

UC Berkeley

UC Berkeley Previously Published Works

Title

Per-Core DVFS With Switched-Capacitor Converters for Energy Efficiency in Manycore Processors

Permalink

<https://escholarship.org/uc/item/06j7h5b8>

Journal

IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 23(4)

ISSN

1063-8210

Authors

Jevtic, Ruzica
Le, Hanh-Phuc
Blagojevic, Milovan
[et al.](#)

Publication Date

2015-04-01

DOI

10.1109/tvlsi.2014.2316919

Peer reviewed

Per-Core DVFS With Switched-Capacitor Converters for Energy Efficiency in Manycore Processors

Ruzica Jevtić, *Member, IEEE*, Hanh-Phuc Le, *Member, IEEE*, Milovan Blagojević, *Member, IEEE*, Stevo Bailey, *Student Member, IEEE*, Krste Asanović, *Fellow, IEEE*, Elad Alon, *Senior Member, IEEE*, and Borivoje Nikolić, *Senior Member, IEEE*

Abstract—Integrating multiple power converters on-chip improves energy efficiency of manycore architectures. Switched-capacitor (SC) dc–dc converters are compatible with conventional CMOS processes, but traditional implementations suffer from limited conversion efficiency. We propose a dynamic voltage and frequency scaling scheme with SC converters that achieves high converter efficiency by allowing the output voltage to ripple and having the processor core frequency track the ripple. Minimum core energy is achieved by hopping between different converter modes and tuning body-bias voltages. A multicore processor model based on a 28-nm technology shows conversion efficiencies of 90% along with over 25% improvement in the overall chip energy efficiency.

Index Terms—Dynamic voltage and frequency scaling (DVFS), multicore processors, switched capacitor (SC).

I. INTRODUCTION

THE growing need for energy efficiency while utilizing increased transistor densities has led to the development of manycore architectures. To maximize the energy efficiency of a processor when using dynamic voltage and frequency scaling (DVFS) [1], it is highly desirable to independently control the supply and the clock frequency for each core [2], [3].

Manuscript received September 15, 2013; revised February 10, 2014; accepted March 19, 2014. Date of publication May 14, 2014; date of current version March 18, 2015. This work was supported by DARPA under Award HR0011-12-2-0016 and Intel ARO, in part by the Marie Curie International Outgoing Fellowship through the FP7 Programme, and in part by the Berkeley Wireless Research Center.

R. Jevtić is with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94704 USA, and also with the Technical University of Madrid, Madrid 28040, Spain (e-mail: ruzica@die.upm.es).

H.-P. Le is with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94704 USA, and also with Lion Semiconductor Inc., Berkeley, CA 94704 USA (e-mail: phucle@berkeley.edu).

S. Bailey, K. Asanović, E. Alon, and B. Nikolić are with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94704 USA (e-mail: stevo.bailey@eecs.berkeley.edu; krste@eecs.berkeley.edu; elad@eecs.berkeley.edu; bora@eecs.berkeley.edu).

M. Blagojević is with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94704 USA, also with ST Microelectronics, Crolles 38920, France, and also with the Institut Supérieur d'Electronique de Paris, Paris 75006, France (e-mail: micleon@eecs.berkeley.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.2014.2316919

As the number of cores grows, fine-grained DVFS schemes become prohibitively challenging to implement using off-chip inductor-based converters. In contrast, reconfigurable switched-capacitor (SC) dc–dc converters can be completely integrated, while offering reduced switch V – A stress and reduced overshoot [4]. Their primary disadvantage lies in the inherent SC loss caused by voltage ripple across the flying capacitors and the fact that a conventional digital system is operated based on minimum supply voltage [5]. In this paper, we show that by adapting the clock waveform to the rippling supply voltage through the use of adaptive clock schemes [6], the voltage ripple can be turned into additional performance, resulting in conversion efficiencies of 90% across a wide range of conversion ratios. In addition, by paying a modest penalty in efficiency, and operating the converters with high power densities, the area overhead can be reduced by an additional order of magnitude, thus enabling their practical implementation in fine-grained DVFS schemes with many cores.

Traditional system energy analysis assumes fixed supply voltage. We introduce the analysis of the manycore system energy when operated under changing supply voltage. The analysis helps us to perform a global optimization to find the minimum-energy operating point of a processor core for a desired application performance level. To overcome the limitation of the finite number of conversion ratios in an SC dc–dc converter, we introduce a combined technique exploiting the body-bias voltage tuning applicable to fully depleted silicon-on-insulator (FDSOI) technology [8] together with dc–dc state hopping. This technique reduces the energy per core by up to 25% compared with DVFS schemes with traditional on-chip SC voltage regulators.

II. OPTIMIZED SC DC–DC CONVERTER DESIGN

Switched power converter circuits have an inherent voltage ripple at their output, on top of which are superimposed unpredictable large voltage droops caused by switching activity of dynamic load elements, which are usually largest at the beginning of every clock cycle [7]. Because digital circuits must function correctly at the minimum possible voltage, there is often significant clock period margin added to allow for this supply noise, substantially reducing system energy efficiency.

Adaptive clocking has recently been proposed to cope with both static and dynamic variations by scaling performance with




Fig. 1. Output of the converter and adaptive clock for digital circuits.




Fig. 2. 2:1 step-down converter and the operational waveforms.

supply [6]. In this paper, we introduce an aggressive form of adaptive clocking that tracks voltage ripple at the converter's output (Fig. 1), allowing greater converter efficiency and improving energy efficiency of the entire system. Unlike [6], our adaptive clock is dynamically tracking changes in the supply within each clock cycle, and not with one clock cycle of latency.

A. Loss Optimization in Conventional DC–DC Converter

Our design is built around a reconfigurable SC dc–dc converter using MOS capacitors in a 28-nm FDSOI technology [8], [9], and it has three different configurations: 2–1 and 3–2 topologies operating off a 1 V input and a 2–1 topology operating off a 1.8 V input. To elucidate the key loss mechanisms, we begin by examining the operation of 2:1 step-down reconfigurable converter shown in Fig. 2. The specific details of the SC circuit design are provided in Section II-C. For the 2:1 conversion, the switch S3 is always OFF. The converter operates in two nonoverlapping phases ϕ_1 and ϕ_2 . The equivalent waveforms on the capacitor and at the output are shown in the same figure.

In a fully integrated SC dc–dc converter, multiple switching phases are used to reduce the output ripple [5], [14]. We are referring to this type of converter as the conventional interleaved or just conventional converter.

Optimizing the converter requires selecting the capacitor size C_{fly} , the switch size W_{sw} , and the switching frequency f_{sw} . The capacitor size is usually fixed by the chosen power density. Power density represents the ratio between the converter output power and its area, and is a useful metric for calculating converter area overhead for a given processor power.

The two remaining design parameters are obtained through the optimization of four major loss components [5]:

- 1) intrinsic SC loss $P_{C_{fly}}$, proportional to f_{sw}^{-1} , and independent of the switch size W_{sw} ;
- 2) bottom plate P_{bott} proportional to f_{sw} and independent of the switch size W_{sw} ;
- 3) switching loss P_{gate} , proportional to f_{sw} and W_{sw} ;

TABLE I

POWER BREAKDOWN IN CONVENTIONAL AND PROPOSED CONVERTERS

Approach	P_{out}	P_{cfly}	P_{cond}	P_{gate}	P_{bott}
Conventional	80%	8%	5%	3%	4%
Proposed	95%	$\approx 0\%$	2%	2%	1%

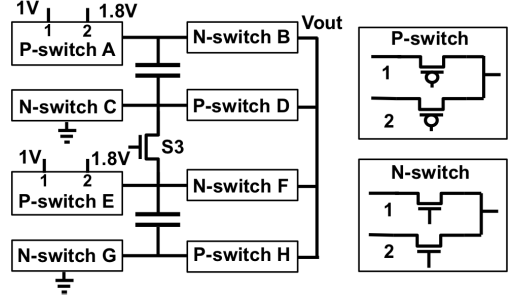


Fig. 3. Reconfigurable SC dc–dc converter.

- 4) conduction loss P_{cond} , independent of f_{sw} and proportional to W_{sw}^{-1} .

To optimize f_{sw} and W_{sw} , the sum of loss terms that are directly or inversely proportional to f_{sw} and W_{sw} is minimized.

B. Loss Optimization in Rippled DC–DC Converter

The SC loss is a consequence of charging and discharging the flying capacitor and is manifested in the ripple at the output. The performance of traditional circuits is typically set by the minimum voltage V_{min} of the supply rail. Any voltage above that will result in power loss since this additional power does not contribute to the increase in performance. By eliminating the SC loss, the converter can operate at the minimum switching frequency necessary to guarantee a maximum allowable ripple at the output (which is typically set by transistor reliability concerns) since the remaining loss terms are either directly proportional or independent of f_{sw} . This results in a much higher converter efficiency.

To illustrate the idea, the breakdown of power losses when the SC converter is optimized for maximum efficiency is presented in Table I for a conventional interleaved converter and the proposed approach. Both converters are assumed to have the same power density of 0.4 W/mm² (i.e., same flying capacitor area and the same load power), and we use 16 interleaved switching phases for the conventional converter. Loss components, inversely proportional to f_{sw} , dominate in interleaved converters. Once the ripple constraint is relaxed, f_{sw} can be scaled down in the proposed approach, resulting in substantially smaller bottom plate and switching losses. Furthermore, smaller switching losses allow for the switches to be larger, resulting also in smaller conduction loss and overall efficiency above 90%.

C. DC–DC Circuit Design

To achieve reconfigurability, we embed two identical sub-converter unit cells into one, as shown in [5]. Two sets of switches are used for better energy efficiency: one set for the configurations operating off a 1 V (set 1) and the other set

TABLE II
EMPLOYED SWITCHES FOR EACH CONFIGURATION

Conf	Input	$\phi 1$	$\phi 2$
Conf1	1.8V	A2, D2, E2, H2	B2, C2, F2, G2
Conf2	1V	A1, D1, E1, H1	B1, S3, G1
Conf3	1V	A1, D1, E1, H1	B1, C1, F1, G1




Fig. 4. Simulation waveforms for three dc–dc configurations.

for the configuration operating off a 1.8 V (set 2). A more detailed dc–dc converter circuit diagram is shown in Fig. 3.

The switches employed for each of the configuration are given in Table II. Instead of single transistors, pass gates are used for some of the switches (B1, D1, F1, and H1) to decrease their ON resistance when the output voltage is low (i.e., configuration 3). All three configurations are achieved by switching the converter in two clock phases, as explained in Section II-A. To illustrate the operation of a representative SC converter, detailed simulation waveforms for all configurations are shown in Fig. 4. The waveforms were generated by choosing 0.36 mm^2 as available area for the flying capacitor and microprocessor power consumption of 150 mW at 1 V. For these settings, we found that the ripple size varies from 100 mV for configuration 3 (bottom waveform) to 300 mV for configuration 1 (top waveform), while the switching frequency varies from 10 MHz for configuration 3 to 100 MHz for configuration 1. Small droop in voltage observed in configuration 3 is due to slower turn-ON time of pMOS switch A1 (E1) with respect to pass-gate switch D1 (H1) and occurs when the converter switches from phase $\phi 2$ to $\phi 1$.

III. SYSTEM MODELING

Detailed analysis of the entire system’s energy is important to determine the method that will be used to optimize the energy efficiency and serve as a guideline for the DVFS hardware implementation. A block diagram of a four-core processor where each core has its own DVFS scheme is shown in Fig. 5. We use MATLAB as the modeling environment to describe the behavior of each DVFS element: the processor core frequency, leakage and dynamic energy, the converter efficiency, and the clock frequency of the adaptive clock scheme. The goal of these models is to quickly analyze the




Fig. 5. DVFS on a manycore processor.




Fig. 6. Flow diagram for power and timing information extraction.

system energy for different performance targets, power supply values, ripple sizes, and converter configurations.

A. DC–DC Converter Model

The efficiencies for both our rippled and the conventional interleaved approaches are obtained using a previously presented analysis that was shown to calculate efficiencies to within a few percent compared with the measured efficiencies [5]. The rippled approach has three loss terms instead of four. The efficiencies are modeled as a function of the reference voltage, V_{ref} (marked in Fig. 1), for a particular load I_{load}

$$\eta = f(V_{\text{ref}}, I_{\text{load}}). \quad (1)$$

We also model the output waveform of the converter for a given configuration and V_{ref} to match the transistor-level schematic simulation. Equivalent RC circuits for converter switching phases are modeled with differential equations.

B. CPU Energy and Frequency Model

We use a vector processor core based on a Reduced Instruction Set Computing – V (RISC-V) Instruction Set Architecture [10] implemented in a 28-nm FDSOI technology. Using Synopsys PrimeTime (PT), we obtain power and timing reports for the core (Fig. 6).

We model the CPU frequency and CPU dynamic and leakage energies per cycle as a function of supply voltage V_{DD} [0.4 V–1.3 V], body-bias voltage V_{BB} [0–1.5 V], and temperature T [–40 °C–125 °C] by fitting PT results into a set of analytical equations. Frequency is modeled as

$$f_{\text{cpu}} = k_f \frac{(V_{\text{DD}} - V_{\text{th}})^\alpha}{V_{\text{DD}}} + k_{\text{Tc}}(V_{\text{DD}}, V_{\text{th}})(1 + k_{\text{Tf}}T) \quad (2)$$

where k_f is a proportionality constant and α is the velocity saturation term that models short-channel effects. Coefficients k_{Tf} and k_{Tc} are introduced to model the frequency's dependence on temperature. For a fixed temperature, this frequency model follows the well-known alpha power law. Drain current I_d has a positive temperature coefficient for near-threshold operation and a negative temperature coefficient for strong saturation [11]. PT simulations confirmed that the CPU performance follows the same temperature behavior.

Dynamic switching energy is given by

$$E_{\text{dyn}} = C_{\text{sw}} V_{\text{DD}}^\beta (1 + k_{\text{bb}} V_{\text{BB}})(1 + k_{T_{\text{dyn}}} T) \quad (3)$$

where C_{sw} is the total effective switched capacitance and β is a constant close to two, but is left as a fitting knob in the equation to achieve better fitting results. In addition, two more constants k_{bb} and $k_{T_{\text{dyn}}}$ are introduced to model how back bias and temperature influence the dynamic energy.

Leakage power is modeled as

$$P_{\text{leak}} = I_0 V_{\text{DD}} T^\gamma \exp\left(-\frac{V_{\text{th}}(V_{\text{DD}}, V_{\text{BB}}, T)}{n V_{\text{therm}} \ln 10}\right) \quad (4)$$

where V_{therm} is the thermal voltage, n is the subthreshold swing coefficient, and I_0 is a constant. We introduce the power supply and temperature dependence through factors V_{DD} and T^γ . While γ usually has a value of two [15], we found that better fitting results for this technology can be obtained by setting it to 1.8. Leakage energy per cycle is obtained as $P_{\text{leak}}/f_{\text{cpu}}$.

C. Adaptive Clock Model

As Fig. 1 shows, the clock period changes dynamically with the rippling supply voltage. A tunable replica circuit models the critical path of the processor to produce the dynamic CPU frequency given by (2), similar to [6]. To obtain an accurate estimate of the per-cycle energy, we model this adaptive clock on cycle-by-cycle basis.

We start with the output voltage waveform $V_{\text{DD}}(t)$, by simulating the operation of dc-dc converter, as described in Section III-A. The load for the converter is obtained using the processor model at $V_{\text{BB}} = 0$ and 25 °C, as described in Section III-B. Detailed steps for obtaining the voltage waveform are given in Section III-D. At a time t_i , the CPU clock period $T_{\text{cpu}}(t_i)$ is calculated from the reciprocal of (2) evaluated at $V_{\text{DD}}(t_i)$. Then, the next CPU clock period should be evaluated at a time $t_{i+1} = t_i + T_{\text{cpu}}(t_i)$. To account for a CPU frequency change with the supply voltage, we split the clock period into 100 time steps and calculate the incremental delay at each step following the same methodology. Averaging the incremental delays gives the total CPU clock period.

D. Per-Core DVFS Model

Given a target speed for the processor core, the complete model calculates the energy per cycle for a system that contains one processor core, the reconfigurable dc-dc converter, and the adaptive clock generator. The energy-delay (E-D) curves for the proposed approach are shown in Fig. 7.

The temperature is set to the room temperature of 25 °C, which corresponds to the temperature value that was used for dc-dc converter model characterization. Each point on the E-D curve for a particular configuration is obtained as follows.

- 1) First, we model the CPU as a voltage-controlled current source at the output of the converter. The current dependence on voltage is calculated using (2)-(4)

$$I_{\text{load}} = \frac{(E_{\text{dyn}} f_{\text{cpu}} + P_{\text{leak}})}{V_{\text{DD}}} \quad (5)$$

- 2) Then, a reference voltage value is chosen. This sets the ripple size and the switching frequency of the converter for a given load. Approximate voltage waveforms are computed based on the chosen converter configuration, V_{ref} , and I_{load} , as described in Section III-A.
- 3) The adaptive clock model is applied, and by knowing the voltage waveform, the processor's core energy is computed in each varying clock cycle. The results are averaged over time to obtain the energy per cycle and average clock period.
- 4) We divide the core energy per cycle by the converter efficiency for that particular setting, and obtain the energy per cycle of the entire system. This allows us to model many different real-world scenarios and analyze the behavior of the proposed system in a short time.

IV. OPTIMIZED DVFS SCHEME

Due to the nonlinearities of processor energy and frequency, there is a small energy overhead when the processor is supplied by a variable voltage compared with the processor supplied by a flat voltage for the same target performance. Still, this penalty is ripple size dependent and is negligible for the converter's most efficient operating point.

This loss can be described by a simple model. We assume that the CPU frequency is linearly dependent on the supply voltage and that the energy is dominated by the dynamic energy (i.e., it follows CV^2 law). We also assume a constant current load. With these assumptions, the interleaved approach has the same performance as the proposed approach if its output voltage is $V_{\text{avg}} = (V_{\text{max}} + V_{\text{ref}})/2$ (Fig. 1). The energy of the interleaved approach is then $E_{\text{interleaved}} = CV_{\text{avg}}^2$, while the energy of the rippled approach is $E_{\text{rippled}} = 1/(V_{\text{max}} - V_{\text{ref}}) \int_{V_{\text{ref}}}^{V_{\text{max}}} CV^2 dV$. Simplifying the previous equations gives an energy difference of $\Delta E = E_{\text{rippled}} - E_{\text{interleaved}} = C(V_{\text{max}} - V_{\text{ref}})^2/12$, which is only around 1% of the interleaved energy, assuming, for example, a 300 mV ripple for $V_{\text{avg}} = 0.9$ (configuration 1).

A. Energy Optimization Algorithm: V_{DD} Hopping

System software schedules tasks onto the cores and applies two different constraints to each of the DVFS blocks: the total clock cycles required to execute the task N and the desired run time for the task execution T_d .

Minimum system energy is usually found using the optimal dc-dc configuration that still meets the target frequency (e.g., the curve for configuration 3 for a 6 ns delay constraint in Fig. 7). However, a more energy-optimized




Fig. 7. E–D curves for three converter configurations and illustration of hopping between two configurations.

execution is possible by hopping between two different configurations.

Assume that N_1 clock cycles are spent in the first and N_2 cycles in the second configuration. Many different average CPU delays t_i ($i = 1, 2$) are achieved by sweeping V_{ref} for each of the configurations. The following must be true:

$$\text{a) } N_1 + N_2 = N; \quad \text{b) } N_1 \cdot t_1 + N_2 \cdot t_2 = T_d. \quad (6)$$

The total system energy is

$$E = N_1 \cdot E_1(t_1) + N_2 \cdot E_2(t_2). \quad (7)$$

To find the minimum energy, the derivatives of (7) with respect to N_1 and N_2 are set equal to zero, resulting in

$$\frac{dE_1}{dt_1} = \frac{dE_2}{dt_2} = \frac{E_1(t_1) - E_2(t_2)}{t_1 - t_2}. \quad (8)$$

The first equality forces the same slope for the tangents at t_1 and t_2 , while the second one positions them on the same line. Thus, minimum energy is achieved if the system is hopping between two states of two different configurations that lie on the common tangent of both E–D curves (Fig. 7). This narrows down the design-space exploration, since there are only two V_{ref} values in each configuration that need to be considered (one for each tangent toward adjacent configurations).

Similarly, it can be proved that the energy for hopping between three or more configurations is larger than the two-state hopping energy and should not be considered. Hopping between states incurs an energy loss during the actual state transition, but it is negligible with a large total time T_d . We found that the highest energy difference occurs when switching from configuration 3 to configuration 1 and is equal to the energy spent by the processor in eight cycles. Since we assume that the applications running on cores contain millions of cycles, this energy overhead is not significant.

B. Combined Use of Hopping and Body Biasing

The optimal point for a given workload can be reached by properly tuning all variables that are available to a designer. State-of-the-art FDSOI technology enables a much wider body-bias range and better leakage-performance tradeoffs compared with bulk CMOS because of the buried oxide that isolates the channel from the backside of the transistor [9] and a body factor of around 80 mV/V. In this paper, body bias is considered together with V_{DD} hopping to perform a sensitivity analysis [12], [13] to reach the minimum energy point.




Fig. 8. Body bias and hopping for minimum energy point.




Fig. 9. Analytical prediction of optimized power density versus efficiency tradeoff for both approaches for a 2:1 SC converter.

Our modeling framework is able to calculate hopping ratios and body-bias values for any given temperature and any performance demand to reach the minimum energy point. An example of an E–D curve for different configurations and different V_{BB} s is shown in Fig. 8. The light shaded region indicates a target performance range where hopping should be applied, while the dark shaded region shows where changing the body bias will give a more optimized design. As expected, body bias has greater influence in lower dc–dc states, where a small energy loss can be traded for a huge speed improvement.

V. RESULTS

First, we evaluate the accuracy of the proposed dc–dc converter, and then we evaluate the models proposed in Section III and calculate the energy savings of the proposed DVFS technique using the modeling framework.

A. DC–DC Converter Evaluation

The efficiency was analytically computed for two converter types: 1) one with 16 interleaved phases and 2) the proposed converter with a rippling output.

The efficiency versus power density is shown in Fig. 9. The curve for the rippled approach has a more moderate slope, resulting in a better efficiency–power density tradeoff. By paying a small penalty in efficiency, e.g., 5%, the power density of the rippled approach can be increased by a factor of two (i.e., the area of the rippled converter is half the area of the interleaved converter). Thus, the proposed converter is a better solution, particularly for manycore processors where the area overhead of the per-core dc–dc converter is an important constraint.




Fig. 10. Relative errors for energy and frequency models.




Fig. 11. E-D curves for lower power density optimization point.

B. DVFS Evaluation

1) *Model Evaluation*: Relative fitting errors for CPU frequency, CPU dynamic energy per cycle, and CPU leakage power at -40°C and 125°C are shown in Fig. 10, in the first, second, and third columns, respectively. Dynamic energy and frequency models are highly accurate with most of the relative errors below 10%. Although relative errors for leakage power go slightly above 30% for some points at -40°C , their impact on the total energy is very small, since leakage energy accounts for less than 5% of the total CPU energy at -40°C .

2) *E-D Comparison*: For this set of experiments, we chose two different power densities for the converter optimization: 1) 0.4 W/mm^2 , the point where the efficiency versus power density curve in Fig. 9 saturates, resulting in a very high efficiency for the dc-dc converter and 2) 4 W/mm^2 , where the area of the converter becomes a small overhead compared with the area of the processor. For each power density, we plot the E-D curves for both the interleaved and the proposed approach for all three dc-dc configurations using the modeling framework presented in Section III.

E-D curves for the 0.4 W/mm^2 optimization point are shown in Fig. 11. The area overhead for this point is around 25%. We assume that V_{BB} equals to 0 V and temperature equals to 25°C . The dashed and dash-dotted curve represent the minimum energy curves that are achieved through V_{DD} hopping for the rippled and interleaved approach, respectively.

The energy savings of the proposed approach for the lower configurations (i.e., 2 and 3) vary between 5% and 25% when compared with the DVFS scheme based on the interleaved approach. Due to the larger ripple size in configuration 1, the nonlinear effects described in Section IV-A are responsible for up to 8% worse energy efficiency of the rippled approach




Fig. 12. E-D curves for high power density optimization point.

in this case. However, this can be mitigated using linear regulators at the input of the rippled converter that will cause the E-D curves for the proposed approach to shift toward lower energy and higher delay values (i.e., they will shift down and to the right in the figure). Therefore, more moderate slope on the V_{DD} hopping curve will create a smaller difference in the energy of both approaches.

Fig. 12 shows E-D curves for the $10\times$ greater power density optimization point. The area overhead of the dc-dc converter is reduced to 2.5% by paying a 6% penalty in efficiency decrease for the proposed converter, as opposed to 12% for the interleaved converter. The energy savings of the proposed approach are between 0.1% and 12% when compared with the interleaved approach over a whole range of performance constraints. At high power densities, the switching frequency of the converter increases as well. However, the increase is less for the rippled approach than for the interleaved approach verifying the analysis from Section II (e.g., 30% difference for high target speed, as shown in Fig. 12). In addition, high switching frequencies are easier to achieve for the rippled converter, since the voltage regulation scheme can be much simpler when there is no need for the reduction of the voltage ripple at the output of the converter. This makes the proposed approach extremely suitable for the fine-grained DVFS schemes.

VI. CONCLUSION

We have presented a novel fine-grained DVFS technique with integrated SC converters for manycore processors. The extra power due to charging and discharging the flying capacitors is turned into additional performance, allowing extreme optimization of the SC converters that results in efficiencies over 90% for a wide range of conversion ratios. The overall system energy minima are obtained through a combined technique of body-bias voltage tuning and dc-dc state hopping, resulting in energy savings between 5% and 25% over a wide range of possible performance constraints. The approach is fully compatible with CMOS processes and can have the area overhead as low as 2.5%, which makes it suitable for practical use in manycore systems.

ACKNOWLEDGMENT

The authors would like to thank a Marie Curie International Outgoing Fellowship within the FP7 Program for the support

and would also like to thank the students, faculty, and sponsors of the Berkeley Wireless Research Center for their contributions.

REFERENCES

- [1] T. D. Burd, T. A. Pering, A. J. Stratakos, and R. W. Brodersen, "A dynamic voltage scaled microprocessor system," *IEEE J. Solid-State Circuits*, vol. 35, no. 11, pp. 1571–1580, Nov. 2000.
- [2] W. Kim, M. S. Gupta, G.-Y. Wei, and D. Brooks, "System level analysis of fast, per-core DVFS using on-chip switching regulators," in *Proc. IEEE 14th Int. Symp. HPCA*, Feb. 2008, pp. 123–134.
- [3] J. Lee and N. S. Kim, "Optimizing total power of many-core processors considering voltage scaling limit and process variations," in *Proc. 14th ACM/IEEE ISLPED*, Aug. 2009, pp. 201–206.
- [4] S. Sanders, E. Alon, H.-P. Le, M. D. Seeman, M. John, and V. W. Ng, "The road to fully integrated DC–DC conversion via the switched-capacitor approach," *IEEE Trans. Power Electron.*, vol. 28, no. 9, pp. 4146–4155, Sep. 2013.
- [5] 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.
- [6] K. Bowman *et al.*, "All-digital circuit-level dynamic variation monitor for silicon debug and adaptive clock control," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 58, no. 9, pp. 2017–2025, Sep. 2011.
- [7] T. Rahal-Arabi, G. Taylor, J. Barkatullah, K. L. Wong, and M. Ma, "Enhancing microprocessor immunity to power supply noise with clock/data compensation," in *Symp. VLSI Circuits, Dig. Tech. Papers*, Jun. 2005, pp. 16–19.
- [8] P. Magarshack, P. Flatresse, and G. Cesana, "UTBB FD-SOI: A process/design symbiosis for breakthrough energy-efficiency," in *Proc. DATE*, Mar. 2013, pp. 952–957.
- [9] B. Pelloux-Prayer *et al.*, "Planar fully depleted SOI technology: The convergence of high performance and low power towards multimedia mobile applications," in *Proc. IEEE FTFC*, Jun. 2012, pp. 1–4.
- [10] A. Waterman, Y. Lee, D. A. Patterson, and K. Asanović, "The RISC-V instruction set manual, volume I: Base user-level ISA," Dept. EECS, Univ. California, Berkeley, CA, USA, Tech. Rep. UCB/EECS-2011-62, 2011.
- [11] H. H. Chen, S. H. Tseng, and J. Gong, "The temperature-dependence of threshold voltage of N-MOSFETs with nonuniform substrate doping," *Solid-State Electron.*, vol. 42, no. 10, pp. 1799–1805, 1998.
- [12] R. W. Brodersen, M. A. Horowitz, D. Markovic, B. Nikolic, and V. Stojanovic, "Methods for true power minimization," in *Proc. Int. Conf. Comput.-Aided Des.*, Nov. 2002, pp. 35–42.
- [13] V. Zyuban and P. Strenski, "Unified methodology for resolving power-performance tradeoffs at the microarchitectural and circuit levels," in *Proc. Int. Symp. Low Power Electron. Des.*, 2002, pp. 166–171.
- [14] D. Somasekhar *et al.*, "Multi-phase 1 GHz voltage doubler charge pump in 32 nm logic process," *IEEE J. Solid-State Circuits*, vol. 45, no. 4, pp. 751–758, Apr. 2010.
- [15] K. Roy S. Mukhopadhyay, and H. Mahmoodi-Meimand, "Leakage current mechanisms and leakage reduction techniques in deep-submicrometer CMOS circuits," *Proc. IEEE*, vol. 91, no. 2, pp. 305–327, Feb. 2003.



Ruzica Jevtić (M'13) received the B.S. degree in electrical engineering from the University of Belgrade, Belgrade, Serbia, and the Ph.D. degree in electrical engineering with the European Ph.D. mention from the Technical University of Madrid, Madrid, Spain, in 2004 and 2009, respectively. Her Ph.D. work was oriented toward CAD tools for high-level modeling, power estimation, measurements, and architecture design for high-speed computational systems in FPGAs.

She was a Post-Doctoral Researcher with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, USA, from 2011 to 2013, where she was involved in energy-efficient microprocessor design. She is currently a Researcher with the Technical University of Madrid.

Dr. Jevtić was a recipient of the FP7 Marie Curie International Outgoing Fellowship.



Hanh-Phuc Le (M'13) received the B.S. degree in electrical engineering from the Hanoi University of Technology, Hanoi, Vietnam, the M.S. degree from the Korea Advanced Institute of Science and Technology, Daejeon, Korea, and the Ph.D. degree from the University of California, Berkeley, CA, USA, in 2003, 2006, and 2013, respectively.

He was with the Vietnam Academy of Science and Technology, Hanoi, Korea Expressway Corporation, Seongnam, Korea, LG Corporation, Seoul, Korea, JDA Technologies, Daejeon, Korea, Sun Microsystems, Inc., Santa Clara, CA, USA, Intel Corporation, Santa Clara, and Rambus Inc., Sunnyvale, CA, USA. He co-founded Lion Semiconductor Inc., Berkeley, where he is currently the CTO. His current research interests include circuit designs for power electronics and telecommunications, focusing on SMPS, fully integrated conversions for high-performance digital ICs, control methodology, and mix-signal integrated circuits.



Milovan Blagojević (M'12) received the B.Sc. and M.Sc. degrees in electrical engineering from the School of Electrical Engineering, University of Belgrade, Belgrade, Serbia, in 2010 and 2012, respectively.

He interned with HW/SW Co-Design Group, Institute for Informatics, Erlangen, Germany, in 2010, and with Intel, Belgrade, in 2011, where he was involved with the image processing team under ultra mobility group. In 2012, he enrolled in a CIFRE Ph.D. program that is realized in cooperation among three institutions, the Berkeley Wireless Research Center, Berkeley, CA, USA, STMicroelectronics, Crolles, France, and Institut Supérieur d'Electronique de Paris, Paris, France. He is involved in research and development of HDR algorithm implemented on a specific VLIW processor. His current research interests include energy-performance optimization in modern digital systems, with emphasis on advantages of new UTBB FDSOI technology.



Stevo Bailey (S'11) received the B.S. degrees in engineering science and physics from the University of Virginia, Charlottesville, VA, USA, in 2012. He is currently pursuing the Ph.D. degree with the University of California, Berkeley, CA, USA.

He interned at Jefferson Laboratories, Newport News, VA, USA, through the Virginia Microelectronics Consortium and Old Dominion University, Norfolk, VA, USA, during his undergraduate studies. He was involved in fault-tolerant reconfigurable adder designs for future nanoelectronic systems. He is currently involved in exploring soft-error resilient logic design techniques and researching automation of these techniques. In 2012, he joined the Berkeley Wireless Research Center, Berkeley. His current research interests include robust and power-efficient processor and ASIC design.



Krste Asanović (S'90–M'98–SM'12–F'14) received the B.A. degree in electrical and information sciences from Cambridge University, Cambridge, U.K., and the Ph.D. degree in computer science from the University of California, Berkeley, CA, USA, in 1987 and 1998, respectively.

He was an Assistant and Associate Professor of Electrical Engineering and Computer Science with the Massachusetts Institute of Technology, Cambridge, MA, USA, from 1998 to 2007. He is currently a Professor with the Electrical Engineering and Computer Sciences Department, University of California, Berkeley. His current research interests include computer architecture, VLSI design, and parallel programming and run-time systems.

Dr. Asanović is an ACM Distinguished Scientist.



Elad Alon (SM'12) received the B.S., M.S., and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, USA, in 2001, 2002, and 2006, respectively.

He joined the University of California, Berkeley, CA, USA, in 2007, where he is currently an Associate Professor of Electrical Engineering and Computer Sciences and the Co-Director of the Berkeley Wireless Research Center. He has held consulting or visiting positions at Wilocity, Sunnyvale, CA, USA, Cadence, San Jose, CA, USA, Xilinx, San Jose,

Oracle, Redwood City, CA, USA, Intel, Santa Clara, CA, USA, Advanced Micro Devices, Sunnyvale, Rambus, Inc., Sunnyvale, Hewlett Packard, Palo Alto, CA, USA, and IBM Research, Armonk, NY, USA, where he was involved in digital, analog, and mixed-signal integrated circuits for computing, test and measurement, and high-speed communications. His current research interests include energy-efficient integrated systems, including the circuit, device, communications, and optimization techniques used to design them.

Dr. Alon was a recipient of the 2008 IBM Faculty Award, the 2009 Hellman Family Faculty Fund Award, and the 2010 UC Berkeley Electrical Engineering Outstanding Teaching Award, and has co-authored papers that received the 2010 ISSCC Jack Raper Award for Outstanding Technology Directions Paper, the 2011 Symposium on VLSI Circuits Best Student Paper Award, and the 2012 Custom Integrated Circuits Conference Best Student Paper Award.



Borivoje Nikolić (S'93–M'99–SM'06) received the Dipl.Ing. and M.Sc. degrees in electrical engineering from the University of Belgrade, Belgrade, Serbia, in 1992 and 1994, respectively, and the Ph.D. degree from the University of California, Davis, CA, USA, in 1992, 1994, and 1999, respectively.

He joined the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, USA, in 1999, where he is currently a Professor. His current research interests include digital, analog, and RF integrated circuit design,

and VLSI implementation of communications and signal processing algorithms. He has co-authored the book *Digital Integrated Circuits: A Design Perspective—2nd Edition* (Prentice-Hall, 2003).

Dr. Nikolić was a recipient of the NSF CAREER Award in 2003, the College of Engineering Best Doctoral Dissertation Prize and the Anil K. Jain Prize for the Best Doctoral Dissertation in Electrical and Computer Engineering from the University of California at Davis in 1999, and the City of Belgrade Award for the Best Diploma Thesis in 1992. He was also a recipient of the Best Paper Awards at the IEEE International Solid-State Circuits Conference, the Symposium on VLSI Circuits, the IEEE International SOI Conference, the European Solid-State Device Research Conference, and the ACM/IEEE International Symposium of Low-Power Electronics, for working with his students and colleagues.