

Tutorial: Design and Optimization of Power Delivery Networks

Woojoo Lee

Intelligent SoC Department, Electronics and Telecommunications Research Institute / Daejeon, Korea space@etri.re.kr

Received August 2, 2016; Revised September 1, 2016; Accepted September 6, 2016; Published October 30, 2016

* Review Paper: This paper reviews the recent progress possibly including previous works in a particular research topic, and has been accepted by the editorial board through the regular reviewing process.

Abstract: The era of the Internet of Things (IoT) is upon us. In this era, minimizing power consumption becomes a primary concern for system-on-chip designers. While traditional power minimization and dynamic power management (DPM) techniques have been heavily explored to improve the power efficiency of devices inside very large-scale integration (VLSI) platforms, there is one critical factor that is often overlooked, which is the power conversion efficiency of a power delivery network (PDN). This paper is a tutorial that focuses on the power conversion efficiency of the PDN, and introduces novel methods to improve it. Circuit-, architecture-, and system-level approaches are presented to optimize PDN designs, while case studies for three different VLSI platforms validate the efficacy of the introduced approaches.

Keywords: Power delivery network, PDN, DC-DC converter, Low power design

1. Introduction

With drastic technology scaling, the number of intellectual property (IP) cores integrated into a chip has been increasing to tens and the future potentially hundreds. Furthermore, growing demand for increased Internet of Things (IoT) device functionality has been driving the trend toward including many high-performance modules (such as high-speed processors, fast wireless interfaces, large- and high-resolution displays, and sophisticated sensors) on the IoT platform. Unfortunately, the accompanying increase in power density and high rates of heat generation have become critical roadblocks to the scalability and integration of chips and platforms.

Consequently, there has been a surge of interest in reducing power consumption of these chips and platforms. Conventional power minimization techniques, such as dynamic voltage and frequency scaling (DVFS) and dynamic power management (DPM), have been explored to improve the power efficiency of the devices. DVFS is a well-known technique to mitigate the power consumption of chip multicore processors (CMPs) by dynamically varying the supply voltage and operating frequency values applied to the process cores in response to the load conditions or workload characteristics [1, 2]. DPM is represented by power gating and clock gating techniques that shut off the power and clock to the unused circuit blocks [3, 4].

While the conventional low power techniques have

been heavily investigated, there is one critical factor that has often been overlooked: the power conversion efficiency of the power delivery network (PDN) [5, 6]. The PDN is an essential part of platforms, which delivers power to all devices in the platform from the power source. Because the voltage level of the power source is fixed (e.g., a lithium-ion battery comprised of a single battery cell provides 3.7 V) whereas the required voltage levels of the devices vary, a PDN consists of DC-DC converters to convert the output voltage level of the power source to adequate input voltage levels for the devices. Furthermore, if the device supports voltage scaling, an adjustable DC-DC converter is indispensable. In reality, a DC-DC converter inevitably dissipates power, and power dissipations from all converters can result in a considerable amount of power loss.

Fig. 1 shows two examples of the power conversion efficiency resulting from (a) a smartphone [5], and (b) a multicore platform [6]. The average power conversion efficiency of the smartphone platform in Fig. 1(a) is 67%, which translates into a 33% total power dissipation in the PDN. Fig. 1(b) shows that power loss while delivering power to a core in a multicore platform is sometimes more than 53%. Therefore, reducing such power losses can appreciably reduce the total power consumption of the platforms.

This paper introduces circuit-, architecture-, and system-level approaches to enhance the power conversion efficiency of various IoT platforms. Starting with detailed




Fig. 1. Power conversion efficiency traces from (a) Qualcomm MDP 8660, (b) Sniper with LTC3618.

models of DC-DC converters, circuit-level methods to improve the efficiency of each single DC-DC converter are presented first. The presented optimization method is validated on a smartphone platform that runs various applications. Architectural approaches are discussed for the multicore platform. Novel DVFS methods and a reconfigurable PDN architecture are introduced to minimize the core and PDN power consumption. A new design for a multicore platform is proposed to support the methods. Finally, a procedure to optimally design organic light-emitting diode (OLED) display systems is introduced. A reconfiguration PDN to support fine-grain OLED dynamic voltage scaling (OLED-DVS) is established to minimize the system power consumption.

The remainder of this paper is organized as follows. Section 2 provides some background on the PDN, including DC-DC converter modeling. In Section 3, an optimization method to improve efficiency of each converter in a smartphone platform is presented. Section 4 describes optimization methods to minimize the power consumption of a multicore platform, while Section 5 introduces design and optimization methods of the PDN in OLED display systems. Finally, Section 6 concludes this paper.

2. PDN model

A PDN typically consists of DC-DC converters that can be classified into three types, low-dropout (LDO), switched-capacitor (SC), and inductive switching converters [5]. Compared to inductive switching converters, LDOs and SCs are relatively small and easy to integrate. However, inductive switching converters achieve higher conversion efficiencies over a wide range of output loads, in general. Additionally, utilizing digitally programmable controllers enables the inductive converters to easily support DVS with fast transient response. Therefore, this paper focuses on inductive switching converters (simply called converters hereafter) and targets




Fig. 2. Circuit schematic of an inductive switching DC-DC converter.

them at optimization objectives.

As shown in Fig. 2, the converter is composed of an inductor, capacitors, two switches (either MOSFET or powerFET), and a pulse width modulation (PWM) controller.

According to the output voltage level that becomes higher or lower than the input voltage level, the converter is a boost- or buck-type, respectively. The power loss model of buck converter P_{buck} can be expressed as follows [7, 8]:

$$P_{buck} = P_{conduction} + P_{switching} + P_{controller} \quad (1)$$

where $P_{conduction}$, $P_{switching}$ and $P_{controller}$ are the power loss from conduction, switching, and PWM control, each of which are determined as follows:

$$P_{conduction} = I_{out}^2 (R_L + DR_p + (1-D)R_n) + (\Delta I)^2 (R_L + DR_p + (1-D)R_n + R_c) / 12 \quad (2)$$

$$P_{switching} = V_{in}^2 f_s (C_p + C_n) = V_{in}^2 f_s C_{ox} (W_p L_p + W_n L_n) \quad (3)$$

$$P_{controller} = V_{in} I_{controller} \quad (4)$$

In Eqs. (2)-(4), R_L , R_C , R_p , and R_n are the resistances of inductor, capacitor, and p-type and n-type switches, respectively; D is the PWM duty ratio of the switches, which can be expressed as V_{out}/V_{in} ; I_{out} is the output current of the converter, while $\Delta I = (1-D)V_{out} / (L f_s)$ is the amplitude of the maximum current ripple, where L and f_s denote the inductance and switching frequency, respectively. In (3), the width and length of the switches are represented by W_p and L_p for the p-type switch, and W_n and L_n for the n-type switch. C_{ox} is the gate capacitance per unit area. In (4), $I_{controller}$ denotes the current used in the control logic section of the converter. Similarly, the power loss model of a boost converter can be modeled. Interested readers can refer to Choi et al. [7] and Lee et al. [8].

From (2) and (3), the power losses due to the switches consist of the conduction loss and switching loss. If the minimum length L_{min} MOSFET switches are used, the power losses from pMOS (P_{pmos}) and nMOS (P_{nmos}) may be reformulated by:

$$P_{pmos} = C_{ox} W_p L_{min} \frac{m}{m-1} V_{in}^2 f_{sw} + \frac{DI_{out}^2}{\mu_p C_{ox} \frac{W_p}{L_{min}} (V_{in} - |V_{pth}|)} \quad (5)$$




Fig. 3. Simulation result of converter efficiency according to I_{out} and V_{out} changes. $P_{switching}$ and $P_{conduction}$ dominant regions are conceptually indicated.

$$P_{nmos} = C_{ox} W_n L_{min} \frac{m}{m-1} V_{in}^2 f_{sw} + \frac{(1-D)I_{out}^2}{\mu_p C_{ox} \frac{W_p}{L_{min}} (V_{in} - V_{nth})} \quad (6)$$

where μ_p and μ_n are the hole mobility in the pMOS and the electron mobility in the nMOS, respectively; V_{pth} and V_{nth} denote the pMOS and nMOS threshold voltages. m is the tapering factor for the (super buffer-like) gate driver of the switches. As I_{out} becomes bigger, the conduction loss becomes the dominant source of the power loss. An efficiency graph of V_{out} and I_{out} changes is shown in Fig. 3, where $P_{switching}$ and $P_{conduction}$ dominant regions are indicated conceptually.

Eqs. (5) and (6) reveal an important characteristic of P_{pmos} and P_{nmos} —they are convex functions of the change in gate width. A smaller gate width reduces the switching loss, but increases the conduction loss, and vice versa, for a larger gate width. Furthermore, for a given I_{out} , the function to find the optimal widths of the switches can be derived by $dP_{pmos}/dW_p = 0$ and $dP_{nmos}/dW_n = 0$ as follows [9, 10]:

$$W_{p,opt}(I_{out}) = \frac{I_{out}}{C_{ox} V_{in}} \sqrt{\frac{D(m-1)}{\mu_p (V_{in} - |V_{pth}|) f_{sw} m}} \quad (7)$$

$$W_{n,opt}(I_{out}) = \frac{I_{out}}{C_{ox} V_{in}} \sqrt{\frac{(1-D)(m-1)}{\mu_n (V_{in} - |V_{nth}|) f_{sw} m}} \quad (8)$$

Meanwhile, the power conversion efficiency of a converter, η , is calculated as:

$$\eta = \frac{V_{out} I_{out}}{V_{in} I_{in}} = \frac{V_{in} I_{in} - P_{buck}}{V_{in} I_{in}} \quad (9)$$

3. Dynamic switch modulation scheme in a smartphone platform

As mentioned about Eqs. (7) and (8), the optimal sizes of the switches that minimize the power loss of a converter




Fig. 4. Circuit schematic for DSM scheme.




Fig. 5. Example of DSM operation with two pMOS switches.

can be derived when I_{out} is given a priori. Unfortunately, the load conditions of various applications give rise to dynamically changing I_{out} . To tackle this problem, adaptively turning on or off some of the multiple parallel-connected switches was introduced [11, 10]. However, the different gate voltages required for each switch set by Abdel-Rahman et al. [11] need additional DC-DC converters, which likely causes control and area overhead. Furthermore, although the number of switches and their sizes should be determined carefully to achieve the maximum efficiency under the given design specifications, the fixed number of switches in Abdel-Rahman et al. [11] and Kudva and Harjani [10] limits their availability. Instead, the multiple switch scheme presented by Lee et al. [5] (which is called dynamic switch modulation, or DSM) effectively solves the problems.

Fig. 4 shows a converter schematic for the DSM scheme. In the figure, N pairs of switches are supposed to be connected in parallel. These switches are arranged such that the first switch has the minimum width (denoted by W_{p1} and W_{n1}), and the last switch has the maximum width (denoted by W_{pN} and W_{nN}). Depending on I_{out} , a different on/off combination of the switches can be used to achieve maximum conversion efficiency. More precisely, when an effective width is defined as the sum of widths of all turned-on switches of the same type, the effective width of each type should be nearest to $W_{p,opt}$ or $W_{n,opt}$ in (7) and (8), according to I_{out} . Fig. 5 is an example of DSM operation with two pMOS parallel-connected switches. In the figure, the i^{th} smallest effective width and the boundary current of switch type $type$ are defined as $W_{eff,type,i}$ and $I_{bd,type,i}$. The boundary current is the condition for turning the switches on/off when I_{out} is over/under a certain current value. For instance, from (7), $I_{bd,p,i}$ may be calculated as




Fig. 6. Conversion efficiency and general current distribution of one module group in the Qualcomm platform.

$$I_{bd,p,i} = C_{ox} V_{in} \sqrt{\mu_p f_{sw} W_{eff,p,i} W_{eff,p,i+1} \frac{m}{D(m-1)}} \quad (10)$$

Similarly, $I_{bd,n,i}$ can be derived from (8).

To validate the efficacy of DSM for enhancing the conversion efficiency, the Qualcomm MDP 8660 smartphone platform was used. There are 35 modules in the platform, and they can be classified into seven groups such that each group includes the modules that require the same voltage level. For example, a group that requires 1.1 V has the following modules: internal memory, audio DSP, the GPU, and modems.

Fig. 6 shows the current distribution of this group, which is derived from running various smartphone applications and exploiting a general smartphone usage pattern [8]. The red line in the figure indicates the conversion efficiency of a converter equipped within the platform. In the figure, the best conversion efficiency region is mismatched to the current distribution. In other words, the best conversion efficiency region should be located in the 150–180 mA region, wherein the current distribution has the highest probability. To make them well-matched, the widths of both nMOS and pMOS switches should be sized down by using (7) and (8), respectively. More precisely, the switch sizes should be 0.3793 times smaller than the original sizes.

Although the switch sizes are tuned to match the current distribution, these sizes are only available for the current distribution derived from a certain usage pattern. Again, DSM is thus necessary for the varying current conditions. Simulation results in Table 1 show that applying DSM to the target module group achieves the high-efficiency enhancement for a wide load current range, when efficiency enhancement G_η is defined as

$$G_\eta = \left(\frac{\eta_{tuned}}{\eta_{original}} - 1 \right) 100(\%) \quad (11)$$

where η_{tuned} and $\eta_{original}$ are the conversion efficiencies with/without applying a sizing method. In the table, there are two types of DSM, each of which uses three or four parallel switches (i.e., the normalized size of each switch is

Table 1. Efficiency Enhancement Results of the DSM [5].

Application	$G_{\eta, fixed}$	$G_{\eta, DSM1}$	$G_{\eta, DSM2}$
System Setting	4.4993	4.4576	4.6029
Call	4.8733	4.9012	4.9499
Skype-videochat	-0.0035	1.5326	1.5531
Facebook	3.2841	3.4753	3.5499
Neocore(game)	-3.0773	0.4821	0.4636
$I_{out} = 200\text{mA}$	2.7330	2.8415	8.3743
$I_{out} = 250\text{mA}$	0.7476	4.7393	4.8130
$I_{out} = 300\text{mA}$	-5.8153	2.1294	2.2558
DSM1: three switches {0.2970, 0.3945, 0.8983}			
DSM2: four switches {0.2970, 0.3717, 0.4122, 0.8983}			

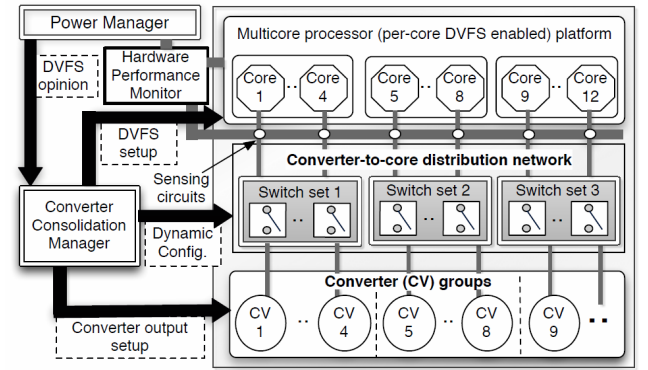


Fig. 7. A new multicore platform architecture presented by Lee et al. [6] to enable per-core DVFS with a reconfigurable PDN.

indicated in the table). $G_{\eta, fixed}$, $G_{\eta, DSM1}$ and $G_{\eta, DSM2}$ are the efficiency enhancements from the fixed sizing method and DSMs with three and four parallel switches, respectively. The negative values of $G_{\eta, fixed}$, in the table mean the conversion efficiency decreases, because the fixed-switch sizing method works less well under some workload conditions, whereas applying DSM results in positive G_η in all cases.

4. Converter Consolidation in a Multicore Processor Platform

One of the most effective techniques for reducing the power consumption of CMPs is to dynamically vary the supply voltage and operating frequency values applied to the process cores in response to load conditions or workload characteristics (also known as DVFS). The conventional approach is to perform DVFS for all cores in a processor (per-chip DVFS). This approach has not been able to take full advantage of the power-saving that DVFS potentially achieves. For instance, some of the cores may not need a high voltage/frequency level, but cannot be lowered because of the other cores. To overcome this shortcoming, applying DVFS to each individual core (per-core DVFS) was suggested, and has given rise to excellent




Fig. 8. Converter consolidation results from a four-core multicore platform that runs SPLASH2-Barnes in two cores and PARSEC-Streamcluster in two cores.

flexibility in controlling power [12-14]. Unfortunately, this approach can still have the inevitable drawbacks, such as a larger footprint, higher power conversion loss, and higher control complexity due to the more complicated PDN requirement.

To support per-core DVFS, at least the same number of converters as the number of cores should be equipped in the platform, each of which inevitably dissipates power, and power dissipations from all the converters can result in a considerable amount of power loss. Fig. 1(b) is an example to show that PDN power loss is sometimes more than 53% in such platforms [15, 6]. Of course, the previously introduced method to optimize a single converter is helpful in decreasing PDN power loss, but here comes another novel method that effectively reduces the power loss of multiple converters. Furthermore, the introduced method ultimately minimizes the total power consumption of the system (converters and cores) by collaborating with the per-core DVFS scheme.

The idea starts from the concept of combining some cores that operate at the same voltage level and drive relatively small amounts of load current so they are powered by a single converter. This approach can significantly reduce the converter power loss in a multicore platform for the following reasons: i) the converter used to power multiple cores has relatively high load current, and has higher efficiency, according to the converter characteristics, and ii) the unused converters are turned off to save power. Based on this concept of converter consolidation, a new design for the multicore platform can be proposed, which exploits (multiple) sets of network switches that enable reconfiguration of the PDN. Fig. 7 shows the proposed architecture; a detailed description will be discussed soon.

Along with the reconfigurable PDN architecture, two optimization methods were presented by Lee et al. [6] to minimize the converter power loss, and thus, maximize the total energy savings. The first proposed method is a reactive approach that configures the PDN based on the sensed voltage/current level of each core. Next, a proactive method is presented to decide the optimal voltage/frequency level of each core when considering

Table 2. Simulation results of Lee et al. [6] from applying converter consolidation methods to multicore processor platforms.

Benchmark	# of cores	$\beta = 5\%$		$\beta = 10\%$	
		G_{PDN}	G_{total}	G_{PDN}	G_{total}
Streamcluster	16	28.81	7.28	23.19	5.95
Swaption	16	24.34	6.75	24.34	6.75
Barnes	8	32.21	7.86	31.30	7.98
FFT	8	6.40	1.16	6.59	1.25
Ocean	4	19.11	5.04	19.74	5.28
Raytrace	4	18.09	3.40	22.96	4.71
Cholesky	4	18.99	4.70	21.54	5.68
FMM	4	16.04	3.57	17.73	4.20

maximized consolidation opportunities for converters in order to minimize the energy consumption of the whole system.

Now, let us go into the details of the presented architecture and the two methods from Lee et al. [6]. The platform in Fig. 7 has a number of converters and multiple cores. There are several groups of reconfigurable converter-to-core connection networks supported by network switches implemented with power MOSFET switches. The converter-to-core network delivers power for each core from any converter in the same group. The power manager (PM) in a conventional CMP platform controls the processor's operating condition by using the DVFS technique. Compared to the conventional architecture, the newly added converter consolidation manager (CCM) finally controls the core's frequency and voltage levels, as well as the operations of the converters and the configurations of the converter-to-core network (by controlling on/off states of the network switches).

Normally, the network switch on/off time, T_{NS} , is much smaller than the voltage-level transition time of converters T_{CV} [16]. As a consequence, the DVFS setting and network reconfiguration can be treated as global and local power management, respectively, of the consolidation method. The reactive method (as local management) applies only to cores operating at the same voltage level determined by the power manager. The detailed way to find the connections between converters and cores is as follows: the CCM first sorts the cores that have the same voltage levels and a lower amount of input current than the maximum driving capability of a converter. Then, based on the current levels, the CCM finds the two cores by merging, in which the converter energy savings is maximized. After consolidation of those two cores, the CCM keeps repeating this procedure until there are no cores available, or the converter energy savings from the consolidation method for the remaining cores is less than the power loss of the network switch transition.

As global power management, the proactive method exploits DVFS techniques to minimize the power consumption of all the cores, network switches, and converters in the decision period T_{CV} . In the proposed method, a trade-off exists between the energy savings by DVFS (which is initially determined by the PM) and the

reduced energy loss by adaptively turning off the VRs and using fewer VRs at higher conversion efficiencies. If the CCM determines that the latter option is more beneficial, the CCM will not decrease the frequency/voltage levels of some cores to the minimum possible level. Instead, it will adjust the frequency/voltage levels of the cores to increase the opportunities to apply the converter consolidation procedure. Note that in order to guarantee that the performance (i.e., total execution time of applications) is not degraded by the modification of the DVFS schedule, an important constraint of the CCM is that the original DVFS levels suggested by the PM should be kept the same or increased (but never decreased).

To demonstrate the efficacy of the presented converter consolidation methods, a multicore processor simulator, Sniper [17], was used, and various PARSEC [18] and SPLASH2 [19] benchmarks were performed in the simulator. An ILP-based algorithm presented by Kim et al. [12] was adopted to derive the original DVFS levels of cores that are supposed to be obtained from the PM. Finally, Fig. 8 is an example of the simulation results from applying the presented methods. The simulation is set to have four cores, each of which has two cores to run *Barnes* and *Streamcluster*. The histogram in the figure indicates that converter consolidation can be applied in many cases (e.g., Case 0, where four cores are connected to only one converter was almost 20% when running the benchmarks). Representative simulation results are introduced in Table 2 [6]. In the table, β is a performance penalty from applying the original DVFS levels, in that a higher β means applying DVFS more aggressively. The energy loss reduction (%) from all the converters and the whole platform are defined as G_{PND} and G_{total} , respectively.

5. Reconfigurable PDN for OLED displays

As in dealing with the smartphone platform in Section 2, one interesting factor is a power consumption breakdown of the Qualcomm MDP 8660, which is shown in Fig. 9. Consuming much of the power in the figure is a display system, which is sometimes more than 40% of the total. In this section, the display system is targeted to effectively reduce its power consumption by exploring PDN optimization.

Among the various types of panel display systems, this paper focuses on OLED display systems, which have emerged as a promising light source. OLED is a surface-emitting lighting source, with each pixel comprised of red, green, and blue cells. From a power consumption perspective, OLED cells with different displayed colors have different power efficiencies and different power consumption at a given luminance level. As a result, to display black, an OLED pixel (with red, green, and blue cells) consumes less than 40% of the power consumed by a liquid crystal display (LCD) pixel, whereas displaying white consumes almost three times as much power as an LCD pixel [20].

To tackle the power efficiency issue in OLED displays, many power management methods have been proposed, which mainly focus on controlling pixel color composition.




Fig. 9. Power breakdown results of Qualcomm mobile development platform.

Some examples are the local dimming method presented by Betts-LaCroix [21], and color remapping methods [22, 23]. Furthermore, OLED-DVS was proposed [24] to minimize power dissipation in OLED pixel drivers. Given that the luminance of the OLED pixel is proportional to its driving current, this OLED-DVS method can maintain the image quality as long as the driving current of the OLED pixels can be maintained regardless of the voltage scaling.

Recently, a more aggressive approach for the OLED-DVS has been investigated [20, 25], which partitions a panel into several zones (sub-panels) and applies different possible voltage levels to the different zones. Applying OLED-DVS to each zone can take full advantage of the power-saving the DVS method can offer. Note that, similar to the previous discussion about per-core DVFS vs. per-chip DVFS in Section 4, if DVS is applied to the whole panel, some regions of the panel may not need a high voltage level, but their voltage level cannot be lowered due to the requirements of other regions. This method is called zone-specific OLED-DVS.

In order to realize zone-specific OLED-DVS, multiple converters are indispensable, which gives rise to inevitable power dissipation due to the converters. Furthermore, the finer the OLED panel is sub-panelized, the more implementation overhead is required to equip the multiple converters [26]. For example, based on the converter component prices shown in Table 3, a converter with an LT3791 buck-boost LED driver controller, along with one inductor, three capacitors, and four powerFETs, costs US\$19.60. Besides, this converter occupies at least a 172 mm² printed circuit board area, which results in significant area overhead. To overcome the problem due to the multiple converters, let us exploit the reconfigurable PDN concept [6] again.

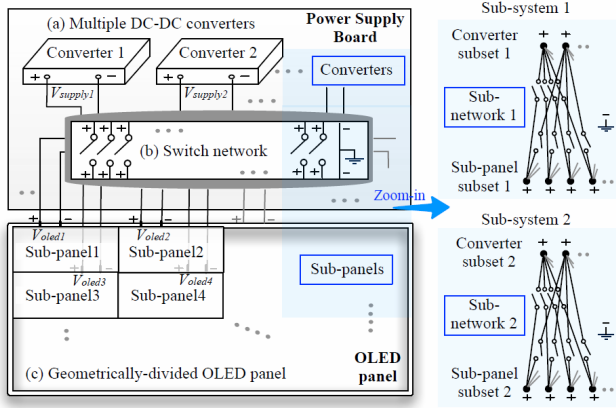


Fig. 10. OLED display systems with a reconfigurable PDN and sub-panels.

Although the concept is similar to the converter consolidation method presented in Section 4, the method introduced in this section is inherently different because i) the number of converters equipped in the OLED display is less than the total number of sub-panels, so the sub-panels need to be grouped, and ii) when considering the area overhead of converters and the switch network, powerFET switches should be adopted instead of MOSFET switches because MOSFET switches used in the multicore platform cannot drive high current. Hence, for the OLED panel, it is necessary to use powerFET switches with a larger footprint and higher current driving capability.

Fig. 10 shows the presented reconfigurable PDN architecture for the zoned OLED display [26]. In spite of the advantages of the switch network, in terms of conversion efficiency and area/cost overhead of the multiple converters, the complexity of the switch network needs to be controlled, and thus, the number of sub-panels that one converter can be connected to should be limited. For each converter, the required number of powerFET switches linearly increases with the number of sub-panels. If the number of subpanels and converters is large, the area/cost overhead of switches becomes significant. Moreover, using too many switches gives rise to a power dissipation increase from the unused switches. Therefore, the presented switch network in Fig. 10 is divided into sub-networks, and each converter (and sub-panel) exclusively belong to its own sub-network (i.e., the sub-network forms a complete bipartite graph). To determine the sub-network size at design time, designers should consider i) area/cost overhead of powerFET switches, ii) the maximum current that a single converter will inject into the sub-network, and iii) the power conversion efficiencies of the converters. A sub-network should be designed to be neither too large, owing to the requirement for a large number of powerFET switches, nor too small, owing to the limited freedom in reconfiguration and the low conversion efficiency under low load current conditions.

To demonstrate the effectiveness of the presented framework, a 65-inch 4K ultra high-definition OLED display is used. Four 4K images, namely Balloons, Bridge, Leopard, and Heidelberg, are explored to apply OLED-DVS to the target panel, which is divided into 4 by 4 (i.e.,

Table 3. Components in a converter/switch network.

Component	Spec.	Product	Manufacturer	Price
Inductor	10uH	7447709100	Würth Electronics	\$3.10
Capacitor	10uF	1EA100WR	Panasonic	\$0.50
Regulator	Buck-boost	LT3791	Linear Technology	\$11.80
PowerFET	N-type	Si1470DH	Vishay Siliconix	\$0.80

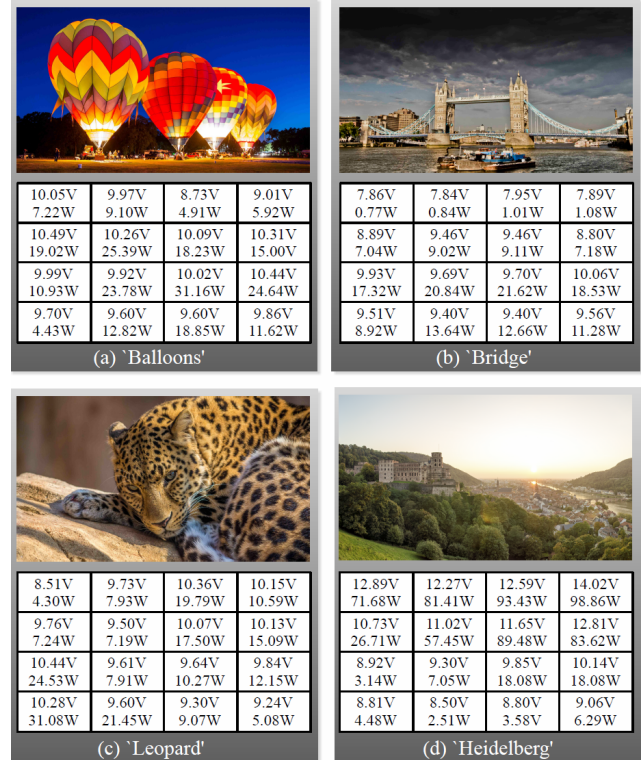


Fig. 11. Experimental results from applying OLED-DVS to 4K images in a 4x4-zoned 65-inch OLED panel.

a total of 16 sub-panels.) The derived voltage level of each panel is listed in Fig. 11. Then, three different sub-network setups are determined such that each setup delivers power to the (upper or lower) eight sub-panels from four, three, and two converters. According to the number of converters and sub-panels in a sub-network, the proposed methods are notated as follows: *DVS* 8:4, 8:3, and 8:2 imply that there are eight sub-panels with four, three, and two converters in a sub-network. Table 4 shows the simulation results for five different methods, including i) DVS from Shin et al. [24] applied to a whole panel that is denoted by *DVS* 16:1, ii) DVS from Chen et al. [20] and Chen and colleagues [25] applied to each sub-panel denoted by *DVS* 16:16, and iii) the proposed methods from Lee and colleagues [26]. As a reference point, the *DVS* NO column lists the power consumption values without DVS. For the results of the proposed methods, the power consumption of the powerFET switches (Si1470D) are calculated and included

Table 4. Simulation results of a 65-inch 4K OLED display system that is divided into 4-by-4 sub-panels. According to the applied method, the power consumption of the panel for each image is denoted by P_{method} . Power saving (%) as well as the additional cost for each method and image are provided.

Image	No DVS	Previously presented methods		Proposed methods		
	$P_{DVS\ NO}$ (W)	$P_{DVS\ 16:1}$ (W)	$P_{DVS\ 16:16}$ (W)	$P_{DVS\ 8:4}$ (W)	$P_{DVS\ 8:3}$ (W)	$P_{DVS\ 8:2}$ (W)
'Balloon'	395.2	285.6 (27.7%)	256.8 (35.0%)	255.2 (35.4%)	258.9 (34.5%)	261.7 (33.7%)
'Bridge'	268.4	184.8 (31%)	173.6 (35.3%)	169.5 (36.8%)	170.4 (36.5%)	172.2 (35.8%)
'Leopard'	343.8	246.6 (28.2%)	225.2 (34.5%)	222.6 (35.3%)	224.3 (34.8%)	227.5 (33.8%)
'Heidelberg'	985.9	977.9 (0.8%)	696.3 (29.4%)	730.5 (25.9%)	751.1 (23.8%)	799.4 (18.9%)
Additional Cost		-	\$~294	\$~188.4	\$~136.4	\$~84.4

in the table. Price information on each component in a converter, seen in Table 3, is used to estimate the cost for each method.

As expected, *DVS* 16:16 saves more power than *DVS* 16:1. In particular, *DVS* 16:16 achieves remarkable power saving with Leopard and Heidelberg, which consume high power with pixels at the highest luminance. However, implementing *DVS* 16:16 costs an extra \$294, which is expensive. On the other hand, compared to *DVS* 16:16, the proposed methods can achieve similar power-saving levels at much less expense. Furthermore, if images do not require many pixels to have high luminance, the proposed methods can save more power than *DVS* 16:16. For example, *DVS* 8:4 saves 37% for Bridge, but *DVS* 16:16 saves 35% thanks to one of the benefits of the proposed reconfigurable PDN, i.e., fewer converters, which lowers power consumption. In addition, each converter may have higher efficiency than the converter used in *DVS* 16:16. Also note that the cost of implementing *DVS* 8:4 is 36% lower than *DVS* 16:16. Finally, the results in

Table 4 prove that the proposed framework consistently achieves high power-conversion efficiency and significant energy savings while minimizing the overhead of the converters.

6. Conclusion

This paper addressed the problem of power conversion efficiency in various very large-scale integration platforms, where significant power is dissipated by PDNs. To accomplish the conversion-efficiency improvement, thereby maximizing power savings on the platforms, circuit-, architecture-, and system-level approaches were introduced. Each presented method was validated on a real platform, such as smartphone, multicore processor and display system.

Summarizing this tutorial, the reader should retain at least the following fundamental concepts for the design and optimization of PDNs.

- The best efficiency region of a converter should match the load current condition in order to maximize conversion efficiency. The best efficiency region can be tuned by sizing the widths of the equipped switches.
- For dynamically varying load current conditions, a multiple (parallel connected) switch scheme can be

adopted.

- To achieve the full power-saving potential of a DVFS in a CMP, per-core DVFS should be supported. A reconfigurable PDN is a solution to minimize the power dissipation from multiple converters, which is indispensable for per-core DVFS.
- The reconfigurable PDN operates with two algorithms: reactive and proactive converter consolidation algorithms.
- A similar concept for per-core DVFS can be applied to OLED display systems, such as a fine-grain (sub-panelized) OLED-DVS. And multiple converters give rise to a similar problem: power dissipation and cost/area overhead.
- To tackle the problem in the OLED display system presented, the concept of reconfigurable PDNs can be utilized.

Acknowledgement

This research was supported by grants from the IT R&D Program of MSIP/IITP [Near-Zero-Voltage Micro-Grain Architecture for Ultra-Low-Energy Processor].

References

- [1] J. Henkel and S. Parameswaran, *Designing Embedded Processors—A Low Power Perspective*. Dordrecht, The Netherlands: Springer, 2007. [Book \(CrossRef Link\)](#)
- [2] A. Alimonda, S. Carta, A. Acquaviva, A. Pisano, and L. Benini, "A feedback-based approach to DVFS in data-flow applications," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 28, no. 11, pp. 1691–1704, Nov. 2009. [Article \(CrossRef Link\)](#)
- [3] L. Benini, A. Bogliolo and G. De Micheli, "A survey of design techniques for system-level dynamic power management," in *IEEE Trans. on Very Large Scale Integr. Syst.*, vol. 8, no. 3, pp. 299–316, June 2000. [Article \(CrossRef Link\)](#)
- [4] K. Roy, S. Mukhopadhyay and H. Mahmoodi-Meimand, "Leakage current mechanisms and leakage reduction techniques in deep-submicrometer CMOS circuits," in *Proc. of the IEEE*, vol. 91, no. 2, pp. 305–327, Feb 2003. [Article \(CrossRef Link\)](#)
- [5] W. Lee, Y. Wang, D. Shin, N. Chang and M. Pedram, "Optimizing the Power Delivery Network in a

- Smartphone Platform," in IEEE Trans. on Comp.-Aided Design of Integr. Circuits and Syst., vol. 33, no. 1, pp. 36-49, Jan. 2014. [Article \(CrossRef Link\)](#)
- [6] W. Lee, Y. Wang and M. Pedram, "Optimizing a Reconfigurable Power Distribution Network in a Multicore Platform," in IEEE Trans. on Comp.-Aided Design of Integr. Circuits and Syst, vol. 34, no. 7, pp. 1110-1123, July 2015. [Article \(CrossRef Link\)](#)
- [7] Y. Choi, N. Chang, and T. Kim, "DC-DC converter-aware power management for low-power embedded systems," IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst., vol. 26, no. 8, pp. 1367-1381, Aug. 2007. [Article \(CrossRef Link\)](#)
- [8] W. Lee, Y. Wang, D. Shin, N. Chang, and M. Pedram, "Power conversion efficiency characterization and optimization for smartphones," in Proc. Int. Symp. Low Power Electron. Design, 2012, pp 103-108. [Article \(CrossRef Link\)](#)
- [9] S. Musunuri and P. L. Chapman, "Optimization of CMOS transistors for low power dc-dc converters," in Proc. Power Electron. Specialist Conf., 2005, pp. 2151-2157 [Article \(CrossRef Link\)](#)
- [10] S. Kudva and R. Harjani, "Fully-integrated on-chip DC-DC converter with a 450X output range," IEEE J. Solid-State Circuits, vol. 46, no. 8, pp. 1940-1951, Aug. 2011. [Article \(CrossRef Link\)](#)
- [11] O. Abdel-Rahman, J. A. Abu-Qahouq, L. Huang, and I. Batarseh, "Analysis and design of voltage regulator with adaptive FET modulation scheme and improved efficiency," IEEE Trans. Power Electron., vol. 23, no. 2, pp. 896-906, Mar. 2008. [Article \(CrossRef Link\)](#)
- [12] W. Kim, M. Gupta, G.-Y. Wei, and D. Brooks, "System level analysis of fast, per-core DVFS using on-chip switching regulators," in Proc. Int Symp. High-Perform. Comput. Archit., Salt Lake City, UT, USA, Feb. 2008, pp. 123-134. [Article \(CrossRef Link\)](#)
- [13] T. Kolpe, A. Zhai, and S. S. Sapatnekar, "Enabling improved power management in multicore processors through clustered DVFS," in Proc. Design Autom. Test Europe, Grenoble, France, Mar. 2011, pp. 1-6. [Article \(CrossRef Link\)](#)
- [14] K. Wang, H. Yu, B. Wang, and C. Zhang, "3D reconfigurable power switch network for demand-supply matching between multi-output power converters and many-core microprocessors," in Proc. Design Autom. Test Europe, Grenoble France, Mar. 2013, pp. 18-22. [Article \(CrossRef Link\)](#)
- [15] W. Lee, Y. Wang, and M. Pedram, "VRCon: Dynamic reconfiguration of voltage regulators in a multicore platform," in Proc. Design Autom. Test Europe, Dresden, Germany, Mar. 2014, pp. 1-6. [Article \(CrossRef Link\)](#)
- [16] J. Park, D. Shin, N. Chang, and M. Pedram, "Accurate modeling and calculation of delay and energy overheads of dynamic voltage scaling in modern high-performance microprocessors," in Proc. Int. Symp. Low-Power Electron. Design, Austin, TX, USA, 2010, pp. 419-424. [Article \(CrossRef Link\)](#)
- [17] T. E. Carson, W. Heirman, and L. Eeckhout, "Sniper: Exploring the level of abstraction for scalable and accurate parallel multi-core simulation," in Proc. Int. Conf. High Perform. Comput. Netw. Storage Anal., Seattle, WA, USA, 2011, pp. 1-12. [Article \(CrossRef Link\)](#)
- [18] C. Bienia and K. Li, "PARSEC 2.0: A new benchmark suite for chipmultiprocessors," in Proc. 5th Workshop Model. Benchmark. Simulat., New York, NY, USA, Jun. 2009, pp. 1-9. [Article \(CrossRef Link\)](#)
- [19] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta, "The splash-2 programs: Characterization and methodological considerations," in Proc. Int. Symp. Comput. Archit., Santa Margherita Ligure, Italy, 1995, pp. 24-36. [Article \(CrossRef Link\)](#)
- [20] X. Chen, J. Zeng, Y. Chen, W. Zhang and H. Li, "Fine-grained dynamic voltage scaling on OLED display," Asia and South Pacific Design Autom. Conf., Sydney, NSW, 2012, pp. 807-812. [Article \(CrossRef Link\)](#)
- [21] J. Betts-LaCroix, "Selective dimming of OLED displays," 2010. U.S. Patent 0149223 A1.
- [22] M. Dong, Y. Choi, and L. Zhong, "Power-saving color transformation of mobile graphical user interfaces on OLED-based displays", IEEE/ACM Int. Symp. on Low Power Elec. and Design, 2009, pp 339-342 [Article \(CrossRef Link\)](#)
- [23] M. Dong and L. Zhong, "Power Modeling and Optimization for OLED Displays," in IEEE Transactions on Mobile Computing, vol. 11, no. 9, pp. 1587-1599, Sept. 2012. [Article \(CrossRef Link\)](#)
- [24] D. Shin, Y. Kim, N. Chang and M. Pedram, "Dynamic voltage scaling of OLED displays," ACM/IEEE Design Autom. Conf., New York, NY, 2011, pp. 53-58. [Article \(CrossRef Link\)](#)
- [25] X. Chen, J. Zheng, Y. Chen, M. Zhao and C. J. Xue, "Quality-retaining OLED dynamic voltage scaling for video streaming applications on mobile devices," ACM/IEEE Design Autom. Conf., San Francisco, CA, 2012, pp. 1000-1005. [Article \(CrossRef Link\)](#)
- [26] W. Lee, Y. Wang, D. Shin, S. Nazarian and M. Pedram, "Design and optimization of a reconfigurable power delivery network for large-area, DVS-enabled OLED displays," IEEE/ACM Int. Symp. on Low Power Elec. and Design, Rome, 2015, pp. 159-164. [Article \(CrossRef Link\)](#)



Woojoo Lee is a senior engineer at the Electronics and Telecommunications Research Institute (ETRI), Daejeon, Korea. He received his BSc in Electrical Engineering from Seoul National University (SNU), Korea, in 2007, and his MSc and PhD in Electrical Engineering from the University of Southern California (USC), USA, in 2010 and 2015, respectively. Dr. Lee worked at NEXON, Korea, as a programmer, and at Broadcom as a mobile ASIC designer. He has served, or is currently serving, as a reviewer and on the Technical Program Committee for many important journals, conferences, symposiums, and workshops for the Circuits and Systems Society and the Computer Society. His research interests include low-power VLSI design, cross-layer power and thermal management, ultra-low power designs, and hardware/software co-designs.