Impact of Low-Impedance Substrate on Power Supply Integrity

Rajendran Panda and Savithri Sundareswaran
Motorola, Austin

David Blaauw
University of Michigan, Ann Arbor

Editor's note:
Although it is tempting to think of the power grid as an independent medium of the transfer of energy from the package to the devices in the IC, some second-order technology-related effects can sometimes cause unforeseen problems. This article focuses especially on the relationship of the power delivery system to the silicon substrate properties, and shows how a low-impedance substrate can make a substantial difference in the noise generated by the power grid.
—Sani R. Nassif, IBM Austin Research Laboratory

In modern VLSI designs, the substrate consists of either a lightly or highly doped p+ material. Several digital designs use highly doped substrates to take advantage of their very low resistance (typically a few ohms per square).

These designs use low-resistance substrates with an epitaxial process in which a thin, lightly doped layer resides atop the highly doped substrate. Lightly doped substrates have a much higher resistance, and these have been typically favored for analog designs. Here, we consider only highly doped, low-resistance substrates, which many Motorola designs use. You can extend the modeling approach presented here to highly resistive substrates as well, although the simulation results you obtain will differ significantly.

The substrate impacts the power distribution network in two ways. First, the substrate provides an alternate path for the current to reach devices from the ground pads, and hence reduces the ground distribution network's dc voltage drop. Second, the parasitic capacitance between the substrate and the n-wells acts as a decoupling capacitance between the power (V_{DD}) and ground (V_{SS}) supply networks, and will reduce the V_{DD} and V_{SS} supplies' ac voltage swings. The substrate's presence therefore improves both the dc and ac voltage drops. Consequently, a power distribution analysis without modeling the substrate can lead to an overdesigned distribution network and wasted chip resources.[1]

Because the substrate aids the power distribution network's integrity, voltage fluctuations in the power distribution network also affect the integrity of voltages in the substrate. The current that either the V_{DD} or V_{SS} distribution networks inject into the substrate causes the substrate's voltage to vary in time and location on the die. This voltage fluctuation results in a variation of the bulk-to-source voltage (V_{BS}) of the MOS devices, which in turn causes these devices' threshold voltage to change. This threshold voltage fluctuation has both a temporal and spatial form. In digital circuit analysis, the small threshold voltage fluctuations are negligible. However, they can significantly impact analog circuits that rely on accurately matched devices with identical threshold voltages. Therefore, in mixed-signal designs, accurate substrate voltage simulation and spatial and temporal behavior analysis of these fluctuations is critical.

This article presents a new analysis approach, combining detailed power distribution network and substrate models. On-chip, package, and board-level models represent the power distribution network. The substrate model consists of a detailed extracted model for the design's analog portions and a simplified model for the digital portions. Because the size of this com-

0740-7475/03/$17.00 © 2003 IEEE Copublished by the IEEE CS and the IEEE CASS

bined power distribution and substrate model can be very large (more than 100 million RLC components for large designs), we propose an iterative simulation approach.

Substrate and power grid models

An earlier article proposed detailed models for the power distribution network in the chip and package, the switching currents, the decoupling capacitance, and techniques for simulating the combined model.[2] Here, we adopt these models and extend them to include the substrate. We describe the substrate model in detail, while only briefly reviewing the power network, current, and decoupling capacitance models.

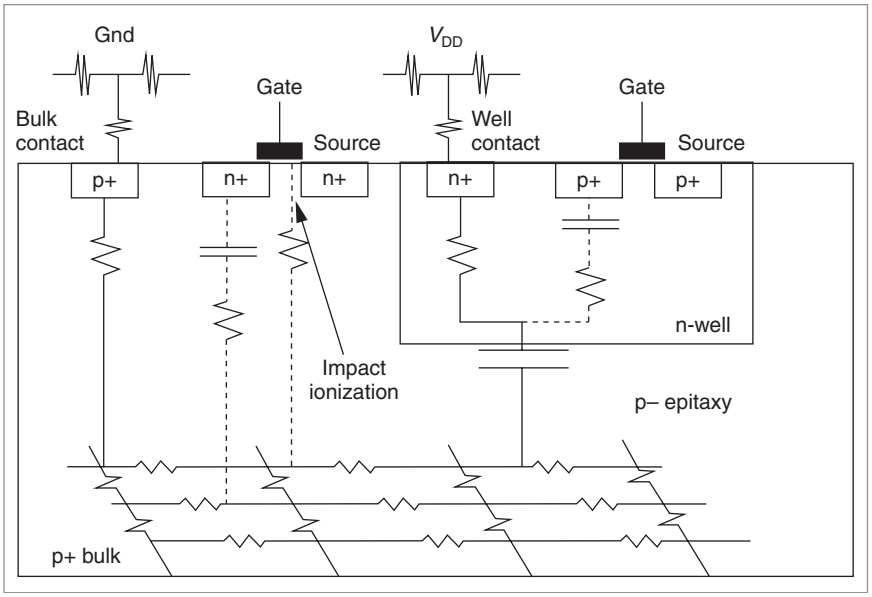

Figure 1. Electrical model of the substrate.

Substrate

The substrate model that we use assumes an n-well process. However, you can easily extend the model to other processes, such as the p-well or twin-well processes. Figure 1 shows a cross section of the substrate. The substrate consists of a p-epitaxial layer and a low resistive p+ buried layer. We tie the bulk to the V_{SS} network through p+ bulk contacts. We likewise tie the n-well to the V_{DD} network through the n+ well ties.

We inject noise into the bulk through three mechanisms:[3]

- resistive coupling,
- capacitive coupling, and
- impact ionization.

We inject V_{SS} network power supply noise into the bulk through resistive coupling from bulk contacts to the bulk. Power supply noise in the V_{DD} network couples resistively to the well through well ties, which in turn couple to the bulk through the well barrier capacitance. The source and drain terminals of devices inject noise through capacitive coupling through the source/drain diffusion capacitance. Finally, impact ionization causes current flow from the NMOS pinch-off point to the bulk. Of these mechanisms, the impact ionization is the least significant in magnitude. The noise injected from power rails is much stronger than the noise from the source and drain, because power rails couple resistively to the bulk and well, and the well couples strongly to the bulk due to the large well capacitance. On the other hand, the source and drain couple to the bulk only through very small capacitances. In comparison, for example, a chip's total diffusion capacitance is typically five times smaller than the total well capacitance, and its admittance at the typical gate switching speeds is 100 times smaller than the conductance between the bulk contact and the bulk or between the well contact and the well. Although the source and drain terminals can have voltage swings as large as the supply voltage, the amount of noise injected remains far smaller. In view of this, our substrate noise model considers only the power supply noise injection.

Figure 1 shows the electrical model for the substrate. The p+ buried layer acts as a conducting plane and is modeled as a 2D or 3D mesh of resistors. We employ the 2D model for the digital circuit section of the substrate using a constant bulk profile. We use a more accurate, 3D model for the substrate's analog circuit section using detailed substrate profile information at various depths. The lateral conduction through the highly resistive epitaxial is smaller than that of the p+ bulk by several orders of magnitude, and hence the effect of epitaxial is significant only as far as its vertical conduction from the bulk contacts to the p+ buried layer. Because the epitaxial layer's thickness is far smaller than the distance between two adjacent bulk contacts, the vertical resistances dominate the effective resistance between two bulk contacts. As such, the epitaxial layer is modeled by vertical resistances between the bulk contacts and

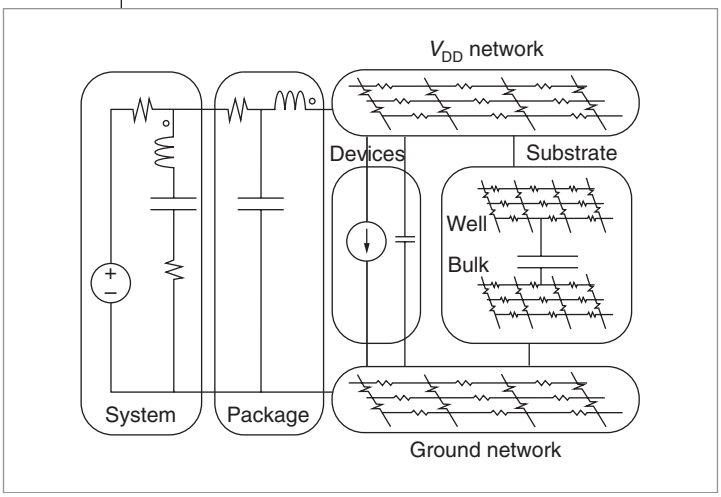

Figure 2. Simplified view of the combined model.

the bulk. We ignore the noise injected from source and drain terminals, as well as the noise due to impact ionization, for the reasons we discussed earlier, and Figure 1 uses dash lines to show the RC components corresponding to them.

In this study, we constructed the bulk's 2D resistive mesh in the chip's digital circuit section using the bulk's sheet resistance parameter, and the ties' vertical resistances using characteristic measurements of the process. Although we could have used commercial extraction tools for this task, we've found that our simpler model based on the bulk's characteristic measurements works well and is adequate for estimating the noise that the chip's digital section injects. For the substrate's analog circuit section, however, we constructed a more accurate 3D model using a commercial extraction tool.

Power network and substrate

We model the V_{DD} and V_{SS} networks on the chip as resistive networks, using an accurate extraction tool. We then supplement this network with distributed decoupling capacitances, which the devices, interconnects, and explicit decoupling structures contribute. We estimate the devices' intrinsic decoupling capacitance (when not switching) using Spice simulation of the input impedance for representative circuit blocks set at random quiescent (stationary) states. We model the switching currents by statistically distributing triangular gate current pulses, which together produce the specified total current profile at the chip's input supply pins. We also extract an RLC model of the power supply system and package from the board and package layouts using commercial extraction tools.

Figure 2 shows a simplified view of the complete model (power network and substrate) that we used in this work for simulation. The resulting network is an extremely large (10^6 to 10^8 nodes) linear network. That network consists of R, L, and C elements representing the power grid, package, and substrate modeling the power supply. It also consists of independent, time-varying, current sources modeling the device-switching currents.

Simulation methodology

We present several techniques for efficient time domain simulation of very large RLC models, such as the one in Figure 2.[2,4] We also introduce new techniques for addressing simulation requirements of the substrate's inclusion in the model.

For very large power networks, it is customary to analyze the V_{DD} and V_{SS} networks individually; this limits the size of the network to be simulated. However, this approach causes some difficulty when the simulation includes a substrate. Because the substrate couples to both the V_{DD} and the V_{SS} networks (through bulk ties, well ties, and well capacitance), we must simulate the substrate and both power networks simultaneously. This is possible when the model is small, as is the case for analog power network and substrate models for analog circuits. But, given the limited memory and computing resources, this approach is infeasible for simulating the digital section. We therefore propose an iterative approach in which we simulate the substrate first with V_{DD} and then with V_{SS}, and we repeat this procedure until the voltages in the network converge. We base this approach on waveform relaxation techniques.[5] We use this approach only for simulating the digital power supplies and the digital circuit sections of the substrate noise. We simulate the analog power supplies and the substrate's analog circuit section as a whole. We solve the grid using a direct solver that relies on Cholesky factorization techniques.[4] For designs where the grid, including both the power distribution and the substrate networks, exceeds a few tens of millions of nodes, we use a hierarchical technique.[6]

When the V_{DD} network is simulated with the substrate, the noise from the V_{SS} network (obtained during the previous simulation) is injected at the bulk tie nodes. Likewise, the noise voltages obtained from the V_{SS} network simulation are injected at the well tie locations during the V_{DD} network simulations (with substrate). At the beginning of the iterative procedure, we set the other power rail's noise to 0, and terminate the iteration when the noise voltages do not change above

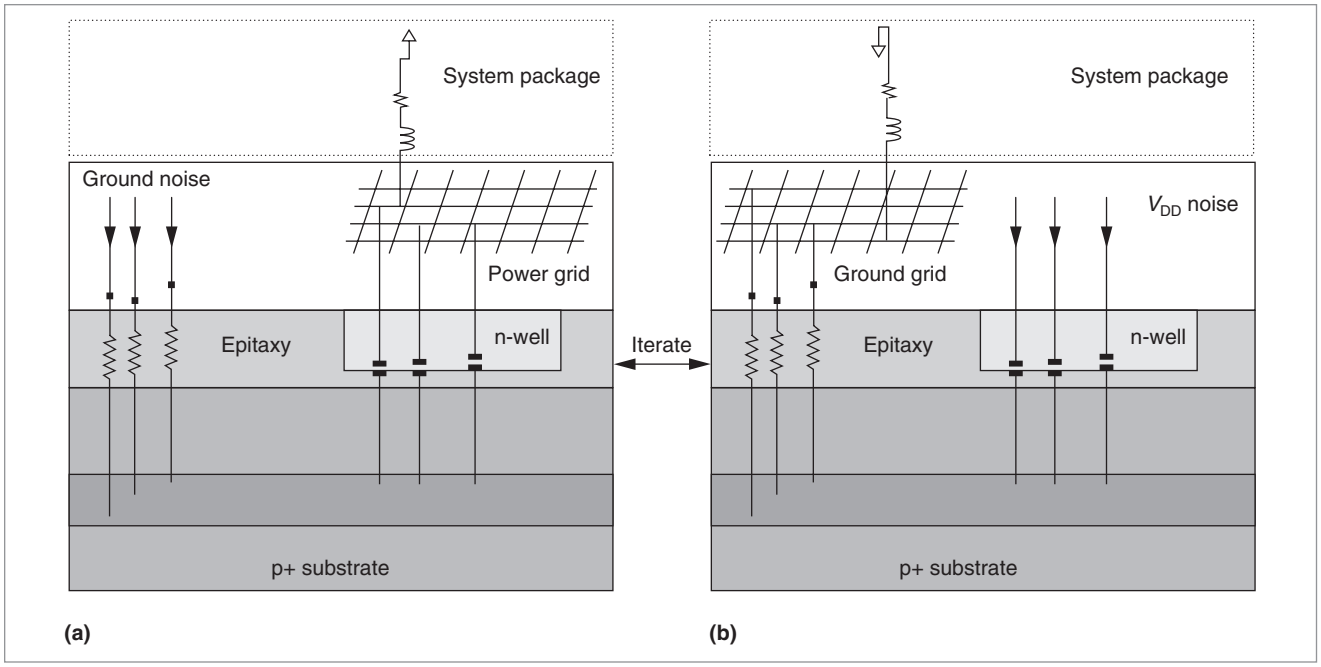

Figure 3. Iterative power (a) and ground (b) simulations with substrate grid.

a specified threshold. We observed that the iteration quickly converges within one or two iterations in all cases. Figure 3 illustrates these iterative simulations with a substrate grid, for digital circuits.

When the V_{DD} and V_{SS} networks of the design's analog block are simulated, the digital substrate noise must be propagated to the digital-analog substrate interface. We achieve this by defining various observation ports at this interface. During the digital V_{DD} and V_{SS} network simulations with the substrate, we observe the voltages at these ports. We use these voltages as bounding conditions for the analog block during the analog block's V_{DD} and V_{SS} network simulations (with substrate).

Table 1. Power grid and tie details for three processors.

Design	Average current (mA)	No. of nodes in the V_{DD} and V_{SS} grids (millions)	No. of well ties (millions)	No. of bulk ties (millions)
Chip 1	1501	8.1	0.3	0.2
Chip 2	300	8.2	0.6	1.3
Chip 3, digital	312	13.9	1.6	0.6
Chip 3, analog	NA	2.4	0.2	0.5

Table 2. Power supply noise due to the substrate.

	Maximum dc voltage drop (mV)		Maximum ac voltage variation (mV)	
Design	Without substrate	With substrate	Without substrate	With substrate
Chip 1	155	131	220	37
Chip 2	333	91	300	100
Chip 3, digital	102	20	119	23
Chip 3, analog			51	24

Substrate effect on power supply noise

We implemented and tested the proposed power distribution and substrate analysis approach on three processor designs. Table 1 gives information on the number of nodes in the power network and the substrate, as well as the number of ties for each test case. The first two designs are communication processors, and the third design is a mixed-signal design. Table 2 compares the voltage drops in the power network with and without the substrate's inclusion. For the dc analysis (columns 2 and 3 of Table 2), we report only the voltage drop in the V_{SS} network, because the substrate has minimal impact on dc voltage drop in the V_{DD} network. For ac analysis

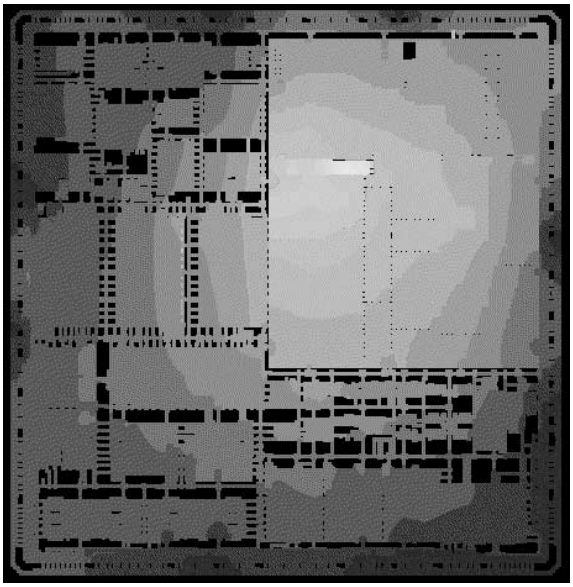

Figure 4. Voltage distribution in the ground network for chip 1 without the substrate.

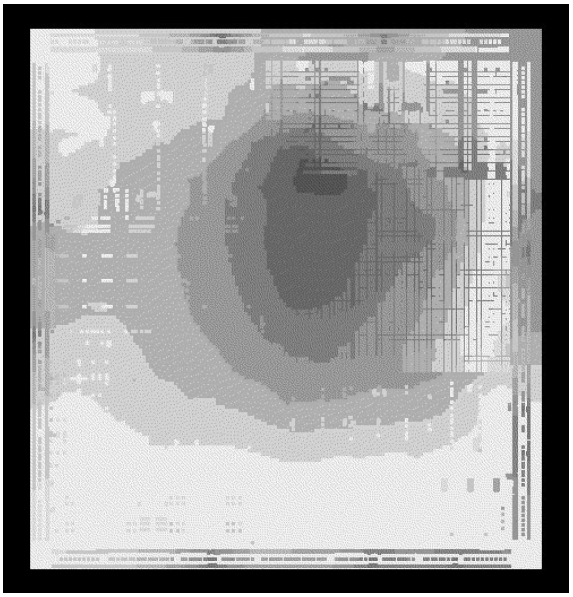

Figure 6. Maximum temporal voltage variation in the substrate for chip 1.

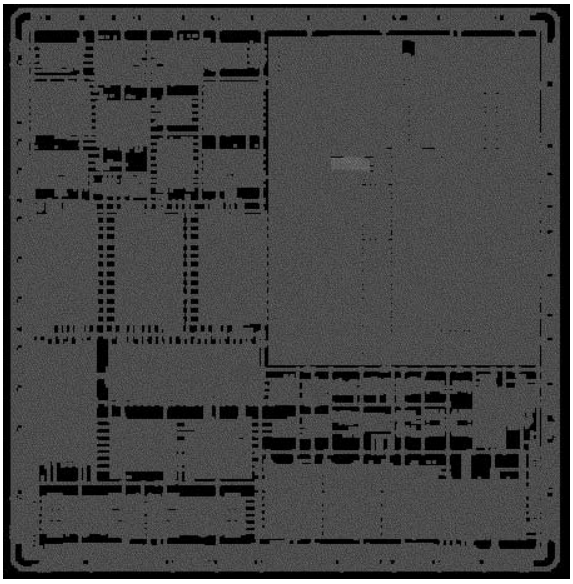

Figure 5. Voltage distribution in the ground network for chip 1 with the substrate.

(columns 4 and 5 of Table 2), we included a package RLC model in the simulation. The voltages shown are the maximum voltage fluctuation in time of a selected point in the V_{DD} and V_{SS} networks. We selected the point displaying the worst dc drop for tracking the ac voltage fluctuation. The results show that the substrate substantially improves the dc and ac voltage drops (15% to 80% for dc, and up to 83% for ac). In fact, if we performed the analysis without a substrate, some of these designs would have falsely appeared to fail their power distribution integrity requirements.

Figures 4 and 5 show the voltage distribution in one of the test case's V_{SS} networks at a typical time instance during ac simulation, without and with the substrate, respectively. In these pictures, a lighter shade indicates a higher voltage drop. As expected, Figure 4 shows that the worst voltage drop occurs toward the design's center. Figure 5 shows that the substrate substantially reduces this voltage drop and also creates a relatively uniform voltage drop across the die. This is particularly important for issues such as clock skew, which relies on a small spatial variation of the supply voltage. The significant power supply voltage improvement allowed the chip integrators to move approximately 10% of V_{SS} grid metal lines to the V_{DD} network. The result was a stronger power grid and an improved voltage in the power distribution network. Placing more metal lines in the V_{DD} grid compared to the V_{DD} and V_{SS} grid is now a common practice.

Substrate coupled noise

The data here pertains to the noise in the substrate nodes, and we obtained the data from the same simulations as in the previous section for the substrate's effect on power supply noise. Figure 6 shows the maximum temporal voltage variation (that is, the variation between the maximum and minimum voltage) for all

substrate nodes in chip 1. We collected the temporal variation data over a simulation time period of five clock cycles.

Table 3 shows the maximum temporal and spatial voltage variations for all three test cases. The temporal variation data corresponds to a selected point's voltages, and the spatial variation corresponds to a single time instance. Table 3 shows chip 3's digital and analog portions individually.

Table 3. Temporal and spatial noise in the substrate.

Design	No. of nodes in substrate (millions)	Maximum temporal variation (mV)	Maximum spatial variation (mV)
Chip 1	0.7	40.0	14.2
Chip 2	0.3	10.0	18.7
Chip 3, digital	0.3	2.0	30.0
Chip 3, analog	0.06	3.3	8.0

Although the temporal variation for chip 3 is relatively small, the spatial variation is large, possibly causing considerable quality degradation in the analog circuit's behavior.

The substrate's temporal variation for chips 2 and 3 is significantly smaller than for chip 1. Also, improvement in the dc voltage drop due to the substrate's inclusion is significantly higher for these designs. Two factors cause this: the number of ties, and their placement and distribution.

Chip 1 has significantly less substrate and well ties (see Table 1) than chips 2 and 3. Therefore, the impedance between the substrate and power supply network is relatively higher, resulting in more (temporal) substrate noise. However, the spatial voltage shows significant variation in chip 3. The substrate ties' placement is also significantly different between these designs. Figures 7 and 8 show the tie placement (shown as tie density per unit square) for chips 1 and 3. These two cases showed extreme values of temporal substrate voltage variation.

As for the placement and distribution of ties, the tie concentration for chip 3 is high near the design's periphery, close to the supply pads. These ties are very effective in maintaining good temporal substrate voltage control. In chip 1, the tie placement is more distributed in the center of the design (away from the power supply pads). However, we distributed the ties more spatially in chip 1, resulting in smaller spatial substrate noise than in chip 3. Therefore, the number of substrate ties and their placement play an important role in the stability of both the power distribution network and the substrate.

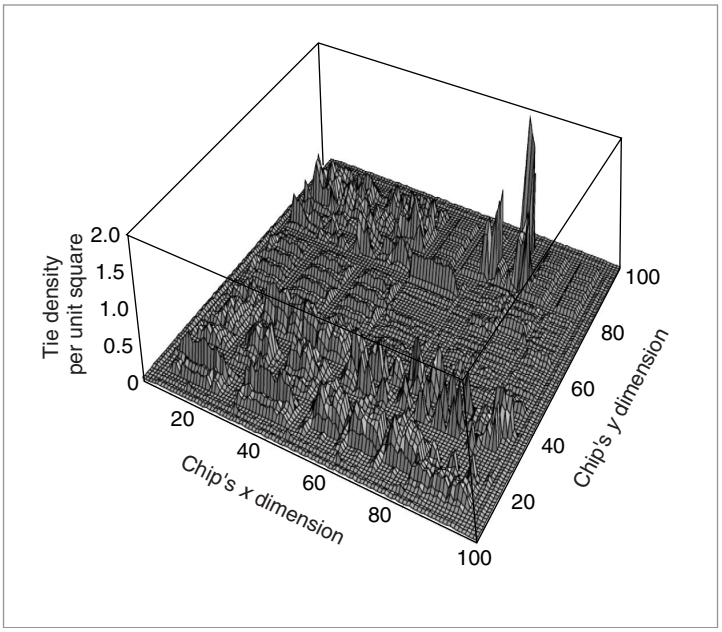

Figure 7. Tie placement for chip 1.

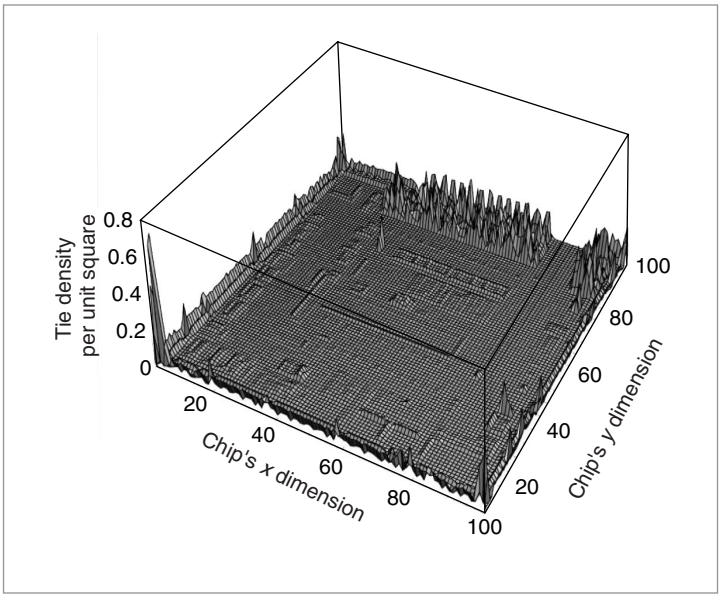

Figure 8. Tie placement for chip 3.

THE SIMULATION RESULTS we presented here demonstrate our claim that chip designers must consider the substrate's effect, to estimate power supply noise more realistically and thus to avoid overdesigning the power

distribution network. The proposed methodology also enables estimation of noise in the substrate, a parameter that is especially crucial in mixed-signal designs. Future work is needed in the area of proactively utilizing the substrate to improve power supply noise. This would involve adding additional substrate and well ties and determining their optimal placement. ■

■ References

1. R. Panda et al., "On the Interaction of Power Distribution Network with Substrate," *Proc. Int'l Symp. Low Power Electronics and Design* (ISLPED 01), ACM Press, 2001, pp. 388-393.
2. R. Panda et al., "Model and Analysis for Combined Package and On-Chip Power Grid Simulation," *Proc. Int'l. Symp. Low Power Electronics and Design* (ISLPED 00), ACM Press, 2000, pp. 179-184.
3. M. Felder and J. Ganger, "Analysis of Ground-Bounce Induced Substrate Noise Coupling in a Low Resistive Bulk Epitaxial Process: Design Strategies to Minimize Noise Effects on a Mixed-Signal Chip," *IEEE Trans. Circuits and Systems II: Analog and Digital Signal Processing*, vol. 46, no. 11, Nov. 1999, pp. 1427-1436.
4. A. Dharchoudhyry et al., "Design and Analysis of Power Distribution Networks in PowerPC Microprocessors," *Proc. 35th Design Automation Conf.* (DAC 98), ACM Press, 1998, pp. 738-743.
5. L.T. Pillage, R.A. Rohrer, and C. Visweswariah, *Electronic Circuit and System Simulation Methods*, McGraw Hill Professional, 1995.
6. M. Zhao et al., "Hierarchical Analysis of Power Distribution Networks," *Proc. 37th Design Automation Conference* (DAC 00), ACM Press, 2000, pp. 150-155.

Rajendran Panda manages the High Performance Design Technology group for Motorola in Austin, Texas. His research interests include work on low-power design, power supply and signal integrity issues, and circuit optimization. Panda has a BE (bachelor of engineering) in electrical engineering from Madurai University, an LLB (bachelor of laws) from Bangalore University, and a PhD in computer science from the University of Illinois at Urbana-Champaign.

Savithri Sundareswaran is a principal member of the technical staff in the High Performance Design Technology Group at Motorola in Austin, Texas. Her research interests include impact of process variations in VLSI and ultralarge-scale integration (ULSI) design analysis for power/timing/reliability, and CAD for VLSI. Sundareswaran has a BE in electrical engineering and an MS in physics from Birla Institute of Technology and Sciences, India.

David Blaauw is an associate professor at the University of Michigan, Ann Arbor. His research interests include VLSI design and CAD, with particular emphasis on circuit analysis and optimization problems for high-performance and low-power designs. He has a BS in physics and computer science from Duke University, and an MS and PhD in computer science from the University of Illinois at Urbana-Champaign.

■ Direct questions and comments about this article to David Blaauw, 2214C EECS, 1301 Beal Ave., Dept. of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109-2122; blaauw@umich.edu.