Efficient Placement of Distributed On-Chip Decoupling Capacitors in Nanoscale ICs

Mikhail Popovich, Eby G. Friedman
Department of Electrical and Computer Engineering
University of Rochester
Rochester, New York 14627
[nhlover, friedman]@ece.rochester.edu

Radu M. Secareanu, Olin L. Hartin
MMSTL
Freescale Semiconductor
Tempe, Arizona 85284
[r54143, lee.hartin]@freescale.com

Abstract— **Decoupling capacitors are widely used to reduce power supply noise. On-chip decoupling capacitors have traditionally been allocated into the white space available on the die based on an unsystematic or *ad hoc* approach. In this way, large decoupling capacitors are often placed at a significant distance from the current load, compromising the signal integrity of the system. This issue of power delivery cannot be alleviated by simply increasing the size of the on-chip decoupling capacitors. To be effective, the on-chip decoupling capacitors should be placed physically close to the current loads. The area occupied by the on-chip decoupling capacitor, however, is directly proportional to the magnitude of the capacitor. The minimum impedance between the on-chip decoupling capacitor and the current load is therefore fundamentally affected by the magnitude of the capacitor.**

A distributed on-chip decoupling capacitor network is proposed in this paper. A system of distributed on-chip decoupling capacitors is shown to provide an efficient solution for providing the required on-chip decoupling capacitance under existing technology constraints. In a system of distributed on-chip decoupling capacitors, each capacitor is sized based on the parasitic impedance of the power distribution grid. Various tradeoffs in a system of distributed on-chip decoupling capacitors are also discussed. Related simulation results for typical values of on-chip parasitic resistance are also presented. An analytic solution is shown to provide accurate distributed system. The worst case error is 0.003% as compared to SPICE. Techniques presented in this paper are applicable not only for current technologies, but also provide an efficient placement of the on-chip decoupling capacitors in future technology generations.

Index Terms— **Power distribution systems, power distribution grids, decoupling capacitors, power noise**

I. INTRODUCTION

Decoupling capacitors are widely used to manage power supply noise [1]. Decoupling capacitors are an effective way to reduce the impedance of power delivery systems operating at high frequencies [2], [3]. A decoupling capacitor acts as a local reservoir of charge, which is released when the power supply voltage at a particular current load drops below some tolerable level. Since the inductance scales slowly [4], the location of the decoupling capacitors significantly affects the design of power/ground (P/G) networks in high performance integrated circuits (ICs) such as microprocessors. At higher frequencies, a distributed system of decoupling capacitors will need to be placed on-chip to effectively manage the power supply noise [5].

The efficacy of decoupling capacitors depends upon the impedance of the conductors connecting the capacitors to the current loads and power sources. As described in [6], a maximum parasitic impedance between the decoupling capacitor and the current load (*or* power source) exists at which the decoupling capacitor is effective. Alternatively, to be effective, an on-chip decoupling capacitor should be placed such that both the power supply and the current load are located inside the appropriate effective radius [6]. The efficient placement of on-chip decoupling capacitors in nanoscale ICs is the subject of this paper. Unlike the methodology for placing a single lumped on-chip decoupling capacitor presented in [6], a system of *distributed* on-chip decoupling capacitors is proposed. A design methodology to estimate the parameters of the distributed system of on-chip decoupling capacitors is also presented, permitting allocation of the required on-chip decoupling capacitance under existing technology constraints.

The paper is organized as follows. The problem of placing on-chip decoupling capacitors in nanoscale ICs under technology constraints is formulated in Section II. The design of an on-chip distributed decoupling capacitor network is presented in Section III. Various design tradeoffs are discussed in Section IV. Related simulation results for typical values of the on-chip parasitic resistance are discussed in Section V. Some specific conclusions are summarized in Section VI.

II. PLACING ON-CHIP DECOUPLING CAPACITORS IN NANOSCALE ICS

Decoupling capacitors have traditionally been allocated into the white space (those areas not occupied by the circuit elements) available on the die based on an unsystematic or *ad hoc* approach [7], [8], as shown in Fig. 1. In this way, decoupling capacitors are often placed at a significant distance from the current load. Conventional approaches for placing on-chip decoupling capacitors result in oversized capacitors. The

811

conventional allocation strategy, therefore, results in increased power noise, compromising the signal integrity of an entire system, as illustrated in Fig. 2. This issue of power delivery cannot be alleviated by simply increasing the size of the on-chip decoupling capacitors. Furthermore, increasing the size of more distant on-chip decoupling capacitors results in wasted area, increased power, reduced reliability, and higher costs. A new design methodology is therefore required to account for technology trends in nanoscale ICs, such as increasing frequencies, larger die sizes, higher current demands, and reduced noise margins.

White spaces

Fig. 1.   Placement of on-chip decoupling capacitors using a conventional approach. Decoupling capacitors are allocated into the white space (those areas not occupied by the circuits elements) available on the die using an unsystematic or *ad hoc* approach. As a result, the power supply voltage drops below the minimum tolerable level for remote blocks (shown in dark grey). Low noise regions are light grey.

To be effective, a decoupling capacitor should be placed physically close to the current load. This requirement is naturally satisfied in board and package applications, since large capacitors are much smaller than the dimensions of the circuit board (or package) [9]. In this case, a lumped model of the decoupling capacitor provides sufficient accuracy [10].


Fig. 2.   A conventional on-chip decoupling capacitor. Typically, a large decoupling capacitor is placed at a greater distance from the current load due to physical limitations. Current flowing through the long power/ground lines results in large voltage fluctuations across the terminals of the current load.

The size of an on-chip decoupling capacitor, however, is directly proportional to the area occupied by the capacitor and can require a significant portion of the on-chip area. The

minimum impedance between an on-chip capacitor and the current load is fundamentally affected by the magnitude (and therefore the area) of the capacitor. Systematically partitioning the decoupling capacitor solves this issue. A system of on-chip distributed decoupling capacitors is illustrated in Fig. 3.


Fig. 3.   A network of distributed on-chip decoupling capacitors. The magnitude of the decoupling capacitors is based on the impedance of the interconnect segment connecting a specific capacitor to a current load. Each decoupling capacitor is designed to only provide charge during a specific time interval.

In a system of distributed on-chip decoupling capacitors, each decoupling capacitor is sized based on the impedance of the interconnect segment connecting a capacitor to a current load. A particular capacitor only provides charge to a current load during a short period. The rationale behind the proposed scheme can be explained as follows. The capacitor closest to the current load is engaged immediately after the switching cycle is initiated. Once the first capacitor is depleted, the next capacitor is activated, providing a large portion of the total current drawn by the load. This procedure is repeated until the last capacitor becomes active. Similar to the hierarchical placement of decoupling capacitors presented in [3], [11], the proposed technique provides an efficient solution for providing the required on-chip decoupling capacitance based on specified capacitance density constraints. A system of distributed on-chip decoupling capacitors should therefore be utilized to provide a low impedance, cost effective power delivery network in nanoscale ICs.

III. DESIGN OF AN ON-CHIP DISTRIBUTED DECOUPLING CAPACITOR NETWORK

As described in Section II, a system of distributed on-chip decoupling capacitors is an efficient solution for providing the required on-chip decoupling capacitance based on the maximum capacitance density available in a particular technology. A physical model of the proposed technique is illustrated in Fig. 4. For simplicity, two decoupling capacitors are assumed to provide the required charge drawn by the current load. Note that as the capacitor is placed farther from the current load, the magnitude of an on-chip decoupling capacitor increases due to relaxed constraints. In the general case, the proposed methodology can be extended to any practical number of on-chip decoupling capacitors. Note that Z_1 is typically limited by a specific technology (determined by the impedance of a single metal wire) and the magnitude of C_1 (the area available in the vicinity of a circuit block).

A circuit model of the proposed system of distributed on-chip decoupling capacitors is shown in Fig. 5. The impedance of the metal lines connecting the capacitors to the current

$$V_{C_1}|_{t=t_r} = \frac{1}{2(C_1+C_2)^3 t_r}\left[2C_1^3 t_r + C_1^2 t_r\left(6C_2 - I_{max}t_r\right) - C_2^2 t_r\left(2C_2\left(I_{max}R_2 - 1\right)\right)\right. \tag{3}$$
$$\left. + \quad I_{max}t_r\right) + 2C_1C_2\left(C_2^2\left(1 - e^{-\frac{(C_1+C_2)t_r}{C_1C_2R_2}}\right)I_{max}R_2^2 + C_2\left(3 - I_{max}R_2\right)t_r - I_{max}t_r^2\right)\right],$$

$$V_{C_2}|_{t=t_r} = \frac{1}{2(C_1+C_2)^3 t_r}\left[2C_1^3 t_r + C_2^2 t_r\left(2C_2 - I_{max}t_r\right) + 2C_1C_2 t_r\left(C_2\left(3 + I_{max}R_2\right)\right)\right. \tag{4}$$
$$\left. - \quad I_{max}t_r\right) + C_1^2\left(2C_2^2\left(e^{-\frac{(C_1+C_2)t_r}{C_1C_2R_2}} - 1\right)I_{max}R_2^2 + 2C_2\left(3 + I_{max}R_2\right)t_r - I_{max}t_r^2\right)\right].$$


Fig. 4. A physical model of the proposed system of distributed on-chip decoupling capacitors. Two capacitors are assumed to provide the required charge drawn by the load. Z_1 and Z_2 denote the impedance of the metal lines connecting C_1 to the current load and C_2 to C_1, respectively.

load is modeled as resistors R_1 and R_2. Modeling the power and ground paths as simple resistors is a tradeoff, providing sufficient accuracy to estimate the required on-chip decoupling capacitance, while dramatically reducing the model complexity. A triangular current source is assumed to characterize the current load. The magnitude of the current source increases linearly, reaching the maximum current I_{max} at rise time t_r, i.e., $I_{load}(t) = I_{max}\frac{t}{t_r}$. The maximum tolerable ripple at the current load is 10% of the power supply voltage.



Fig. 5. A circuit model of an on-chip distributed decoupling capacitor network. The impedance of the metal lines is modeled as R_1 and R_2, respectively.

Note from Fig. 5 that since the charge drawn by the current load is provided by the on-chip decoupling capacitors, the voltage across the capacitors during discharge drops below the initial power supply voltage. The required charge during the entire switching event is thus determined by the voltage drop across C_1 and C_2.

The voltage across the decoupling capacitors at the end of the switching cycle ($t = t_r$) can be determined from Kirchhoff's laws [12]. Writing KVL and KCL for each of the loops (see Fig. 5), the system of differential equations for

the voltage across C_1 and C_2 at t_r is

$$\frac{dV_{C_1}}{dt} = \frac{V_{C_2} - V_{C_1}}{R_2C_1} - \frac{I_{load}}{C_1}, \tag{1}$$

$$\frac{dV_{C_2}}{dt} = \frac{V_{C_1} - V_{C_2}}{R_2C_2}. \tag{2}$$

Simultaneously solving (1) and (2) and applying the initial conditions, the voltage across C_1 and C_2 at the end of the switching activity is determined by (3) and (4), respectively. I_{max} is the maximum magnitude of the current load and t_r is the rise time.

Note that the voltage across C_1 and C_2 after discharge is determined by the magnitude of the decoupling capacitors and the parasitic resistance of the metal line(s) between the capacitors. The voltage across C_1 after the switching cycle, however, depends upon the resistance of the P/G paths connecting C_1 to a current load and is

$$V_{C_1} = V_{load} + I_{max}R_1, \tag{5}$$

where V_{load} is the voltage across the terminals of a current load. Assuming $V_{load} \geq 0.9V_{dd}$ and $V_{C_1}^{max} = V_{dd}$ (meaning that C_1 is infinitely large), the upper bound for R_1 is

$$R_1^{max} = \frac{V_{dd}(1-\alpha)}{I_{max}}, \tag{6}$$

where α is the ratio of the minimum tolerable voltage across the terminals of a current load to the power supply voltage ($\alpha = 0.9$ in this paper). If $R_1 > R_1^{max}$, no solution exists for providing sufficient charge drawn by the load. In this case, the circuit block should be partitioned, reducing the current demands (I_{max}).

Note that expressions for determining the voltage across the decoupling capacitors are transcendental functions. No closed-form solution, therefore, exists. From (3) and (4), the design space can be graphically obtained for determining the maximum tolerable resistance R_2 and the minimum magnitude of the capacitors, maintaining the voltage across the load equal to or greater than the minimum allowable level. The voltage across C_1 after discharge as a function of C_1 and R_2 is depicted in Fig. 6.

Observe from Fig. 6 that the voltage across capacitor C_1 increases exponentially with capacitance, saturating for large C_1. The voltage across C_1, however, is almost independent

(a)

(b)

Fig. 6. Voltage across C_1 during discharge as a function of C_1 and R_2: $I_{max} = 0.01$ mA, $V_{dd} = 1$ volt, and $t_r = 100$ ps. a) Assuming $C_1 = C_2$ and $R_1 = 10\,\Omega$, the minimum tolerable voltage across C_1, resulting in $V_{load} \geq 0.9V_{dd}$, is 0.91 volts (shown as a black equipotential line). b) The design space for determining C_1 and R_2 resulting in the voltage across C_1 equal to 0.91 volts.

of R_2, decreasing slightly with R_2 (see Fig. 6(a)). This behavior can be explained as follows. As a current load draws charge from the decoupling capacitors, the voltage across the capacitors drops below the initial level. The charge released by a capacitor is proportional to the capacitance and the voltage drop. A larger capacitance therefore results in a smaller voltage drop. From Fig. 5, note that as resistance R_2 increases, capacitor C_2 becomes less effective (a larger portion of the total current is provided by C_1). As a result, the magnitude of C_1 is increased to maintain the voltage across the load above the minimum tolerable level. Similarly, a larger C_2 results in a smaller C_1. As C_2 is increased, a larger portion of the total current is provided by C_2, reducing the magnitude of C_1. This phenomenon is well pronounced for small R_2, diminishing

with larger R_2, as illustrated in Fig. 6(b).

In general, to determine the parameters of the system of distributed on-chip decoupling capacitors, the following assumptions are made. The parasitic resistance of the metal line(s) connecting capacitor C_1 to the current load is known. R_1 is determined by technology constraints (the sheet resistance) and by design constraints (the maximum available metal resources). The minimum voltage level at the load is $V_{load} = 0.9V_{dd}$. The maximum magnitude of the current load I_{max} is 0.01 A, the rise time t_r is 100 ps, and the power supply voltage V_{dd} is one volt. Note that the voltage across C_2 after discharge as determined by (4) is also considered a design parameter. Since the capacitor C_2 is connected directly to the power supply (a shared power rail), the voltage drop across C_2 appears on the global power line, compromising the signal integrity of the overall system. The voltage across C_2 at t_r is therefore based on the maximum tolerable voltage fluctuations on the P/G line during discharge (the voltage across C_2 at the end of the switching cycle is set to 0.95 volts).

The system of equations to determine the parameters of an on-chip distributed decoupling capacitor network as depicted in Fig. 5 is

$$V_{load} = V_{C_1} - I_{max}R_1, \tag{7}$$
$$V_{C_1} = f(C_1, C_2, R_2), \tag{8}$$
$$V_{C_2} = f(C_1, C_2, R_2), \tag{9}$$
$$\frac{I_{max}t_r}{2} = C_1\left(V_{dd} - V_{C_1}\right) + C_2\left(V_{dd} - V_{C_2}\right), \tag{10}$$

where V_{C_1} and V_{C_2} are the voltage across C_1 and C_2 and determined by (3) and (4), respectively. Equation (10) states that the total charge drawn by the current load is provided by C_1 and C_2. Note that in the general case with the current load determined *a priori*, the total charge is the integral of $I_{load}(t)$ from zero to t_r. Solving (7) for V_{C_1} and substituting into (8), C_1, C_2, and R_2 are determined from (8) to (10), as discussed in the following section.

IV. DESIGN TRADEOFFS IN A DISTRIBUTED ON-CHIP DECOUPLING CAPACITOR NETWORK

To design a system of distributed on-chip decoupling capacitors, the parasitic resistances and capacitances should be determined based on design and technology constraints. As shown in Section III, in a system composed of two decoupling capacitors (see Fig. 5) with known R_1, R_2, C_1, and C_2 are determined from the system of equations, (7)–(10). Note that since this system of equations involves transcendental functions, a closed-form solution cannot be determined. To determine the system parameters, the system of equations, (7)–(10), is solved numerically [13].

Various tradeoff scenarios are discussed in this section. The dependence of the system parameters on R_1 is presented in Section IV-A. The design of a distributed on-chip decoupling capacitor network with minimum magnitude C_1 is discussed in Section IV-B.

A. Dependence of System Parameters on R_1

The parameters of a distributed on-chip decoupling capacitor network for typical values of R_1 are listed in Table I. Note that the minimum magnitude of R_2 exists for which the parameters of the system can be determined. If R_2 is sufficiently small, the distributed decoupling capacitor network degenerates to a system with a single capacitor (where C_1 and C_2 are merged). For the parameters listed in Table I, the minimum magnitude of R_2 is four ohms, as determined from numerical simulations.

TABLE I

DEPENDENCE OF THE PARAMETERS OF THE DISTRIBUTED ON-CHIP DECOUPLING CAPACITOR NETWORK ON R_1

R_1 (Ω)	$R_2 = 5\,(\Omega)$		$R_2 = 10\,(\Omega)$	
	C_1 (pF)	C_2 (pF)	C_1 (pF)	C_2 (pF)
1	1.35	7.57	3.64	3.44
2	2.81	5.50	4.63	2.60
3	4.54	3.64	5.88	1.77
4	6.78	1.87	7.56	0.92
5	10.00	0	10.00	0
$V_{dd} = 1$ V, $V_{load} = 0.9$ V, $t_r = 100$ ps, and $I_{max} = 0.01$ A				

Note that the parameters of a distributed on-chip decoupling capacitor network are determined by the parasitic resistance of the P/G line(s) connecting C_1 to the current load. As R_1 increases, the capacitor C_1 increases substantially (see Table I). This increase in C_1 is due to R_1 becoming comparable to R_2, and C_1 providing a greater portion of the total current. Alternatively, the system of distributed on-chip decoupling capacitors degenerates to a single oversized capacitor. The system of distributed on-chip decoupling capacitors should therefore be carefully designed. Since the distributed on-chip decoupling capacitor network is strongly dependent upon the first level of interconnection (R_1), C_1 should be placed as physically close as possible to the current load, reducing R_1. If such an allocation is not practically possible, the current load should be partitioned, permitting an efficient allocation of the distributed on-chip decoupling capacitors under specific technology constraints.

B. Minimum C_1

In practical applications, the size of C_1 (the capacitor closest to the current load) is typically limited by technology constraints, such as the maximum capacitance density, and design constraints, such as the available area. The magnitude of the first capacitor in a distributed system is therefore typically small. In this section, the dependence of the distributed on-chip decoupling capacitor network on R_1 is determined for minimum C_1. A target magnitude of 1 pF is assumed for C_1. The parameters of a system of distributed on-chip decoupling capacitors as a function of R_1 under the constraint of a minimum C_1 are listed in Table II. Note that V_{C_2} denotes the voltage across C_2 after discharge.

TABLE II

DISTRIBUTED ON-CHIP DECOUPLING CAPACITOR NETWORK AS A FUNCTION OF R_1 UNDER THE CONSTRAINT OF A MINIMUM C_1

R_1 (Ω)	$V_{C_2} \neq$ const				$V_{C_2} = 0.95$ volt	
	R_2 (Ω)	C_2 (pF)	R_2 (Ω)	C_2 (pF)	R_2 (Ω)	C_2 (pF)
1	2	5.59	5	8.69	4.68	8.20
2	2	6.68	5	11.64	3.46	8.40
3	2	8.19	5	17.22	2.28	8.60
4	2	10.46	5	31.70	1.13	8.80
5	2	14.21	5	162.10	–	–
$V_{dd} = 1$ V, $V_{load} = 0.9$ V, $t_r = 100$ ps, $I_{max} = 0.01$ A, and $C_1 = 1$ pF						

Note that two scenarios are considered in Table II to evaluate the dependence of a distributed system of on-chip decoupling capacitors. In the first scenario, the distributed on-chip decoupling capacitor network is designed to maintain the minimum tolerable voltage across the terminals of a current load. In this case, the magnitude of C_2 increases with R_1, becoming impractically large for large R_2. In the second scenario, an additional constraint (the voltage across C_2) is applied to reduce the voltage fluctuations on the shared P/G lines. In this case, as R_1 increases, C_2 slightly increases. In order to satisfy the constraint for V_{C_2}, R_2 should be significantly reduced for large values of R_1, meaning that the second capacitor should be placed close to the first capacitor. As R_1 is further increased, R_2 becomes negligible, implying that capacitors C_1 and C_2 should be merged to provide the required charge to the distant current load.

V. CASE STUDY

The dependence of the system of distributed on-chip decoupling capacitors on the current load and the parasitic impedance of the power delivery system is described in this section to quantitatively illustrate the previously presented concepts. The load is modeled as a ramp current source with a 100 ps rise time. The minimum tolerable voltage across the load terminals is 90% of the power supply level. The magnitude of the on-chip decoupling capacitors for various parasitic resistances of the metal lines connecting the capacitors to the current load is listed in Table III. The parameters of the distributed on-chip decoupling capacitor network listed in Table III are determined for two amplitudes of the current load.

The parameters of the system of distributed on-chip decoupling capacitors are analytically determined from (7)–(10). The resulting power supply noise is estimated using SPICE and compared to the maximum tolerable level (the minimum voltage across the load terminals V_{load}^{min}). The maximum voltage drop across C_2 at the end of the switching activity is also estimated and compared to $V_{C_2}^{min}$. Note that the analytic solution produces accurate estimates of the on-chip decoupling capacitors for typical parasitic resistances of a power distribution grid. The maximum error in this case study is 0.003%.

TABLE III

THE MAGNITUDE OF ON-CHIP DECOUPLING CAPACITORS AS A FUNCTION OF THE PARASITIC RESISTANCE OF THE POWER/GROUND LINES CONNECTING

THE CAPACITORS TO THE CURRENT LOAD

R_1 (Ω)	R_2 (Ω)	I_{max} (A)	C_1 (pF)	C_2 (pF)	V_{load} (mV)		Error (%)	V_{C_2} (mV)		Error (%)
					V_{load}^{min}	SPICE		$V_{C_2}^{min}$	SPICE	
0.5	4.5	0.01	0	9.99999	900	899.999	0.0001	950	949.999	0.0001
0.5	6	0.01	1.59747	6.96215	900	899.986	0.002	950	949.983	0.002
0.5	8	0.01	2.64645	4.97091	900	899.995	0.0006	950	949.993	0.0004
0.5	10	0.01	3.22455	3.87297	900	899.997	0.0003	950	949.996	0.0004
0.5	12	0.01	3.59188	3.17521	900	899.998	0.0002	950	949.997	0.0003
0.5	14	0.01	3.84641	2.69168	900	899.998	0.0002	950	949.997	0.0003
0.5	16	0.01	4.03337	2.33650	900	899.999	0.0001	950	949.998	0.0002
0.5	18	0.01	4.17658	2.06440	900	899.998	0.0002	950	949.998	0.0002
0.5	20	0.01	4.28984	1.84922	900	899.999	0.0001	950	949.998	0.0002
0.5	1.5	0.025	0	24.99930	900	899.998	0.0002	950	949.998	0.0002
0.5	2	0.025	4.25092	17.56070	900	899.999	0.0001	950	949.999	0.0001
0.5	3	0.025	7.97609	11.04180	900	899.999	0.0001	950	949.999	0.0001
0.5	4	0.025	9.67473	8.06921	900	899.999	0.0001	950	949.999	0.0001
0.5	5	0.025	10.65000	6.36246	900	899.999	0.0001	950	949.999	0.0001
0.5	6	0.025	11.2838	5.25330	900	899.999	0.0001	950	949.999	0.0001
0.5	7	0.025	11.72910	4.47412	900	899.999	0.0001	950	949.999	0.0001
0.5	8	0.025	12.05910	3.89653	900	899.999	0.0001	950	949.999	0.0001
0.5	9	0.025	12.31110	3.44905	900	899.980	0.002	950	949.973	0.003
$V_{dd} = 1$ V, and $t_r = 100$ ps										

From Table III, note that in the case of a large R_2, the distributed decoupling capacitor network degenerates to a system with a single capacitor. Capacitor C_1 is therefore excessively large. Conversely, if C_2 is placed close to C_1 (R_2 is small), C_2 is excessively large and the system behaves as a single capacitor. An optimal ratio $\frac{R_2}{R_1}$ therefore exists for specific characteristics of the current load that results in the minimum budgeted on-chip decoupling capacitance. Alternatively, in this case, both capacitors provide an equal portion of the total charge (see Table III for $R_1 = 0.5\,\Omega$ and $R_2 = 10\,\Omega$).

VI. CONCLUSIONS

A system of distributed on-chip decoupling capacitors has been proposed in this paper. A distributed on-chip decoupling capacitor network is an efficient solution for providing the required on-chip decoupling capacitance under existing technology constraints. In a system of distributed on-chip decoupling capacitors, each capacitor is sized based on the parasitic impedance of the power delivery system. Initially, the capacitor closest to the current load provides the required charge. Once the first capacitor is depleted, the next decoupling capacitor is activated, providing a large portion of the total current drawn by the load. This procedure is repeated until the last decoupling capacitor becomes active. Hierarchically allocating the on-chip decoupling capacitors greatly relaxes the technology constraints for physically distant capacitors, making the distant on-chip decoupling capacitors more effective.

REFERENCES

[1] M. Popovich, A. V. Mezhiba, and E. G. Friedman, *Power Distribution Networks with On-Chip Decoupling Capacitors*, Springer Publishers, 2008 (in press).

[2] M. Popovich and E. G. Friedman, "Impedance Characteristics of Decoupling Capacitors in Multi-Power Distribution Systems," *Proceedings of the IEEE International Conference on Electronics, Circuits and Systems*, pp. 160–163, December 2004.

[3] M. Popovich and E. G. Friedman, "Decoupling Capacitors for Multi-Voltage Power Distribution Systems," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 14, No. 3, pp. 217–228, March 2006.

[4] A. V. Mezhiba and E. G. Friedman, "Scaling Trends of On-Chip Power Distribution Noise," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 12, No. 4, pp. 386–394, April 2004.

[5] M. Popovich and E. G. Friedman, "Noise Aware Decoupling Capacitors for Multi-Voltage Power Distribution Systems," *Proceedings of the IEEE International Symposium of Quality Electronic Design*, pp. 334–339, March 2005.

[6] M. Popovich, E. G. Friedman, M. Sotman, A. Kolodny, and R. M. Secareanu, "Maximum Effective Distance of On-Chip Decoupling Capacitors in Power Distribution Grids," *Proceedings of the ACM Great Lakes Symposium on VLSI*, pp. 173–179, April/May 2006.

[7] H. Su, S. S. Sapatnekar, and S. R. Nassif, "Optimal Decoupling Capacitor Sizing and Placement for Standard-Cell Layout Designs," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 22, No. 4, pp. 428–436, April 2003.

[8] S. Zhao, K. Roy, and C.-K. Koh, "Decoupling Capacitance Allocation and Its Application to Power-Supply Noise-Aware Floorplanning," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 21, No. 1, pp. 81–92, January 2002.

[9] J. Kim, *et al.*, "Separated Role of On-Chip and On-PCB Decoupling Capacitors for Reduction of Radiated Emission on Printed Circuit Board," *Proceedings of the IEEE International Symposium on Electromagnetic Compatibility*, pp. 531–536, August 2001.

[10] T. Hubing, "Effective Strategies for Choosing and Locating Printed Circuit Board Decoupling Capacitors," *Proceedings of the IEEE International Symposium on Electromagnetic Compatibility*, pp. 632–637, August 2005.

[11] L. D. Smith, R. E. Anderson, D. W. Forehand, T. J. Pelc, and T. Roy, "Power Distribution System Design Methodology and Capacitor Selection for Modern CMOS Technology," *IEEE Transactions on Advanced Packaging*, Vol. 22, No. 3, pp. 284–291, August 1999.

[12] M. E. Van Valkenburg, *Network Analysis*, Prentice-Hall, 1974.

[13] Mathematica 5.2, Wolfram Research, Inc.