

Distributed On-Chip Switched-Capacitor DC–DC Converters Supporting DVFS in Multicore Systems

Pingqiang Zhou, Ayan Paul, Chris H. Kim, *Senior Member, IEEE*, and Sachin S. Sapatnekar, *Fellow, IEEE*

Abstract—Dynamic voltage and frequency scaling (DVFS) is a powerful technique to reduce power consumption in a chip multiprocessor. To support DVFS in the multicore power delivery network, we integrate on-chip switched-capacitor (SC) dc–dc converters that can work with multiple conversion ratios to provide varying levels of V_{dd} supplies. We study the application of such SC converters in multicore chips by simulation. Our results show that distributed SC converters can significantly reduce the voltage droop seen by the local core loads by providing better localized power regulation. Considering the fact that the current distribution in a multicore chip is unbalanced, we further develop computer-aided design techniques to automate the design (size) and distribution (number and location) of these SC converters, using the efficiency of the whole power delivery system as the optimization metric. This is a major concern, but has not been addressed at the system level in prior research. We develop models for the power loss of such a system as a function of size and distribution of the SC converters, then propose an approach to optimize the SC converters to maximize the efficiency of the system, while considering all the possible conversion ratios an SC converter can work with. We verify the accuracy of our models for the power loss in the power delivery system, and demonstrate the efficiency of our techniques to optimize the SC converters on both homogenous and heterogenous multicore chips.

Index Terms—Chip multiprocessor (CMP), dynamic voltage and frequency scaling (DVFS), efficiency, MILP, on-chip switched-capacitor dc-dc converter.

I. INTRODUCTION

IN recent years, the chip industry has migrated toward chip multiprocessors (CMPs), with the purpose of maximizing computation while remaining with an affordable power envelope [1]. In this multicore era, larger numbers of smaller, more power-efficient cores are being integrated onto a single die to build CMPs. This change has resulted in major challenges to the design of power delivery networks. Individual cores may run different kinds of applications and this application mix can change over time and hence power delivery hotspots may move to different parts of a chip. Therefore, temporal

and spatial variations in power demands are particularly acute in multicore processors. Such issues are complex even for homogeneous multicores due to the spatial variations in power demands within each core, which consists of heterogeneous function units such as central processing units (CPUs), memory units ($L1$ and $L2$ caches), and communication units (I/Os). The integration of heterogeneous cores onto a single die further aggravates the spatial and temporal variations in power demands of the chip. This is because: 1) heterogeneous cores are designed with different capabilities and performance levels, and therefore have different core sizes and power densities and 2) heterogeneous CMPs can dynamically switch workloads between the cores at runtime to take full advantage of the heterogenous architecture when executing a program [2].

Multicore systems can benefit very significantly from the use of dynamic voltage and frequency scaling (DVFS), which enables power management while conducting computations under stringent power considerations [3]–[5]. It is broadly acknowledged that DVFS is one of the most effective techniques to reduce power consumption in CMPs. The variations in the power demands over all the cores in a CMP can be best met if DVFS is supported by providing multiple levels of V_{dd} supplies from either off-chip or on-chip voltage regulators (dc–dc converters) that are essential components of the power delivery network.

There are two kinds of dc–dc converters: switching converters and linear converters. Current-day dc–dc converters are mostly implemented by linear regulators, such as LDOs [6]–[9], but only switching converters can provide a wide range of output voltage at high efficiency which is critical for the application of DVFS in CMPs [10]. Switching converters may be built using either inductors or capacitors. The inductors or capacitors used to build the off-chip switching converters at the board level are costly and bulky, and this limits the use of off-chip voltage regulators in CMPs to ensure supply integrity and serve diverse loads [10], [11]. Therefore, to enable effective DVFS in a multicore chip, it is essential to build fully integrated on-chip switching converters. Capacitors have advantages over inductors for building on-chip switching converters because they can achieve higher quality factors while incurring lower cost overheads than inductors, including area and the number of fabrication steps [10].

Historically, on-chip capacitive switching converters have only been used for low-power applications (in the order of μW) primarily due to the limited power density they can provide [12]. Recent progress [13], [14] shows that through the use of deep trench capacitors, switched-capacitor (SC)

Manuscript received January 17, 2013; revised May 18, 2013 and August 8, 2013; accepted August 12, 2013. This work was supported in part by NSF under Grant CCF-0903427 and in part by SRC under Grant 2009-TJ-1990.

P. Zhou is with the School of Information Science and Technology, Shanghai Tech University, Shanghai 200031, China (e-mail: zhoupq@shanghaitech.edu.cn).

A. Paul, C. H. Kim, and S. S. Sapatnekar are with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455 USA (e-mail: paul0661@umn.edu; chriskim@umn.edu; sachin@umn.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TVLSI.2013.2280139




Fig. 1. Schematic of a power delivery system.

converters can provide high current density up to 2.3 A/mm², high energy transfer efficiency ($\approx 90\%$) and minimal parasitic losses. This implies that now SC converters are feasible for high-performance applications such as CMPs. In addition, SC converters have been demonstrated to support DVFS with low overheads, providing a wide range of output voltages by dynamically reconfiguring the internal structure of SC converters (Section II). This reconfiguration allows the converter to provide different voltage conversion ratios (i.e., from the same input voltage, they can generate different levels of voltage supplies) at runtime [11].

This work studies the application and optimization of SC converters for DVFS in multicore power delivery system that may have multiple power/voltage domains. Since each domain has to be optimized separately, we present an approach for optimizing a single voltage domain in this paper. Fig. 1 shows a simplified power delivery system including the global V_{dd} supply, an SC converter to translate the input V_{dd} to required voltage supply level in a power domain, a power grid to distribute the power to local core loads, and a core load. The output voltage of the converters is V_{out} , but the exact voltage supply seen by the cores is downgraded to V_{core} due to voltage losses such as voltage droop (e.g., due to IR drop) in the power delivery network. To overcome these losses and ensure correct core operation, the ideal value of V_{out} must be set to V_{ideal} , the specification of supply voltage in the power domain, as given by

$$V_{ideal} = V_{vdd,core} + V_{droop} + \Delta V \quad (1)$$

where $V_{vdd,core}$ is the minimum voltage specified at the core load, V_{droop} is the peak voltage droop between V_{out} and V_{core} , and ΔV is the peak-to-peak output voltage ripple of the converter. For a core that draws current I_{out} , the power supplied to the converters is

$$P_{cvt} = I_{out} V_{ideal}. \quad (2)$$

However, the power drawn by the core load is smaller

$$P_{load} = I_{out} V_{vdd,core}. \quad (3)$$

The remainder of the power, $I_{out}(V_{droop} + \Delta V)$, is wasted in various parts of the power delivery network. Note that there is additional wasted power from the energy transfer process within the converter.

There has been limited prior work on the optimization of on-chip SC dc-dc converters in a multicore system. The work in [10] has focused primarily on optimizing the internal design of the converter to reduce wasted power within the converter (“SC converter” box in Fig. 1) by controlling the voltage ripple ΔV , and choosing the optimal switch width and switching

frequency. Under this paradigm, the burden of optimizing the other term for the voltage droop, V_{droop} (corresponding to the “power grid” box in Fig. 1) in the system, is placed on conventional means for power grid optimization, e.g., grid topology selection and wire widening.

In this paper, we address this problem from two aspects.

- 1) First, we suggest the use of distributed SC converters in a multicore system. Our simulation results show that the voltage droop seen by the core loads is affected by both the number and location, i.e., distribution, of the converters. Compared with a single lumped converter, distributed converters with the same total amount of capacitance can significantly reduce the voltage droop by providing better localized voltage regulation. With the same number of converters, the voltage droop is also dependent on the locations of the converters on the chip.
- 2) Second, we consider a holistic optimization of the SC converters at the system level to minimize the power loss in the whole system. Due to the fact that the current distribution in a CMP system is spatially imbalanced, using SC converters with identical size and evenly distributing them over a chip area is not the best choice. Therefore, we develop a computer-aided design (CAD) approach to automate the design and distribution of the SC converters for DVFS, with the aim of maximizing the efficiency of the whole system.

We begin with the development of models for the power loss in the power delivery system as a function of the size and distribution of the SC converters, and verify the accuracy of our models by simulation. Before work [10], [11] presented related models for the loss inside the converters that have only one single interleaving stage. In contrast, our loss analysis applies to the whole power delivery system, and we consider converters with multiple interleaving stages.

We then show that the efficiency optimization problem with SC converters supporting DVFS can be formulated as an mixed-integer nonlinear program (MINLP) problem, and we propose a two-step approach to solve the MINLP problem. In particular, we show that by optimizing the distribution of the converters for the chip, it is possible to control the power loss in the power grid and enhance the efficiency of the whole power delivery system. Our results also show that the optimal solution for one conversion ratio can be suboptimal for another, with up to 10% difference in efficiency results.

To the best of our knowledge, our work is the first to study the application of SC converters that can support DVFS in a CMP system, and to optimize both the design (size) and distribution (number and location) of the SC converters to minimize the power loss at the system level.

II. SC DC-DC CONVERTERS

A block diagram of a general SC converter system is shown in Fig. 2(a). The system consists of N_{phase} interleaving stages (typical values of N_{phase} are 16 and 32), which reduce the ripple voltage by $1/N_{phase}$ compared with an SC converter without any interleaving.




Fig. 2. SC dc-dc converter. (a) Block diagram of an SC DC-DC Converter. (b) Topology of a 2:1 SC converter. (c) Output waveform.

At the core of the system is the switch matrix, one for each phase [11]. This matrix is a reconfigurable arrangement of switches and flying (charge-transfer) capacitors, that provides the ability to produce a different voltage conversion ratio, allowing the converter to generate one of the several output voltage levels from the same input global V_{dd} supply [11] to support DVFS in a CMP. The conversion ratio of the converter, $ratio_{cvt}$, is defined as the ratio between the input supply voltage, V_{dd} , and the desired output voltage $V_{vdd,dom}$. The control circuit generates the nonoverlapping clock signals Φ_1 and Φ_2 for the switches in the switch matrix.

A switch matrix topology is shown in Fig. 2(b), with a conversion ratio $ratio_{cvt}$ of 2:1.¹ Fig. 2(c) (top) shows that during Φ_1 , the flying capacitor C_{fly} is connected to the input global V_{dd} to get charged, and during Φ_2 , the charge stored in C_{fly} is transferred to the load and its voltage drops by ΔV as it is discharged. This is reflected as the output voltage at the output, V_{out} of the converter, as shown in Fig. 2(c) (bottom) [10]. Note that the signals Φ_i are generated by a relatively low-frequency clock ($f_{sw} \approx 100$ MHz), which is distinct from the multi-GHz clock used by the multicore processor.

III. APPLICATION OF SC CONVERTERS IN MULTICORE POWER DELIVERY SYSTEM

In this section, we explore the application of on-chip SC dc-dc converters in the context of CMPs. Before work has not adequately studied the layout implications of on-chip power supply design. In particular, when SC converters are integrated into an on-chip power delivery network, they may be built in either lumped or distributed form, as shown in Fig. 3.

For the lumped case, a large central converter delivers power to all the blocks in the whole chip. In contrast, for the distributed case, several smaller converters can be distributed across the chip and each load can absorb power from the nearby converters. It is well known that power delivery is most efficient if the power sources are close to the utilization points (it is for this reason that decoupling capacitors—which deliver power based on stored charge—are placed close to large noise sources [15]). In this paper, we quantitatively

¹More complex matrices are used for a larger set of voltage levels [11]. For simplicity, we stay with a simple converter topology here, but the switch matrices used for our experiments are more complex and deliver more diverse voltage conversion ratios.




Fig. 3. Lumped versus distributed on-chip dc-dc converters. (a) Lumped. (b) Distributed.




Fig. 4. Model of power delivery network used in our simulations.

TABLE I
SUMMARY OF SC DC-DC CONVERTERS

Capacitance density	0.2 μ F/mm ²
Total area	23.04mm ²
Total capacitance	4.608 μ F
N_{phase}	16
Duty cycle	50%
f_{sw}	100 Mhz
Switch resistance	20m Ω

compare the lumped and distributed designs of on-chip SC converters by simulations using realistic power profiles from CMP applications.

A. Simulation Setup

Fig. 4 shows a detailed model of the power delivery network for the CMP used in this paper. The package and C4 bump contacts are modeled as RL pairs. The on-board power supply is modeled as a dc voltage source. The on-chip power delivery network consists of a global V_{dd} grid, lumped or distributed on-chip dc-dc converters, a local power grid, a global GND grid, core or decoupling capacitors, and current loads. The global sparse V_{dd} grid supplies power to on-chip SC converters. The local power grid distributes power to the local core loads, and its voltage is controlled by the lumped or distributed on-chip SC converters. Note that in this paper, the converters are shared by all the cores on-chip, although one core may mainly draw power from its nearby converters.

In our simulations in this section, we show a realistic instance where the lumped and distributed designs of SC converters have significantly different performance. we consider a test case with three cores, whose floorplan is shown in Fig.6(a).

In our simulations, we model each core as a single current source and generate the current profiles for the cores by




Fig. 5. Power trace for three cores obtained from the simulation of a typical multicore workload ($V_{dd} = 1.2$ V).

TABLE II
SIMULATION CONFIGURATION

DC voltage source	$V_{dd}=1.2V$
Package	$L_{pkg} = 15pH, R_{pkg} = 1m\Omega$
C4 bump	#bumps = 768, $L_{bump} = 7.2pH, R_{bump} = 1.5m\Omega$
Core load	Capacitance = 1 nF




Fig. 6. Case 1 with one single converter. (a) Floorplan. (b) Simulation results, min voltage = 774 mV.




Fig. 7. Case 2 with three distributed converters. (a) Floorplan. (b) Simulation results, min voltage = 823 mV.

simulating a typical SPEC OMP2001 [16] workload using an accurate full system multicore simulator GEMS [17].

Fig. 5 shows a typical power trace, we obtained from the workload. From this figure, we can clearly see that there are both temporal and spatial variations in the power demands of these cores in the test case.

For the SC converters, we use the structures shown in [11]. The switches are modeled as resistors when they are turned on. In accordance with common practice, as outlined at the beginning of Section II, 16-phase interleaving is used to reduce the output ripple of the converters. The parameters for the SC converters studied here are summarized in Table I, and the other parameters for the power grid and the CMP are listed in Table II.




Fig. 8. Case 3 with three distributed converters at different locations compared with Case 2. (a) Floorplan. (b) Simulation results, min voltage = 816 mV.

B. Simulation Results

We now compare the performance of the lumped and distributed designs of the on-chip SC converter. For this experiment, we assume that the SC converter(s) works with a 4:3 conversion ratio, i.e., the nominal V_{dd} supply to the cores is 0.9 V. We then compare the following three cases.

- Case 1: Lumped design with one single SC converter in the center of the test chip that delivers power to all three cores as shown in Fig. 7(a).
- Case 2: Distributed design with three SC converters whose floorplan on the chip is shown in Fig. 8(a).
- Case 3: Distributed design with three SC converters placed differently compared with Case 2, as shown in Fig. 9(a).

For fair comparison: 1) the same amount of total available flying capacitance is used for these three cases and 2) 16-phase interleaving is used in all the converters.

We exercised these three designs by applying the power trace shown in Fig. 5, and the results are, respectively, shown in Figs. 7–9. Compare Case 1 with Case 2, we can see that, for a nominal voltage of 900 mV, the minimum voltage seen by the cores can be improved from 774 to 823 mV, and the maximum IR drop can be reduced by up to 52% if we move from the lumped design to the distributed design. Note that in Case 2, the IR drops of three cores are different due to the spatial variation in the power demands of these cores, as discussed in the previous section. Compare Case 2 with Case 3, we can see that although these two cases use the same number of converters, the IR drop and actual voltage seen by the core loads are different due to the different floorplan of these converters. Therefore, the voltage droop seen by the core loads is dependent on both the number and location (i.e., distribution) of the converters on chip.

IV. ANALYSIS OF POWER LOSS IN THE POWER DELIVERY SYSTEM USING SC CONVERTERS

In Section III, we have shown an example that illustrates that distributed converters can significantly reduce the voltage droop seen by the local core loads by providing better localized voltage regulation, and the voltage droop is affected by the distribution of the converters. Therefore, in the rest of this paper, we develop a CAD solution to find the optimal size and distribution of SC converters for a given CMP.

We begin with the development of models for the SC converter, which will be used within an optimization framework. As will be described in further detail in Section V,

TABLE III

TOPOLOGY-DEPENDENT PARAMETERS [11]. α IS THE RATIO OF THE PLATE CAPACITANCE TO ITS EFFECTIVE CAPACITANCE

Conversion ratio	Nominal V_{dd}	N_{sw}	M_{sw}	γ	M_p	M_{topo}
1:1	1.2V	2	1	1	0	1/2
4:3	0.9V	10	7/3	2/3	$3/8\alpha$	8/9
3:2	0.8V	7	2	1	$1/3\alpha$	9/8
2:1	0.6V	4	2	2	$1/4\alpha$	2
3:1	0.4V	7	2	3/4	0.2775α	9/8

we use efficiency, one of the key design metrics for the on-chip dc-dc converters [10], [18] as an optimization objective. Since the efficiency of a multicore power delivery system is determined by the total power loss in the system, from a modeling standpoint, we analyze various components of power loss in a multicore power delivery system in this section. We present models for various components of the power loss in Section IV-A, as a function of the size and distribution of the SC converters, and then discuss the verification of our loss models in Section IV-B.

A. Power Loss Analysis

We now analyze the inefficiency and power loss in the power delivery system using SC converters. Our analysis borrows extensively from previous work as well as on conversations with designers. Before work [10], [11] has presented models only for the loss inside the converters, and they only consider converters with one single interleaving stage. In contrast, our loss analysis applies to the whole power delivery system, and we consider converters with multiple interleaving stages.

For each converter, let f_{sw} be the switching frequency of the converter, $C_{sw} = C_{fly} \times N_{phase}$ be the total amount of flying capacitance, and ΔV be the output ripple of the converter. Our model description will use the parameters described in Table III, which shows how some key parameters vary with the conversion ratio. These parameters are defined as follows.

- 1) N_{sw} is the number of switches used in one topology.
- 2) M_{sw} is topology-related constant that models conduction loss.
- 3) γ is topology-related constant that models switch width.
- 4) M_p is topology-related constant that models parasitic loss.
- 5) M_{topo} is topology-related constant that models the amount of current a converter can provide.

The second column in Table III shows the levels of ideal V_{dd} supplies provided by the converter under different conversion ratios when the input V_{dd} supply to the converter is 1.2 V.

Note that most of the loss components described here are dependent on the particular conversion ratio for a converter corresponding to a specific level of V_{dd} supply to the loads, i.e., on the internal topology of the converters. This is because: 1) as shown in Table III, the values of several major parameters are different for different converter ratios (topologies) and 2) when the cores are working at different levels of V_{dd} supply under DVFS, they have different demands on the current I_{out}

drawn from the converters.

The components of power loss can be categorized as follows.

1) *Conduction Loss*: This corresponds to the power loss in the switches as the flying capacitors are charged. Before work [10] presents a model for conduction loss with one single interleaving stage ($N_{phase} = 1$), we extend it for the general case with multiple interleaving stages ($N_{phase} \geq 2$) here. For each converter, the conduction loss is modeled as

$$P_{cond} = M_{sw} \frac{I_{out}^2}{N_{phase}} \frac{R_{on}}{W_{sw}} \quad (4)$$

where M_{sw} is a constant determined by the converter topology (Table III), I_{out} is the total current delivered by the converter, R_{ON} is the switch resistance density measured in $\Omega \cdot m$, and W_{sw} is the switch width. For a given topology, W_{sw} is proportional to f_{sw} and C_{sw} [11]

$$W_{sw} = \sigma \gamma f_{sw} \frac{C_{sw}}{N_{phase}} \quad (5)$$

where σ is a fitting coefficient, and γ is topology dependent (Table III).

2) *Gate-Drive Loss of the Switches*: Similarly, we generalize the model presented in [10] for special case with $N_{phase} = 1$, to model the power loss in driving the gate nodes of transistors (switches in the converter) for multiple interleaving stages ($N_{phase} \geq 2$) as

$$P_{sw} = N_{phase} \cdot N_{sw} \cdot f_{sw} \cdot (C_{gate} W_{sw}) \cdot V_{dd}^2 \quad (6)$$

where C_{gate} is the per-unit-width gate capacitance of the switches and N_{sw} is topology dependent (Table III).

3) *Parasitic Loss*: This loss, from the bottom-plate parasitic capacitance of the flying capacitors, can be estimated as [10]

$$P_{para} = M_p f_{sw} C_{sw} V_{dd}^2 \quad (7)$$

where M_p is a topology-related parameter (Table III).

4) *Load Loss*: The load power loss $I_{out}(V_{droop} + \Delta V)$, described in Section I, can be separated into two parts.

- a) The part determined by the voltage ripple, ΔV , is

$$P_{L1} = I_{out} \frac{\Delta V}{2}. \quad (8)$$

When switching at frequency f_{sw} , the current a converter can provide is $I_{out} = M_{topo} \cdot f_{sw} \cdot C_{sw} \cdot N_{phase} \cdot \Delta V$

$$\Delta V = \frac{I_{out}}{M_{topo} f_{sw} C_{sw} N_{phase}}. \quad (9)$$

From (9), with the same output current I_{out} , the voltage ripple ΔV varies inversely with the charge-transfer capacitance C_{sw} .

- b) The power loss associated with the voltage droop, V_{droop} , is

$$P_{L2} = I_{out} V_{droop}. \quad (10)$$

Note that the voltage droop changes as we alter the number and locations of the converters on the chip, since the distance between the converters and the utilization points (cores) changes.




Fig. 9. Experimental setup for verification. (a) Simplified power delivery system with one SC converter driving one lumped load. (b) Topologies for five different voltage conversion ratios [11].

5) *Loss from the Control Circuitry and Clock*: The power losses from the control circuitry P_{ctrl} and clock P_{clock} [Fig. 2(a)] are both dependent on the number of used converters. We use a penalty term for these two items in the objective formulation, as stated in Section VI-B.

B. Verification of Our Power Loss Model

In this section, we verify the accuracy of our SC-converter-specific loss models presented in Section IV-A.

In this paper, we verify the loss components [1) to 4a)] in Section IV-A, which are the key converter-topology-specific components of loss and are complicated to model in a power delivery system. For the remaining components of power loss, we have used standard models. Therefore, we build a simplified power delivery system with a single SC converter delivering power to a lumped current load representing the core loads in the chip, as shown in Fig. 9. As discussed in Section II, this SC converter is capable of reconfiguring its internal structure to produce different voltage conversion ratios (Fig. 9 shows four of them used in this paper), therefore delivering a wide range of supply voltage to the load.

Table IV summarizes the design parameters for our simulation-based experimental validation. The converter can work with four different conversion ratios; therefore, with a global V_{dd} supply of 1.2 V, the nominal voltage supplied by the converter ranges from 0.4 V (3:1 conversion ratio) to 0.9 V (4:3 conversion ratio).

In our experiments, we compare the efficiency numbers obtained in the following two different ways.

- 1) Using the analytical loss model presented in Section IV-A: For each load voltage, we use the loss models to calculate each loss component, and then estimate the efficiency number from the calculated total power loss and the actual load power.

TABLE IV
DESIGN PARAMETERS

Global V_{dd}	1.2V
Voltage conversion ratios	4:3, 3:2, 2:1 and 3:1
Load current I_{load}	0.025Amp – 0.4Amp
Load voltage V_{out}	0.25V – 0.9V
Switching frequency f_{sw}	200Mhz
Number of interleaving stages N_{phase}	16

- 2) By simulation of the power delivery system shown in Fig. 9 in HSPICE: We implement the converter with five possible voltage conversion ratios. For each conversion ratio, we sweep the output load to obtain the efficiency plot.

Fig. 10 shows the results for the comparison over a wide range of output supply voltage from 0.25 to 0.9 V (0.9 V is the maximum output voltage supported by the industrial 32-nm SoI process used in our experiments). The red curve shows the efficiency plot created by analytical analysis, and the blue curve shows the plot generated by simulation. We can see that the efficiency plot predicted by our analysis closely matches the simulated efficiency values. Therefore, our loss model is accurate and good enough for the efficiency optimization in our later work presented in Section V.

The maximum efficiency for each conversion ratio can also be seen from the peaks in Fig. 10. For each conversion ratio, with a fixed global V_{dd} supply and a given current load, there is an optimal load voltage at which the efficiency of the system is maximized. This is because, as can be seen in Section IV-A, conduction loss increases as ripple ΔV (the voltage difference between ideal and actual output voltage of the converter, Section II) increases. However, other loss components (e.g., gate-drive loss, parasitic loss) decrease with ΔV , and therefore, for a given conversion ratio, there is an optimum ΔV where the sum of the two losses is minimized.

In a multicore chip design, for a certain level of operating V_{dd} , the minimum voltage for the core load is determined by the circuit specification, such as the working clock frequency, providing a hard constraint that must be satisfied. However, the actual voltage supplied to the load is optimizable, and is determined by the global V_{dd} supply, the converter design and its conversion ratio, and the voltage loss in the power delivery network connecting the converter output to the load (refer to Fig. 1). Therefore, in this paper, we optimize the global V_{dd} supply, together with both the design (size) and the distribution (number and location) of the converters on the chip, so as to find the optimal load voltage for a given chip to maximize the efficiency of the whole power delivery system, while meeting the minimum voltage constraints for the core loads.

V. OPTIMIZATION OF SC CONVERTERS IN THE POWER DELIVERY SYSTEM

In this section, we propose the formulation for the optimization of efficiency in the power delivery system using SC dc–dc converters that can support DVFS by providing multiple voltage conversion ratios. In the scenario studied here, it is safe to assume that the switching frequency f_{sw} and interleaving stages N_{phase} are fixed for the converters.




Fig. 10. Comparison of efficiency plots with change in load voltage.

With the analysis in Section IV-A, when converters are working with a certain voltage conversion ratio l , the components of power loss can be divided into three categories. We extend the notation in Section IV-A with a superscript (l) , which denotes the corresponding power loss at a conversion ratio of l .

The first component of power loss, $P_1^{(l)}$, includes the conduction loss $P_{\text{cond}}^{(l)}$, gate-drive loss of the switches $P_{\text{sw}}^{(l)}$, parasitic loss $P_{\text{para}}^{(l)}$, and part of load loss $P_{L1}^{(l)}$. $P_1^{(l)}$ is determined by the C_{sw} and the global V_{dd} . The second component, $P_2^{(l)}$, is part of load loss. The third component, $P_3^{(l)}$, is the sum of the power loss from the control circuitry and clock. Both $P_2^{(l)}$ and $P_3^{(l)}$ are determined by the number and distribution of the converters.

At the system level, the efficiency of the power delivery system $\eta^{(l)}$ is defined as the ratio between power delivered to the load and total power extracted from the input V_{dd} supply

$$\eta^{(l)} = \frac{P_{\text{load}}^{(l)}}{P_{\text{load}}^{(l)} + P_1^{(l)} + P_2^{(l)} + P_3^{(l)}} \quad (11)$$

where $P_{\text{load}}^{(l)}$ is defined in (3). To improve the overall efficiency of the power delivery system using SC converters for the given conversion ratio l , we should minimize the total loss in the power delivery system, i.e., $P_1^{(l)} + P_2^{(l)} + P_3^{(l)}$.

Further, for SC converters that can provide N voltage conversion ratios, we optimize the weighted sum of normalized power loss for each possible conversion ratio l as

$$\text{minimize } \sum_{l=1}^N w_l \cdot \frac{P_1^{(l)} + P_2^{(l)} + P_3^{(l)}}{P_{\text{load}}^{(l)}} \quad (12)$$

where w_l is the weighting factor for ratio l . In general, this factor can be chosen to provide additional weight to some conversion ratios over others, although our experimental evaluation sets equal weights for all conversion ratios. In the real design, w_l may also be user specified.

The optimization variables are:

- 1) the number of converters used;
- 2) the locations of the used converters;
- 3) the capacitance of each used converter C_{sw} .

Which are common to all the N possible conversion ratios. The optimization is subject to the following constraints.

- 1) For each conversion ratio $l = 1, \dots, N$, the supply voltage at each core load must meet a lower bound

$$V_{\text{core}}^{(l)} \geq V_{\text{vdd,core}}^{(l)} \quad (13)$$

where $V_{\text{vdd,core}}^{(l)}$ is the minimum voltage specified at the core load. Note that $V_{\text{vdd,core}}^{(l)}$ is different when the cores are working at different levels of V_{dd} supply under DVFS.

- 2) Since in reality the voltage ripple constraint must limit $\Delta V^{(l)} \leq \Delta V_{\text{max}}^{(l)}$, where $\Delta V_{\text{max}}^{(l)}$ is the maximum allowable voltage ripple associated with conversion ratio l , (9) provides a bound on C_{sw} for each ratio l

$$C_{\text{sw}} \geq \frac{I_{\text{out}}^{(l)}}{f_{\text{sw}} N_{\text{phase}} M_{\text{topo}}^{(l)} \Delta V_{\text{max}}^{(l)}}, \quad l = 1, \dots, N. \quad (14)$$

- 3) To control the capacitance resource used, we require that

$$\sum C_{\text{sw}} \leq C_{\text{max}} = C_{\text{unit}} \cdot \text{Area}_{\text{max}} \quad (15)$$

where C_{unit} is the capacitance density, and Area_{max} is the maximum available area for the converters.

We present our solution to the above efficiency optimization problem in Sections VI for a special case with $N = 1$, i.e. with one single voltage conversion ratio, then provide solution to the more generalized case with $N \geq 2$ in Section VII.

VI. SOLUTION FOR SPECIAL CASE WITH ONE SINGLE VOLTAGE CONVERSION RATIO ($N = 1$)

In this section, we show that the efficiency optimization problem described in Section V can be formulated as an MINLP, and then propose a two-step-based approach to solve it. Note that in this section, to simplify the notations in the formulas, we drop the superscripts “ (l) ” in the variables and constants associated with a certain voltage conversion ratio l .

Fig. 11 shows a simplified schematic of the on-chip power delivery network for a multicore processor, which is part of the power delivery system showed in Fig. 4. The voltage supplied to the power grid is controlled by a set of on-chip SC converters, which can be placed at a list of predefined candidate locations on the chip.

We now show an optimization formulation for the problem defined in Section V with $N = 1$ as an MINLP, by introducing 0-1 integer variables z_i s, with $z_i = 1$ denoting a placed converter at candidate location i . We first macromodel the power grid in Section VI-A, and then present the complete MINLP formulation in Section VI-B.




Fig. 11. (a) Model of power delivery network. (b) Network macromodel with m candidate converters and n observation nodes.

A. Macromodeling of the Power Grid

We build a macromodel of the power grid with only: 1) the set of selected n observation ports of the core loads, denoted as OBS and 2) the set of m predefined candidate ports for the converters, denoted as Src, and abstract away all the other nodes in the grid using the approach in [19], as shown in Fig. 11.

By partitioning the ports into sets Src and OBS, the transfer characteristics of the macromodel are

$$\begin{bmatrix} I_{\text{Src}} \\ I_{\text{OBS}} \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} V_{\text{Src}} \\ V_{\text{OBS}} \end{bmatrix} + \begin{bmatrix} S_{\text{Src}} \\ S_{\text{OBS}} \end{bmatrix} \quad (16)$$

where $(I_{\text{Src}}, V_{\text{Src}})$ and $(I_{\text{OBS}}, V_{\text{OBS}})$ are the (current, voltage) values at the Src and OBS ports. $A_{11}, A_{12}, A_{21}, A_{22}$ are conductance matrixes, $(S_{\text{Src}}, S_{\text{OBS}})$ are constant vectors of current from the ports to the reference node depending on the conversion ratio l . The reader is referred to [19] for the details about the derivation of (16).

Since $I_{\text{OBS}} = 0$, we have

$$V_{\text{OBS}} = T \cdot V_{\text{Src}} + B \quad (17)$$

where $T = -A_{22}^{-1}A_{21}$ and $B = -A_{22}^{-1}S_{\text{OBS}}$. Further

$$I_{\text{Src}} = A_{11}V_{\text{Src}} + A_{12}V_{\text{OBS}} + S_{\text{Src}} = A'V_{\text{Src}} + S'_{\text{Src}} \quad (18)$$

where $A' = A_{11} + A_{12}T$ and $S'_{\text{Src}} = S_{\text{Src}} + A_{12}B$.

From (17) and (18), we can see that the current vector of the Src ports I_{Src} and voltage vector of the OBS ports V_{OBS} are linear functions of the voltage vector of the Src ports V_{Src} .

B. MINLP Formulation

Using the macromodel shown in Fig. 11, the optimization problem described in Section V is equivalent to finding the optimal z_i assignments, and for each used converter i (with $z_i = 1$), determining its size C_i .

Based on (4), (5), (8), and (9), P_1 (Section V), the power loss associated with the converter and the global V_{dd} supply, can be written as

$$P_1 = \sum_{i=1}^m \left(e_1 e_3 I_{\text{Src}}^i \Delta V_i + e_2 V_{\text{ideal}}^2 C_i \right) \quad (19)$$

where

$$\begin{aligned} e_1 &= \left(\frac{M_{\text{sw}} R_{\text{on}}}{\sigma \gamma} + \frac{1}{2M_{\text{topo}} N_{\text{phase}}} \right) \frac{1}{f_{\text{sw}}} \\ e_2 &= f_{\text{sw}} (N_{\text{sw}} C_{\text{gate}} f_{\text{sw}} \sigma \gamma + M_p) \cdot \text{ratio}_{\text{cvt}}^2 \\ e_3 &= M_{\text{topo}} f_{\text{sw}} N_{\text{phase}}. \end{aligned}$$

Using (17), P_2 , the power loss in the grid, and P_3 are

$$\begin{aligned} P_2 &= \underbrace{\sum_{i=1}^m (V_{\text{Src}}^i (I_{\text{Src}}^i - S_{\text{Src}}^i))}_{\text{Power supplied to the macromodel}} \\ &\quad - \underbrace{\sum_{j=1}^n (V_{\text{OBS}}^j S_{\text{OBS}}^j)}_{\text{Power delivered from the macromodel}} \\ &= \sum_{i=1}^m \left(V_{\text{Src}}^i (I_{\text{Src}}^i - S_{\text{Src}}^i) \right) - \sum_{j=1}^n (B^j S_{\text{OBS}}^j) \quad (20) \end{aligned}$$

$$P_3 = P_{\text{ctrl}} + P_{\text{clock}} = c \cdot \sum_{i=1}^m z_i \quad (21)$$

where c is penalty weight for control circuit and clock network, $V_{\text{ideal}}, V_{\text{Src}}^i, I_{\text{Src}}^i, C_i, \Delta V_i$ are the continuous variables and z_i s are the 0–1 integer variables in the optimization problem.

We then transform the problem in Section V into the MINLP

$$\begin{aligned} \text{minimize } P_1 + P_2 + P_3 &= \sum_{i=1}^m \left(e_1 e_3 I_{\text{Src}}^i \Delta V_i + e_2 V_{\text{ideal}}^2 C_i \right) \\ &\quad + \sum_{i=1}^m \left(V_{\text{Src}}^i (I_{\text{Src}}^i - S_{\text{Src}}^i) \right) - \sum_{j=1}^n (B^j S_{\text{OBS}}^j) + c \sum_{i=1}^m z_i \quad (22) \end{aligned}$$

subject to $\forall j \in \text{OBS}$

$$V_{\text{OBS}}^j = \sum_{i=1}^m (T_{ji} \cdot V_{\text{Src}}^i) + B^j \geq V_{\text{th}}^j \quad (23)$$

$\forall i \in \text{Src}$:

$$I_{\text{Src}}^i = \sum_{k=1}^m (A'_{ik} \cdot V_{\text{Src}}^k) + S'_{\text{Src}}^i \quad (24)$$

$$0 \leq I_{\text{Src}}^i \leq M \cdot z_i \quad (25)$$

$$I_{\text{Src}}^i = e_3 \cdot \Delta V_i \cdot C_i \quad (26)$$

$$0 \leq C_i \leq M \cdot z_i \quad (27)$$

$$0 < \Delta V_i \leq \Delta V_{\text{max}} \quad (28)$$

$$V_{\text{Src}}^i + \Delta V_i \leq V_{\text{ideal}} \quad (29)$$

and

$$\sum_{i=1}^m C_i \leq C_{\text{max}} \quad (30)$$

where V_{th}^j is the minimum required voltage at the observation nodes of each core and M is a large positive number.

Constraints (23) are transformed from (13), to specify the minimum voltage for each core load. Constraints (24) are from (18), and Constraints (26) from (9). Constraints (25) are structured to ensure that the current I_{Src}^i is zero when no converter connected to candidate port i , while Constraints (27) ensure that converter size C_i is zero when I_{Src}^i is zero, both through the use of M . Constraints (28) and (30) are from (14) and (15), and Constraints (29) set the bound for the Vdd supply.

We can observe that there are nonlinear (actually nonconvex) terms in the objective function (22) and constraints (26) are also nonlinear. Therefore, the above optimization problem is an MINLP.

C. Two-Step Optimization Approach

It is well known that MINLP problems are difficult to solve [20]. Therefore, in this paper we develop a two-step approach to solve the MINLP optimization problem presented in Section VI-B. For the objective function in (22).

- 1) $P_2 + P_3$ is determined by the number/location of the converters.
- 2) P_1 is determined by the converter design, i.e, the size of converters C_i , and V_{ideal} , the V_{dd} supply. From (1), we can see that V_{ideal} is determined by the voltage droop in the power grid and the ripple in the converters.

Therefore, we may optimize the power loss in two steps. We first optimize $P_2 + P_3$, the power in the distribution network, by finding the optimal number and location of the converters. We present an MILP-based approach for this step. Next, we optimize P_1 to determine the optimal size of each used converter C_i .

1) *Approximation for the Voltage Ripple*: We introduce the approximation that all converters have the same voltage ripple, implying that the current delivered by a converter i is proportional to its capacitance C_i (26) when working with a conversion ratio l . We justify this approximation as follows. In (19), let P_1^i be the contribution of the i th converter to P_1 . If $z_i = 1$

$$P_1^i = e_1 e_3 I_{\text{Src}}^i \Delta V_i + e_2 V_{\text{ideal}}^2 C_i. \quad (31)$$

According to (26), P_1^i is equivalent to

$$P_1^i = e_1 \frac{(I_{\text{Src}}^i)^2}{C_i} + e_2 V_{\text{ideal}}^2 C_i. \quad (32)$$

If we minimize P_1^i locally by setting $\partial P_1^i / \partial C_i = 0$, we obtain

$$C_i = \frac{I_{\text{Src}}^i}{V_{\text{ideal}}} \sqrt{\frac{e_1}{e_2}}. \quad (33)$$

Therefore, according to (26) we can see that

$$\Delta V_i = \frac{I_{\text{Src}}^i}{e_3 C_i} = \frac{V_{\text{ideal}}}{e_3} \sqrt{\frac{e_2}{e_1}}. \quad (34)$$

Since e_1 , e_2 , and e_3 are constants, and V_{ideal} is common to all the converters, ΔV_i s can be assumed to be the same among the used converters if they are locally optimized. Therefore, in the following discussion, we assume $\Delta V_i = \Delta V$ for each used converter.

If all C_i s were free variables, allowed to take any value, this would not be an approximation. However, according to (30), the C_i s are not unconstrained, therefore this is an approximation.

2) *Optimizing Converter Number/Location*: As stated earlier, the number and location of the converters also affects the efficiency of the power delivery system. Distributing the converters with finer granularity and optimized floorplan over the chip can help improve the efficiency loss by reducing the voltage droop seen by the local core loads, when placing the converters closer to the utilization points. However, there is an overhead associated with the power loss in the control circuitry and clock network. In this paper, we ignore the area effect of the converters when optimizing the distribution of the converters. This is because we consider the SC converters fabricated with deep-trench capacitors, and the size of these SC converters is small compared with the size of cores in a CMP due to the high power density of deep-trench capacitors.

a) *MILP-based approach*: In this section, we present an MILP-based approach by reducing the MINLP problem in Section VI-B through a natural approximation and relaxation process.

We proceed under the assumption that for each used converter, $\Delta V_i = \Delta V$, and define

$$V_{\text{loc}} = V_{\text{ideal}} - \Delta V. \quad (35)$$

From (29), we can see that

$$V_{\text{Src}}^i \leq V_{\text{loc}}. \quad (36)$$

The loss due to voltage droop, P_2 (20), can be relaxed as

$$P_2 \leq V_{\text{loc}} \sum_{i=1}^m I_{\text{Src}}^i - \sum_{i=1}^m (S_{\text{Src}}^i V_{\text{Src}}^i) - \sum_{j=1}^n (B^j S_{\text{OBS}}^j). \quad (37)$$

In the above expression, $\sum_{i=1}^m I_{\text{Src}}^i$ is the total current delivered to the cores, and therefore, a constant. We can see that by relaxation we can transform the nonlinear cost function P_2 to be linear. In our experiments using all approaches, we find that V_{Src}^i is nearly equal for every converter i , so that (36) is in practice an equality, confirming the validity of the minimizing the relaxed P_2 .

Since $\sum_{j=1}^n (B^j S_{\text{OBS}}^j)$ is a constant, it is unchanged under any optimization. Then the relaxed power loss ($P_2 + P_3$), denoted by $P_{23,\text{rlx}}$, can be minimized by solving the following MILP problem:

$$\text{minimize } V_{\text{loc}} \sum_{i=1}^m I_{\text{Src}}^i - \sum_{i=1}^m (S_{\text{Src}}^i V_{\text{Src}}^i) + c \sum_{i=1}^m z_i \quad (38)$$

subject to the linear constraints in (23), (25), and (36).

Note that I_{Src}^i is substituted with V_{Src}^i according to (24), hence this MILP formulation has m 0-1 integer variables (z_i s), $m + 1$ continuous variables (V_{loc} and V_{Src}^i s), and $3m + n$ constraints.

3) *Optimization of Converter Size*: After determining the number and location of converters using the MILP-based approach, the second step is to determine C_i for each converter i by optimizing P_1 .

Let $I_{\text{total}} = \sum_{i=1}^m I_{\text{Src}}^i$ and $C_{\text{total}} = \sum_{i=1}^m C_i$. From (34)

$$\Delta V = \frac{I_{\text{Src}}^i}{e_3 C_i} = \frac{I_{\text{total}}}{e_3 C_{\text{total}}}. \quad (39)$$

Minimizing P_1 in (19) is thus equivalent to minimizing

$$P_1 = e_1 I_{\text{total}}^2 \frac{1}{C_{\text{total}}} + e_2 V_{\text{ideal}}^2 C_{\text{total}}. \quad (40)$$

Using (35), (40) can be further transformed to

$$P_1 = e_2 V_{\text{loc}}^2 C_{\text{total}} + I_{\text{total}}^2 (e_1 + \frac{e_2}{e_3}) \frac{1}{C_{\text{total}}} + \frac{e_2}{e_3} V_{\text{loc}} I_{\text{total}} \quad (41)$$

where I_{total} is a constant and V_{loc} can be found after solving the MILP problem (38). The constraints for the above problem are (30), and [from (28) and (39)]

$$C_{\min} = \frac{I_{\text{total}}}{e_3 \Delta V_{\max}}. \quad (42)$$

Since P_1 is a convex function of C_{total} , the optimal solution to the unconstrained problem defined in (41) is given by

$$C_0 = \frac{I_{\text{total}}}{V_{\text{loc}}} \sqrt{\frac{e_1 + \frac{e_2}{e_3}}{e_2}}. \quad (43)$$

However, this value of C_0 may fall outside the bounding constraints (30) and (42). If so, from a convexity argument, we can conclude that the optimum must be at the extreme point of the allowable C_{total} interval that is closer to C_0 . The optimal value of C_{total} , C_{opt} , is

$$C_{\text{opt}} = \begin{cases} C_{\min}, & \text{if } C_0 < C_{\min} \\ C_0, & \text{if } C_{\min} \leq C_0 \leq C_{\max} \\ C_{\max}, & \text{if } C_0 > C_{\max}. \end{cases} \quad (44)$$

We now calculate the voltage ripple ΔV using (39) and C_{opt} , and the optimal size of each used converter C_i by (39) since I_{Src}^i is known after solving the MILP problem (38).

VII. SOLUTION FOR GENERALIZED CASE WITH MULTIPLE CONVERSION RATIOS ($N \geq 2$)

The previous section considered the simplistic case where the chip is operated at a single supply voltage, and laid the basis for the solution for the general case where DVFS is used. To support DVFS, an SC converter must work with multiple conversion ratios by reconfiguring its internal topology, as presented in Section II. In this section, we discuss the solution to the efficiency optimization problem for more practical case with multiple voltage conversion ratios ($N \geq 2$), based on our discussion in Section VI for the case with single conversion ratio ($N = 1$).

A. MINLP for Multiple Conversion Ratios With $N \geq 2$

The MINLP formulation stated in Section VI-B is for the case with a single voltage conversion ratio. The formulation is modified so that each conversion ratio l has its own individual set of:

- 1) topology-dependent parameters presented in Table III, and therefore topology-dependent constants e_1 , e_2 , and e_3 in objective function (22);
- 2) constant vectors from the macromodeling of the power grid: B , S'_{Src} , and S_{OBS} that are dependent on the load current when the cores are working at a certain V_{dd} level;
- 3) design specification for the converters and core loads: ΔV_{\max} and V_{th} , that are dependent on the specific level of V_{dd} supply;

- 4) optimization variables: V_{ideal}^i , V_{Src}^i , I_{Src}^i , and ΔV_i that are also dependent on the specific level of V_{dd} supply.

For SC converters that can provide N voltage conversion ratios, we optimize the following problem:

minimize objective defined in (12)

where the loss $P_1^{(l)} + P_2^{(l)} + P_3^{(l)}$ for each conversion ratio l is given by (22).

The optimization is subject to the following.

- 1) One individual set of constraints (23)–(26) and (28) and (29) for each conversion ratio $l \in \{1, \dots, N\}$, because these constraints have either constants or variables that are dependent on the specific conversion ratio l .
- 2) Common constraints (27) and (30) for all the conversion ratios, because the size and number/location of the converters are determined at design time, and are therefore independent on the particular voltage conversion ratios. In other words, the MINLP formulation for each ratio l in Section VI-B has the same variables z_i s, that determine the number/location of the converters, same variables C_i s, that determine the size of all used converters, and same constant C_{\max} , the upperbound for total amount of usable capacitance for all the converters.

It is easy to verify that the resulting optimization problem is still a MINLP, and we can also use the two-step approach presented in Section VI-C to break it down into two subproblems. In the first, we optimize the number/location of the converters by solving an MILP problem, and then in the second, we optimize the size of each used converters using a closed-form solution. We will present the details in Sections VII-B and C.

In summary, the MINLP formulation for the generalized case with multiple conversion ratios can be derived from the MINLP for one single conversion ratio in Section VI-B by: 1) expanding the objective function to consider multiple conversion ratios and 2) then replicating part of the variables and constraints, once for each conversion ratio. After solving the resulting MINLP problem, we can find the size and number/location of used converters over all the possible conversion ratios. In reality, it is also possible for the designers to choose different weighting factors w_l s in (12) to obtain different optimal solutions of interest.

B. Optimizing Converter Number/Location

The approximation and relaxation process presented in Section VI-C can also be used for the MINLP problem defined in Section VII-A. For each voltage conversion ratio l , we first relax its power loss $P_2^{(l)}$ as shown in (37), by introducing an individual variable $V_{\text{loc}}^{(l)}$ (Section VI-C.2). Then, the part in the objective function shown in (12) that is only determined by the number/location of converters could be relaxed to be

$$\text{minimize } \sum_{l=1}^N \frac{w_l}{P_{\text{load}}^{(l)}} \cdot P_{23,\text{rlx}}^{(l)}$$

where $P_{23,\text{rlx}}^{(l)}$ is the relaxed sum of $P_2^{(l)}$ and $P_3^{(l)}$ as described in and around (38). This is still a linear objective function of $V_{\text{loc}}^{(l)}$ s, $V_{\text{Src}}^{(l)}$ s, and z_i s. The constraints can be obtained by

replicating the linear constraints in (23), (25), and (36), once for each conversion ratio l .

Then, the MILP optimization problem for N conversion ratios will have m 0–1 integer variables z_i s, $N \cdot (m + 1)$ continuous variables (one $V_{loc}^{(l)}$ and m $V_{Src}^{i(l)}$ s for each ratio l), and $N \cdot (3m + n)$ linear constraints.

C. Optimizing Converter Size

We then optimize the part in the objective function shown in (12) that is mainly determined by the size of converters as

$$\text{minimize } \sum_{l=1}^N \frac{w_l}{P_{load}^{(l)}} \cdot P_1^{(l)} \quad (45)$$

where P_1 for converter ratio l is defined as stated in (41). As before, the objective here is also a convex function of the single variable C_{total} .

The upperbound for C_{total} is still C_{max} (15), while the lower bound for C_{total} is updated to

$$C_{min}^{multi} = \max\{C_{min}^{(1)}, \dots, C_{min}^{(N)}\} \quad (46)$$

where $C_{min}^{(l)}$ is minimum total size of converters for ratio l given by (42).

Let $e_1^{(l)}, e_2^{(l)}, e_3^{(l)}, I_{total}^{(l)}$, and $V_{loc}^{(l)}$ be the coefficients and constants for ratio l as stated earlier, then the unconstrained solution to unconstrained problem defined in (45) is given by

$$C_0^{multi} = \sqrt{\frac{\sum_{l=1}^N w_l \frac{I_{total}^{(l)2}}{P_{load}^{(l)}} (e_1^{(l)} + \frac{e_2^{(l)}}{e_3^{(l)2}})}{\sum_{l=1}^N w_l \frac{e_2^{(l)} V_{loc}^{(l)2}}{P_{load}^{(l)}}}} \quad (47)$$

This is a generalized expression for the solution presented in (43).

The optimal total size of C_{total}^{multi} for all the used converters, C_{opt}^{multi} , over all the conversion ratios, is

$$C_{opt}^{multi} = \begin{cases} C_{min}^{multi}, & \text{if } C_0^{multi} < C_{min}^{multi} \\ C_0^{multi}, & \text{if } C_{min}^{multi} \leq C_0^{multi} \leq C_{max} \\ C_{max}, & \text{if } C_0^{multi} > C_{max}. \end{cases} \quad (48)$$

Then, the size for each used converter can be calculated using the same approach presented in Section VI-C.

VIII. EXPERIMENTAL RESULTS

Our two-step approach described in Sections VI and VII are implemented in C++. The MILP problem is solved using CPLEX [21].

A. Test Cases

Our approaches were exercised on two chips, one of which is a homogeneous multicore while the other is a heterogenous multicore processor.




Fig. 12. Two test cases with 16 homogeneous cores (left) and 32 heterogeneous cores (right), together with the distribution of the converters used in the results of heuristic-MILP shown in Table VI.

TABLE V
CONFIGURATIONS OF THE TWO TEST CHIPS FOR THE CASE WITH ONE SINGLE CONVERSION RATIO

Individual parameters	Homo16	Hete32	Common parameters	
$Ratio_{cvt}$	2:1	3:2	f_{sw}	100Mhz
I_{total}	16A	3.14A	N_{phase}	16
ΔV_{max}	20mV	40mV	C_{unit}	200nF/mm ²
$Area_{max}$	28.8mm ²	1.056mm ²	C_{gate}	3fF/ μ m
C_{max}	5.76 μ F	0.21 μ F	R_{on}	130 $\Omega \cdot \mu$ m
			c	10mW
			α	0.1%
			σ	512 μ m/(μ F·MHz)

1) *Homogeneous Chip*: Our homogeneous test case consists of a chip with one power domain of 16 identical cores, as shown in Fig. 12 (left), which follows the tile-based design for multicore chip [22]. Each core consists of a CPU, L1 I/D cache, and L2 cache with area ratio of 2:1:2. The core is 3×3 mm² with a peak current of 1 A at 0.6 V. In our simulations, we model the current ratio among CPU, L1 cache and L2 cache inside each core using guidelines consistent with [23].

2) *Heterogeneous Chip*: We also consider a heterogeneous test case consisting of a set of ARM Cortex cores [24]. Simpler versions of such heterogeneous cores are already on the market today [25]. This test case has one power domain of 32 cores as shown in Fig. 12 (right). Core types A through E are, respectively, the A9, A8, A5, M4, and M0 cores.

B. Effectiveness of Our Two-Step Optimization Approach

In this section, we present results to show the effectiveness of our approach presented in Section VI-C on optimizing the size and distribution of converters. For the purpose of this initial comparison, we assume that the converters are working with one single conversion ratio.

Table V shows our experimental parameters in the 32-nm technology node based on the published literature and PTM [26]. We assume the available converter area to be up to 20% of the total core area.

We have presented a MILP-based heuristic approach for the optimization of the number and location of the converters in Section VI-C. Because there is no before similar work we can compare with, we compare this approach with the following.

- 1) Manual design approach that distributes the converters over the chip at different levels of granularity with total number of converters set to be 2^k , $k = 0, 1, 2, \dots, \lfloor \log_2^m \rfloor$, where m is the numbers of candidate locations for the converters.

TABLE VI

COMPARISON OF THREE APPROACHES, WITHOUT LIMITATION ON THE NUMBER OF USABLE CONVERTERS

Chip	m	n	Manual		Greedy			Heuristic		
			#cvt	η	#cvt	η	CPU	#cvt	η	CPU
Homo16	56	208	32	84.5	26	84.8	5.8	44	85.7	353.1
Hetero32	76	203	16	83.8	11	87.2	7.3	13	88.2	362.7




Fig. 13. Comparison of power loss for three approaches, without limitation on the number of usable converters. (a) Homogeneous. (b) Heterogeneous.

- Greedy approach which explores the number and location of converters at different levels of granularity: from one converter at each candidate location, to a single lumped converter for all the cores in the chip.

For the greedy approach, we begin with a design with one individual converter at each candidate location, then at each iteration we greedily merge two neighboring converters with minimum possible increase of power loss at the next level of granularity. The increase in the power loss from combining two converters V_i and V_j into a single converter V_{ij} , is the total change in the power loss $P_2 + P_3$, which includes: 1) the change in power loss from the change in voltage droop ΔV_{droop} [(1), (2), and (10)] as $\Delta P_{L2} = \Delta V_{\text{vdd, dom}} \cdot \sum I_{\text{core}}$; 2) the change in power loss from the control circuit ΔP_{ctrl} ; and 3) the change in power loss from the clock network ΔP_{clock} . With m candidate locations, our approach will repeat the merging process $m - 1$ times to evaluate all possible levels of converter granularity.

These three approaches differ in the way to explore the distribution (number and location) of the converters over the chip. For each approach, once the best number/location of converters is found, we further optimize the size of converters using a closed-form solution as presented in Section VI-C. The results of these approaches are shown in Table VI and Fig. 13.

Table VI shows m , the numbers of candidate locations for the converters, and n , the number of observation nodes for the cores. For each approach, it shows #cvt, the total number of used converters in the solutions for each approach, and η , the system-level efficiency of the power delivery system. It also shows CPU, the runtime of greedy and heuristic approaches in seconds (on a 64-bit 2.5-GHz Intel Quad-core platform). Fig. 13 shows the breakdown of total power loss (Section V), P_1 , P_2 , and P_3 , in mW.

On average, compared with the manual design, the greedy approach can reduce P_2 (the power loss due




Fig. 14. Comparison of power loss for three approaches, with the same limitation on the number of usable converters. (a) Homogeneous. (b) Heterogeneous.

TABLE VII

COMPARISON OF OPTIMIZATION EFFICIENCY, WITH SAME LIMITATION ON NUMBER OF CONVERTERS

Chip	Max. #cvt	Manual		Greedy			Heuristic		
		#cvt	η	#cvt	η	CPU	#cvt	η	CPU
Homo16	16	16	80.3	16	81.7	2.9	16	82.1	360.4
Hetero32	8	8	84.8	8	86.6	1.7	8	87.6	374.4

to voltage droop) by 16% and total power loss by 15% with higher system-level efficiency. The heuristic approach based on MILP can reduce P_2 by $\sim 50\%$ and total power loss by 21%. The system-level efficiency is improved from 84.5% to 85.7% for the homogeneous chip and from 83.8% to 88.2% for the heterogeneous chip. The runtime of the MILP problem is tractable, it takes only few minutes for CPLEX to solve these two chips.

As stated before, the manual design has limited search space with respect to the number of converters, as compared with the greedy and heuristic approaches. For a comparison that is more favorable to the limited search space of manual design, and to explore the quality of our approach under stringent constraints, we perform another set of experiments by setting the same upperbound for the available number of converters for these three approaches. The results are presented in Table VII and Fig. 14.

Column 2 in Table VII shows the upper bound for number of usable converters. From the results, we can see that compared with manual design, on average, greedy and heuristic can still improve the results, respectively, by 12% and 17% in terms of the total power loss. This is because, for purposes of fairness, with the same number of converters, the heuristic approaches can search different combinations of the converters. Even for the homogeneous chip, there is still room for improvement because of the unevenly distribution of current within each core and the asymmetry in the power pads shared by different power domains in a single chip.

C. Optimization Over Multiple Conversion Ratios

In the previous section, we had made the temporary assumption that each converter uses a single conversion ratio. While this is useful in determining the effectiveness of our optimization methods, in a practical DVFS scenario, the assumption of a single conversion ratio is clearly invalid.




Fig. 15. Results of optimization over multiple conversion ratios on homogeneous chip.




Fig. 16. Results of optimization over multiple conversion ratios on heterogeneous chip.

In this section, we present the results for the optimization of SC converters for DVFS, over multiple voltage conversion ratios: 1:1, 4:3, 3:2, and 3:1. The values for most parameters used in the experiments are taken from Table V. The core current and maximum voltage ripple ΔV_{\max} are scaled appropriately for each conversion ratio.

Fig. 15 shows the results of optimization for the homogeneous test case shown in Fig. 12 (left). The first four bars of each ratio present the results evaluated for the solution optimized exclusively for one single conversion ratio. In other words, in objective function (12) we set all the weighting factors w_{iS} to be 0 except for the particular ratio we are interested in. As an example, the red bar of ratio 1:1 shows that if we only optimize the number/location of converters for ratio 1:1, then 56 converters are used and the peak efficiency of the whole system with the converters working under conversion ratio of 1:1 is 92%. The red bars for the other ratios show that if we use these 56 converters in the design, then the efficiency numbers of the system with converters working under other three ratios 4:3, 3:2, and 3:1 are, respectively, 83%, 82%, and 69%.

The bars represented by different colors in Fig. 15 also show that the optimal solutions for different conversion ratios are different. As we change the conversion ratio from 1:1 to 4:3, 3:2, and 3:1, the optimal number of converters used in the design reduces from 56 to 22. This is because with the same global V_{dd} supply, as we reduce the domain V_{dd} by downgrading the conversion ratios, the load power in the domain decreases, which cause the loss from voltage droop in the power grid also to decrease because of the reduced current through the power grid, therefore less converters are used in the design.

The blue bars in Fig. 15 shows that if we optimize the distribution of converters over all the four ratios [with all w_{iS}

set to be 1 in objective function (12)], then 38 converters are used. This presents a clear tradeoff among the optimization over all the four conversion ratios.

Fig. 16 shows the results for the heterogeneous test case shown in Fig. 12 (right). We can observe similar results to the homogeneous case as presented in Fig. 15. The main difference is that for the heterogeneous test case, the current load is much less than the homogeneous case, and therefore, the solutions use a much smaller number of converters.

IX. CONCLUSION

In this paper, we have studied the application and optimization of SC converters that can support DVFS in a multicore power delivery system. We first suggest distributing the SC converters over the chip to achieve better localized voltage regulation, and then develop a CAD approach to automate the design and distribution of the SC converters. We develop models for the power loss in the power delivery system as a function of size and distribution of the SC converters, and verify the accuracy of our models by simulation. We then optimize the size and distribution of SC converters to maximize the efficiency of the whole power delivery system using these converters. We show that the efficiency optimization problem for converters supporting DVFS can be formulated as an MINLP, and we propose a two-step approach to solve the MINLP to maximize efficiency over a variety of converter conversion ratios that are invoked during DVFS. The effectiveness of our approaches are demonstrated on both homogenous and heterogenous multicore chips.

ACKNOWLEDGMENT

The authors would like to thank W. H. Choi, B. Kim, and D. Jiao at the University of Minnesota, Minneapolis, MN, USA, for discussions on the verification of the power loss models in this paper.

REFERENCES

- [1] S. Borkar, "Thousand core chips: A technology perspective," in *Proc. ACM/IEEE Design Autom. Conf.*, Jun. 2007, pp. 746–749.
- [2] R. Kumar, D. M. Tullsen, N. P. Jouppi, and P. Ranganathan, "Heterogeneous chip multiprocessors," *Computer*, vol. 38, no. 11, pp. 32–38, 2005.
- [3] J. Shin, D. Huang, B. Petrick, C. Hwang, K. Tam, A. Smith, H. Pham, H. Li, T. Johnson, F. Schumacher, A. Leon, and A. Strong, "A 40 nm 16-core 128-thread SPARC SoC processor," *IEEE J. Solid-State Circuits*, vol. 46, no. 1, pp. 131–144, Jan. 2011.
- [4] J. Hart, S. Butler, H. Cho, Y. Ge, G. Gruber, D. Huang, C. Hwang, D. Jian, T. Johnson, G. Konstantinidis, L. Kwong, R. Masleid, U. Nawathe, A. Ramachandran, Y. Sheng, J. L. Shin, S. Turullois, Z. Qin, and K. Yen, "3.6 GHz 16-core SPARC SoC processor in 28 nm," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2013, pp. 48–49.
- [5] H. David, C. Fallin, E. Gorbatov, U. R. Hanebutte, and O. Mutlu, "Memory power management via dynamic voltage/frequency scaling," in *Proc. 8th ACM Int. Conf. Autonomic Comput.*, Jun. 2011, pp. 31–40.
- [6] G. Patounakis, Y. Li, and K. L. Shepard, "A fully integrated on-chip DC-DC conversion and power management system," *IEEE J. Solid-State Circuits*, vol. 39, no. 3, pp. 443–451, Mar. 2004.
- [7] R. J. Milliken, J. Silva-Martinez, and E. Sanchez-Sinencio, "Full on-chip CMOS low-dropout voltage regulator," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 54, no. 9, pp. 1879–1890, Sep. 2007.
- [8] J. Bulzacchelli, Z. Toprak-Deniz, T. Rasmus, J. Iadanza, W. Bucossi, S. Kim, R. Blanco, C. Cox, M. Chhabra, C. LeBlanc, C. Trudeau, and D. Friedman, "Dual-loop system of distributed microregulators with high DC accuracy, load response time below 500 ps, and 85-mV dropout voltage," *IEEE J. Solid-State Circuits*, vol. 47, no. 4, pp. 863–874, Apr. 2012.

- [9] S. Lai, B. Yan, and P. Li, "Stability assurance and design optimization of large power delivery networks with multiple on-chip voltage regulators," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design*, Nov. 2012, pp. 247–254.
- [10] H.-P. Le, S. R. Sanders, and E. Alon, "Design techniques for fully integrated switched-capacitor DC-DC converters," *IEEE J. Solid-State Circuits*, vol. 46, no. 9, pp. 2120–2131, Sep. 2011.
- [11] Y. K. Ramadass, "Energy processing circuits for low-power applications," Ph.D. dissertation, Dept. Electr. Eng. Comput. Sci., Massachusetts Institute of Technology, Cambridge, MA, USA, 2009.
- [12] Y. Ramadass and A. Chandrakasan, "Voltage scalable switched capacitor DC-DC converter for ultra-low-power on-chip applications," in *Proc. IEEE Power Electron. Specialists Conf.*, Jun. 2007, pp. 2353–2359.
- [13] H.-P. Le, M. Seeman, S. Sanders, V. Sathé, S. Naffziger, and E. Alon, "A 32 nm fully-integrated reconfigurable switched-capacitor DC-DC converter delivering 0.55 W/mm² at 81% efficiency," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2010, pp. 210–211.
- [14] L. Chang, R. Montoye, B. Ji, A. Weger, K. Stawiasz, and R. Dennard, "A fully-integrated switched-capacitor 2:1 voltage converter with regulation capability and 90% efficiency at 2.3 A/mm²," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2010, pp. 55–56.
- [15] S. S. Sapatnekar and H. Su, "Analysis and optimization of power grids," *IEEE Design Test Comput.*, vol. 20, no. 3, pp. 7–15, May/Jun. 2003.
- [16] (2001). *SPEC OMP2001* [Online]. Available: <http://www.spec.org/omp/>
- [17] M. M. K. Martin, D. J. Sorin, B. M. Beckmann, M. R. Marty, M. Xu, A. R. Alameldeen, K. E. Moore, M. D. Hill, and D. A. Wood, "Multifacet's general execution-driven multiprocessor simulator (GEMS) toolset," *ACM SIGARCH Comput. Archit. News*, vol. 33, no. 4, pp. 92–99, 2005.
- [18] Z. Zeng, X. Ye, Z. Feng, and P. Li, "Tradeoff analysis and optimization of power delivery networks with on-chip voltage regulation," in *Proc. 47th ACM/IEEE Design Autom. Conf.*, Jun. 2010, pp. 831–836.
- [19] M. Zhao, R. Panda, S. Sapatnekar, T. Edwards, R. Chaudhry, and D. Blaauw, "Hierarchical analysis of power distribution networks," in *Proc. ACM/IEEE 37th Design Autom. Conf.*, Jun. 2000, pp. 150–155.
- [20] M. R. Bussieck and A. Pruessner, "Mixed-integer nonlinear programming," *SIAG/OPT Newsltt., Views News*, vol. 14, no. 1, pp. 1–7, 2003.
- [21] (2011). *IBM ILOG CPLEX Optimization Studio v.12* [Online]. Available: <http://www-01.ibm.com/software/integration/optimization/cplex-optimization-studio/>
- [22] S. Bell, B. Edwards, J. Amann, R. Conlin, K. Joyce, V. Leung, J. MacKay, M. Reif, L. Bao, J. Brown, M. Mattina, C.-C. Miao, C. Ramey, D. Wentzlaff, W. Anderson, E. Berger, N. Fairbanks, D. Khan, F. Montenegro, J. Stickney, and J. Zook, "TILE64—Processor: A 64-core SoC with mesh interconnect," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2008, pp. 88–598.
- [23] S. Vangal, J. Howard, G. Ruhl, S. Dighe, H. Wilson, J. Tschanz, D. Finan, A. Singh, T. Jacob, S. Jain, V. Erraguntla, C. Roberts, Y. Hoskote, N. Borkar, and S. Borkar, "An 80-tile sub-100-W TerafLOPS processor in 65-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 43, no. 1, pp. 29–41, Jan. 2008.
- [24] (2011). *ARM Cortex Processors* [Online]. Available: <http://arm.com/products/processors/index.php>
- [25] (2011). ARM Holdings PLC. *big.LITTLE Processing*, Cambridge, U.K. [Online]. Available: <http://www.arm.com/products/processors/technologies/biglittleprocessing.php>
- [26] (2008). Predictive Technology Model. *Device Group at Arizona State University*, Tempe, AZ, USA [Online]. Available: <http://www.eas.asu.edu/~ptm>



Pingqiang Zhou received the B.E. degree from the Nanjing University of Posts and Telecommunications, Nanjing, China, in 2005, the M.E. degree from Tsinghua University, Beijing, China, in 2007, and the Ph.D. degree from the University of Minnesota, Minneapolis, MN, USA, in 2012.

He has been an Assistant Professor with the School of Information Science and Technology, ShanghaiTech University, Shanghai, China, since July 2013. Prior to joining ShanghaiTech, he was with the IBM T. J. Watson Research Center, Yorktown Heights, NY, USA, as a research intern, in 2011. He was with the University of Minnesota as a Post-Doctoral Researcher from 2012 to 2013. His current research interests include computer-aided design of VLSI circuits, multicore processors, 3-D integration circuits, and the smart grid.



Ayan Paul received the B.E. degree in electronics and telecommunication engineering from Jadavpur University, Kolkata, India, in 2005, and the M.S. degree in electrical engineering from the University of Michigan, Ann Arbor, MI, USA, in 2008. He is currently pursuing the Ph.D. degree in electrical engineering with the University of Minnesota, Minneapolis, MN, USA.

He is involved in modeling of nanoscale devices. He was with Atrenta India Pvt. Ltd., India, as an Applications Engineer, from 2006 to 2007. He was with Broadcom Corporations, Irvine, CA, USA, as an intern in 2011, where he was involved in high-speed SRAM design. His current research interests include designing high power density, high efficiency dc-dc converters for microprocessor applications.



Chris H. Kim (M'04–SM'10) received the B.S. and M.S. degrees from Seoul National University, Seoul, Korea, and the Ph.D. degree from Purdue University, West Lafayette, IN, USA.

He was with Intel Corporation, where he performed research on variation-tolerant circuits, on-die leakage sensor design, and crosstalk noise analysis. He joined the Electrical and Computer Engineering Faculty, University of Minnesota, Minneapolis, MN, USA, in 2004, where he is currently an Associate Professor. He is the author or co-author of 100+

journal and conference papers. His current research interests include digital, mixed-signal, and memory circuit design in silicon and nonsilicon (organic TFT and spin) technologies.

Prof. Kim is the recipient of the National Science Foundation CAREER Award, the McKnight Foundation Land-Grant Professorship, the 3M Non-Tenured Faculty Award, the DAC/ISSCC Student Design Contest Awards, the IBM Faculty Partnership Awards, the IEEE Circuits and Systems Outstanding Young Author Award, the ISLPED Low Power Design Contest Awards, and the Intel Ph.D. Fellowship. He has served as the Technical Program Committee Chair for the 2010 International Symposium on Low Power Electronics and Design.



Sachin S. Sapatnekar (S'86–M'93–F'03) received the B.Tech. degree from the Indian Institute of Technology, Bombay, India, the M.S. degree from Syracuse University, Syracuse, NY, USA, and the Ph.D. degree from the University of Illinois at Urbana-Champaign, Urbana, IL, USA.

He was on the faculty of the Department of Electrical and Computer Engineering, Iowa State University, Ames, IA, USA, from 1992 to 1997. Since 1997, he has been with the University of Minnesota, Minneapolis, MN, USA, where he currently holds

the Distinguished McKnight University Professorship and the Robert and Marjorie Henle Chair of electrical and computer engineering. He is the author or editor of eight books, and has published widely in the area of computer-aided design of VLSI circuits.

Dr. Sapatnekar is a recipient of the National Science Foundation CAREER Award, six conference Best Paper Awards and the Best Poster Award, the Semiconductor Research Corporation Technical Excellence Award, and the Semiconductor Industry Association University Researcher Award. He has served as a General Chair and Technical Program Chair of the ACM/EDAC/IEEE Design Automation Conference, the ACM International Symposium on Physical Design, and the IEEE/ACM International Workshop on the Specification and Synthesis of Digital Systems (Tau). He has served on the editorial boards of several publications, including the IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS (currently as an Editor-in-Chief), IEEE DESIGN AND TEST OF COMPUTERS, the IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS, and the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—PART II: EXPRESS BRIEFS.