

A Fully Integrated Reconfigurable Switched-Capacitor DC-DC Converter With Four Stacked Output Channels for Voltage Stacking Applications

Tao Tong, Sae Kyu Lee, Xuan Zhang, David Brooks, *Fellow, IEEE*, and Gu-Yeon Wei

Abstract—This work presents a fully integrated 4-to-1 DC-DC symmetric ladder switched-capacitor converter (SLSCC) for voltage stacking applications. The SLSCC absorbs inter-layer load power mismatch to provide minimum voltage guarantees for the internal rails of a multicore system that implements four-way voltage stacking. A new hybrid feedback control scheme reduces the voltage ripple across stacked voltage layers for high levels of current mismatch, a condition that exacerbates voltage noise in conventional SC converters. Furthermore, the proposed SLSCC dynamically allocates valuable flying capacitor resources according to different load conditions, which improves conversion efficiency and supports more power mismatch between the layers. Implemented in TSMC's 40G process, the SLSCC converts a 3.6 V input voltage down to four stacked output voltage layers, each nominally at 900 mV.

Index Terms—DC-DC converter, fully integrated voltage regulator, hybrid feedback, switched-capacitor, voltage stacking.

I. INTRODUCTION

POWER delivery has been a challenging issue for multicore SoC applications. The decreasing supply voltages as well as the increasing supply currents of the processors exacerbate losses in the off-chip voltage regulator modules and power delivery network. On the other hand, on-chip voltage regulators typically have low efficiencies at high conversion ratios (e.g., 4-to-1) unless ultra-high-quality on-chip capacitors or inductors are used [7-18].

Recent work has proposed *voltage stacking* as an alternative on-chip power delivery solution [1-6]. Rather than delivering current to all cores in parallel, voltage stacking vertically connects the cores in serial layers. For the same power, a single high voltage supply supplies a proportionally lower level of current to the chip, recycled through the cores in the stacked voltage layers. If all stacked layers consume the same current, internal rail voltages should distribute evenly. Unfortunately,

any load power mismatch between layers directly translates to inter-layer voltage noise. This motivates using a fully integrated voltage regulator to compensate for load power mismatch between the stacked layers.

Prior work has proposed several fully integrated voltage regulators for differential power processing in voltage stacking applications. For example, push-pull linear regulators have been used to provide voltage regulation for stacked outputs [1-2]. Although linear regulators have small area overhead and are easy to integrate, their inherently low conversion efficiency limits the power delivery efficiency of the overall stacking system. Alternatively, a 2-to-1 switched-capacitor (SC) converter was implemented to regulate the intermediate voltage between two stack layers [4]. To support more than two stack layers, multiple 2-to-1 SC converters can be used to regulate the internal rails [5]. For an N -layer voltage stacking system, this multi-stage solution needs a total of $N-1$ 2-to-1 SC converters, resulting in many switches on the power train and complicating the design of the control loop. Lastly, inductive converters have also been proposed as off-chip solutions for differential power processing [6]. However, it is more difficult to integrate high-quality on-chip inductors than on-chip capacitors.

This paper describes a fully integrated 4-to-1 SC converter that absorbs inter-layer load power mismatches to maintain minimum voltage levels for the stacked voltage domains of a multicore system that implements four-layer voltage stacking in TSMC's 40G CMOS process. The integrated converter implements a symmetric ladder SC converter (SLSCC) topology [7]. By tapping into the internal rails of the symmetric ladder (as shown in Fig. 1), the SLSCC can neutralize mismatched load currents between layers. Thanks to the ladder topology, none of the power switches or flying capacitors is exposed to high voltages and can, therefore, be implemented with native thin-oxide devices, which improves conversion efficiency and power density. The SLSCC operates off a 3.6 V input voltage and the nominal voltage of each voltage stack layer is 900 mV, which is the nominal operating voltage of the transistors. While this paper presents results for a symmetric ladder SC topology with four-layer stacking, many of the conclusions and findings of this work also applies to other SC converter designs

Manuscript received November 4, 2015; revised April 27, 2016; accepted May 31, 2016. Date of publication July 19, 2016; date of current version September 1, 2016. This paper was approved by Associate Editor Pavan Kumar Hanumolu.

T. Tong, S. K. Lee, D. Brooks, and G. Y. Wei are with Harvard University, Cambridge, MA 02138 USA.

X. Zhang is with Washington University in St. Louis, St. Louis, MO 63130 USA.

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/JSSC.2016.2580598

0018-9200 © 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

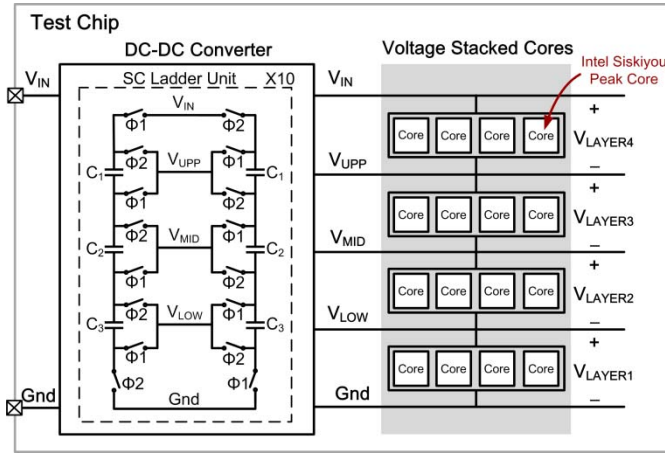


Fig. 1. Voltage stacking system diagram showing the SLSCC and the 16 four-layer stacked cores.

with different topologies and different numbers of stacked layers.

In the test-chip prototype, the SLSCC connects to four stacked layers of 16 Siskiyou Peak microcontroller cores from Intel [19]. These cores were codesigned with the SLSCC to operate in a voltage-stacked system and leverage adaptive frequency clocking to maximize throughput and energy efficiency. Detailed analysis of the SLSCC show that charge flow depends on layer-to-layer load conditions. Hence, the proposed SLSCC dynamically allocates valuable flying capacitor resources depending on different load conditions. This improves conversion efficiency and allows for larger amounts of power mismatch between the layers. The SLSCC also employs a hybrid feedback scheme to ensure minimum voltage guarantees across the four stack layers simultaneously and to reduce voltage ripple for high levels of power mismatch—a condition that exacerbates voltage ripple in conventional SC converters. With voltage stacking, conversion losses only apply to inter-layer mismatched power provided by the SLSCC while recycled current flows efficiently through the entire stack. Therefore, measurement results show the average power delivery efficiency with voltage stacking is as high as 87% across a wide range of layer-to-layer load conditions.

The remainder of this paper is organized as follows. Section II describes the basic operations of the SLSCC in the voltage stacking application and discusses the optimization of the SLSCC. Section III presents several design techniques, such as flying capacitor reconfiguration, that improve the performance of the converter. Section IV goes through the proposed hybrid feedback control scheme. Finally, measurement results are presented in Section V.

II. A DC-DC SLSCC IN A VOLTAGE-STACKED SYSTEM

A. System Overview

Fig. 1 presents an overview of the voltage-stacked system implemented in a 40 nm digital CMOS process. A total of 16 Intel Siskiyou Peak microcontroller cores are configured in a 4x4 stacked array. The cores can operate in one of two clocking modes: 1) fully synchronous global fixed-frequency clocking mode (FFClk) or 2) per-layer adaptive frequency clocking

mode (AFCIk). The FFCIk is generated from a single external clock source. A digitally configurable free-running ring oscillator on each layer generates the AFCIk, wherein the oscillator frequency tracks per-layer voltage fluctuations.

Connected to a 3.6 V input voltage V_{IN} , the four stack layers nominally divide to 900 mV each. However, inter-layer power mismatch can lead to large inter-layer voltage mismatch. Therefore, a fully integrated SLSCC, implemented in the same chip with the cores, ensures minimum voltage levels across the stack layers of the stacked system. Depending on which layer consumes more or less current, the SLSCC either pushes current to the stacked cores or shunts current away from them through the internal V_{UPP} , V_{MID} , and V_{LOW} rails. The SLSCC consists of ten switched-capacitor ladder units, each controlled by one of ten interleaved switching signals. Interleaving reduces output voltage ripple.

A limitation of the SLSCC in voltage stacking applications, however, is that it cannot “regulate” the intermediate rails to exactly even distributions of V_{IN} . This is because switched-capacitor converters rely on a voltage difference between the flying capacitors and the output to deliver charge. Therefore, instead of guaranteeing a regulated stack layer voltage, the SLSCC implements single-bound feedback control that keep each stack layer voltage above a prescribed reference voltage. This control scheme guarantees a minimum performance level for the cores and prevents minimum voltage reliability issues such as SRAM instability. However, since the layer voltages must sum up to V_{IN} (e.g., 3.6 V), this control scheme cannot prevent voltage droop in one stack layer setting the worst-case condition for a globally shared clock frequency, which translates to excess voltage margin for other stack layers. To address this inefficiency, the test chip implements per-layer adaptive frequency clocking, which can convert excess voltage margins into higher performance. The overall voltage stacking system shown in Fig. 1, and its merits are explored in more detail in [19]. This paper focuses on the implementation and measurement results of the SLSCC itself.

B. Losses and Optimizations of the SLSCC

The way the SLSCC delivers power to the load in this voltage stacking application is very different from a conventional power delivery system, because the load circuits are spread across multiple stacked voltage layers. Charge flow in the SLSCC depends heavily on the layer-to-layer load conditions and the SLSCC has different performance characteristics, such as efficiency and maximal supported power mismatch, for delivering power to different output layers. Fig. 2 presents examples of charge flow analyses for different load conditions. For a fair comparison, the total load current in all scenarios is the same (4q). Based on this charge flow analysis, we can conclude that: 1) the charge flow through the capacitors and power switches depend on which layer the SLSCC delivers power to and 2) distributing the load current from a single layer to multiple layers reduces losses. Understanding charge transfer flow also informs guidelines for optimizing the efficiency and power delivery range of the SLSCC. Fig. 3 presents the optimized flying capacitor size and the optimized power

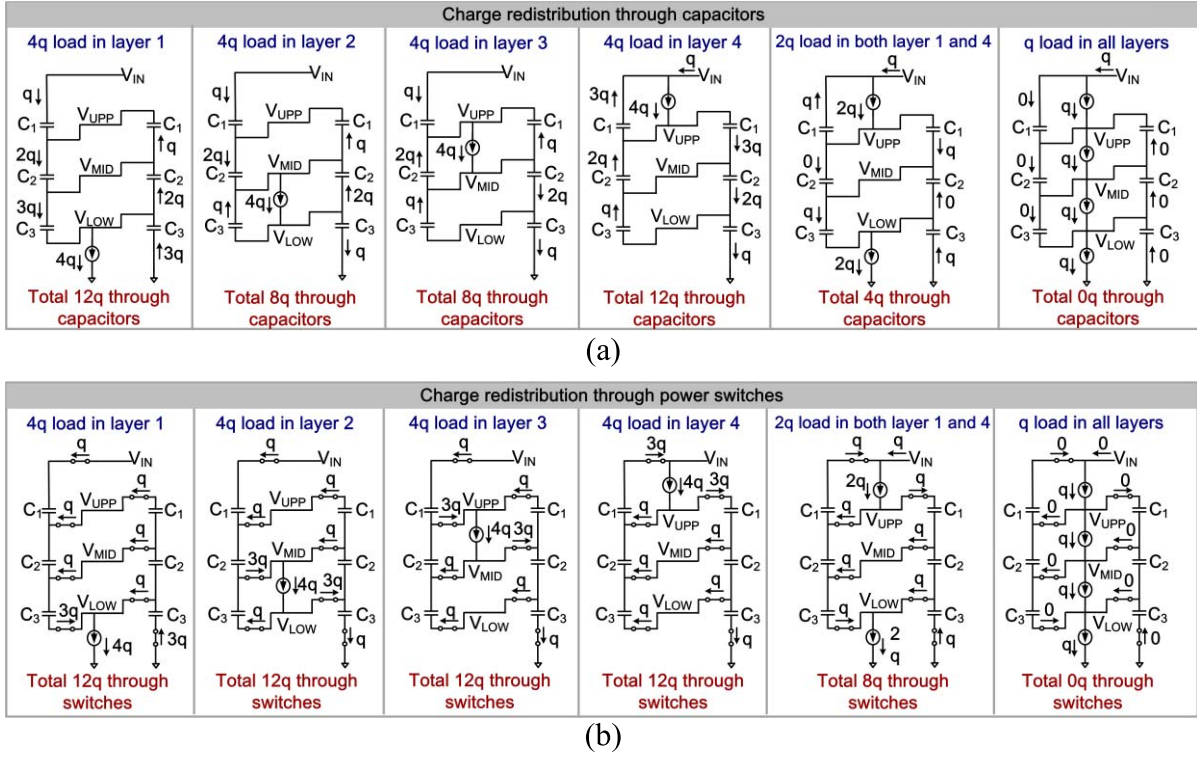


Fig. 2. Internal charge flow diagrams in the SC ladder for different load conditions. (a) Charge flow through flying capacitors. (b) Charge flow through power switches.

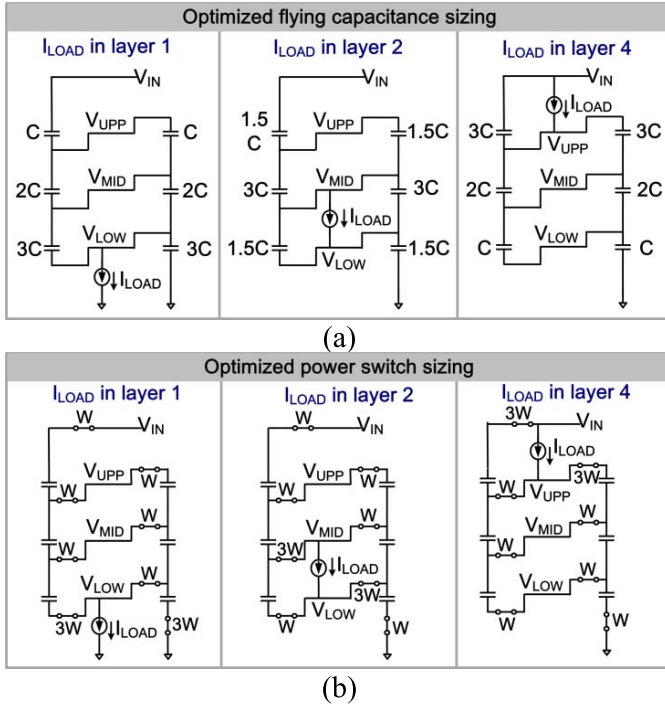


Fig. 3. Optimized SLSCC for different load conditions with (a) optimized flying capacitors and (b) optimized power switches

switch size for three different load conditions.

In this design, the SLSCC operates in the “slow-switching limit” (SSL) mode [7], where the flying capacitors, rather than the power switches, dominate conductive loss. The switching




Fig. 4. Transistor implementation of the SC ladder.

loss in this SLSCC is similar to that of a typical SC converter in conventional applications where only the bottom layer consumes load current. Switching loss has been well studied in [7, 17-18].

III. IMPLEMENTATION OF THE SC LADDER

A. Implementation of the Building Blocks in the SLSCC

Fig. 4 shows the implementation of the SC ladder unit. The power switches in the main SC ladder are implemented with thin-oxide NMOS or PMOS transistors. The flying capacitors rely on thin-oxide PMOS transistors, which have 20% smaller capacitor density, but only 1/6 of the leakage current of NMOS




Fig. 5. Implementation of capacitor-coupled level shifter.

transistors. Similar to [7], triple-well is used to isolate the body of NMOS transistors in different stacked layers and to alleviate voltage breakdown issues.

The thin-oxide power switches in the SLSCC also operate in multiple stacked voltage domains. Hence, level shifters are required to connect signals across different voltage domains. Fig. 5 presents the implementation of a capacitor-coupled level shifter that translates a signal from low to high voltage domains. Since the nominal voltage across the inverters is approximately 0.9V, thin-oxide CMOS transistors can be used. However, the voltage across the coupling capacitor depends on the difference between the two voltage domains. To go from the lowest ($Gnd \sim V_{LOW}$) to the highest voltage domain ($V_{UPP} \sim V_{IN}$), the voltage across the capacitors can be up to $\sim 2.7V$, much higher than the gate-oxide breakdown voltage of the transistors available. Hence, the capacitors in the level shifters use metal-oxide-metal (MOM) capacitors.

B. Flying Capacitance Reconfiguration Scheme

Since the best allocation of flying capacitance depends on layer-to-layer load conditions in this voltage stacking system, it is preferable for the SLSCC to be able to dynamically modify the sizes of its flying capacitors according to load current conditions. The power switch width is not reconfigured in this prototype, since SLSCC conductive loss is dominated by the flying capacitors. Based on simulation results, the conversion efficiency improves by less than 3% if the power switches are also reconfigurable.

Fig. 6 shows the implementation of the reconfigurable SC ladder, which consists of one nonconfigurable main SC ladder (shown with a gray background in Fig. 6) and four sets of reconfigurable cap-bank units that are connected to the main SC ladder. Each cap-bank set contains three identical cap-bank units that can individually configure their connections. Each of the reconfigurable capacitors in the cap-bank units can be connected in parallel with different capacitors in the main ladder by closing either SW1 or SW2. In this way, the capacitance is reallocated dynamically according to the load conditions. The switches (SW1 and SW2) in the cap-bank units only switch ON/OFF when the capacitors need to be reconfigured.

Implementation of the switch network in the cap-bank units requires careful balance between numerous tradeoffs. These additional switches in the cap-bank units add conductive loss

and switching loss to the converter. To minimize losses, the design relies on a pair of thin-oxide flying inverters to implement the reconfigurable switches SW1 and SW2, as shown in the right half of Fig. 6. In each paired SW1 and SW2, the gates are connected together, driven by another small flying inverter. By using this design, rather than connecting the gates to a fixed voltage to turn SW1 and SW2 ON/OFF, the gate voltage switches together with V_H , V_M , and V_L when the main SC ladder is switching. Thus, SW1 and SW2 can be implemented using thin-oxide transistors to reduce associated conductive and switching losses. Since load conditions typically fluctuate at a lower rate than the main switching frequency of the converter, the switching losses associated with SW1 and SW2 are small and justified by the efficiency improvements that configurability offers. A total of 12C is used in each of the 10 interleaved SC ladder units, where 1 C equals 37.5 pF, for a total capacitance of 4.5 nF.

In this prototype, the load condition detection is performed offline. A MATLAB algorithm, based on the charge flow analysis, determines the optimal flying capacitor configuration. The switch network in the cap-bank is then programmed from off-chip.

IV. RIPPLE-REDUCED HYBRID FEEDBACK CONTROL

Fig. 7 shows the system block diagram for the proposed SLSCC. Hybrid feedback control circuitry operates off of a clock from a local voltage-controlled ring oscillator (VCRO). The feedback circuitry monitors the voltages across each output layer and generates a 10-phase interleaved switching signal SW_{HYBRID} . The power switch control signal generator then creates the switching signals, Φ_1 and Φ_2 , for the ten interleaved reconfigurable SC ladder units. The level shifters shift these switching signals to the correct voltage domains and drive the corresponding switches in the SC ladder.

The hybrid feedback control loop in this converter is composed of a primary single-bound control loop that shuffles charge around to maintain a minimum voltage level for each of the stacked layers and a secondary proactive loop, which helps reduce voltage ripple for heavily mismatched load conditions.

A. Primary Single-Bound Control loop

In this design, the primary feedback loop tries to keep all output layer voltages above a predefined reference voltage level, as opposed to regulating the layer voltages to the reference voltage level. In a voltage stacking application, all of the output layer voltages must add up to the input voltage, which is 3.6 V in this design. If the load is identical for all layers, per-layer voltage evenly divides to 900 mV across all layers. On the other hand, mismatch in load activity between layers leads to per-layer voltage deviations away from 900 mV due to Kirchhoff's voltage and current laws. If powered independently, rising load activity would translate to higher load current for a particular supply voltage. With voltage stacking, however, the same level of current must flow through each layer, which therefore translates to voltage fluctuations in each layer. If the mismatch is sufficiently large, per-layer voltages will fall below the reference voltage, the

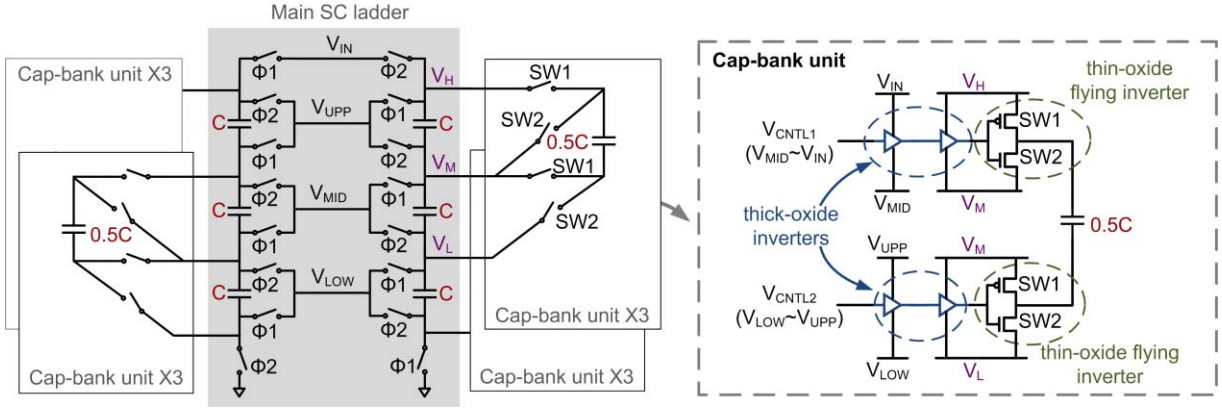


Fig. 6. Implementation of the reconfigurable SC ladder

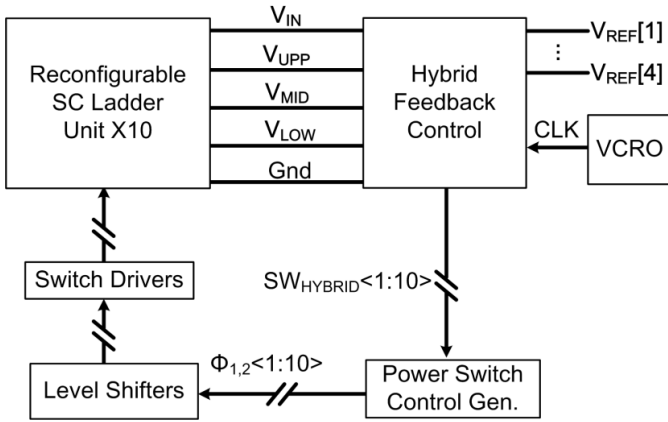


Fig. 7. Block diagram of the SLSCC.

feedback loop in the SLSCC detects that one of the layer voltages is lower than the reference, and the SC units switch to restore voltage levels above the reference level. In other words, the switching behavior of the SLSCC redistributes charge across the stack to keep all layer voltages above the reference voltage.

Fig. 8 illustrates the implementation of the primary feedback control loop. Four 2.5 GHz clocked comparators compare each layer voltage to corresponding per-layer reference voltages generated on-chip. If the voltage of any layer falls below the reference, the associated comparator generates a pulse. The primary feedback control logic combines the outputs of all the comparators in different voltage domains and generates a single high frequency switching signal, $COMP_{TRIG}$. When any of the four comparators generates a pulse, $COMP_{TRIG}$ also generates a pulse. $COMP_{TRIG}$ is further processed by the secondary proactive loop (discussed below) and is turned into 10 interleaved slow switching signals by a barrel shifter. The interleaved switching signals eventually drive the switches in the SC ladder, as shown in Fig. 7. Each interleaved SC ladder unit switches at a maximum frequency of 250 MHz.

The per-layer reference voltage can be set by application requirements and presents a tradeoff between SLSCC power delivery capability and tolerable voltage droop. Increasing the

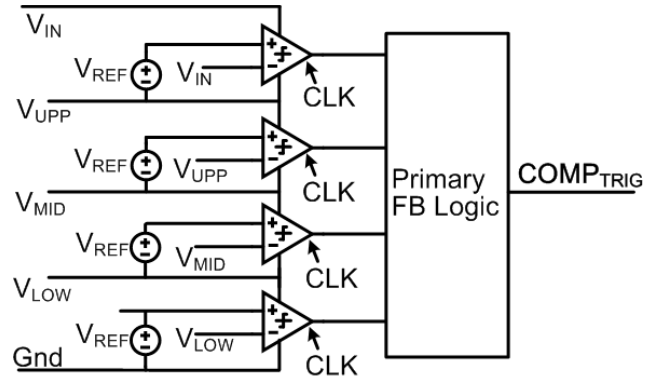


Fig. 8. Implementation of the primary single-bound control.

reference voltage reduces voltage droop, but it also reduces the maximal supported mismatch power, with diminishing returns as the reference voltage approaches 900 mV [4], [5]. In contrast, if the reference voltage is set too low, inter-layer voltage differences can be correspondingly large. Since our SLSCC only keeps layer voltages above the lower bound, there is also the possibility that one or more layer voltages exceed maximum voltage limits. One straightforward solution to resolve this issue would be to also impose an upper bound to the control loop such that the SC ladder switches whenever layer voltages exceed upper or lower bounds, thereby keeping each layer voltage within a predetermined range.

B. Secondary Ripple-Reduced Proactive Loop

One of the major advantages of single-bound control is its fast response for handling large current steps. It can change the effective switching frequency of an SC converter from a very low frequency to its maximum frequency within a few nanoseconds [11], [13], [15]. However, due to nonzero feedback latency and pulse-skipping nature of the control loop, single-bound control loops typically result in much larger static voltage ripple compared to voltage-controlled-oscillator (VCO) based pulse frequency modulation (PFM) loops [10], [15].




Fig. 9. Diagram of typical voltage noise in a SC converter using single-bound control.

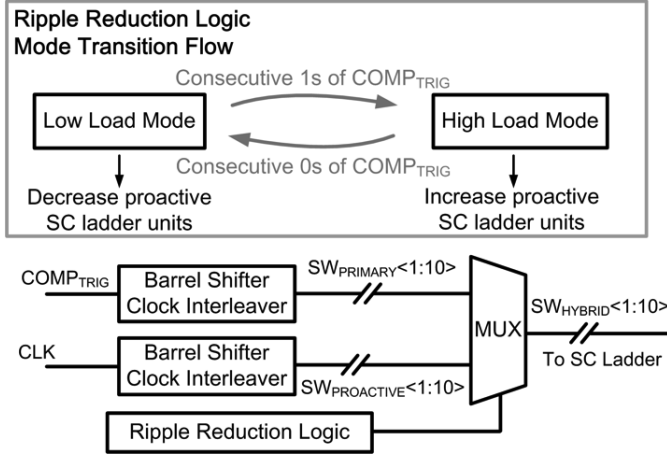


Fig. 10. Implementation of the proactive feedback control.

Techniques such as switch conductance modulation [16], [20] have been proposed to improve efficiency and reduce voltage ripple at light loads. Interleaved designs also reduce voltage ripple at light loads. However, the ripple at heavy loads can also be very large [13]. Fig. 9 illustrates a typical output voltage waveform for an SC converter that relies solely on single-bound control. Under heavy load conditions, the load current quickly discharges V_{OUT} before the loop can detect and react, resulting in large ripple. The larger the delay is, the larger the ripple will be. In this design, the feedback delay is about 1.5 ns.

To address the issue of large ripple associated with single-bound control, we added a secondary proactive loop to reduce voltage ripple for heavily mismatched load conditions. To support the ripple reduction feature, all of the SC ladder units are dynamically divided into two groups. One group is controlled by the primary single-bound control loop. The other ladder units, called proactive units, are controlled by the secondary loop. By detecting load conditions, the secondary proactive loop tells the proactive units to always switch at the maximum rate (250 MHz). These proactive units that constantly switch reduce the amount of load power that the primary feedback loop has to handle, thus reducing voltage droop. As shown in Fig. 10, ripple reduction logic monitors consecutive 1's and 0's in $COMP_{TRIG}$ to dynamically allocate SC ladder units between single-bound and proactive control. If several consecutive 1's are detected, more SC ladder units become proactive units. If several consecutive 0's are detected, the number of proactive units is reduced. This detection scheme adds hysteresis to the secondary loop.




Fig. 11. Implementation and characterization of the on-chip load generator.

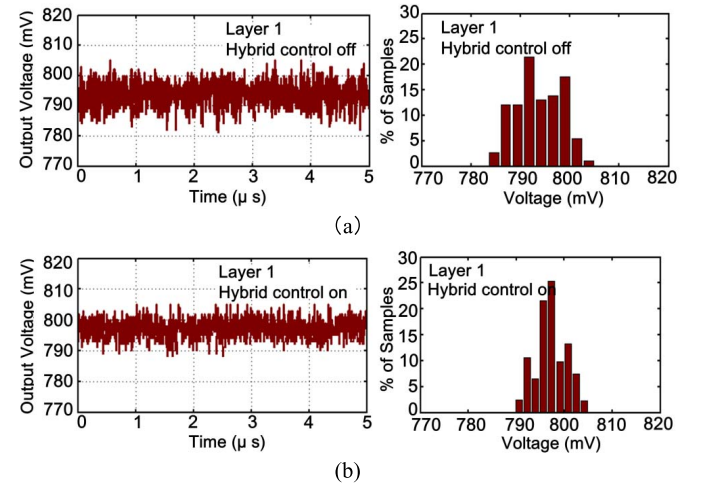


Fig. 12. Measured transient waveforms and histograms of V_{LAYER1} with load current only in Layer 1 with (a) hybrid control OFF and (b) hybrid control ON ($P_{OUT} \approx 16$ mW, all reference voltages are set to 800 mV).

In a very heavy load condition, most of the SC ladder units are proactive units switching at peak frequency to deliver the needed power while also reducing voltage ripple. All the proactive units are interleaved to further reduce voltage ripple.

V. MEASUREMENT RESULTS

This section presents measurement results for the test-chip prototype implemented in TSMC's 40 G triple-well CMOS process. The test setup has off-chip capacitors to bypass V_{IN} , but no external capacitors connect between the internal rails. To fully characterize the SLSCC, we primarily rely on on-chip load current generators to measure conversion efficiency and transient response across a broad range of different load conditions in the stacked layers. We use the 4x4 voltage-stacked array of Intel Siskyu Peak processor cores to demonstrate the




Fig. 13. Measured transient responses with dynamic load currents in multiple layers. ($V_{IN} = 3.6$ V, all reference voltages are set to 800 mV.)

advantages of adaptive frequency clocking in voltage-stacked systems.

Fig. 11 shows the implementation of the load generators and measured electrical characteristics. Each layer has an identical but individually configurable load generator array, consisting of six binary weighted NMOS transistors. These transistors create load currents in the stacked layers when they are ON, by driving the transistor gate with the local supply voltage level. Each array is controlled by a programmable linear feedback shift register (LFSR), also implemented on-chip. The current of the load generator is a function of the transistors that are turned ON as well as the voltage across them. The bottom part of Fig. 11 plots the measured load current of a 4x transistor versus supply voltage.

A. Voltage Ripple and Transient Response

Fig. 12 plots the transient waveforms and histograms of V_{LAYER1} when only the bottom layer (Layer 1) consumes about 20 mA of static current, with the proposed hybrid control turned ON and OFF. In both plots, V_{LAYER1} stays around 800 mV, regulated by the feedback loop. When the hybrid control is turned OFF, there is peak-to-peak voltage ripple of 25 mV [Fig. 12(a)]. The hybrid feedback control scheme reduces the voltage ripple by 30%, to 18 mV [Fig. 12(b)].

Fig. 13 presents the transient response of the SLSCC, further verifying functionality of the feedback control loop. Fig. 13 plots output voltage fluctuations of all four layers due to load current transients as NMOS transistors across the layers turn ON and OFF. In this plot, each layer current is labeled with respect to the nominal current that was measured at 900 mV. With the reference voltage set to 800 mV, the SLSCC ensures a minimum voltage of about 800 mV for all layers. From 0 to 5 μ s, all load generators are initially

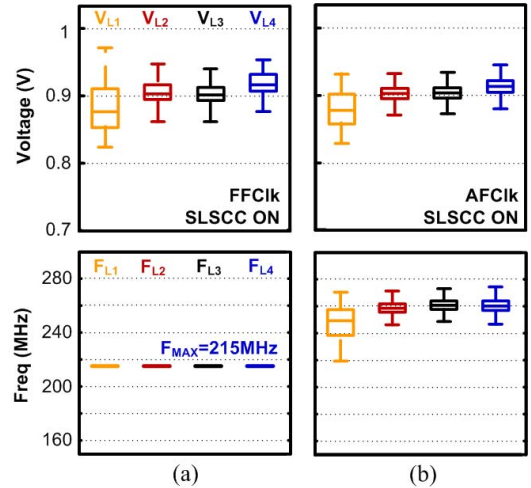


Fig. 14. Measured voltage and frequency distribution of an unbalanced workload scenario with higher load in Layer1, for two operation modes: (a) cores running on fixed-frequency clock with SLSCC ON and (b) cores running on adaptive frequency clock with SLSCC ON

OFF and all layer voltages settle roughly settle to 900 mV, one quarter of the 3.6 V input voltage, as expected. Small layer-to-layer voltage differences can be attributed to leakage current mismatch. Then, at $t = 5$ μ s, load current in the bottom layer increases to 25 mA. Voltage stacking redistributes the layer voltages with the Layer 1 voltage drooping below 900 mV, and the SLSCC maintains a minimal voltage of V_{LAYER1} around 800mV. As load currents change over time, the SLSCC works to always ensure a minimum voltage of ~ 800 mV for all layers.

Since the SLSCC only guarantees the lower bound of the layer voltages, however, Fig. 13 also shows that voltage droop in one layer leads to elevated voltage levels in other stack layers. For global fixed-frequency operation, wherein the

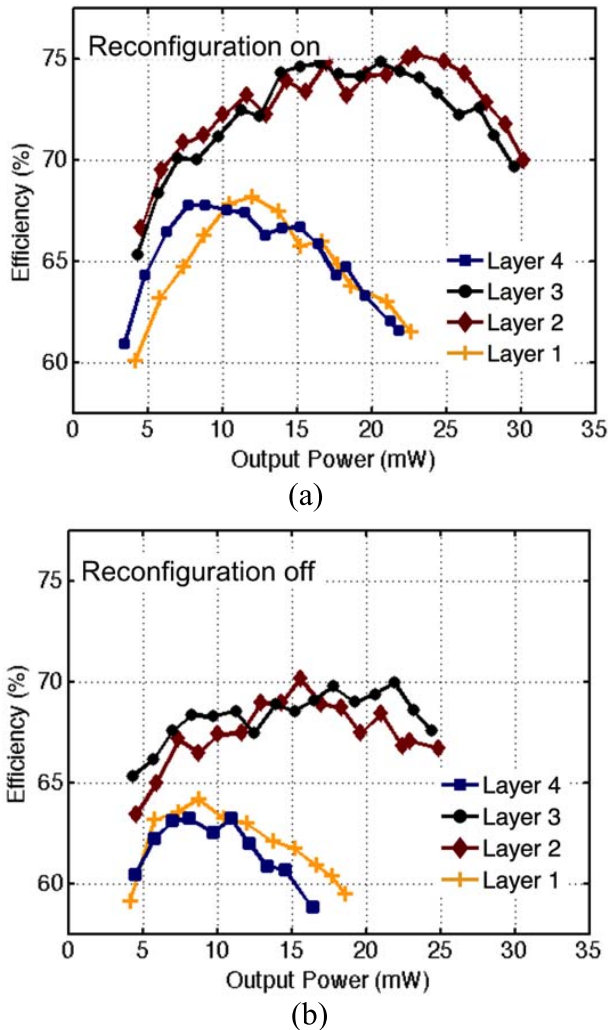


Fig. 15. Measured efficiencies of the SLSCC when only one layer consumes current, with reconfiguration (a) ON and (b) OFF ($V_{IN} = 3.6V$, all reference voltages are set to 800 mV).

clock frequency must be set with respect to the worst-case voltage droop observed in all stack layers, the cores in layers with elevated voltages end up operating with higher voltage margin for the same performance, resulting in inefficient energy utilization. One way to make better use of the higher margins is to allow cores to operate in per-layer adaptive frequency clocking (AFClk) mode. Fig. 14 presents box plots of the output layer voltage distribution as well as the core operating frequency distribution for each layer for a voltage noise scenario using active cores as loads. The cores execute an unbalanced workload scenario, with all cores in Layers 2 through 4 running a lower-power, control-intensive, string-matching kernel, while the four cores in Layer 1 execute a higher-power, compute-intensive, molecular dynamics kernel. The distribution was collected over a 1ms execution window. This scenario leads to voltage droop in Layer 1, but the SLSCC (with reference voltage set to 850 mV) guarantees a minimum layer voltage close to the reference voltage. Fig. 14(a) shows that with fixed-frequency clocking, worst-case voltage droop limits the clock frequency to 215 MHz for all cores in all layers. This means cores in Layers 2 through 4 operate with excess voltage margins. In contrast, Fig. 14(b) shows

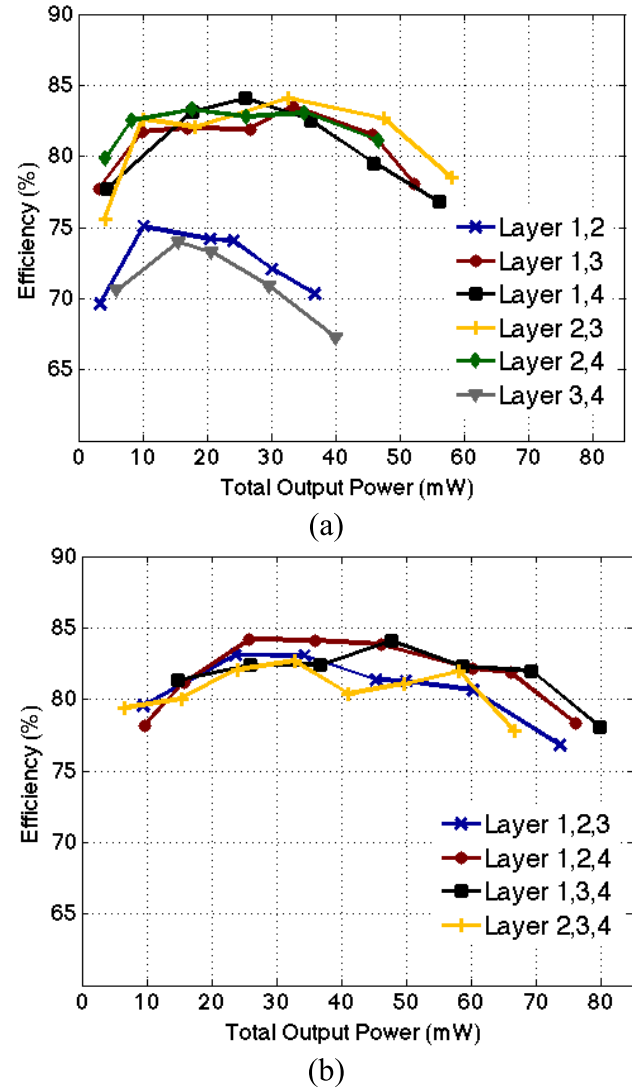


Fig. 16. Measured efficiencies of the SLSCC when multiple layers consume currents ($V_{IN} = 3.6V$, all reference voltages are set to 800 mV).




Fig. 17. A histogram of measured power delivery efficiency of the overall voltage stacking system ($V_{IN} = 3.6V$, all reference voltages are set to 800 mV).

how adaptive frequency clocking improves average per-layer clock frequency, since clock frequency tracks per-layer voltage levels.

	[4] VLSI 10	[8] ISSCC 14	[9] ISSCC 13	[10] ISSCC 13	[15] JSSC 14	This Work
Technology	45nm	0.25 μ m	180nm	65nm	22nm	40nm
Capacitor technology	Trench	MIM	On-Chip	MOS	MIM	MOS
Total capacitance	—	3nF	2.24nF	3.88nF	—	4.5nF
V_{IN}	2V	2.5V	3.4V-4.3V	3V-4V	1.23V	3.6V
V_{OUT}	0.95V-1.05V	0.1V-2.18V	0.9V-1.5V	1V	0.45V-1V	0.8V-1V
Voltage stacking	2-Way	None	None	None	None	4-Way
Quoted efficiency (η)	90%	60%	72%	73%	70%	65% *
Conv. ratio @ η	2:1	4:1	4:1	3:1	2:1	4:1
P_{OUT} @ η	—	1mW	0.27mW	122mW	6.4mW	17.5mW *
Power density (mW/mm ²)@ η	2185	0.215	0.16	190	—	21.1 *

* For consistency, the SLSCC only delivers power to the bottom layer (Layer 1).

Fig. 18. Performance summary and comparison to prior work.

B. Conversion Efficiency

As discussed previously, the SLSCC only processes differential power consumed by the load. In the worst case, the load circuitry in only one of the four stacked output layers consumes current. In such cases, all the power that is consumed must be delivered by the SLSCC. Fig. 15 plots the measured conversion efficiency of the converter when only one layer consumes current. These results take into account the power consumption of the control logic with hybrid control turned ON. In Fig. 15(a), the proposed SLSCC (with dynamic flying capacitor allocation) achieves higher efficiency and supports higher mismatched power (output power) for the two middle layers. Analysis of the internal SLSCC charge flow in Fig. 2 confirms that the losses are smaller (i.e., lower conductive loss) when delivering current to the middle layers. In the measurements presented in Fig. 15(b), the flying capacitor reconfiguration is OFF and the flying capacitance resource is equally distributed. Comparison of Fig. 15(a) with Fig. 15(b) confirms the benefits of reconfiguring capacitor allocations. Reconfiguration improves conversion efficiency as well as the range of power delivery. Given its clear benefits, flying capacitor reconfiguration is always ON for all remaining measurement results unless stated otherwise.

As shown in Fig. 15, conversion efficiency of the SLSCC is not very high when delivering power to only one output layer. However, conversion efficiency improves significantly when more than one layer consumes current, which is the more common case for the intended voltage stacking application. Fig. 16 presents the measured efficiencies of the SLSCC when more than one layer consumes currents. In Fig. 16(a), two layers consume current. The load generators are set in a way that the same size load-creating transistors are turned ON in the two layers. Generally, the efficiencies are higher compared with those in Fig. 15. The output power delivery range also increases over the results in Fig. 15. These results are consistent with the charge-flow analysis in Section II. Fig. 16(a) also shows that efficiency differs depending on which two layers consume current. This is because charge flow depends on layer-to-layer load conditions. Fig. 16(b) shows the efficiencies when three output layers consume currents. Both efficiency and power delivery range show




Fig. 19. Die micrograph of the SLSCC

much improvement over the scenarios where only one layer consumes current.

Fig. 17 presents the power delivery efficiency of the overall voltage stacking system for a diverse collection of load current conditions. Each output layer consumes a random amount of current. Efficiency is computed as the total power consumed by all loads versus the total power supplied by V_{IN} at 3.6 V. The average efficiency is as high as 87%, confirming the benefits of the voltage stacking system. SLSCC losses only apply to inter-layer power mismatches. All of the converter's reference voltages are 800mV, and the SLSCC might not need to switch at all unless one or more of the layer voltages fall below 800 mV.

For consistency, Fig. 18 compares this work to prior work assuming power delivery to a single layer, but note that both conversion efficiency and power density improve when power is delivered to multiple layers, as required by voltage stacking. The on-chip area occupied by the converter is used when computing the power density. The area occupied by off-chip decoupling capacitors is not included when computing the power density.

A chip micrograph of the 0.829 mm² SLSCC is shown in Fig. 19.

VI. CONCLUSION


This paper demonstrates a fully integrated switched capacitor converter for a 4-way voltage stacking application. The symmetric ladder topology, with reconfigurable flying capacitance, supports 4 stacked output layers simultaneously with improved conversion efficiency and power delivery range for the high 4-to-1 conversion ratio. The proposed hybrid feedback control scheme reduces the static peak-to-peak voltage ripple and at the same time provides fast transient response to handle large current steps.

ACKNOWLEDGMENTS


The authors would like to thank the TSMC's university shuttle program for chip fabrication and Intel Corporation for Siskiyou Peak IP.

REFERENCES

- [1] S. Rajapandian, K. L. Shepard, P. Hazucha, and T. Karnik, "High-voltage power delivery through charge recycling," *IEEE J. Solid-State Circuits*, vol. 41, no. 6, pp. 1400–1410, Jun. 2006.
- [2] S. Rajapandian, Z. Xu, and K. L. Shepard, "Implicit DC–DC downconversion through charge-recycling," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 846–852, Apr. 2005.
- [3] S. K. Lee, D. Brooks, and G.-Y. Wei, "Evaluation of voltage stacking for near-threshold multicore computing," in *Proc. ACM/IEEE ISLPED*, Jul. 2012, pp. 373–378.
- [4] L. Chang *et al.*, "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.
- [5] K. Mazumdar and M. Stan, "Breaking the power delivery wall using voltage stacking," in *Proc. Great Lakes Symp. VLSI*, May 2012, pp. 51–54.
- [6] K. Kesarwani, C. Schaef, C. R. Sullivan, and J. T. Stauth, "A multi-level ladder converter supporting vertically-stacked digital voltage domains," in *Proc. 28th Annu. IEEE Appl. Power Electron. Conf. Expo. (APEC)*, Mar. 2013, pp. 429–434.
- [7] M. Seeman, "A design methodology for switched-capacitor DC-DC converters," Dept. Elect. Eng. Comput. Sci., Univ. California, Berkeley, Berkeley, CA, USA, Tech. Rep. UCB/EECS-2009-78, 2009.
- [8] L. G. Salem and P. P. Mercier, "An 85%-efficiency fully integrated 15-ratio recursive switched-capacitor DC-DC converter with 0.1-to-2.2 V output voltage range," in *ISSCC Dig. Tech. Papers*, Feb. 2014, pp. 88–89.
- [9] S. Bang, A. Wang, B. Giridhar, D. Blaauw, and D. Sylvester, "A fully integrated successive-approximation switched-capacitor DC-DC converter with 31 mV output voltage resolution," in *ISSCC Dig. Tech. Papers*, Feb. 2013, pp. 370–371.
- [10] H.-P. Le, J. Crossley, S. R. Sanders, and E. Alon, "A sub-ns response fully integrated battery-connected switched-capacitor voltage regulator delivering 0.19 W/mm² at 73% efficiency," in *ISSCC Dig. Tech. Papers*, Feb. 2013, pp. 372–373.
- [11] T. M. Andersen *et al.*, "A sub-ns response on-chip switched-capacitor DC-DC voltage regulator delivering 3.7 W/mm² at 90% efficiency using deep-trench capacitors in 32 nm SOI CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2014, pp. 90–91.
- [12] S. R. Sanders, "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.
- [13] T. Tong, X. Zhang, W. Kim, D. Brooks, and G.-Y. Wei, "A fully integrated battery-connected switched-capacitor 4:1 voltage regulator with 70% peak efficiency using bottom-plate charge recycling," in *Proc. IEEE CICC*, Sep. 2013, pp. 1–4.
- [14] T. M. Andersen *et al.*, "A 4.6 W/mm² power density 86% efficiency on-chip switched capacitor DC-DC converter in 32 nm SOI CMOS," in *Proc. IEEE Appl. Power Electron. Conf.*, Mar. 2013, pp. 692–699.
- [15] R. Jain *et al.*, "A 0.45–1 V fully-integrated distributed switched capacitor DC-DC converter with high density MIM capacitor in 22 nm tri-gate CMOS," *IEEE J. Solid-State Circuits*, vol. 49, no. 4, pp. 917–927, Apr. 2014.
- [16] S. S. Kudva and R. Harjani, "Fully integrated capacitive DC–DC converter with all-digital ripple mitigation technique," *IEEE J. Solid-State Circuits*, vol. 48, no. 8, pp. 1910–1920, Aug. 2013.
- [17] T. M. Van Breussegem and M. S. J. Steyaert, "Monolithic capacitive DC-DC converter with single boundary–multiphase control and voltage domain stacking in 90 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 46, no. 7, pp. 1715–1727, Jul. 2011.
- [18] 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.
- [19] S. K. Lee, T. Tong, X. Zhang, D. Brooks, and G.-Y. Wei, "A 16-core voltage-stacked system with an integrated switched-capacitor DC-DC converter," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2015, pp. C318–C319.
- [20] R. Jain *et al.*, "Conductance modulation techniques in switched-capacitor DC-DC converter for maximum-efficiency tracking and ripple mitigation in 22 nm tri-gate CMOS," *IEEE J. Solid-State Circuits*, vol. 50, no. 8, pp. 1809–1819, Aug. 2015.
- [21] Y. K. Ramadass, A. A. Fayed, and A. P. Chandrakasan, "A fully-integrated switched-capacitor step-down DC-DC converter with digital capacitance modulation in 45 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 45, no. 12, pp. 2557–2565, Dec. 2010.



systems.



Tao Tong received the B.E. degree from Tsinghua University, Beijing, China, the M.S. degree from Oregon State University, Corvallis, OR, USA, and the Ph.D. degree from Harvard University, Cambridge, MA, USA.

He worked at MediaTek Wireless Inc. and Lion Semiconductor Inc., designing analog-to-digital converters and fully integrated DC-DC converters for mobile applications. His research interests include integrated voltage regulators and their applications in energy-efficient computing

Sae Kyu Lee received the B.S. degree in electrical engineering from Seoul National University, Seoul, Korea, in 2006, and the M.S. degree in electrical and computer engineering from The University of Texas at Austin, Austin, TX, USA, in 2008. He is currently working toward the Ph.D. degree at Harvard University, Cambridge, MA, USA.

He previously worked at Intel Corporation and AMD, working on mobile microprocessor designs. His research focuses on VLSI design for efficient on-chip power delivery solutions.



Xuan Zhang received the B.Eng. degree in electrical engineering from Tsinghua University, Beijing, China, in 2006, and the Ph.D. degree in electrical and computer engineering from Cornell University, Ithaca, NY, USA, in 2011.

She is an Assistant Professor with the Preston M. Green Department of Electrical and Systems Engineering, Washington University, St. Louis, MO USA. Prior to joining Washington University, she was a Postdoctoral Fellow with the Harvard School of Engineering and Applied Sciences, working on the design of the "brain" for an insect-scale micro-robot. She works across the fields of robotics, system engineering, VLSI, and computer architecture. Her research focus has been on miniaturization and optimization of autonomous systems for performance, reliability, security, and energy efficiency, with diverse applications in micro-robotics, Internet-of-Things (IoT), wearable/implantable devices, ubiquitous computing, and resilient cyber-physical systems.

Dr. Zhang was the recipient of an Intel Fellowship in 2008–2009 and won the Design Contest Award at IEEE ISLPED in 2013.



David Brooks (F'16) received the B.S. degree from the University of Southern California, Los Angeles, CA, USA, and the M.A. and Ph.D. degrees from Princeton University, Princeton, NJ, USA, all in electrical engineering.

He is the Haley Family Professor of Computer Science with the School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA. Prior to joining Harvard, he was a Research Staff Member with the IBM T.J. Watson Research Center. His research interests include resilient and power-efficient computer hardware and software design for high-performance and embedded systems.

Prof. Brooks was the recipient of several honors and awards including the ACM Maurice Wilkes Award, ISCA Influential Paper Award, the National Science Foundation CAREER Award, IBM Faculty Partnership Award, and DARPA Young Faculty Award.



Gu-Yeon Wei received the B.S., M.S., and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, USA, in 1994, 1997, and 2001, respectively.

He is the Gordon McKay Professor of Electrical Engineering and Computer Science with the John A. Paulson School of Engineering and Applied Sciences (SEAS), Harvard University, Cambridge, MA, USA. His research interests span multiple layers of a computing system: mixed-signal integrated circuits, computer architecture, and design tools for efficient hardware. His research efforts focus on identifying synergistic opportunities across these layers to develop energy-efficient solutions for a broad range of systems from flapping-wing microrobots to large-scale servers.