

Maximum Effective Distance of On-Chip Decoupling Capacitors in Power Distribution Grids

Mikhail Popovich
Eby G. Friedman

Dept. of Electrical and
Computer Engineering
University of Rochester
Rochester, New York 14627
USA

nhlover@ece.rochester.edu
friedman@ece.rochester.edu

Michael Sotman
Avinoam Kolodny

Dept. of Electrical Engineering
Technion – Israel Institute of
Technology
Haifa 32000, Israel

michael.sotman@intel.com
kolodny@ee.technion.ac.il

Radu M. Secareanu

MMSTL
Freescale Semiconductor
Tempe, Arizona 85284, USA
r54143@freescale.com

ABSTRACT

Decoupling capacitors are widely used to reduce power supply noise. On-chip decoupling capacitors have traditionally been allocated into the available white space on a die. The efficacy of on-chip decoupling capacitors depends upon the impedance of the power/ground lines connecting the capacitors to the current loads and power supplies. A maximum effective radius exists for each on-chip decoupling capacitor. Beyond this effective distance, a decoupling capacitor is completely ineffective. Two effective radii determined by the target impedance (during discharge) and charge time are presented in this paper. Depending upon the parasitic impedance of the power distribution system, the maximum voltage drop as seen at the current load is achieved either at the first droop or at the end of the switching activity (the second droop). Two criteria to estimate the minimum required on-chip decoupling capacitance are developed based on the critical parasitic impedance. To be effective, the decoupling capacitor has to be fully charged before the next switching event. A design space is described that characterizes the tolerable parasitic resistances and inductances, while restoring the charge on the decoupling capacitor within a target charge time. An overall design methodology for placing on-chip decoupling capacitors is presented in this paper. It is shown that for an on-chip decoupling capacitor to be effective, both effective radii criteria should be simultaneously satisfied.

This research is supported in part by the Semiconductor Research Corporation under Contract No. 2003-TJ-1068 and 2004-TJ-1207, the National Science Foundation under Contract No. CCR-0304574, the Fulbright Program under Grant No. 87481764, a grant from the New York State Office of Science, Technology & Academic Research to the Center for Advanced Technology in Electronic Imaging Systems, and by grants from Intel Corporation, Eastman Kodak Company, Manhattan Routing, and Intrinsix Corporation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GLSVLSI'06, April 30–May 2, 2006, Philadelphia, PA, USA.
Copyright 2006 ACM 1-59593-347-6/06/0004 ...\$5.00.

Categories and Subject Descriptors

B.7.m [Integrated Circuits]: Miscellaneous—*power distribution grids, power distribution systems, decoupling capacitors*

General Terms

Design

Keywords

Power distribution systems, power distribution grids, decoupling capacitors

1. INTRODUCTION

The feature size of integrated circuits (IC) has been aggressively reduced in the pursuit of higher speed, lower power, and less cost. Semiconductor technologies with a feature size of several tens of nanometers are currently in development. The scaling of CMOS is expected to continue for at least another decade. Future nanometer circuits will soon contain more than a billion transistors and operate at clock speeds well over 10 GHz [1]. Distributing robust and reliable power and ground levels in such a high speed, high complexity environment is, therefore, a highly challenging task [2].

Decoupling capacitors are widely used to manage power supply noise. A decoupling capacitor acts as a reservoir of charge, which is released when the power supply voltage at a particular current load drops below some tolerable level. Alternatively, decoupling capacitors can be used to reduce the impedance of power delivery systems operating at high frequencies. Since the inductance scales poorly [3], the location of the decoupling capacitors significantly affects the design of the power/ground (P/G) networks in high performance ICs such as microprocessors. With increasing frequencies, a distributed system of on-chip decoupling capacitors is needed.

The efficacy of decoupling capacitors depends upon the impedance of the conductors connecting the capacitors to the current loads and power sources. During discharge, the current flowing from the decoupling capacitor to the current load results in resistive noise (IR drops) and inductive noise ($L \frac{di}{dt}$ drops) due to the parasitic resistance and inductance of the power delivery network. The resulting voltage drop

at the current load is, therefore, always greater than the voltage drop at the decoupling capacitor. Once the switching event is completed, a decoupling capacitor has to be fully charged before the next clock cycle begins. During the charging phase, the voltage at the decoupling capacitor rises exponentially. The charge time of a capacitor is determined by the parasitic resistance and inductance of the interconnect between the capacitor and the power supply. A design space that characterizes a tolerable resistance and inductance exists, permitting the restoration of the charge at the decoupling capacitor within the target charge time. The maximum frequency at which the decoupling capacitor is effective is determined by the parasitic resistance and inductance of the metal lines and the size of the decoupling capacitor.

A maximum parasitic impedance between the decoupling capacitor and the current load or power supply exists at which the decoupling capacitor is effective. Alternatively, to be effective, a decoupling capacitor should be placed within the maximum effective distance from the current load or power supply, as shown in Fig. 1. The optimal allocation of the on-chip decoupling capacitors is, therefore, an important issue in high performance ICs. The effective radii of the on-chip decoupling capacitors are the primary subject of this paper. Design techniques to estimate the minimum required on-chip decoupling capacitance are also presented.




Figure 1: The effective radii of an on-chip decoupling capacitor. The on-chip decoupling capacitor should be placed such that both the current load and the power supply are located inside the respective effective radius. If the power supply is located outside the effective radius d_{ch1}^{max} , the current load should be partitioned, resulting in a smaller decoupling capacitor and, therefore, an increased effective distance d_{ch2}^{max} .

The paper is organized as follows. The effective radius of an on-chip decoupling capacitor determined by the target impedance is presented in Section 2. Design techniques to estimate the minimum magnitude of the required on-chip decoupling capacitance are discussed in Section 3. The effective radius of an on-chip decoupling capacitor based on the charge time is presented in Section 4. A model of a power distribution network is introduced in Section 5. Simulation results for typical values of on-chip parasitic resistances and

inductances are discussed in Section 6. Some specific conclusions are summarized in Section 7.

2. EFFECTIVE RADIUS BASED ON TARGET IMPEDANCE

Only a small fraction of the power supply voltage is typically permitted as a ripple voltage (power noise). A target impedance of a power distribution system is [4]

$$Z_{target} = \frac{V_{dd} \times Ripple}{I}, \quad (1)$$

where V_{dd} is the power supply voltage, $Ripple$ is the allowable ripple voltage, and I is the current. According to general scaling theory [5], the current I will further increase and the power supply voltage will continue to drop, at least for the next ten years. The impedance of a power distribution system should therefore be reduced to satisfy target power noise constraints. The target impedance of a power distribution system is falling at an alarming rate, a factor of two per technology generation [4]. The specific impedance must be satisfied not only at DC, but also at all frequencies where current transients occur [6].

Neglecting the parasitic capacitance, the impedance of a unit length wire can be modeled as $Z'(\omega) = r + j\omega l$, where r and l are the resistance and inductance per length, respectively, and ω is an equivalent frequency, determined by the rise time of the current load. The unit length inductance l is the effective inductance per unit length in the power distribution grid, incorporating the partial self inductance and mutual coupling among the lines. The target impedance of a metal line with a particular length is therefore

$$Z(\omega) = Z'(\omega) \times d, \quad (2)$$

where $Z'(\omega)$ is the impedance of a unit length metal line, and d is the distance between the decoupling capacitor and the current load. Substituting (1) into (2), the maximum effective radius (or distance) d_Z^{max} between the decoupling capacitor and the current load is

$$d_Z^{max} = \frac{Z_{target}}{Z'(\omega)} = \frac{V_{dd} \times Ripple}{I \times \sqrt{r^2 + \omega^2 l^2}}, \quad (3)$$

where $\sqrt{r^2 + \omega^2 l^2}$ denotes the magnitude of the impedance of a unit length wire. Note that the maximum effective radius as determined by the target impedance is inversely proportional to the magnitude of the current load and the impedance of a unit length line. Also note, in a meshed structure, multiple impedance paths between any two points add in parallel. The maximum effective distance corresponding to Z_{target} is therefore larger as compared to a single line. In this paper, the effective radius of an on-chip decoupling capacitor is defined as follows.

Definition 1: The effective radius of an on-chip decoupling capacitor is the maximum distance between the current load or power supply and the decoupling capacitor, at which the capacitor is capable of providing sufficient charge to the current load in order to maintain the overall power distribution noise below the maximum tolerable level.

3. MINIMUM REQUIRED ON-CHIP DECOUPLING CAPACITANCE

To estimate the on-chip decoupling capacitance required to support the current demand, the current load is modeled

as a triangular current source. The magnitude of the current source increases linearly, reaching the maximum current I_{max} at peak time t_p . The magnitude of the current source decays linearly, becoming zero at t_f , as shown in Fig. 2. The on-chip power distribution network is modeled as a series RL circuit. A single decoupling capacitor with a single current load is assumed to support the voltage fluctuations at the P/G terminals. The maximum tolerable ripple at the current load is 10% of the power supply voltage.




Figure 2: Linear approximation of the current demand of a power distribution network by a current source. The magnitude of the current source reaches the maximum current value I_{max} at peak time t_p . t_r and t_f denote the rise and fall times of the current load, respectively.

The total charge Q_{dis} required to support the current demand during the switching event is modeled as the sum of the areas of two triangles (see Fig. 2). Since the required charge is provided by an on-chip decoupling capacitor, the voltage at the capacitor during discharge drops below the initial (power supply) voltage. Thus, the required charge during the entire switching event is

$$Q_{dis} = \frac{I_{max} \times (t_r + t_f)}{2} = C_{dec} \times (V_{dd} - V_C^f), \quad (4)$$

where I_{max} is the maximum magnitude of the current load, t_r and t_f are the rise and fall times, respectively, C_{dec} is the decoupling capacitance, V_{dd} is the power supply voltage, and V_C^f is the voltage across the decoupling capacitor after the switching event. Note that, since no current exists after switching, the voltage at the current load is equal to the voltage on the decoupling capacitor.

The voltage at the P/G terminals of a power delivery system should not exceed the maximum level (typically 10% of the power supply voltage) to guarantee fault-free operation. Thus,

$$V_C^f \equiv V_{load}^f \geq 0.9 V_{dd}. \quad (5)$$

Substituting (5) into (4) and solving for C_{dec} , the minimum on-chip decoupling capacitance required to support the current demand during a switching event is

$$C_{dec}^f \geq \frac{I_{max} \times (t_r + t_f)}{0.2 V_{dd}}. \quad (6)$$

Note that (6) is only applicable to the case when the voltage drop at the end of the switching event (the second droop) is larger than the voltage drop at the peak time t_p (the first droop). This phenomenon can be explained as follows. The

voltage drop as seen at the current load is caused by the current flowing through the parasitic resistance and inductance of the on-chip power distribution system. The resulting voltage fluctuations are the sum of the ohmic IR voltage drop, inductive $L \frac{dI}{dt}$ voltage drop, and voltage drop at the decoupling capacitor at t_p . A critical parasitic RL impedance, therefore, exists for any given set of rise and fall times. Beyond this critical impedance, the first droop becomes larger than the second droop, as shown in Fig. 3. The decoupling capacitor should therefore be increased to reduce the voltage drop at the capacitor V_C^r , lowering the magnitude of the first droop.




Figure 3: Power distribution noise during discharge of an on-chip decoupling capacitor: $I_{max} = 100$ mA, $V_{dd} = 1$ V, $t_r = 20$ ps, $t_f = 80$ ps, $R = 100$ m Ω , $L = 15$ pH, and $C_{dec} = 50$ pF. For these parameters, the parasitic impedance of the metal lines connecting the decoupling capacitor to the current load is larger than the critical impedance. The first droop is larger than the second droop and (6) underestimates the required decoupling capacitance. The resulting voltage drop on the power terminal is therefore larger than the maximum tolerable noise.

The charge Q_r required to support the current demand during the rise time of the current load is equal to the area of the triangle formed by I_{max} and t_r . The required charge is provided by the on-chip decoupling capacitor. The voltage at the decoupling capacitor therefore drops below the power supply level by ΔV_C^r . The required charge during t_r is

$$Q_r = \frac{I_{max} \times t_r}{2} = C_{dec} \times \Delta V_C^r, \quad (7)$$

where Q_r is the charge drawn by the current load during t_r and ΔV_C^r is the voltage drop at the decoupling capacitor at t_p . From (7),

$$\Delta V_C^r = \frac{I_{max} \times t_r}{2 C_{dec}}. \quad (8)$$

By time t_p , the voltage drop as seen from the current load (the first droop) is the sum of the ohmic IR drop, the inductive $L \frac{dI}{dt}$ drop, and the voltage drop at the decoupling capacitor. Alternatively, the power noise is further increased

by the voltage drop ΔV_C^r . In this case, the voltage at the current load is

$$V_{load}^r = V_{dd} - I \times R - L \frac{dI}{dt} - \Delta V_C^r, \quad (9)$$

where R and L are the parasitic resistance and inductance of the P/G lines, respectively. Note that the last term in (9) accounts for the voltage drop ΔV_C^r at the decoupling capacitor during the rise time of the current load.

Assuming that $V_{load}^r \geq 0.9 V_{dd}$ and substituting (8) into (9) and solving for C_{dec}^r , the minimum on-chip decoupling capacitance to support the current demand during t_r is

$$C_{dec}^r \geq \frac{I_{max} \times t_r}{2 (0.1 V_{dd} - I \times R - L \frac{dI}{dt})}. \quad (10)$$

Note that if the first droop is larger than the second droop, C_{dec} is oversized. Hence, the voltage drop at the end of the switching event is always smaller than the maximum tolerable noise.

Also note that, as opposed to (6), (10) depends upon the parasitic impedance of the on-chip power distribution system. Alternatively, in the case of the first droop, the required charge released by the decoupling capacitor is determined by the parasitic resistance and inductance of the P/G lines connecting the decoupling capacitor to the current load. Assuming the impedance of a single line, the critical line length d_{crit} can be determined by equalizing C_{dec}^r and C_{dec}^f ,

$$\frac{I_{max} \times t_r}{(0.1 V_{dd} - I r d_{crit} - l d_{crit} \frac{dI}{dt})} = \frac{I_{max} \times (t_r + t_f)}{0.1 V_{dd}}, \quad (11)$$

where r and l are the per length parasitic resistance and inductance of a single line, respectively. Solving (11) for d_{crit} ,

$$d_{crit} = \frac{0.1 V_{dd} \left(1 - \frac{t_r}{t_r + t_f}\right)}{I r + l \frac{dI}{dt}}. \quad (12)$$

For a single line connecting the current load to the decoupling capacitor, the minimum required on-chip decoupling capacitor is determined by (6) for lines shorter than d_{crit} and by (10) for lines longer than d_{crit} . Note that for a line length equal to d_{crit} , (6) and (10) result in the same required capacitance. Also note that the maximum length of a single line is determined by (3). In the case of multiple current paths existing between the current load and a decoupling capacitor, no closed form solution for the critical line length exists. The impedance of the power grid connecting a decoupling capacitor to a current load is compared to the critical impedance. Either (6) or (10) is utilized to estimate the required on-chip decoupling capacitance.

4. EFFECTIVE RADIUS DETERMINED BY CHARGE TIME

Once discharged, the decoupling capacitor must be charged to support the current demands during the next switching event. If charge at the capacitor is not fully restored during the relaxation time between two consecutive switching events (the charge time), the decoupling capacitor will gradually be depleted, becoming ineffective after several clock cycles. A maximum effective radius, therefore, exists for an

on-chip decoupling capacitor determined during the charging phase for a target charge time. Similar to the effective radius based on the target impedance presented in Section 2, an on-chip decoupling capacitor should be placed in close proximity to the power supply (power pins) to be effective.

To determine the current flowing through a decoupling capacitor during the charging phase, the parasitic impedance of a power distribution system is modeled as a series RL circuit between the decoupling capacitor and the power supply, as shown in Fig. 4. When the discharge is completed, the switch is closed and the charge is restored on the decoupling capacitor. The initial voltage V_C^0 across the decoupling capacitor is determined by the maximum voltage drop during discharge.




Figure 4: Circuit diagram of charging an on-chip decoupling capacitor.

For the circuit shown in Fig. 4, the KVL equation for the current in the circuit is [7]

$$L \frac{di_{ch}}{dt} + R i_{ch} + \frac{1}{C_{dec}} \int i_{ch} dt = V_{dd}. \quad (13)$$

Differentiating (13),

$$L \frac{d^2 i_{ch}}{dt^2} + R \frac{di_{ch}}{dt} + \frac{1}{C_{dec}} i_{ch} = 0. \quad (14)$$

Assuming an overdamped power distribution system with an on-chip decoupling capacitor (which is typical for lines longer than several micrometers in a 90 nm CMOS technology) and applying the initial conditions, the current through the decoupling capacitor during the charging phase is

$$i_{ch}(t) = \frac{I_{max} (t_r + t_f)}{4 L C_{dec} \sqrt{\left(\frac{R}{2L}\right)^2 - \frac{1}{LC_{dec}}}} \times \left\{ \exp \left[\left(-\frac{R}{2L} + \sqrt{\left(\frac{R}{2L}\right)^2 - \frac{1}{LC_{dec}}} \right) t \right] - \exp \left[\left(-\frac{R}{2L} - \sqrt{\left(\frac{R}{2L}\right)^2 - \frac{1}{LC_{dec}}} \right) t \right] \right\}. \quad (15)$$

The voltage at the decoupling capacitor during the charging phase can be determined by integrating (15) from zero to the charge time,

$$V_C(t) = \frac{1}{C} \int_0^{t_{ch}} i_{ch}(t) dt, \quad (16)$$

where t_{ch} is the charge time, and $V_C(t)$ and $i_{ch}(t)$ are the voltage at the decoupling capacitor and the current flowing through the decoupling capacitor during the charging phase, respectively. Substituting (15) into (16) and integrating from zero to t_{ch} , the voltage across the decoupling capacitor during the charging phase is given by (17).

From (17), the design space is determined for the maximum tolerable resistance and inductance, permitting the

$$V_C(t_{ch}) = \frac{I_{max}(t_r + t_f)}{4C_{dec}^2 L \sqrt{\left(\frac{R}{2L}\right)^2 - \frac{1}{LC_{dec}}}} \times \left(\frac{\exp\left[\left(-\frac{R}{2L} + \sqrt{\left(\frac{R}{2L}\right)^2 - \frac{1}{LC_{dec}}}\right)t_{ch}\right] - 1}{-\frac{R}{2L} + \sqrt{\left(\frac{R}{2L}\right)^2 - \frac{1}{LC_{dec}}}} + \frac{1 - \exp\left[\left(-\frac{R}{2L} - \sqrt{\left(\frac{R}{2L}\right)^2 - \frac{1}{LC_{dec}}}\right)t_{ch}\right]}{-\frac{R}{2L} - \sqrt{\left(\frac{R}{2L}\right)^2 - \frac{1}{LC_{dec}}}} \right). \quad (17)$$

decoupling capacitor to be recharged within a given t_{ch} , as shown in Fig. 5. The parasitic resistance and inductance should be maintained below the maximum tolerable values, permitting the decoupling capacitor to be charged during the relaxation time. Note that as the parasitic resistance of the power delivery network decreases, the voltage across the decoupling capacitor increases exponentially. In contrast, the voltage across the decoupling capacitor during the charging phase is almost independent of the parasitic inductance, slightly increasing with inductance. This phenomenon is due to the behavior that an inductor resists sudden changes in the current. Alternatively, an inductor maintains the charging current at a particular level for a longer time. Thus, the decoupling capacitor is charged faster.




Figure 5: Design space for maximum tolerable parasitic resistance and inductance of a power distribution grid: $I_{max} = 100$ mA, $t_r = 100$ ps, $t_f = 300$ ps, $C_{dec} = 100$ pF, $V_{dd} = 1$ volt, and $t_{ch} = 400$ ps. For a target charge time, the maximum resistance and inductance can be determined, resulting in the voltage across the decoupling capacitor being greater or equal to the power supply voltage (dark line).

An overall design methodology for placing on-chip decoupling capacitors is as follows. The maximum effective radius based on the target impedance is determined from (3) for a particular current load (circuit block), power supply voltage, and allowable ripple. The minimum required on-chip decoupling capacitance is estimated to support the required current demand. If the second droop is greater than the first droop, (6) is used to determine the required on-chip decoupling capacitance. If the first droop dominates, the

on-chip decoupling capacitance is determined from (10). In the case of a single line connecting the decoupling capacitor to the current load, the critical wire length is determined from (12). The maximum effective distance based on the charge time is determined from (17). Note that (17) results in a range of tolerable parasitic resistance and inductance of the metal lines connecting the decoupling capacitor to the power supply. Also note that the on-chip decoupling capacitor should be placed such that both the power supply and the current load are located inside the respective effective radius, as shown in Fig. 6. If this placement is not possible, the current load (circuit block) should be partitioned into several blocks and the on-chip decoupling capacitors should be allocated for each block, satisfying both effective radii. Note that the effective radius determined by the target impedance does not depend on the decoupling capacitance. The effective radius determined by the charge time however is inversely proportional to C_{dec}^2 . Finally, the on-chip decoupling capacitors should be distributed across the circuit to provide the required charge for each functional unit.




Figure 6: Effective radii of an on-chip decoupling capacitor. For a power distribution system modeled by a distributed RL mesh, the maximum effective radius is the manhattan distance between two points. The overall effective radius is therefore shaped like a diamond.

5. MODEL OF A POWER DISTRIBUTION NETWORK

On-chip power distribution networks in high performance ICs are commonly modeled as a mesh. Early in the de-

sign process, minimal physical information characterizing the P/G structure is available. A simplified model of a power distribution system is therefore appropriate. For simplicity, equal segments within a mesh structure are assumed. The current demands of a particular module are modeled as current sources with equivalent magnitude and switching activities. The current load is located at the center of a circuit module and determines the connection point of the circuit module to the power grid.

Typical effective radii of an on-chip decoupling capacitor is in the range of several hundreds micrometers. In order to determine the location of an on-chip decoupling capacitor, the size of each RL mesh segment should be much smaller than the effective radii. In modern high performance ICs such as microprocessors with die sizes approaching 1.5 inches by 1.5 inches, a fine mesh is infeasible to simulate. In the case of a coarse mesh, the effective radius is smaller than the size of each segment. The location of each on-chip decoupling capacitor, therefore, cannot be accurately determined. To resolve this dilemma, the accuracy of the capacitor location can be traded off with the complexity of the power distribution network. A hot spot (an area where the power supply voltage drops below the minimum tolerable level) is first determined based on a coarse mesh. A fine mesh is used next inside each hot spot to accurately estimate the effective radius of the on-chip decoupling capacitor.

In modern high performance ICs, up to 3000 I/O pins can be available [1]. Only half of the I/O pads are typically used to distribute power. The other half is dedicated to signaling. Assuming an equal distribution of power and ground pads, a quarter of the total number of pads is typically available for power *or* ground delivery. For high performance ICs with die sizes of 1.5 inches by 1.5 inches inside a flip chip package, the distance between two adjacent power or ground pads is about 1300 μm . Note that the proposed approach to model a power distribution system is applicable to ICs with both conventional low cost and advanced high performance packaging.

6. CASE STUDY

The dependence of the effective radii of an on-chip decoupling capacitor on a power distribution system is described in this section to quantitatively illustrate the concepts previously presented in this paper. A flip chip package is assumed. An on-chip power distribution system with a flip chip pitch (the area formed by the four closest power pins) is modeled by an RL mesh of forty by forty equal segments to accurately determine the maximum effective distance of an on-chip decoupling capacitor. The load is modeled as a triangular current source with a 100 ps rise time and 300 ps fall time. The maximum tolerable ripple at the load is 10% of the power supply voltage. The relaxation time between two consecutive switching events (the charge time) is set to 400 ps. The maximum effective radii determined by the target impedance and charge time for three sets of on-chip parasitic resistances and inductances are listed in Table 1. Three scenarios listed in Table 1 are chosen to represent typical parasitic resistances and inductances of the top, intermediate, and bottom layers of an on-chip power distribution grid.

From (12), for the parameters listed in Table 1, the critical voltage drop is 75 mV. If the voltage fluctuations at the current load do not exceed the critical voltage, the magni-

tude of the second droop dominates and the required on-chip decoupling capacitance is determined by (6). Note that for the three scenarios with a 10 mA current load, the maximum effective radii of the on-chip decoupling capacitor based on the target impedance and charge time are larger than forty cells (the longest distance within the mesh from the center of the mesh to the corner). The maximum effective radii of the on-chip decoupling capacitor is therefore larger than the pitch size. The decoupling capacitor can therefore be placed anywhere inside the pitch. Note that the required on-chip decoupling capacitance is estimated for the longest distance (40 cells). For a 100 mA current load, the voltage fluctuations at the current load exceed the critical voltage drop. The first droop dominates and the required on-chip decoupling capacitance is determined by (10).

Table 1: Maximum effective radii of an on-chip decoupling capacitor

Metal Layer	Res. ($\Omega/\mu\text{m}$)	Ind. (pH/ μm)	I_{load} (A)	C_{dec} (pF)	d_{max} (cells)	
					Z	t_{ch}
Top	0.007	0.5	0.01	20	>40	>40
	0.007	0.5	0.1	357	2	>40
	0.007	0.5	1	–	<1	–
Int.	0.04	0.3	0.01	20	>40	>40
	0.04	0.3	0.1	227	1	<1
	0.04	0.3	1	–	<1	–
Bot.	0.1	0.1	0.01	20	>40	>40
	0.1	0.1	0.1	–	<1	–
	0.1	0.1	1	–	<1	–

$V_{dd} = 1\text{ V}$, $V_{ripple} = 100\text{ mV}$,
 $t_r = 100\text{ ps}$, $t_f = 300\text{ ps}$, $t_{ch} = 400\text{ ps}$

Note that C_{dec}^r as determined by (10) increases rapidly with the effective radius based on the target impedance, becoming infinite at d_Z^{max} . As the decoupling capacitor is moved closer to the current load, the required on-chip decoupling capacitance, estimated by (10), decreases. A tradeoff, therefore, exists between the maximum effective distance determined by the target impedance and the size of the minimum required on-chip decoupling capacitance (if the overall voltage drop at the current load is primarily caused by the first droop).

The effective radii of an on-chip decoupling capacitor decreases linearly with current load. The optimal size of an RL distributed mesh should therefore be determined for a particular current demand. If the magnitude of the current requirements is low, the mesh can be coarser, significantly decreasing the simulation time. For a 10 mA current load, the effective radii determined both by the target impedance and charge time are longer than the pitch size. Thus, the distributed mesh is overly fine. For a current load of 1 A, the effective radii are shorter than one cell, meaning that the distributed RL mesh is overly coarse. A finer mesh should therefore be used to accurately estimate the maximum effective radii of the on-chip decoupling capacitor.

Note that the maximum effective radius as determined by the charge time becomes impractically short for large decoupling capacitors, making the capacitors ineffective. In this case, the decoupling capacitor should be placed closer

to the current load, permitting the decoupling capacitance to be decreased. Alternatively, the current load should be partitioned into several blocks, lowering the required on-chip decoupling capacitance. The parasitic impedance between the decoupling capacitor and the current load or power supply should be reduced, increasing the maximum effective radii of the on-chip decoupling capacitors.

7. CONCLUSIONS

A design methodology for placing on-chip decoupling capacitors is presented in this paper. It is shown that a maximum effective distance exists between the current load or power supply and the decoupling capacitor. Beyond this distance, the decoupling capacitor becomes ineffective.

The maximum effective radii of an on-chip decoupling capacitor based on the target impedance (discharge) and charge time have been determined. Depending upon the parasitic impedance of the P/G lines, the maximum voltage drop is caused either by the first droop (the rise time) or by the second droop (the fall time). Design expressions to estimate the minimum on-chip decoupling capacitance required to support expected current demands based on critical line lengths are proposed. The critical length of the interconnect between the decoupling capacitor and the current load is also determined. The on-chip decoupling capacitor should be placed such that both the power supply and the current load are located within the effective radius. If this placement is not possible, the current load should be partitioned into several functional units. The on-chip decoupling capacitors should be allocated for each block, satisfying both effective radii criteria. The on-chip decoupling capacitors placed within effective radii should be allocated on the die to satisfy local transient current demands.

8. REFERENCES

- [1] *International Technology Roadmap for Semiconductors, 2004 Update*, 2004. Semiconductor Industry Association.
- [2] A. V. Mezhiba and E. G. Friedman, *Power Distribution Networks in High Speed Integrated Circuits*, Kluwer Academic Publishers, 2004.
- [3] 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, Issue 4, pp. 386–394, April 2004.
- [4] 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.
- [5] S. R. Nassif and O. Fakhouri, "Technology Trends in Power-Grid-Induced Noise," *Proceedings of the ACM International Workshop on System Level Interconnect Prediction*, pp. 55–59, April 2002.
- [6] W. D. Becker *et al.*, "Modeling, Simulation and Measurement of Mid-Frequency Simultaneous Switching Noise in Computer Systems," *IEEE Transactions on Components, Packaging, and Manufacturing Technology, Part B*, Vol. 21, Issue 2, pp. 157–163, May 1998.
- [7] M. E. Van Valkenburg, *Network Analysis*, Prentice-Hall, 1974.