

# On-Chip Decoupling Capacitor Optimization Using Architectural Level Prediction

Mondira Deb Pant, *Member, IEEE*, Pankaj Pant, *Member, IEEE*, and Donald Scott Wills, *Senior Member, IEEE*

**Abstract**—Switching activity-generated power-supply grid-noise presents a major obstacle to the reduction of supply voltage in future generation semiconductor technologies. A popular technique to counter this issue involves the usage of decoupling capacitors. This paper presents a novel design technique for sizing and placing on-chip decoupling capacitors based on activity signatures from the microarchitecture. Simulation of a typical processor workload (SPEC95) provides a realistic stimulation of microarchitecture elements that is coupled with a spatial power grid model. Evaluation of the proposed technique on typical microprocessor implementations (the Alpha 21264 and the Pentium II) indicates that this technique can produce up to a 30% improvement in maximum noise levels over a uniform decoupling capacitor placement strategy.

**Index Terms**—Decoupling capacitors, ground bounce, signal integrity.

## I. INTRODUCTION

**S**IGNAL INTEGRITY is emerging as an important issue in today's deep submicron design. To preserve signal integrity, every circuit must have an adequate noise margin to allow for signal degradation. Rapid changes in supply current resulting from fast-switching circuits generate voltage fluctuations in the power-distribution system (commonly known as ground bounce), thereby limiting performance. In this paper, ground bounce is used to refer to changes in either supply voltage and ground voltage due to external-switching activities. Previously, ground bounce noise [1], [2], has been small compared to typical CMOS circuit noise margins. However, as designs begin to exploit deep submicron CMOS technologies with smaller feature size, faster switching speed, higher circuit density, and lower supply voltages, the problem of ground bounce is expected to increase significantly [3].

Typically, the amount of variation in the power supply level is modeled as  $\Delta V = L\delta I/\delta t$ , where  $\delta I$  is the current change during the transition,  $\delta t$  is the rise or fall time and  $L$  is the effective wire inductance of the power buses between the power supply and the current source. Excessive ground-bounce noise may not only introduce additional signal delay, but may also cause incorrect switching of logic gates. To maintain signal reliability there is a need for developing a very robust power-distribution network.

Manuscript received July 13, 2000; revised August 20, 2001. This work was supported by MARCO under 98-IT-674 of the DARPA Grant MDA-972-99-1 at the Georgia Institute of Technology, Atlanta.

M. D. Pant and P. Pant are with the Massachusetts Microprocessor Design Center, Intel Corporation, Shrewsbury, MA 01545 USA.

D. S. Wills is with the Microelectronics Research Center, Electrical & Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0250 USA.

Publisher Item Identifier S 1063-8210(02)06303-5.




Fig. 1. Projections of worst case  $di/dt$  noise on future microprocessors.

Projections made on worst case on-chip  $\delta I/\delta t$  using numbers from the international technology roadmap for semiconductors (ITRS'99) [4] indicate an exponential increase in the current slew rate as progress is made toward year 2014 (see Fig. 1). The maximum projected  $\delta I/\delta t$  was estimated as

$$\max(\delta I/\delta t) \approx \frac{I_{\max}}{T_c} \approx \frac{(\alpha P_{\max})/V_{dd}}{T_c}$$

where, the maximum activity factor  $\alpha$  ( $\approx 10\%$ ) indicates the fraction of the total power being drawn on the average.

To counter the effect of switching noise, decoupling capacitors are added near the switching current sources [1], [5]. These capacitors act as local reservoirs of charge for switching circuits and reduce the effect of the power-supply glitches on neighboring circuits. Optimal value and placement of decoupling capacitors is essential to ensuring a robust power distribution network. For low-frequency ground bounce, it may be adequate to use only off-chip decoupling capacitors to alleviate the problem. However, at high frequencies, the excessive power supply voltage swings  $\Delta V$  necessitates the use of on-chip decoupling capacitors due to their proximity to the switching units.

Estimation of the optimal configuration of decoupling capacitors is difficult due to the simulation challenges posed by large nonlinear power distribution networks that need to be analyzed. Furthermore, it also depends on the actual instantaneous current distribution drawn from the power distribution network. This in turn depends on the layout of the IC and the instruction workload being executed.

A large amount of research has been conducted to develop efficient models for fast power-distribution network simulation [6], [7]. Previous work on on-chip capacitor placement used abstract circuit-level models for generating switching activities and calculating decoupling requirements [8], [9]. Work has also been done to predict the worst-case current profiles [10].

In this paper, a methodology is presented to determine the value and placement of on-chip decoupling capacitors for general purpose microprocessor architectures with the goal of reducing on-chip power supply noise. The uniqueness of the proposed work stems from use of architectural level current signatures for obtaining the switching activities to determine the extent of on-chip decoupling required. The current signatures have been obtained by running SPEC95 benchmarks on the Alpha 21 264 and the Pentium II architectures. Since on-chip decoupling can occupy as much as 10% of the total chip area, it is imperative to capture accurate current signatures of the circuitry for determining not only the amount of decoupling required, but also where to place them on the chip. The goal is to be able to predict decoupling requirements to attenuate the problem of ground bounce in the early phases of a design cycle itself. In the sections that follow, detailed descriptions will be provided on how this target may be realized. Most of the discussion revolves around the Alpha 21 264 processor. Similar studies were done on the Pentium II architecture and experimental results on both these architectures clearly highlight the benefits of the proposed scheme. Preliminary results were presented at [11].

## II. ARCHITECTURAL MODELING

### A. Power-Distribution Modeling

In order to analyze the noise distribution and estimate the value and placement of decoupling capacitors, there is a need to develop an on-chip power-distribution model which includes the power grid as well as the power sources and drains [12]. Typically, power distribution within an integrated circuit is done from the top-level metal layer, which is connected to the package, down through interlayer vias and finally, to the active devices. The metal wires and vias are well modeled as a linear, time invariant and passive network consisting of resistive, capacitive and inductive elements. For modern integrated circuits such as microprocessors, this type of network can easily include millions of nodes. The models of power sources and drains, can be quite complex. Power-source models often include sophisticated package and board models. Power-drain models can account for the complex interaction between the power grid, the underlying nonlinear circuit and the time-varying signals propagating across the chip. However, huge grid sizes make it unfeasible to include any but the simplest models for power sources and drains.

For the experiments conducted in this paper, the power-distribution grid was modeled as a rectangular mesh network of resister-inductor-capacitor (RLC) elements, with each segment represented by an equivalent SPICE RLC model, as shown in Fig. 2. Each link in the network was assumed to be of the same length with identical electrical characteristics. The power sources were modeled as simple constant-voltage sources while




Fig. 2. Equivalent RLC power-bus network.

the power drains were represented by time-varying current sources. The number of links in each dimension were selected to conform with the aspect ratios of the Alpha 21 264 and the Pentium II dies, using published dimensions for the two processors [13], [14]. Rectangular grids of dimensions  $15 \times 20$  and  $23 \times 24$  were used for the Alpha 21 264 and the Pentium II processors, respectively.

Once the number of links in the power-distribution network have been decided, the length of each link can be determined (again using the published die dimensions), which allows the computation of the required RLC values for the network [5]. The external power-supply connections were assumed to be present along all four edges of the grid.

Next, the power network was mapped onto the layout of the processor (Figs. 3 and 4). Each functional block was represented by a set of current sources connected to those nodes in the power network that mapped onto it. The currents drawn by these sources directly reflect the switching activity of the functional blocks. The estimation of these values is detailed in Section II-B. The current sources have been assumed to be uniformly distributed among all the grid nodes covered by the functional block.

It should be noted that this simulation model was chosen to demonstrate the effectiveness of the methodology. In real world situations, the power-distribution strategy is more likely to be irregular, being more dense in the areas that are expected to draw higher currents. Moreover, accurate models for the power and current sources could be used. The algorithms presented should be easily extendible to all these more complex cases.

### B. Current Prediction

To model the switching activities of the functional units, the SPEC95 benchmark suite was simulated on the processors (modeled using the Alpha 21 264 and Pentium II architectural

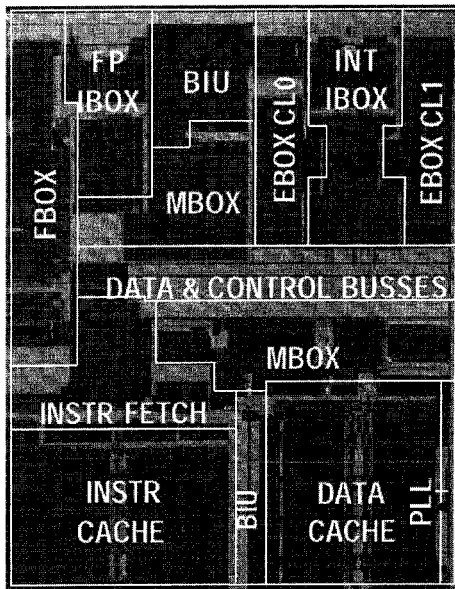


Fig. 3. Floorplan of the Alpha 21 264 processor.

TABLE I  
AVERAGE CURRENT DISTRIBUTION  
FOR THE ALPHA 21 264

Unit ( <i>FU</i> )	<i>I<sub>FU</sub></i>
FP Ibox	3.79A
MBox	8.02A
BUI	2.31A
EBox CL0	3.315A
EBox CL1	3.315A
INT Ibox	5.15A
Data/Control buses	2.67A
Instr Fetch	4.33A
Instr Cache	2.76A
Data Cache	5.25A

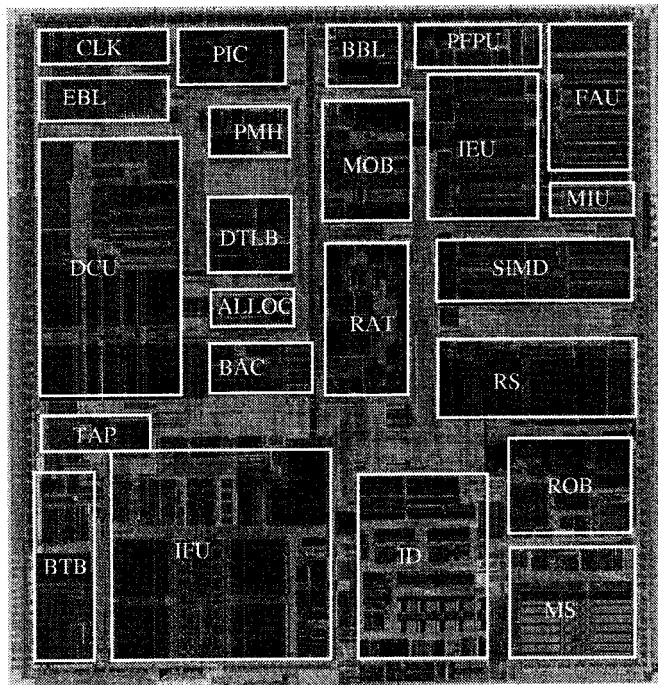


Fig. 4. Floorplan of the Pentium II processor.

TABLE II  
AVERAGE CURRENT DISTRIBUTION FOR THE PENTIUM II

Unit ( <i>FU</i> )	<i>I<sub>FU</sub></i>	Unit ( <i>FU</i> )	<i>I<sub>FU</sub></i>
CLK	0.910A	EBL	0.910A
DCU	1.274A	TAP	0.182A
BTB	0.546A	IFU	2.548A
BAC	0.364A	ALLOC	0.364A
DTLB	0.364A	PMH	0.364A
PIC	0.182A	BBL	0.182A
MOB	0.728A	RAT	0.728A
ID	2.548A	MS	0.546A
ROB	1.274A	RS	0.910A
SIMD	0.546A	IEU	0.910A
FEU	1.092A	MIU	0.182A
PFPU	0.546A		

specifications). The simulations were performed using the SimpleScalar tool set [15] developed at the University of Wisconsin-Madison. This tool set provides a fast, flexible, and accurate simulation of a processor that implements the SimpleScalar architecture (a close derivative of the MIPS architecture [16]). The advantage of using this tool is that standard benchmarks (SPEC95, etc.) can be compiled for the SimpleScalar instruction set and evaluated against any specified architecture.

The simulations provided the average number of cycles during which each functional unit was active. Using published actual power consumption values of the functional units [13],

the average power consumed by each unit ( $P_{FU_i}$ ), under typical workloads, was determined. The associated current for each functional unit was then approximated as  $I_{FU_i} = P_{FU_i}/V_{dd}$ , where  $P_{FU_i}$  is the average power consumed and  $V_{dd}$  is the power supply voltage.

Table I shows the current values of the functional units of the Alpha 21 264 estimated as described above using 2.2 V as the power-supply voltage. Similar results obtained for the Pentium II processor (using a supply voltage of 2.0 V) are reflected in Table II.

The average current obtained for each unit over a number of SPEC95 benchmarks was used to derive the triangular current




Fig. 5. Power-distribution impedance for Alpha 21 264.

waveform which reflects the current signature of that functional unit. With a goal to demonstrate the effectiveness of the scheme, this model was considered to be sufficient.

The following section describes how the value of the decoupling capacitors is estimated and discusses the experiments conducted to evaluate the effectiveness of various placement strategies for the Alpha 21 264 and the Pentium II.

### C. Decoupling Capacitance Optimization

The optimal amount of decoupling capacitance required to maintain noise within acceptable limits was estimated using simple back-of-the-envelope calculations [17]. The derivation of these expressions is provided below.

The average power supply current is the time integral of charge transferred throughout a clock cycle. The charge drawn during a burst of switching activity is

$$\delta Q = \frac{I}{2f_c}$$

where  $\delta Q$  is the charge per burst,  $I$  is the current drawn and  $f_c$  is the frequency of operation. The factor of 2 comes from the assumption that most logic activity occurs at both edges of the clock. The charge drawn during the transitions will come from the nearby decoupling capacitors, reducing the voltage across the capacitors by

$$\delta V = \frac{\delta Q}{C}$$

where  $C$  is the amount of decoupling capacitance. If  $n$  is the fraction of power-supply ripple that the circuit can tolerate, i.e., the allowable voltage swing is  $\delta V = nV_{dd}$ , then the amount of




Fig. 6. Power-distribution impedance for Pentium II.

decoupling capacitance required to maintain the power supply within a given ripple specification can be calculated as

$$C = \frac{\delta Q}{\delta V} = \frac{I}{2f_c n V_{dd}}. \quad (1)$$

In subsequent sections, experimental results demonstrate that these simple decoupling capacitance estimates can be very effectively used to determine the values and placement of the capacitors over the entire die.

### III. EXPERIMENTS AND RESULTS

The following four cases were evaluated to study the effect of possible decoupling capacitance selection strategies.

- Case 1) No decoupling capacitors were added. This situation was used as a benchmark to compare the various decoupling capacitor options.
- Case 2) A single decoupling capacitor was added at the center of each functional block. The value of the capacitor was proportional to the value predicted by (1) using the average current values for the functional units from Tables I and II.
- Case 3) Equal valued decoupling capacitors were placed at all the nodes of the power grid.
- Case 4) As in Case 3, decoupling capacitors were placed at each node of the power grid. However in this case, for each functional unit, the value of the capacitance used in Case 2 was distributed equally over all the

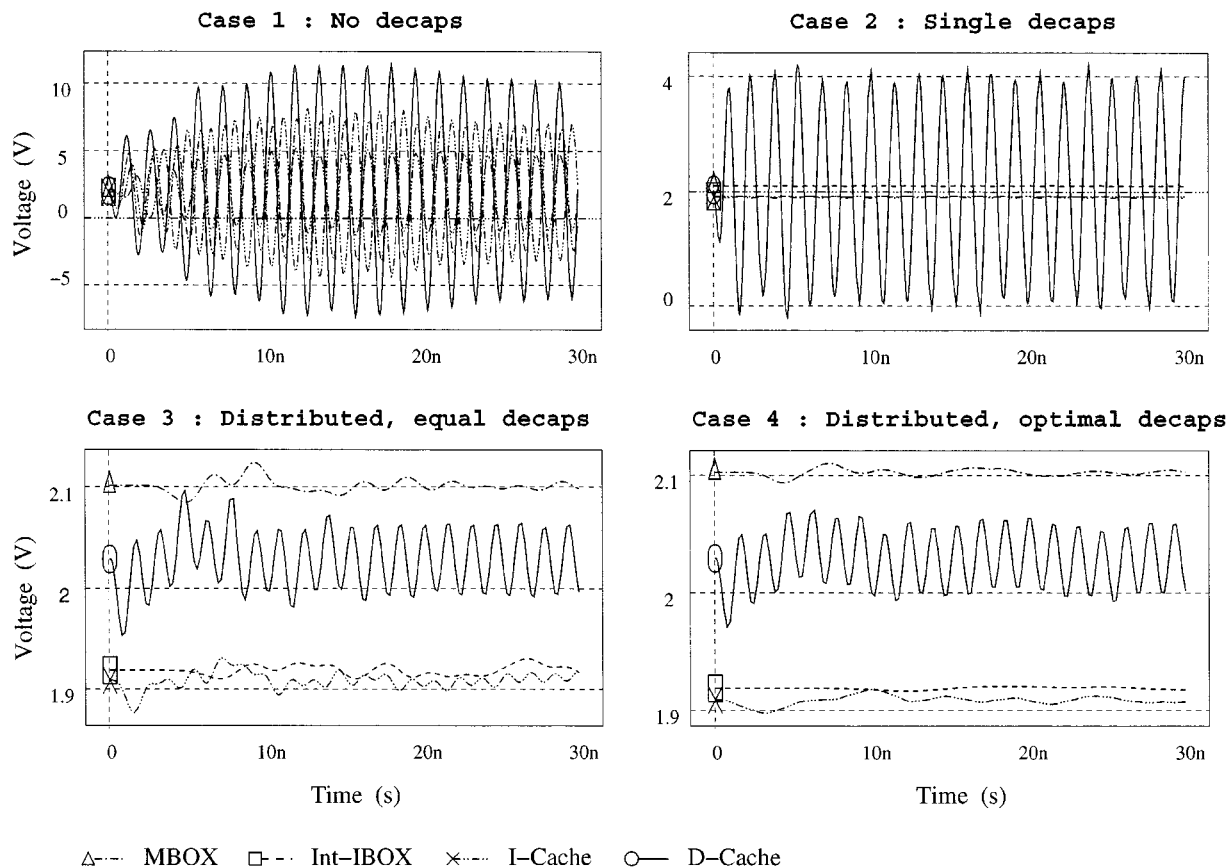


Fig. 7. Measured ground bounce for Alpha processor.

power-grid nodes that were mapped onto that functional unit. This is shown to provide the best noise attenuation of all the four cases.

In Cases 2, 3, and 4, the capacitance values were scaled to ensure that the total amount of decoupling capacitance used over the entire die was constant. The total capacitance was 320 nF for the Alpha 21264 and 180 nF for the Pentium II. These choices reflect the actual total decoupling capacitance used on the Alpha 21264 and the Pentium II dies. Published values of the frequency,  $f_c$ , and the power supply voltage,  $V_{dd}$ , were used in (1). For the Alpha 21264 [13] these were 575 MHz and 2.2 V respectively while for the Pentium II [14] they were 450 MHz and 2.0 V, respectively. All the simulations were performed using HSPICE [18].

#### A. Frequency-Domain Studies

The first set of experiments were conducted to investigate the effect of the various decoupling capacitance optimization strategies on the impedance of the power distribution network.

Fig. 5 shows the power supply impedance at the center of the data cache of the Alpha 21264 processor plotted as a function of frequency. A 1 A ac current source was placed at the point of interest and the voltage across it at varying frequencies is shown in Fig. 5. This directly reflects the variation of the power grid impedance over the frequency range. Peaks in the impedance plot represent resonant frequencies where the ground bounce

problem is the worst. At these frequencies, a small voltage glitch in the power supply gets amplified and may easily affect neighboring circuits. Decoupling capacitors should be chosen to reduce any resonance peaks in the impedance plots for the frequency range of interest.

For the case when no decoupling capacitors are used (Case 1), the plot indicates a resonance peak at about 450 MHz, which is below the maximum frequency of operation of 575 MHz. If single monolithic decoupling capacitors were placed at the centers of each functional block (Case 2), there are no resonances in the range of interest, but one could encounter problems as one moves to higher frequencies. Using smaller distributed capacitors over the entire power grid provides much better attenuation. Upon further investigation, it can be seen that making the capacitors proportional to the currents drawn by the functional units (Case 4) is more efficient than using a uniform distribution of capacitors (Case 3).

Fig. 6 shows a similar plot for a node in the instruction fetch unit (IFU) of the Pentium II processor. The trend is seen to be the same as that of the Alpha 21264. In this example, one may clearly observe that calculated distribution of capacitors (as opposed to equal distribution) yields much better noise attenuation.

#### B. Time-Domain Studies

The next set of experiments quantifies the amount of ground-bounce for the decoupling capacitance schemes.

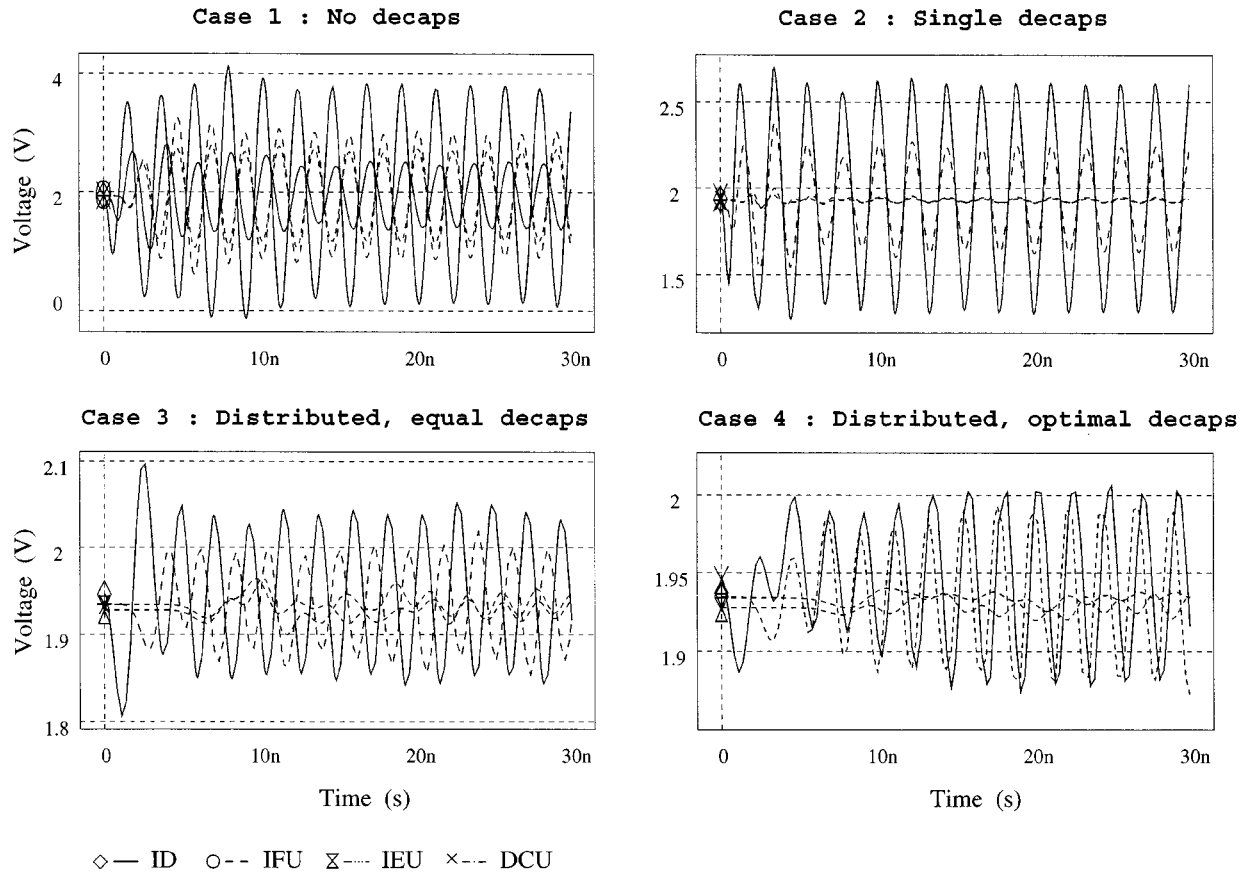


Fig. 8. Measured ground bounce for Pentium II processor.

For the Alpha 21264 processor, the data-cache (D-cache) was assumed to be switching at the resonant frequency of 450 MHz (Fig. 5). The current sources representing the D-cache (Section II-A) were assumed to be drawing triangular current waveforms that were scaled to ensure that the total average current for the D-cache was 5.25 A (from Table I). All the other units were kept in an inactive state during this experiment, i.e., they were drawing steady DC currents (again, the values were chosen to satisfy Table I). For each of the capacitance distribution strategies, ground bounce was measured at power grid points in the data cache (D-cache), the instruction cache (I-cache), the integer instruction unit (int-IBOX) and the lower-memory unit (MBOX), Fig. 3.

The primary aim of the experiment was to measure the effect of the switching currents in the D-cache on the neighboring units. Fig. 7 shows the ground bounce for the four-points selected. The ground bounce noise was predictably the most in the switching node itself. It can be seen that as the distance from the switching unit increased, the effect of the switching activity decreases and the voltage fluctuation diminish. Note that the different average current levels reflect the varying resistive (IR) drops due to the currents drawn by each unit.

The plots clearly indicate the resonance effects observed by running the D-cache at 450 MHz. Both Case 1, which does not have any decoupling, and Case 2, which uses a single decoupling capacitor per functional block, show significant voltage swings

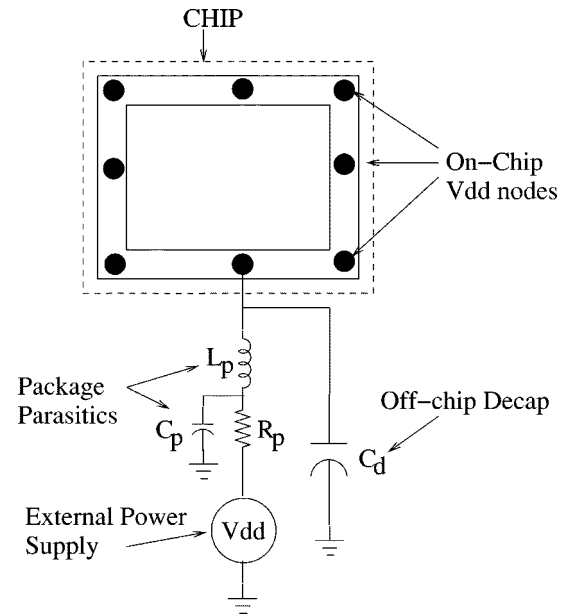


Fig. 9. Simple package level model.

in the D-cache. While the use of monolithic capacitors do suppress the glitches in the neighboring units, the scheme is still not very effective in containing the fluctuations in the D-cache. From the plots, it is clear that the best results are obtained when

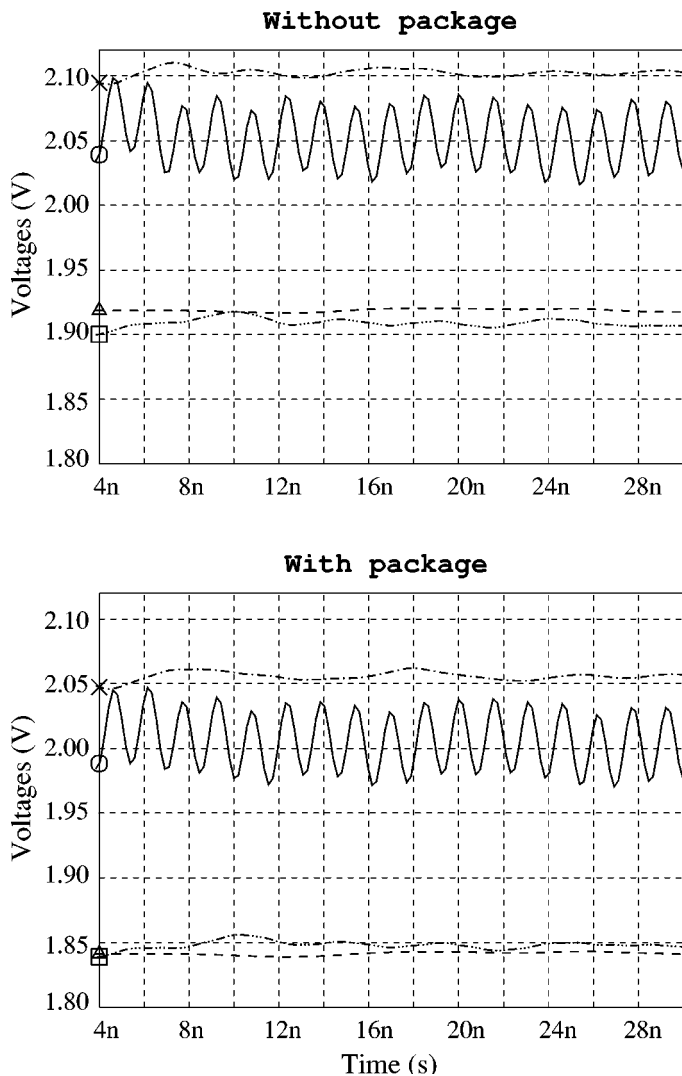


Fig. 10. Effect of adding off-chip decoupling capacitors.

the capacitors are distributed throughout the grid. Studying the last two cases in more detail indicates that selecting the values of the decoupling capacitors depending on the local switching activity profiles yields better noise suppression both in the time and frequency domains than uniformly distributing equal value capacitors throughout the die. The improvement for the D-cache is as much as 30 mV in the worst-case bounce, which is about 15%. Even for the nonswitching units, one can see that the optimal capacitance placement strategy performs significantly better.

In the case of the Pentium II, the instruction decode (ID) unit was kept switching at the resonant frequency of 260 MHz (Fig. 6), which is below its operating frequency of 450 MHz. As before, the current values for this experiment were selected to conform with values from Table II. Effects of ground bounce were observed at points in the instruction decode (ID), integer fetch (IFU), data-cache (DCU) and the integer execute (IEU) units and are plotted in Fig. 8. Once again the effectiveness of using an optimal distributed decoupling strategy to counter ground bounce is clearly demonstrated. An improvement of about 45 mV in the worst case ground bounce is observed for the ID unit which reduces from 2.045 to 2.00 V. With the

DC voltage for the unit being about 1.90 V, this improvement approximately equals 33%.

Since the same amount of total capacitance was used in all the capacitance distribution schemes, our decoupling strategy does not incur any additional area penalty. The scheme could also be used to optimize the amount of decoupling capacitance necessary for a prespecified noise level and given activity profiles.

The current signatures allow us to identify the potential hot spots where the most significant  $V_{dd}$  drop may occur. Hence, one can explore hybrid capacitance placement options where distributed decoupling capacitors are used in the highly active regions and centralized capacitors in other functional blocks.

### C. Packaging Effects

In order to assess the effects that the package may have, the experiments were repeated using package parasitics along with an off-chip decoupling capacitor, using the setup shown in Fig. 9. Unlike the previous experiments, the voltages at the on-chip power nodes are no longer constant. It was seen that even when good low  $L_{eff}$  packages with decoupling capacitors are used, there is a discernible ground-bounce noise introduced within the chip due to high on-chip current slew rates. Fig. 10 shows the ground-bounce observed for the same four points in the Alpha 21264 as in Fig. 7, with and without the package model. In both cases, the on-chip decoupling capacitors were distributed according to our placement strategy. In the presence of on-chip decoupling capacitors, off-chip decoupling capacitors hardly provide any further improvements as observed in Fig. 10.

It has always been known that low-effective inductance is crucial for controlling packaging switching noise levels. However, the trend for VLSI is to place higher-speed circuits at greater densities on a chip. Specifically, with the advent of the SOC philosophy, off-chip decoupling is no longer going to be able to control the high levels of switching inside the ICs. From our experiments it is evident that to augment this effect, as chips get more complex and increase in density, it will be necessary to place high-frequency capacitors in the chip as well as off-chip.

## IV. CONCLUSION

With the advent of higher clock speeds and smaller circuit geometry, on-chip noise problem is growing to be a major concern for high performance VLSI circuit design. Our methodology describes a unique technique of using architecturally obtained activity signatures to determine optimal values and placement of on-chip decoupling capacitors on typical microprocessor dies to alleviate power supply noise issues. Even though the models used in this research are very simplistic, the aim of our experiments has been to show the potential of the scheme. It is hoped that this method of predicting activity signatures will get incorporated in layout, placement, and parasitic determining tools where more elaborate models are used. As we continue to scale down the feature size and power-supply voltage in deep submicron circuits, our methodology of predicting the amount of decoupling required will be able to play a vital role in preserving the reliability, reducing the cost and achieving performance targets very early on in the design cycle of future VLSI circuits.

## REFERENCES

- [1] H. B. Bakoglu, *Circuits, Interconnections and Packaging for VLSI*. Reading, MA: Addison-Wesley, 1990.
- [2] Y.-S. Chang, S. K. Gupta, and M. A. Breuer, "Analysis of ground bounce in deep sub-micron circuits," in *Proc. VLSI Test Symp.*, 1997, pp. 110–116.
- [3] V. Tiwari, D. Singh, S. Rajagopal, G. Mehta, R. Patel, and F. Baez, "Reducing power in high-performance microprocessors," in *Proc. Design Automation Conf.*, June 1998, pp. 732–737.
- [4] "The 1999 International Technology Roadmap for Semiconductors," Semiconductor Industry Assoc., San Jose, Calif., Tech. Rep., 1999.
- [5] *Microelectronics Packaging Handbook*, 2nd ed., vol. I, Chapman & Hall, London, U.K., 1997.
- [6] M. Zhao, R. Panda, S. Sapatnekar, T. Edwards, R. Chaudhry, and D. Blaaw, "Hierarchical analysis of power distribution networks," presented at the Proc. Design Automation Conf., June 2000.
- [7] J. Choi, S. Chun, N. Na, M. Swaminathan, and L. Smith, "A methodology for the placement and optimization of decoupling capacitors for gigahertz systems," in *Proc. Int. Conf. VLSI Design*, Jan. 2000, pp. 156–161.
- [8] H. H. Chen and S. E. Schuster, "On-chip decoupling capacitor optimization for high-performance VLSI design," in *Proc. Int. Symp. VLSI Technology, System Applications*, 1995, pp. 99–103.
- [9] D. L. Smith, R. E. Anderson, D. W. Forehand, T. J. Pelc, and T. Roy, "Power distribution system design methodologies and capacitor selection for modern CMOS technology," *IEEE Trans. Comp., Packag., Manufact. Technol.*, vol. 22, pp. 284–291, Aug. 1999.
- [10] A. Krstic and K. Cheng, "Vector generation for maximum instantaneous current through supply lines for CMOS circuits," in *Proc. Design Automation Conf.*, June 1997, pp. 383–388.
- [11] M. D. Pant, P. Pant, and D. S. Wills, "On-chip decoupling capacitor optimization using architectural level current signature prediction," *Proc. IEEE ASIC/SOC Conf.*, pp. 286–292, Sept. 2000.
- [12] H. H. Chen and J. S. Neely, "Interconnect and circuit modeling techniques for full-chip power noise analysis," *IEEE Trans. Comp., Packag., Manufact. Technol. Part B: Advanced Packaging*, vol. 21, pp. 209–215, Aug. 1998.
- [13] D. Grunwald, T. Mudge, and S. Manne, "Cool chips tutorial," presented at the Proc. Symp. MicroArchitecture, Haifa, Israel, Nov. 1999.
- [14] L. Gewnnap, "Microprocessor Report," vol. 11, Pentium II debuts at 300 MHz, May 1997.
- [15] D. Burger and T. M. Austin, "The SimpleScalar Tool Set, Version 2.0," Univ. of Wisconsin-Madison Computer Science Department, Tech. Rep. 1342, June 1997.
- [16] C. Price, *MIPS IV Instruction Set, Revision 3.1*. Mountain View, CA: MIPS Technologies Inc., Jan. 1995.
- [17] L. Smith, "Decoupling capacitor calculations for CMOS circuits," in *IEEE 3rd Topical Meeting of Electrical Performance Electronic Packaging*, Nov. 1994, pp. 101–105.
- [18] *HSPICE User's Manual*, 1996.

**Mondira Deb Pant** (S'96–M'00) received the B.Tech. degree in computer science and engineering from the Indian Institute of Technology, Kharagpur, and the M.S. and Ph.D. degrees in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, GA, in 1995, 1997, and 2000, respectively.

She is currently a Senior Circuit Design Engineer with the Massachusetts Microprocessor Design Center at Intel Corporation Shrewsbury, MA. Her research interests include inductive noise analysis, low-power VLSI circuit design, and computer architectures.

Dr. Pant is a member the Eta Kappa Nu Electrical Engineering honor society.



**Pankaj Pant** (S'96–M'00) received the B.Tech. degree in computer science and engineering from the Indian Institute of Technology, Kharagpur, and the M.S. and Ph.D. degrees in electrical engineering from the Georgia Institute of Technology, Atlanta, in 1995, 1997, and 2000, respectively.

He is currently a Senior DFT Engineer with the Massachusetts Microprocessor Design Center at Intel Corporation, Shrewsbury, MA. His research interests include test and diagnosis of digital circuits and design of low-power VLSI circuits.

Dr. Pant is a Member of the Eta Kappa Nu Electrical Engineering honor society.



**Donald Scott Wills** (S'79–M'90–SM'98) received the B.S. degree in physics from Georgia Institute of Technology, Atlanta, in 1983, and the S.M., E.E., and Sc.D. degrees in electrical engineering and computer science from Massachusetts Institute of Technology (MIT), Cambridge, in 1985, 1987, and 1990, respectively.

He is currently an Associate Professor of Electrical and Computer Engineering at the Georgia Institute of Technology. His research interests include short wire VLSI architectures, high throughput portable processing systems, architectural modeling for gigascale (GSI) technology, and high efficiency image processors.

Dr. Wills is a Member of the Computer Society and an Associate Editor of IEEE TRANSACTIONS ON COMPUTERS.