2 used two pad pairs on the top side a little bit further apart, which approximates a transfer impedance through part of the pinfield. For Option 1 the bridging capacitors on the back side of the board were moved out to nearby locations. Option 3 used connections on the bottom side (shown with top view), one

probe on a capacitor pad (the capacitor was moved away) and the other probe connected on top of the capacitor (the capacitor was left in place).

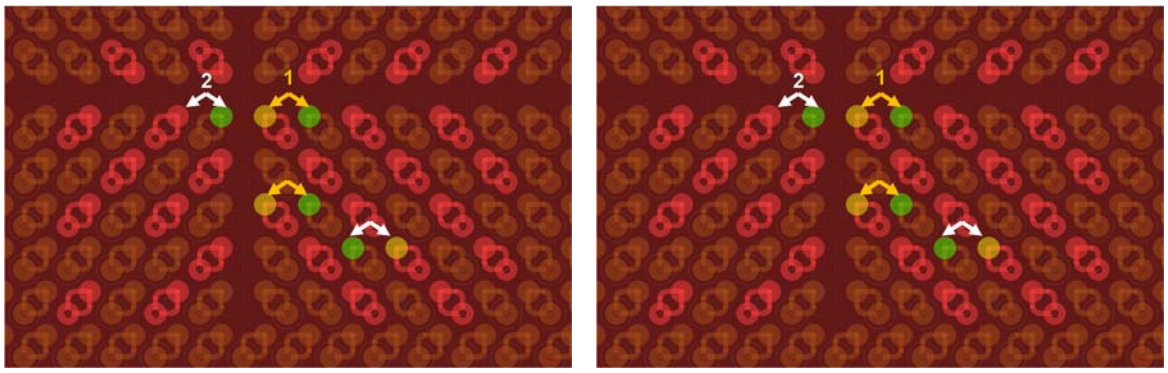

Figure 20: *Close-up sketch marking the three connection point pairs in the ASIC pinfield. Options 1 and 2 pads are on the top side, marked by yellow and white arrows, respectively. Option 3 pads are on the bottom (shown as top view).*

The layout is symmetric and identical in the pinfields of the two ASICs. This was first verified by measuring all connection permutations in the two pinfields.

For the measurements ferrite-covered flexible coax cables were used with semirigid coax probes. When bottom pads were involved, to avoid the coloring of the measurement results, the bridging 4.7uF MLCC was removed from that site and was soldered on top of the adjacent capacitor. Due to the large number of bridging capacitors in each of the ASIC pinfields, this preserved the total capacitance in the area but otherwise had a negligible effect on the impedance profile. A top-side connection is shown in *Figure 21*. The close-up photo of the two semirigid probes connected to the DUT.

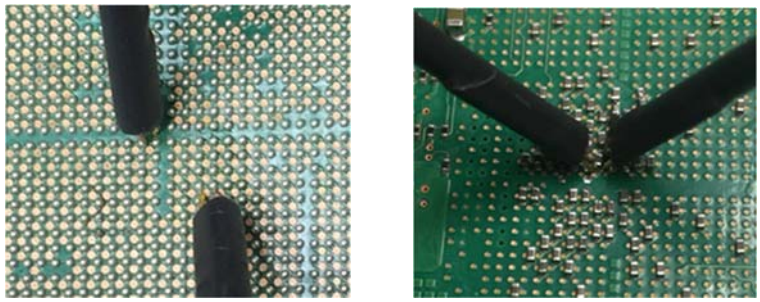

Figure 21: *Semirigid probes landed in the ASIC pinfield on the top side of the board (left photo) and bottom side of the board (right photo).*

The following three figures show the measured impedance values for the three options, each summary charts containing data for the four cases described above for the sense point connection.

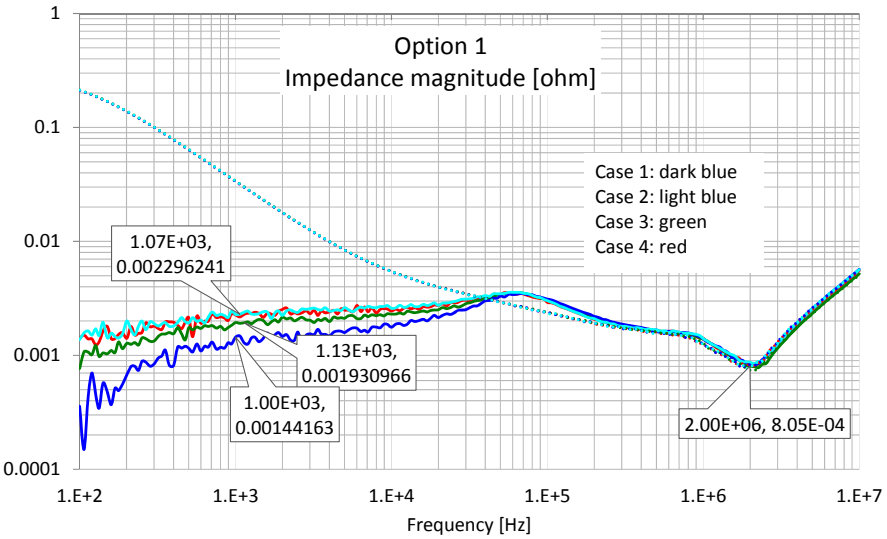

***Figure 22**: Impedance magnitude for the four sense-point connection cases for Option 1.*

Figure 22 refers to Option 1, when we measure at ASIC-1 across adjacent pad pairs, with the bridging capacitors removed on the bottom. The impedance minimum at 2 MHz is 0.8 mOhm, which is the cumulative effect of the bridging ceramic capacitors on the back. Above 100 kHz the sense-point connection does not make any difference. Below 100 kHz the impedance curves deviate and we get the lowest impedance when the sense point connection is closest to the measurement point. The difference is the spatial plane resistance among the various components.

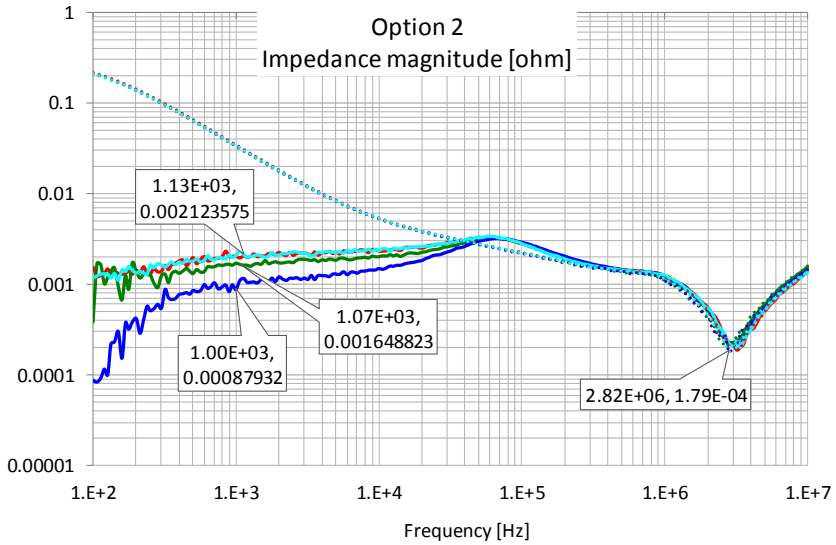

***Figure 23**: Impedance magnitude for the four sense-point connection cases for Option 2.*

Figure 23 shows Option 2, which is a transfer impedance measured on the top, across part of the pinfield's diagonal. The cumulative series resonance frequency of the bridging capacitors has shifted from 2 MHz to 2.82 MHz and the impedance dip dropped to 0.18 mOhm. Below 100 kHz we see a separation of the impedance curves similar to *Figure 22*.

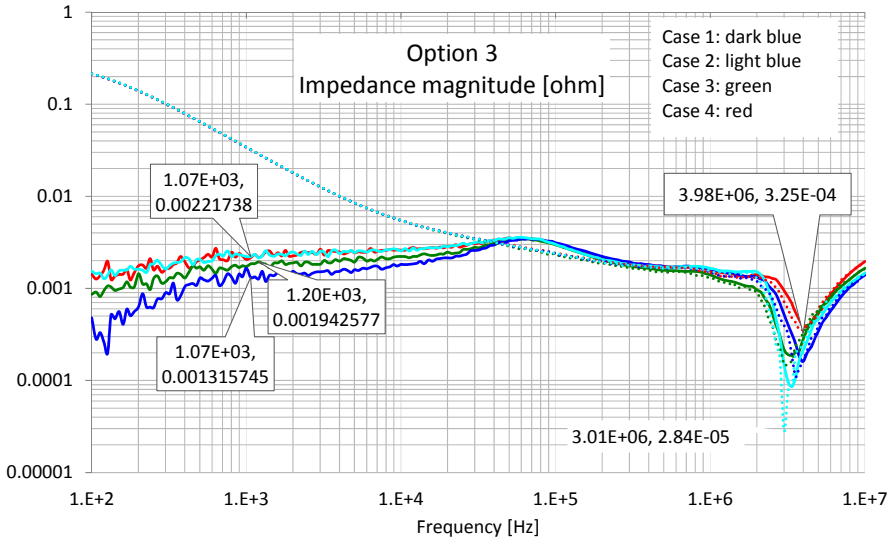

***Figure 24**: Impedance magnitude for the four sense-point connection cases for Option 3.*

Figure 24 refers to Option 3, which represents a 'bad', not preferred connection, when at least one of the connections is across a capacitor. Note that in this case the impedance minimum at the series resonance of the bridging capacitors is very sensitive to the sense-point location. The impedance minimum has a spread of 28 microOhm to 325 microOhm. The low-frequency variation is similar to the variation in the other two options.

Conclusions

In this paper we documented that DC-DC converters with just one ceramic capacitor on the output can operate linearly for light and moderate transient loads to a degree that rising and falling step responses when mirrored and superimposed will line up within measurement error. It was also illustrated on multiple scenarios that in such cases time and frequency domain measurements are equivalent, namely time domain step response can be converted into frequency domain impedance profile and vice versa and the results overlap within measurement error. For sufficiently linear (and time invariant) systems we can apply the *Reverse Pulse Technique* to calculate the worst-case transient noise.

It was shown that for any given impedance limit the worst-case transient noise is minimized by minimizing the deviation of the impedance profile from a flat (resistive) impedance. It was also shown that driving minima (notches) into the impedance profile will also increase the worst-case transient noise even if the maximum bound of the impedance profile stays the same. A single dip increases the worst-case noise by up to 200% as the max/min ratio of the impedance profile gets large. If impedance dips are separated in frequency enough that in between the dips the impedance comes up to the maximum bound, the noise increase will be additive and each impedance dip further increases the worst-case noise.

Impedance dips can occur in PDN impedance profiles either by design (intentionally) or by measurement artifacts (unintentionally). The paper showed two scenarios when an impedance dip becomes seemingly deeper due to measurement error. One scenario was due to the filtering effect when a measurement connection happens across a low-ESR capacitor (MLCC). Another scenario for unintentional change (increase or decrease) of impedance dip value is when converter sense lines get coupled to active measurement cables. This can happen during tests or bringup or when open pigtail connections of measurement cables get close to surface traces carrying sense signals. In either case the predicted worst-case noise will have an error due to the measurement artifact, and it happens whether we measure in the time or frequency domain. Variation of impedance spatially with fixed sense-point connections and when sense points are moved can occur when the resistance of connecting horizontal metal structure's resistance is not negligible compared to the impedance of capacitors on the rail. Low output impedance forced by high-gain converter loops are counter-balanced by series plane resistance: at locations close to the sense connections the control loop dictates the impedance and the location of low-frequency bulk capacitors will have little influence on the impedance. Further away of the sense points the impedance will also be influenced by the plane resistance and bulk-capacitor impedance. In perforated plane areas in larger pin arrays with bringing capacitors on the back side of the board the impedance variation can be substantial.

References

[1] Larry D. Smith, Raymond E. Anderson, Douglas W. Forehand, Thomas J. Pelc, and Tanmoy Roy, ‘‘Power distribution system design methodology and capacitor selection for modern CMOS technology’’, IEEE Transactions on Advanced Packaging, vol. 22, no. 3, pp. 284-291, Aug.1999.

[2] Drabkin, et al, “Aperiodic Resonant Excitation of Microprocessor power Distribution Systems and the Reverse Pulse Technique,” Proceedings of EPEP 2002, p. 175.

[3] “Bypass Capacitor Selection Based on Time Domain and Frequency Domain Performances,” in TecForum TF-MP3 "Comparison of Power Distribution Network Methods" available at http://www.electrical-integrity.com/Paper_download_files/DC06_TF-MP3_SUN.pdf

[4] “Systematic Estimation of Worst-Case PDN Noise: Target Impedance and Rogue Waves,” QuietPower column, 2015, available at http://www.electrical-integrity.com/Quietpower_files/QuietPower-34.pdf

[5] Smith, Bogatin: Principles of Power Integrity for PDN Design. Prentice Hall, 2017.

[7] “Impact of Regulator Sense-point Location on PDN Response,” DesignCon 2015, Santa Clara, CA, January 27 - 30, 2015

[8] “Overview and Comparison of Power Converter Stability Metrics,” DesignCon 2017, Santa Clara, CA, January 31 - February 2, 2017

[9] Non-invasive Stability Measurement, https://www.picotest.com/non-invasive-stability-measurement.html

[10] LM20143 evaluation board: www.ti.com/lit/ug/snva277b/snva277b.pdf

[11] “Connections to Evaluation Boards,” http://www.electrical-integrity.com/Quietpower_files/QuietPower-38.pdf

[12] “Do not measure PDN noise across capacitors!,” http://www.electrical-integrity.com/Quietpower_files/QuietPower-23.pdf

[13] Non-invasive Stability Measurement, https://www.picotest.com/products_NISM_software.html