DesignCon 2016

Electrical and Thermal Consequences of Non-Flat Impedance Profiles

Jae Young Choi, Oracle
Jae.young.choi@oracle.com

Ethan Koether, Oracle
Ethan.koether@oracle.com

Istvan Novak, Oracle
Istvan.novak@oracle.com

Disclaimer: This presentation does not constitute as an endorsement for any specific product, service, company, or solution.

Abstract

It has been long known that high-Q antiresonances between capacitors in power distribution networks (PDN) create harmful impedance peaks, which can impair the delivery of clean power, signal integrity and electromagnetic compatibility. The patterns of creating worst-case or almost-worst-case time-domain response has also gain interest in recent years. With the constant rise of power density in modern electronics, the thermal aspects of antiresonance may need to be considered as well. Capacitors have specified maximum allowed AC current flowing through them so that the internal temperature rise is within safe limits. This paper will analyze these aspects and suggest design guidelines.

Author(s) Biography

Jae Young Choi is a Senior Hardware Engineer at Oracle Corporation. He is engaged in the design and characterization of high-speed signal interconnects and power distribution networks for high-performance server systems. He received his Ph.D. in Electrical and Computer Engineering from the Georgia Institute of Technology.

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.

Istvan Novak is a Senior Principle Engineer at Oracle. Besides signal integrity design of high-speed serial and parallel buses, he is 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

It has been long known that high-Q antiresonances among any power distribution network elements create harmful impedance peaks, which can impair not only the delivery of clean power, but also signal integrity and electromagnetic compatibility [1], [2], [3]. A very typical scenario is between two different-value capacitors or capacitor banks: at a frequency where one capacitor is below its series resonance frequency and therefore it behaves as a capacitive (C) reactance whereas another capacitor is already beyond its series resonance frequency and therefore behaves as an inductance (L). Note that as it was pointed out in [4] and [5], the antiresonance peak depends on the ratio of imaginary and real parts of impedances. If the Z_0 equivalent characteristic impedance calculated from L and C, Z_0 = sqrt(L/C) is bigger than the real parts (in case of capacitors: ESRs) of the paralleled impedances, a peak is developed in the resulting impedance profile. The impedance profile becomes flat with no peaking if the equivalent characteristic impedance matches the real part [5]. *Figure 1* illustrates these scenarios.

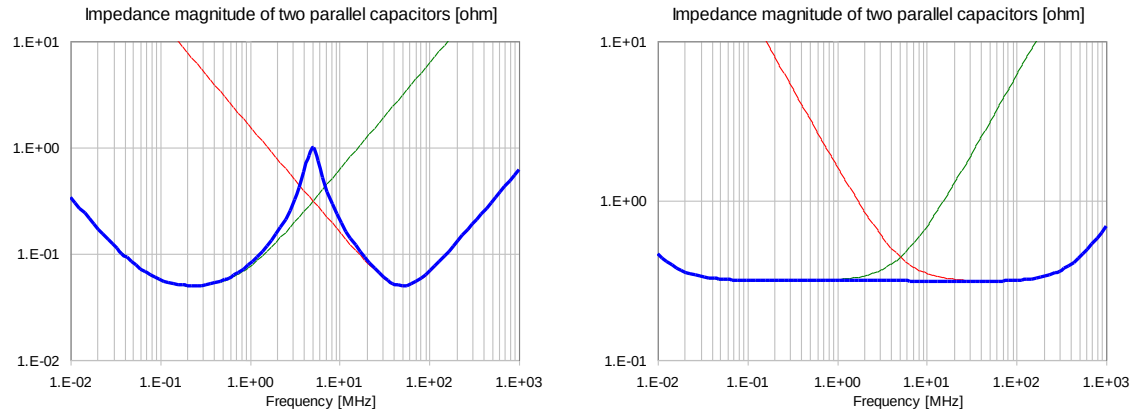

Figure 1: *Illustration of potential antiresonance peaking between different-value capacitors. Large peaking occurs with Z_0 > ESR (on the left), whereas there is no peaking when Z_0 = ESR (on the right).*

Note that the antiresonance peak can develop from reactances originated from any component or structure, not only capacitors. Typical such cases at low frequencies are for instance the inductance created by the band limitation of DC-DC converter control loops. At high frequencies the static capacitance of PCB planes can create antiresonances with the cumulative inductance of bypass capacitors.

With the constant rise of power density in modern electronics, the thermal aspects of antiresonance may need to be considered as well. Capacitors have a specified maximum allowed AC current flowing through them so that the internal temperature rise is within safe limits. A typical situation when this matters is the input capacitors of switching DC-DC converters (see *Figure 2*): in step-down buck converters the input ripple current fluctuates between zero and the full output current of the converter plus half of the inductor ripple current.

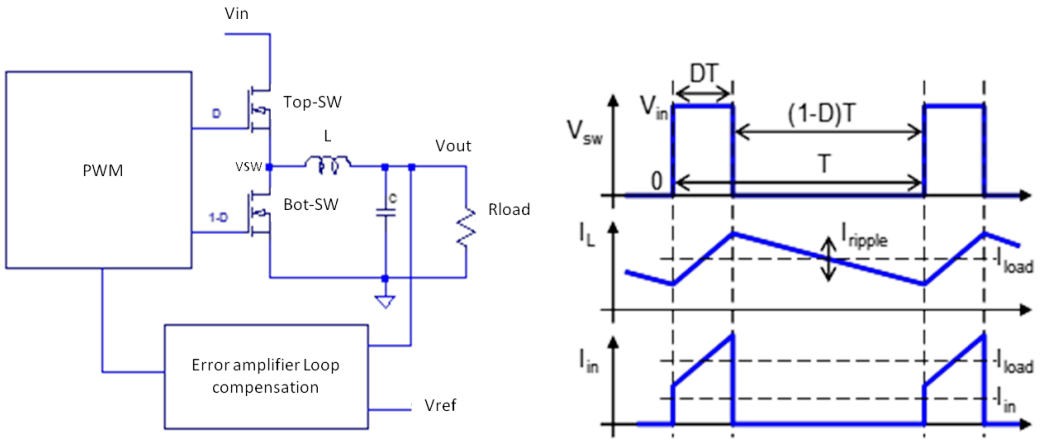

Figure 2: *Simplified non-isolated buck converter schematics on the left, typical current waveforms on the right. Note that the I_{in} current, which has to return through the input capacitors, has full swings between zero and the load current plus half of the inductor ripple current.*

When different valued capacitors are paralleled, either in front of the same converter, isolated by some intentional or unintentional inductance from the rest of the input supply rail, or by having different DC-DC converter inputs tightly connected to the same input rail, at the antiresonance frequency there is a current multiplication, increasing the dissipation and therefore also increasing the temperature in each capacitor. *Figure 3* illustrates this scenario. The figure shows the simulated performance of two capacitors in parallel, which may represent a bulk (C_1) and a ceramic (C_2) capacitor at the input of a DC-DC converter. To illustrate the point, we use two different sets of values, but we change only L_1 and C_2; the rest of the parameters remain unchanged. With C_1=470uF, R_1=5mohm, L_1=10nH and C_2=10uF, R_2=5mohm, L_2=1nH the resulting parallel circuit has a Q of 3.16. The plot shows the impedance magnitude, V[n001], and the magnitudes of the currents in the capacitors in the 10kHz to 10MHz frequency range. Below 100kHz most of the current is carried by C_1 and very little current flows in C_2. Above 1MHz most of the current flows in C_2. At the 500kHz resonance frequency both capacitors carry 10dB (3.16x) higher current. For the second scenario we change for instance L_1 and C_2 to reduce Q. With C_1=470uF, R_1=5mohm, L_1=3.16nH and C_2=31.6uF, R_2=5mohm, L_2=1nH the resulting parallel circuit has a Q of 1. The plot indicates that we still have noticeable peaking in the impedance profile and a smaller, but still noticeable peaking in the currents. As Q reduces further, the peaking disappears and there is no multiplication in the current.

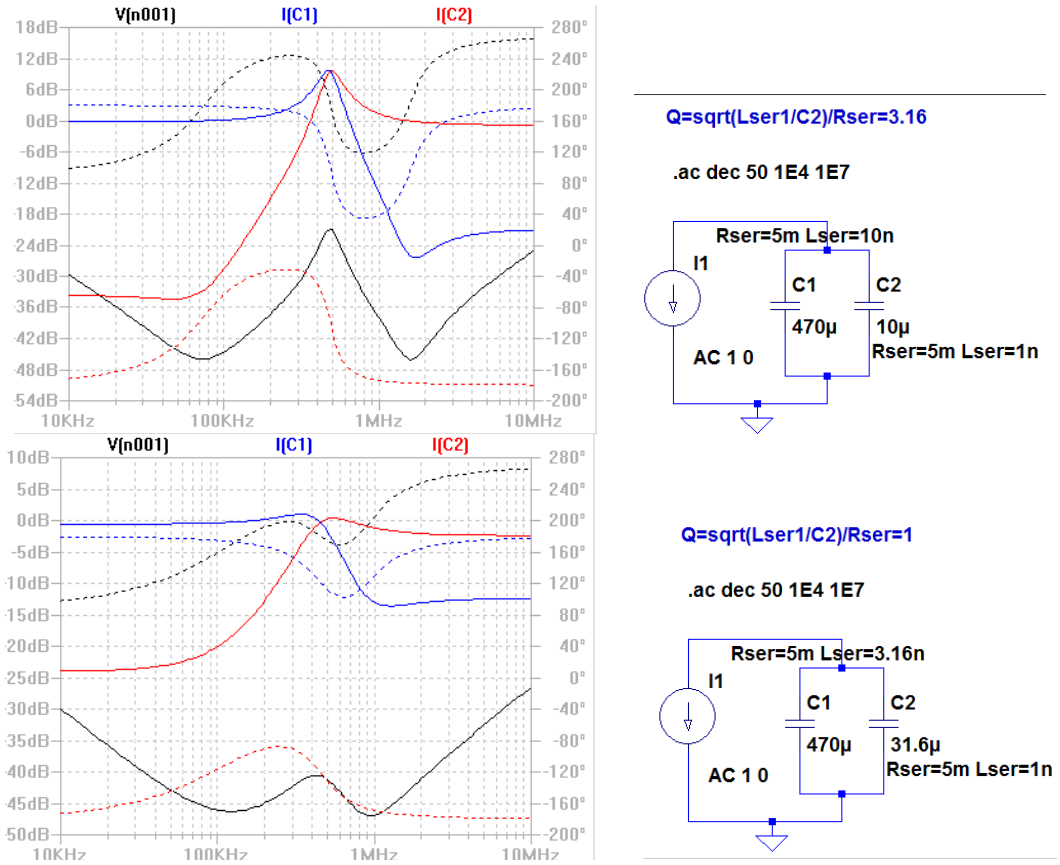

Figure 3: *Current multiplication at parallel resonance.*

The patterns of creating worst-case or almost-worst-case time-domain response has also gained interest in recent years [4]. The peak value of step response, the response to a repetitive excitation at a resonance peak as well as the absolute worst-case time-domain response are all different. For a power distribution network that can be approximated with a linear network, these are all related to and can be calculated from the step response of the network [6], which in turn can be calculated from the impedance profile. *Figure 4* and *Figure 5* explain the process.

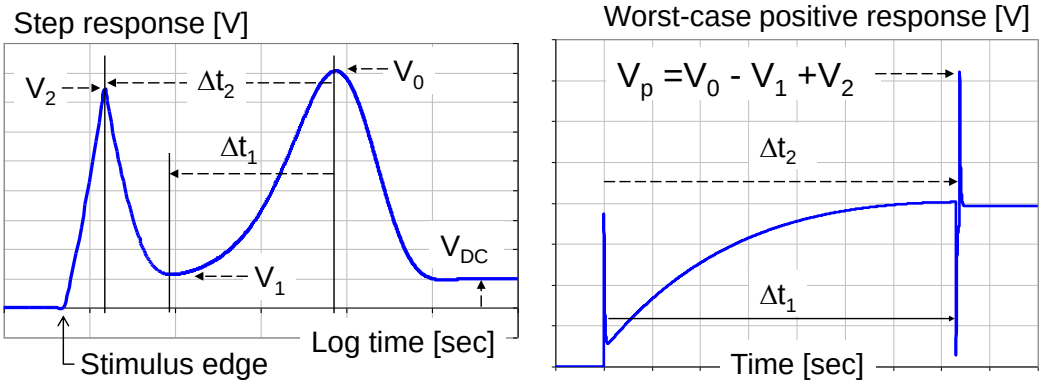

Figure 4: *Calculating the worst-case transient from the Step Response of the network.*

Assuming that the power distribution network is linear, its corresponding network functions in the time and frequency domain are Fourier pairs. The swept-frequency response function in the frequency domain is a Fourier pair of the Impulse Response (response to a Dirac-delta excitation) in the time domain. In a typical time-domain approach we assume that the PDN is hit by a current load step and we look at the resulting transient voltage at that location, which is called the Step Response. The Impulse Response is the time derivative of the Step Response, or vice versa, the Step Response is the time integral of the Impulse Response. This means that if we take the integral of the inverse Fourier transform of a simulated or measured impedance profile, it gives us the Step Response.

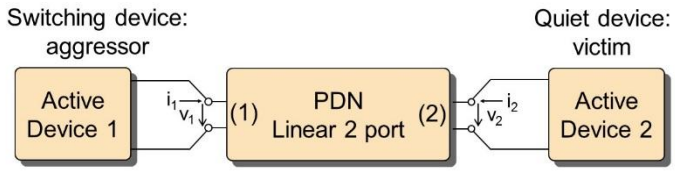

Transfer function for self-inflicted noise: $v_1/i_1 = Z_{11}$
Transfer function for propagated noise: $v_2/i_1 = Z_{21}$

$$h(t) = IFFT\{Z(f)\} \qquad s(t) = \int h(t)dt$$

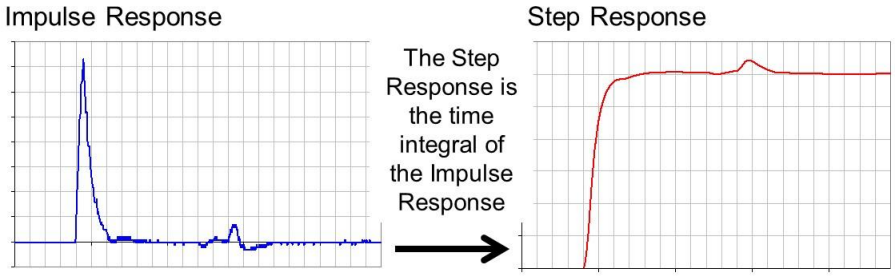

***Figure 5**: Calculating the Step Response from the impedance profile.*

II. Thermal impact of antiresonance

We will show the effect on a filter circuit that we may typically find at the input of DC-DC buck converters. *Figure 6* shows the equivalent schematics. The filter has a series power inductor or lossy ferrite, and in the parallel path a bulk capacitor with ceramic capacitor(s). The filter may serve multiple purposes, among others to prevent high-frequency burst noise from getting out back to the main power rail. Unless we can afford large filter components, the inductor will have ferromagnetic material, the ceramic capacitors will have ferroelectric material, both exhibiting potentially strong dependence on DC bias conditions and/or temperature [7] and therefore it is very important to consider the electrical parameters of the components under actual bias conditions:

circuits, which may have no peaking with no bias, might develop significant peaking when full DC voltage and DC current bias conditions are applied.

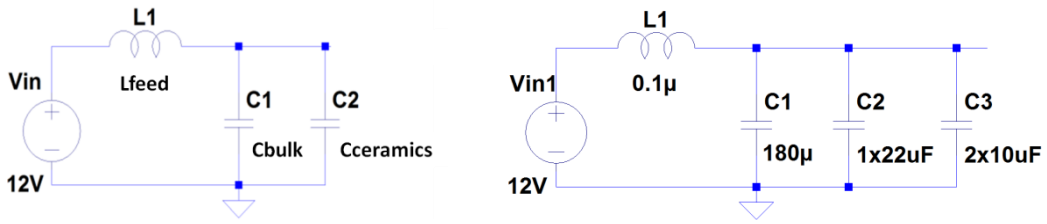

Figure 6: *Simplified schematics of a typical filter circuit at the input of a buck DC-DC converter (left) and one actual implementation (right).*

When we characterize the filter circuit, we need to remember to apply the correct DC voltage and DC current to bias the components. This is especially important when different kinds of capacitors are paralleled up, because as opposed to ceramic capacitors, electrolytic and tantalum capacitors exhibit little or no voltage dependence. The ESR and ESL of all capacitors are relatively independent of DC bias voltage. Under the working operating conditions all the above may result in a bigger antiresonance peaking between the bulk and ceramic capacitors compared to what we see without bias. *Figure 7* shows the measured impedance profile of an input filter that we will use as an example for the later figures. Note that while there is only a small antiresonance between the bulk and ceramic capacitors with no bias, under the normal operating conditions, there is much more pronounced peaking.

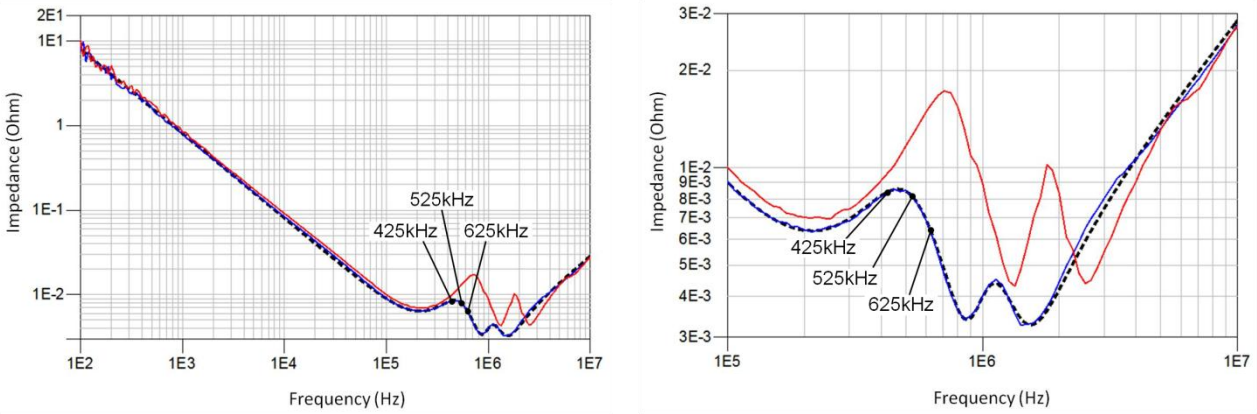

Figure 7a: *Measured impedance profile and model correlation of a DC-DC converter input filter with (red trace) and without (blue trace) DC bias. The black dashed line shows the impedance magnitude of the simulation model without DC bias. Left: full 100Hz – 10MHz frequency scale; right: zoomed 100kHz – 10MHz frequency scale. Note the frequency markers at the three switching frequencies used for later measurements.*

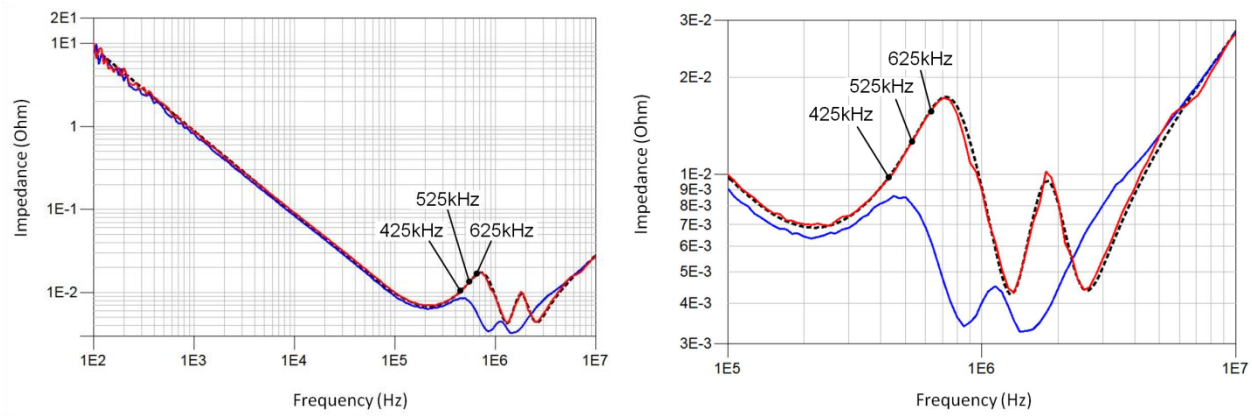

Figure 7b: *Measured impedance profile and model correlation of a DC-DC converter input filter with (red trace) and without (blue trace) DC bias. The black dashed line shows the impedance magnitude of the simulation model with 12V DC bias. Left: full 100Hz – 10MHz frequency scale; right: zoomed 100kHz – 10MHz frequency scale. Note the frequency markers at the three switching frequencies used for later measurements.*

Table I shows the actual model parameters used for *Figure 7* and for the subsequent simulations. Note that values referring to C_3 in the table are for the two 10uF ceramic capacitors in parallel.

DC Bias	Cc1	Lc1	Rc1	Cc2	Lc2	Rc2	Cc3	Lc3	Rc3
0V	167uF	2.7nH	7mOhm	15.16uF	2.2nH	4.8mOhm	15.6uF	0.74nH	4mOhm
12V	167uF	2.7nH	7mOhm	7.78uF	1.9nH	4.9mOhm	5.3uF	0.77nH	5mOhm

Table I*: Simulation model values.*

The DC-DC converter behind this filter was a single-phase non-isolated 12V to 1V buck converter, producing up to 30A DC current. The converter was tuned over a range of frequencies such that the switching frequency and its harmonics span over some of the filter resonances. *Figure 8* shows the simulated current waveforms that the input filter capacitors have to handle. The input current was simulated with 30A DC load current and 10A inductor ripple current at 425kHz. As the switching frequency increases, proportionally the ripple current in the output inductor goes down.

At all three frequencies, the plots on the left show the waveforms with no DC bias; the waveforms on the right show the current waveforms with the nominal 12V DC bias. From these waveforms we can calculate the rms current separately in each of the capacitors.

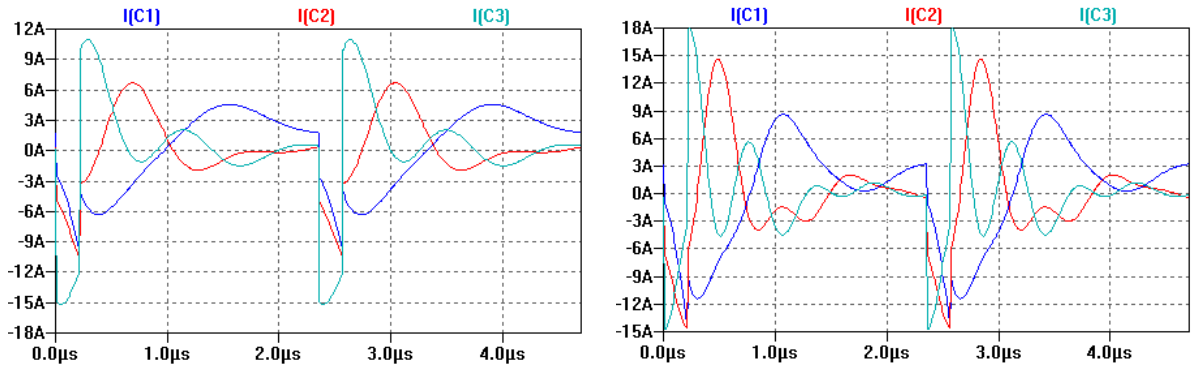

Figure 8a: *Simulated current waveforms flowing through the capacitors at 425kHz. Left: no DC bias; right: with 12V DC bias.*

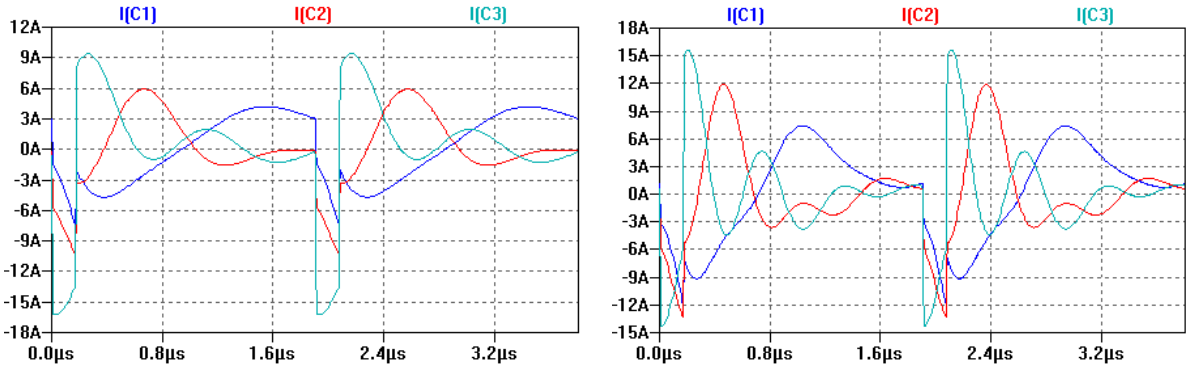

Figure 8b: *Simulated current waveforms flowing through the capacitors at 525kHz. Left: no DC bias; right: with 12V DC bias.*

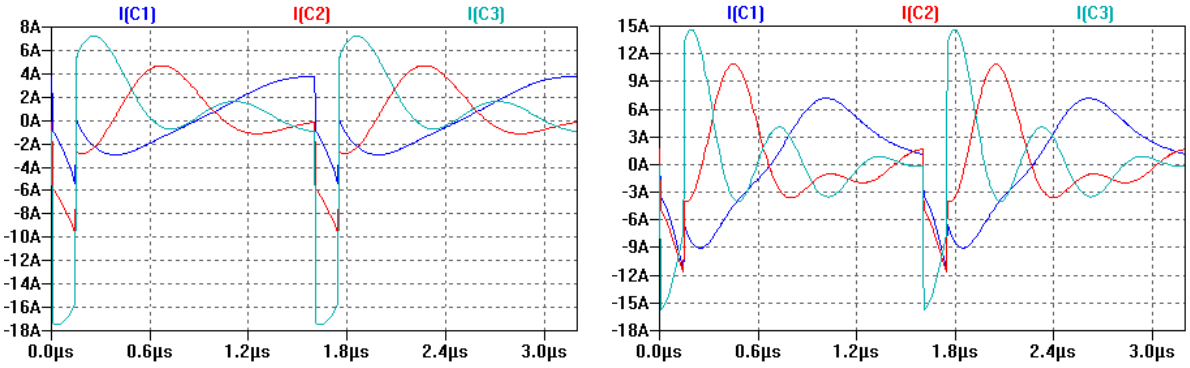

Figure 8c: *Simulated current waveforms flowing through the capacitors at 625kHz. Left: no DC bias; right: with 12V DC bias.*

The component values in Table *I* were fitted to the measured impedance of the input filter. The actual values in volume production, and assuming components from alternate sources, may show a considerable spread. Using a given case size and connection

geometry for each capacitor, the least variation is expected from the inductance associated with the components. With a few exceptions, the ESR of capacitors is not guaranteed as a range by the vendors: tantalum and polymer electrolytic capacitors specify only the maximum ESR; for ceramic capacitors only typical values are given. Under large DC bias power-filter applications, the biggest spread is expected in capacitance: DC and AC bias curves are provided as typical values only and they show large variations across vendors [7]. The temperature dependence of capacitance is also specified only as a range; the shape of curves is not specified. Lastly, the frequency of interest is high enough that the loss of effective capacitance due to dielectric loss and internal constructions will be significant; it can amount to an additional 40-60% of the capacitance already diminished by other factors. *Table II* shows the value ranges that were assumed for the simulations assessing the impact of tolerances under DC bias conditions. The impedance curves, the time-domain capacitor currents for the three operating frequencies as well as the calculated dissipation values are shown in *Figures 8d*, *8e* and *8f*, respectively. Note that the red and yellow lines are with the parameter combinations from *Table II*, which create the biggest first and second peak, respectively.

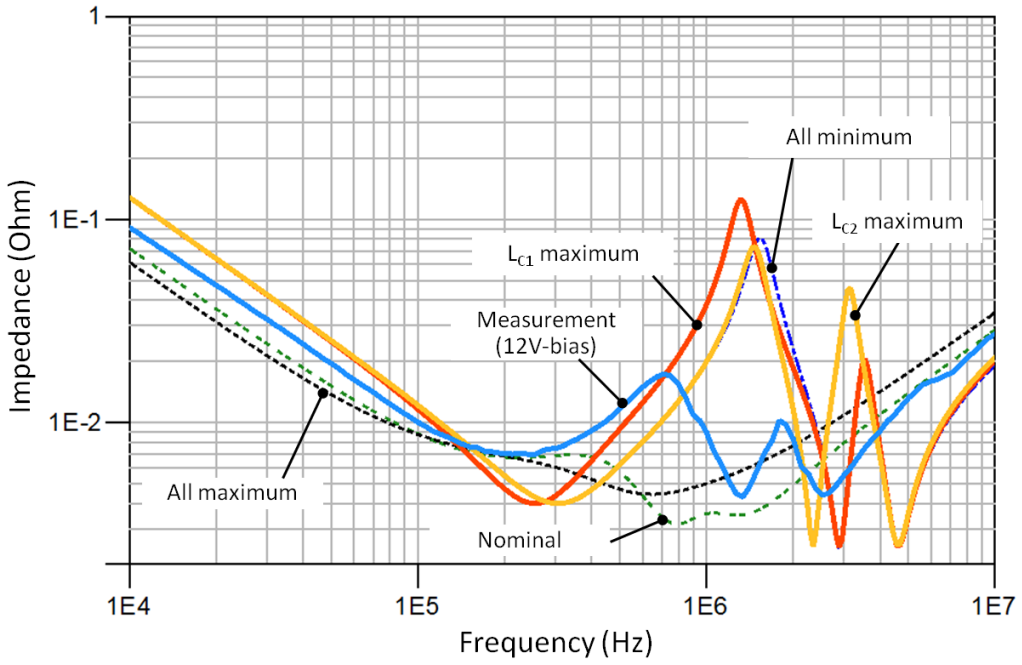

Figure 8d: *Simulated impedance curves with component values from Table II.*
Green dashed line: all nominal values (middle row from Table II)
Black dotted line: all maximum values (bottom row from Table II)
Blue dashed line: all minimum values (top row from Table II)
Blue solid line: measured impedance with 12V DC bias.
Red solid line: all minimum values except L_{c1} is maximum from Table II
Yellow solid line: all minimum values except L_{c2} is maximum from Table II

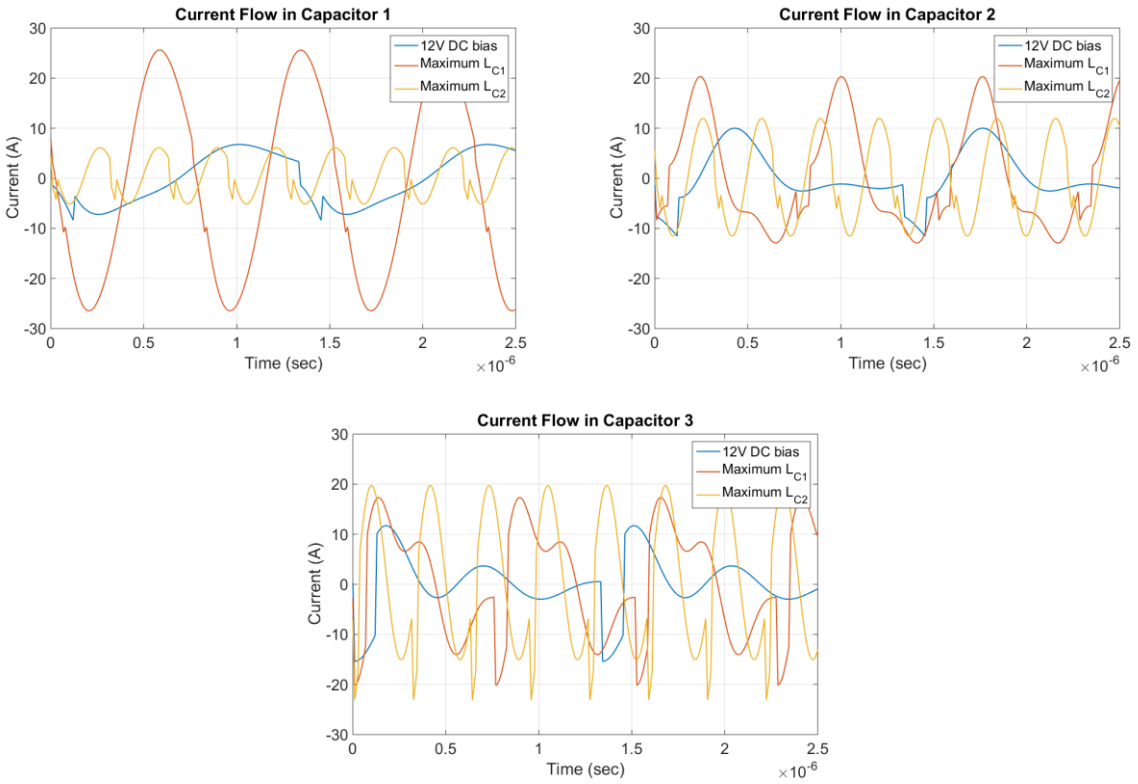

Figure 8e: *Simulated capacitor current waveforms with component values from Table II. Blue line: measured impedance with 12V DC bias. (Fsw=750kHz) Red line: all minimum values except L_{C1} is maximum from Table II. (Fsw=1.318MHz) Yellow line: all minimum values except L_{C2} is maximum from Table II. (Fsw=3.162MHz)*

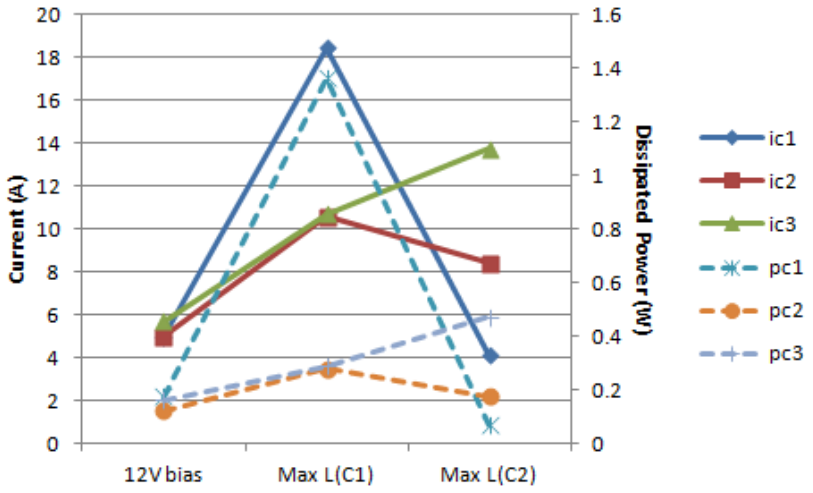

Figure 8f: *Simulated rms current and dissipation in the capacitors with component values from Table II. The three cases (horizontal labels) correspond to the three cases in Figure 8e. Note the significant increase of ripple current under resonance conditions.*

In *Figure 8d* the blue line is the measured impedance under 12V DC bias. Though there is no guarantee for it, we may consider this the typical response under 12V bias. It has two peaks: 17mOhm at 750kHz and 10mOhm at 1.7MHz. The first peak reaches its maximum value when all parameters in *Table II* are minimum, except L_{C1} is maximum: 126mOhm at 1.3MHz. The peak shifted to higher frequencies and the peak value increased seven folds. The second peak reaches its maximum when all parameters in *Table II* are minimum, except L_{C2} is maximum: 46mOhm at 3.2MHz.

	C_{C1}	L_{C1}	R_{C1}	C_{C2}	L_{C2}	R_{C2}	C_{C3}	L_{C3}	R_{C3}
Min.	120uF	2.2nH	4mOhm	2uF	1.5nH	2.5mOhm	2uF	0.6nH	2.5mOhm
Nominal	180uF	2.7nH	8mOhm	22uF	1.9nH	5mOhm	20uF	0.77nH	5mOhm
Max.	200uF	3.2nH	10mOhm	30uF	2.3nH	10mOhm	30uF	0.93nH	10mOhm

***Table II**: Simulation model values for worst-case tolerance analysis. Values for the C3 are for the two parts in parallel.*

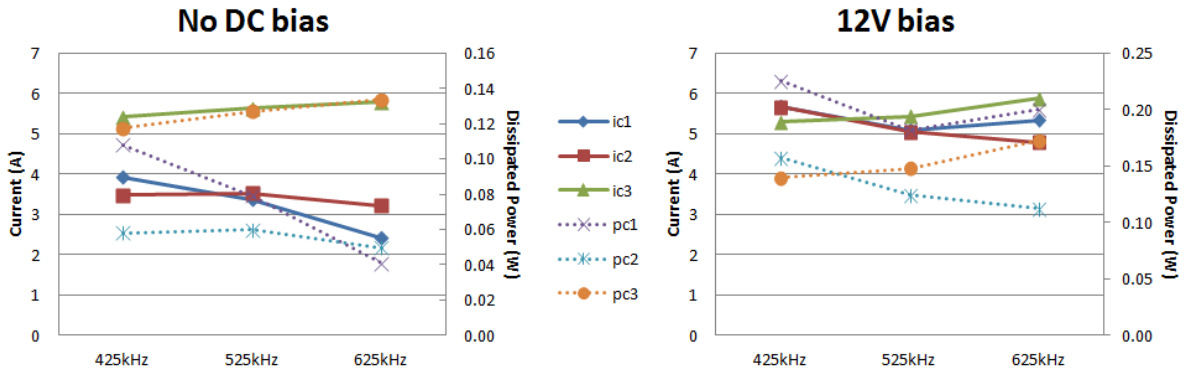

***Figure 9a:** Simulated rms current and power dissipation in each capacitor bank versus switching frequency of the DC-DC converter. No DC bias on the left, 12V DC bias on the right.*

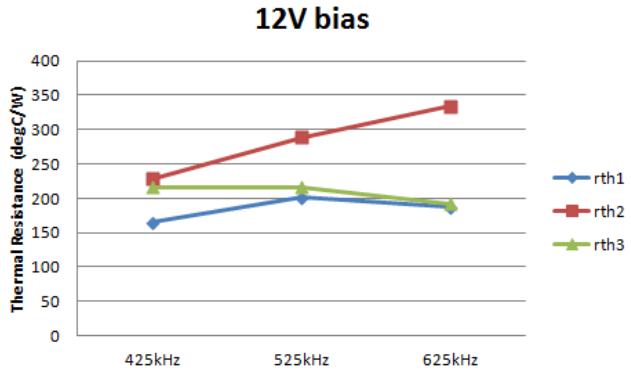

***Figure 9b:** Calculated thermal resistance of each capacitor.*

Figure 9 shows the calculated rms current (solid lines) and power dissipation (dashed lines) in the capacitors as a function of the converter's switching frequency. The thermal resistance of each capacitor is calculated as

$$R_{th} = \frac{T_{capacitor} - T_{ambient}}{Dissipated\_power}$$

and is shown in *Figure 9b* with the assumption of $T_{ambient}$ = 20 degC. When the frequency was changed in the DC-DC converter, the cooling and geometry was not changed and therefore we would expect a constant thermal resistance for each of the capacitors. Contrary to this expectation, the calculated thermal resistance (when the ambient temperature is taken as the ambient temperature to the board) is not constant. The reason for this will be explained below.

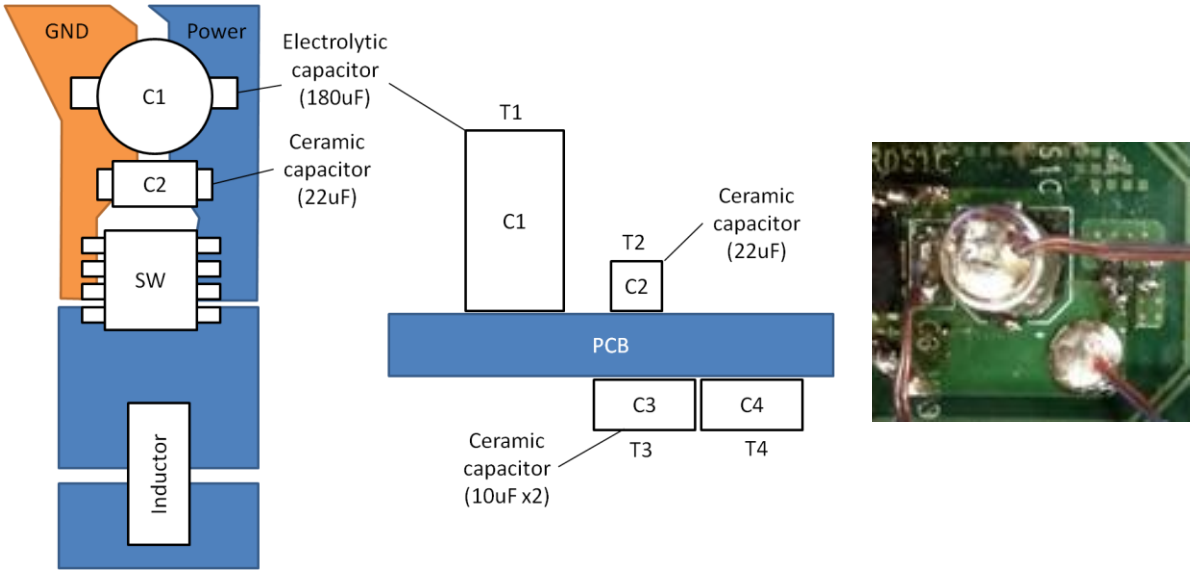

Figure 10: *Simplified sketch of the filter layout. Top view on the left, side view in the middle. T1 through T4 indicate thermocouple locations on capacitors. On the right: photo of thermocouplers on capacitors C1 and C2.*

Figure 10 shows the top and the side views of the filter. The 180uF electrolytic capacitor (C1) and the 22uF ceramic capacitor (C2) are on the top side of the board. Two pieces of the 10uF ceramic capacitors (C3) are on the bottom side of the board.

The capacitors had thermocouplers attached to their top surfaces and the temperature was monitored while the switching frequency was changed (T1 through T4). Also, three additional thermocouplers were attached to the PCB surface to measure the temperature on the top side of the board (T5 ~ T7). They are respectively located right next to C1,

half inches from C1, and two inches from C1. After about 30 minutes of settling time the temperature readings were taken. *Figure 11* shows the measured results.

As can be seen on *Figures 9* and *11*, the dissipation inside the capacitors does change considerably when the switching frequency changes, yet their top body temperature changes much less. This is mostly because of the heat-spreading effect of the printed circuit board, which tends to equalize the temperature across the closely spaced capacitors.

Note also that the test setup had no forced air flow; the large printed circuit board was placed horizontally approximately half inches above the bench surface. The only circuit enabled on the large printed circuit board was the switching regulator under study. With a 20 degC ambient temperature, the board surface around the switching regulator was close to 60 degC, a few degrees warmer than any of the capacitors' top surfaces. This happens for multiple reasons. For one, the input capacitors are not the only components dissipating heat in this setup. The major dissipating components are: input filter inductor (which has to carry the DC input current of the regulator, approximately 3A in this case), the top and bottom-side switching FETs and the output inductor (which have to carry the DC output current plus the inductor's ripple current).

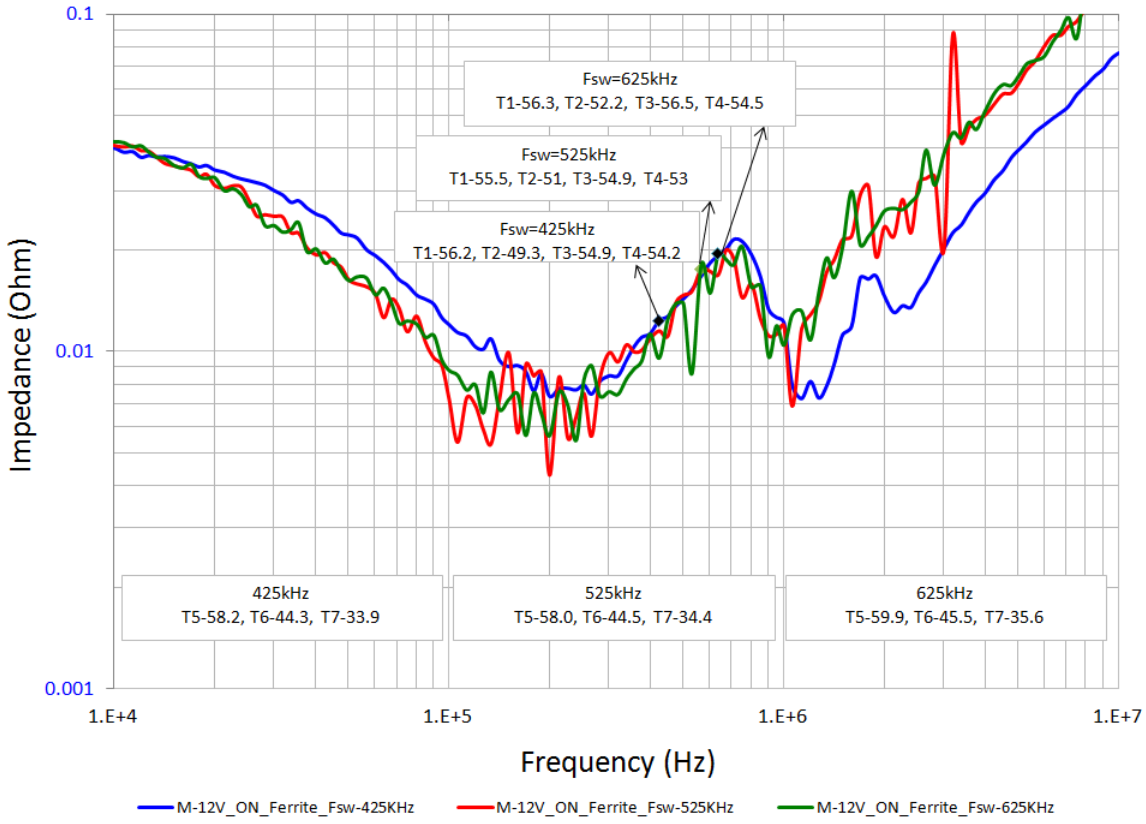

Figure 11: *Measured impedance and temperature of filter components.*

Altogether the static and dynamic losses of the power components in the regulator add up to about 5W. When we compare this value with the less than quarter of a watt total dissipation in the input capacitors, we see that in this particular case the temperature of the capacitors is set by the dissipation of the nearby power components and the overall thermal resistance of the geometry and much less by the dissipation in the capacitors themselves.

During the design phase, when we do not have the stackup and layout details yet, the input filter capacitor selection can only rely on the estimated ripple current through the part compared against the allowed maximum ripple current. Bulk capacitors that we typically use for switching regulators have had the allowed ripple current on their detailed data sheets for a long time.

Recently similar data sheet entries have become available also for larger ceramic capacitors. However, there is no industry-wide agreed-upon process or standard to report the maximum allowed ripple current in these parts. Vendors not knowing the user geometry, usually fall back to a safe minimum assumption, measuring the part in free space mounted on conductors that do not count as significant heat sinks for the part [10]. While this gives a conservative safe bound, it does not take the additional heat removal in typical user geometry into account. Working out a standard procedure would nevertheless be welcome.

For multi-layer medium-size and large-size rigid boards the measured data shows that capacitors directly soldered on copper patches with a thermal mass exceeding the thermal mass of the capacitors, the rated ripple-current specification can be exceeded by at least a factor of two. This is because most of the generated heat escapes into the PCB metal instead of through the air from the capacitor body surface. During the design process, however, the ambient temperature for the capacitors' thermal design has to be adjusted to include the local temperature rise caused by the power components in the immediate vicinity. Also, the detailed simulated waveforms proved that at the capacitor antiresonance frequencies the injected ripple current gets multiplied by the Q of the resonance, hence further increasing the dissipation in the capacitors. In the particular example, the local board temperature rose by 40 degC (with no air flow) merely from the dissipation of the power components of this single DC-DC converter.

III. Impedance flatness and worst-case transient noise

III.1 The Target Impedance and the Reverse Pulse Technique

Though DC-DC converter design processes and stability criteria were developed in the 70s and 80s, systematic PDN design processes have not been established until the target impedance concept was developed [1]. The target impedance concept assumes and is valid for, linear PDNs, where by definition a superposition of separate stimulus current excitations results in the superposition of individual responses. The original target

impedance concept also assumed and referred to flat impedance profiles, in other words, resistive PDN impedances. The design process requires the knowledge of allowed voltage fluctuations, ΔV, and the maximum stimulus step current magnitude, ΔI, and from these establishes a Z_{target} impedance requirement:

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

It is assumed that the ΔI current stimuli have a shortest rise or fall times of t_r, resulting in an approximate bandwidth of

$$BW = \frac{1}{\pi\, t_r}$$

Flat impedance response can be created, but it takes extra effort. A measured illustration from [8] is shown in *Figure 12*.

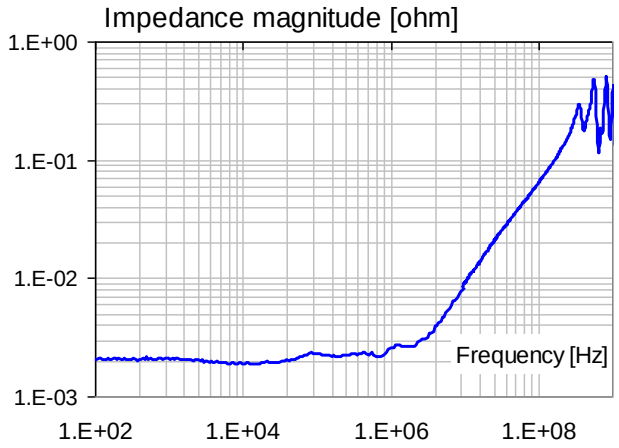

Figure 12: *Measured illustration of a flat impedance profile on a CPU core rail of a server computer board.*

At low frequencies it requires Voltage Positioning [9] in the DC source; at medium frequencies it requires capacitors with controlled ESR [10] and at high frequencies it may require techniques to suppress plane resonances [11]. In low-cost designs or when creating a flat impedance profile is not critical, the PDN has an impedance profile that changes with frequency within the BW bandwidth. In such cases a few systematic impedance synthesis approaches are listed in [12]. A simulated Multi-pole and 'Big-V' PDN synthesis result is shown on the left and right of *Figure 13*, respectively.

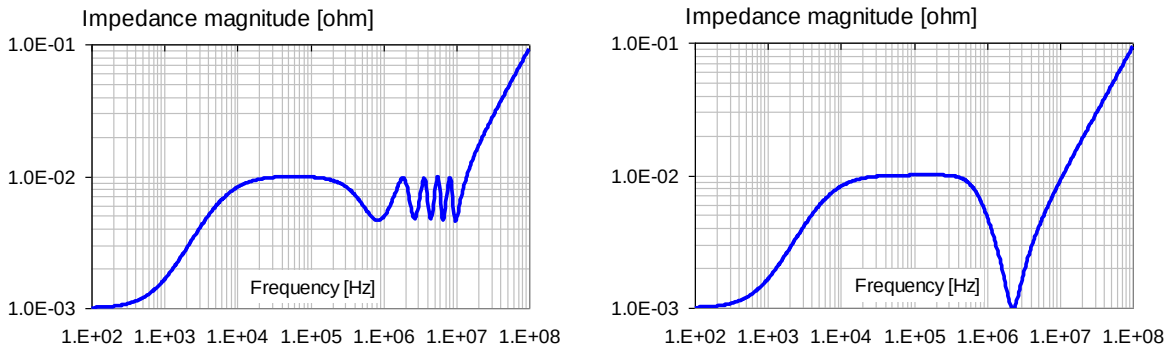

Figure 13: *Left: simulated illustration of a multi-pole impedance profile. Right: Simulated illustration of a Big-V impedance profile.*

When the impedance profile within the excitation bandwidth is not flat, the target impedance is still a useful guidance as long as the appropriate correction factor is applied [12]. The actual worst-case time-domain noise for an arbitrary sequence of current steps with a maximum ΔI magnitude and shortest rise/fall times of t_r can be obtained through the *Reverse Pulse Technique* [6]. With simulated waveforms of a Big-V design, the process is illustrated and summarized in *Figures 14* and *15*.

The process assumes that we start from the steady-state value of the *Step Response* and work backwards towards the moment of excitation. Each maxima and minima are noted and identified by a time stamp and a voltage value. To obtain the worst-case deviation of voltage droop on the supply rail, the excitation starts with steady low and after arbitrarily toggling up and down, stays high. To obtain the worst-case voltage overshoot, the excitation starts at steady high and after arbitrary toggling, stays low. With linear PDN, the two responses must be mirror images. The worst-case peak-to-peak transient noise is twice the difference of the sum of maxima and minima, minus the steady-state DC offset.

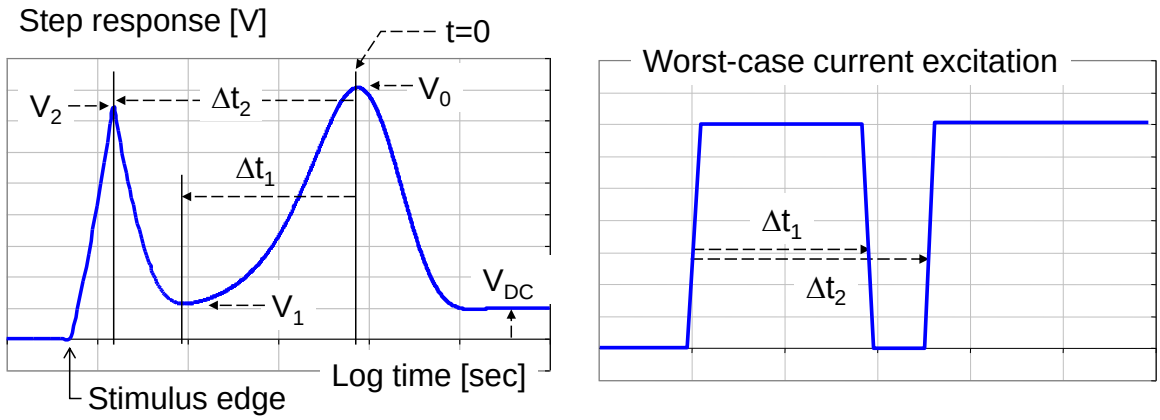

Figure 14: *Illustration and summary of Reverse Pulse Technique. On the left: Step Response and peaks and valleys time stamped. On the right: worst-case excitation waveform based on the time stamps.*

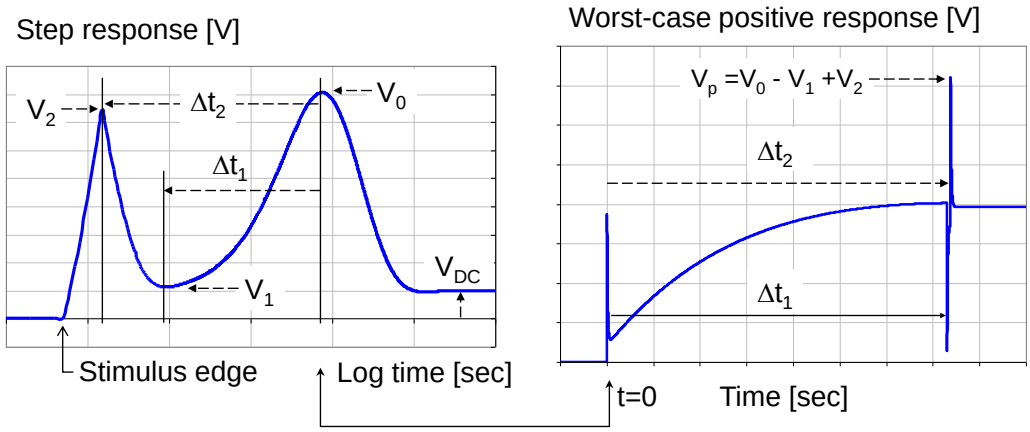

Figure 15: *On the left: Step Response and peaks and valleys time stamped. On the right: worst-case positive transient noise simulated with the waveform from Figure 14.*

Note that on the left plots of *Figures 14* and *15* the horizontal time scale is logarithmic. This is needed to clearly show both fast and slow signatures in the response. The horizontal time scale on the right of *Figure 15* is linear, which shows that the actual time difference between the second and third excitation steps is actually very small.

III.2 Obtaining the PDN Step Response

We can obtain the *Step Response* in several different ways. *Figure 16* shows measured time-domain waveforms obtained with a transient current source connected to the DUT.

The DUT was a multi-phase DC-DC converter with a maximum current capability in excess of 100A DC load current. The measurements were taken with medium-speed high-resolution real-time oscilloscopes from two different vendors [13]. The transient load current was locally generated by a fast power FET driven from an arbitrary waveform generator and it was asynchronous with respect to the DC-DC converter switching: this allows the suppression of the switching ripple by time-domain averaging (#256 in this case). The sketch of the connection geometry is shown on the bottom of *Figure 16*. Impedance profiles were measured based on Two-port Shunt-through method [14]. The response of the DUT was measured on the top side of the load board, where the remote sense of the DC-DC converter was connected. To guarantee consistency, the time and frequency domain measurements used the same cables, probes and probe connection points.

When DC-DC converters are involved, the different loop responses in load attack versus load release may show up. This is more obvious when we have a step-down converter with a duty cycle very different from 0.5: either very low or very high. In such situations

the control has more headroom to make adjustment in the duty cycle in one way than the other.

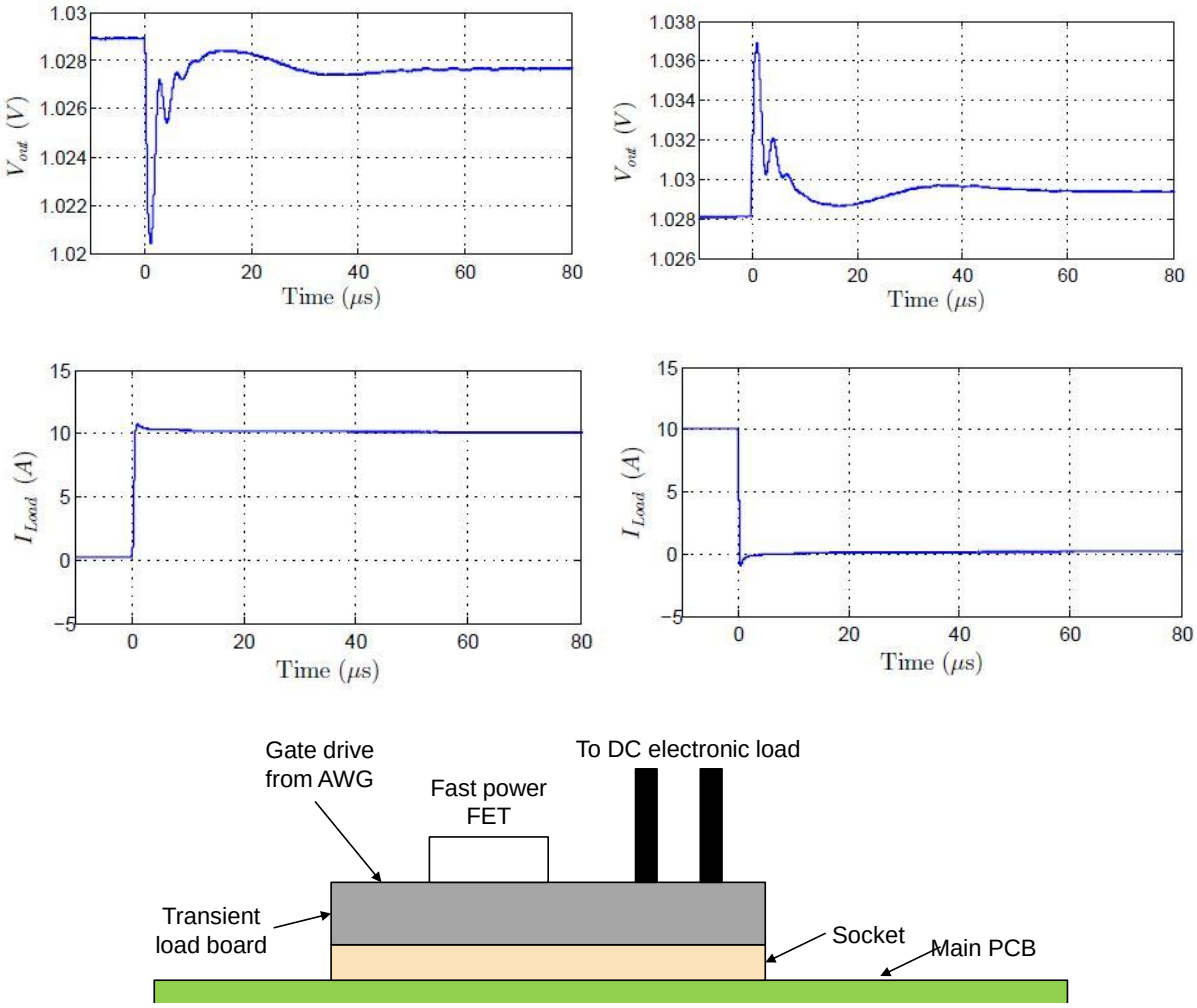

Figure 16: *Top row: measured Step Response waveforms for positive and negative going current stimulus. Middle row: corresponding current stimulus. Bottom: sketch of the connection geometry.*

Figure 17 shows an example of large rising and falling transients when the duty cycle was less than 10%: under such circumstances there is little overhead for the control loop when the load current suddenly decreases. The dashed vertical lines on the right plot identify the time instances of peaks and valleys on the green (load attack) waveform. Note that if we try to use these waveforms for the *Reverse Pulse Technique*, the peaks and valleys on the rising and falling responses do not line up and therefore the process would fail to generate the worst-case transient noise.

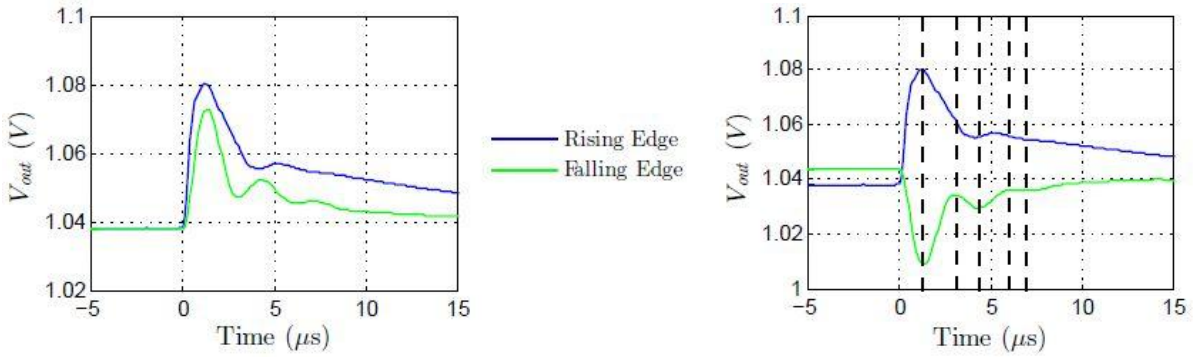

Figure 17: *Different response signatures on the rising and falling edges under large-signal transient conditions. The same data is plotted as measured (on the left) and mirrored (on the right).*

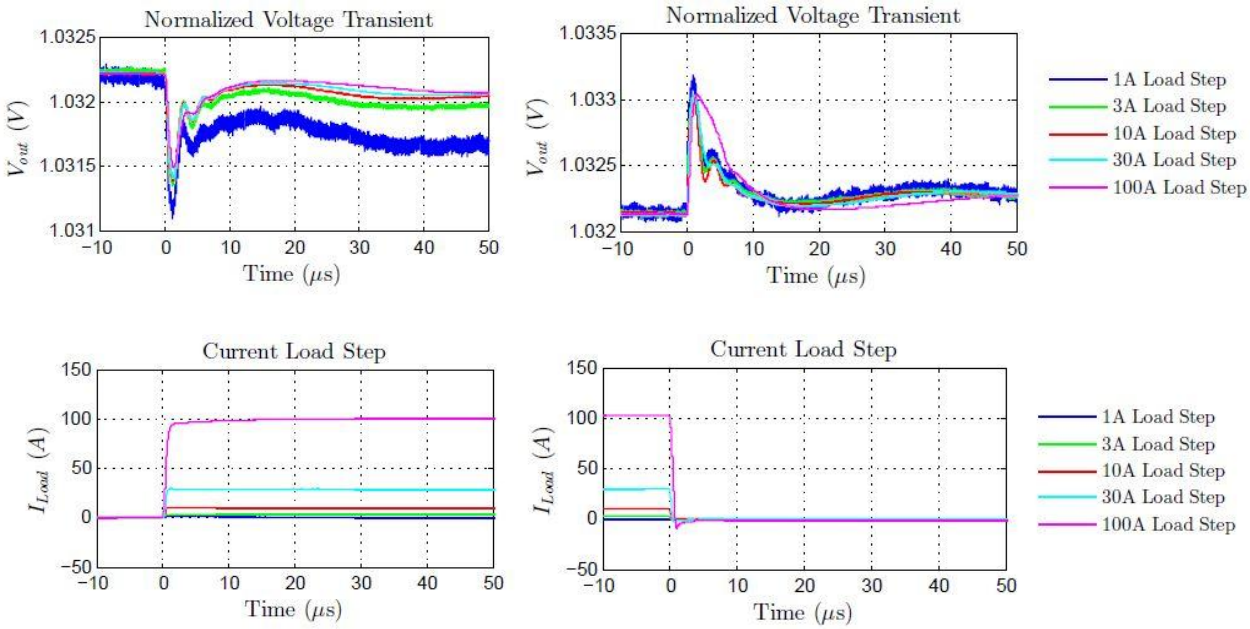

Figure 18: *Normalized load attack and load release Step Response and current excitation waveforms.*

Other nonlinearities can be observed as we change the load step magnitude over a wide range. *Figure 18* is an illustration of load attack and load release waveforms with normalized responses. All response waveforms were scaled to the response of a 1A step and the initial lead-in baselines were manually lined up. While large-signal excitations create challenges with possible nonlinear behavior, small time-domain responses may be degraded by noise and thermal drifts. The 1A trace on the upper left plot in *Figure 18* is a good example of the latter, shifting the response downward by a couple of hundred microvolts. Note that for sake of clarity, the excitation step waveforms are also included,

because it is challenging to maintain the same current wave shape under large variations of step current magnitudes.

As long as we can assume reasonable linearity, we can measure impedance in the frequency domain and calculate the *Step Response* from it, or we can measure the *Step Response* in the time domain and calculate the impedance profile from it. *Figure 19* is an illustration for the latter: the green trace is the measured frequency-domain impedance profile; the blue trace is the impedance profile calculated from the *Step Response* measured in the time domain. *Figure 20* illustrates the agreement between *Step Responses* measured in the time domain vs. calculated from the impedance profile. *Figure 21* shows a series of impedance measurements with various excitations.

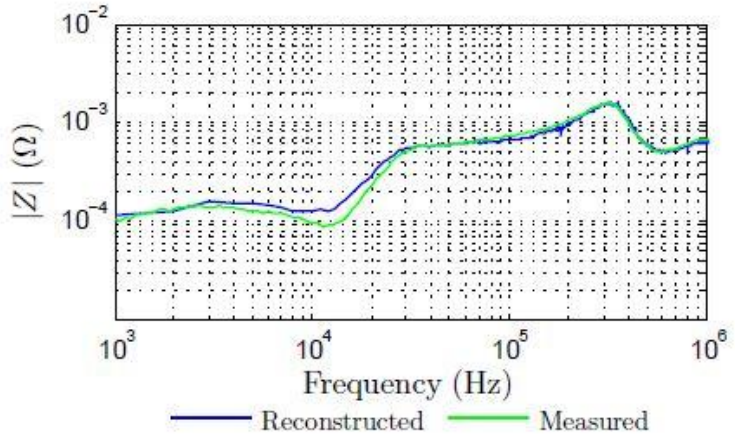

Figure 19: *Impedance profile calculated from the measured Step Response.*

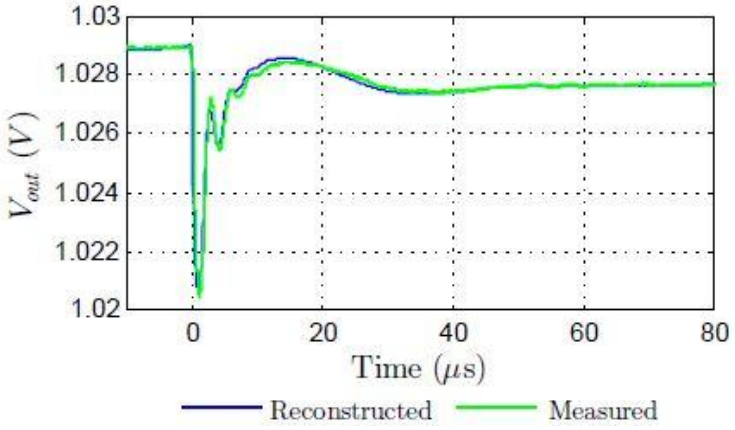

Figure 20: *Calculating Step Response from the measured impedance profile.*

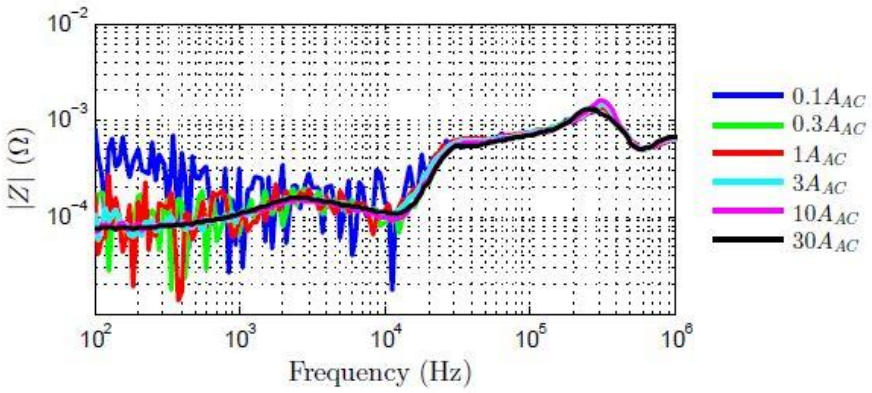

Figure 21: *Various measured impedance profiles measured with sine-wave excitations ranging from 0.1App to 30App.*

.

Note that except the instrument noise floor showing up at low frequencies when we measure a fraction of milliohm impedance with small test signal magnitude, all response traces run on top of each other for most of the frequency range. The only noticeable sign of nonlinearity becomes visible around 300kHz, when we have 30App excitation: the peak value gets reduced by about 15%.

Whether we measure the *Step Response* directly in the time domain with a current-step excitation or calculate it from measured impedance profiles, measurement errors and DUT noise can impair the quality of data. High-resolution oscilloscopes and vector network analyzers today provide sufficient vertical resolution and low enough noise floor that the measurement limitations are not significant. More often the limitation comes from the DUT itself: in addition to the possibility of inherent large-signal nonlinearity, advanced control loops may use nonlinear regulation features and sometimes intentional modulation of the switching frequency to reduce electromagnetic interference. Some converter topologies and component combinations may create a tendency to subharmonic oscillations: this is a known concern in current-mode regulation loops but can also happen in multi-phase regulators as a result of phase imbalance. Subharmonic or beat components show up as bimodal jitter on the time-domain signal and as a single-frequency narrow peak in the frequency response. Finally, random noise on the switching edges show up as jitter in the time domain and elevated noise floor in the frequency domain. Some converter topologies and control loops, such as hysteretic and constant ON-time converters as well as regulators with very aggressive loop bandwidth are typical examples where switching jitter may be a limiting factor. In this regard the three-phase DC-DC converter used for our measured examples represent a middle ground: some jitter was present on the switch nodes, but as *Figures 19, 20 and 21* illustrate, with low bandwidth in the frequency domain and averaging in the time domain the collected data was sufficiently clean and stable.

III.3 Getting the worst-case transient response

Using a DUT setup similar to what was used for the earlier figures, we now demonstrate the full process. *Figure 22* shows the time-domain measured *Step Response* for the rising and falling edges for medium-size current steps. The left plot shows the load-release response mirrored and laid over the load-attack response. The plot on the right shows the same data with the vertical dashed lines identifying the time stamps for the *Reverse Pulse Technique*. Note that the peak and valley timings as well as the ringing signatures show good agreement between the rising and falling waveforms.

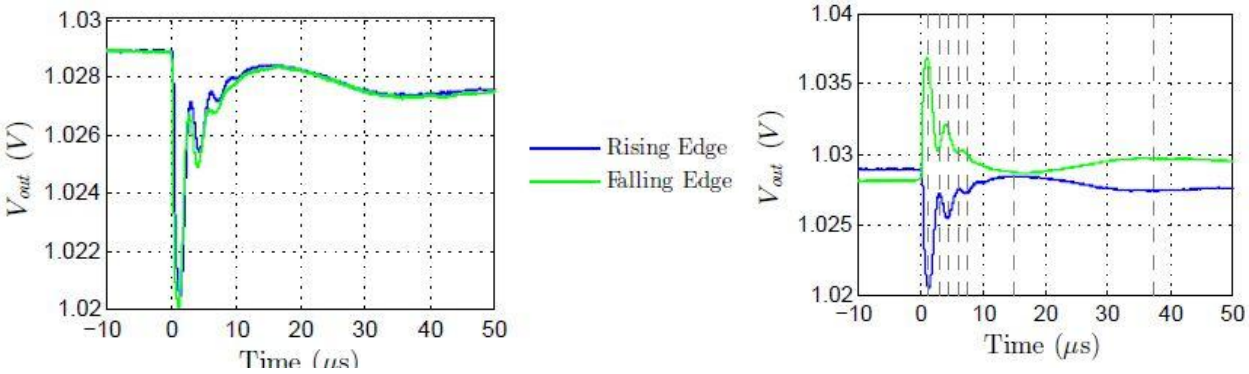

Figure 22: *Measured time domain rising and falling step responses.*

To get this level of agreement even for reasonably linear DUTs, care must be taken to ensure the same bandwidth (transition times and waveshapes) for the rising and falling current excitation edges; moreover the test-point connections have to be chosen carefully.

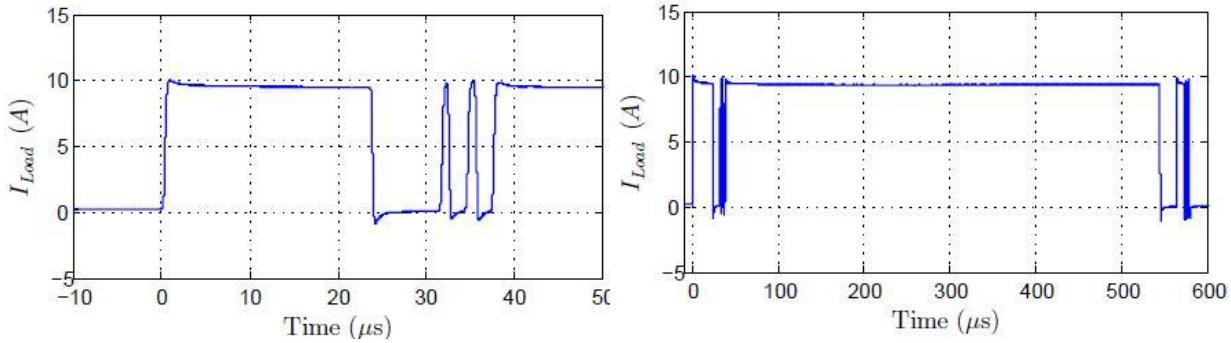

Figure 23: *Calculated worst-case excitation pattern using the time stamps from Figure 22. Step-up sequence on the left, full step-up and step-down sequences on the right.*

Once we have the current time stamps from the *Step Response*, we can put together the worst-case excitation waveform (see *Figure 23*). Using that waveform, we can now excite the DUT with the predicted worst-case pattern. For correlation purposes, we can also get the predicted worst-case waveforms by simulation. We take the measured impedance profile, use a causal rational polynomial approximation [15] of the complex

impedance (to suppress measurement noise) and use the created broadband SPICE model to simulate the worst-case response with a circuit simulator. The correlation result is shown in *Figure 24*. The worst-case one-sided transient noise for an arbitrary sequence of 10A current steps is 10.2 mVpp. The worst-case peak-to-peak transient noise for an arbitrary sequence of 10A steps is 20.2mVpp.

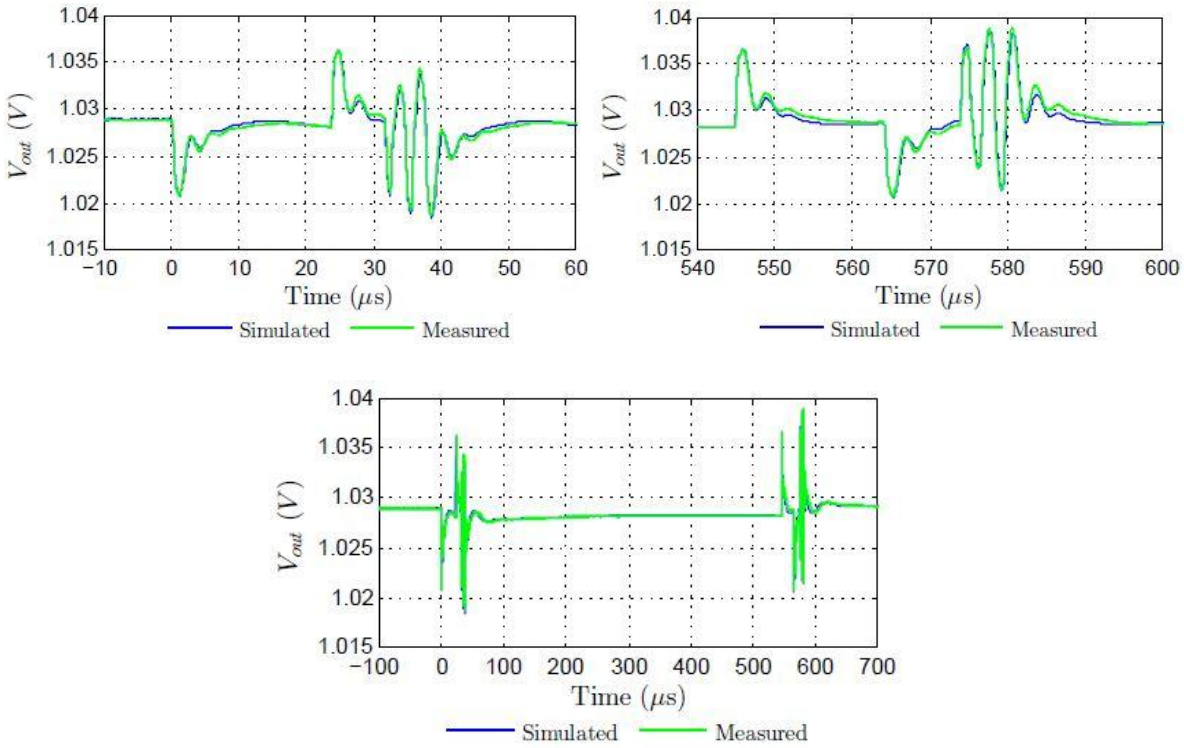

Figure 24: *Correlation of one-sided worst-case transient noise simulation and measurement data based on the Reverse Pulse Technique and 10A steps. Step-up sequence on top left, step-down sequence on top right. Bottom: two-sided worst-case noise starting and ending low.*

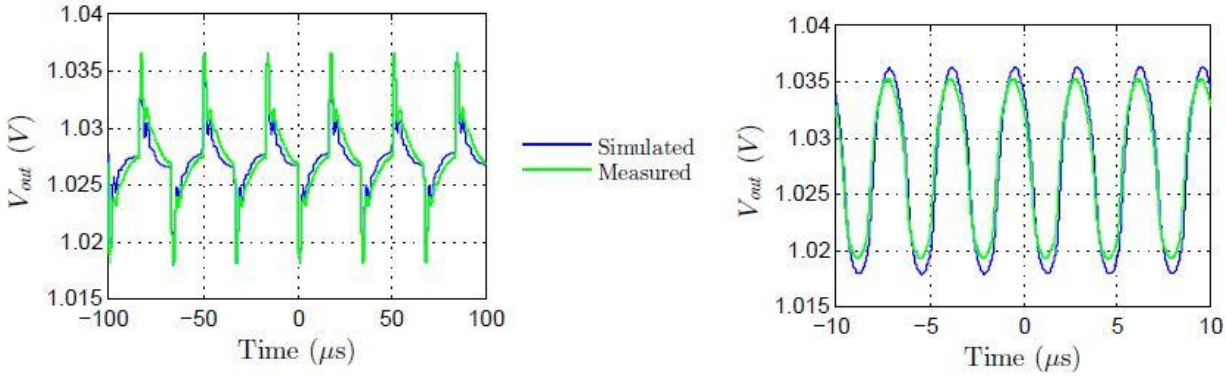

Figure 25: *Transient noise resulting from periodical current steps of 10A magnitude with 50% duty cycle. 30kHz repetition frequency on the left, 300kHz repetition frequency on the right.*

Lastly we show the result on the same DUT with different approximations for the noise voltage. By looking at the impedance profile, we can identify two local peaks: one around 30 kHz and another around 300 kHz. If the testing process uses only periodical waveforms with fixed or variable duty cycle, it is customary to tune the switching frequency around the peaks until the maximum peak-to-peak deviation is reached. *Figure 25* shows the result for these two frequencies.

The measured peak-to-peak transient noise with 30 kHz and 300 kHz repetitive edges is 18.5 mVpp and 16.1 mVpp, respectively. Note that the transient noise measured this way underestimates the worst case noise by 8.4% and 20%, respectively.

Conclusions

Non-flat impedance profiles in power distribution networks carry several penalties. In this paper we looked at two distinct scenarios: the case of an input filter for DC-DC converters and the case of board distribution network impedances composed of DC-DC converters and board capacitors.

In case of the particular input filter used in this study, it was found that on the multi-layer large-size rigid board with capacitors directly soldered on copper patches with a thermal mass exceeding the thermal mass of the capacitors, the rated ripple-current specification can be exceeded by at least a factor of two. This is because most of the generated heat escapes into the PCB metal instead of through the air from the capacitor body surface. During the design process, however, the ambient temperature for the capacitors' thermal design has to be adjusted to include the local temperature rise caused by the power components in the immediate vicinity. It was also shown that at the antiresonance frequency of the capacitors the injected ripple current gets multiplied by the Q of the resonance, hence further increasing the dissipation in the capacitors. DC bias effect can reduce the capacitance of ceramic capacitors, potentially increasing further the impedance peaks in worst case combinations, and in extreme corner cases the ripple current can significantly increase, potentially exceeding the rated limit.

For board decoupling applications, the Reverse Pulse Technique can be used to estimate the worst-case transient noise, assuming that the PDN, including the DC-DC converter, is reasonably linear and time invariant. Measured and simulated data shows that the target impedance concept can still be used for non-flat impedance profiles: for our practical range of non-flatness values, a conservative correction factor of three can be applied. It was demonstrated that for reasonably linear PDNs the time-domain *Step Response* and frequency domain impedance profile can be used interchangeably and can be transformed from one to the other. It was shown that for reasonably linear PDNs the small-signal and large signal impedance profiles do agree to a large degree. It was also shown that even with minor non flatness of the impedance profile, the popular approach of estimating the worst-case noise by tuning a periodical waveform to yield maximum noise can underestimate the worst case noise by up to 20%, which error increases as the non-flatness increases.

Future work is necessary to establish procedures to obtain guaranteed worst-case transient noise under medium and heavy nonlinearities.

Acknowledgement

The authors wish to thank Kavitha Narayandass for her support with lab instrumentation and measurements, Bill Couillard, Gene Whitecomb, Seyla Lang for their lab support and the Oracle SPARC Enterprise design team for providing their transient test fixture.

References

[1] Larry Smith, et al, "Chip−Package Resonance in Core Power Supply Structures for a High Power Microprocessor," Proceedings of IPACK 2001

[2] Christian Schuster, "Developing a Physical Model for Vias," Proceedings of Design 2006

[3] Istvan Novak, "Reducing EMI Noise by Suppressing Power-Distribution Resonances," Proceedings of Ansoft EMC Seminar, Boston, 2007.

[4] Steven Sandler, "Target Impedance Limitations and Rogue Wave Assessments on PDN Performance," Proceedings of DesignCon 2015

[5] Novak, et al, "Distributed Matched Bypasssing for Board-Level Power Distribution Networks," IEEE CPMT, August 2002

[6] Drabkin, et al, "Aperiodic Resonant Excitation of Microprocessor Power Distribution Systems and the Reverse Pulse Technique," Proceedings of EPEP, 2002, pp. 175- 178.

[7] "DC and AC Bias Dependence of Capacitors Including Temperature Dependence," DesignCon East 2011, Boston, MA, September 27, 2011

[8] "Controlled-ESR Bypass Capacitors Have Arrived" in TecForum TF-MP3 at DesignCon 2007, January 29 – February 1, 2007, Santa Clara, CA

[9] "Designing Minimum Cost VRM8.2/8.3 Compliant Converters," Proceedings of High Frequency Power Conversion Conference, November 7-13, 1998, Santa Clara, CA.

[10] "Measuring the Power Dissipation Capability of High-Voltage, Low-Capacitance Ceramic Capacitors," CARTS USA 2005, Palm Springs, CA, March 2005.

[11] "Reducing Simultaneous Switching Noise and EMI on Ground/Power Planes by Dissipative Edge Termination," Proceedings of EPEP 1998, October 26-28, 1998, West Point, New York.

[12] Bypass Capacitor Selection Based on Time Domain and Frequency Domain Performances in TecForum TF-MP3 "Comparison of Power Distribution Network Methods", DesignCon 2006, Santa Clara, CA, February 6-9, 2006.

[13] http://literature.cdn.keysight.com/litweb/pdf/5991-4087EN.pdf?id=2456396 and http://teledynelecroy.com/oscilloscope/oscilloscopeseries.aspx?mseries=458

[14] Frequency-Domain Characterization of Power Distribution Networks. Artech House, 1996.

[15] www.simberion.com