Optimal Selection of Voltage Regulator Modules in a Power Delivery Network

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
High efficiency low voltage DC-DC conversion is a key enabler to the design of power-efficient integrated circuits. Typically a star configuration of the DC-DC converters, where only one converter resides between the source and each load, is used to deliver currents with appropriate voltage levels to different loads in the circuit. In this paper we show that using a tree topology of suitably chosen voltage regulators between the power source and loads yields higher power efficiency in the power delivery network. We formulize the problem of selecting the best set of regulators in a tree topology as a dynamic program and efficiently solve it. Experimental results demonstrate the efficacy of proposed problem formulation and solution.

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 International Technology Roadmap for Semiconductors (ITRS) has projected an increase in the power consumption of microprocessors for future technology nodes [1]. For example, for complex designs done in 2007 with a feature size of 65nm and a supply voltage of 1.1V, the power dissipation is 104 Watts. The power delivery network (PDN) provides the power supply to the processor. If improperly designed, this network can be a major source of noise, such as ground bounce and IR drop [2].

The power delivery network is a critical design component in large designs, especially for high-speed electronic systems [3]. A robust PDN is required to achieve a high level of signal integrity. In particular, PDN design comprises of three steps:

- Establishing PDN target impedance,
- Designing a proper system-level decoupling network,
- Selecting the right voltage regulator modules.

A methodology for designing a good PDN is to define a target impedance for the network that should be met over a broad frequency band [4]. This parameter can be computed by assuming a 5% allowable ripple in the voltage supply and a 50% switching current in the rise and fall time of the processor clock. The target impedance can then be calculated as:

$$Z_{target} = \frac{0.05 \times V_{dd}}{50\% \times I} = 0.1 \times \frac{V_{dd}}{I} \quad (1)$$

where V_{dd} is the core voltage of the processor and I is the current drawn by the microprocessor from the PDN. For the 65nm node, I=104/1.1=94A and Z_{target}=1.2mΩ. The decoupling capacitors play an important role in the PDN as they act as charge reservoirs for the switching circuits. The PDN target impedance has to be met over a broad frequency band; low frequency, mid-frequency and high frequency capacitors need to be suitably placed to meet this requirement. It is difficult to provide sufficient decoupling in the mid-frequency range of 200-300 MHz to 2-3 GHz. This presents a challenge to designers if they are to meet the impedance requirement over the entire frequency range [5]. This problem however falls outside the scope of the present paper.

Every electronic circuit is designed to operate off of some supply voltage, which is usually assumed to be constant. A voltage regulator module (VRM) provides this substantially constant DC output voltage regardless of changes in load current or input voltage (this statement assumes that the load current and input voltage are within the specified operating range for the part). A switching power supply is a device transforming the voltage from one level to another. Typically voltage is taken from the AC power lines or unregulated DC power lines and transformed to the regulated DC levels that logic circuits require. A switching-mode power supply (SMPS) is a power supply that provides the power supply function through low-loss components such as capacitors, inductors, and transformers -- and the use of switches that are in one of two states, ON or OFF. The advantage 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 efficiency. Usually a SMPS operates in a closed loop system to regulate the power supply output, for example through pulse-width modulation (PWM) or pulse-frequency modulation (PFM).

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., 5V with +/- 2% ripple) be specified. The

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

DAC 2007, June 4–8, 2007, San Diego, California, USA

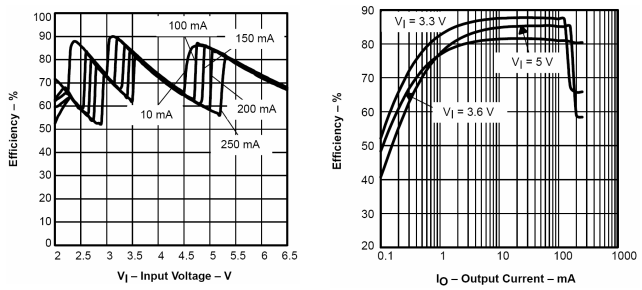

Figure 1: The efficiency of TPS60503 as a function of input voltage and output current [6].

regulator'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.,

$$\eta = \frac{V_{out} I_{out}}{V_{in} I_{in}} \quad (2)$$

Power efficiency is one of the most important figures of merit for a voltage regulator and is a function of the input voltage and output current of the VRM. Figure 1 shows the efficiency of a commercial VRM as a function of input voltage and output current.

For DC-DC type conversion, there are many design choices. One option is to use *regulated charge-pump* (*switched capacitor*) DC-DC converter that utilize capacitors as energy storage elements. They are often used when the load current demand is rather low (in Amperes or less). Regulation is achieved by sensing the output voltage through a resistor divider and modulating the charge pump output current based on an error signal. The other option is to use a *regulated inductor-based* (*switched-mode*) DC-DC converter, which utilize inductors as energy storage elements. These regulators are often used when the load current demand is high (in tens or even hundreds of Amperes). Finally, one may use a *linear regulator* (and its most efficient form a low-dropout regulator, LDO), which operates by using a voltage-controlled current source to force a fixed voltage to appear at the regulator output terminal. The control circuitry must monitor (sense) the output voltage, and adjust the current source (as required by the load) to hold the output voltage at the desired value. The design limit of the current source defines the maximum load current the regulator can source and still maintain regulation. The dropout voltage of a linear regulator is defined as the minimum voltage drop required across the regulator to maintain output voltage regulation. The lower the dropout voltage is, the higher the power efficiency of the linear regulator is since the maximum power delivered to the load is simply $(V_{IN}-V_{dropout})\times I_{load}$ whereas the power extracted from the input source is $V_{IN}\times(I_{load}+I_{quies})$. Here, I_{quies} denotes the quiescent current in the internal circuitry of the LDO. To have a high efficiency LDO regulator, the dropout voltage and the quiescent current must be minimized. In addition, the voltage difference between input and output must be minimized since the internal power dissipation of LDO regulators, which is $(V_{OUT}-V_{IN})\times I_{load}$, accounts for the loss of power efficiency.

In recent years, PWM DC-DC converters integrated in standard foundry-available digital CMOS processes have been demonstrated. Although analog components, such as a bandgap voltage reference, amplifiers, and oscillators, are required to implement the PWM and/or PFM control functions, power dissipation due to digital logic is becoming increasingly important to the overall power budget [7].

Each IC specifies its voltage regulator configuration in its datasheets or comes with a companion document that defines the power delivery feature set necessary to support that IC within a larger electronic system. For example, the Intel's VRM version 10.2 describes the Intel® processors' V_{cc} power delivery requirements for desktop computer systems using socket 478. This includes design recommendations for DC-DC regulators which convert the 12 V supply to the processor consumable V_{cc} voltage along with specific feature set implementation such as thermal monitoring and Dynamic Voltage Identification.

In a large PCB design or equivalently 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 or interface blocks, MPEG encoder/decoder blocks, RF front-end, on-chip memory, and various controllers. The V_{cc} regulator design on a specific platform (PCB or SoC) must meet the specifications of all FB's supported in that platform.

Another low power design trend is emerging that makes the design of the *VRM tree*[1] even more important. More precisely, multiple voltage domains are being introduced on the same SoC in order to meet a performance constraint while minimizing the overall power dissipation of the system. This means that it is possible to have multiple relatively-small logic blocks operated at different, yet fixed, voltages [8] (the question of VRM tree design to support dynamic voltage scaling based on workload monitoring falls outside the scope of the present paper). This is also known as the multiple voltage island approach [9].

Figure 2 depicts the role of the regulators in providing appropriate voltage levels to different FB's on a single chip.

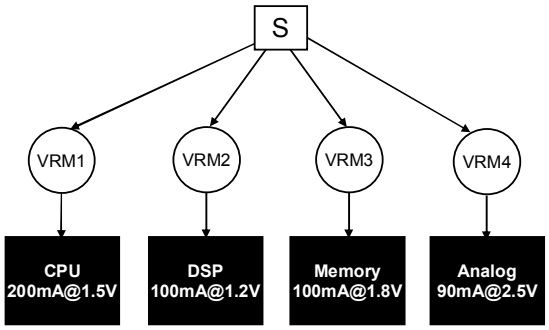

Figure 2: The role of VRM tree in providing appropriate voltage level for each FB.

2. VRM TREE OPTIMIZATION FOR MINIMUM POWER LOSS

The VRM tree optimization (RMTO) problem is defined next.

RMTO Problem: Given is

- A library $\mathcal{R}$ of VRM's and for each $r \in \mathcal{R}$, its output voltage $v_{r,out}$, the minimum and maximum input voltages $v_{r,in}^{\min}$ and $v_{r,in}^{\max}$, the maximum load current $\iota_{r,out}^{\max}$, and its efficiency η_r as a function of load current and input voltage,
- A power source S, with the nominal voltage of V_S,
- A set $\mathcal{L}$ of N loads, and for each $l \in \mathcal{L}$ its required voltage V_l and average current demand I_l.

[1] The graph representation of the VRM network will have a tree structure, that is, no VRM can be driven by more than one other VRM.

The goal is to build a tree topology of VRM's that connects S to all loads and minimizes the PDN power loss from the power source to the loads while meeting the voltage and current constraints.[2]

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

$$P_{FBs} = \sum_{l \in \mathcal{L}} V_l I_l \quad (3)$$

Therefore, to minimize the power loss in the PDN from the power source to the loads, one needs to minimize the power drawn from the power supply. Given that the voltage of the power supply is assumed to be fixed, the objective of RMTO problem is to minimize the current drawn from the power supply. We assume that each VRM can provide only one output voltage (multi-output VRM's are considered as multiple VRM's, each with its own fixed voltage output).

Although the problem definition does not put any constraints on the depth of the VRM tree that drives the loads, in practice, such a constraint is useful. The reason is that utilizing a VRM tree with a large number of internal levels tends to increase the number of regulators, which in turn increases their cost and chip area overhead with little (if any) benefit in terms of improving the power efficiency of the PDN. For this reason, in this work, we only consider up to two levels of regulators in the VRM tree, i.e., the (node) depth of the tree is 4, with one corresponding to the power source, one corresponding to the loads and up to two internal levels dedicated to VRM's. Our solution, however, can be easily extended to handle VRM trees with higher depth.

To improve the efficiency of our solution technique by implicitly considering a large class of tree topologies under one class representative, it is convenient to introduce an *ideal VRM* whose efficiency is 100% and whose output voltage and thus output current are equal to its input voltage and current, respectively. This ideal VRM (really a lossless buffer) is added to library $\mathcal{R}$ of VRM's. Note that ideal VRM's are inserted on every path from the tree root to a leaf node in the tree so that the logical depth of each such path is exactly four (c.f. Figure 3).

Definition: A VRM satisfies *monotone input current property* if its input current is a monotone increasing function of its output current independent of the input voltage.

Notice that this property may hold in spite of the non-monotone power efficiency characteristics for a VRM. This is because of the way that power efficiency is defined and its relation to input and output voltages and currents.

If the tree topology is *fixed* (-F option) and the *monotone input current property* holds for all VRM's in the library (-M option), then the selection of the appropriate regulator for each node can be done optimally by using *dynamic programming* starting from the leaf nodes. This algorithm, called RMTO-FM, is detailed in Figure 4. Before providing details, we introduce some notation and definitions.

$\mathcal{R}$: Set of all VRM's including the ideal VRM
$\mathcal{L}$: Set of all loads
$\mathcal{U}$: Set of all output voltages of the VRM's
$\mathcal{V}_n$: Set of candidate input voltages for node n
$\mathcal{V}_{n,r}$: Set of candidate input voltages for n when r is the VRM of n
$\mathcal{C}_n$: Set of candidate VRM's for internal node n of the tree
T : Topology of VRM tree
$\pi(n)$: Optimum VRM selection for node n
L_i : Set of all level i internal nodes, i=1,2
V_l and I_l : Voltage level and current demand of FB$_l$
$v_{r,out}$: Output voltage level of regulator r
$v_{r,in}^{\min}$ and $v_{r,in}^{\max}$: Minimum and maximum input voltage levels of regulator r
$i_{c_n,out}^{\max}$: Maximum output current of regulator r
$V_{out}(n)$: Output voltage of a node n
$I_{out,r}(n)$ and $I_{in,r}(n)$: Output and input current of node n given that regulator r is assigned to this node
$f_r(v_{in}, i_{out})$: Efficiency of regulator r as a function of its input voltage v_{in} and output current i_{out}
$\Psi_n(v_{in})$: One dimensional table in node n with the key v_{in} and the value of input current of node.

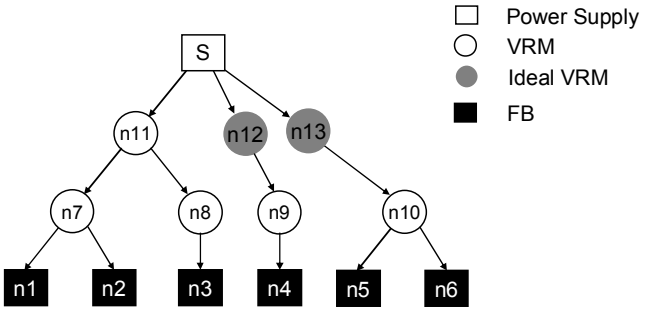

Figure 3: A VRM tree after inserting ideal VRM's.

The algorithm starts with the nodes in the second internal level of the tree T. If any such node is connected to two FB's with different input voltage requirements, then the tree will not be a feasible VRM tree (a precise definition is provided later) and the algorithm terminates; otherwise, the output current of the node is calculated as the sum of the current demands of all leaf nodes (FB's) that are connected to it. Next all candidate VRM's with compatible output current and voltage are evaluated. Since the input voltage of the node is not known at this time, the power efficiency of each candidate VRM for the node in question cannot be calculated directly. Based on the fact that this second-level node is driven by any first-level VRM node, all voltage values in $\mathcal{U}$ must be enumerated. Thus the power efficiency of the candidate second-level VRM is obtained from the efficiency curves for each regulator. This information is then used to compute the input current of the second-level node as the minimum of the input currents of the candidate VRM's which take the specific input voltage level for the second-level node. The calculated input current is stored in a one dimensional table with the key set to the input voltage of the second-level node and the value set to the input current of that same node.

The first-level nodes are visited next. For each such node n, all candidate output voltages $v_{out}(n)$ (defined as the voltages in the intersection of all $\mathcal{V}_m$'s, where m denotes a fanout of n) are considered. Next a set of output voltages are identified where each of these output voltages show up in every input current vs. input voltage table stored at each fanout of n. For every such output voltage, the sum of the input currents of the driven second-level nodes is computed and set as the target output current of the first-level node. Next based on the output current of that first-level node

[2] In this paper, we focus on this RMTO problem statement. An interesting variant of the problem, which we do not address here, may be defined as follows. Given a cost δ_r associated with each regulator r, minimize the power loss in the PDN while ensuring that the total cost of the VRM tree does not exceed a cost budget.

RMTO-FM ($\mathcal{R}$, $\mathcal{L}$, T, V_S)
Begin
1. **For** each node $n \in L_2$
2. $V_{out}(n) = V_l : l \in FO(n)$
3. $I_{out}(n) = \sum_{l \in FO(n)} I_l$
4. $\mathcal{C}_n = \{r \in \mathcal{R} \mid v_{r,out} = V_{out}(n)\}$
5. **For** each $r \in \mathcal{C}_n$
6. $\mathcal{V}_{n,r} = \{u \in \mathcal{U} \mid v_{r,in}^{\min} \le u \le v_{r,in}^{\max}, I_{out}(n) \le i_{r,out}^{\max}\}$
7. **For** each $v_{in}(n) \in \mathcal{V}_n$
8. $\eta_r = f(v_{in}(n), I_{out}(n))$
9. $I_{in,r}(v_{in}(n)) = \dfrac{v_{out}(n) \times I_{out}(n)}{v_{in}(n) \times \eta_r}$
10. **End**
11. **End**
12. $\mathcal{V}_n = \bigcup_{r \in \mathcal{C}_n} \mathcal{V}_{n,r}$
13. **For each** $v_{in} \in \mathcal{V}_n$
14. $\Psi_n(v_{in}) = \min_{r: v_{in}(n) \in \mathcal{V}_{n,r}} (I_{in,r}(v_{in}(n)))$
15. **End**
16. **End**
17. **For** each node $n \in L_1$
18. $\mathcal{C}_n = \{r \in \mathcal{R} \mid v_{r,out} \in \bigcap_{i \in FO(n)} \mathcal{V}_i,\ v_{r,in}^{\min} \le V_S \le v_{r,in}^{\max}\}$;
19. **For** each $r \in \mathcal{C}_n$
20. $V_{out}(n) = v_{r,out}$
21. $I_{out,r}(n) = \sum_{m \in FO(n)} \Psi_m(V_{out}(n))$
22. **End**
23. $(\pi(n), \pi(m) : m \in FO(n)) = \arg\min_{r \in \mathcal{C}_n} (i_{out,r}(n))$
24. **End**
End

Figure 4: RMTO-FM algorithm for VRM tree optimization when tree topology is fixed.

and the known input voltage of the same node (which is the same as the output voltage of the power source for the VRM tree), the optimum VRM assignment for the first-level node is determined by enumerating all possible VRM matches at that node, i.e., a VRM assignment is chosen that minimizes the input current of the first-level node (and hence the output current demand on the power source along the edge that leads to that node) while providing the output current needed by the driven second-level nodes under the selected output voltage assignment for the first-level node. The optimal solution of VRM tree problem when the tree topology may be varied (-V option) is found by enumerating all *feasible* trees.

Definition: A VRM tree topology is *feasible* when (a) it has an exact depth of 4, i.e., every path from the root to a leaf node comprises of the zeroth level node corresponding to the tree root, the third-level node corresponding to the leaf node, with two levels of internal nodes in between; (b) the leaf nodes under any second-level internal node in the tree have the same voltage assignments.

Since each VRM can only provide one output voltage level, the number of VRM's in a feasible VRM tree topology cannot be less than the number of distinct voltage levels of the FB's. The number of possible combinations for the first level of the tree is the power set of the number of second-level nodes in that tree. After generating each feasible tree instance T, the RMTO-FM algorithm is used to find the optimum solution for the corresponding T.

One issue with this procedure is that the number of feasible trees with n leaves appears to be quite large; fortunately, in the RMTO problem, many of the generated trees are isomorphic (cf. Figure 5).

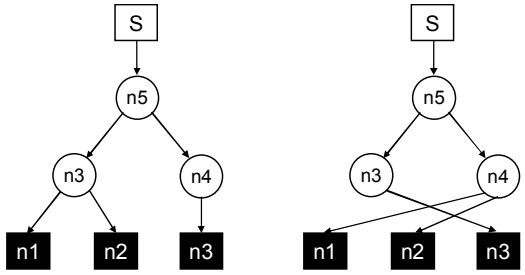

Figure 5: Two inter-isomorphic trees.

Definition: Two VRM trees T_1 and T_2 are called *inter-isomorphic* if by a change of labeling in the intermediate vertices of one tree, it becomes equal to the other; otherwise, they are called *non-inter-isomorphic*. The set of all non-inter-isomorphic trees comprising of exactly two internal levels and n leaf nodes is denoted by $\mathcal{T}_2(n)$.

Lemma 1: The number of all non-inter-isomorphic trees with exactly 2 internal levels and n leaf nodes is obtained from

$$|\mathcal{T}_2(n)| = \sum_{k_2=1}^{n} \sum_{k_1=1}^{k_2} \begin{Bmatrix} n \\ k_2 \end{Bmatrix} \cdot \begin{Bmatrix} k_2 \\ k_1 \end{Bmatrix} \tag{4}$$

where $\begin{Bmatrix} n \\ k \end{Bmatrix}$ is the Stirling number of the second kind [10].

For each n and m, the Stirling number of second kind is the number of ways of partitioning a set of n elements into m nonempty sets. These numbers can be computed from the following sum [10],

$$\begin{Bmatrix} n \\ k \end{Bmatrix} = \frac{1}{k!} \sum_{i=0}^{k} (-1)^i \binom{k}{i} (k-i) \tag{5}$$

Table 1 shows the number of non-inter-isomorphic trees with 2 internal levels and n leaves. From the table data, it is seen that by using only non-inter-isomorphic trees, the number of enumerations required to find the optimal solution becomes more manageable.

Table 1: Number of non-inter-isomorphic trees with *n* leaves

n	1	2	3	4	5
$\lvert\mathcal{T}_2(n)\rvert$	1	3	12	60	358

An algorithm for solving the RMTO-VM problem is presented in Figure 6. It should be noted that although the time complexity of RMTO-VM algorithm is exponential in the number of leaf nodes, because the number of different voltage domains is small, the runtime of the algorithm is quite reasonable.

RMTO-VM ($\mathcal{R}$, $\mathcal{L}$, V_S)
Begin
1. **For** each $T \in \mathcal{T}_2(n)$
2. **If** T is feasible
3. **RMTO-FM** ($\mathcal{R}$, $\mathcal{L}$, T, V_S)
4. **End**
5. **End**
6. **Return** best **RMTO-FM** ($\mathcal{R}$, $\mathcal{L}$, T, V_S)
End

Figure 6: RMTO-VM algorithm for VRM tree optimization.

3. PRACTICAL ISSUES

3.1 Non-Monotone Input Current

The monotone input current property holds as long as the VRM has a single mode, where the basic feedback loop in the regulator which performs the output and line regulation does not change its parameters (reference voltage levels, sensing network parameters, switch configuration, etc) in response to applied input voltages. There are, however, VRM's that may operate as say 2X charge pump or 1.5X charge pump or even an LDO depending on the applied input voltage. Such VRM's tend to exhibit a non-monotone input current vs. output current behavior, which will then break the principle of dynamic programming and require an exhaustive search mechanism to produce the optimum VRM tree solution.

Two changes in the RMTO-FM algorithm are needed to solve the RMTO-FN problem (-*N* option means some of the VRM's have non-monotone input current property). The first is that current look-up tables Ψ that are generated and stored in level-2 nodes should be made 2-D, where the key into the table entries is a pair of values: input voltage of the second-level node and the candidate VRM for that node. The second change occurs when level-1 nodes are traversed. In this case, for each candidate set of a level-1 node, all candidate VRM's in its fanouts should also be enumerated in order to find the best assignment of VRM's. The RMTO-VN algorithm is the same as the RMTO-VM problem except that it calls RMTO-FN in line 3.

3.2 Effect of the Current Profile of the Loads

Current profiles of the loads play a key role in the design of an efficient VRM tree. To motivate the need for considering the load profile of the FB's, consider the following example. Assume that to provide a FB with a desired voltage level, a buck converter is needed and the only candidate converters are those shown in Figure 7. Now, if the load profile of the FB is $\{(200mA,90\%),(100mA,10\%)\}$, i.e., in 90% of the time the FB consume 200mA and in 10% it consumes 100mA current, then using the VRM (b) is more efficient whereas with a load profile of $\{(200mA,10\%),(100mA,90\%)\}$ VRM (a) is a better choice.

In the following, we describe how to account for the effect of load profiles in the RMTO-FM algorithm. To begin with, for simplicity, we assume that the profiles of different FB's are independent of one another. In the next section, we show how to account for the correlations among load profiles.

Assume that *m* FB's, X_1, X_2 , X_m, with the same required voltage level *V* are connected to a node *n*. The current profiles of the FB's are expressed as $\{(I_i^j,\alpha_i^j)\}$ where I_i^j and α_i^j are the current demand and the probability of FB *i* being in its j^{th} state. Notice that

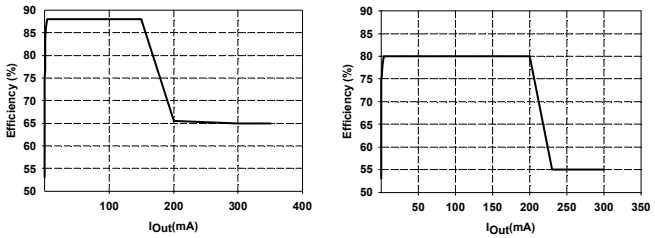

Figure 7: The efficiency curves of two commercial buck VRM (TPS60502 [11] and TPS60503 [6]).

for every *i*, $\sum_{j\in S(i)}\alpha_i^j=1$, where $S(i)$ is the set of states of the load profile of FB *i*. When calculating the efficiency and input current of a candidate regulator c_n for *n* (line 10 and 22 From Figure 4) $i_{out}(n)$ becomes a piecewise-linear function; so, instead of having a constant value for the efficiency and input current of node *n*, we need to model both of them as piecewise-linear functions. That is,

$$\eta^{k_1,k_2,\ldots,k_m}(c_n,v_{in}(n))=f\left(v_{in}(n),I_1^{k1}+I_2^{k2}+\ldots+I_m^{km}\right)$$
$$i_{in}^{k_1,k_2,\ldots,k_m}(c_n,v_{in}(n))=\frac{v_{out}(n)\times\left(I_1^{k1}+I_2^{k2}+\ldots+I_m^{km}\right)}{v_{in}\times\eta^{k_1,k_2,\ldots,k_m}(c_n,v_{in}(n))} \quad (6)$$
$$\Pr(S(k_1,k_2,\ldots,k_m))=\alpha_1^{k1}\alpha_2^{k2}\ldots\alpha_m^{km}\text{ , for } k_i\in S(i),1\le i\le m$$

where $\eta^{k_1,k_2,\ldots,k_m}$ and $i_{in}^{k_1,k_2,\ldots,k_m}$ are the efficiency and input current when FB X_i is in state k_i and $\Pr(S(k_1, k_2,\ldots,k_m)$ is the probability of such an event. Notice that the number of states in node *n* is the product of the number of states in its fanout nodes. An example of generating the piecewise linear input current for the fanin node is shown in Figure 8.

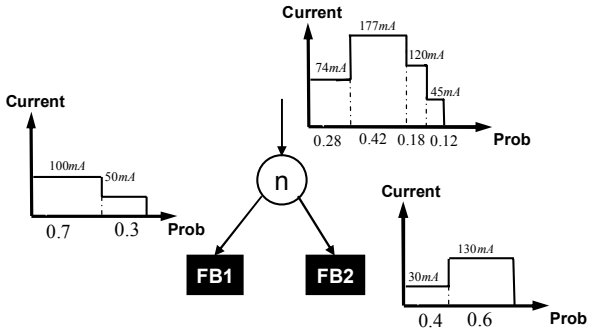

Figure 8: Piecewise-linear modeling of the input current of a VRM, assuming the VRM shown in Figure 7(a) is used and V_{out}/V_{in}=0.5.

The average input current of node *n*, which is used in optimization, can be obtained from

$$\eta^{k_1,k_2,\ldots,k_m}(c_n,v_{in}(n))=f\left(v_{in}(n),I_1^{k1}+I_2^{k2}+\ldots+I_m^{km}\right)$$
$$i_{in}^{avg}=\sum_{k_i\in S(i),1\le i\le m} i_{in}^{k_1,k_2,\ldots,k_m}(c_n,v_{in}(n))\times\Pr(S(k_1,k_2,\ldots,k_m)) \quad (7)$$

The candidate VRM c_n at node *n* should satisfy the constraint that,

$$\iota_{c_n,out}^{\max}\ge\max_{k_i\in S(i),1\le i\le m}\left(I_1^{k1}+I_2^{k2}+\ldots+I_m^{k2}\right) \quad (8)$$

3.3 Effect of Correlations among Current Profiles

The correlation between the load profiles of FB's could be used to design a more efficient VRN tree. To motivate the problem, consider two corner case examples. In the first case, the load currents of the FB's are *positively correlated* in the sense that both FB's have the same peak and off-peak load intervals. An example of such a case is two processor cores that work in parallel. In this case both processors achieve their minimum and maximum currents at the same intervals (c.f. Figure 9(a)). On the other hand, in some cases, the load profiles of the FB's are *negatively correlated*, i.e., when one FB is in its low-load state, the other one is in the high-load state and vice versa (cf. Figure 9(b)). An instance of such a scenario occurs by using activity migration technique for dynamic thermal management in which the peak junction temperature is controlled by moving computation between multiple replicated units [12].

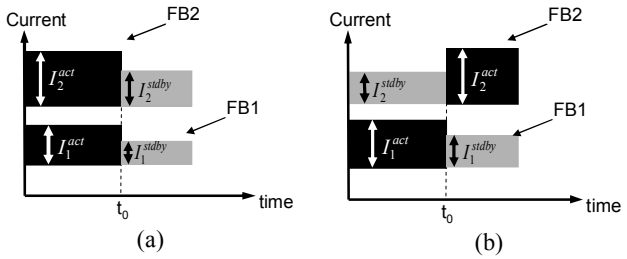

Figure 9: (a) Positively correlated FB's (b) negatively correlated FB's.

It is clear that these two scenarios put different constraints on the VRM tree design. For example, when two FB's are negatively correlated, it is more likely that by sharing a single VRM for both of them, a more power-efficient VRM network can be achieved. Rather minor changes need to be made to the RMTO-FM algorithm so that it can handle the effect of load profile correlations. These changes are similar to those that have been discussed in Section 3.2; so for the sake of space, we do not provide their details here.

4. SIMULATION RESULTS

The algorithms proposed in this paper have been implemented in C++ and evaluated on a set of test-benches. A set of thirty DC-DC commercial regulators from Texas Instruments and National Semiconductors were used to create a library of VRMs. This set consists of ten variants of each type of buck, boost, and LDO regulators. Two of these regulators are those shown in Figure 7. 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.

We compared the results of our RMTO-VM with the results of the optimal VRM assignment in a star topology. Table 2 summarizes the specifications of our benchmarks along with the reduction of power loss in the VRM tree achieved by applying our algorithm. In Table 2, TB is the name of the testbench, V_S is the voltage of power supply, N is the number of FB's in the problem statement, I_{min} and I_{max} denote the minimum and maximum required current by any FB's, while V_{min} and V_{max} are the minimum and maximum required voltages of any FB's. To illustrate how RMTO-VM algorithm selects the topology of the VRM tree, we depict the VRM tree for the first test-bench in Figure 10. In this figure, VRM1 and VRM5 are two LDO's used at the output of two buck regulators to decrease their voltage levels.

Table 2: Simulation results for a few test cases

TB	V_S	N	I_{min}	I_{max}	V_{min}	V_{max}	VRM Tree Power Loss Reduction (%)
1	2.5	6	50m	100m	1.1	1.8	21.9
2	2.5	7	50m	200m	1.3	1.8	23.6
3	2.5	5	60m	200m	1.3	3.0	14.5
4	2.5	8	50m	200m	1.3	3.3	9.4
5	3.3	6	30m	100m	1.2	1.8	14.6
6	3.3	10	50m	300m	1.1	2.7	21.5
7	3.3	12	30m	350m	1.1	3.0	12.3
8	3.3	8	50m	200m	1.3	3.3	17.9

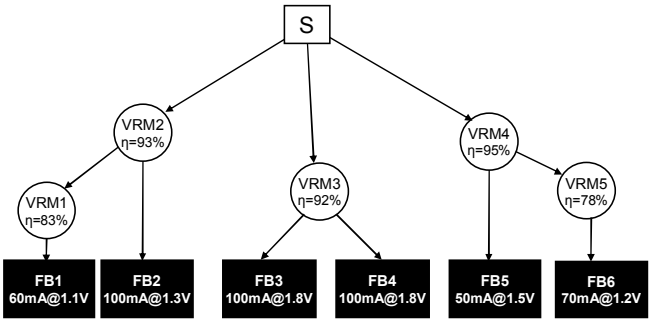

Figure 10 : VRM tree topology in testbench 1.

5. CONCLUSION

In this paper we showed that by using a tree topology of suitably chosen voltage regulators between the power source and loads, one can achieve higher power efficiency in the power delivery network. We formulated the problem of optimizing the VRM tree as a dynamic program and solved it efficiently. The experimental results demonstrate the efficacy of proposed problem formulation and solution. Experimental results showed that by using the proposed technique, the power loss in the VRM tree can be reduced by an average of 17%.

REFERENCES

[1] Semiconductor Industry Association, International Technology Roadmap for Semiconductors, 2003 edition, http://public.itrs.net/.

[2] S. Chun, "Methodologies for Modeling Simultaneous Switching Noise in Multi-Layered Packages and Boards," Ph.D. Dissertation, Georgia Institute of Technology, 2002.

[3] W. Dally and J. Poulton, *Digital Systems Engineering*. Cambridge, U.K.: Cambridge Univ. Press, 1998.

[4] 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.

[5] J. Kim, B. Choi, H. Kim, *et al.*, "Separated role of on-chip and on-PCB decoupling capacitors for reduction of radiated emission on printed circuit board," in *Proc. of International Symposium on Electromagnetic Compatibility*, 2001, pp. 531-536.

[6] Texas Instruments, "TPS60503 Datasheet," available at http://www.ti.com/lit/gpn/tps60503

[7] A. Stratakos, "High-efficiency low-voltage DC-DC conversion for portable applications," Ph.D. Dissertation, University of California, Berkeley, 1998.

[8] M. Pedram and J. Rabaey, *Power Aware Design Methodologies*. Boston, MA: Kluwer Academic Publishers, 2002.

[9] 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.

[10] R. L. Graham, D. E. Knuth, and O. Patashnik, *Concrete Mathematics*. Reading, MA: Addison-Wesley, 1990.

[11] Texas Instruments, "TPS60502 Datasheet," available at http://www.ti.com/lit/gpn/tps60502

[12] S. Heo, K. Barr, and K. Asanovic, "Reducing power density through activity migration," in *Proc. of International Symposium on Low Power Electronics and Design* 2003, pp. 217-222.