FOCUSED ISSUE FEATURE

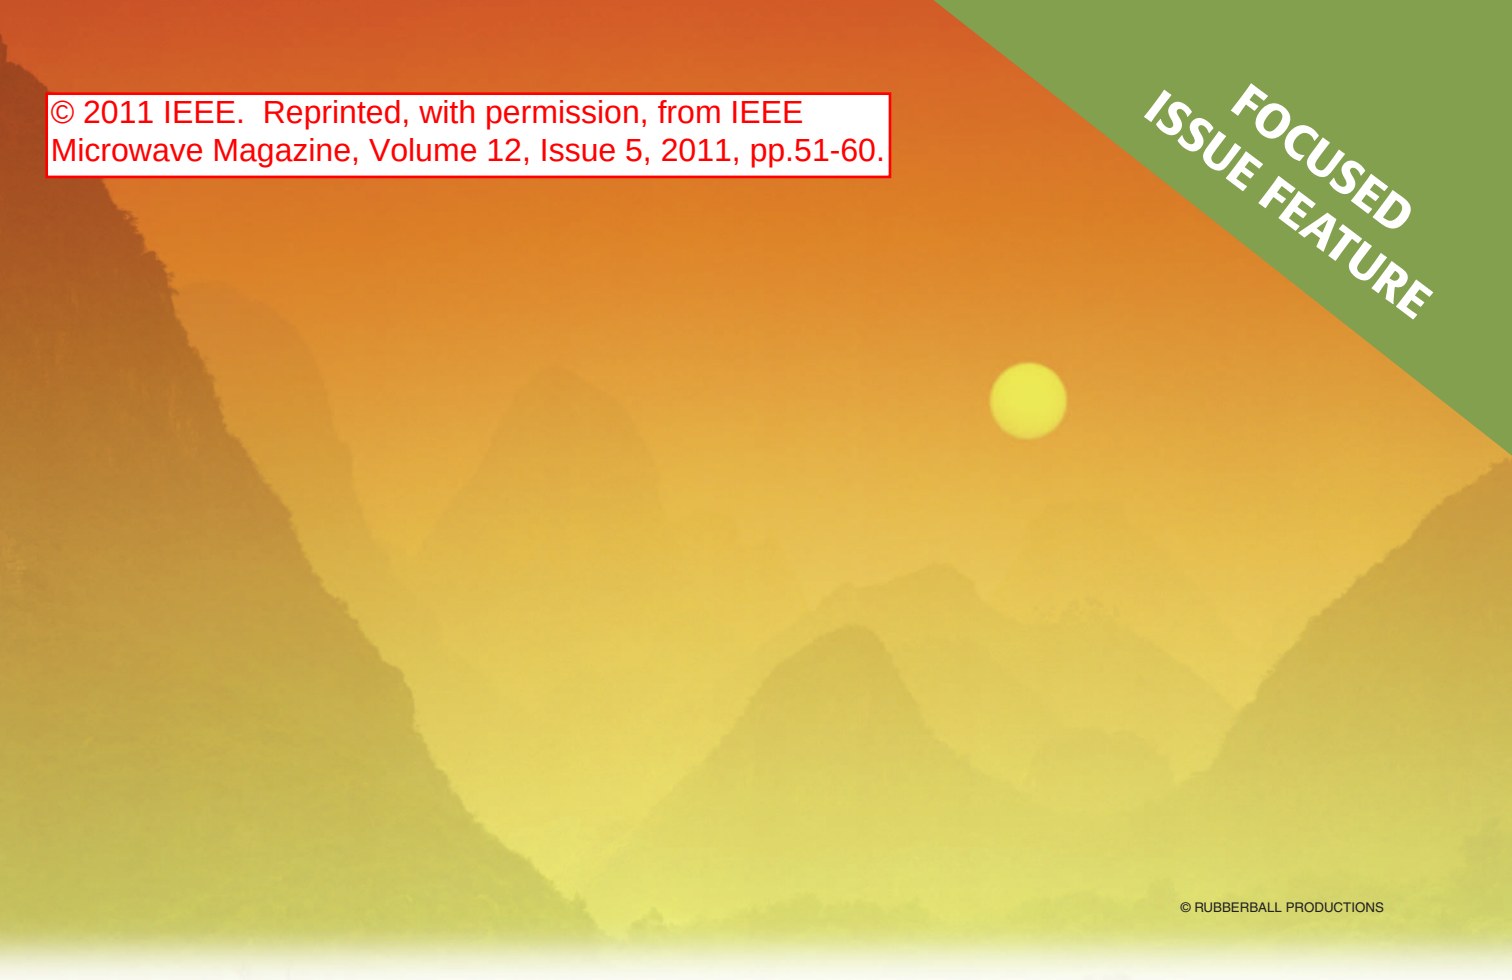

A New Dawn in R&D

Istvan Novak

The design of power distribution networks in modern high-speed electronic systems has become an ever-growing challenge. From boards with literally no bypass capacitors in the old days, through the one-capacitor-per-pin rule twenty years ago, the power distribution design has reached a state when it is equally as important as the signal-integrity design. The challenges stem from the simultaneous functions that the power distribution network must provide: 1) clean power, 2) return path for high-speed signals, and 3) a resonance-free power-distribution network. Sometimes these functions are interrelated and contradictory. There are additional challenges such as increasing power density pushing the required impedance levels into the mΩ range and the number of independent supply rails increasing. The average designer, using off-the-shelf components, still may find it hard to obtain the necessary silicon data for a proper power distribution design.

Istvan Novak (istvan.novak@ieee.org) is with Oracle.

Digital Object Identifier 10.1109/MMM.2011.941413
Date of publication: 7 July 2011

Today, well-established solutions exist to synthesize the self-impedance of the power distribution network.

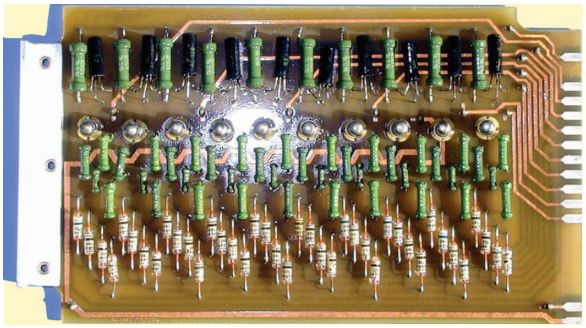

Figure 1. *Old computer card based on diode transistor logic with no bypass capacitors.*

This article covers the major steps of the power distribution network design flow: dc-drop analysis and ac self-impedance and transfer function requirements. Today, well-established solutions exist to synthesize the self-impedance of the power distribution network; the most common approaches are described. The front end of the design process uses simulations to help the initial "what-if" process. Simulation options are shown for dc-drop analysis and ac simulations. For transfer-function simulations, optimization routines help the designer to meet target parameters. In validation, very low impedance values can be measured accurately by four-point ac Kelvin connections.

Smaller component sizes come with new challenges, some capacitors today exhibit strong dc and ac bias dependence. Moreover, switching power-converter circuits may create large high-frequency burst noise, potentially interfering with sensitive signaling.

The article finishes with a brief outlook to the anticipated new technology trends, new challenges, and their potential solutions.

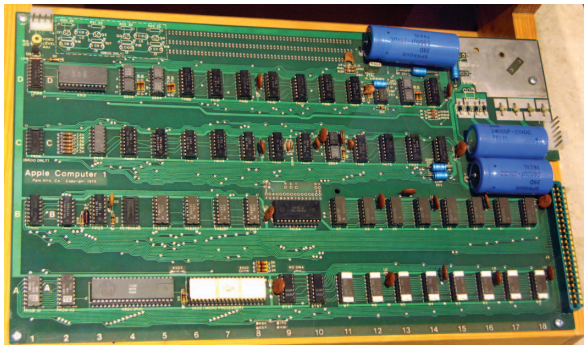

Figure 2. *One of the first Apple computer boards with integrated circuits, bulk, and ceramic bypass capacitors.*

Background

Historical Overview

During the second half of the 20th century, the power distribution network grew from an almost nonexistent after-thought to an important subsystem. Figures 1–3 illustrate this on computers. Figure 1 shows an old computer card made in the 1960s. The discrete components are on a 11 cm × 18 cm (4.3 in × 7 in) two-sided board. There are 40 germanium diodes, 20 transistors, and 50 resistors on the board. The power connections are at the lower right edge of the board, and there are no bypass capacitors at all. The operating frequency was a few hundred kilohertz. Power distribution, signal integrity, and electromagnetic (EM) radiation were not a concern.

In the 1970s, personal computers emerged, using microprocessors and integrated circuits. I took the photo (with permission) in Figure 2 at the Computer Museum in Mountain View, California, showing one of the first Apple computer boards from 1976. The white integrated circuit in the middle bottom is a Mostek 6502 microprocessor, and we see around it only one ceramic capacitor. Wide traces zigzag through the integrated circuits, carrying power and ground around. Several ceramic capacitors and a few large bulk capacitors complete the power network.

Jump forward two decades. Figure 3 shows a CPU module from a SUN V890 server. The 20 cm × 50 cm (8 in × 20 in) module has two UltraSparcIV CPUs and 16 memory slots. In the 22-layer board, there are four thin laminates to support eight supply rails with seven dc-dc converters and 1907 bypass capacitors. The parallel busses ran with subnanosecond edges. The CPUs operate at a clock frequency above 1 GHz. The power distribution was designed to ensure not only clean power, but to also serve as signal reference and to ensure that EM compatibility radiation limits are not violated.

Difficulties of Power Distribution Design

Lower-speed digital electronics were schematic based, and the physical implementation, PCB stackup, component placement, and layout, only mattered a little. At higher speeds, first signal-integrity issues emerge. We make sure that digital signals arrive on time and with low distortion. Signal integrity work checks time of flight, skew, timing, reflections, matching, crosstalk, attenuation, and dispersion of signals. Increasing the speed and system density even further, power-integrity issues also become challenging.

Exploding Number of Supply Rails

Systems used to have one, maybe two or three supply rails. Large computer boards today have up to several dozens of independent supply rails. This is a big challenge, but the number of supply rails is still much lower than the number of high-speed interconnects.

2-D and 3-D Problem Instead of 1-D

Signal integrity design is mostly a one-dimensional (1-D) problem. The main signal goes along traces that we design. The traces determine the path of the signal. A major contributor to the power integrity challenge is the fact that noise on planes can propagate anywhere in the X-Y directions and vias can carry the power distribution network noise in the Z direction as well. We don't design the power distribution network to intentionally propagate noise, just to the contrary: a well-behaved power distribution network should block the propagation of noise. To figure out how noise spreads, we need broadband two-dimensional (2-D) or three-dimensional (3-D) power distribution network models, which are harder to create and manage than a 1-D or localized 3-D signal integrity model.

Unknown Excitation and Requirements

In signal integrity, we know what the signal levels and waveforms should be. In power integrity, the main signal is the power rail noise, but its exact nature is mostly unknown. Noise excitation could come from several sources: power converter ripple and burst noise, shoot-through current of active devices, and return currents of signals. With the exception of the power-converter output ripple, which is typically fairly steady, all other noise components heavily depend on the system activity, creating a potentially huge parameter space. Also, supply-rail noise tolerances of the various active devices are only crudely specified; we usually have absolute limits for the supply voltage, but very little knowledge about the sensitivity of the device in different frequency ranges.

Lack of Standardization

Signal integrity gets support from component and CAD vendors. Manufacturers of active and passive devices are willing to provide simulation models and characterization services or create reference designs. There are standard signal integrity requirements to adhere to: eye masks, jitter specifications, attenuation, reflection, and crosstalk limits. Vendors can come up with solutions to simulate and measure our designs against specifications. Power integrity is left without this support; there are hardly any standards for power distribution network noise and testing. This leaves vendors without guidance for their offerings. The lack of standards also leaves power integrity designers on their own to come up with the design requirements. Some application notes and reference designs are available with power integrity content, but their applicability is usually limited to a narrow choice of parameters.

Power Distribution Network Design Flow

The design flow should start with the understanding of the functional, mechanical, thermal, and cost targets. We collect data about the various power consumers in the system, followed by a preliminary design of the printed circuit boards in terms of overall size, total layer count, and power-layer assignment. We also need to understand the static and dynamic requirements of the power consumers, and we need to catalogue our noise sources.

In the 1970s, personal computers emerged, using microprocessors and integrated circuits.

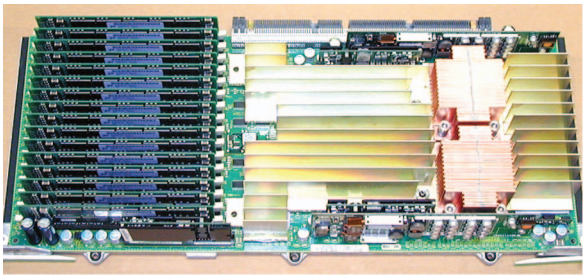

Figure 3. *CPU module, approximately 25 cm × 50 cm (10 in × 20 in) in size, designed in the early 2000s.*

The preliminary selection of components could be made on simplified assumptions, but later we may find that changes become necessary. Since the power distribution network has to serve the electronics in tight integration with its electrical and other functions, we simply cannot design an optimum power distribution network in isolation without taking into account the numerous other requirements of the system. A good power distribution design is iterative and cooperative; we do power-domain partitioning, component selection, and component placement in close cooperation with the signal-integrity, thermal, and mechanical designs.

The Functions of Power Distribution Network

We first have to understand the functions of power distribution networks [1].

- Provide clean power to the active devices.
- Optionally, serve as return path for signals.
- Ensure that radiation related to power distribution network does not violate legal limits.

In a complex system, the above three major functions can also be interrelated, or one or two of the major functions may not be present or may be irrelevant for a particular rail. The sketch in Figure 4 illustrates a board with one supply rail. In reality, we usually have multiple supply rails, each having different requirements, but being interrelated by shared plane layers with cross coupling.

There are two major classes of power distribution networks [1]: core and input/output power distribution networks. A third major class is also emerging: filtered supply rails.

In its simplest form, a core power distribution network feeds only one active device. If the power plane is

Large computer boards today have up to several dozens of independent supply rails.

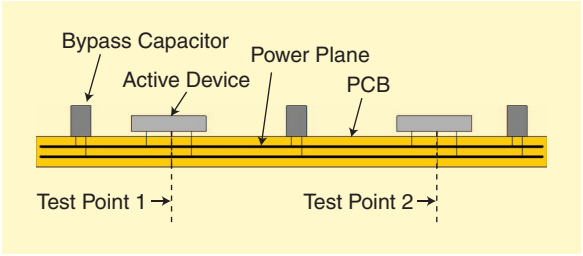

Figure 4. *Sketch of board-level power distribution network and its major components [1].*

sandwiched between neighboring ground planes, the core plane does not serve as signal return, and being between ground planes, EM radiation may be of no concern. Cost-effective designs, however, many times result in shared solutions: planes with the primary function of feeding cores or distributed input/output networks, may also serve as reference plane for 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 than what is allowed for high-speed reference planes. Similarly, since a core power distribution through a package behaves like a low-pass filter, high-frequency board resonances will have little influence on the power delivery, but it may be detrimental for EM interference (EMI) radiation and/or signal-return functions.

Most requirements can be captured by the maximum transient noise due to worst-case system activity. With the exception of low-end systems, where occasional failure may be acceptable, we usually keep the worst-case peak-to-peak transient noise within predefined limits. Since most power distribution network components can be fairly accurately described by linear and time invariant models, the worst-case peak-to-peak transient noise can be obtained from the step response of the power distribution network [9].

Due to rising CPU clock frequencies and signaling rates, more features on our boards and packages become electrically long. Quarter-wavelength or longer structures may become effective radiators, therefore, the power distribution network has become one of the primary EMI risk factors. It is generally understood that the best way to avoid radiation problems is to ensure a resonance-free impedance profile. It also helps to minimize simultaneous switching noise and jitter of high-speed signals when the power distribution network serves as signal reference.

We want to ensure that the power distribution network will properly function over the entire foreseeable range of parameter, 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, lower sensitivity is preferred. In high-power networks, a uniform stress distribution among components is important as well.

Cost-effective solutions optimize the margins and remove the extra padding of performance in the power distribution. This may result in designs where changing one element has a ripple effect and everything has to be realigned and redesigned. Therefore, portability of the design is becoming increasingly valuable.

We may also have a set of important nonelectrical parameters and targets

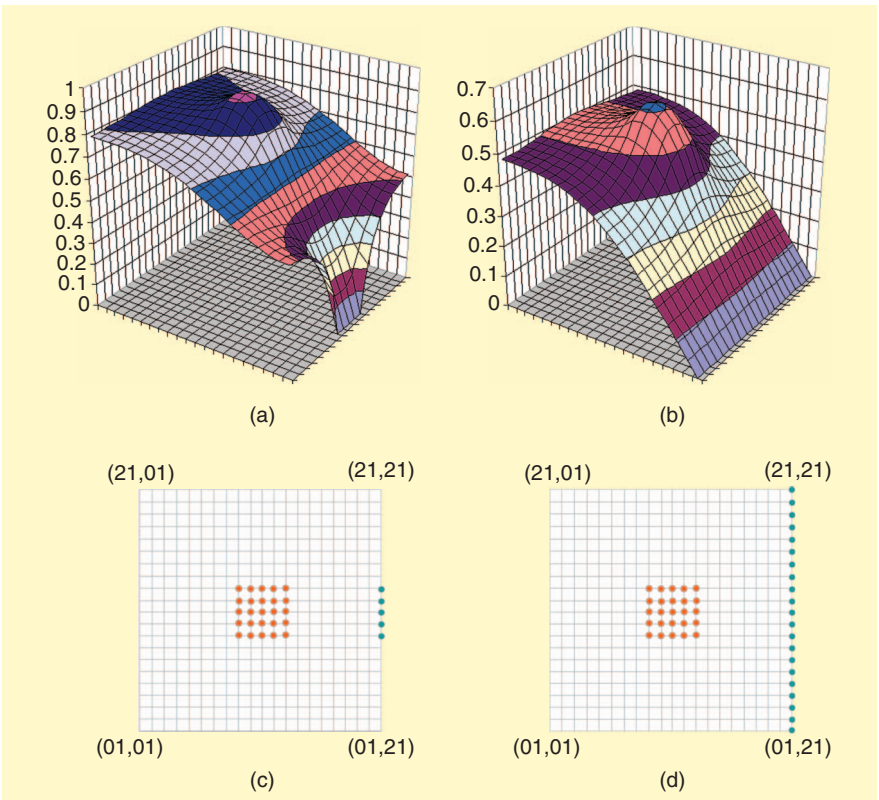

Figure 5. *Illustration of dc-drop simulations on a square plane with a large current sink in the middle. (a) and (b) The normalized voltage drop on the vertical axes. (c) and (d) The node connections.*

to meet, such as cost, size (area, height), weight, component connection style (through hole, surface mount), placement restrictions, simplicity of bill of material, etc.

DC Voltage Drop

Once considered necessary only for the highest-current applications, dc drop analysis is becoming a requirement on more boards. Increasing density forces us to use the narrower traces, planes, and cables, whereas the dropping supply voltages come with lower absolute tolerance limits. The well-known formula for dc resistance multiplies resistivity with the length and divides by the cross section of conductor. We can use this approximation for traces and cables because their typical aspect ratio guarantees that current redistribution takes up negligible length. For any application where the current path is not much longer than the cross-sectional dimensions, we have to rely on detailed meshing of the conductors.

Figure 5 shows a square plane supplying current to a large load in the middle. The plane is simulated with a uniform square resistive grid. The load is connected in the center of the plane through a 5×5 array of nodes (orange dots). The source is different: in Figure 5(c), the source has only five nodes in the middle of the right side (green dots), whereas in Figure 5(d) the source is connected through 21 nodes along the entire right side. The 3-D plots in Figure 5(a) and (b) show the voltage at each node. Feeding a large load through more nodes stretched out over a larger area can reduce the voltage drop significantly (in this normalized case from 0.9 units to 0.6 units).

For simulating more complex shapes, hand-carved resistive grids are not efficient. We use the meshing power of commercial tools, which output also the current density and power dissipation. Figure 6 is from a commercial CAD tool, showing the voltage gradient and dissipated power in a multilayer computer board [3]. Note that voltage drop is cumulative, whereas dissipation depends on the local current density and resistance, therefore, the locations with the highest voltage drop and highest dissipation do not coincide.

Impedance

Most of our electronics expect constant supply voltage. We make sure that the voltage fluctuation (noise) is within predefined limits. Though our target is in the time domain, the power distribution network design is customarily done in the frequency domain. The allowed noise is typically expressed in percentage around the nominal voltage; $+-10\%$ used to be sufficient. It later changed to $+-5\%$, then to $+-3\%$, and, in some applications, we see $+-1\%$. A good power distribution network is stiff. The power rail impedance has to be much less than the dynamic

It is generally understood that the best way to avoid radiation problems is to ensure a resonance-free impedance profile.

impedance of the loads. The assumption that load impedances are much higher than the power distribution network impedance leads us to an impedance matrix description of the power distribution network. The Z matrix (as opposed to, for instance, the admittance or scattering matrices), requires open termination (or nothing) on the ports, which makes the measurements and simulations of the Z matrix easier. At the location of any active device, the noise will be the sum of self-inflicted noise (here we need the self impedance) and noise injected by other devices, propagated through the power distribution network (here we need transfer impedance). Simplified methods deal with only the self impedance of a power distribution network, and, if there is any concern about noise sources affecting each other, separating the loads to individual supply rails may be the answer. As long as we assume that the power distribution network is a linear, time invariant circuit, we can apply the complex Ohm's Law

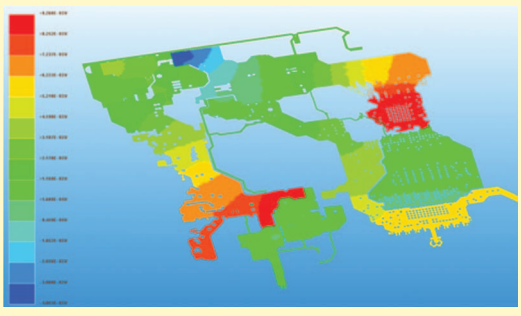

Figure 6. *(a) DC voltage drop and (b) power dissipation on a single power rail in a computer board. Colors toward red indicate larger voltage drop and higher dissipation. (Courtesy of ANSYS, Inc.)*

Portability of the design is becoming increasingly valuable.

$$v(t) = invFFT\{\mathbf{Z}(\mathbf{f})\ FFT(i(t))\}. \quad (1)$$

If the power distribution network impedance is flat and resistive in the frequency range of interest, (1) reduces to the simple form of $v(t) = Z_{PDN}\, i(t)$, where we call Z_{PDN} the target impedance [4].

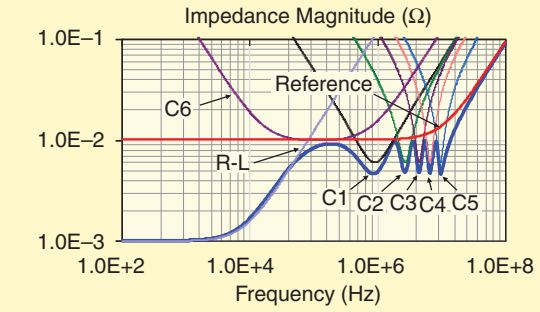

Figure 7. *Self-impedance profile of a multipole power distribution network design for 10 mΩ target impedance.*

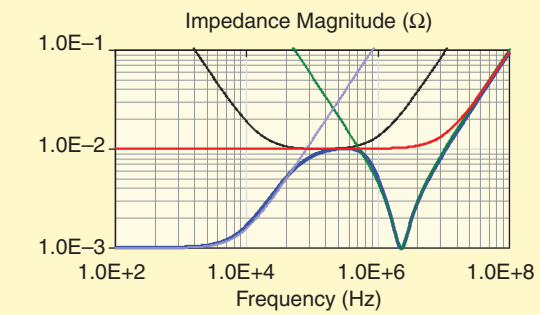

Figure 8. *Self-impedance profile of a Big-V power distribution network design for 10 mΩ target impedance.*

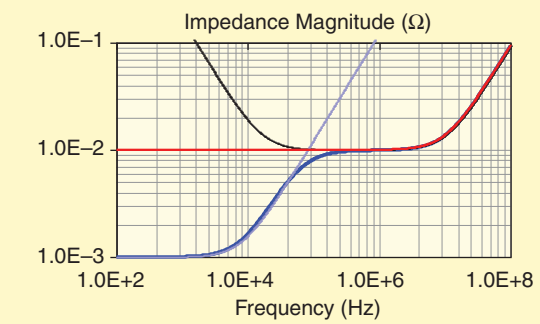

Figure 9. *Self-impedance profile with distributed matched bypassing for 10 mΩ target impedance.*

Synthesizing Lumped Power Distribution Network Impedance

For many years, it was sufficient to follow the advice "place a bulk capacitor where the power enters the board and place one 0.1 μF ceramic capacitor next to each active device's power pins." [4] was one of the pioneering works, which defined target impedance and its approximation with the multipole power distribution network design process, also called the "frequency-domain target-impedance method (FDTIM)" [2]. In this approach, the target impedance is approximated by lining up the series resonance frequencies of different-valued bypass capacitors along the frequency axis. The number of capacitors for each value is chosen such that their cumulative effective series resistance (ESR) values match the target impedance. The flat red line in Figure 7 shows the target impedance, the thin colored lines illustrate each capacitor bank, and the heavy blue line is the cumulative impedance.

Another popular approach is the "Big-V," named after the shape of its impedance profile [2]. The mid-frequency impedance target is matched with very-low-Q bulk capacitors, and the low impedance at higher frequencies is generated by a number of same-valued capacitors, driving a notch in the impedance profile. This is shown in Figure 8, where the flat red line shows the target impedance, the thin colored lines show the capacitor banks, and the heavy blue line is the cumulative impedance. No matter how unscientific the old advice may sound today, for all practical purposes, it created a Big-V impedance profile.

Wideband flat impedance can be synthesized with a few very-low-Q bypass capacitors. It was realized early on that smooth or flat impedance profiles minimize transient noise. The solution was first developed for the interface between the dc-dc converter and bulk capacitors. Adaptive voltage positioning [5] cuts the converter transient-response noise in half. Later, the approach was expanded to the capacitor-to-capacitor interfaces, and the extended adaptive voltage positioning technique was born [6]. Eventually, it was combined with matched power-plane impedances, resulting in the distributed matched bypassing approach [7]. This is illustrated in Figure 9.

Each synthesis approach has its merits and limitations. A fair and objective comparison proves to be very difficult, primarily because of the fuzziness of the metric [8]. Using the worst-case instantaneous time-domain noise as the basis for the comparison, the reverse pulse technique [9] can help. From the step response of the power distribution network (assuming it is linear and time invariant), we can construct the worst-case transient noise. This would be the result of a random series of current steps with bounded magnitudes and rise/fall times. If we apply the reverse pulse

technique to the power distribution network impedances in Figures 7–9, we get the step responses shown in Figure 10.

From the Reverse Pulse Technique calculations, the Distributed Matched Bypassing, the multipole and the Big-V designs yield 19, 25 and 27.5 mVpp/A worst-case transient noise, respectively. This gives a clear guidance: the 'smoother' the impedance profile, the lower is the worst-case time-domain noise.

Distributed Effects

When the power distribution network spreads out more than a fraction of the wavelength of the highest frequency of interest, distributed effects need to be considered. Similar to unterminated signal traces, power and ground planes do resonate. For simple shapes, the modal resonance frequencies and the resultant self and transfer impedance profiles can be analytically calculated [10]. For complex contours and perforated planes, numerical techniques are best suited. Similarly, local details, such as via transitions, plane perforations in via arrays, are best handled by numerical analysis. Figure 11 shows a silicon and package distributed model.

Avoid Resonances

The worst-case time-domain transient noise can be minimized by flattening out the power distribution network impedance profile. This is in line with the EM compatibility guidelines, telling us to minimize resonances. In a power distribution network, any two components can create an antiresonance peak; dc-dc converter loops can resonate with bulk capacitors, different valued capacitors can create resonant peaks, capacitors can resonate with planes, and the plane shapes and plated through holes may also resonate. At low frequencies, where lumped approximations hold, we can follow the process outlined in [6] and [7]. To suppress plane resonances, we can either terminate the plane's edges in its characteristic impedance [11] or use very thin dielectrics, which attenuate modal resonances [12]. Figure 12 shows the self impedance in the middle of a 25 × 25 cm power/ground plane pair, with the dielectric thickness as a parameter. As dielectric thickness drops below 10 μm, modal resonances almost completely disappear. The resonance suppression also depends on the horizontal size of plane; smaller-sized planes resonate stronger [13].

The Front-End and Back-End

The design and/or the validation can be done either in the frequency or time domain. We complete the front-end tasks component selections, what-if simulations, and detailed design simulations. After the circuit is built, we need to bring up the system and validate the correctness of the design.

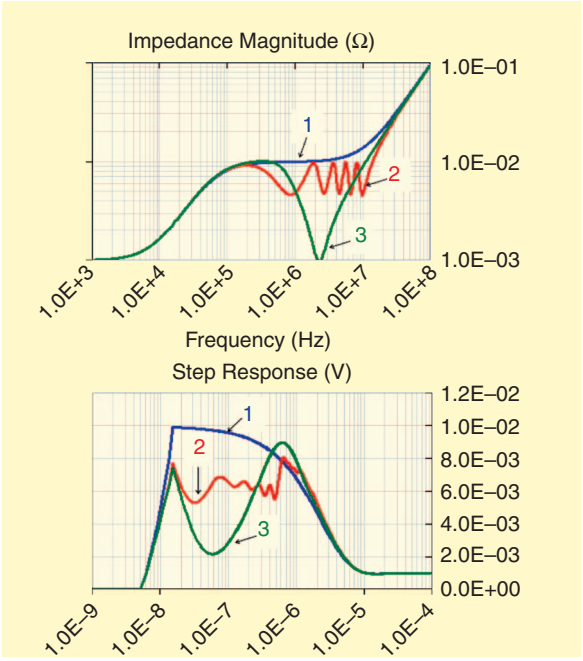

Figure 10. *Impedance profiles and step responses for the distributed matched bypassing (line 1), multipole (line 2), and Big-V (line 3) design methodologies, all targeting 10 mΩ impedance.*

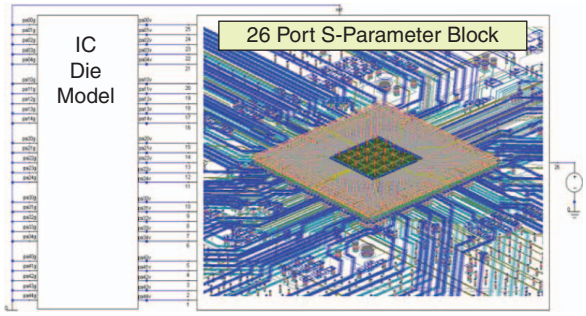

Figure 11. *Distributed silicon model connected to a multiport package power distribution network model [3]. (Courtesy of ANSYS, Inc.)*

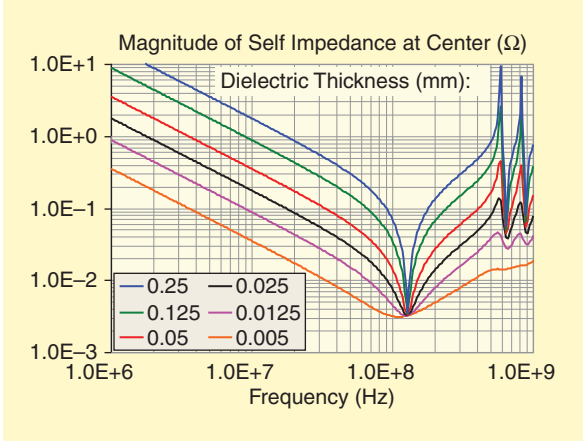

Figure 12. *Illustration of resonance suppression by thin laminates.*

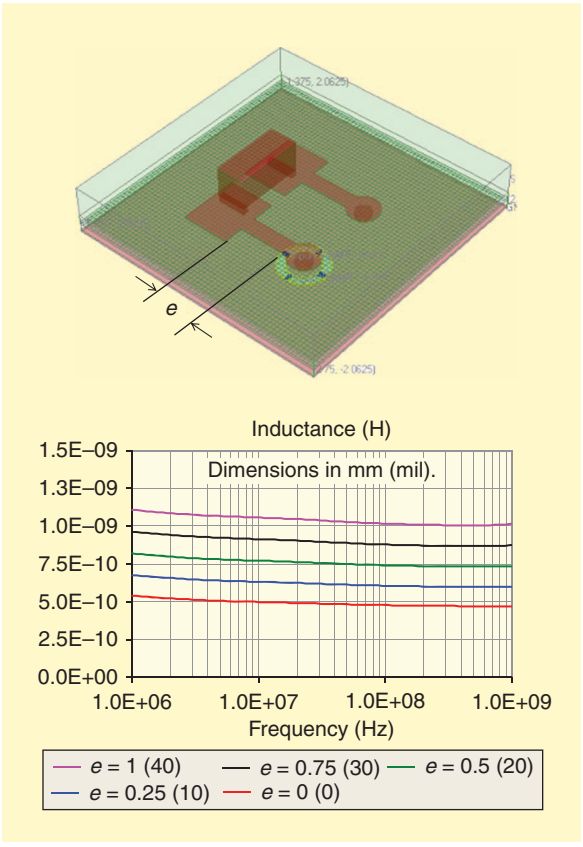

Figure 13. *Inductance of a 0402-size bypass capacitor (a standard case size 1 mm by 0.5 mm) connected to a pair of power-ground planes by two vias. The parameter is the length of the escape trace.*

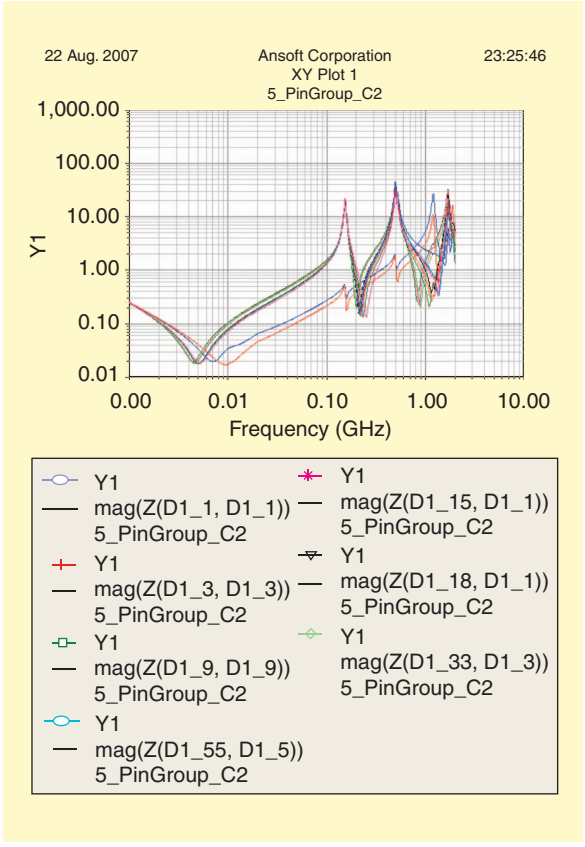

Figure 14. *Optimization of a power distribution network impedance profile. The traces show the impedance profile of a board with bypass capacitors at various locations. (Courtesy of ANSYS, Inc.)*

Frequency or Time Domain?

Because digital signaling is defined by its voltage or current levels in the time domain, it may seem obvious to characterize the behavior of the power distribution network also in the time domain. Unfortunately validation of a real system in the time domain by comparing measured and simulated behavior is not easy. Oscilloscopes triggered from the noise signal are prone to erroneously picking up noise from the environment, especially when we need to measure millivolt levels. To validate the design in the time domain, we have to know both the impedance matrix and the transient noise vector. Except for the active voltage sources, the power distribution network is built of passive components. Semiconductor devices show more change of their electrical parameters due to process, voltage, and temperature (PVT) than passive resistor, inductor, and capacitor (RLC) components, so the impedance of a power distribution network is usually more stable than the semiconductor network it feeds. With well-behaving components, it is certainly not significantly time varying. In contrast, the transient noise in a complex system comes from many active devices, and the sources have their own activity schedule.

If we wanted to find the transient load current by measurement, the difficulty is practical; it would require a shunt element in each current path or a current-measuring loop around current-carrying conductors [14]. Even if we could measure the current, it does not remove the highly statistical nature of the system activity that creates the excitation current.

Recent publications have described ways to indirectly measure impedance profiles and/or transient currents [15] with active devices exercised in a controlled manner. Dedicated transient exercisers are also available [16]. However, time-domain instruments have fewer effective bits, which limits the achievable dynamic range in comparison to narrowband frequency-domain instruments. Eventually, it is easier to segment the task and do the design and validation separately for the impedances. This is then followed by time-domain measurement and validation if it becomes necessary and if transient-current data is available.

The Front End: Selecting Stack-Up, Doing Prelayout Simulations

Selecting the proper stack-up for the packages and boards is crucial: the copper weight of plane layers

assures low dc drop, and the order of layers has to satisfy the return-path function. Via patterns of power and ground connections may resonate, which could create long-reaching crosstalk interactions between aggressor and victim signals [17]. Around layer-transitioning vias, the return path of high-speed signals goes through the power plane cavities and, therefore, the correct modeling of the cavity impedance, including the frequency-dependent dielectric and conductor parameters, is important [18].

A critical piece is finding the appropriate escape patterns for bypass capacitors and deciding on the locations of components. Figure 13 illustrates the impact of bypass-capacitor escape pattern on its effective inductance [19].

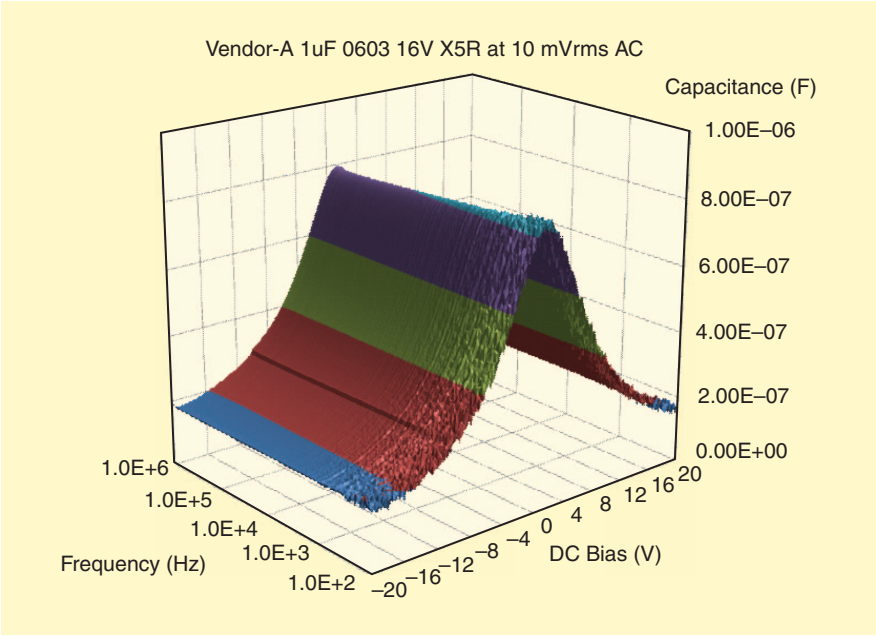

Figure 15. *Surface plot of capacitance versus frequency and bias voltage. (Reprinted from [22] with permission.)*

The Back End: Post-Layout Simulation, Hardware Validation

Once a package or board layout is available, the entire design can be simulated and, if necessary, optimized. The impedance-synthesis approaches described in the "Synthesizing Lumped Power Distribution Network Impedance" section can be described algorithmically, but they still rely on a series of manual decision points. As a result of the conflicting requirements, there is still no unified systematic and straightforward design process that could be implemented in software. Some CAD packages offer optimization of the impedance profile against user-defined targets. Figure 14 shows an example [3].

Once hardware is available, we validate its performance. The most straightforward validation is to correlate the measured impedance profiles against simulated estimates. At high frequencies, the impedance matrix is not preferred because it requires open termination, where fringing fields create errors. The solution is to use scattering parameters for measurements followed by a conversion to Z matrix.

Measuring low impedances is challenging. Connection discontinuity, instrumentation dynamic range, and error floor as well as potential cable-braid loop errors may degrade our data. When our goal is to measure mΩ impedance values, a proven approach is to use a two-port shunt-through connection [20]. For measuring dc-dc converters with a sub-mΩ impedance, the dynamic range of the instrument has to exceed 120 dB.

A built-in dc source with a scripting language to control the measurement setup allows for automated dc and ac bias sensitivity measurement of capacitors and inductors [21]. Figure 15 shows the measured capacitance of a 1 μF 0603 16 V multilayer ceramic capacitor as a function of dc bias voltage [22] (0603 is a standard case size of 1.5 × 0.57 mm.)

Finally, Figure 16 compares the measured impedance profile of a partially and a fully populated working mother board.

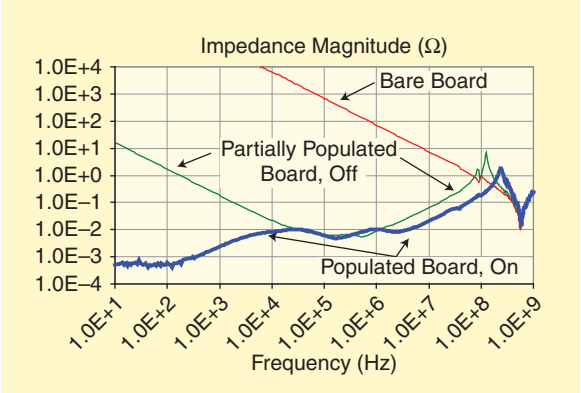

Figure 16. *Impedance magnitude of a bare board, partially populated board, and fully populated, powered board. The power rail supplied the core of a medium-power chip, where the goal was to maintain approximately 10 mΩ or less impedance up to 3 MHz.*

A critical piece is finding the appropriate escape patterns for bypass capacitors and deciding on the locations of components.

Looking into the Crystal Ball

In recent years, digital power has gained traction. The term "digital power" could be used loosely. It may simply refer to the availability of monitoring registers for digital polling, all the way to full-blown digital signal processing implementation of the regulator loop and control circuitry. The digital telemetry part is a very convenient function, but it does not necessarily affect the dynamic response of the converter. Digital loop control, on the other hand, makes it possible to change the feedback loop parameters on the fly and apply nonlinear loop control and timing adjustments, especially in multiphase converters, which can significantly improve the transient response. The digital power feature is currently available in high-end converters. As the technology matures, we can expect it to become available at lower-tier converters as well.

As system densities continue to rise, we can expect more parameter interdependencies. Some of the popular ceramic bypass capacitors already show a large bias dependence, as illustrated in Figure 15 and [21]. We can anticipate that, eventually, to help alternate-source selections, various classes of bias dependence may emerge, making it necessary to add this parameter to the component specification sheet.

The ongoing miniaturization helps to balance interconnect resistance and inductance, making it possible to reduce pulse edge rise time and increase operation speeds. Direct-attached stacked devices, 3-D packaging, and embedded components, both passive and active, as well as band-gap filter structures embedded in multilayer printed-circuit boards, will eventually become more readily available, further reducing parasitics [23] and enhancing power distribution network performance [24].

References

[1] R. Kaw, M. Swaminathan, and I. Novak, "Towards developing a standard for data input/output format for PDN modeling and simulation tool," in *Proc. of the 2005 IEEE Symp. on Electromagnetic Compatibility*, Chicago, IL, Aug. 8–12, 2005, pp. 644–649.

[2] S. Weir, "Bypass filter design considerations for modern digital systems, a comparative evaluation of the big 'V', multi-pole, and many pole bypass strategies," in *Proc. DesignCon East*, Worcester, MA, 20–21 September 2005.

[3] Ansys, Inc. Application Bried: Automated Decoupling Capacitor Analysis Using SiWave PI Advisor. Ansys Inc., 2011. Available: http://www.ansoft.com/products/si/siwave

[4] L. D. Smith, R. E. Anderson, D. W. Forehand, T. J. Pelc, and T. Roy, "Power distribution system design methodology and capacitor selection for modern CMOS technology," *IEEE Trans. Adv. Packag.*, vol. 22, no. 3, p. 284–291, Aug. 1999.

[5] R. Redl, B. P. Erisman, J. M. Audy, and G. Reizik, "Voltage regulator compensation circuit and method," *U.S. Patent 6 229 292*, 2001.

[6] A. Waizman and C. Chung, "Extended adaptive voltage positioning (EAVP)," in *Proc. EPEP Conf.*, Scottsdale, AZ, Oct. 23–25, 2000, pp. 65–68.

[7] I. Novak, L. Noujeim, V. St. Cyr, N. Biunno, J. Howard, A. Patel, G. Korony, and A. Ritter, "Distributed matched bypassing for board-level power distribution networks," *IEEE Trans. Adv. Packag.*, vol. 25, no. 2, pp. 230–243, May 2002.

[8] I. Novak, "Comparison of power distribution design methods," in *Proc. DesignCon 2006, TecForum TF-MP3*, Santa Clara, CA, Feb. 6–9, 2006.

[9] V. Drabkin, C. Houghton, I. Kantorovich, and M. Tsuk, "Aperiodic resonant excitation of microprocessor power distribution systems and the reverse pulse technique," in *Proc. EPEP Conf.*, Monterey, CA, 21–23 October 2002, p. 175–178.

[10] M. Carver, "Microstrip antenna technology," *IEEE Trans. Antennas Propagat.*, vol. AP-29, no. 1, pp. 2–24, 1981.

[11] I. Novak, "Reducing simultaneous switching noise and EMI on ground/power planes by dissipative edge termination," *IEEE Trans. Comp. Packag. Manufact. Technol.*, vol. 22, no. 3, pp. 274–283, Aug. 1999.

[12] I. Novak, L. Smith, T. Roy, and R. Anderson, "Lossy power distribution networks with thin dielectric layers and/or thin conductive layers," *IEEE Trans. Comp. Packag. Manufact. Technol.*, vol. 23, no. 3, pp. 353–360, Aug. 2000.

[13] J. R. Miller, G. Blando, K. B. Williams, and I. Novak, "Impact of PCB laminate parameters on suppressing modal resonances," in *Proc. DesignCon 2008*, Santa Clara, CA, Feb. 4–7, 2008.

[14] J. Kim, E. Song, J. Fan, J. Kim, and J. Drewniak, "IC noise source for dynamic PDN assessment," in *Proc. DesignCon 2011*, Santa Clara, CA, Jan. 31–Feb. 3, 2011.

[15] G. Taylor, C. Deutschle, T. Arabi, and B. Owens, "An approach to measuring power supply impedance of microprocessors," in *Proc. 10th Topical Meeting on Electrical Performance of Electronic Packaging*, Cambridge, MA, Oct. 29–31, 2001, pp. 211–214.

[16] S. Chickamenahalli, K. Aygun, M. J. Hill, K. Radhakrishnan, K. Eilert, and E. Stanford, "Microprocessor platform impedance characterization with VTT tool," in *Proc. Applied Power Electronics Conf.*, Austin, TX, Mar. 6–10, 2005, pp. 1,466–1,469.

[17] J. R. Miller, R. Dame, G. Blando, I. Novak, S. McMorrow, A. Rebelo, A. Lacap, and X. Zheng, "Examining the impact of power structures on EM model accuracy," in *Proc. DesignCon 2011*, Santa Clara, CA, Jan. 31–Feb. 3, 2011.

[18] J. Bell, S. McMorrow, M. Martin, and A. Neves, "Unified methodology of 3D-EM/channel simulation/jitter separation," in *Proc. DesignCon 2011*, Santa Clara, CA, Jan. 31–Feb. 3, 2011.

[19] Y. O. Shlepnev, "Building advanced via-hole models for analysis of PCB interconnects," PCB Design Conference East, Durham, NC, 24 October 2007. Available: www.simberian.com

[20] I. Novak and J. R. Miller, *Frequency-Domain Characterization of Power Distribution Networks*. Norwood, MA: Artech House, 2007.

[21] I. Novak, Y. Mori, and M. Resso, "Accuracy improvements of PDN impedance measurements in the low to middle frequency range," in *Proc. DesignCon 2010*, Santa Clara, CA, Feb. 1–4, 2010.

[22] I. Novak, K. B. Williams, J. R. Miller, G. Blando, and N. Shannon, "DC and AC bias dependence of capacitors," in *Proc. DesignCon 2011*, Santa Clara, CA, Jan. 31–Feb. 3, 2011.

[23] P. Viklund, "Embedded active parts: How, when and who?" in *Proc. DesignCon 2011*, Santa Clara, CA, Jan. 31–Feb. 3, 2011.

[24] A. Ciccomancini, T. Wu, and A. Orlandi, "Power noise suppression in mixed signal circuits using a GSPL structure," in *Proc. DesignCon 2011*, Santa Clara, CA, Jan. 31–Feb. 3, 2011.