

Decoupling Capacitance Allocation and Its Application to Power-Supply Noise-Aware Floorplanning

Shiyou Zhao, Kaushik Roy, *Fellow, IEEE*, and Cheng-Kok Koh, *Member, IEEE*

Abstract—We investigate the problem of decoupling capacitance (decap) allocation for power supply noise suppression at floorplan level. First, we assume that a floorplan is given and consider the decap placement as a postfloorplan step. Second, we consider the decap placement as an integral part of a floorplanning methodology (noise-aware floorplanning). In both cases, the objective is to minimize the floorplan area while suppressing the power supply noise below the specified limit. Experimental results on MCNC benchmark circuits show that, for postfloorplan decap placement, the white space allocated for decap is about 6%–9% of the chip area for the 0.25- μm technology. The power-supply noise is kept below the specified limit. Compared to postfloorplan approach, the peak power-supply noise can be reduced by as much as 40% and the decap budget can be reduced by as much as 21% by using noise-aware floorplanning methodology. The total area is also reduced due to the reduced total decap budget gained from reduced power supply noise.

Index Terms—Decoupling capacitance, floorplan, power supply noise.

I. INTRODUCTION

SIGNAL integrity is emerging as an important issue as very large scale integration (VLSI) technology advances to nanoscale regime. Of particular importance among the signal integrity issues is the power-supply noise. In today's deep submicrometer complementary metal–oxide–semiconductor (CMOS) technology, devices are of smaller feature size, faster switching speed, and higher integration density. Large current spikes due to a large number of “simultaneous” switching events in the circuit within a short period of time can cause considerable current–resistance (IR) drop and Ldi/dt (delta-I) noise over the power-supply network [1]. Power-supply noise degrades the drive capability of transistors due to the reduced effective supply voltage seen by the devices. Power-supply noise may also introduce logic failures and jeopardize the reliability of high performance VLSI circuits, since the noise margin gets lower as the supply voltage scales with the tech-

nology. Recently, many research efforts [2]–[7] have been directed toward power-supply noise analysis and power-supply network optimization. Topology optimization [8], wire sizing [9], onchip voltage regulation [10], and decoupling capacitance (decap) deployment [2], [11] are the most widely used techniques to relieve power-supply noise.

Wire sizing is an important factor in power/ground (P/G) network optimization. A robust P/G network is, however, hard to achieve purely by adjusting the sizes of the wire segments as inductive noise becomes more pronounced as CMOS technology scales. Since inductance scales poorly with sizing [1], decap placement becomes an indispensable technique for robust P/G network design for high-performance microprocessors. In the past, decap optimization has been investigated at circuit level or system level [11], [12] with the assumption that there is always white (empty) space available for decap.

In this paper, we investigate the problem of onchip decap deployment for global (P/G) mesh at floorplan level. Given a floorplan with the placement information and the worst case switching-activity profile of each circuit module, we want to find an area efficient scheme to deploy the decap such that the power-supply noise at each module is suppressed to below a specified limit. We estimate the worst case noise in the power-supply network experienced by each module according to the placement information and switching profiles. Based on the worst case power-supply noise, we calculate decap budget for each circuit module. We allocate white space (WS) for metal–oxide–semiconductor (MOS) decoupling capacitors in two steps. Existing WS is first allocated to the neighboring blocks using a linear programming (LP) technique to maximize the utilization of the existing WS in the floorplan. Additional WS, if needed, is inserted into the floorplan in an incremental fashion to meet the total decap demand of the whole circuit.

Based on the proposed decap placement algorithm, we also propose a power-supply noise-aware floorplanning methodology. Compared to conventional floorplanning algorithm, the noise-aware floorplanning methodology takes the power-supply noise, as a factor of the cost function, into consideration. The noise-aware floorplanning algorithm arranges modules based on the switching activities and the spatial correlations between modules such that the peak power-supply noise is lowered and, therefore, the decap budget is reduced. For example, a cluster of high switching-activity modules can overload specific power pins and generate a noisy spot in the floorplan, while a scattered distribution of high switching-activity modules can lead to

Manuscript received April 20, 2001; revised July 12, 2001. This research is supported in part by the Semiconductor Research Corporation under Grant C 99-TJ-689, in part by the National Science Foundation under Grant CCR-9984553, and in part by the Intel Corporation. This paper was presented in part at the International Symposium on Physical Design, Sonoma, CA, April 2001. This paper was recommended by Guest Editor S. S. Sapatnekar.

S. Zhao was with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907-1285 USA. He is now with Micron Semiconductor, Boise, ID 83707 USA.

K. Roy and C.-K. Koh are with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907-1285 USA.

Publisher Item Identifier S 0278-0070(02)00098-2.

reduced peak power-supply noise and decap budget. Similar ideas have been applied to thermal placement [13], [14] to smooth out the hot spots and to substrate aware mixed-signal macrocell placement [15] to reduce the substrate coupling.

The rest of the paper is organized as follows. Problem formulation is given in Section II. Power-supply noise estimation is presented in Section III. Decap budget and white-space allocation are discussed in Sections IV and V. Power-supply noise-aware floorplanning is addressed in Section VI. Experimental results are presented in Section VII. Finally, conclusions are drawn in Section VIII.

II. PROBLEM FORMULATION

Decap deployment is an important issue in high-performance VLSI design. Onchip decoupling capacitors are usually fabricated as MOS capacitors. The unit area capacitance of a MOS capacitor is given by $C_{ox} = \epsilon_{ox}/t_{ox}$, where t_{ox} is the oxide thickness and ϵ_{ox} is the permittivity of SiO_2 . The decap budget can be converted to the area of silicon required for the decap fabrication as follows:

$$S_{\text{decap}} = \frac{C_{\text{decap}}}{C_{ox}} \quad (1)$$

where C_{decap} is the decap budget and S_{decap} is the silicon area required to fabricate C_{decap} .

In practice, decoupling capacitors are placed wherever there is WS available on the chip. Since the decoupling capacitors are placed blindly, there is no guarantee that they are placed at the right places with the right amount of capacitance to suppress the power-supply noise below the specified noise limit.

In this paper, two different scenarios of the decap deployment problem are investigated. In the first scenario, decap allocation is treated as a postfloorplanning step. Given a floorplan, we determine where to place the decoupling capacitors and how to allocate the WS to the modules such that the final chip area is minimized and the power-supply noise is suppressed below the specified limit. The proposed methodology determines the decap budgets based on the power-supply noise experienced by each circuit module and then allocates the right amount of WS to each module in its vicinity for decap fabrication. The postfloorplanning decap allocation problem is addressed in Sections IV and V.

In the second scenario, the decap allocation is handled as an integral part in the floorplanning process. Given the worst case switching-activity profiles of the circuit modules, we seek a floorplanning methodology to produce the optimal floorplan such that the floorplan area and wire length are minimized and the decaps are properly placed in the floorplanning process so that the power-supply noise is kept below the specified limit. We propose a power-supply noise-aware floorplanning methodology in which the power-supply noise is incorporated into the cost function. The noise-aware floorplanning algorithm evaluates the merit of a candidate floorplan based on the power-supply noise at each circuit module and the associated cost due to decap. Since there is more freedom for floorplan optimization, power-supply noise-aware floorplanning methodology can result in lower peak power-supply noise




Fig. 1. Power-supply network-mesh structure.




Fig. 2. Model of power-supply network.

and, consequently, smaller decap budget and chip area. The details about the power-supply noise-aware floorplanning methodology are discussed in Section VI.

III. POWER-SUPPLY NOISE ESTIMATION

Decap is allocated to each module based on its switching profile and the power-supply noise it experiences. To determine the decap demand of each module, we must estimate the power-supply noise at each module in the floorplan. In the following subsections, we will discuss power-supply network modeling, switching current distribution and noise estimation.

A. Power-Supply Network Modeling

In today's VLSI technology, most power-supply networks are of a mesh structure, as illustrated in Fig. 1. This research targets at decap deployment for global P/G mesh since there is no detailed information regarding lower level P/G structure at the floorplanning phase. We make the following assumptions: 1) all the segments of the mesh grids are of the same physical dimensions and 2) the connection points of the circuit modules to the power grids are determined by the locations of the centers of the modules. We model each segment of the power grids as a lumped resistance-inductance-capacitance (RLC) element and the whole mesh as a pseudodistributed RLC network, as illustrated in Fig. 2. The ground (GND) node in Fig. 2 should be regarded as a ground network with a similar mesh structure, as illustrated in Fig. 2. The unit length parasitics r , l , and c are technology dependent. The package parasitics of the power pins are R_p and L_p . The

unit length inductance l should be regarded as the effective inductance per unit length in the power-supply grids. The circuit blocks are modeled as time-varying current sources that draw current from the power-supply (VDD) sources through their connection points in the power-supply grids. Since the circuit should operate correctly even under the worst case scenario, we use the worst case switching-activity profiles of the circuit blocks to deduce the current waveforms for power-supply noise estimation. We also assume a worst case timing alignment of the current sources for all the modules.

B. Current Distribution

Given the mesh topology and the switching current waveforms of the circuit modules, we can approximately determine the distribution of those switching currents among the power-supply network. A key observation is that currents follow the least-impedance paths when flowing from the VDD source to the destination sink. In other words, if there are multiple paths from a VDD source to a destination sink, the current flowing along each path is inversely proportional to the impedance of the path. Based on this observation, we make the following assumption: the switching current drawn by a sink comes from only the neighboring VDD sources; the contributions from remote VDD sources are negligibly small and therefore can be ignored. This assumption significantly simplifies the current distribution analysis without compromising the validity of the results. The direct consequence of this assumption is that currents flowing along the neighboring grids of the sinks are slightly overestimated and consequently the power-supply noise at the connection points will be overestimated. With that assumption in mind, the question comes down to how the current drawn by a sink is split among the neighboring VDD sources or, in other words, how much current each neighboring VDD pin is contributing.

Suppose that there are N ($N = 4$ in most cases) neighboring VDD sources surrounding a sink. Let Z_1, Z_2, \dots, Z_N be the impedances between the current sink and the N neighboring VDD pins, respectively. Let I be the current a sink is sourcing from the power network. Let I_1, \dots, I_N be the currents contributed by the N neighboring VDD pins, respectively. I_1, \dots, I_N are given by the following equations:

$$I_1 + I_2 + \dots + I_N = I \quad (2a)$$

$$Z_1 I_1 = Z_2 I_2 = \dots = Z_N I_N \quad (2b)$$

$$Y_j = \frac{1}{Z_j}, \quad j = 1, 2, \dots, N \quad (2c)$$

$$\Rightarrow I_j = \frac{Y_j}{\sum_{i=1}^N Y_i} I, \quad j = 1, 2, \dots, N \quad (2d)$$

where Y_j is the admittance from the sink to VDD source j . Equation (2a) states that the contributions from the neighboring VDD pins sum up to the total current the sink is sourcing. Equation (2b) states that the voltage differences from the sink to different neighboring VDD pins are the same. Solving (2a)–(2c) gives the solution to I_1, \dots, I_N , as shown in (2d).

The impedance between a sink and a VDD pin in the power mesh is mainly determined by the least impedance paths that link them. The impedance of a path can be calculated based on




Fig. 3. Point-to-point paths classification.

its length and the unit length parasitics of the wire segments in the power grids. The equivalent impedance of the shortest paths, the second shortest paths, and so on connected in parallel will be a reasonable estimate of the impedance between the two points.

The shortest paths refer to those paths that are of least impedance. Similar interpretations apply to second shortest paths, third shortest paths, and so on. With the uniform sizing assumption, the geometrically shortest paths are the paths of least impedance. Similar arguments apply to second shortest paths and third shortest paths, etc. Fig. 3 illustrates the classification of the paths between point (A) and point (3). Thick lines are the shortest paths. The thinnest dash lines are the third shortest paths and the solid thin lines are the second shortest paths. The paths may overlap each other. When we calculate the equivalent impedance between the two points, we treat all the paths as separate nonoverlapping paths connected in parallel. This treatment of overlapping paths will partly compensate for those longer paths ignored in our model.

The impedance of a unit length wire is $Z(\omega) = r + j\omega l$. Suppose there are two paths of length w_1 and w_2 , respectively, connected in parallel. The equivalent impedance will be $((w_1 * w_2)/(w_1 + w_2))Z(\omega)$. The equivalent impedance between any two points in the P/G mesh can be expressed as a constant times $Z(\omega)$. In (2), all the impedances can be normalized to $Z(\omega)$. The ratios between the impedances are determined by the geometrical dimensions of those paths since $Z(\omega)$ will eventually be canceled out. So, the current distribution will not be affected no matter at what frequency the impedances are evaluated.

Clearly, the accuracy of the approximation improves as more paths are considered. Experimental results show that it is sufficient to consider only the shortest paths and the second shortest paths. Table I show some data on impedances calculated by using the shortest paths and second shortest paths approximation. The experiment is performed on a power mesh that has the exact structure as shown in Fig. 1. The impedance values are normalized to the impedance of a wire segment in the power grids. We calculated the impedances between the connection point of Module A [denoted as (a)] and its four neighboring VDD pins [labeled (1), (2), (3), (4)]. We did the same for the

TABLE I
IMPEDANCES CALCULATED USING SHORTEST AND SECOND SHORTEST
PATHS APPROXIMATION

Method	Z ₍₁₎ (a)	Z ₍₂₎ (a)	Z ₍₃₎ (a)	Z ₍₄₎ (a)	Z ₍₃₎ (c)	Z ₍₄₎ (c)	Z ₍₅₎ (c)	Z ₍₆₎ (c)
Shortest	1.25	1.50	1.00	1.25	1.00	1.76	2.00	1.96
Shortest+ 2nd Shortest	1.02	1.16	0.70	0.86	0.60	1.04	0.92	1.32
Hspice	1.00	1.10	0.68	0.84	0.54	1.02	0.90	1.28




Fig. 4. Current paths in power-supply mesh.

connection point of Module C [denoted as (c)]. Impedance estimates in the first row are obtained by considering the impedances of shortest paths only. Estimates in the second row are obtained by considering the impedances of both the shortest paths and the second shortest paths.

We compare the estimated impedances with the results obtained from SPICE simulation. As shown by the data, a good estimate of the impedance can be obtained by considering only the shortest paths and the second shortest paths. The error (compared with SPICE results) is less than 10%.

Once the component currents I_j ($j = 1, 2, \dots, N$) from the neighboring VDD sources are determined, we distribute I_j among the dominant paths from VDD source j to the sink, as illustrated in Fig. 4.

Let $\{P_1, P_2, \dots, P_w\}$ denote the set of the shortest paths and the second shortest paths under consideration. Let $Y_{P_1}, Y_{P_2}, \dots, Y_{P_w}$ be the admittance of these paths. By a similar derivation used in (2), the current I_j can be distributed among these paths, denoted by $i_{P_1}, i_{P_2}, \dots, i_{P_w}$, as follows:

$$\begin{aligned} i_{P_1} + i_{P_2} + \dots + i_{P_w} &= I_j \\ i_{P_k} &= \frac{Y_{P_k}}{\sum_{i=1}^w Y_{P_i}} I_j \\ k &= 1, 2, \dots, w. \end{aligned} \quad (3)$$

C. Noise Estimation

To estimate the power-supply noise that a circuit block experiences, we calculate the voltage variation at the connection point of the block in the power-supply grids, which is the voltage difference between the connection point and its neighboring power-supply pins [3]. Suppose P_k is a dominant current path between the connection point of circuit module k and the VDD pin closest to it. Let $T^{(k)} = \{P_j; P_j \cap P_k \neq \emptyset\}$ be a collection of the current paths in the power-supply mesh that overlap with

path P_k (including P_k itself). Let $P_{jk} = P_j \cap P_k$ denote the overlapping part between path P_j and path P_k , $R_{P_{jk}}$ denote the resistance of P_{jk} , and $L_{P_{jk}}$ denote the inductance of P_{jk} . Let $V_{\text{noise}}^{(k)}$ denote the power-supply noise at module k . $V_{\text{noise}}^{(k)}$ can be calculated using Kirchhoff's voltage law

$$V_{\text{noise}}^{(k)} = \sum_{P_j \in T^{(k)}} \left(i_j R_{P_{jk}} + L_{P_{jk}} \frac{di_j}{dt} \right) \quad (4)$$

where i_j is the current flowing along path P_j . One should note that not only the switching current of module k contributes to $V_{\text{noise}}^{(k)}$, other modules that draw current from the same VDD pins as module k contribute as well, as long as their current distribution paths overlap with P_k . Since there are potentially several paths leading to a module from a VDD pin, we choose the path of the worst current load to calculate the noise.

Since we only consider the neighboring VDD/GND pins and dominant current paths in the power mesh, the complexity of estimating the power-supply noise for a circuit module is bounded by a finite constant. The overall complexity for power-supply noise estimation is $O(n)$ for a floorplan of n modules. Since we ignore the effect of remote power pins and potential current paths of large impedances, we overestimate the currents flowing along the dominant paths and, consequently, the power-supply noise at each module. The power-supply noise is about 10%~15% percent overestimated compared with results from SPICE, which leads to a slightly conservative design.

Remark 1: The power-supply noise has a negative feedback effect since a reduced effective power-supply voltage leads to reduced drive currents in the switching gates. The power-supply noise calculated in this section should be interpreted as the “power-supply noise” experienced by the modules if all circuits (devices) are driving at the designed capacity. Since the goal of placing decap is to maintain the designed performance, we should consider all modules driving at the noise-free capacity when we calculate the “power-supply noise” as an indicator for decap budgeting. In other words, we calculate the “power-supply noise” under feedback-free scenarios.

IV. DECOUPLING-CAPACITANCE BUDGET

In this section, we estimate the decap budget for each circuit module in the floorplan based on: 1) the power-supply noise the module experiences and 2) the upper limit of the power-supply noise, denoted as $V_{\text{noise}}^{(\text{lim})}$, that the circuit can tolerate.

A. Decoupling-Capacitance Estimation

Suppose there are M modules in the floorplan and the switching current of module k is $I^{(k)}$, $k = 1, 2, \dots, M$. Let $C^{(k)}$ be the decap required for circuit module k . Let $Q^{(k)}$ be the total charge that module k will draw from the power-supply network during the worst case switching process. $Q^{(k)}$ is given by the following equation:

$$Q^{(k)} = \int_0^\tau I^{(k)}(t) dt$$

where τ is the duration that the switching process lasts. The upper limit of $C^{(k)}$ is $Q^{(k)} / V_{\text{noise}}^{(\text{lim})}$, which assumes that $C^{(k)}$ can provide all of the switching current of module k if the voltage

across it is lowered from Vdd to $\{Vdd - V_{noise}^{(lim)}\}$. The decoupling effect will diminish when $C^{(k)}$ is increased beyond the limit.

A greedy budget scheme is $C^{(k)} = Q^{(k)} / V_{noise}^{(lim)}$, $k = 1, 2, \dots, M$. This scheme is suboptimal in the sense that it will result in a larger decap budget than required. We refer to a solution produced by this scheme as “greedy solution.” Although the greedy solution method is not optimal, it is commonly used in practice and cited in research literatures [1], [11].

In this paper, we take a different approach to compute $C^{(k)}$. The decap required for each circuit module can be initially estimated as follows:

$$\theta = \max \left(1, \frac{V_{noise}^{(k)}}{V_{noise}^{(lim)}} \right)$$

$$C^{(k)} = \frac{(1 - \frac{1}{\theta})Q^{(k)}}{V_{noise}^{(lim)}}, \quad k = 1, 2, \dots, M. \quad (5)$$

Suppose the estimated power-supply noise (before considering decap) of module k is θ times the tolerable noise limit $V_{noise}^{(lim)}$. In order to reduce the power-supply at module k to $V_{noise}^{(lim)}$, we need to scale the noise at module k by a factor of θ , which is achievable if we scale down all the currents that contribute to $V_{noise}^{(k)}$ by a factor of θ according to (4). The current flowing through the network can be reduced to $1/\theta$ of its value by adding enough decap to buffer $(1 - 1/\theta)$ portion of the current load. Since the decap at module k is only responsible for providing the switching current of module k , the decap $C^{(k)}$ should be such that when its voltage is lowered from Vdd to $\{Vdd - V_{noise}^{(lim)}\}$, it will release $(1 - 1/\theta)Q^{(k)}$ amount of charge to supply the demand of module k during the switching process, which leads to $C^{(k)}V_{noise}^{(lim)} = (1 - 1/\theta)Q^{(k)}$. When $V_{noise}^{(k)} \leq V_{noise}^{(lim)}$, no decap is required.

When $C^{(k)}$ is added to module k , we update the power-supply noises at module k and all the modules that draw currents from the same VDD pins as module k according to (4). Since the switching current at module k also contributes to the power-supply noise at those modules, when the current drawn by module k is reduced due to decoupling effect of $C^{(k)}$, the noise at those affected modules will also be relieved to some extent as dictated by (4). Due to the contributions by the switching currents of the neighboring modules, the updated $V_{noise}^{(k)}$ may still be above $V_{noise}^{(lim)}$ after adding decap $C^{(k)}$. However, $V_{noise}^{(k)}$ will be further relieved as we add decap to the neighboring modules.

After the initial decap budgets are calculated for all the modules in the floorplan, we verify the updated power-supply noise at each module to make sure it is indeed below $V_{noise}^{(lim)}$. If $V_{noise}^{(k)}$ is still above $V_{noise}^{(lim)}$ for some module k , we will increase $C^{(k)}$ by an adequate amount (without exceeding its upper limit) such that $V_{noise}^{(k)}$ goes below $V_{noise}^{(lim)}$. If $C^{(k)}$ is increased to the limit and $V_{noise}^{(k)}$ is still above the limit, we need to increase the decap of its neighboring modules until $V_{noise}^{(k)}$ goes below $V_{noise}^{(lim)}$. This process is guaranteed to converge (since the greedy solution is the worst case solution of this approach) and it usually converges in two to three runs.

Decoupling Capacitance (decap) Budget()

Input: Floorplan with placement information, power supply noise of all circuit modules.

Sort circuit modules according to power supply noise;

For each module in the sorted list—starting with the module with the worst noise — **do**

Calculate its decap budget using Eqn. (5);

Update power supply noise of the modules affected due to the added decap using Eqn. (4);

For each module in the sorted list (after initial run) **do**

Check to see if its power supply noise is below $V_{noise}^{(lim)}$;

If power supply is not below $V_{noise}^{(lim)}$ **then**

Increase its decap until noise goes below limit or the decap reaches its limit;

If the power supply noise is still above $V_{noise}^{(lim)}$ **then**

Increase the decap of neighboring modules until noise goes below limit;

Output: Decoupling capacitance budget for each module.

Fig. 5. Iterative procedure for decap budget calculation.

The decap budgets generated with our “iterative” procedure can be significantly smaller than the greedy solution (please refer to Table III). The procedure for decap budgets calculation is summarized in Fig. 5.

Remark 2: The added onchip capacitance may change the resonance condition of the chip. If the clock frequency (or its harmonics) coincides with the resonance frequency of the chip, a large voltage fluctuation can build up in the power-supply network and cause circuit failure. Simulation must be performed to identify the potential resonance frequencies [1] and the power-supply network may need to be redesigned to prevent resonance.

V. WHITE-SPACE ALLOCATION FOR DECOUPLING CAPACITANCES

The decap budget for each circuit module is converted to the area of silicon required to fabricate the decap as follows:

$$S^{(k)} = \frac{C^{(k)}}{C_{ox}}, \quad k = 1, 2, \dots, M \quad (6)$$

where $S^{(k)}$ is the WS required to fabricate $C^{(k)}$.

Decaps need to be placed in the close neighborhood of switching activities to effectively relieve the power-supply noise. Decaps located far from the noisy spot are not effective due to the longer resistance–capacitance (RC) delay time and the IR drop [2]. The total area required for decap fabrication, denoted as S_{decap} , is given as follows:

$$S_{decap} = \sum_{k=1}^M S^{(k)}.$$

The decap’s allocation problem really boils down to WS allocation in the existing floorplan. Due to timing and routing constraints, it is best not to make dramatic changes to the given floorplan. Decap allocation can be done as a postfloorplan refinement to the existing floorplan in an incremental manner [16]. There are two issues in decap allocation. First, we must

allocate $S^{(k)}$, $k = 1, 2, \dots, M$ amount of WS to module k . Second, the amount of WS $S^{(k)}$ must be in the vicinity of module k in order for the decap to be effective. The WS allocation are carried out using a two-step approach as follows.

A. Allocation of Existing White Space

The isolated WSs in the original floorplan are treated as WS modules and can be used for decap fabrication. Since decap (or equivalent WS) must be placed close to the target circuit module, WS modules located far from a circuit module are considered inaccessible. When we allocate an existing WS module to its neighboring circuit modules, it is possible that after the WS demands of all its neighboring circuit modules has been met, there is still some WS left and the remaining WS is not neighboring to any circuit blocks and therefore considered as inaccessible WS. We must allocate the existing WS judiciously such that the inaccessible WS is minimized.

The problem can be solved using the LP technique. Suppose there are H isolated WS modules with area A_k , $k = 1, 2, \dots, H$, in the existing floorplan.

Let $N_k = \{j: \text{module } j \text{ is adjacent to WS module } k\}$ $k = 1, 2, \dots, H$ denote a set of circuit modules neighboring WS module k . Let $x_k^{(j)}$ be the amount of WS allocated to circuit module j from WS module k . The WS allocation problem can be formulated as follows:

$$\begin{aligned} \text{maximize } S &= \sum_{k=1}^H \sum_{j \in N_k} x_k^{(j)} \\ \text{subject to } \sum_{j \in N_k} x_k^{(j)} &\leq A_k, \quad k = 1, 2, \dots, H \\ \sum_{k=1}^H x_k^{(j)} &\leq S^{(j)}, \quad j = 1, 2, \dots, M \\ x_k^{(j)} &\geq 0, \quad \forall k, \forall j \end{aligned} \quad (7)$$

where S is the total WS allocated. The first set of constraints guarantee that the total WS allocated from a WS module k is less than or equal to its area A_k . The second set of constraints guarantee that the WS allocated to a circuit module j is less than or equal to its WS demand $S^{(j)}$ because there is no need to oversupply its WS demand. The third set of constraints guarantee that all the allocations are positive.

After we solve the LP problem, we know exactly how the existing WS modules are allocated to the circuit modules and how much WS is inaccessible. We compute the updated WS demand $\tilde{S}^{(j)}$, $j = 1, 2, \dots, M$ for all circuit modules after the WS allocation as follows:

$$\tilde{S}^{(j)} = S^{(j)} - \sum_{k=1}^H x_k^{(j)}, \quad j = 1, 2, \dots, M. \quad (8)$$

The additional amount of WS δA that needs to be inserted into the floorplan is determined as

$$\delta A = \sum_{j=1}^M \tilde{S}^{(j)} = \sum_{j=1}^M S^{(j)} - S. \quad (9)$$

If $\delta A = 0$, allocation process is complete; otherwise, we need to insert δA into the floorplan such that the WS can be used for




Fig. 6. (a) Original floorplan. (b) Moving modules in y direction in the order $\{(A, B), (C, D), (E, F), (G)\}$ to make WS for decap.

decap allocation. The δA is the area penalty in the cost function associated with power-supply noise discussed in Section VI.

Remark 3: The effectiveness of the decap diminish gradually as the distance increases. We restrict the placement of decaps to the neighborhood of the target modules to make our presentation easier. A metric η ($0 < \eta < 1$) can be defined for the effectiveness of the decap. For example, one feasible definition can be $\eta(d) = V_N(0)/V_N(d)$, where $V_N(0)$ is the power-supply noise when the decap is placed in the neighborhood and $V_N(d)$ the power-supply noise when the decap is placed at a distance d . The effectiveness of a decap placed in the neighborhood is one (normalized) and the effectiveness of a decap (same size) placed at a distance d is $\eta(d)$. Depending on the switching speed of the current waveform, the $\eta(d)$ can be characterized with respect to distance d based on simulations. With the help of η , a decap of size C at a distance d from a target module can be allocated to the module as a decap with an effective size of $\eta(d)C$. A general placement scheme without the neighborhood restriction can be obtained by directly extending the placement scheme employed in this paper.

B. Insertion of Additional White Space Into Floorplan

We use a heuristic to insert δA into the floorplan. The WS is inserted by extending the floorplan dimensions in both x direction and y direction. Suppose α portion of the additional WS δA is obtained by extending the floorplan in y direction and $(1 - \alpha)$ portion of δA is obtained by extending the floorplan in x direction. Let $\text{Layout}X$ and $\text{Layout}Y$ be the width and height of the original floorplan. The extensions of the floorplan in x direction and y direction, denoted by $\text{Ext}X$ and $\text{Ext}Y$, are given as follows:

$$\text{Ext}Y = \frac{\alpha \delta A}{\text{Layout}X}; \quad \text{Ext}X = \frac{(1 - \alpha) \delta A}{(\text{Layout}Y + \text{Ext}Y)}.$$

The heuristic works as follows. The modules in the floorplan are arranged into rows according to their levels in the constraint graph [2], [17] with the level of the source node in the graph set to zero. First, we move the circuit modules in y direction row by row. We move the modules in the top row by $\text{Ext}Y$, then the rows below it will be moved subsequently as illustrated in Fig. 6. We insert WS bands between the rows by shifting the adjacent rows by different amounts in y direction. The width of the WS

band is determined by the WS demand of the circuit modules in the previous row. The width of the WS band inserted between row $j - 1$ and row j , denoted by $B_{WS}^{(j-1)}$, is given as follows:

$$B_{WS}^{(j-1)} = \frac{\sum_{i \in \text{row } (j-1)} \alpha \tilde{S}^{(i)}}{\text{Layout}X}.$$

The inserted WS band provides α portion of the WS demanded by the circuit modules in row $j - 1$.

Similarly, WS bands are inserted between columns by moving the modules in x direction

$$B_{WS}^{(k-1)} = \frac{\sum_{i \in \text{column } (k-1)} (1 - \alpha) \tilde{S}^{(i)}}{\text{Layout}Y + \text{Ext}Y}.$$

Our heuristic inserts the additional WS required into the existing floorplan in an incremental manner. Since modules in the same row (or column) are shifted by the same amount of distance in our algorithm, our heuristic preserves the topology of the original floorplan.

Remark 4: The decap budgets may be slightly changed when inserting additional WS into the floorplan since module positions are changed. However, the additional WS inserted is no larger than 8.1% of the chip area from the experimental results and the additional WS is inserted between the rows and columns of modules by extending the original floorplan both horizontally and vertically. The dimensions of the floorplan are increased by less than 4% in both directions. The relative change of the module positions is about 4% since the increase is distributed between the rows and columns of modules. The current distribution and consequently the noise and decap budgets is only changed slightly. In the worst case, the modification can be taken care of by iteration and the extension is straightforward.

VI. APPLICATION TO POWER-SUPPLY NOISE-AWARE FLOORPLANNING

In conventional floorplanning, area and wire length are the main objectives and the optimality of a floorplan is measured based on the following cost function, which is a weighted sum of the chip area and total wire length

$$\Psi = A + \lambda W$$

where A is the total area, W is the total wire length, and λ is the weight parameter. The decap deployment is considered as an afterthought and addressed in a postfloorplanning step as illustrated in previous sections.

In this section, we directly address the issue of power-supply noise in the floor planning process so that the peak power-supply noise and, therefore, the decap budget, is minimized. Fig. 7 illustrates the rationale for noise-aware floorplanning methodology. The circuit blocks in dark grey are highly active modules, while the blocks in light grey are lightly active modules. The floorplan in Fig. 7(a) is extremely unbalanced and the power pin 1 is overloaded compared to other power pins. As a result, the spot around power pin 1 is very noisy and, therefore, requires a large decap to relieve the noise. On the other hand, the floorplan in Fig. 7(b) is more balanced as the highly active modules are scattered across the floorplan. Consequently, the peak power-supply noise is reduced and so is the decap. While the two floorplans have the same area and may look equally good in the conventional floorplan-




Fig. 7. Correlation between power-supply noise and floorplanning. A rationale for noise-aware floor planning. (a) Unbalanced. (b) Balanced.

ning, it does make a difference in the noise-aware floorplanning and floorplan (b) will be chosen over floorplan (a). In typical high-performance VLSI circuits, the switching activities are quite different for different circuit modules.

It is very important to take the variations of switching activities into consideration during the floorplanning process.

To consider the power-supply noise during floorplanning process, we redefine the cost function Ψ as follows:

$$\Psi = A + \lambda_1 W + \lambda_2 V_N \quad (10)$$

where V_N is the cost associated with the power-supply noise and λ_1 and λ_2 are the weight parameters used in the cost function for balancing the three factors.

The power-supply noise must be suppressed below a given specified limit by placing decap in the vicinity of each module. Decaps that cannot be deployed in the existing WS incurs area and wire length penalties δA and δW , respectively. The cost associated with power-supply noise V_N can be converted to the area penalty δA and the wire length penalty δW . The cost function Ψ can be rewritten accordingly as follows:

$$\Psi = (A + \delta A) + \lambda_1 (W + \delta W). \quad (11)$$

Due to the many research efforts directed toward floorplanning, several significant advancements in floorplan representation, namely, slicing tree [18], sequence pair [19], BSG [20], O-tree [21], and B^* -tree [22], have been made. For the implementation of the proposed noise-aware floorplanning methodology, we use sequence pair to represent the floorplan. We use longest common subsequence (LCS) computation, an efficient algorithm of complexity $O(n \log \log n)$ proposed in [23], for fast sequence-pair evaluation.

The proposed power-supply noise-aware floorplanning methodology is implemented using a simulated annealing technique [18]. The simulated annealing procedure is detailed in Fig. 8. Current floorplan is perturbed by performing one of the legal movement operations defined in [19], such as switching the order of two modules in the sequence pair or rotating a module by 90° . The merit of the perturbed floorplan is evaluated according to the cost function given in (11). If the perturbed floorplan has a smaller cost, the movement is accepted. Otherwise, the perturbed floorplan is accepted with a probability of $e^{-\Delta\Psi/T}$. The temperature scheduling ratio r is a constant ($0 < r < 1$). Simulated annealing will continue until the frozen temperature T_{FROZEN} is reached.

Algorithm Simulated Annealing
 Initial floorplan; $T_0 = INIT_T$
while $T_k > T_{Frozen}$
 Perturb_Current_Floorplan()
 Estimate_Power_Supply_Noise()
 if $T_k > T_{LOW}$
 Evaluate_Cost_Function $T_{HIGH}()$
 else
 Evaluate_Cost_Function $T_{LOW}()$
 $\Delta\Psi = \Psi_{new} - \Psi_{old}$
 if $\Delta\Psi < 0$
 Accept the perturbed floorplan
 $\Psi_{old} = \Psi_{new}$
 else
 Accept the floorplan with probability $e^{-\Delta\Psi/T_k}$
 $k = k + 1$; $T_k = rT_{k-1}$
END Simulated Annealing

Fig. 8. Simulated-annealing algorithm for power-supply noise-aware floorplanning.

The wire length for a net is calculated as half the perimeter of the bounding box. Since the LCS algorithm calculates the module positions as the sequence pair is evaluated, the area A of the floorplan and the total wire length W are easy to calculate.

The difficult part of the cost function evaluation is to determine the cost associated with power-supply noise: the area penalty δA and the wire penalty δW . The exact δA and δW can be determined only when the existing WS in the floorplan is allocated with a LP technique. Although the LP formulation is efficient as a postfloorplanning step, it is too computationally expensive to solve for every run in the simulated annealing process. To resolve this, LP is solved only at low temperature to determine the exact δA , as illustrated in Sections IV and V, and δW as follows. Since the modules are pushed further apart after the additional WS insertion, the total wire length should be recalculated to determine the wire length penalty δW . The positions of the modules are updated after the additional WS insertion. New wire lengths are calculated based on the updated module positions. The wire length penalty is given by

$$\delta W = W_{updated} - W_{old}.$$

Function *Evaluate_Cost_Function* $T_{LOW}()$ evaluates the cost function of each intermediate floorplan at low simulated temperature following exactly the procedures outlined above.

At high temperature, we estimate δA and δW as follows. Suppose the total area required for decap fabrication is S_{decap} . Let S_{exist} denote the existing WS in the floorplan. We assume that $\gamma(0 < \gamma \leq 1)$ portion of S_{exist} is accessible for decap fabrication, then the additional WS that needs to be added to the floorplan is given by:

$$\delta A = \max(0, S_{decap} - \gamma S_{exist}).$$

δA is the area penalty due to power-supply noise (or decap) in the cost function. If δA is 0, there is no penalty to wire length; Otherwise, the additional δA WS is inserted into the floorplan as WS bands between the levels of circuit modules as illustrated in Fig. 6. Since we do not know exactly how the existing WS is al-

TABLE II
TECHNOLOGY PARAMETERS

Parameters	Description	Value
r	wire resistance per unit length ($\Omega/\mu m$)	0.0125
l	wire inductance per unit length ($pH/\mu m$)	0.4
c	wire capacitance per unit length ($fF/\mu m$)	20
L_P	package inductance per VDD pin (nH)	0.2
R_P	package resistance per VDD pin (Ω)	0.5

TABLE III
COMPARISON OF DECAP BUDGETS—ITERATIVE VERSUS GREEDY SOLUTION

Circuit	decap budget ("Iterative") (nF)	decap budget ("Greedy Solution") (nF)	percentage (%)
apte	20.72	23.92	86.6
xerox	6.74	10.41	64.7
hp	4.45	6.50	68.5
ami33	0.085	0.6	14.2
ami49	9.34	20.68	45.2

located to the modules, we assume the additional WS δA is distributed evenly between rows of modules in the floorplan. Then the width of the WS band, denoted by B_{WS} , can be calculated based the total module levels, denoted by d , the dimensions of the floorplan and δA

$$B_{WS} = \frac{\delta A^*}{d \cdot \text{Layout}X}$$

where $\text{Layout}X$ is the width of the floorplan. Module positions are updated after WS insertion. Wire length is recalculated. The change of the wire length is the wire length penalty.

The *Evaluate_Cost_Function* $T_{HIGH}()$ function performs the cost function evaluation at high temperature, as illustrated above.

VII. EXPERIMENTAL RESULTS

The proposed power-supply noise-aware floor planning methodology is implemented in C. The LP part of the algorithm is solved using Matlab by invoking a system call to Matlab in our C program. Experiments are performed on five MCNC benchmark circuits implemented in a 0.25- μm technology. The pitch for the metal lines in the power-supply mesh is 333.3 μm and the pitch for VDD pins is 1000 μm . The power supply V_{dd} is 2.5 V. The parameters such as unit length parasitics of the metal grids in the power-supply network are provided by a leading semiconductor company. The technology parameters are listed in Table II. The worst case switching current profiles for the circuit modules are generated as follows. The worst case current density j_s is estimated for 0.25- μm technology based on the technology parameters, such as integration density and transistor channel length, obtained from [24]. The peak switching current for a circuit module k is $I^{(k)} = \text{factor}[k] \cdot j_s A_k$, where A_k is the area of module k and $\text{factor}[k]$ is either one or two depending on the random number generated. If the $\text{factor}[k]$ is two, the module k is a highly active module; otherwise, module k is a low-activity module. The overall switching current waveform of module k is approximated with a triangular waveform with peak value $I^{(k)}$ and duration of the switching current waveform (τ) is assumed to be half the clock cycle. Our method, however, is not limited to the

TABLE IV
EXPERIMENTAL RESULTS FOR MCNC BENCHMARK CIRCUITS FROM POSTFLOORPLANNING

Circuit	Modules	Existing WS (μm^2) (%)	decap Budget (nF)	Inaccessible WS (μm^2) (%)	Added WS (μm^2) (%)	Peak Noise (V)(before)	Peak Noise (V)(after)
<i>apte</i>	9	363220 (0.8)	20.72	0 (0)	3780910 (8.1)	2.05	0.24
<i>xerox</i>	10	1288356 (6.2)	6.74	152872 (0.7)	211753 (1.0)	1.63	0.23
<i>hp</i>	11	1525272 (14.7)	4.45	1085162 (10.5)	449892 (4.3)	1.61	0.21
<i>ami33</i>	33	97510 (7.8)	0.085	N/A	0	0.38	0.19
<i>ami49</i>	49	1452752 (3.9)	9.34	448723 (1.2)	867844 (2.4)	1.70	0.20

TABLE V
COMPARISON OF EXPERIMENTAL RESULTS—NOISE-AWARE VERSUS POSTFLOORPLAN

Circuit	peak noise (post) (V)	peak noise (noise-aware) (V)	percentage improved (%)	decap (post) (nF)	decap (noise-aware) (nF)	percentage improved (%)	time (post) (s)	time (noise-aware) (s)
<i>apte</i>	2.05	1.23	40.0	20.72	16.36	21.0	12	119
<i>xerox</i>	1.63	1.18	27.7	6.74	5.85	13.2	18	193
<i>hp</i>	1.61	1.42	11.8	4.45	4.12	7.4	16	215
<i>ami33</i>	0.38	0.35	7.9	0.085	0.084	1.2	45	956
<i>ami49</i>	1.7	1.45	14.7	9.34	8.00	14.3	57	1582

TABLE VI
EXPERIMENTAL RESULTS FOR MCNC BENCHMARK CIRCUITS ($\lambda = 1$)

Circuit	Modules	area (noise-aware) (μm^2)	area (post) (μm^2)	wire length (noise-aware) (μm)	wire length (post) (μm)	A+W (traditional) (no decap)	A+W (noise-aware)	A+W (post)
<i>apte</i>	9	50235794	50705710	595920	830445	46250577	50831714	51536155
<i>xerox</i>	10	20581079	20850453	545615	556625	21193275	21126694	21407078
<i>hp</i>	11	10559300	10876803	186952	184126	1060695	10746252	11060929
<i>ami33</i>	33	1253960	1254350	85884	87359	1341709	1339844	1341709
<i>ami49</i>	49	37548000	37766000	1125750	1277880	38173024	38673750	39043880

triangular waveform assumption and more sophisticated piecewise linear waveforms can be used to represent the switching current waveforms of the circuit modules. In our experiments, j_s is set to $0.2 \mu A/\mu m^2$ and τ is set to 1 ns. The power-supply noise limit $V_{noise}^{(lim)}$ is set to be 0.25 V.

The typical value of γ ranges from 0.3 to 0.8. The solution quality is sensitive to the value of γ , but there is no general trend for all circuits. In the experiments, the γ value is adjusted around 0.5. The λ is set to 1.


To compare our iterative method with the greedy solution method (see Section IV), the decap budgets obtained with the two methods are listed in Table III. It is clear that our method generates significantly smaller decap budgets. For circuit *ami49*, the decap generated with our method is only 45.2% of that generated using greedy solution method. The average reduction of the decap budget is 27.5% for the five benchmark circuits.

The floorplans of the benchmark circuits used for the post-floorplanning decap placement are generated with conventional floorplanning algorithm (similar to noise-aware floorplanning except that no noise associated cost is included in the cost function). Power-supply noise and decap placement are considered in a postfloorplan process. The experimental results obtained from postfloorplan decap placement are presented in Table IV. The total decap budgets for the benchmark circuits vary significantly depending on the size of the modules and the dimensions of the floorplan. Benchmark *ami33* has small chip area and small circuit modules and its decap budget is only 0.085 nF.


Large circuits such as *hp* and *apte* suffer serious power-supply noise and require considerable amount of WS for decap fabrication. The WS used for decap is about 9% of its chip area for *apte* and about 8.5% of its chip area for *hp*. Data on the existing WS in the original floorplan, the inaccessible WS, and the added WS are also collected for each benchmark circuit as shown in Table IV. The percentage in the parentheses is the percentage of the total chip area for WS. To determine the effectiveness of decap placement, the peak-noise data before and after decap's deployment are collected and compared. It is evident that peak noise is indeed suppressed to within the noise limit of 0.25 V.

The peak noise in Table IV is calculated with our algorithm. Simulations have been run for *apte* and *ami49* with HSPICE to validate the calculations. The peak noise before decap placement is 1.81 V for *apte*, 1.54 V for *ami49*. The peak noise after decap placement is 0.21 V for *apte* and 0.17 V for *ami49*. The results are close to our calculations. The overestimation of the peak noise is due to the conservative approximations we have made in our model.

The experimental results from noise-aware floorplanning are presented in Tables V and VI in comparison with the results from postfloorplanning approach. Compared to conventional floorplanning, the peak power-supply noise and the total decap are reduced for all the five circuits as shown in Table V. For *apte* and *xerox*, the peak power-supply noise is reduced by 40% and 27.7% and the decap is reduced by 21.0% and 13.2%, respectively. On average, the peak power-supply noise is reduced by



(a)



(b)

Fig. 9. Floorplans of benchmark circuit *xerox*. (a) Prior and (b) post decap placement.

20.4% and the decap budget is reduced by 11.5%. The reason that the decap is not reduced as much as the peak power-supply noise is that while the noise-aware floorplanning approach can reduce the peak power-supply noise by scattering highly active modules across the floorplan, it does, in the meantime, increase the power-supply noise at other quiet spots. The overall decap is reduced and the distribution of power-supply noise becomes

more even across the floorplan. For circuit *ami33*, the floorplan generated with conventional floorplanning is very close to the floorplan generated with power-supply noise-aware floorplanning. There is not much room for improvement for both peak supply noise and the decap.

The peak power-supply noise presented in Table V is the peak power-supply noise calculated before decap placement.




Fig. 9. (Continued) Floorplans of benchmark circuit *xerox*. (c) Noise aware.

The peak-noise reduction (compared to postfloorplanning) for the noise-aware floorplanning comes from the optimized configuration of the circuit modules based on the switching profiles.

The area of the final floorplans (after decap placement) of the five benchmark circuits are shown in Table VI. The floorplans produced by power-supply noise-aware floorplanning algorithm have smaller area than the corresponding floorplans from post-floorplanning. The area reduction for circuit *hp* is about 2.9% of its floorplan area. The area save for circuits *apte* and *xerox* is 0.93% and 1.3% of its floorplan area, respectively. The average area reduction of the benchmark circuits is 1.2%.

As for the wire length, most of the benchmark circuits have improved total wire length due to the reduced decap gained from noise-aware floor planning. The total wire length for *hp* is, however, increased. This is due to the fact that the gain from decap outweighs the loss to wire length and the overall cost of the floorplan is improved. For comparison purpose, the sums of the area and wire length are also listed in Table VI. The cost of the traditional floorplanning without decap placement is also presented as a baseline for comparison.

Fig. 9 shows the floorplans of circuit *xerox*. The floorplan in Fig. 9(a) is obtained from conventional floorplan and the floorplan in Fig. 9(b) is Fig. 9(a) after decap placement. Since most of the decap comes from the existing WS in Fig. 9(a), only a small amount of additional WS ($211\,753\,\mu\text{m}^2 \sim 1\%$ of the floorplan area) is inserted and changes made to Fig. 9(a) are small and incremental. Fig. 9(c) is produced by noise-aware floorplanning and it has lower peak power-supply noise and less area than Fig. 9(b). Note that the highly active blocks (in dark) are placed differently in Fig. 9(b) and Fig. 9(c). Fig. 9(c) is $269\,374\,\mu\text{m}^2$ ($\sim 1.3\%$ of its floorplan area) smaller than Fig. 9(b) and its total wire length is also reduced by 2.0%. All these gain come

from the optimization of floorplan configuration. Postfloorplan decap placement methodology is, however, easy to implement and modifies the original floorplan incrementally. Circuit designers may benefit from both the postfloorplan decap placement methodology and the noise-aware floorplanning methodology depending on the targeted designs.

VIII. CONCLUSION

The decap allocation problem is investigated at floorplan level. Two decap placement methodologies are proposed for two variants of the problem. The postfloorplan decap placement methodology is proposed to deploy decap for a given floorplan in an incremental manner as a postfloorplan refinement. The power-supply noise-aware floorplanning methodology, which incorporates power-supply noise into the cost function, is proposed to handle decap placement in the floorplanning process. The proposed power-supply noise-aware floorplanning algorithm can reduce the peak power-supply noise and, therefore, decap budget by judiciously arranging circuit modules based on their switching activities and spatial correlations.

Experimental results on MCNC benchmark circuits show that our iterative method also produces significantly smaller decap budgets than the greedy solution method commonly used in practice and research. The algorithm implemented for postfloorplan decap placement modifies the floorplan incrementally without dramatically changing the topology of the original floorplan. Power-supply noise-aware floorplanning methodology can further reduce the peak power-supply noise by as much as 40%. The total decap budget and the total area of the floorplan are also reduced due to the gain from reduced power-supply noise.

For future work, we will consider wire sizing and decap placement simultaneously for optimal P/G network design. Due to the wire length penalty (namely, timing penalty), the high-power modules may not always be pushed apart. In that case, grid sizing may become an option for power-supply noise suppression without incurring timing penalty.

ACKNOWLEDGMENT

The authors would like to thank Dr. M. Wong, Mr. X. Tang, and Mr. R. Tian at the University of Texas, Austin, for providing us the code of their LCS-based fast sequence-pair evaluation floorplanning program.

REFERENCES

- [1] H. B. Bakoglu, *Circuits, Interconnects and Packaging for VLSI*. MA: Addison-Wesley, 1990.
- [2] H. H. Chen and D. D. Ling, "Power supply noise analysis methodology for deep-submicron VLSI chip design," in *Proc. Design Automation Conf.*, June 1997, pp. 638–643.
- [3] S. Zhao, K. Roy, and C.-K. Koh, "Estimation of inductive and resistive switching noise on power supply network in deep sub-micron CMOS circuits," in *Proc. Int. Conf. Computer Design*, Sept. 2000, pp. 65–72.
- [4] M. Zhao, R. Panda, S. Sapatnekar, T. Edwards, R. Chaudhry, and D. Blaauw, "Hierarchical analysis of power distribution networks," in *Proc. Design Automation Conf.*, June 2000, pp. 150–155.
- [5] J. Oh and M. Pedram, "Multi-pad power/ground network design for uniform distribution of ground bounce," in *Proc. Design Automation Conference*, June 1998.
- [6] H. Su, K. Gala, and S. Sapatnekar, "Fast analysis and optimization of power/ground networks," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 2000, pp. 477–480.
- [7] J. M. Wang and T. Nguyen, "Extended Krylov subspace method for reduced order analysis of linear circuits with multiple sources," in *Proc. Design Automation Conf.*, June 2000, pp. 247–252.
- [8] K.-H. Erhard, F. M. Johannes, and R. Dachauer, "Topology optimization techniques for power/ground networks in VLSI," in *Proc. Eur. Design Automation Conf.*, Mar. 1992, pp. 362–367.
- [9] R. Dutta and M. M. Sadowska, "Automatic sizing of power/ground networks in VLSI," in *Proc. Design Automation Conf.*, June 1989.
- [10] M. Ang, R. Salem, and A. Taylor, "An on-chip voltage regulator using switched decoupling capacitors," *Proc. Int. Solid-State Circuits Conf. Dig. Tech. Papers*, pp. 438–439, Feb. 2000.
- [11] L. Smith, "Decoupling capacitor calculations for CMOS circuits," in *Proc. IEEE 3rd Topical Meeting of Electrical Performance of Electronic Packaging*, Nov. 1994, pp. 101–105.
- [12] G. Bai, S. Bobba, and I. N. Hajj, "Simulation and optimization of the power distribution network in VLSI circuits," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 2000, pp. 481–486.
- [13] K. Y. Chao and D. F. Wong, "Thermal placement for high performance multiple-chip modules," in *Proc. Int. Conf. Computer Design*, Oct. 1995, pp. 218–230.
- [14] C.-H. Tsai and S.-M. Kang, "Cell-level placement for improving substrate thermal distribution," *IEEE Trans. Computer-Aided Design*, vol. 19, pp. 253–266, Feb. 2000.
- [15] S. Mitra, R. A. Rotenbar, L. R. Carley, and D. J. Allstot, "Substrate-aware mixed-signal macrocell placement in WRIGHT," *IEEE J. Solid-State Circuits*, vol. 30, pp. 269–278, Mar. 1995.
- [16] J. Cong and M. Sarrafzadeh, "Incremental physical design," in *Proc. Int. Symp. Physical Design*, Apr. 2000, pp. 84–92.
- [17] R. Otten, "Graphics in floorplan design," *J. Circuit Theory Applicat.*, vol. 16, no. 4, pp. 391–410, Oct. 1988.
- [18] D. F. Wong and C. L. Liu, "A new algorithm for floorplan design," in *Proc. Design Automation Conf.*, June 1986, pp. 101–107.
- [19] H. Murata, K. Fujiyoshi, S. Nakatake, and Y. Kajitani, "VLSI module placement based on rectangle-packing by the sequence pair," *IEEE Trans. Computer-Aided Design*, vol. 15, pp. 1518–1524, Dec. 1996.
- [20] S. Nakatake, H. Murata, K. Fujiyoshi, and Y. Kajitani, "Module placement on BSG-structure and IC layout applications," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 1996, pp. 484–491.
- [21] P. N. Guo, C. K. Cheng, and T. Yoshimura, "An O-tree representation of nonslicing floorplans and its applications," in *Proc. Design Automation Conf.*, June 1999, pp. 268–273.
- [22] Y. C. Yang, Y. W. Chang, G. M. Wu, and S. W. Wu, " B^* -trees: A new representation for nonslicing floorplans," in *Proc. Design Automation Conf.*, June 2000, pp. 458–463.
- [23] X. Tang, R. Tian, and D. F. Wong, "Fast evaluation of sequence pair in block placement by longest common subsequence computation," in *Proc. Design Automation Test Eur.*, Mar. 2000, pp. 106–111.
- [24] *International Technology Roadmap for Semiconductor*, Semiconductor Industry Assoc., 1997.
- [25] B. Krauter and S. Mehrotra, "Layout based frequency dependent inductance and resistance extraction for on-chip interconnect timing analysis," in *Proc. Design Automation Conf.*, June 1998, pp. 303–308.
- [26] S. R. Nassif and J. N. Kozhaya, "Fast power grid simulation," in *Proc. Design Automation Conf.*, June 2000, pp. 156–161.



Shiyou Zhao received the M.S. degree in physics from Fudan University, Shanghai, China, in 1993 and the M.S.E.E. and Ph.D. degrees from the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, in 1998 and 2001, respectively.

He is currently with Micron Semiconductor, Boise, ID, where he is a Signal Integrity Engineer. His current research interests include power-supply noise analysis, interconnect modeling, signal integrity, and robust circuitry.



Kaushik Roy (S'83–M'90–SM'95–F'02) received the B.Tech. degree in electronics and electrical communications engineering from the Indian Institute of Technology, Kharagpur, India, in 1983 and the Ph.D. degree in electrical and computer engineering from the University of Illinois at Urbana-Champaign in 1990.

He was with the Semiconductor Process and Design Center of Texas Instruments, Dallas, where he worked on FPGA architecture development and low-power circuit design. He has been with the Electrical and Computer Engineering Faculty of Purdue University, West Lafayette, IN, since 1993, where he is currently a Professor and a University Faculty Scholar. His current research interests include VLSI/CAD design with particular emphasis on low-power electronics for portable computing and wireless communications, VLSI testing and verification, and reconfigurable computing.

Dr. Roy received the National Science Foundation Career Development Award in 1995, IBM Faculty Partnership Award, and Best Paper Award at the 1997 International Test Conference and the 2000 International Symposium on Quality of Integrated Circuit Design. He is on the editorial board of IEEE DESIGN AND TEST, IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS, and IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS. He was a Guest Editor for the IEEE DESIGN AND TEST SPECIAL ISSUE ON LOW-POWER VLSI in 1994 and IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS in June 2000.



Cheng-Kok Koh (S'92–M'98) received the B.S. and M.S. degrees in computer science from the National University of Singapore in 1992 and 1996, respectively, and the Ph.D. degree in computer science from the University of California, Los Angeles, in 1998.

He is currently an Assistant Professor of Electrical and Computer Engineering with Purdue University, West Lafayette, IN. His current research interests include physical design of high-performance low-power VLSI circuits, with an emphasis on VLSI interconnect layout optimization.

Dr. Koh received the Lim Soo Peng Book Prize for Best Computer Science Student from the National University of Singapore in 1990, the Tan Kah Kee Foundation Postgraduate Scholarship in 1993 and 1994, the GTE Fellowship and the Chorasfas Foundation Prize from the University of California at Los Angeles in 1995 and 1996, respectively, the ACM Special Interest Group on Design Automation Meritorious Service Award in 1998, the Chicago Alumni Award from Purdue University in 1999, and the National Science Foundation CAREER Award in 2000.