Manuscript for TF-MP3 “Comparison of Power Distribution Network Design Methods” at DesignCon 2006, February 6-9, 2006, Santa Clara, CA

DesignCon 2006

Comparison of Power Distribution Network Design Methods: Bypass Capacitor Selection Based on Time Domain and Frequency Domain Performances

Istvan Novak

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

Author Biography

Istvan Novak is signal-integrity senior staff 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 Fellow of IEEE for his contributions to signal-integrity and RF measurement and simulation methodologies.

Abstract

Power Distribution Network (PDN) design of high-speed and high-power systems appears to be black magic: there are contradictory design philosophies, component selection and layout rules. Suboptimum or not proper power distribution increases the supply-rail noise, but this impacts system performance in a statistical manner, convoluted with many other variables. For this reason, over a wide range, different PDN designs may appear to perform equally well, giving rise to several misconceptions. This paper compares some of the popular board-level PDN design methods by their impedance profiles and worst-case transient noise. It is shown that smooth R-L type self-impedance profiles result in the lowest transient noise. Component placement is also analyzed, and it is concluded that bypass-capacitor location matters mostly in PDNs, where the impedance of the power/ground planes does not match the target impedance. Finally, measured impedance profiles are shown for some of the popular SUN servers, illustrating some of the possible PDN design philosophies.

I. Introduction

With the growing number of different supply rails, continuously increasing currents, shrinking supply voltages and supply-rail noise allowances, the Power Distribution Network (PDN) has become one of the most challenging subsystems in today's high-end systems. Figure 1 shows three illustrations from SUN servers. The photo on the top left shows the CPU module from a SUN Enterprise 450 server. The 4"x6" board, designed in the mid 1990's, had one CPU, two major supply rails, and a total of 351 bypass capacitors. The photo on the top right shows the CPU module from a V890 server. The 10"x20" board, designed in the early 2000's, has two quad-core CPUs, eight major supply rails, and a total of 1907 bypass capacitors. The bottom photo shows a T2000 server, with a 8"x10" CPU board, recently designed. It has one UltraSparc T1 CPU, six major supply rails, and a total of 596 bypass components.

Figure 1. Three illustrations from SUN servers. Top left: E450 CPU module, approximately 4"x6" in size, designed in the mid 90's. Top right: V890 CPU module, approximately 10"x20" in size, designed in the early 2000's. Bottom: T2000 2U server, with 8"x10" CPU module, recent design. Note that the scales are different for the three photos.

Part of the challenge is that PDNs possibly have to serve a set of diverse functions in the systems [1]:

- First and foremost, to provide clean power to the active devices
- Optionally to serve as return path for signals
- And to ensure that radiation related to PDN does not violate legal limits

The simple sketch in Figure 2 is an illustration of a board-level PDN with one supply rail. In reality, most of the time we have multiple supply rails in a system, each having different requirements, but being inter-related by shared planes and cross couplings among elements. In a complex system, the above three major functions can also be inter-related, and dependent on the nature of the PDN, one or two of the major functions may not be present or may be irrelevant for a particular rail.

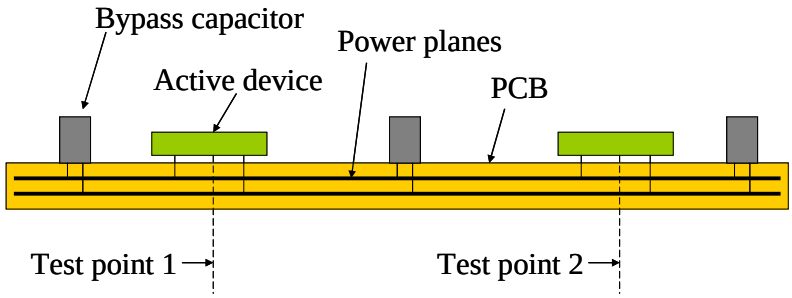

Figure 2. Sketch of board-level PDN and its major components.

[1] outlined two major classes of PDNs: core PDN and IO PDN. In its simplest form, a core PDN feeds only one active device with a dedicated network. If, moreover, in the implementation the core plane is sandwiched between neighboring ground planes, the core plane does not serve as signal return path, and being between ground planes, EMI radiation may also be less of a concern. The push for more cost-effective designs, however, many times results in shared PDNs: a PDN, with its primary function of feeding one or more cores or distributed IO networks, may also serve as reference plane for slow or high-speed signals. In such situations the functions have to be looked at separately, and the most restrictive will have to drive the design. For instance, a core may be able to live with more noise on the rail than what is being allowed for some high-speed reference planes. Similar considerations apply to the EMI limits: since core distributions through packages behave like low-pass filters, resonances of board PDN at several hundred MHz or above will have very little influence on the power delivery to the core, but it may be detrimental for EMI radiation and/or for signal-return functions.

The requirements of supplying clean power to the chips and the function of reference path for high-speed signals both can be captured by the maximum transient noise as a result of worst-case activity of the system. With the exception of low-end systems, where occasional failure may be acceptable, we usually have to keep the *worst-case peak-to-peak transient noise* under control, so that we can maintain the proper operating conditions of the chips and limit simultaneous switching noise, jitter and inter-symbol interference introduced by the rail noise through the return path. Since most PDN components can be fairly accurately described by linear models, the worst-case peak-to-peak transient noise can be obtained from the impedance matrix of the PDN. Assuming an arbitrary sequence of step-like excitation current with a fixed rise time, the worst-case excitation pattern and peak-to-peak transient noise can be obtained by processing the network functions of the PDN [2].

CPU clock frequencies and signaling rates used today push more features on the board into the red zone, such that their relative dimensions with respect to the wavelength of clock signals approach or exceed

the quarter-wave limit. Quarter-wave or longer structures may become effective primary or secondary radiators; therefore the PDN has become one of the primary EMI risk factors. It is generally understood that the best way to avoid EMI radiation problems from PDNs is to ensure their *resonance-free* impedance profile. The resonance-free impedance profile also helps to minimize SSN and jitter of high-speed signals, when the PDN serves as signal reference.

In a good industrial design, we want to ensure that the system, and within the system the PDN, will properly function over the entire foreseeable range of parameter variations, such as component tolerances and aging. It is an added bonus if the design is less sensitive to missing and/or broken components. In short, low *sensitivity* is preferred.

To improve long-term reliability of the PDN, we also need to ensure that we have as uniform *stress distribution* among the PDN components as possible.

The constant push for more cost-effective solutions results in optimizing the margins and removing the extra padding of performance in inter-related PDN subsystems. This may result in designs, where changing one dimension or one component value may have a ripple effect and a whole chain of components and dimensions have to be re-aligned and re-designed. Therefore *portability* of the design is becoming increasingly valuable and important.

In addition to the above electrical requirements, we may have a set of important non-electrical parameters and targets to meet, among others: cost, size (area, height), weight, component connection style (through hole, surface mount), placement restrictions, simplicity of Bill of Material, etc.

When we design the board PDN, we have to select the DC sources, bypass capacitors, PCB planes and stackup, and more importantly, we need to find the proper location for all of the parts. The primary questions about the bypass capacitors of a board-level PDN:

- What values of bypass capacitors should we use and how many of them?
- Where should we place the components?

Note that selecting the proper DC sources (DC-DC converters) is becoming a challenge by itself, but we do not address that question in this paper.

II. So what is the metric anyway

Because we usually do not know the transient noise current excitation very accurately, it is customary instead to design the PDN to meet a required impedance profile [3]. For single-load and small-size core PDNs, not serving the additional task of reference path for high-speed signals, the requirements can be eventually distilled to the self-impedance profile provided to the silicon core. In this case the transfer functions among various points of the PDN are less important. For large, distributed core and IO PDNs, possibly with multiple noise sources, the description has to be done by full impedance matrices: both self and transfer-impedance terms at the key locations are important.

Since we can derive the worst-case transient noise from the impedance matrix, it may appear that the impedance profile of PDN is the key parameter. While the impedance profile is very important, we should not loose sight of the original requirement, which was the *worst-case transient noise*. As we will show soon, impedance profiles which *may look better* than others may in fact be *worse and generate more noise*.

In this paper we will analyze some of the popular PDN design methods by first comparing their performances based on lumped self-impedance profiles. Later we look at some of the important distributed aspects, such as component placement and modal resonances. The analyzed approaches will be illustrated by simulated and measured performance data.

III. Comparison of popular methods based on lumped self-impedance

In the early days of PDN design, the focus was on the number of 0.1uF capacitors to be placed close to the active devices' leads. This approach was sufficient as long as the active devices did not generate significant transient noise spectrum above the Series Resonance Frequency (SRF) of the bypass capacitors. Higher system speeds later created the need to design the PDN systematically, to meet certain impedance requirements. Here we compare four possible ways of synthesizing a lumped self-impedance profile:

- "Multi-pole" (MP) [3]
- "Capacitors-by-the-decade" (CBD) [4]
- "Big-V" [4]
- "Distributed Matched Bypassing" (DMB) [5]

Note that in terms of impedance profile, CBD and Big-V are variations of MP and DMB. As it will be shown, there is no clear boundary between these design approaches; we can gradually transform one into the other by varying parameters and there are an infinite number of variants in between. The following illustrations in this section depict simulated scenarios with hypothetical component values. Implementation details will be given later together with measured illustrations.

We start with DMB, which calls for a flat impedance profile from DC up to a certain corner frequency, beyond which frequency the impedance may rise linearly, following the reactance of the PDN inductance. For the sake of simplicity, we assume that we want to synthesize 10 milliohms of impedance, and the corner frequency is 10MHz, corresponding to 150pH of PDN inductance. The impedance profile and the corresponding step response are shown in Figure 3.

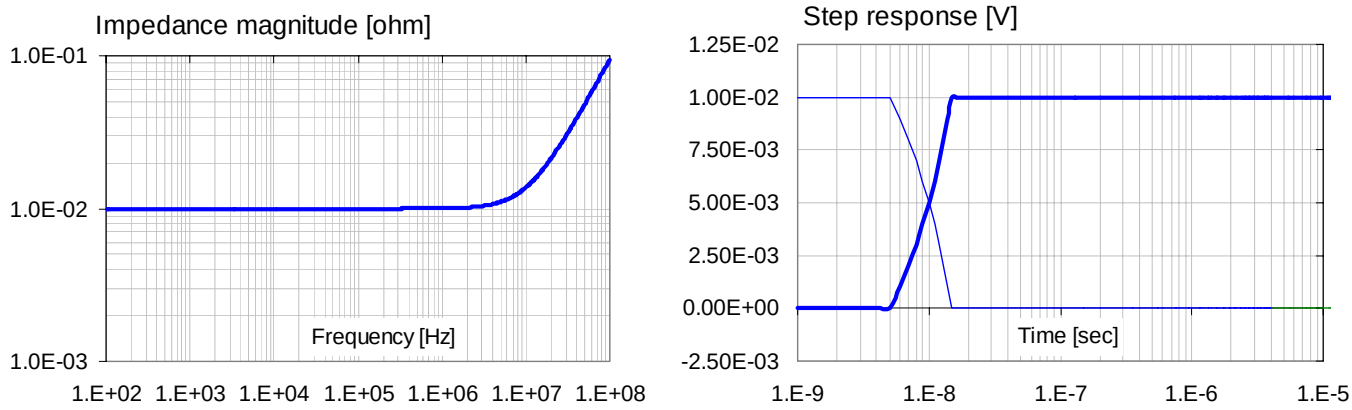

Figure 3. Impedance profile and step response of a 10-milliohm PDN based on Distributed Matched Bypassing with AVP.

The horizontal scales both in the frequency and time domains are logarithmic. On a logarithmic frequency scale we cannot show the DC value, but we assume that it is the same 10 milliohms value.

The finite DC output resistance is generated by Adaptive Voltage Positioning (AVP) [6] in the DC source. We further assume that we generate the step response with a 1A 10-ns rise-time step excitation after a 10MHz single-pole low-pass filter. The 10MHz high-pass and 10MHz low-pass filter functions cancel each other, which restores the 10 ns piece-wise-linear response. This is what we see on the right in Figure 3. The graph shows the positive and negative step responses. Note again that the horizontal scale is logarithmic, which is used here and on subsequent step-response graphs to enable us to see details close to the initial edge as well as changes later in the step response. The step-response is monotonic, there is no overshoot/undershoot or ringing, therefore the worst-case transient noise for any arbitrary sequence of step excitation with 10ns or slower rise time with 10 ns or slower step excitation will be 10 mVpp for each ampere of current-step magnitude. This 10 mVpp/A number serves as a baseline in comparing other designs.

In case we decide to use high DC gain in the DC source instead of AVP, the DC output resistance will be low. Figure 4 shows the impedance profile with 1 milliohm DC resistance and the step response with the same filtered 10 ns step excitation.

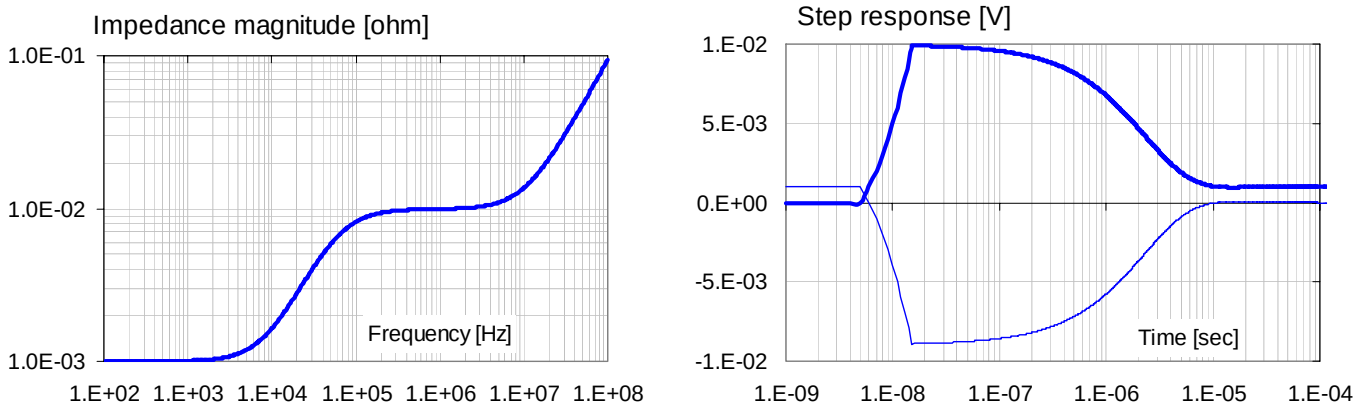

Figure 4. Impedance profile and step response of a 10-milliohm PDN based on Distributed Matched Bypassing with no AVP.

The step response is still smooth, with no ringing, but with respect to its steady state, it has a 9 mV overshoot. The worst-case peak-to-peak transient noise thus becomes two times the overshoot plus the steady-state swing: 2 * 9 + 1 = 19 mVpp.

Next, we look at an MP design with AVP. We assume five different capacitor values, covering one decade of frequency in the range of 1 – 10 MHz. We have four capacitance values per decade, and each part is assumed to have a Q of approximately 3. As shown in Figure 5, the impedance and step responses in this case have multiple peaks and ringing. We need to follow the procedure in [2] to get the worst-case peak-to-peak transient noise, which comes out as 15.7 mVpp/A. The assumed component values and the detailed impedance profile, showing the impedance magnitudes of each contributing capacitor banks as well, are shown in Figure 6. Figure 7 shows the impedance profile and step response for the same Multi-pole PDN design with high DC gain in the DC source. The worst-case peak-to-peak transient noise increases to 25 mVpp/A. Figure 8 shows the component values and the detailed impedance profile. By comparing the component values in Figures 6 and 8, we can see that only the DC source (R-L_ and the bulk capacitor (C1) changed, the other component values are identical in the two designs.

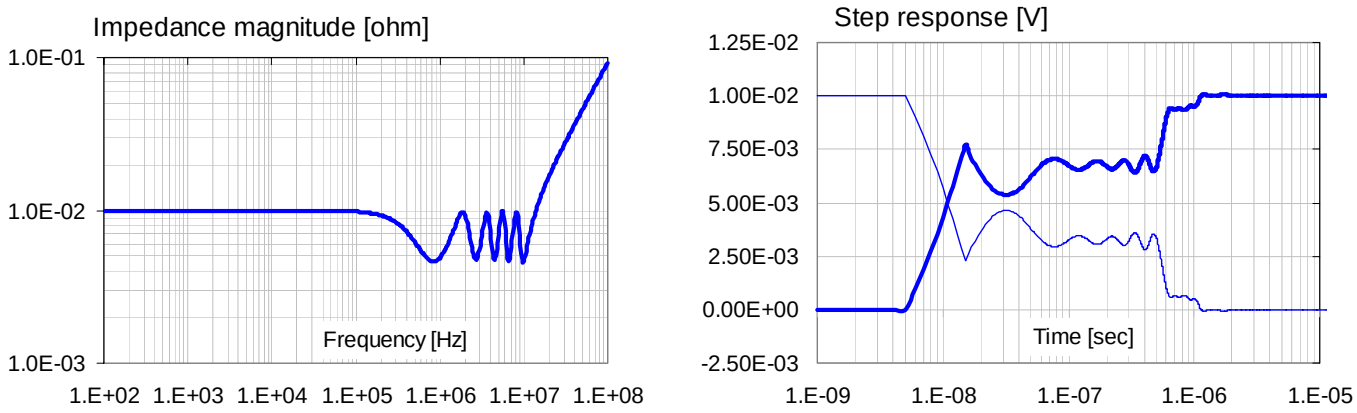

Figure 5. Impedance profile and step response of a 10-milliohm MP PDN with AVP.

	C [uF]	ESR [ohm]	L [nH]
R-L	-	0.01	2
C1	33	0.006	1
C2	3.6	0.006	1
C3	1.3	0.006	1
C4	0.75	0.006	0.8
C5	0.56	0.006	0.5

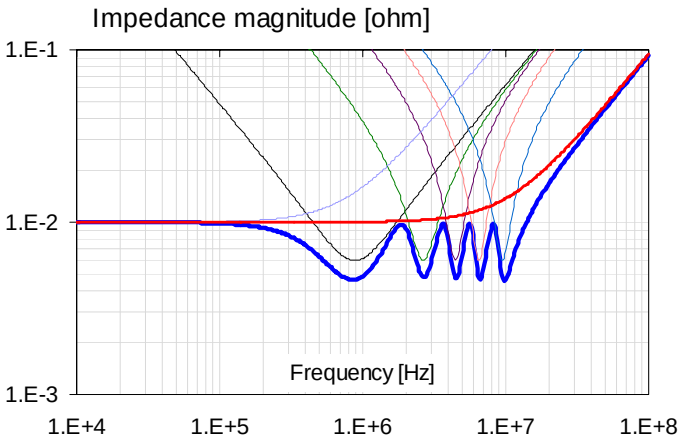

Figure 6. Component values and detailed impedance profile for the PDN shown in Figure 5.

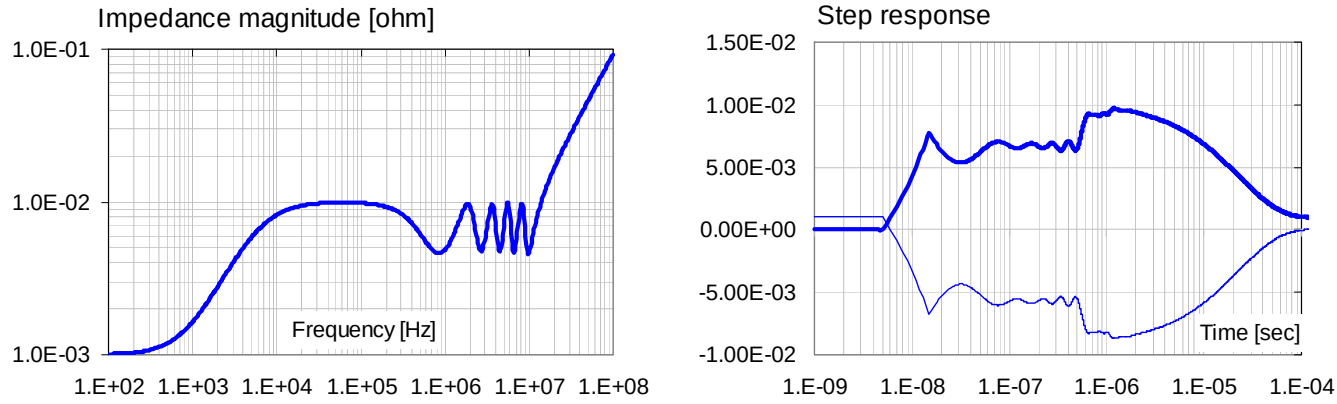

Figure 7. Impedance profile and step response of a 10-milliohm MP PDN with high DC gain.

	C [uF]	ESR [ohm]	L [nH]
R-L	-	0.001	20
C0	1000	0.01	2
C1	33	0.006	1
C2	3.6	0.006	1
C3	1.3	0.006	1
C4	0.75	0.006	0.8
C5	0.56	0.006	0.5

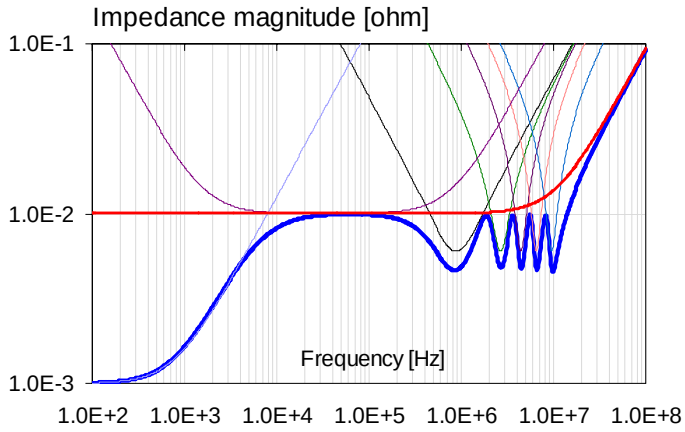

Figure 8. Component values and detailed impedance profile for the PDN shown in Figure 7.

We continue with a Big-V illustration, and skip the CBD as it is a variant of the MP design approach, and it will be illustrated with its measured performance. Figure 9 shows the impedance profile and step response of a 10 milliohm Big-V design with AVP. This impedance profile assumes a 10 milliohm 1.5 nH DC source, in parallel to a multitude of bypass capacitors, cumulatively generating a C = 30 uF, ESR = 1E-3 Ohm, ESL = 165 pH impedance. The step response has one pronounced dip and one peak, resulting in a worst-case transient peak-to-peak noise of 21.4 mVpp/A.

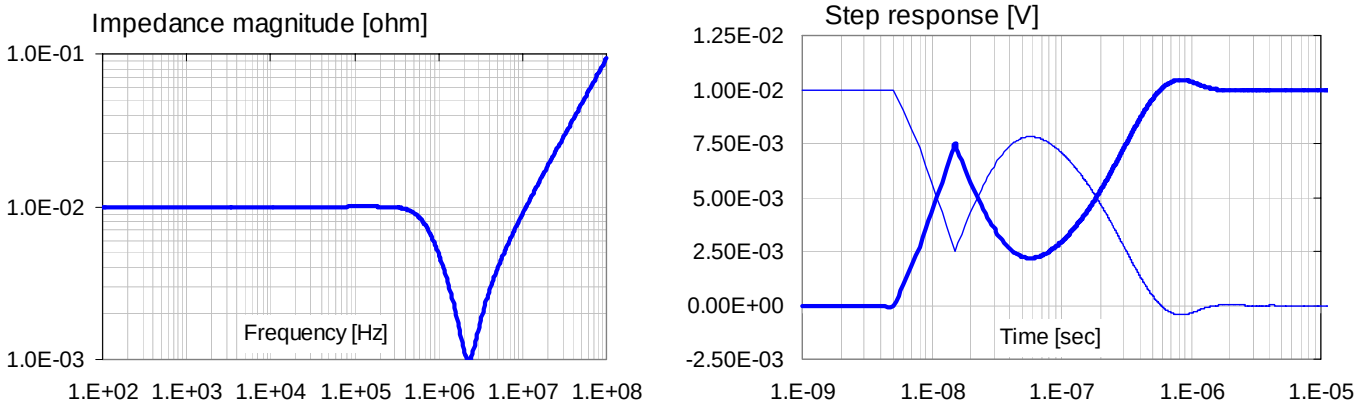

Figure 9. Impedance profile and step response of a 10-milliohm Big-V PDN with AVP.

If we use high DC gain with the “Big-V” impedance profile instead of AVP, we get the impedance and step responses shown in Figure 10. This impedance profile can be obtained with the same set of bypass capacitors, by replacing the 10 milliohm 1.5 nH DC source with the combination of a 1 milliohm 100nH DC source and C = 10,000uF, ESR = 1E-2 Ohm, ESL = 1.5 nH bulk capacitor. This impedance profile results in a 27.5 mVpp/A worst-case transient noise.

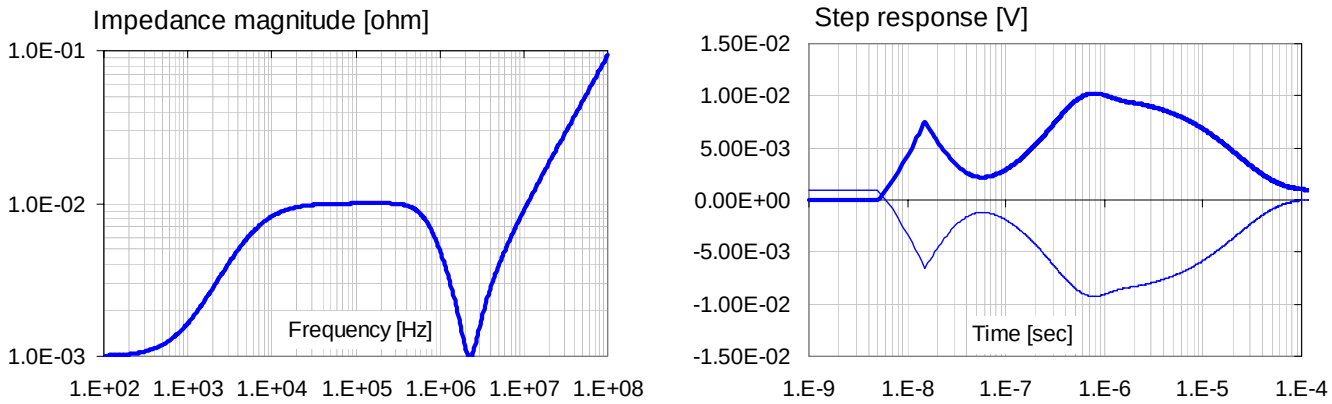

Figure 10. Impedance profile and step response of a 10-milliohm Big-V PDN with no AVP.

When we compare the impedance and step responses with the worst-case noise performance numbers, we can notice that the 10 milliohm 150pH DMB impedance profile of Figure 3 resulted in the lowest noise number, even though this impedance response is an upper bound to all consecutive impedance profiles. This suggests that making the impedance profile ‘*better*’ by driving dips to make the impedance profile meet an impedance target over a wider frequency range, creates more noise. We can clearly show this by comparing the impedance and transient response profiles on one graph. Figure 11 compares the MP, Big-V, and DMB designs with AVP.

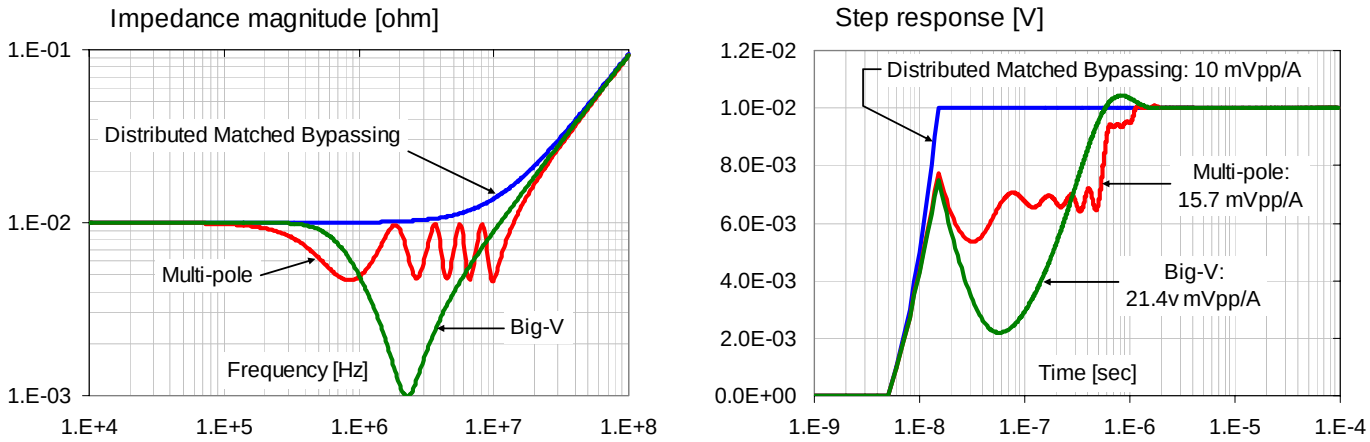

Figure 11. Comparison of impedance profiles and step responses of Multi-pole, Big-V, and Distributed Matched Bypassing design approaches with AVP.

From the comparison in Figure 11 we may be inclined to conclude that the Big-V design is generally worse than MP. This generalization, however, is again not correct; eventually the worst-case transient noise depends on the number and depth of dips and peaks in the impedance profile. For the comparison of Figure 11 the MP design performs better than the Big-V, because we assumed a 2:1 ripple in the MP impedance profile (5 milliohms minima and 10 milliohms peaks) as opposed to a 10:1 ratio with the Big-V design. If we can manage to reduce the impedance variation of the Big-V design, the worst-case transient noise will improve. Figure 12 first shows the result of ‘partial’ AVP, where we change the DC resistance of the source. As the DC resistance approaches the mid-frequency AC impedance, the overshoot in the transient step response becomes smaller.

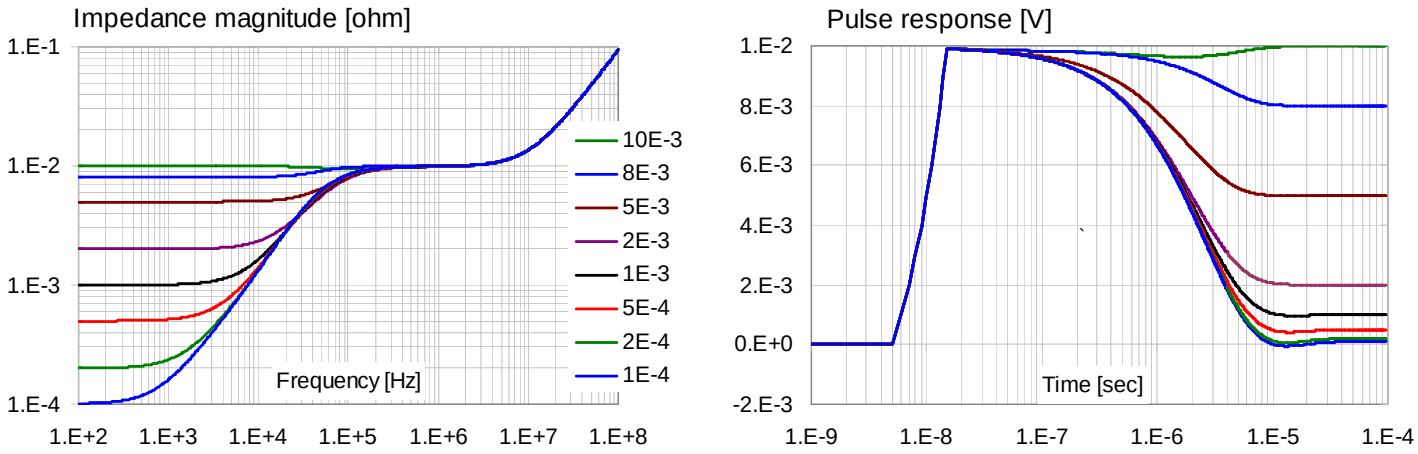

Figure 12. DMB design with different DC resistances. The color scale on the impedance plot also refers to the traces on the step-response plot.

We can also look at impact of changing the value of the impedance minimum in the Big-V design. The impedance and step responses of Figure 13 span the minimum impedance in a 100:1 ratio. Note that the deeper the minimum, the more fluctuation we get in the step response, hence we get larger worst-case transient noise.

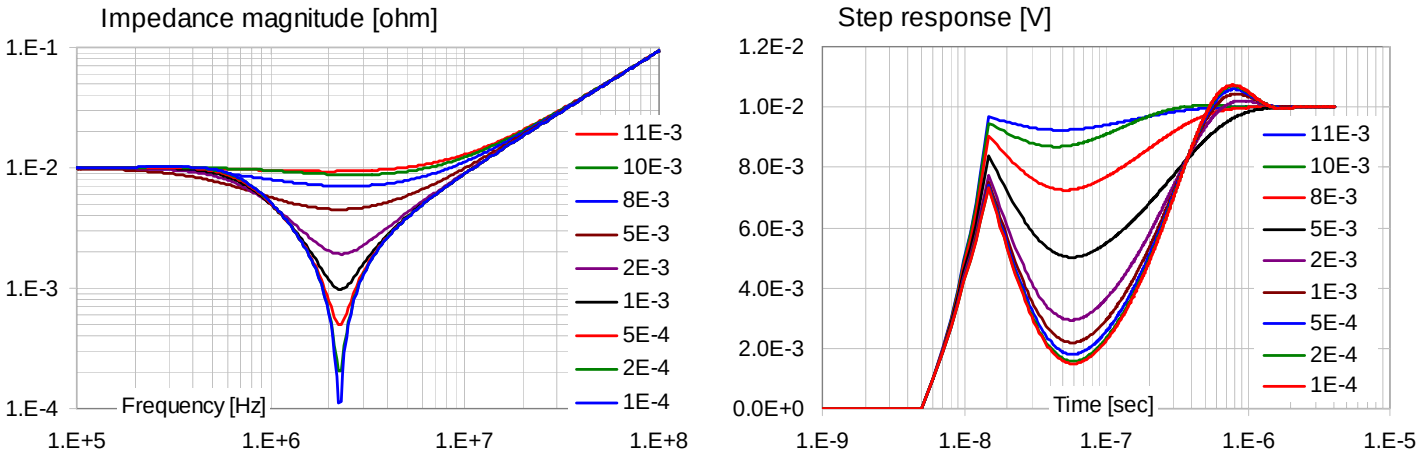

Figure 13. Big-V design of Figure 9, except ESR of the capacitor bank generating the impedance dip varies between 0.1 to 11 milliohms.

We can ask similar questions about the MP design. How can we control the fluctuations in the step response; by increasing the number of capacitance values per decade (N), or by reducing Q of each capacitor bank? The correct answer is not one or the other, but a combination of the two, because eventually what matters is the peak/valley ratio in the impedance profile. We can use lower N with lower-Q parts, or we can afford higher-Q parts if N is higher. This is illustrated in Figure 14. The charts assume the circuit of Figure 6, with five capacitor banks. ESR in each bank was held constant, and the capacitance and inductance values were varied to get different Q and N values. The series resonance frequencies of capacitor banks were allocated logarithmically on the frequency scale; this, for a given set of N and Q, results in equal-height peaks and dips in the impedance response. The left plot in Figure 14 shows the impedance peak values normalized to the common ESR as a function of Q and N. The plot

on the right shows the impedance dip values normalized to the common ESR. Note that the values on the right are less than one, indicating that the impedance minimum of parallel-connected capacitors with closely spaced SRF is below their respective ESR values. We can see from these plots that for a given ratio of impedance maximum and minimum we can trade N for Q and vice versa.

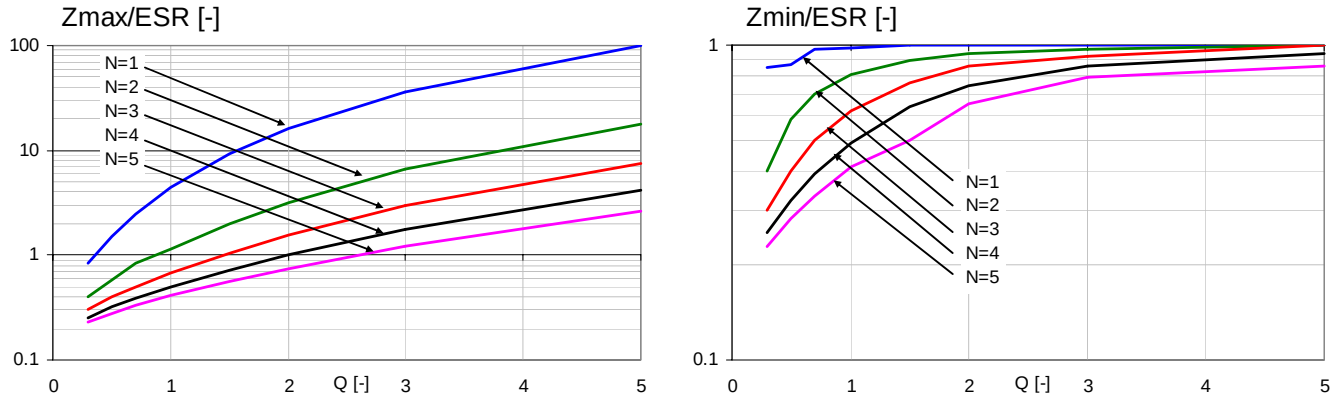

Figure 14. Normalized impedance maximum and minimum values of the Multi-pole design as a function of the number of capacitance values per decade (N) and Q of capacitors.

IV. Component placement

The previous simulated examples assumed lumped component connections: delays and impedance among the components were neglected. In addition to the proper selection of bypass capacitor values, another hotly debated question is where to place them. Most application notes recommend placing bypass capacitors close to the active device’s supply pins. With high-power area-connection packages (PGA, BGA, LGA) implementing this advice is becoming very hard. So the designer faces the question: how close should we put the bypass capacitors? Conventional wisdom argues that we need to place bypass capacitors close to the active device to have charge available quickly when the device needs it for fast current transients [7]. This argument, however, ignores the fact that charge can be *stored* also by inductors, moreover, charge can also be *supplied* by transmission lines. Ultimately all what matters is that the active device should be connected to an impedance low enough that the current transient will generate noise voltages within our specified limits. Therefore, if we assume a matched transmission line with suitable characteristic impedance feeding the active device, what matters is that this transmission line should reach the active device; in which case the locations of charge-storage elements do not matter.

In power-distribution networks, the power planes and/or puddles behave like two-dimensional transmission lines. The approximate impedance of a rectangular plane pair is [5]:

$$Z_p = \frac{532}{\sqrt{\varepsilon_r}} \frac{h}{P} \qquad (1)$$

where Z_p is the approximate plane impedance in ohms, h and P are the plane separation and plane periphery, respectively, in the same but otherwise arbitrary units. For instance, a 10”x5” plane pair with 12-mil plane separation yields approximately 0.1 ohms impedance. The impedance can be set for instance by 60 pieces of RC components placed around the plane periphery, one at every half-inch interval. If we select R=5.6 ohms and C=0.1 uF, the sixty pieces in parallel will give a flat impedance profile with about 0.1 ohms. The 12-mil plane separation is very easy to ensure in a stackup; in fact we have room also to place one or even two signal layers between these power/ground planes. If we need

lower than 0.1 ohms impedance, we can place the planes closer. Without extra cost we can get down to 3-mil plane separation, yielding approximately 25 milliohms impedance on the 10"x5" planes. For added cost we can get plane separations down to 1-mil or 0.5-mils ([8], [9], [10]), resulting in a plane impedance of 12.5 and 6 milliohms, respectively.

But what happens if we need a 1-milliohm supply-rail impedance? To get it with matched planes, the plane separation would need to shrink to thin-film dimensions, not practical for large rigid boards today. Or what happens if we need just 10 milliohms, but we want to maintain for instance 10-mil plane separation (for cost or maybe other reasons). In these cases, we can still create the necessary low and mid-frequency impedance with bypass capacitors, but if these bypass capacitors are connected to the active device with higher-impedance planes, *now* the capacitor placement does matter.

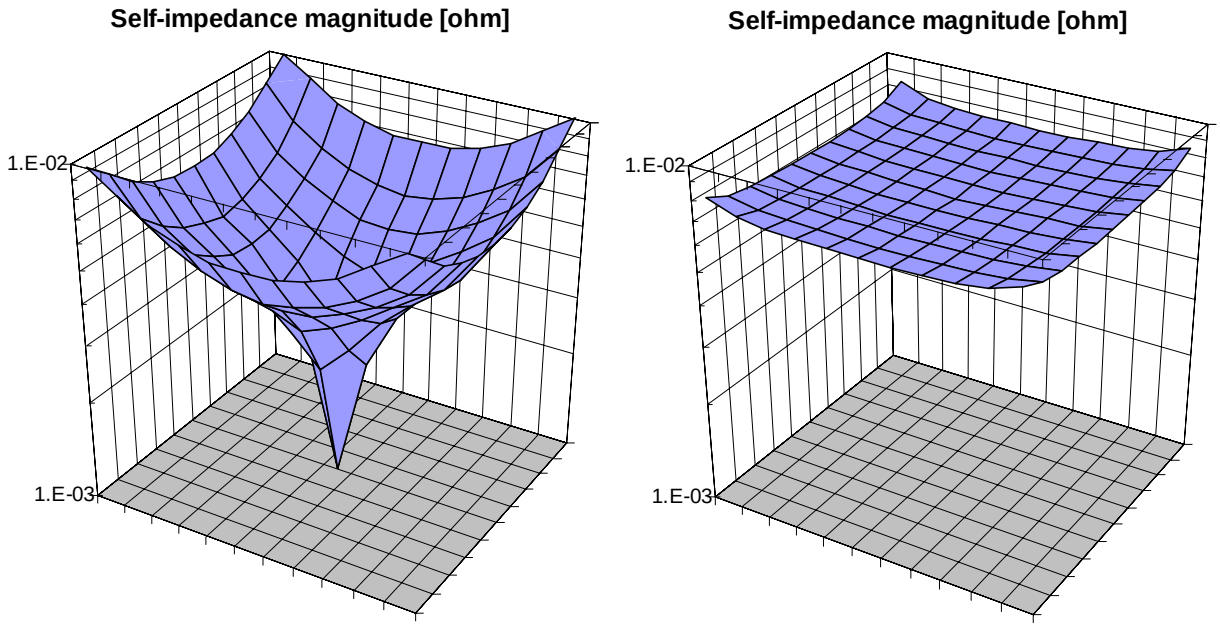

Figure 15. Self-impedance at 0.5 MHz on a 2"x 2" plane pair with one capacitor placed in the middle. Left: 50 mils dielectric separation, with a 100 uF, 1 milliohm, 1 nH capacitor. Right: 2 mils dielectric separation, with 100 uF, 7 milliohm, 1 nH capacitor.

As shown in [11], if we have terminated planes, the impedance on the planes will not strongly depend on where we are on the planes, therefore component placement is largely irrelevant. If, however, we need impedance lower than what we can conveniently create with power planes, component placement does matter.

The above points are illustrated by impedance plots generated by a spreadsheet illustration tool [12]. We assume one piece of capacitor located in the middle of a pair of 2"x2" planes. For this example, we use C = 100 uF, ESR = 0.001 Ohm, ESL = 1 nH. The SRF of this part is 0.5 MHz. The surface plots of Figure 15 show the variation of self-impedance magnitude over the plane at 0.5 MHz. The gray bottom area of the graph represents the top view of the planes. The grid on the bottom area shows the locations where the impedance was calculated: the granularity was 0.2 inches. The logarithmic vertical scale shows the impedance magnitude between 1 and 10 milliohms.

The impedance surface on the left has a sharp minimum in the middle; here the capacitor forces its ESR value over the plane impedance. However, as we move away from the capacitor, the impedance rises very sharply. At 0.2 inches away, the impedance is approximately 50% higher; 0.4 inches away, the impedance magnitude doubles. At the corners of the 2"x2" plane shape, the impedance magnitude is almost 5 milliohms. When changing either the plane impedance or the ESR of capacitor so that their values are closer, the variation of impedance over the plane shape gets smaller. The plot on the right shows the impedance surface of the same plane shape and same capacitor in the middle, except we increased its ESR from 1 to 7 milliohms and decreased the plane separation from 50 mils to 2 mils. Now the impedance surface at SRF varies only about 10% over the plane area.

V. Implementation examples

The implementation examples are taken from the server modules shown in Figure 1. The MP PDN design is illustrated in Figure 16. On the left the percentage usage versus capacitance values shows an approximately uniform distribution, characteristic to MP PDNs. The measured impedance profile on one of the supply rails is shown on the right. The finite granularity of available capacitance, ESR and ESL values results in unequal peaks and valleys. ESR of MLCC parts with several uF and higher capacitance is approaching a couple of milliohms, making the synthesis of 10 milliohms and higher target impedances difficult. Also, having just one or two capacitors covering each narrow frequency range comes with higher sensitivity to component tolerances, to missing or broken components, and results in high possible component dissipation.

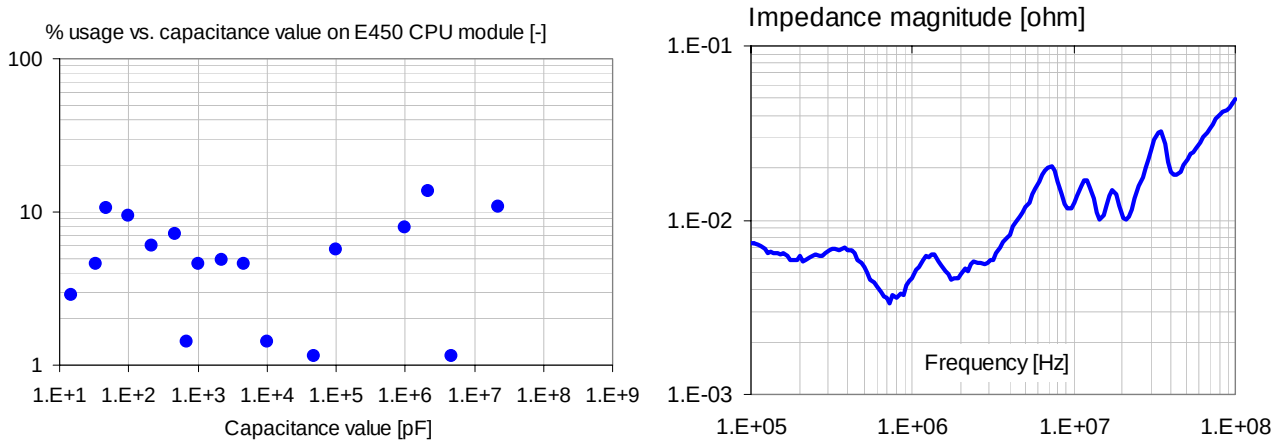

Figure 16. Distribution of capacitance values and measured self-impedance profile on the E450 CPU module.

Figure 17 is an illustration of the CBD PDN design from one of the V890 supply rails. The left graph shows the capacitance distribution over the entire CPU module. The impedance plot on the right is the measured self impedance of a high-current rail. Low-Q bulk capacitors were complemented by a set of 10uF and 1uF MLCC parts. The choice of ceramic capacitors was dictated by the required low cumulative inductance requirement, and the restrictions on component placement. Note that the minima from MLCCs do not go deep; this minimizing the transient noise.

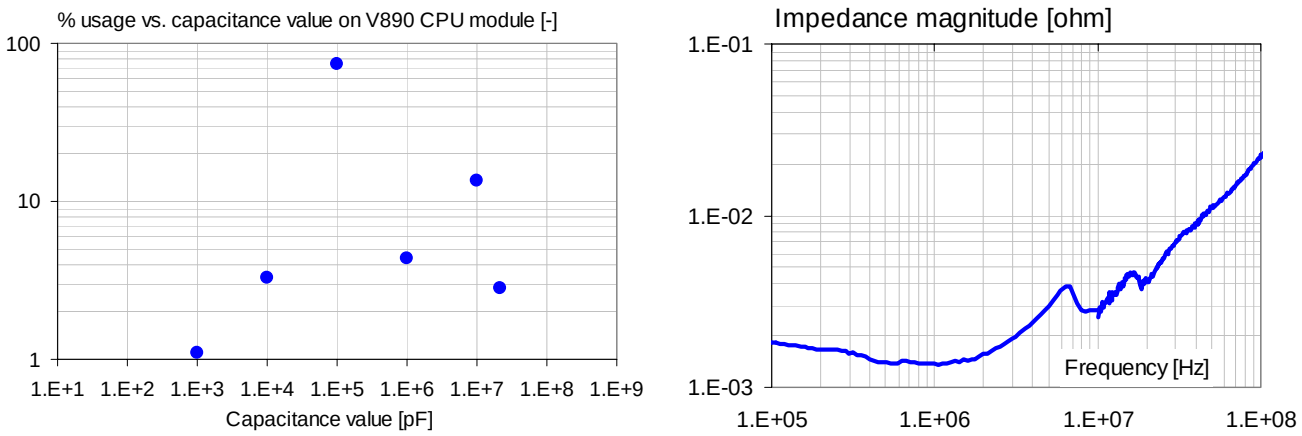

Figure 17. Distribution of capacitance values and measured self-impedance profile on the V890 CPU module.

Figures 18 and 19 are illustrations from the T2000 CPU module. The graph on the left of Figure 18 shows the capacitance distribution over the entire module. The lowest value is 0.1uF MLCC used for plane termination. Otherwise each supply rail has just one MLCC value. The impedance plot on the right of Figure 18 is an illustration of the Big-V PDN design for a medium-current rail.

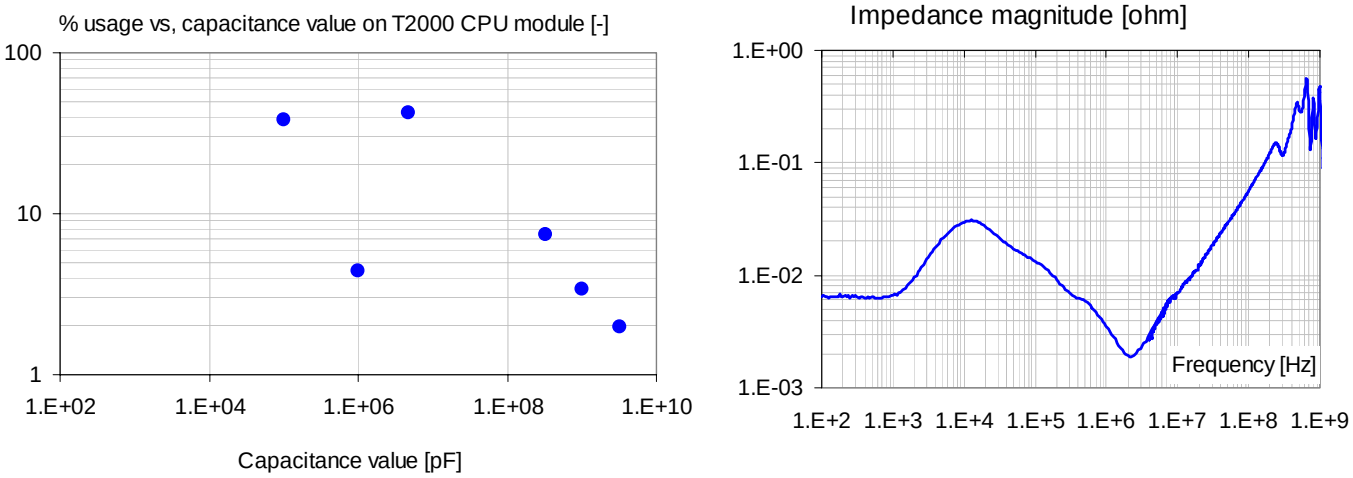

Figure 18. Distribution of capacitance values and measured self-impedance profile on the T2000 CPU module, low-current rail with Big-V design.

Figure 19 is an illustration of a high-current rail. In addition to low-Q bulk capacitors, only one value of MLCC was used, formally creating a Big-V design. The high-current rail had multiple power planes in the stackup. A combination of horizontal plane resistance, component placement and plane allocation made it possible to almost completely eliminate the impedance minimum at the SRF of the MLCC parts, effectively creating a flat R-L type response, characteristic to DMB. With zero horizontal plane resistance between capacitors, the simulated impedance minimum would be 0.06 milliohms. This supply rail was also a signal reference, the modal plane resonances were suppressed by RC termination along the plane edge. The graph on the right compares the measured impedance profiles with and without Dissipative Edge Termination (DET) components. Note the significant reduction of impedance peaks in spite of the fact that the Big-V design results in a very low cumulative ESR.

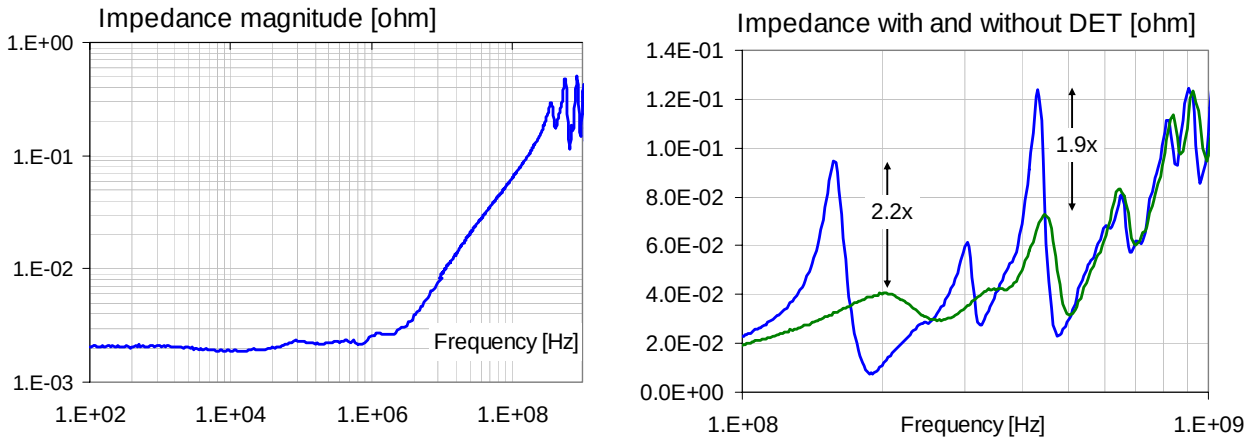

Figure 19. Left: measured self-impedance profile on the T2000 CPU module, high-current rail with Big-V design. Right: high-frequency impedance profile of the same rail with and without Dissipative Edge Termination.

Conclusions

It was shown that, if done properly, each of the popular PDN design methodologies can be made to work. The Multi-pole approach has high flexibility in component selection, but results in higher sensitivity and higher transient noise. The Big-V approach has simple BOM, low sensitivity, but also higher transient noise. The Distributed Matched Bypassing results in very low sensitivity, simple BOM and the lowest possible noise.

References

[1] Ravi Kaw, et al, “Towards Developing a Standard for Data Input/Output Format for PDN Modeling & Simulation Tools,” IEEE EMC 2005 Symposium, August 8-12, 2005, Chicago, IL.

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

[3] Larry D. Smith, et-al, "Power Distribution System Design Methodology and Capacitor Selection for Modern CMOS Technology," IEEE Tr. AdvP, Vol. 22, No. 3, August 1999, pp284-291

[4] Steve Weir, "Bypass Filter Design Considerations for Modern Digital Systems, A Comparative Evaluation of the Big “V”, Multi-pole, and Many Pole Bypass Strategies," DesignCon East 2005

[5] Istvan Novak, Leesa M. Noujeim, Valerie St. Cyr, Nick Biunno, Atul Patel, George Korony, Andy Ritter, "Distributed Matched Bypassing for Board-Level Power Distribution Networks," IEEE Transactions on Advanced Packaging, Vol. 25, No.2, May 2002, pp.230-243

[6] Alex Waizman, Chee-Yee Chung, “Extended Adaptive Voltage Positioning (EAVP),” Proc. of EPEP conference, October 23-25, 2000, Scottsdale, AZ

[7] Bruce Archambeault, “The Effect of Decoupling Capacitor Distance on Printed Circuit Boards Using both Frequency and Time Domain Analysis,” Proc. of IEEE EMC Symposium, 2005

[8] http://www.mcm.dupont.com/MCM/en_US/Products/embedded_passives/emb_pass_pcl.html

[9] http://www.oakmitsui.com/pages/advancedTechnology/faradFlex.asp

[10] http://www.3m.com/us/electronics_mfg/microelectronic_packaging/materials/

[11] I. Novak, “Tolerance Calculations in Power Distribution Network,” XCell Summer 2004, Xilinx

[12] http://home.att.net/~istvan.novak/tools/Caprange_rev10.xls