

Design of an Efficient Power Delivery Network in an SoC to Enable Dynamic Power Management

Behnam Amelifard
Department of EE-Systems
University of Southern California
Los Angeles, CA
(213) 740-9481
amelifar@usc.edu

Massoud Pedram
Department of EE-Systems
University of Southern California
Los Angeles, CA
(213) 740-4458
pedram@ceng.usc.edu

ABSTRACT

Dynamic voltage scaling (DVS) is known to be one of the most efficient techniques for power reduction of integrated circuits. Efficient low voltage DC-DC conversion is a key enabler for the design of any DVS technique. In this paper we show how to design an efficient power delivery network for a complex system-on-a-chip (SoC) so as to enable dynamic power management through assignment of appropriate voltage level (and the corresponding clock frequency) to each function block in the SoC. We show that the proposed technique reduces the power loss of the power delivery network by an average of 34% while reducing its cost by an average of 8%.

Categories and Subject Descriptors

B.8.2 [Performance and Reliability]: Performance Analysis and Design Aides

General Terms

Algorithms, Design, Performance

Keywords

Low-power design, power delivery network, DC-DC converter, voltage regulator

1. INTRODUCTION

The power delivery network (PDN) is a critical design component in large designs, especially for high-performance systems [1]. A robust PDN is required to achieve a high level of system signal integrity. If improperly designed, this network could be a major source of noise, such as IR-drip, ground bounce and electromagnetic interference [2].

Emerging low-power design solutions [3] have made the design of PDN an even more challenging task. More precisely, multiple voltage domains (also known as voltage islands [4]) and dynamic voltage scaling (DVS) [5, 6] are being introduced on the System-on-a-chip (SoC) in order to minimize the overall power dissipation of the system while meeting a performance constraint. This means

This research was sponsored in part by a grant from the National Science Foundation.

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, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ISLPED '07, August 27–29, 2007, Portland, Oregon.

Copyright 2007 ACM 978-1-59593-709-4/07/0008...\$5.00.

that it is possible to have multiple relatively-small logic blocks operated at different and dynamically changing voltages based on workload monitoring. In these systems, it is required that the PDN delivers power at appropriate voltage levels to different functional blocks (FB's) while incurring the minimum power loss when the voltage level of a FB is changed in response to a change in its workload. From a system-level perspective, PDN design for a high-performance SoC comprises of three steps: (a) establishing PDN target impedance, (b) designing a proper system-level decoupling capacitance network, and (c) selecting the right voltage regulator modules.

The PDN *target impedance* can be calculated by assuming ΔV allowable ripple in the voltage supply and a 50% switching current in the rise and fall time of the processor clock [7], i.e., $Z_{target} = \Delta V / (0.5 \times I)$, where I is the current drawn by the microprocessor from the PDN. Since the current drawn by digital circuits can change suddenly with different frequencies, the target impedance should be met over a broad frequency range to guarantee the ripple on the voltage supply does not exceed the allowable value. To meet this requirement, low-frequency, mid-frequency, and high-frequency *decoupling capacitors* need to be suitably placed in the design.

Every FB is designed to operate off of some supply voltage, which is usually assumed to be constant. A *voltage regulator module* (VRM) should be used to provide this substantially constant DC output voltage regardless of changes in load current or input voltage. A dynamic programming technique was recently proposed to address the problem of optimal selection of VRM's in a power delivery network [8]. That work, however, did not address the question of VRM selection to enable dynamic voltage scaling.

In this paper we address the problem of designing an efficient PDN to support DVS. The remainder of this paper is organized as follows. Section 2 provides the background on voltage regulator modules. Our idea to efficient PDN for SoC to enable dynamic power management is presented in Section 3. Section 4 shows the simulation results, while Section 5 concludes the paper.

2. BACKGROUND

A voltage regulator module (VRM) is an electrical device designed to automatically maintain a constant voltage level, regardless of changes in input voltage or output current (This statement assumes that the load current and input voltage of the VRM are both within the specified operating range for it). The output voltage of a VRM may not be equal to the DC of the input voltage. If the output voltage of the VRM is smaller than the input voltage, the VRM is called *step-down (buck)* and if the output voltage is greater than the input voltage, it is called *step-up (boost)* [9].

Let the range of input voltages and load currents over which a regulator can maintain a target voltage level within the specified

tolerance band (e.g., 1.3V with $\pm 5\%$ ripple) be specified. The VRM's *power efficiency* may be calculated as the ratio of the power that is delivered to the load to the power that is extracted from the input source, i.e.,

$$\eta = \frac{V_{out} I_{out}}{V_{in} I_{in}} \quad (1)$$

Power efficiency is one of the most important figures of merit for a VRM and is a function of the input voltage and output current of the VRM.

Based on how voltage conversion is achieved, VRM's are classified into two main categories: *linear regulators* and *switching regulators*. A linear regulator is based on an active device, such as a BJT or a MOSFET, which operates as a voltage-controlled resistor, continuously adjusting a voltage divider network to maintain a constant output voltage. A switching regulator is a device transforming the voltage from one level to another with utilizing low-pass components such as capacitors, inductors, or transformer and switches that are in one of two states, ON or OFF. In *charge-pump* switching regulators (also known as switched-capacitors), capacitors are utilized as energy storage elements, whereas in *inductor-based* switching regulators, inductors are the energy storage components. The advantage of using a switching regulator is that the switch dissipates very little power in either of these two states and power conversion can be accomplished with minimal power loss, which equates to high power efficiency.

Each VRM has an associated cost which depends on its complexity, silicon area, and passive element costs. For example, because of their inductors, regulated inductor-based VRM's are usually the most expensive type of DC-DC converters. Linear regulators, on the other hand, are typically the least expensive ones.

3. VRM NETWORK OPTIMIZATION

In a complex SoC design, there are many functional blocks (FB's) providing various functionalities. Examples of processing elements are DSP or CPU cores. Examples of other FB's are random logic blocks, custom signal (audio or video) signal processing blocks, RF front-end, on-chip memory, and various controllers. A VRM design must meet the requirements of all FB's that are powered by it. In an SoC with DVS option, the supply voltage level of some of the FB's is dynamically adjusted in order to minimize the total power consumption while meeting the performance demands [6]. An on-chip power manager decides when to switch the *SoC power-performance state (PPS)*, where each PPS corresponds to a particular combination of voltage level (and associated clock frequency) assignments to various FB's in the SoC.

In the conventional technique to support dynamic voltage scaling for different FB's, which is depicted in Figure 1, each FB has its own VRM with multiple output voltage levels [5, 6]. The power manager selects the supply level that VRM_i provides to the FB_i.

This architecture, despite its simplicity, has several shortcomings: i) the number of VRM's used in the PDN is equal to the number of FB's i.e., when the number of FB's that can accept multiple voltage levels becomes large, the number of VRM's increases, which in turn increases the chip area and cost, ii) design of variable output voltage VRM is quite challenging and its cost is correspondingly higher than that of a fixed output voltage VRM, iii) unlike the VRM's with fixed- V_{out} where the power conversion efficiency is highly optimized for a specific output voltage level,




Figure 1: The role of VRM tree in providing appropriate voltage level for each FB. The output voltage of each VRM is changed dynamically

the power conversion efficiency of the multiple- V_{out} VRM varies as a function of the chosen V_{out} and may sometimes degrade severely from one V_{out} to next [9].

Based on these observations, we propose a new technique to address the problem of PDN design to support dynamic voltage scaling. In our technique, which is depicted in Figure 2, the PDN is composed of two layers. In the first layer of PDN, which is called the *power conversion network (PCN)*, VRM's are used to generate all voltage levels that may be needed by different FB's in the SoC design. This is accomplished by using fixed- V_{out} VRM's; so, if \mathcal{U} is the set of all voltage levels required by any FB's, then there must be at least $|\mathcal{U}|$ VRM's in the PCN. Usually this number is small since many of the FB's share the same set of allowed voltage levels. In the second layer of PDN, a *power switch network (PSN)* is used to dynamically connect the power supply terminals of each FB to the appropriate VRM output in the PCN.

In our system modeling framework, it is assumed that the transition of the system into different PPS's can be described as a time-homogenous Markov chain, and hence, PPS transitions can be captured by a stationary time-independent transition matrix $[p_{ij}]$ (c.f., Figure 3). In each state of this Markov chain, the supply voltage level of all FB's is specified. Clearly, no two states will have the same supply voltage assignments. Let π_i denote the probability of being in state i of this Markov chain. In vector $\pi = [\pi_i]$ entries π_i sum to one and satisfy

$$\pi_i = \sum_{j \in S} \pi_j p_{ji} \quad (2)$$

Additionally, for simplicity, in this section it is assumed that the

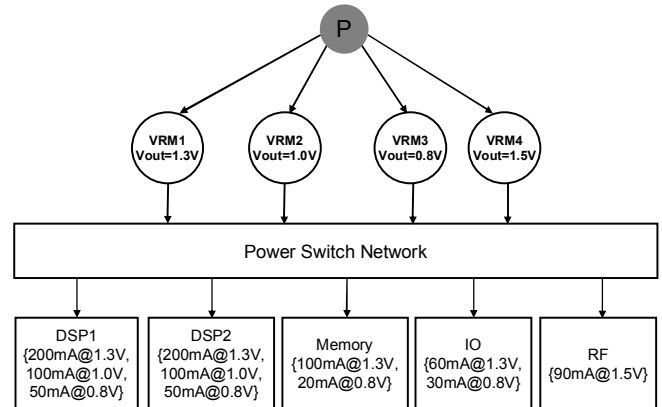


Figure 2: The proposed architecture of PDN to support dynamic voltage scaling. The output voltage of each VRM is fixed




Figure 3: Operating states and state transition of a system

current demands of every FB when it is working with each of its voltage levels is specified and is constant. In the next section it will be shown how to change the problem formulation to handle the general case when the current demands of FB's follow some probability distribution function around a mean value. Moreover, it is assumed that level shifters have been included in the SoC to enable communication among FB's operating on different supply voltages. Now, the question becomes how to design the PCN to achieve minimum power loss in the power distribution network, and how to design the PSN to make sure that all FB's receive the desired supply voltage levels.

3.1 Power Conversion Network Optimization

The PCN optimization supporting dynamic voltage scaling (PCODS) problem is defined next.

PCODS Problem

Given is:

- A library \mathcal{R} of VRM's and for each $r \in \mathcal{R}$, its cost c_r , output voltage $v_{r,out}$, the minimum and maximum input voltages $v_{r,in}^{\min}$ and $v_{r,in}^{\max}$, the maximum load current $i_{r,out}^{\max}$, and the VRM's power conversion efficiency η_r as a function of the load current and input voltage,
- A power source P , with the nominal voltage of V_P ,
- A set \mathcal{F} of FB's, and for each $f \in \mathcal{F}$, the required voltages and the corresponding current demands,
- A Markov chain model \mathcal{S} of the system, where in each state of the Markov chain the supply voltage level of each FB is specified.

The objective is to build a network of VRM's that connects P to all FB's and minimizes a weighted sum of total power consumption and total cost of the VRM's used in the PCN, i.e.,

$$V_P I_P + \lambda \sum_{r \in PCN} c_r$$

while meeting the voltage and current constraints.

In PCODS problem, λ is a parameter which defines the tradeoff between power-efficiency and cost of the PCN. For example, if $\lambda=0$, then PCODS optimizes the power efficiency while $\lambda=\infty$ results in the lowest-cost PCN.

Before giving details of how PCODS can be solved, we define the notation used in the remainder of the paper.

\mathcal{R}	Set of all VRM's, r
\mathcal{F}	Set of all FB's, f
\mathcal{S}	Set of all states of the Markov chain model of the system

\mathcal{V}_f	Set of required voltage levels by FB $f \in \mathcal{F}$
\mathcal{U}	Set of voltage levels required by all FB's; i.e., $\mathcal{U} = \bigcup_{f \in \mathcal{F}} \mathcal{V}_f = \{V_1, V_2, \dots, V_m\}$
$V_{f,s}$	Required voltage of FB $f \in \mathcal{F}$ in state $s \in \mathcal{S}$
$I_{f,s}$	Required current of FB $f \in \mathcal{F}$ in state $s \in \mathcal{S}$
$I_{f,v}$	Required current of FB $f \in \mathcal{F}$ when its required voltage level is $v \in \mathcal{U}$ ($I_{f,v} = I_{f,s} : V_{f,s} = v$)
$\eta_r(V, I)$	Power conversion efficiency of regulator $r \in \mathcal{R}$, with the input voltage V and output current I
$I_{r,s}^{in}$	Input current of regulator r in state $s \in \mathcal{S}$
$I_{avg,r}$	Average input current of regulator r over all states

We assume that if a FB requires the same voltage V in two different states, it is always powered up by an identical VRM. This assumption implies that the number of power switches in PSN to deliver power to FB $f \in \mathcal{F}$ is exactly $|\mathcal{V}_f|$ and hence reduces not only the complexity of PSN, but also its energy loss.

It should be noted that the power delivered to the FB's is independent of the topology of PCN and can be calculated as,

$$P_{FBs} = \sum_{f \in \mathcal{F}} \sum_{s \in \mathcal{S}} \pi_s V_{f,s} I_{f,s} \quad (3)$$

The voltage domain \mathcal{D}_i is defined as the set of all FB's that require voltage level V_i in some state, i.e.,

$$\mathcal{D}_i = \{f \in \mathcal{F} : V_i \in \mathcal{V}_f\} \quad (4)$$

Since each FB may have more than one voltage level, \mathcal{D}_i 's may be overlapping. Assume that the topology of the VRM tree delivering power to \mathcal{D}_i is known. In this case, when the system is in state s , the output current of a VRM r that delivers power to a subset $\mathcal{D}_i^j \subseteq \mathcal{D}_i$ can be computed as,

$$I_{r,s}^{out} = \sum_{f \in \mathcal{D}_i^j, V_{f,s} = V_i} I_{f,s} \quad (5)$$

Therefore, the input current of VRM r in state s is obtained as,

$$I_{r,s}^{in} = \frac{V_r \times I_{r,s}^{out}}{V_P \times \eta_r(V_P, I_{r,s}^{out})} \quad (6)$$

and the average input current of r which is drawn from the power supply is,

$$I_{avg,r} = \sum_{s \in \mathcal{S}} \pi_s I_{r,s}^{in} \quad (7)$$

The average current drawn from the power supply by \mathcal{D}_i is computed as,

$$I_{avg}(\mathcal{D}_i) = \sum_{r \in \mathcal{R}_i} I_{avg,r} \quad (8)$$

where \mathcal{R}_i is the set of all VRM's used to power up \mathcal{D}_i . The total cost of the VRM's used in this topology to deliver power to \mathcal{D}_i is,

$$C_{\mathcal{D}_i} = \sum_{r \in \mathcal{R}_i} c_r \quad (9)$$

Therefore, the average current drawn from the power supply by this PCN and the total cost of VRM's in the PCN can be written as,

$$I_{avg} = \sum_i I_{avg}(\mathcal{D}_i) \quad (10)$$

$$C_{PCN} = \sum_i C_{\mathcal{D}_i} \quad (11)$$

To deliver power to FB's in each \mathcal{D}_i , different options are available (c.f., Figure 4 for a pictorial elaboration). In the first option, which is the lowest-cost one, only one VRM is used to deliver power to all FB's in each \mathcal{D}_i . The other option is to use one VRM per FB. The drawback of this solution is that the number of VRM's

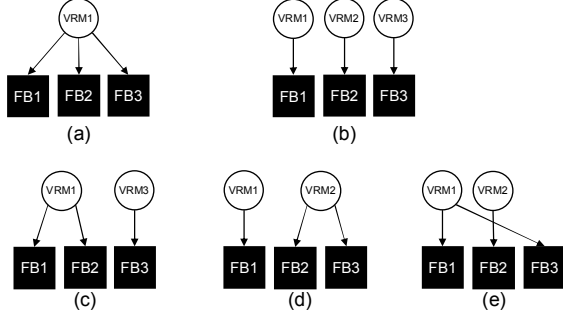


Figure 4: Different options for delivering power to three FB's which require the same voltage at some states. The output voltages of all VRM's are the same.

increases with the number of FB's. Because of the non-monotone dependency of power conversion efficiency on the delivered output current, neither solution may be the best from the power-efficiency point of view and a design in between the two extremes may be the best one. Furthermore, because objective function in the general formulation of the PCODS problem is a weighted sum of the power consumption and the cost of the PCN, by enumerating other solutions a better tradeoff between power-efficiency and cost may be achieved. Therefore, all possible solutions should be enumerated when searching for the optimal VRM assignment to \mathcal{D}_i .

Definition 1: A partition of set \mathcal{D}_i is a collection of disjoint subsets of \mathcal{D}_i whose union is \mathcal{D}_i . Each of these subsets is called a *part*.

The number of partitions of a set with n elements is the n 'th Bell number which can be computed from the following recurrence [10],

$$B_n = \sum_{k=0}^{n-1} \binom{n-1}{k} B_k, B_0 = 1 \quad (12)$$

Definition 2: In a partition of \mathcal{D}_i the *required voltage* of each part is V_i . The *current demand* of a part in a state is the summation of the current demands of all FB's in that part in the specified state.

Definition 3: A *valid VRM assignment* to a partition of \mathcal{D}_i is the assignment of one VRM to each part such that the constraints of each VRM are satisfied, i.e., for each VRM r the input voltage of VRM is between $v_{r,in}^{\min}$ and $v_{r,in}^{\max}$, the required voltage of the part is $v_{r,out}$, and the maximum current demand of the part over all states is lower than $i_{r,out}^{\max}$.

Definition 4: An *optimum VRM assignment* to a partition of \mathcal{D}_i such as $\{\mathcal{D}_i^1, \dots, \mathcal{D}_i^n\}$ is a valid VRM assignment which minimizes $\sum_j V_P I_{avg,j} + \lambda \sum_j c_j$, where $I_{avg,j}$ and c_j are the input current and associated cost of designated VRM to part \mathcal{D}_i^j , respectively.

Theorem 1: A valid VRM assignment to a partition of \mathcal{D}_i is optimum, if and only if in each of its parts such as \mathcal{D}_i^j , $V_P I_{avg,j} + \lambda c_j$ is minimized.

Proof: Assume \mathcal{D}_i is partitioned into n nonempty subsets such as $\{\mathcal{D}_i^1, \dots, \mathcal{D}_i^n\}$. Each valid VRM assignments to a part is shown as a pair of input current of the corresponding VRM and its associated cost, i.e., (I_{avg}, c) . The set of all valid VRM assignments to part \mathcal{D}_i^j is shown as $\mathcal{Z}_j = \{(I_{avg,j}, c_j)\}$. Optimum VRM assignment to partition \mathcal{D}_i is the selection of one tuple $(I_{avg,j}, c_j)$ from each \mathcal{Z}_j such that $\sum_j V_P I_{avg,j} + \lambda \sum_j c_j$ is minimized. It can be seen that

$\sum_j V_P I_{avg,j} + \lambda \sum_j c_j$ is minimized if and only if for each tuple $(I_{avg,j}, c_j)$, the value of $V_P I_{avg,j} + \lambda c_j$ is minimum over all tuples in \mathcal{Z}_j . ■

The result of Theorem 1 is that to find the optimum VRM assignment to set \mathcal{D}_i , all partitions of \mathcal{D}_i should be enumerated. In each partition, the best VRM r that satisfies the constraints and minimizes $V_P I_{avg} + \lambda c$ for every part is found. The partition that results in the minimum $\sum_j V_P I_{avg,j} + \lambda \sum_j c_j$ is the optimum one.

Based on the above discussion, Figure 5 shows *optPCN* algorithm to solve PCODS problem. Basically it starts by constructing \mathcal{D}_i sets and for each \mathcal{D}_i it finds the best VRM assignment by using Theorem 1.

Theorem 2: The *optPCN* algorithm described in Figure 5 finds the optimum solution to the PCODS problem.

Proof: The optimality of *optPCN* algorithm is immediate from Theorem 1 and the fact that for each \mathcal{D}_i , all partitions are enumerated. ■

Theorem 3: The worst-case running time of *optPCN* algorithm is $O(|\mathcal{R}| \cdot |\mathcal{S}| \cdot |\mathcal{F}| \cdot B_{|\mathcal{F}+1})$, where $|\mathcal{R}|$, $|\mathcal{S}|$, and $|\mathcal{F}|$ are the cardinalities of corresponding sets. The worst case happens when for each $f \in \mathcal{F}$, the set of required voltage levels is equal to \mathcal{U} , i.e., $\mathcal{V}_f = \mathcal{U}$.

Proof: It is removed for brevity. ■

Algorithm *optPCN*($\mathcal{R}, \mathcal{F}, \mathcal{S}, V_P$)

Begin

1. **For** each $V_i \in \mathcal{U} = \{V_1, \dots, V_m\}$
 2. $\mathcal{D}_i = \{f \in \mathcal{F} : V_i \in \mathcal{V}_f\}$
 3. $\psi(V) = \text{sub-optPCN}(\mathcal{R}, \mathcal{F}, \mathcal{S}, V_P, V_i, \mathcal{D}_i)$
 4. **End**
- End**

Algorithm *sub-optPCN*($\mathcal{R}, \mathcal{F}, \mathcal{S}, V_P, V_i, \mathcal{D}_i$)

Begin

1. $\text{optCost} = \infty$
2. $\text{optVRMs} = \{\}$
3. **For** each non-empty partition of \mathcal{D}_i such as $\{\mathcal{D}_i^1, \dots, \mathcal{D}_i^n\}$
4. **For** each $\mathcal{D}_i^j, 1 \leq j \leq n$
5. Select the best VRM r that minimizes $V_P I_{avg,r} + \lambda c_r$
6. $\text{cost}_j = V_P I_{avg,r} + \lambda c_r$
7. $\text{VRMs}_j = r$
8. **End**
9. $\text{newCost} = \sum_j \text{cost}_j$
10. **If** ($\text{newCost} < \text{optCost}$)
11. $\text{optCost} = \text{newCost}$
12. $\text{optVRMs} = \{\text{VRMs}_j\}$
13. **End**
14. **End**
15. **Return** ($\text{optCost}, \text{optVRMs}$)

End

Figure 5: The *optPCN* algorithm for solving PCODS




Figure 6: Approximating the continuous distribution with a discrete one

From Theorem 3, one can see that *optPCN* algorithm has exponential complexity in the number of FB's; however, since the number of FB's is small, in practice the runtime of the algorithm is quite reasonable.

3.1.1 Effect of non-constant current

In the formulation of PCODS problem, it is assumed that the current demand of each FB is a constant value independent of the system PPS. In this section it is shown how to modify the problem formulation to handle the case when the current demands of various FB's follow some probability density function (pdf).

We assume that the current demands of different FB's can be modeled as independent Gaussian distribution functions (the case that the demands follow some other probability distribution function can be addressed in a similar manner). In this case, because the output current of a VRM which is connected to a number of FB's is a sum of independent Gaussian random variables (c.f., Equation (5)), it will also be a Gaussian random variable, whose mean and variance respectively are the sum of means and sum of variances of the current demand distributions in the corresponding FB's. This continuous-time random variable is approximated with a discrete-time random variable function which has the probability $\Pr(i)$ in interval $[I_{\min} + i \times \Delta I, I_{\min} + (i+1) \times \Delta I)$ (for $0 \leq i < (I_{\max} - I_{\min}) / \Delta I$) as shown in Figure 6.

Since the efficiency of the VRM and hence its input current are functions of the output current, Equation (6) should be modified to account for this dependency,

$$I_{r,s}^{in} = \sum_{i=0}^L \Pr(i) \frac{V_r \times (I_{\min} + i \times \Delta I)}{V_S \times \eta_r(V_P, I_{\min} + i \times \Delta I)} \quad (13)$$

where $L = (I_{\max} - I_{\min}) / \Delta I - 1$. Selecting a smaller value for ΔI results in a better approximation for input current of the VRM, but also increases the algorithm runtime.

3.2 Power Switch Network Optimization

Power switch network (PSN) performs the function of switching the supply voltage level of the FB's when a new PPS is commanded by the power manager. Figure 7 depicts a PSN for delivering three different voltage levels to an FB. The switches in the PSN are controlled by a *power switch controller* (PSC) which is zero-hot coded, i.e., at any given time only one of its outputs is zero, and hence, only one PMOS transistors in ON.

The number of PMOS transistors needed for each FB f in the PSN is $|\mathcal{V}_f|$. The PMOS transistor which is required to deliver voltage level $v \in \mathcal{V}_f$ to an $f \in \mathcal{F}$ and its width are respectively denoted as $M_{f,v}$ and $W_{f,v}$. This PMOS transistor should be large enough so that the voltage-drop between its drain and source does not exceed a tolerable value.

In the steady state, when FB f is supplied with $v \in \mathcal{V}_f$, the current that flows through the ON PMOS transistor $M_{f,v}$ is the current




Figure 7: A PSN for delivering three different voltage levels to an FB

demand of f at voltage v , i.e., $I_{f,v}$. Since this transistor is in triode region, its current can be derived from the alpha-power model [11] as,

$$I_{f,v} = I_{M_{f,v}} = k \frac{W_{f,v}}{L_{eff}} \left(\frac{V_{GS} - V_{th}}{v - V_{th}} \right)^{\alpha/2} V_{DS} \quad (14)$$

where L_{eff} is the effective length of the transistor, V_{GS} , V_{DS} , and V_{th} are the gate-to-source, drain-to-source, and threshold voltage of the transistor, respectively. Note that k and α are technology parameters with α being 2 for long channel devices and about 1.3 for short channel devices. Now, if the maximum tolerable voltage-drop at the supply of the FB is ΔV , the minimum required width for $W_{f,v}$ will be computed as,

$$W_{f,v}^{\min} = \frac{I_{f,v} L_{eff}}{k \Delta V} \quad (15)$$

3.2.1 PSN Power Consumption

When the state of the system changes from PPS i to j , some energy is consumed to turn ON/OFF some of the PMOS switches. Assume that the power manager changes the state of the system at regular time intervals with a frequency of f_{PM} . If C_{PMOS} is the total capacitance which is charged or discharged during this transition, then the power consumption for this transition is calculated from

$$P_{dyn,i \rightarrow j} = p_{i \rightarrow j} V_{DD}^2 f_{PM} C_{PMOS} \quad (16)$$

where $p_{i \rightarrow j}$ denotes the transition probability from PPS i to j which can be computed as,

$$p_{i \rightarrow j} = \pi_i p_{ij} \quad (17)$$

So, the power consumption of the PMOS switches is calculated as

$$P_{overhead} = \sum_{i,j} \left(\frac{1}{2} p_{i \rightarrow j} V_{DD}^2 f_{PM} \left(\sum_{f: \mathcal{V}_f \neq \mathcal{V}_{f,j}} \left(C_{f,\mathcal{V}_f,i} + C_{f,\mathcal{V}_f,j} \right) \right) \right) \quad (18)$$

where $C_{f,v}$ is the input capacitance of $M_{f,v}$, i.e., $C_{f,v} = W_{f,v} L C_{ox}$.

Equation (18) is the power consumption overhead of our solution compared to the conventional one, where one multiple-output VRM is used for each FB to provide it with appropriate voltage levels.

4. SIMULATION RESULTS

The algorithms described earlier in this paper have been implemented in C++ and evaluated on a set of test-benches. A collection of thirty DC-DC commercially available regulators from Texas Instruments and National Semiconductors were chosen to create the library of VRM's. The power conversion efficiency of each VRM was modeled as a piecewise-linear function of input voltage and output current based on the data sheets for the VRM. The cost of each VRM was assumed to be its dollar cost for a

1000-unit purchase. Note that we did not have access to the efficiency curves and cost of the unpackaged DC-DC converters.

We performed two experiments to compare the performance of the proposed technique with the conventional VRM assignment to support dynamic voltage scaling in a system. In the first experiment, we used *optPCN* algorithm with $\lambda=0$ to find the most power-efficient PCN based on our solution. The best multiple-output VRM assignment to minimize the power consumption of the system based on the conventional solution was also generated for comparison purposes. The results of this experiment are reported in Table 1, where the first column gives the name of the test-bench (Details of the first test-bench are provided in Figure 8.), the second column gives the number of FB's in the problem, and the third column gives the number of states in the Markov chain model of the system. Column 4 and 5 show PDN power loss and cost reduction in the proposed solution compared to those of the conventional solution (power loss in the PDN is the difference between the power delivered to FB's and the power drawn from the power source P). Finally, the last column shows the runtime of *optPCN* algorithm for finding the optimal set of VRM in the PCN. From Table 1, one can see that the proposed technique reduces the power loss of PDN by an average of 34%. Additionally, in most cases it also reduces the PDN cost. The average PDN cost reduction is 8%. Finally, from Table 1 one can see that the runtime of *optPCN* algorithm is quite reasonable.

In the second experiment, we studied the tradeoff between the power-efficiency of the PDN and its cost. More precisely, in addition to designing the optimal PCN for $\lambda=0$ by running *optPCN* algorithm, the algorithm was invoked for other values of λ for which the PCN power loss does not increase beyond 10% of its optimal value. The cost reduction of the PDN for this set of test-benches is reported in Table 2. It is seen that on average by allowing about 8.6% increase in the PDN power loss, the cost of PDN can be lowered by 47%.

5. CONCLUSION


In this paper we presented a new technique to design an efficient power delivery network for systems with dynamic voltage scaling capability. In this technique, the PDN is composed of two layers: PCN and PSN. In PCN, fixed- V_{out} VRM's are used to generate all voltage levels that may be needed by different FB's in the system. PSN is used to dynamically connect the power supply terminals of each FB to the appropriate VRM output in the PCN. We showed that this technique not only reduces the cost of the power conversion network, but also results in a more power-efficient power delivery network. We further described an algorithm to select the best VRM's to achieve a design target in the new PDN. By means of simulation results, it was demonstrated that the proposed technique reduces the power loss of PDN by an average of 34% while reducing its cost by an average of 8%.

Table 1: Power and cost reduction of PDN in the proposed technique compared to those of the conventional technique

	$ \mathcal{F} $	$ \mathcal{S} $	PDN Power Reduction (%)	PDN Cost Reduction (%)	Runtime (sec)
C1	5	4	38.5	1.1	<1
C2	6	4	40.4	5.0	<1
C3	8	5	34.2	-2.8	<1
C4	10	10	30.1	29.7	13
C5	12	10	27.9	8.1	70

Table 2: Trading off power for cost of PDN in the proposed technique

	PDN Power Increase (%)	PDN Cost Reduction (%)
C1	10.0	53.0
C2	4.3	46.9
C3	8.9	57.9
C4	9.6	26.1
C5	10.0	52.9



S1: $\{V_{DSP1}=1.3, V_{DSP2}=1.3, V_{MEM}=1.3, V_{IO}=1.3, V_{RF}=1.5\}$
S2: $\{V_{DSP1}=1.0, V_{DSP2}=1.3, V_{MEM}=1.3, V_{IO}=1.3, V_{RF}=1.5\}$
S3: $\{V_{DSP1}=0.8, V_{DSP2}=1.0, V_{MEM}=1.3, V_{IO}=0.8, V_{RF}=1.5\}$
S4: $\{V_{DSP1}=0.8, V_{DSP2}=0.8, V_{MEM}=0.8, V_{IO}=0.8, V_{RF}=1.5\}$

Figure 8: Test-bench C1. The current demands of FB's are similar to those in Figure 1.

6. REFERENCES

- [1] D. Blaauw, R. Panda, and R. Chaudhry, "Design and analysis of power distribution networks," in *Design of High-Performance Microprocessor Circuits*, A. Chandrakasan, W. J. Bowhill, and F. Fox, Eds. Piscataway, NJ: IEEE, 2001, pp. 499-522.
- [2] S. Chun, "Methodologies for Modeling Simultaneous Switching Noise in Multi-Layered Packages and Boards," Ph.D. dissertation, Georgia Institute of Technology, 2002.
- [3] M. Pedram, "Power minimization in IC design: principles and applications," *ACM Transactions on Design Automation of Electronic Systems*, vol. 1, no. 1, Jan. 1996, pp. 3-56.
- [4] D. E. Lackey, P. S. Zuchowski, T. R. Bednar, et al., "Managing power and performance for System-on-Chip designs using Voltage Islands," in *Proc. of International Conference on Computer Aided Design*, 2002, pp. 195-202.
- [5] T. D. Burd and R. W. Brodersen, "Design issues for dynamic voltage scaling," in *Proc. of International Symposium on Low Power Electronics and Design*, 2000, pp. 9-14.
- [6] K. J. Nowka, G. D. Carpenter, E. W. MacDonald, et al., "A 32-bit PowerPC system-on-a-chip with support for dynamic voltage scaling and dynamic frequency scaling," *IEEE Journal of Solid-State Circuits*, vol. 37, no. 11, Nov. 2002, pp. 1441-1447.
- [7] L. Smith, R. Anderson, D. Forehand, et al., "Power distribution system design methodology and capacitor selection for modern CMOS technology," *IEEE Transactions on Advanced Packaging*, vol. 22, no. 3, Aug. 1999, pp. 284-291.
- [8] B. Amelifard and M. Pedram, "Optimal selection of voltage regulator modules in a power delivery network," in *Proc. of Design Automation Conference*, 2007, pp. 168-173.
- [9] A. Stratakos, "High-efficiency low-voltage DC-DC conversion for portable applications," Ph.D. dissertation, University of California, Berkeley, 1998.
- [10] R. L. Graham, D. E. Knuth, and O. Patashnik, *Concrete Mathematics*. MA: Addison-Wesley, 1990.
- [11] T. Sakurai and A. R. Newton, "A simple MOSFET model for circuit analysis," *IEEE Transactions on Electron Devices*, vol. 38, no. 4, Apr. 1991, pp. 887-894.