DesignCon 2007

Panel discussion:
Emerging Challenges of DC-DC Converters

Panelists:

Don J. Nguyen	Intel
Shaun L. Harris	Hewlett Packard
Istvan Novak*	SUN

* panel organizer and moderator

Panelist biographies:

Don J. Nguyen joined Intel in 1996 as a Power Technologist, responsible for designing and architecting current and future mobile power delivery subsystem including the Intel Mobile Voltage Positioning (IMVP). Prior to joining Intel, Don worked for Image Quest Technology designing electronic driver circuits for active-matrix Thin-Film-Transistor Liquid Crystal Display. Prior to this, he worked for Poqet Computer Inc. designing and testing various battery-operated DC-to-DC Converters battery chargers and display drivers. He received 25 patents in the field of power and display systems.

Shaun L. Harris is a Hewlett Packard Power System Architect for the Business Critical Systems in the High Performance Systems Lab. His responsibilities include driving technology direction in HP's servers products as well as drive industry initiatives. Prior to HP, Shaun worked at Texas Instruments where he was an analog and power supply designer on Electro-Optical weapon systems. Shaun received an MS from Texas Tech University and BS in Electrical Engineering from The University of Oklahoma. He has been granted 28 patents.

Abstract:

Have you ever wondered how much bulk capacitance you need to support a DC-DC converter in a given application? Have you had problems with converter stability with certain input or output capacitor combinations? Have you had signal-integrity and EMC issues because of DC-DC converters? This panel brings together the views of experts from large OEMs (Hewlett Packard, Intel, SUN) and gives you their laundry list of potential problems and challenges with today's off-the-shelf DC-DC converters. Following the wish-list presentation of OEM experts, the floor will be open for discussion.

DC-DC converters: What is wrong with them?

Istvan Novak

Sun Microsystems, Inc.
Tel: (781) 442 0340, e-mail: istvan.novak@sun.com

Author Biography

Istvan Novak is signal-integrity engineer at SUN Microsystems, Inc. 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 Distinguished Engineer of SUN Microsystems and Fellow of IEEE for his contributions to signal-integrity and RF measurement and simulation methodologies.

Abstract

DC-DC converters are widely used in power distribution networks. While they can conveniently generate any regulated voltage, off-the-shelf converters come with several inevitable compromises. The compromises stem from the fact that the converter designer does not know the exact nature of the load: the bypass capacitors and their interconnects are different from design to design. This paper illustrates four major concern areas associated with off-the-shelf converters:

- Limited control-loop bandwidth may require large output capacitance
- Non-flat and/or non-stable output impedance profile results in increased transient noise
- Unspecified high-frequency ringing may risk sensitive clock signals
- Non-linear control loops may make validation more tedious

I. Introduction

DC-DC converters are the staples in today's power distribution networks. The reason for the popularity of DC-DC converters is multi-fold:

- various kinds of semiconductor pieces (memory, IO, core) continue to follow their own supply-voltage scaling and therefore the number of independent supply rails in the systems is on the rise,
- the continued push for higher efficiency in both high-power and low-power systems calls for generating the supply voltages near their destination,
- Distribution of raw high voltage reduces distribution losses.

DC-DC converters are switching circuits, which chop the incoming DC voltage, and by controlling the duty cycle of the pulse stream, the output voltage can be translated and tightly controlled [1]. There are numerous topologies and circuit solutions; with the proper control and output circuit, one can create isolated or non-isolated converters, and output voltages can be in any relations with respect to the input voltage. The same-polarity step-down converter is commonly referred to as buck converter. The same-polarity step-up converter is called boost converter. Finally the converter, which changes the polarity between input and output voltages is called buck-boost converter.

Figure 1 shows the simplified switching scheme of a synchronous buck converter. The upper and lower FETs switch in opposite phase, creating a pulse stream with steady levels approximately equaling zero and the input voltage. The LC low-pass filter on the output suppresses the main switching frequency and its harmonics, and produces a DC average voltage equaling the input voltage times the duty cycle. This results in the output-input voltage relationship following approximately the duty cycle. Static and dynamic switching losses slightly modify this ratio.

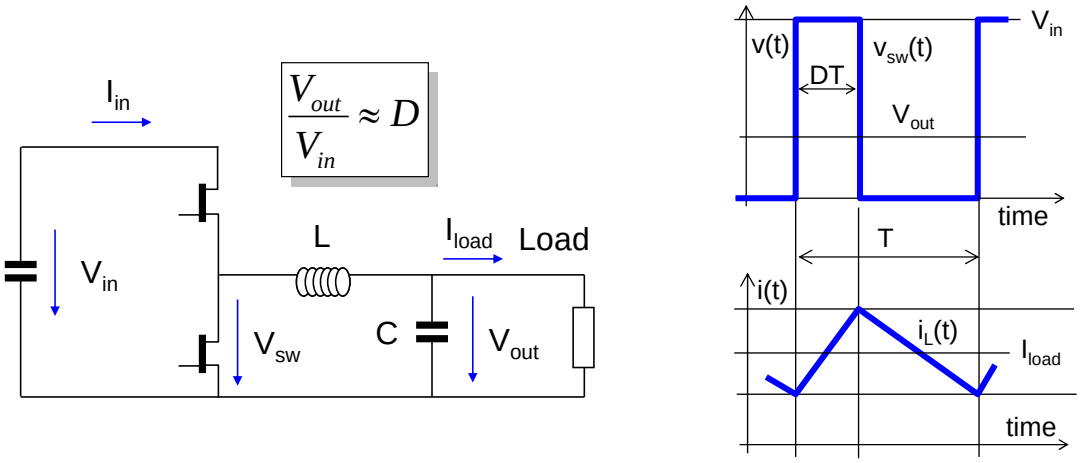

Figure 1: *Switching scheme and characteristic waveforms of non-isolated buck converter.*

By controlling the duty cycle in closed loop, one can keep the average output voltage constant in spite of (slowly) varying input voltage, load current and losses. Assuming ideal lossless components, no power is dissipated in the converter circuit, the entire input power is transported to the output, and the conversion efficiency is 100%. In real circuits, there are always finite losses and the efficiency with close-to-nominal output power is usually in the 75…95% range. As output power decreases, efficiency inevitably goes down due to the current necessary to feed the standby and control circuits.

The switching circuit theoretically can work with a large range of frequencies. As frequency increases, the circuit requires less inductance and capacitance, which results in smaller and lighter reactive components, however dynamic losses increase. This balance creates a sweet spot of switching frequency with any given available technology, which is a balance between losses and size/weight of reactive components. Today the main switching frequency (per phase) typically falls between 200 kHz and 2 MHz. The equivalent switching frequency can be increased, and thus the required inductance and capacitance can be reduced, by using multiple interleaved phases, running at the same frequency.

2. DC-DC converter parameters related to signal integrity

The control loop monitors the voltage across the load and adjusts the duty cycle accordingly. The feedback loop is a sampled system, leading to band limitation, which is proportional to the switching frequency. In today's complex applications, the single load in the equivalent circuit may represent multiple integrated circuits connected by a complicated network of printed circuit board planes, bypass capacitors and optional inductors. *Figure 2* shows the block schematics of the sampled feedback-loop.

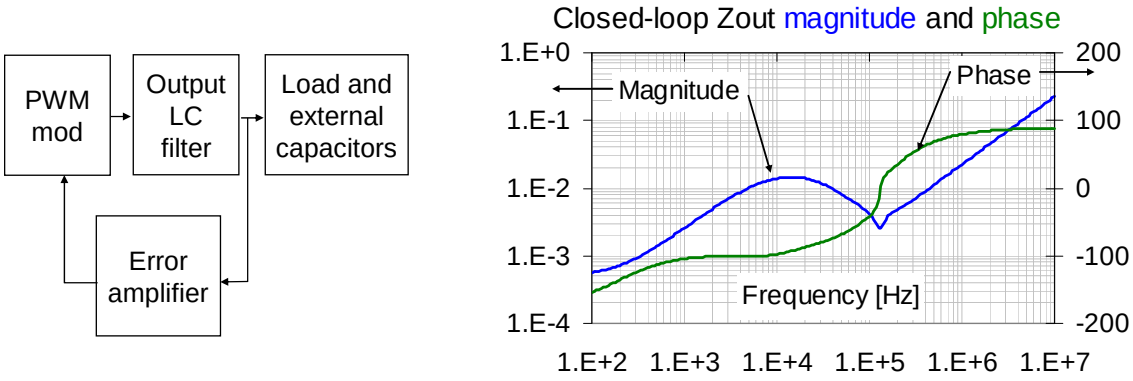

Figure 2: *Block schematics of a DC-DC converter on the left, simulated small-signal output impedance on the right.*

The simulated small-signal output impedance plot on the right of *Figure 2* shows that this control loop maintains constant output voltage in spite of load-current changes, which manifests itself in a low output impedance; less than a milliohm at 100 Hz in this case. However, as frequency increases, the output impedance increases due to the finite bandwidth of the control loop. Above a few hundred kHz, the output impedance follows an inductive slope, dictated by the parasitics of the output capacitor network; beyond the switching frequency the control loop has a diminishing effect on the loop parameters.

Important quality parameters of the DC-DC converter are the transient noise due to load-current and input-voltage changes as well as due to its own internal switching. *Figure 3* illustrates two major signatures of the stead-state switching noise: switching ripple and switching ringing. The switching ripple is an inevitable consequence of how the converter works. The peak-to-peak ripple can be reduced by increasing the inductance and/or the capacitance in the output stage. The tradeoff is that larger values will result in slower transient response. The switching ringing, on the other hand, is a side effect, generated by the main switching signal through the layout and component parasitics. As the main

switching edge gets faster, larger and potentially higher-frequency ringings are induced. The tradeoff is between efficiency, cost and ringing.

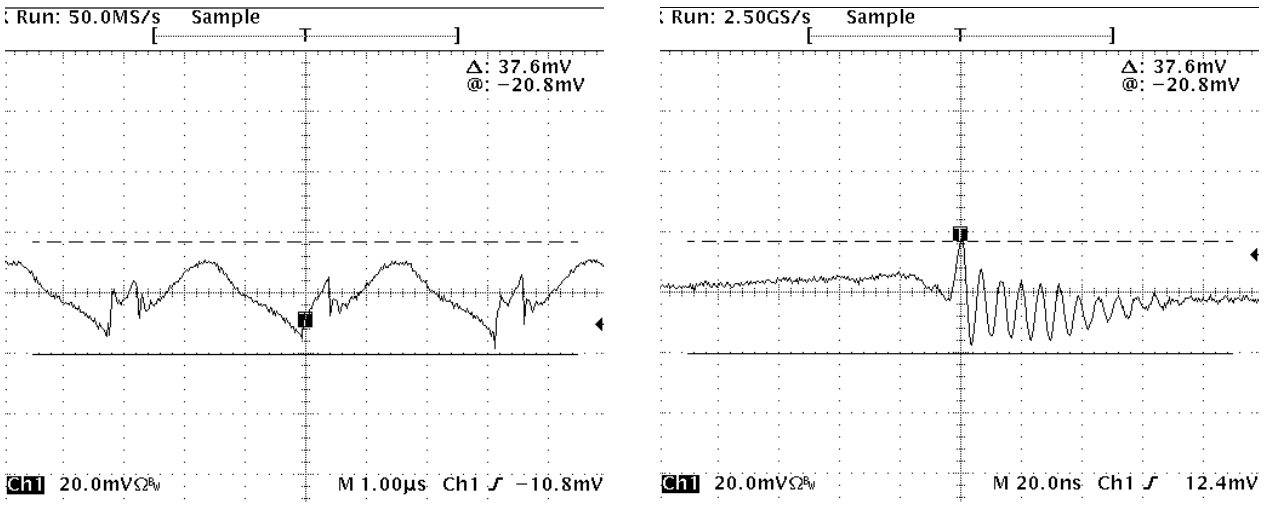

Figure 3: *Steady state transient noise signatures of a 5-A point-of-load buck converter. Switching ripple measured with 20MHz bandwidth on the left, high-frequency ringing measured with 500MHz bandwidth on the right.*

The load transient noise behavior of the converter is determined by the switching elements, control-loop topology and components and the output capacitors. We need to realize that the user load is inside the control loop, determining its dynamic performance. If the load comes with a set of its own bypass capacitors, those will become part of the control loop, hence having an impact on the transient noise. The converter designer has no way of knowing what bypass capacitors the end users will attach to the converters. On the flipside, users of off-the-shelf converters have very little or no information about the control loop parameters, so it is hard to determine how the externally attached bypass capacitors impact the dynamic performance.

Because the converters are based on switched waveforms, they are inherently nonlinear. However, well-designed analog control loops result in a stable averaged response, and when the output has the proper set of capacitors, they can be well approximated as linear circuits. For linear circuits, determining the worst-case transient noise created by step-like excitations is a straightforward process based on the impedance profile [2]. The *Reverse Pulse Technique* takes the step response of the PDN and composes the worst-case transient response from the time stamps and values of subsequent maxima and minima. It is easy to show that the worst-case transient noise is minimized if the frequency-domain response, namely the output impedance, is flat [3].

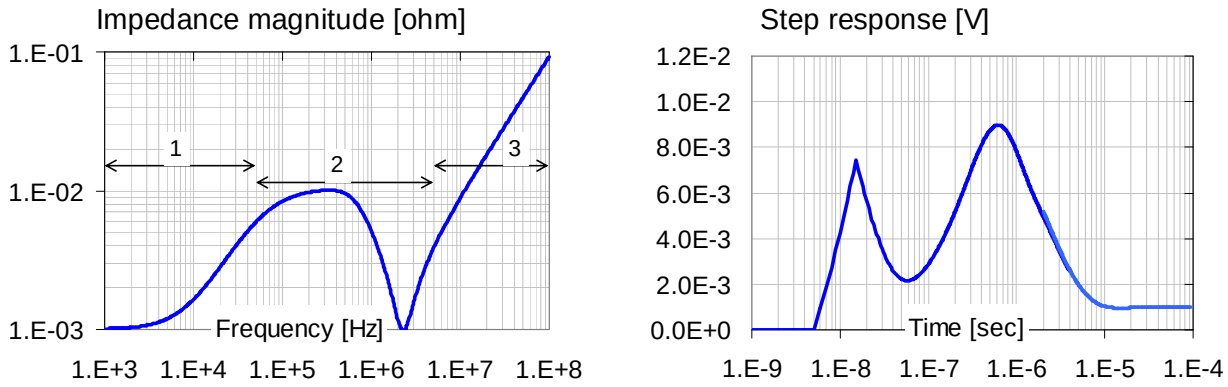

Figure 4: *Simulated output impedance and step response of a DC-DC converter with external capacitors. Note that both horizontal scales are logarithmic.*

Figure 4 illustrates a pair of output impedance and its corresponding step response. In the impedance profile, we can identify three frequency ranges. *Range 1* extends from DC to a few times ten kHz. In this frequency range the components and the control loop of the DC-DC converter determine the impedance; at such low frequencies we have hardly any chance to influence the impedance with external capacitors. *Range 3*, which starts a few times of the switching frequency, behaves the opposite way: way above the switching frequency the control loop has a diminishing influence on the output impedance, and therefore in this frequency range the performance is primarily determined by the passive components. Finally *Range 2* is an interim grey zone, where the output impedance is determined by the interaction of the control loop and output capacitors.

3. Potential SI problems from the user's perspective

Off-the-shelf converters are generic products, aiming for a wide range of applications. Some converters have a wide input-voltage range; some also have a wide adjustable output voltage range. The load current may also vary in a wide range, ultimately between zero and the full rated current. The wide parameter ranges, together with a huge set of potential input and output capacitors make the converter design a real challenge. It is equally challenging though to take an off-the-shelf converter and design a PDN network around it for a given set of requirements. The difficulties mostly come from the middle frequency range, *Range 2*, in *Figure 4*, where two blocks of components, the DC-DC converter and the external PDN interact, but both are designed with very minimal if any information about the other.

3.1 Low loop bandwidth

Figure 5 shows small-signal output impedance profiles, measured on a point-of-load buck converter. The converter had a wide input voltage range (3.0 to 13.2V) and adjustable output voltage range (0.7 to 3.6V). The maximum rated current is 17A. The impedance profiles were measured in small increments along all three parameters: input voltage, output voltage and DC load current. The figure shows two representative summary plots at 1.0V output voltage. The left plot shows the minimum, average and maximum impedance with 3V input voltage, as DC load current steps through 0 to 17A. The right plot shows the same with 10V input voltage. The min/max range is fairly tight, indicating a stable loop, and showing the possibility that the small-signal impedance profile can be used to get the load-step transient response through IFFT.

Instead of doing IFFT on the measured impedance profile, we can make simple approximations to evaluate the usability of the converter. Let us assume that the load current can have 5-A steps, which for a 17-A converter is a modest assumption. If the 5-A steps happen to occur with a periodicity aligning with the 0.1-ohm impedance peak at 16-kHz, the resulting transient noise with 3V input is 5App*0.1ohm = 0.5Vpp. (Actually because the impedance profile will significantly attenuate the harmonics of the assumed 16-kHz square-wave load current, the PDN will pick out the fundamental waveform of the square wave, which is $4/\pi = 1.27$ times the fundamental.) On a 1.0V supply rail we can very seldom afford 0.5Vpp transient noise, and most probably the control loop would become nonlinear for this large excursion. On a 1.0V rail a more realistic noise target could be 50mVpp. To achieve 50mVpp noise with 5-A load step changes, the impedance should stay below 10 milliohms. The output impedance can be reduced by adding low-impedance capacitors. The average impedance curve of the converter crosses the 10-milliohm line at 3 kHz with 3V input voltage. To get 10 milliohm capacitive reactance at 3 kHz, we need approximately 5,000 uF capacitance. To provide a peak-free transition in the impedance profile, we would need to use at least twice of this capacitance, or more than 10,000 uF.

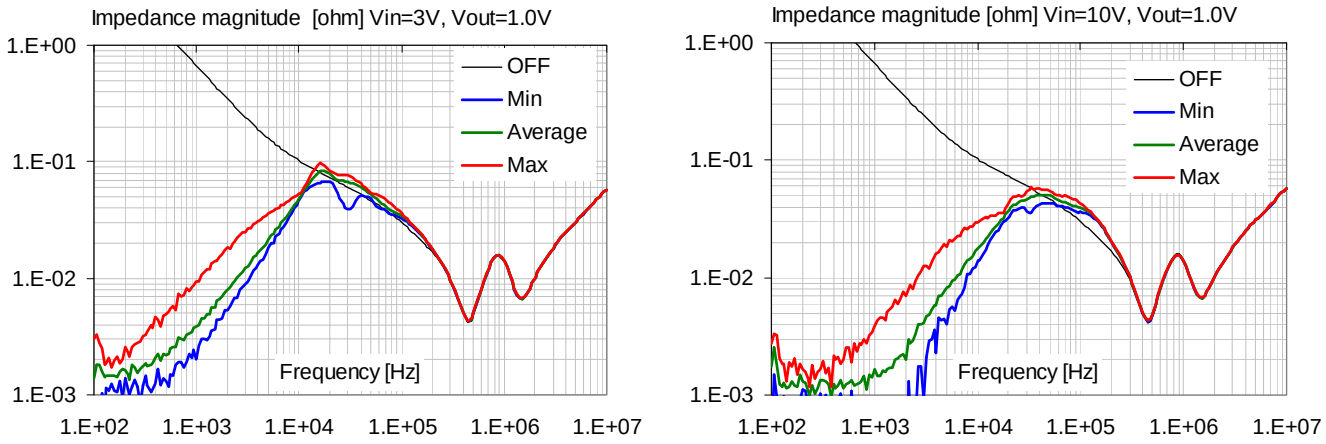

***Figure 5**: Measured small-signal impedance-profile summaries of a 17-A point-of-load converter.*

This simple illustration shows that often times the converter's advertised rating can not be fully utilized without a large amount of external capacitance, which requires extra board space and increases the total cost of ownership.

3.2 Stability and impedance profile

Traditionally the converter stability is described with the phase and magnitude margins of the transfer function of the control loop. Usually 45 degree phase margin is required. While this ensures that the control loop will not exhibit a large class of self oscillations, we need to keep in mind that for the user the ultimate performance parameter is the worst-case transient noise. As long as the small-signal output impedance does not change much with the DC load current, the transient noise can be minimized with a flat impedance profile. So for practical purposes, the small-signal impedance surface is a good indicator of quality. *Figures 6 and 7* illustrate the inter-relation of phase margin and impedance-profile flatness. The example uses the measured impedance profiles of a multi-phase non-isolated 100-A converter with two different loop compensations. Both converters had the same external capacitors. The converter in

Figure 6 had 70 degrees of phase margin, while the same converter in *Figure 7* had only the usual 45 degrees. Note that the impedance response in *Figure 7* has a 1.4 milliohm peak at 67 kHz, which is 40% higher than the broad-band flat plateau exhibited in *Figure 6*.

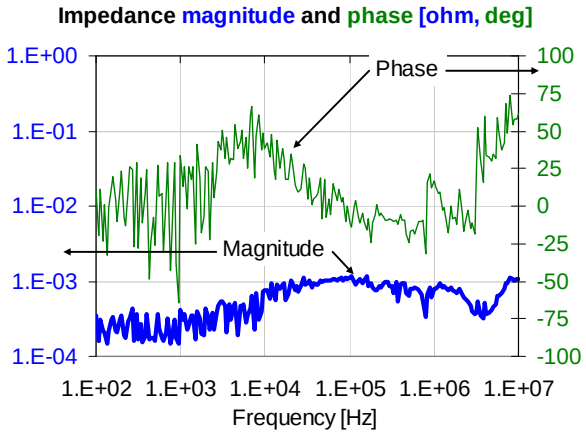

***Figure 6**: Impedance profile of a 100-A multi-phase converter with 70-degree phase margin.*

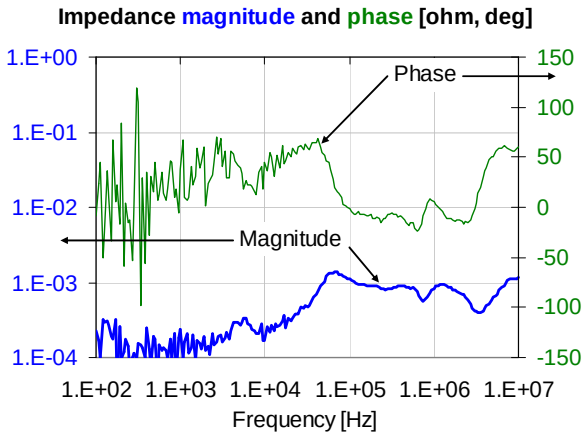

***Figure 7**: Impedance profile of the same 100-A multi-phase converter with 45-degree phase margin.*

The converters in *Figures 6 and 7* had very stable impedance response with the given external capacitors; within the specified ranges of input voltage and DC load current the impedance profile changed very little.

If the control loop is not properly designed, and/or if the converter layout allows the various circuit loops of the converter to interact, the impedance profile may show large variations. *Figure 8* shows the measured impedance plots of a converter with stable impedance profile, whereas *Figure 9* illustrates an unstable impedance profile. Both converters were point-of-load 20-A buck converters with no external output capacitor.

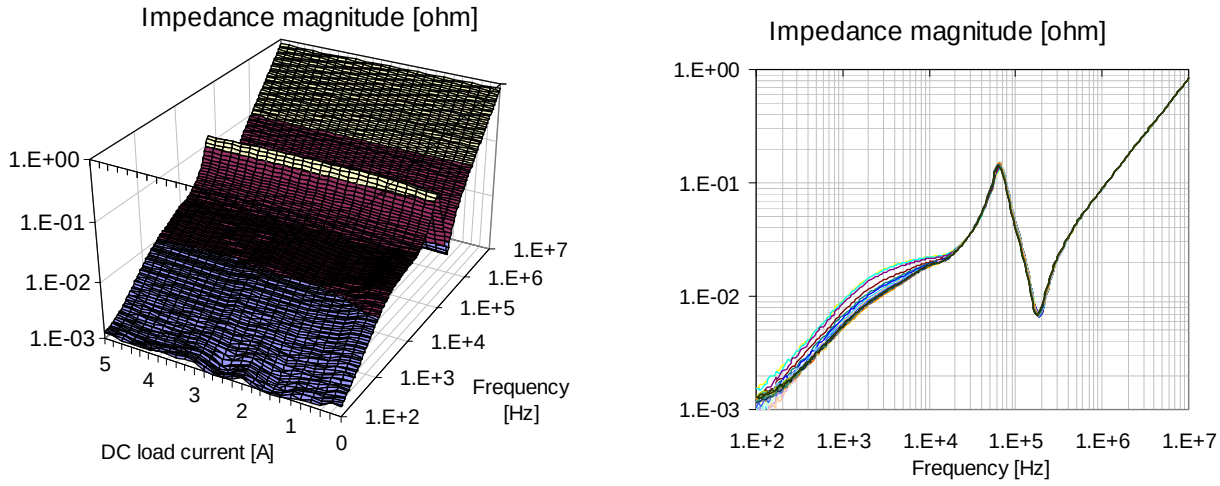

Figure 8: *Impedance surface (on the left) and cumulative traces (on the right) of a point-of-load buck converter with stable impedance profile.*

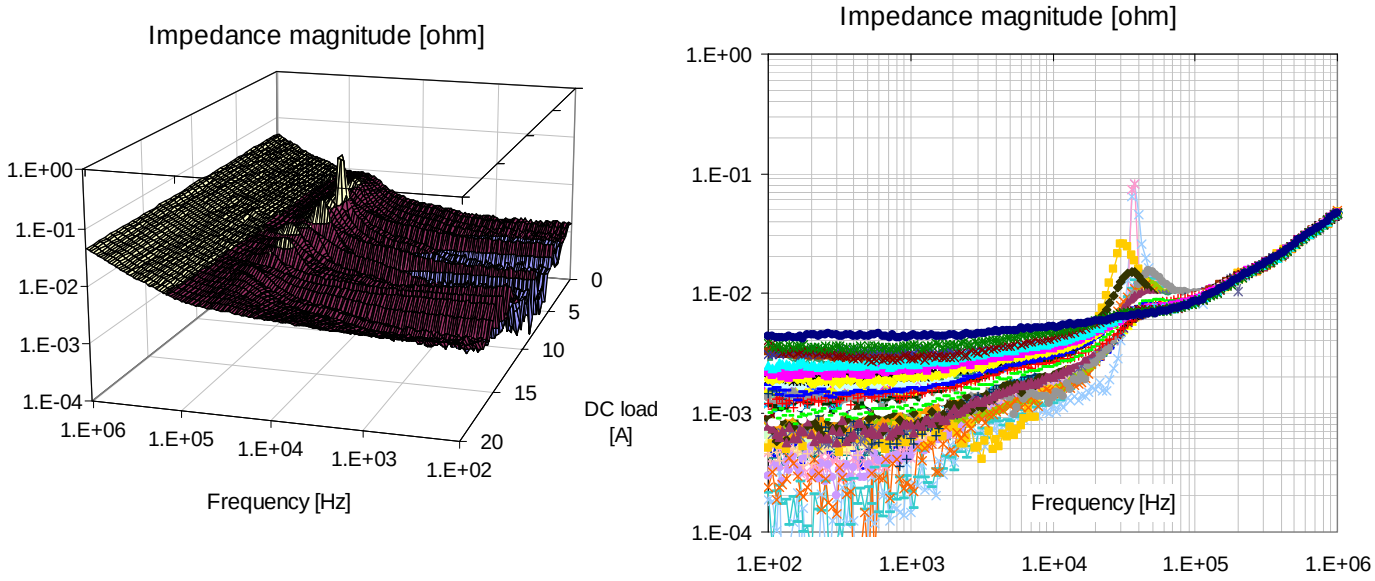

Figure 9: *Impedance surface (on the left) and cumulative traces (on the right) of a point-of-load buck converter with unstable impedance profile.*

While one could argue that the impedance profile of *Figure 9* is very flat *on the average*, and that peaks in the impedance profile occur only in narrow ranges of frequency and load-current parameters, the problem really is how to validate the design and how the performance around the sharp peaks may further degrade with time or due to part-to-part tolerances. Experience shows that converters exhibiting

such anomalies in their impedance profile also exhibit larger unit-to-unit variations, raising questions about the quality and stability of design.

3.3 High-frequency ringing

Figure 3 showed the high-frequency ringing on the output of a 5-A converter. The ringing is 37.6 mVpp in amplitude and 150 MHz in frequency. This is a more subtle problem. The first issue is that most off-the-shelf converters specify the noise unrelated to load-step changes (PARD) with 20 MHz test bandwidth. Clearly, with 20 MHz measurement bandwidth the 150 MHz ringing is completely missed. A second issue is that while the ringing signal is originated in the converter, the ringing magnitude that appears on the user's power distribution network is strongly dependent on the high-frequency impedance of the user PDN, which may also be strongly frequency and location dependent. The ringing magnitude then will be dependent on where it is measured, and it very well may happen that the noise gets *bigger* at hot spots further away from the converter.

Many times the argument to dismiss this concern is that the board should have sufficient high-frequency bypassing and it will suppress this ringing. While this is a valid argument, the ringing current leaving the converter is not specified, so the user does not have numbers to design with.

How the ringing magnitude depends on the input voltage and load current appears to be dependent on circuit and layout design of the converter. Several converters were found where the high-frequency ringing was almost independent of DC load current, but varied linearly with input voltage. The high-frequency ringing in some other designs showed very weak dependence on input voltage, but the DC load current had a strong influence. Data from such a converter is shown in *Figure 10*.

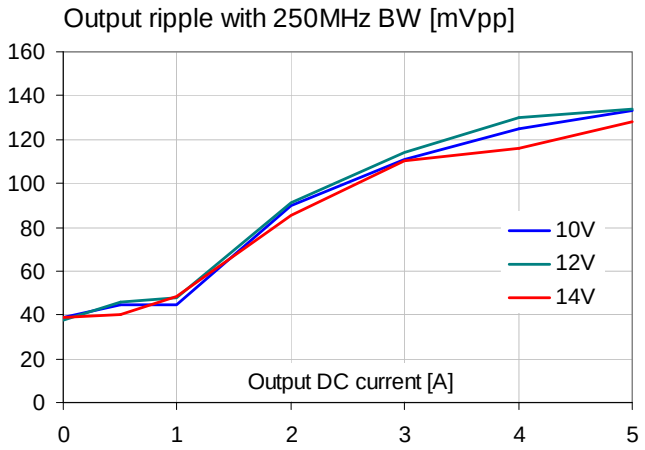

Figure 10: *High-frequency ringing magnitude in a point-of-load DC-DC down-converter as a function of input voltage and load current.*

In today's converters the ringing frequency is in the 100-500 MHz frequency range, which overlaps with several reference clock frequencies of high-performance buses. The clock signals represent a few mA current in each line, whereas the ringing is associated with the switching edge of the converters, which may represent tens of amperes. Having an unspecified noise component related to these high-current

edges that my accidentally line up with clock frequencies, represents a design risk in high-performance systems.

3.4 Challenges with digital power

Digital power is a promising new trend in power conversion. It enables not only a versatile on-the-fly control of the usual parameters like output voltage, turn-on and turn-off delays and ramps, but ultimately enable the implementation of nonlinear control loops. Non-linear control loops make it possible to significantly improve the transient behavior. Though this seems to be all and only good, there is a hidden catch: validation of a system design becomes more difficult. Validation of a high-performance system design should be never skipped, especially when off-the-shelf converters are used. Since the dynamic performance of the converter depends on the external capacitors we connect to the converter, and chances are that the converter during manufacturing is tested with a different set of capacitors, we have to validate our design *in situ*. As it was illustrated above, a well-behaving analog control loop with a properly designed set of external capacitors will show more-or-less constant impedance profile, regardless of the static load current and input voltage. This means that the validation is reduced to checking the small-signal output impedance profile over the entire input-parameter ranges, such as input voltage, output voltage and DC load current. The small-signal output impedance can be obtained in a straightforward manner [4], and because it does not require injecting a test signal in series to the control loop, in general it is easier to perform than measuring the Bode plots for the classical loop stability.

When the control loop is made purposely nonlinear, the validation can not be done any more with small-signal output impedance measurements. Nonlinear loop control may 'tighten' the loop for larger output deviations, and 'loosen' for small deviations. This means that any measured small-signal output impedance may show unrealistically large impedance values. An illustration of this is shown in *Figure 11*. Note the low-frequency impedance magnitude of about 18 milliohms. If this was linear impedance, a 1-A load-current step would produce 18 mV output transient, whereas a 10-A load step would produce 180 mV transients. A quick time-domain scan showed that this was not the case: larger load-step currents did not produce proportionally larger transients, which is the whole purpose of the nonlinear loop.

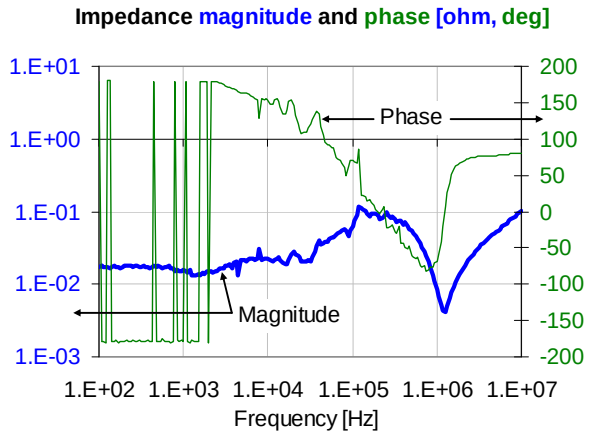

Figure 11: *Measured small-signal output impedance of a DC-DC converter with nonlinear digital loop.*

While the nonlinear control loop helps to improve the transient behavior, now we have two more variables to scan against: DC load current and load-step magnitude. The frequency parameter of the impedance plots is replaced by the slew-rate or rise time of the excitation current step. For a full validation the testing has to be done in the time domain, which due to its wide-band nature, is more prone to external noise contaminating the data.

4. Conclusions

The critical missing link in the proper PDN design is the middle frequency range, where the PDN performance is determined by the interaction of the DC-DC converter and external bypass capacitors. Bypass capacitors and DC-DC converters with well-behaved analog control loops can be modeled as linear circuits; hence their worst-case transient noise can be estimated from the inverse Fourier transform of their combined impedance profile. The worst-case transient noise can be minimized by flattening out the impedance profile. It was also illustrated that traditional phase-margin requirements do not necessarily result in optimum impedance profile.

To help proper system design with off-the-shelf DC-DC converters, the following is necessary:

- Simulation model of the converter's output stage and control loop should be available, which allows the user to properly design the external capacitors to meet user requirements
- Change the PARD specification so that the high-frequency ringing is included

It is understood that a detailed simulation model may reveal proprietary information about the converter design. A usable compromise is to have a sufficiently detailed behavioral model, which first would require the definition of a standard interface for such simulations.

It was also shown that non-linear control loops, though they can improve the noise performance significantly, increase the validation challenges.

References

[1] A. S. Kislovski, R. Redl, N. O. Sokal: Dynamic Analysis of Switching-Mode DC/DC Converters. Van Nostrand Reinhold, 1991.

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

[3] Istvan Novak, "Comparison of Power Distribution Network Methods: Bypass Capacitor Selection Based on Time Domain and Frequency Domain Performances," in TecForum TF-MP3, DesignCon2006, Santa Clara, CA, February 6-9, 2006

[4] "Frequency Domain Power Distribution Measurements - An Overview" in HP-TF1: Measurement of Power-Distribution Networks and Their Elements, at DesignCon2003 East, June 23-25, 2003, Boston, MA