Power Supply Noise in SoCs: Metrics, Management, and Measurement

Karim Arabi

PMC-Sierra

Resve Saleh and Xiongfei Meng

University of British Columbia

Editor's note:

Power integrity is emerging as a major challenge in deep-submicron SoC designs. The lack of predictability is complicating timing closure, physical design, production test, and speed grading of SoCs. This article describes and validates two metrics that quantify the impact of power supply noise.

—*Mohammad Tehranipoor, University of Connecticut*

■ **THE IC INDUSTRY** is moving quickly to adopt new deep-submicron (DSM) technologies that offer unprecedented integration levels and cost benefits. Unfortunately, these advanced technologies pose new and sometimes unexpected challenges to the semiconductor industry. Current DSM problems of excessive power dissipation, voltage-drop effects, coupling noise, and decreasing production yield will continue and will likely get worse. This has led to the development of SoC design methodologies to deal with the problems of complexity and productivity.¹ As a result, the SoC design bottleneck in DSM technologies is shifting away from reusable IP design and integration to physical design and verification of the overall system.

To reduce power dissipation, manufacturers have scaled down supply voltage in each successive technology. However, increasing operating frequencies and power densities in 90-nm and 65-nm ICs lead to a higher percentage of voltage drops in the power grid. Simultaneous switching of I/O blocks and core logic causes large, rapid changes in supply current. The level and change rate of current are the two main sources of power supply noise. At the same time, circuit tolerance to voltage drop in DSM technologies is decreasing as a result of lower logic noise margins. This is because threshold voltages do not scale down as fast as supply voltages as the technology geometry

scales down. A constantly varying V_{DD} or V_{SS} voltage profile introduces unpredictability in the timing behavior of logic gates, making design sign-off difficult. Furthermore, large, instantaneous voltage drops on power supply rails can render standard cells and memory circuits inoperable and lead to unex-

pected behavior.

Power consumption and power supply noise will likely be the two main challenges of semiconductor design over the next few years. Here, we are concerned with maintaining the power integrity of DSM designs. Voltage drop in the power grid is due to two components: IR and Ldi/dt . R represents the resistances of the power mesh network, power pads, and device package. L represents the inductances of the power mesh network, power pads, and device package. Our goal is to provide useful information to properly manage supply noise in large SoCs.

In the past, analysis techniques and design metrics dealing with power supply voltage drop were simplistic. Designers analyzed power supply noise with static voltage drop (SVD) analysis, which might not reflect the true nature of power supply fluctuations, leading to either unnecessary overdesign or risk of timing failures. Dynamic voltage drop (DVD) analysis is emerging as a replacement of SVD analysis for capturing the impact of power supply noise on the timing behavior of logic and memory cells. However, neither approach has been fully justified or validated.

To better understand the impact of power supply noise on performance, it is critical to monitor and measure power supply noise on silicon. The next step is to reduce overall power supply noise, which can be accomplished in many ways. Among physical design

and circuit techniques for reducing power supply voltage drop are improving power supply mesh design, increasing the number of power supply pins, inserting decoupling capacitors (decaps), and integrating voltage regulators. Introducing useful skew to stagger high-activity clocks and buffers across the chip and restructuring the design to reduce rapid changes in the switching profile are also methods designers use to deal with power supply noise. In this article, we present our research results on power supply noise metrics, management, and measurement. We also present experimental results we gathered from chips in production.

Power supply noise analysis metrics

SVD analysis and verification have been an essential part of the physical design and verification flow in the semiconductor industry for more than 10 years. In this approach, the circuit's transistors or blocks draw their average current from the power grid, and the IR drops across the chip are computed using these average currents. The computed fixed values are fed to timing verifiers, which assess the impact of the IR drops on delay. SVD analysis has provided useful feedback in terms of certain glaring errors in the power grid design. However, at 90-nm and smaller technologies, SVD verification is not enough to ensure power integrity. It doesn't take into account the contribution of power density, variations in switching-activity profile, and impact of inductance and decaps, including LC resonance effects.² Therefore, it is inadequate for analyzing and optimizing power delivery networks in SoC designs.

Recently, industry began to use DVD analysis to capture the impact of decaps, inductance, and spatial and temporal switching events in the design. DVD analysis does not return a fixed value and therefore cannot be easily assessed or back-annotated into timing-analysis engines. Currently, there is no sign-off or analysis metric to characterize a DVD profile. This lack of suitable metrics for DVD makes it difficult to compare different DVD profiles and take appropriate design action, in terms of engineering change orders, to improve the DVD profile.

We propose and validate the use of two metrics to qualify a DVD profile and its impact on a design's timing performance. The DVD_{avg} metric is the DVD profile's average value in the timing cycle, and DVD_{max} is the DVD profile's peak value in the timing cycle. These simplified metrics are design optimization indications that accelerate the design process. Users

should add design margins to these metrics to account for their simplifications, or they should perform the final sign-off process with actual DVD profiles. Although these two metrics seem obvious, there has been little analysis to support their use.

Here, we show that DVD_{avg} correlates with the design's timing slowdown and should be set to the design's timing margin (as a percentage of the clock cycle). The DVD_{max} value represents the minimum voltage drop that causes a functional failure in the behavior of memory cells or standard cells. That is, if the voltage drop exceeds DVD_{max} , the behavior of standard cells or memory cells will be unpredictable. Therefore, the value of DVD_{max} depends on the tolerance of individual standard cells or memory cells to voltage drop. In our design environment, we have determined that DVD_{avg} should not exceed 5% of $V_{DD} - V_{SS}$, and DVD_{max} should not exceed 20% of $V_{DD} - V_{SS}$. These limits are pessimistic enough to account for the simplified nature of these metrics. The voltage drop profiles of different instances of critical paths are almost similar to one another, given the behavior of timing-driven placement engines that tend to cluster instances of the same path in close proximity to facilitate timing closure. Therefore, we can safely assume that all instances of a critical path see the same DVD profile. In corner cases, where this assumption is not true, designers must analyze DVD_{avg} and DVD_{max} for all instances of critical paths in the timing window of each instance. The timing window for each instance is where the instance is actually switching.

To support these two metrics, we first establish that the effect of the voltage drop profile on a digital path's timing performance is equivalent to applying a fixed supply voltage of $V_{DD} - DVD_{avg}$ to the same path. In other words, we first simulate a logic path in the presence of the true dynamic voltage profile on the power supply. Then, we use a DC voltage equal to the average of the voltage profile on the power supply and show that the timing behavior of the two cases will match.

Figure 1 illustrates the intuitive reason for this relationship. Consider the delay of the critical path of the circuit between the two flip-flops. Clearly, gate delay decreases when supply voltage overshoots, and increases when supply voltage undershoots. When supply voltage fluctuates, gates that see a voltage drop higher than $V_{DD} - DVD_{avg}$ will accelerate, and gates that see a voltage drop lower than $V_{DD} - DVD_{avg}$ will decelerate, compared with the situation in which all gates see a voltage drop equal to $V_{DD} - DVD_{avg}$.

Figure 1. Dynamic voltage drop profile across a cycle.

Because paths consist of several gates, we assume that the law of averages applies: Delay variations in the path are equalized to the average of the delay variations of each gate in the path. For this assumption to hold, the relationship between gate delays and voltage around the nominal point should be linear, and the path considered should span the entire clock cycle, averaging the entire voltage profile, not just a portion. We find that the assumptions of linearity and path-averaging effects are valid because power supply noise is analyzed up to $DVD_{max} = 20\%$ (keeping gate delays almost linearly proportional to V_{DD}), and we are interested in critical paths by default.

We validated these assumptions on several industrial designs by simulating the extracted layout of their paths in Spice with detailed voltage drop profiles and with $V_{DD} - DVD_{avg}$. We used a commercial DVD analysis tool to obtain detailed supply voltage profiles of the paths.

We selected three critical paths in different physical locations of a communication chip that uses a 90-nm CMOS process with a 1.0-V supply and simulated various data activities to create DVD profiles with different shapes. We simulated the extracted critical paths with the true profile and its equivalent $V_{DD} - DVD_{avg}$. Figure 2 shows the results for one path. For each profile tagged $V_{DD} - DVD_{avg}$, we simulated the path with the actual DVD profile and its equivalent DC voltage of $V_{DD} - DVD_{avg}$. The figure uses points to represent the simulation results for the true profile simulation, and a solid line to represent simulation results with the DC voltage of $V_{DD} - DVD_{avg}$. The difference between timing performances is less than 1% for all simulated cases. In all our critical paths, the DVD profiles for different instances of the critical path

Figure 2. Dynamic voltage profile simulation (solid line) versus $V_{DD} - DVD_{avg}$ (points) for a critical path in a communication chip design.

have similar characteristics. We also simulated each instance with its own DVD profile and observed no noticeable change in results.

The results reveal a strong linear relationship over the range shown and indicate roughly that a 5% to 10% change in V_{DD} leads to a 5% to 10% change in delay. By analyzing different DVD profiles using these metrics, we reached an important conclusion: The Ldi/dt contribution of voltage drop is not as critical as the IR contribution. Ldi/dt affects only DVD_{max} and thus has minimal impact on DVD_{avg} as long as the charge transfer is completed within the cycle. Therefore, Ldi/dt voltage drop doesn't play a major role in the chip's timing performance as long as it doesn't create supply fluctuations that exceed DVD_{max} . The two metrics let us compare and qualify DVD profiles and optimize designs more effectively to improve timing margins. The behavior described by DVD_{avg} is fundamental and originates from the nature of cells responding to voltage variations. This behavior of cells responding to the average power supply voltage should hold at least for 65-nm technology. Armed with this knowledge, the designer's next step is to reduce supply noise.

Reducing power supply noise

Here, we focus on decaps and the new design considerations necessary as technology scales below 90 nm. Among power supply noise reduction techniques, inserting decaps is the most common. Decaps hold a reservoir of charge and are placed around regions of high current demand. When large drivers switch, nearby decaps provide a source of current that reduces IR and Ldi/dt voltage drops to keep the target average and peak supply voltages within their noise budgets (DVD_{avg} and DVD_{max}). White-space decaps,

Figure 3. Decoupling capacitor (decap) modeled as a lumped RC circuit.

which usually consist of NMOS transistors, are placed between blocks in the open areas of the chip. In contrast, standard-cell decaps use both NMOS and PMOS devices and are placed within the logic blocks themselves.

Figure 3 shows a decap formed by an NMOS transistor. This decap is modeled as a lumped RC circuit, in which C determines the charge available and R determines the transient response of charge delivery to the switching circuit.³ Effective capacitance C_{eff} and effective channel resistance R_{eff} at low frequencies are

$$C_{\text{eff}} = C_{\text{ox}}WL + 2C_{\text{ol}}W \quad (1)$$

$$R_{\text{eff}} = L/[12\mu C_{\text{ox}}W(V_{\text{DD}} - V_{\text{TH}})] \quad (2)$$

Here, C_{ox} is the oxide capacitance per unit area, C_{ol} is the sum of overlap and fringing capacitances per unit width, μ is the channel mobility, V_{TH} is the threshold voltage, and W and L are the transistor's width and length.⁴

Both R_{eff} and C_{eff} are functions of the operating frequency, f . This was not an important issue in earlier technologies, but, as f increases, it is important to recognize that these values will not be constant. For example, consider a fixed-area decap with changing W and L . While the transistor area is fixed, the low-frequency capacitances and resistances will change from one layout to another according to Equations 1 and 2.

We used Spice to plot the frequency characteristics of NMOS and PMOS decaps. Typical high-speed clock rates today are in the GHz range, but it is important to study frequency response well beyond clock frequency. Most of the spectral power density of digital signals lies within frequencies of up to $f_{\text{knee}} = 1/(2t_{\text{rise}})$, where

t_{rise} is a signal's rise time (on the order of 50 ps or less), and f_{knee} is the 3-dB cutoff frequency of the spectral power density.⁵ We assume conservatively that $t_{\text{rise}} = 50$ ps, and we carry out the analysis up to 10 GHz. As Figure 4 shows, both R_{eff} and C_{eff} decay in value as frequency increases. The reason for this is that the charge in the MOSFET channel cannot change instantaneously, because mobile carriers from the source and the drain take a finite amount of time to form the inversion layer. The channel charge's response time is controlled by the device's transit time, which is quadratically related to channel length. At high frequencies, the channel charge cannot respond fast enough to match the gate charge, causing an effective degradation in capacitance value.

Both NMOS and PMOS devices with long channels have similar characteristics, but the effect is far more pronounced in PMOS devices. For example, when $L = 32\lambda$, the capacitance drops off noticeably because the device cannot provide the charge in the channel at the rate demanded by voltage changes at the gate. At 10 GHz, the NMOS device's capacitance drops by 20%, whereas the PMOS device's capacitance drops by 80%. This implies that decaps should be designed with smaller L values, requiring the use of multiple parallel devices, commonly known as fingered layouts. For example, one device with $L = 32\lambda$ can be laid out with four devices, each with $L = 8\lambda$. PMOS devices must use roughly twice as many fingers to achieve the same performance as NMOS. You can understand this intuitively by comparing the responses of $L = 32\lambda$ for NMOS with $L = 16\lambda$ for PMOS. Unfortunately, the use of fingers increases the number of contacts and increases the area. The area penalty can be severe, so a suitable L must be selected—one that trades off area and frequency response.

At the 90-nm technology node, two new problems arise because oxide thickness t_{ox} decreases to 2.0 nm or less. A large voltage drop across the thin oxide introduces gate-tunneling leakage⁶ and potential electrostatic discharge (ESD) breakdown.⁷ Higher gate-tunneling leakage through decaps increases the chip's total static-power consumption. An ESD event across the thin gate oxide increases the likelihood of permanent damage to the IC.

We can determine gate leakage levels with BSIM4 Spice models. Using a 90-nm technology with a 2.0-nm oxide thickness and a 1.0 V power supply, we can estimate the corresponding gate leakage current density, J , as 3×10^{-7} A/ μm^2 for NMOS and

Figure 4. Effective capacitance C_{eff} and effective resistance R_{eff} of decaps in 180 nm: NMOS C_{eff} (a), PMOS C_{eff} (b), NMOS R_{eff} (c), and PMOS R_{eff} (d).

$1 \times 10^{-7} \text{ A}/\mu\text{m}^2$ for PMOS. Thus, PMOS gate leakage is roughly three times smaller than NMOS gate leakage for the same-size transistors.

Gate leakage contributes to total static-power consumption, and decaps usually occupy a large on-chip area. White-space decaps (those located outside the standard-cell blocks) can be implemented using a thicker oxide, greatly reducing leakage and ESD, but typically consume three times more area. Standard cells require the use of thin-oxide decaps because the decaps are embedded in the logic blocks. A possible solution is to use only PMOS devices for standard-cell decaps, because they leak less. However, this is not a viable solution for high-performance circuits, because PMOS devices have a poorer frequency response than NMOS devices.

Standard-cell decaps at 90 nm also require modifications for ESD protection. ESD is a transient process of static-charge transfer that can arise from human contact with an IC pin. Typically, about $0.6 \mu\text{C}$ of charge is carried on a body capacitance of 100 pF, generating a potential of 2 kV (or higher) to discharge

from the contacted IC pin to ground for a duration of more than 100 ns.⁷ When running simulations of an ESD event, the designer can measure oxide voltage and compare it with a device's oxide breakdown voltage for a given technology. Oxide breakdown voltage is almost linearly proportional to oxide thickness.⁷ For instance, a 90-nm process using $t_{\text{ox}} = 2.0 \text{ nm}$ has an oxide breakdown voltage of around 5 V. If thickness doubles, oxide breakdown voltage also doubles to around 10 V.

In an ESD protection scheme for decaps, a series resistance R_{in} is required as a protection element. This resistance limits the maximum voltage possible at its gate, where ESD damage would occur. Of course, inserting a large R_{in} dramatically reduces its frequency response and transient response. Consequently, designers must consider the trade-off between ESD protection and decap response time.

To address the issue of ESD reliability, cell library developers have proposed a cross-coupled decap. This design provides better ESD protection by having a larger effective resistance without additional area.

Measuring power supply noise

It is important to correlate power supply noise estimation tools with silicon results to ensure dependable data for reducing DVD.² There are many techniques for monitoring power supply noise on chip.⁸ These techniques are not suitable for production environments, because they require either significant area or complex off-chip data processing. To overcome these limitations, we investigated a simple monitoring technique based on undersampling.⁹ In undersampling, a high-frequency periodic signal is captured from a large number of cycles, using a slower sampling signal to achieve a high-speed sampling rate. A slow, high-precision sampler can operate over many cycles to emulate a high-speed, high-precision sampler. This technique eliminates the need to trade accuracy for speed or vice versa.

In our case, the DVD was not periodic, but we could repeat the same experiment several times, each time skewing the sampling point by a small delta time, T_{SAMP} , which represents the sampling period, resulting in an equivalent sampling frequency of $F_{\text{SAMP}} = 1/T_{\text{SAMP}}$. Our implementation of this technique includes a sample-and-hold circuitry and an A/D converter. The ADC is optional if the analog output can be observed through an analog pin. The sample-and-hold circuitry consists of a regular operational amplifier, which features an offset voltage of less than 1 mV. In our implementation, we did not integrate the ADC on the same chip. We determined that about 30 samples are usually enough to represent the DVD profile, and therefore the measurement must be repeated at least 30 times. We repeated the measurement for each point five times to average and cancel noise effects, resulting in 150 measurements. For a typical clock frequency of 300 MHz, these measurements took about 0.5 μs .

Another simple way to monitor power supply noise is to create an on-chip broadband probe. Figure 5 illustrates this technique. The resistive impedance of the wire trace from the sampled point on the power mesh to the analog pin of the device used to monitor the supply noise must remain at 50 ohms. As Figure 5 shows, during bench measurement, the analog pin should terminate with an external 50-ohm resistor, to prevent ringing and reflections on the probe. The on-chip wire trace must be carefully shielded to prevent any potential for crosstalk with adjacent on-chip signals. The sampling point is just above the standard-cell power rail in the region where power supply noise is high.

Figure 5. Broadband probe for monitoring on-chip power supply noise.

Figure 6 shows the DVD simulation results from a commercial voltage drop analysis tool, and silicon results from a broadband probe. The wave shape is the actual voltage drop on the V_{SS} network at a hot spot during normal mode. The total measured $V_{\text{DD}} - V_{\text{SS}}$ voltage drop was 28 mV, whereas the simulated $V_{\text{DD}} - V_{\text{SS}}$ voltage drop was 25 mV. Achieving this level of correlation required much effort to properly model the equivalent resistances and inductances of the package, pads, and bond wires. Once the correlation was established, we could trust the simulation results in optimizing the power mesh network and our decap insertion design strategy.

Voltage drop effects on scan testing

An area overlooked in the past is the management of supply noise during the testing process. Previous work has attempted to reduce the power or thermal effects of scan-based testing.¹⁰ However, scan test mode is vulnerable to power supply noise because switching activity is typically three to four times higher than in normal mode as a result of the DFT strategy. This leads to excessive voltage drop during scan testing.¹¹ In some cases, voltage drop in scan mode has been so excessive that it has resulted in inadvertent logic value toggling and test result corruption.

The issue is more critical for delay testing because the chip is tested at speed, and speed degradation caused by voltage drop fluctuations creates measurement noise. Performance degradation in scan mode due to power supply noise can be as much as 15%, meaning test engineers should perform delay testing at slower frequencies to ensure that good devices don't fail the test. To address this issue, we adopted a hierarchical scan flow with special treatment of the scan-enable signal. We also investigated a new

Figure 6. Dynamic voltage drop results from an analysis tool (a) and from silicon measurement by a broadband probe (b).

multiple-launch scan test technique to reduce the DVD disparity between delay testing and functional mode.

Hierarchical scan architecture

Research has shown that hierarchical scan can reduce power dissipation in scan mode.¹² To limit the amount of switching activity across the chip during scan vector testing, we chose a hierarchical strategy that divides the chip into DFT regions, as Figure 7 shows. Each DFT region is tested sequentially. The *block_sel* signal selects the DFT region under test. When a given DFT region is tested, the other DFT regions remain in normal mode to ensure that their scan chains don't shift and that their switching profiles stay closer to normal mode.

When a DFT region is selected, its scan-enable signal, *scan_en*, is allowed to toggle, while the *scan_en* signals of other DFT regions are forced to 0 to keep them in functional mode and prevent their scan chains from shifting. Several experiments showed that we need at least eight partitions to keep the switching profile closer to functional mode. We can also steer the ATPG tool to limit the amount of switching activity during shifting to keep it closer to functional mode. However, in many cases, the ATPG tool cannot reduce switching activity to desired levels.

We simulated voltage drop profiles across the chip for all instances of a 90-nm SoC. The design uses a 1-V supply and contains 489,817 instances. Figure 8 presents the distribution of voltage drop for a flat scan

design and the hierarchical scan design. We averaged the voltage profile for each instance in that instance's timing window. The graphs show that the reduced amount of switching activity in our hierarchical scan architecture noticeably improves the voltage drop map across the chip. The average voltage drop across all instances in the flat scan design is 72.6 mV, which decreases to 45.2 mV in the hierarchical design.

Clock-gating techniques to reduce power dissipation can be used in conjunction with the hierarchical scan method. Keeping the inputs of some scan chains at 0 during shifting can also reduce power dissipation during scan testing, but this severely affects test coverage by excluding those scan chains.

Figure 7. Hierarchical DFT architecture for reducing power supply noise during scan testing.

Figure 8. Voltage drop profile change from a flat scan design (a) to a hierarchical scan design (b).

Multiple-launch delay scan testing

The hierarchical approach is acceptable for stuck-at fault testing, but testing for delay faults requires a closer correlation with the normal-mode switching profile. To satisfy this requirement, a multiple-launch technique can keep the block under test in normal mode for more than the traditional launch and capture cycles. Employing multiple launch cycles allows the block's DVD to settle and approach the normal mode profile, thus improving delay-test accuracy. Several experiments showed us that, in most cases, we need at least three launch cycles to create switching activity closer to the normal mode.

We can use this approach for both transition fault testing and path delay fault testing. The initial launch cycle might still trigger higher than normal switching activity, but switching activity will stabilize to normal levels after a few consecutive launch cycles. We perform the timing measurement on the capture cycle, in which the voltage profile should correlate well with the normal mode. This technique's main drawback is that it reduces transition and delay fault coverage because the ATPG tool must deploy a sequential ATPG engine, which generally results in lower test coverage than with combinational ATPG engines.

We simulated voltage drop profiles using a single-launch, single-capture pattern; a multiple-launch, single-capture pattern, and a normal mode. The DVD_{max} and DVD_{avg} of the single-launch, single-capture pattern were respectively 56% and 19% higher than in normal mode at the capture cycle. The DVD_{max} and DVD_{avg} of the multiple-launch, single-capture pattern were respectively 4.5% and 1.6% higher than in normal mode at the capture

cycle, demonstrating a much closer correlation with functional mode.

Our experimental results show a noticeable reduction in power supply noise during scan testing and a closer correlation between scan-based path delay testing and functional speed grading. However, we must point out that it is nearly impossible to recreate the same DVD profile in scan mode, and therefore margin allocation is necessary. Even in functional mode, each data activity creates a distinct DVD profile.

WE WILL CONTINUE to work on improving our decap design techniques, power supply noise-monitoring schemes, and correlations between path delay and functional testing. Other areas we plan to focus on include estimating power supply noise early in the design flow and designing on-chip voltage regulation schemes to work in conjunction with system-level voltage regulation. ■

Acknowledgments

We thank the Natural Sciences and Engineering Research Council (NSERC) and PMC-Sierra for their support of this work. We also thank Joseph Fang and Asad Shayan of PMC-Sierra for their help in generating some of the results presented in this article.

References

1. R. Saleh et al., "System-on-Chip: Reuse and Integration," *Proc. IEEE*, vol. 94, no. 6, June 2006, pp. 1050-1069.
2. S. Pant and E. Chiprout, "Power Grid Physics and Implications for CAD," *Proc. 43rd Ann. ACM/IEEE Design Automation Conf. (DAC 06)*, ACM Press, 2006, pp. 199-204.

3. J.R. Hauser, "Bias Sweep Rate Effects on Quasi-Static Capacitance of MOS Capacitors," *IEEE Trans. Electron Devices*, vol. 44, no. 6, June 1997, pp. 1009-1012.
4. J. Chia, *Design, Layout and Placement of On-Chip Decoupling Capacitors in IP Blocks*, master's thesis, Dept. of Electrical and Computer Engineering, Univ. of British Columbia, 2004.
5. H. Johnson and M. Graham, *High-Speed Digital Design*, Prentice-Hall, 1993.
6. H.H. Chen et al., "On-Chip Decoupling Capacitor Optimization for Noise and Leakage Reduction," *Proc. 16th Symp. Integrated Circuits and Systems Design (SBCCI 03)*, IEEE CS Press, 2003, pp. 251-255.
7. A. Amerasekera and C. Duvvury, *ESD in Silicon Integrated Circuits*, 2nd ed., John Wiley & Sons, 2002.
8. E. Alon, V. Stojanović, and M. Horowitz, "Circuits and Techniques for High-Resolution Measurement of On-Chip Power Supply Noise," *Proc. Symp. VLSI Circuits*, IEEE Press, June 2004, pp. 102-105.
9. B. Kaminska and K. Arabi, "Mixed Signal DFT: A Concise Overview," *Proc. Int'l Conf. Computer-Aided Design (ICCAD 03)*, IEEE CS Press, 2003, pp. 672-680.
10. P. Girard, "Low Power Testing of VLSI Circuits: Problems and Solutions," *Proc. 1st Int'l Symp. Quality Electronic Design (ISQED 00)*, IEEE CS Press, 2000, pp. 173-179.
11. T. Yoshida and M. Watati, "A New Approach for Low-Power Scan Testing," *Proc. Int'l Test Conf. (ITC 03)*, IEEE Press, 2003, vol. 1, pp. 480-487.
12. J. Saxena et al., "Minimizing Power Consumption in Scan Testing: Pattern Generation and DFT Techniques," *Proc. Int'l Test Conf. (ITC 04)*, IEEE Press, 2004, pp. 355-364.

Karim Arabi is a principal engineer at PMC-Sierra, where he is responsible for R&D activities in advanced technology development, design methodology enhancement, and design services. His research interests include low-power design, power management, DFT, and mixed-signal design and test. Arabi has a BSc in electrical engineering from Tehran Polytechnic, and an MSC and a PhD in electrical engineering from Ecole Polytechnique de Montreal. He is a member of the IEEE.

Resve Saleh is a professor and the NSERC/PMC-Sierra Chair in the Department of Electrical and Computer Engineering at the University of British Columbia. His research interests include low-power SoC design, interconnect analysis, and power grid design and analysis. Saleh has a PhD in electrical engineering from the University of California, Berkeley. He is a Fellow of the IEEE.

Xiongfei Meng is pursuing a PhD in the Department of Electrical and Computer Engineering at the University of British Columbia. His research interests include analog and mixed-signal design, with an emphasis on power supply noise control and monitoring. Meng has a BASc and an MASc in electrical engineering from the University of British Columbia. He is a student member of the IEEE.

■ Direct questions and comments about this article to Xiongfei Meng, Dept. of Electrical and Computer Engineering, University of British Columbia, 332-2356 Main Mall, Vancouver, BC, Canada V6T 1Z4; xmeng@ece.ubc.ca.

For further information on this or any other computing topic, visit our Digital Library at <http://www.computer.org/publications/dlib>.

Get access

to individual
IEEE Computer Society
documents online.

More than 100,000 articles and conference
papers available!

\$9US per article for members

\$19US for nonmembers

[www.computer.org/
publications/dlib](http://www.computer.org/publications/dlib)