

A RISC-V Vector Processor With Simultaneous-Switching Switched-Capacitor DC–DC Converters in 28 nm FDSOI

Brian Zimmer, *Member, IEEE*, Yunsup Lee, *Student Member, IEEE*, Alberto Puggelli, *Student Member, IEEE*, Jaehwa Kwak, *Student Member, IEEE*, Ruzica Jevtić, *Member, IEEE*, Ben Keller, *Student Member, IEEE*, Steven Bailey, *Student Member, IEEE*, Milovan Blagojević, *Member, IEEE*, Pi-Feng Chiu, *Student Member, IEEE*, Hanh-Phuc Le, *Member, IEEE*, Po-Hung Chen, *Member, IEEE*, Nicholas Sutardja, *Student Member, IEEE*, Rimas Avizienis, Andrew Waterman, Brian Richards, *Member, IEEE*, Philippe Flatresse, Elad Alon, *Senior Member, IEEE*, Krste Asanović, *Fellow, IEEE*, and Borivoje Nikolić, *Senior Member, IEEE*

Abstract—This work demonstrates a RISC-V vector microprocessor implemented in 28 nm FDSOI with fully integrated simultaneous-switching switched-capacitor DC–DC (SC DC–DC) converters and adaptive clocking that generates four on-chip voltages between 0.45 and 1 V using only 1.0 V core and 1.8 V IO voltage inputs. The converters achieve high efficiency at the system level by switching simultaneously to avoid charge-sharing losses and by using an adaptive clock to maximize performance for the resulting voltage ripple. Details about the implementation of the DC–DC switches, DC–DC controller, and adaptive clock are provided, and the sources of conversion loss are analyzed based on measured results. This system pushes the capabilities of dynamic voltage scaling by enabling fast transitions (20 ns), simple packaging (no off-chip passives), low area overhead (16%), high conversion efficiency (80%–86%), and high energy efficiency (26.2 DP GFLOPS/W) for mobile devices.

Index Terms—Adaptive clock, DC–DC conversion, dynamic voltage and frequency scaling (DVFS), fully integrated converter, integrated voltage regulator, noninterleaved, RISC-V, simultaneous-switching, switched-capacitor.

I. INTRODUCTION

DYNAMIC voltage and frequency scaling (DVFS) is a popular technique to improve energy efficiency in digital systems [1]. As performance requirements change over time, the voltage can be changed appropriately to maximize energy efficiency while meeting performance constraints.

DVFS is commonly implemented using off-chip voltage regulators, but off-chip regulation has a number of disadvantages. Increased parasitics and an on-chip to off-chip feedback loop cause slow mode transitions. Also, each voltage domain must still be supplied through the package separately—limiting the total number of voltages available and increasing packaging costs and complexity. Lastly, off-chip regulators and supporting components, such as inductors, increase total system size and cost.

Integrating regulators on-chip, and tightly connecting power supply control with the microprocessor, offers significant advantages by reducing system cost and supporting much finer grain DVFS in terms of both operating mode period and voltage domain area. Transition times between modes can be reduced, providing additional energy savings through more frequent DVFS to better track changing performance requirements [2]. Supporting many smaller voltage domains provides better isolation between high and low performance regions, and supplying hundreds of independent voltage domains is desirable to improve the energy efficiency of many-core systems [3] [4]. Instead of requiring separate power grids for each domain or needing to support full power delivery requirements through each of a few shared voltage rails, on-chip switched-capacitor DC–DC (SC DC–DC) only requires the delivery of two supplies through the package, which simplifies package design and makes it less expensive. Finally, no off-chip components are necessary, providing significant platform size and cost reductions. Despite these numerous advantages, adoption has been

Manuscript received September 01, 2015; revised December 13, 2015; accepted January 04, 2016. Date of publication March 01, 2016; date of current version March 29, 2016. This paper was approved by Guest Editor Masato Motomura. This work was supported in part by BWRC, in part by ASPIRE, in part by DARPA PERFECT Award Number HR0011-12-2-0016, in part by Intel ARO, in part by AMD, in part by SRC/TxACE, in part by Marie Curie FP7, in part by NSF GRFP, in part by NVIDIA Fellowship, and fabrication donation by STMicroelectronics.

B. Zimmer, Y. Lee, A. Puggelli, J. Kwak, B. Keller, S. Bailey, P.-F. Chiu, N. Sutardja, R. Avizienis, A. Waterman, B. Richards, E. Alon, K. Asanović, and B. Nikolić are with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA (e-mail: bmzimmer@eecs.berkeley.edu).

R. Jevtić is with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA, and also with the Universidad Antonio de Nebrija, Madrid 28040, Spain.

M. Blagojević is with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA, and also with the Institut Supérieur d'Électronique de Paris and STMicroelectronics, Crolles 38920, France.

H.-P. Le is with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA, and also with the University of Colorado, Boulder, CO 80309 USA.

P.-H. Chen is with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA, and also with National Chiao Tung University, Hsinchu 30013, Taiwan.

P. Flatresse is with STMicroelectronics, Crolles 38920, France.

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.2016.2519386

TABLE I
COMPARISON TO PRIOR WORK OF DIGITAL SYSTEMS WITH INTEGRATED VOLTAGE CONVERSION

Metric	[5]	[6]	[35]	[16]	[17]	[18]	This work
Technology	22 nm SOI	22 nm FinFET	65 nm bulk	28 nm FD-SOI	22 nm FinFET	28 nm FD-SOI	28 nm FD-SOI
Digital system	Processor	Processor	Microcontroller	Processor	Graphics core	Processor	Processor
Technique	LDO	Buck	Interleaved SC	Interleaved SC	Hybrid LDO / interleaved SC	SC	Simultaneous SC
Transition time (μs)	~ 300	0.32	–	–	–	–	0.02
Off-chip components	No	Yes	No	No	No	No	No
Input voltage (V)	1.1	1.7	1.2	1.1	1.05	1.55	1, 1.8
Output voltage (V)	0.75–1	–	0.3–0.6	0.33–0.45	0.38–0.92	0.3–0.55	0.5, 0.67, 0.9
Maximum power density ($\frac{\text{W}}{\text{mm}^2}$)	36	~ 32	0.004	0.04	–	0.009	0.35
Maximum load power (mW)	–	–	0.5	5	–	0.2	65
Quoted efficiency (%)	60–90	90	80	65–82	52–84	71–82	80–86

limited because low converter efficiency has negated the many benefits of on-chip regulation.

Previous proposals for on-chip conversion include integrated low-drop-out (LDO) regulators, buck converters with off-chip inductors, and SC DC-DC converters. Wide voltage operation requires a regulator with a high efficiency across the full range of output voltages, and LDO regulators suffer from sub-50% efficiency at low operating voltages [5]. Buck converters with on-chip switches and off-chip inductors offer high efficiency but still require inductors to be integrated into the package or PCB [6]–[8]. Because the quality of integrated inductors is inherently worse than integrated capacitors [9], buck converters with on-chip inductors report lower efficiency [10]. By replacing inductors with capacitors, SC DC-DC converters can be fully integrated on-chip, but achieving high efficiency compared to designs with off-chip passives is challenging. Traditional interleaved SC DC-DC converters stabilize the output voltage to minimize frequency margining for supply variation [11]. Standalone converters have demonstrated high efficiency of 80%–90% [12]–[15]. However, full system implementations that use converters to drive real digital loads have reported limited efficiency of 52%–84% [16]–[18]. Table I provides a summary of prior work.

This paper implements a different switched-capacitor control approach, *simultaneous-switching*, to achieve high efficiency by switching all possible capacitance simultaneously and using an adaptive clock to maximize clock frequency for the resulting voltage ripple. The on-chip SC DC-DC converter powers a RISC-V [19] scalar microprocessor with vector accelerator, enabling improved DVFS with fast transitions between modes (20 ns), low area overhead (16%), simple package requirements (two voltages with no off-chip components), scalability to numerous voltage domains, and high efficiency. Section II describes the reasons for the improved efficiency of simultaneous-switching over interleaved converters. Section III provides details about the design and implementation. Section IV analyzes measurement results from the chip, and discusses different sources of efficiency loss.




Fig. 1. Theoretical efficiency improvement of the simultaneous-switching SC DC-DC converter. The proposed converter achieves higher efficiency than a conventional interleaved design by avoiding charge sharing.

II. SIMULTANEOUS-SWITCHING VERSUS INTERLEAVED SC DC-DC CONVERTERS

Maximizing conversion efficiency of DC-DC converters is essential for on-chip regulation, because low efficiency may cancel the energy efficiency gains of DVFS. Losses in SC DC-DC converters can be categorized into four separate components [15]: charge-sharing SC loss (P_{cfly}), conduction loss (P_{cond}), switching loss (P_{gate}), and bottom-plate loss (P_{bottom}). The contribution of each loss term to total losses for an interleaved and simultaneous-switching SC converter is shown in Fig. 1. After design-time optimization of switch and capacitor size, the only parameter that changes efficiency is the switching frequency (f_{sw}) and the associated ripple size (ΔV). P_{cfly} is inversely proportional to switching frequency, while P_{gate} and P_{bottom} are proportional—therefore, optimizing efficiency requires setting f_{sw} such that the sum of losses is minimized.

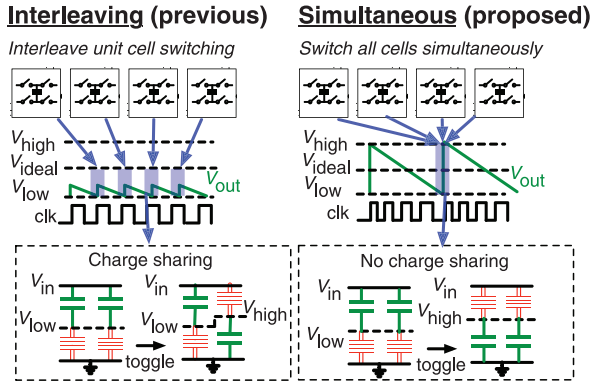


Fig. 2. Proposed simultaneous-switching SC DC–DC converter switches every DC–DC unit cell simultaneously to avoid charge sharing.

The efficiency differences between interleaved and simultaneous-switching SC DC–DC converters arise from the charge-sharing loss term. Fig. 2 compares the overall simultaneous-switching and interleaved approaches. Interleaved converters switch one unit cell at a time to stabilize the output voltage and remove losses due to unnecessarily high voltages for a fixed clock frequency, but unit cells share charge with each other and P_{cflly} remains a significant loss component. Simultaneous-switching operation improves converter efficiency by switching all unit cells simultaneously to avoid charge sharing losses, while an adaptive clock translates the rippling supply voltage into additional performance to eliminate system-level efficiency losses caused by the voltage ripple on the core supply [20]. For simultaneous-switching converters driving an ideal resistive load, perfect frequency adaptation would completely remove all charge-sharing loss ($P_{\text{cflly}} = 0$). By removing the only loss component that is proportional to ripple size, the switching frequency can be decreased to further reduce the other loss terms. In a real implementation, however, nonidealities cause a nonzero P_{cflly} , and Section IV-D analyzes this loss further by using measured results.

III. INTEGRATED SYSTEM IMPLEMENTATION

Fig. 3 shows the chip architecture. The 64 bit scalar core implements the free and open RISC-V instruction set [19]. A high-performance 64 bit vector accelerator improves energy efficiency by amortizing instruction fetch and control overhead for data-parallel operations. The processor boots Linux and executes compiled scalar and vector code. Two voltages, a 1.0 V core and 1.8 V I/O supply, are delivered to the on-chip converters. The SC DC–DC converter is partitioned into twenty-four $90 \mu\text{m} \times 90 \mu\text{m}$ unit cells surrounding the core (16% area overhead) and generates four dynamically reconfigurable average ideal output voltages of 1.0, 0.9, 0.67, and 0.5 V. These fixed ratios were chosen in order to utilize common core and I/O voltages as inputs, and for their low output impedance coefficients [21]. Continuous voltage selection for DVFS is achieved by hopping between discrete SC DC–DC modes [20], [22], and these specific voltages were chosen as a tradeoff between DVFS tuning granularity and implementation complexity. A shared




Fig. 3. System block diagram showing the scalar and vector microprocessor powered by on-chip voltage converters.




Fig. 4. Pipeline diagram of the Rocket scalar core.

SC DC–DC controller switches all of the unit cells simultaneously. An adaptive clock generator adjusts the clock period each cycle based on the instantaneous converter output voltage, and a high-speed receiver is used to provide a 2 GHz reference clock for the clock generator’s DLL. Level shifters and asynchronous FIFOs separate the core and uncore voltage domains. Large random variations in SRAM memory cells typically limit voltage scaling, so custom SRAMs were implemented to enable voltage scaling down to 0.45 V. Each 4 KB SRAM uses 8 T cells and has 512 words of 72 bits with 2:1 interleaving.

A. Scalar Core

The Rocket scalar core, shown in Fig. 4, is a 64 bit 5-stage single-issue in-order pipeline that executes the RISC-V instruction set architecture (ISA). It is carefully designed to minimize the impact of long clock-to-output delays of SRAM macros. For example, the pipeline resolves branches in the memory stage to shorten the critical path through the bypass path, but relies on extensive branch prediction (64 entry branch target buffer, 256 entry two-level branch history table, and a 2 entry return address stack) to mitigate the increased branch resolution penalty. The blocking 16 KB instruction cache is private to the scalar core, while the nonblocking 32 KB data cache is shared between the scalar core and the vector accelerator. The scalar core has a memory-management unit that supports page-based virtual memory. Both caches are virtually indexed and physically tagged, and have separate TLBs that are accessed in parallel with cache accesses. The core has an IEEE 754-2008 compliant floating-point unit that executes single- and double-precision floating-point operations, including fused multiply-add (FMA) operations, with hardware support for

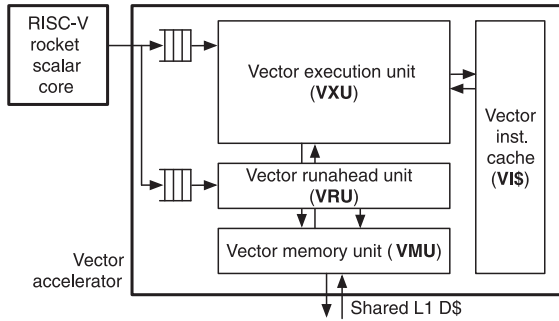


Fig. 5. Block diagram of the Hwacha vector accelerator.

subnormal numbers. The resulting Rocket scalar core is competitive to industrial designs in terms of performance, power consumption, and area [23].

To reduce design complexity, the microprocessor is implemented as a tethered system. Unlike a standalone system, a tethered system depends on a host machine to boot, and lacks any I/O devices such as a console, mass storage, frame buffer, and network card. The host (e.g., an x86 laptop) is connected to the target tethered system via the host-target interface (HTIF), a simple protocol that lets the host machine read and write target memory and control registers. All I/O-related system calls are forward to the host machine using HTIF, where they are executed on behalf of the target. Programs that run on the scalar core are downloaded into the target's memory via HTIF. The resulting system is able to boot modern operating systems, such as Linux, utilizing I/O devices residing on the host machine, and can run standard complex applications such as the Python interpreter.

B. Vector Accelerator

The Hwacha vector accelerator, shown in Fig. 5, is a decoupled single-lane vector unit optimized for ASIC designs. Hwacha executes vector operations temporally (split across subsequent cycles) rather than spatially (split across parallel datapaths), and has a vector length register that simplifies vector code generation and keeps the binary code compatible across different vector microarchitectures with different numbers of execution resources [24].

The Rocket scalar core sends vector memory instructions and vector fetch instructions to the vector accelerator. A vector fetch instruction initiates execution of a block of vector arithmetic instructions. The vector execution unit (VXU) fetches instructions from the private vector instruction cache (VIS), decodes instructions, clears hazards, and then sequences vector instruction execution by sending multiple μ ops down the vector lane. The vector lane consists of a banked vector register file built out of 1R1W SRAM macros, operand registers, per-bank integer ALUs, and long-latency functional units. Multiple operands per cycle are read from the banked register file by exploiting the regular access pattern with operand registers used as temporary space [23]. The long-latency functional units such as the 64 bit integer multiplier, single- and double-precision FMA units are shared between the scalar core and the vector accelerator. The vector memory unit (VMU) supports unit-strided,




Fig. 6. Four switching topologies of the reconfigurable SC DC-DC design.

constant-strided, and gather/scatter vector memory operations to the shared L1 data cache. Vector memory instructions are also sent to the vector runahead unit (VRU) by the scalar core. The VRU prefetches data blocks from memory and places them in the L1 data cache ahead of time to increase performance of vector memory operations executed by the VXU [24], [25].

The resulting vector accelerator is more similar to traditional Cray-style vector pipelines [26] than SIMD units such as those that execute ARM's NEON or Intel's SSE/AVX instruction sets, and delivers high performance and energy efficiency while remaining area efficient.

C. SC DC-DC Unit Cell

This system uses a reconfigurable DC-DC converter unit with a topology similar to [15], where separate networks of switches allow different conversion ratios for the same shared flying capacitor. Due to the availability of two different input voltages in the IO pads, two sets of switches are used: one for the configurations operating off a 1 V input and the other one for configurations operating off a 1.8 V input. Four possible discrete SC DC-DC configurations, shown in Fig. 6, generate voltages between 0.5 and 1 V to enable a wide operating range. The converter has two phases: in the first phase ϕ_1 , the flying capacitor is connected in series with the output, while in the second phase ϕ_2 , the flying capacitor is connected in parallel. The 1 V input is divided with a 2:1 and 3:2 ratio to generate the 0.5 and 0.67 V modes, while the 1.8 input is divided with a 2:1 ratio to generate the 0.9 V mode. All 1 V input switches are implemented as LVT devices to reduce their ON resistance, while the larger 1.8 V input switches are implemented as RVT devices to reduce their leakage. Additionally, the largest switches are forward-body-biased to reduce their ON resistance when they are active (i.e., in 1.8 V 1/2 mode). The flying capacitor is implemented as MOS capacitance with two layers of MOM capacitance above. Parasitic bottom-plate capacitance is reduced by using a series connection of the box, well, and substrate capacitances [27]. SC DC-DC converters are best suited for low-power-density applications where the limited capacitive density of on-chip




Fig. 7. SC DC–DC controller. The circuit simultaneously toggles the phase of all unit cells when converter output voltage falls below specified reference voltage.

capacitors is sufficient and the area overhead of converters is reasonable. While this implementation uses MOS capacitors to reduce cost, area overhead can be further reduced with MIM capacitors. Twenty-four unit cells were used in the design for a total flying capacitance of 2.1 nF. For testing and measurement purposes, the *bypass mode* of the converter uses the 1 V mode to connect the regulator’s 1 V input rail to V_{out} of the microprocessor through power gates in the SC DC–DC unit cells, and the 1 V input rail is supplied by the desired bypass voltage to directly control the voltage of the microprocessor.

D. SC DC–DC Controller

The purpose of the SC DC–DC controller block is to trigger the switching of the converter unit cells in order to guarantee that the converter can provide the required current to the processor at all times. Analytically, the converter output current I_{out} needs to equal the load current I_L , which is assumed constant over one switching cycle T_{sw}

$$I_L = I_{out} = \alpha \times C_{fly} \times \Delta V \times f_{sw}. \quad (1)$$

The topology proportionality constant (α) and the total amount of flying capacitance in the converter C_{fly} are set at design time. During runtime, the SC DC–DC controller needs to maximize efficiency by appropriately controlling the amplitude of the voltage ripple (ΔV) and the converter switching frequency (f_{sw}).

This design implements a lower-bound (hysteretic) controller, shown in Fig. 7, that switches the cells when the output voltage V_{out} drops below a reference voltage V_{ref} —explicitly setting ΔV and implicitly modulating f_{sw} in response to changing load current [28]. Lower-bound control was chosen for quick reaction to changes in the load current I_L and to avoid switching the converter unnecessarily quickly.

The controller is composed of two main components: clocked comparators to detect when V_{out} falls below V_{ref} , and a finite-state machine (FSM) that generates the toggle signal for the unit cells. To guarantee that the toggle signal arrives simultaneously at all cells, the SC DC–DC controller is centralized, and the toggle signal is routed as a clock tree to minimize skew among cells.

Three separate StrongARM [29] comparators are used: the 1 V 2:1 mode uses the PMOS-based-comparator (for the lowest




Fig. 8. State transition of the FSM at the output of the comparator which ensures correct operation in the case where V_{out} remains below V_{ref} .

common mode input voltages), while the other modes use two NMOS-input-based comparators, with one operating on the rising edge of the clock and the other on the falling edge of the clock (for higher common mode input voltages). A multiplexer changes V_{ref} for different conversion ratios. In a lower-bound controller, the shortest achievable time between two switching events ($t_{sw,min}$) is set by the propagation time of the toggling signal from the comparator output to the final power switches. The comparator clock frequency is set to 2 GHz to maximize power density by allowing all unit cells to toggle every $t_{sw,min}$ during high current loads, and to minimize the time that V_{out} remains below V_{ref} before detection triggers a toggle event.

A FSM, shown in Fig. 8, sends the toggle signal to the unit cells based on the comparator output. The rising edge of the comparator output signal `comparator_out` toggles transitions between the two converter phases. If `comparator_out` remains high for multiple cycles (because a large current spike keeps V_{out} below V_{ref} even after a switching event), a counter increments and forces a toggle when it reaches an overflow value. The overflow count is set to be slightly longer than the propagation time from the comparator through the toggle signal clock tree and to the switches, to avoid spurious switching events. The reset state is used during reset and during converter mode transitions.

E. Adaptive Clock

The adaptive clocking scheme, shown in Fig. 9, changes the clock frequency on a cycle-by-cycle basis to ensure that the system operates at the maximum instantaneous frequency obtainable for the instantaneous voltage [30]. The rippling supply voltage from the SC DC–DC converters powers a tunable replica circuit (TRC), adjustable from 4 to 124 FO1 inverter delays with a delay setting register, to mimic the critical path




Fig. 9. Adaptive clock system with a tunable replica path. The system instantaneously changes the clock frequency to track the critical path for constantly changing output voltage.




Fig. 10. Measurement results of the replica timing path. The use of different tuning codes for each DC-DC mode allow the replica path to closely track the critical path.

delay at each instantaneous voltage level. When the TRC generates a pulse, the controller selects one of the 16 DLL phases to send to the core. Separate TRC paths control the high and low clock periods to set the duty cycle. This is a free-running clock, in which nothing determines the average frequency other than the average delay through the TRC.

During operation, the first TRC output pulse asynchronously resets the clock toggler flip-flop to generate the falling edge of the clock output. The second TRC output pulse synchronizes the rising edge of the adaptive clock with the DLL references. Level shifters are located between the TRC and the controller. Since the DLL references and the TRC output pulse are fully asynchronous, a watchdog block monitors the system for metastability. Fig. 10 shows the ability of the adaptive clock to track changes in voltage by using the bypass mode to measure average frequency for different delay settings for the TRC. Annotations above the plot indicate the approximate voltage ranges seen in each SC DC-DC mode. Because the inverter-based replica path delay characteristics do not match the critical paths of the processor, a single delay setting poorly tracks the processor critical path over the entire voltage range. However, manual calibration of specific delay settings for each

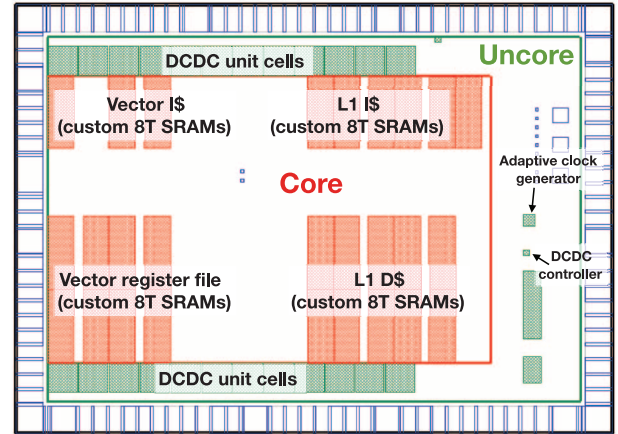


Fig. 11. Annotated floorplan of the design shows the placement of the SC DC-DC switches, controller, and adaptive clock around the RISC-V core.

TABLE II
CHIP SUMMARY

Technology	28 nm FDSOI
Die area	1305 $\mu\text{m} \times$ 1818 μm (2.37 mm^2)
Core area	880 $\mu\text{m} \times$ 1350 μm (1.19 mm^2)
Converter area	24 \times 90 $\mu\text{m} \times$ 90 μm (0.19 mm^2)
Voltage	0.45–1 V (1V FBB)
Frequency	93–961 MHz (1V FBB)
Power	8–173 mW (1V FBB)
SC density	11.0 $\text{fF}/\mu\text{m}^2$
SC power density	0.35 W/mm^2 at 88% efficiency

SC DC-DC mode allows accurate tracking within the small voltage ripple in each mode.

F. Physical Design

A multivoltage and multiclock design flow was used to construct the processor. Fig. 11 shows the processor floorplan, with the dotted red line separating the large core voltage domain at the top from the small uncore voltage domain at the bottom and sides of the chip. The custom SRAMs were manually placed within the core voltage domain. The DC-DC unit cells surround the core to minimize voltage drop. Two layers of thick upper-layer metal were dedicated to a power grid, where V_{out} and GND each utilize 25% of the chip area in each layer, and power rail analysis estimates a 2 mV voltage drop at 1 V and 100 mA (nominal operating condition). Ideally, converter power would come from bumps directly above the converter, but because only wire-bond packaging was available, all of the power is supplied through the pad frame in this implementation. Outside the core, V_{out} rails are not necessary, so the input voltages to the converters ($V_{DD,1.0}$ and $V_{DD,1.8}$) use the majority of the power routing resources to connect power coming from the pad frame to the converters.

IV. EXPERIMENTAL RESULTS

A prototype system was designed and implemented [31] in 28 nm ultra-thin body and BOX fully depleted silicon-on-insulator (UTTB FDSOI) technology [32]. Fig. 23 and Table II show the die micrograph and chip summary, respectively.

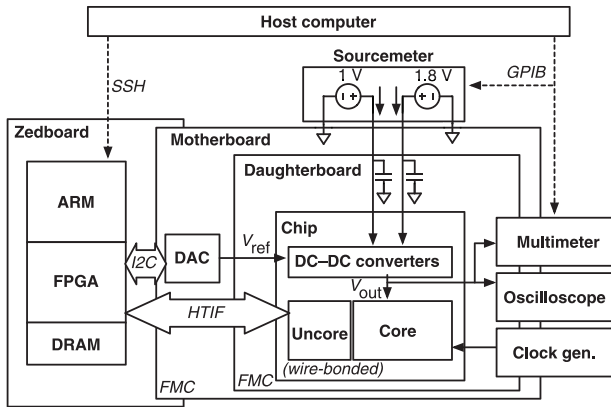


Fig. 12. Block diagram of the test setup for the system.




Fig. 13. Oscilloscope measurements of the core voltage V_{out} through a sense pad for all four on-chip regulation modes.

A. Measurement Setup

The measurement setup is shown in Fig. 12. The die is packaged using chip-on-board wire bonding to a small daughterboard. There is decoupling capacitance for the 1 and 1.8 V inputs to the converter both on the chip and on the daughterboard. A multimeter or oscilloscope connects to sense points on the daughterboard to measure the output voltage rail supplied by the SC DC–DC converter. The daughterboard is connected over FMC to a motherboard which generates the necessary clock, supplies, and reference voltages. Additional testpoints on the motherboard connect to a sourcemeter to measure the input power provided to the SC DC–DC converter. The chip is controlled from a Zedboard, which includes a network-accessible ARM core with FPGA to connect to main memory and emulate system call operations.

B. DVFS for Improved Energy Efficiency

The measured traces of the rippling core voltage domain for all four possible configurations are shown in Fig. 13. The actual average output voltage is lower than the ideal divided output voltages due to charge sharing with the inherent decoupling capacitance of the core. (The relationship between ripple size and average output voltage is further discussed in Section IV-E.) For all possible converter topologies with adaptive clocking, the processor successfully boots Linux and




Fig. 14. Oscilloscope measurements of the core voltage V_{out} transitioning between different DVFS modes, illustrating 20 ns transitions.

runs user applications, demonstrating that complex digital logic operates reliably with an intentionally rippling supply voltage. A small margin on the minimum operating voltage (V_{min}) is required to support operation at V_{ref} instead of V_{avg} .

Tight integration of the on-chip SC DC–DC converter with the processor enables extremely fine-grained DVFS. Fig. 14 shows that the processor can switch between operating modes in approximately 20 ns. These fast mode transitions enable new DVFS algorithms that can operate at much shorter time scales.

The main goal of on-chip conversion is to improve energy efficiency through DVFS. Fig. 15(a) shows the energy efficiency of the system, for both the baseline system with ideal off-chip regulation (bypass mode) and the four topologies. Energy efficiency is measured using a double-precision floating-point matrix multiplication kernel in terms of billions of floating-point operations per watt (GFLOPS/W), which is the inverse of energy per operation. Fig. 15(b) shows how different topologies change the absolute power and delay of the processor. FBB of the microprocessor in FDSOI enables threshold voltage control during runtime to trade off performance and power [33], as shown for this design in Fig. 15(c) and (d). By using the on-chip converter to generate the lowest output voltage, the system achieves a peak efficiency of 26.2 GFLOPS/W.

C. System Efficiency

The efficiency of voltage converters is generally computed by measuring the current and voltage on both the input and output of the converter to measure the ratio of power delivered to power supplied. However, for the proposed system, efficiency defined in this way is not easily measurable. First, it is difficult to measure on-chip voltage and current, because the voltage is rippling very quickly. Second, even if power output of the converter could be measured, this metric would ignore the impact of the adaptive clock, which is an important loss component. Therefore, a different method is required to measure the efficiency of the implemented system.

This paper defines system efficiency with a metric that fairly accounts for the adaptive clock and does not require measuring on-chip voltage and current. To characterize the processor load, the bypass mode is used to directly supply the core with an ideal off-chip voltage source. A self-checking benchmark is run

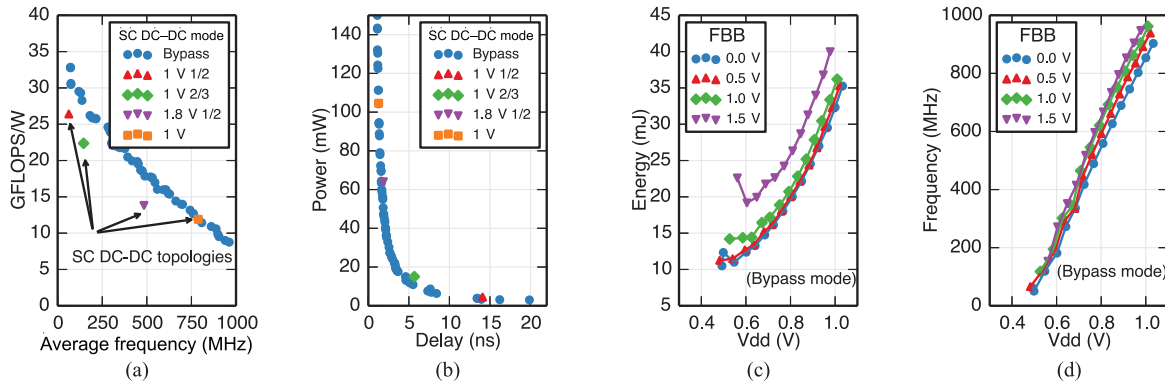


Fig. 15. Energy efficiency and performance versus voltage characteristics for the vector accelerator. (a) Energy efficiency of double-precision, floating-point, matrix multiplication kernel running on the vector accelerator. (b) Power versus delay tradeoff for bypass mode and different DC–DC topologies. (c) Impact of forward body bias (FBB) on system energy. (d) Impact of FBB on operating frequency.

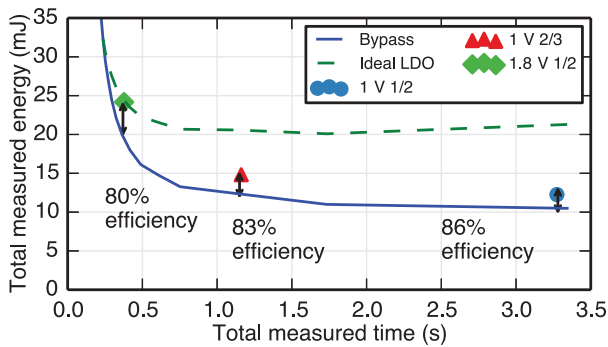


Fig. 16. Measured system efficiency (including overhead of nonideal adaptive clocking).

for a fixed number of cycles at different voltages, and a binary search is performed at each voltage point to find the maximum frequency. At the maximum frequency, the total elapsed time and total energy to run the fixed-length benchmark is measured, where the energy is computed by measuring the current drawn from the off-chip supply and the delivered voltage is measured from sense points on V_{out} , to remove the voltage drop across the on-chip bypass-mode power gates from the efficiency calculation. This provides the blue curve in the figure of energy versus time, and represents a 100%-efficient off-chip regulator.

Then, for each DC–DC mode, the same benchmark is run for the same number of cycles, and the total elapsed time and energy is measured. Due to nonidealities of the converter, it takes more energy to perform the same task in the same amount of time. Therefore, system efficiency is defined as the ratio of energy required to finish the same workload in the same time. This metric includes all sources of overhead, including nonidealities in the adaptive clock. Fig. 16 shows the measured voltage conversion efficiency ranges from 80%–86% for different output voltage modes.

D. Loss Analysis

The 14%–20% system efficiency losses are attributed to both converter losses and nonideal adaptive clocking based on measured results.




Fig. 17. Approximate converter efficiency measured by numerically integrating voltage waveforms with precharacterized load current.

1) *Standalone Converter Losses:* The efficiency of the converter alone is estimated by characterizing the power at each voltage using a repetitive microbenchmark and numerically integrating the waveform at V_{out} to determine the ratio of input to output power. These results are an approximate measure of efficiency, because the ripple measured from off-chip will not perfectly match the true on-chip voltage waveform. Fig. 17 shows that the converter alone achieves a maximum efficiency above 90%, and compares this efficiency to the measured system efficiency (and the corresponding power density of the benchmark) and the hypothetical efficiency for a system running with a fixed frequency clock at the minimum observed voltage. A wide range of power densities was measured by changing the proportion of 24 SC DC–DC unit cells that are enabled, which contributes to the discontinuities in the data.

2) *Adaptive Clocking Losses:* Analytical modeling of the adaptive clock, based on measured results, predicts a 5%–10% efficiency loss due to nonideal adaptive clocking. A simple experiment, illustrated in Fig. 18, shows how clock frequency margins, required to compensate for imperfect adaptive clock





Fig. 18. Simulated impact of an increase in timing margin on system efficiency.

Numerical simulation:



Measurement-based estimate: ~7% delay overhead




Fig. 19. Numerical simulation based on measured results that estimates the efficiency loss due to nonideal adaptive clocking.

generation, translate to system efficiency losses. First, the characteristic total energy versus total runtime of the core is plotted based on measured results. A hypothetical converter with 90% efficiency would require more total energy to complete the same workload in the same amount of time. If the hypothetical converter also increases the critical path delay by 5% due to any nonidealities, the curve shifts to the right due to the increase in runtime, and shifts slightly up due to increased leakage integration time. These shifts correspond to a decrease in efficiency, because an increase in runtime can also be interpreted as requiring a higher operating voltage to achieve the same overall runtime. In this case, a 5% increase in average delay would equate to an approximately 5% decrease in system efficiency. The exact translation from delay increase to efficiency depends on the slope of the energy–delay curve for a particular design and technology.

The quantitative effect of nonideal adaptive clocking can be estimated with numerical simulation based on measured results. The simulation, shown in Fig. 19, divides a voltage ripple into small time steps and tracks the progression of a signal through the replica, clock, and critical path based on the delay at each instantaneous voltage. The voltage ripple and




Fig. 20. Simulated energy efficiency improvement for interleaved and simultaneous-switching converters with differing loads. The frequency at maximum efficiency is annotated for each conversion method.

voltage versus frequency characteristics of the replica and critical path are supplied by measured results, while the insertion delay is supplied by back-annotated timing analysis. Two main effects cause nonideal adaptive clock tracking. First, each path has different characteristic delay versus voltage tradeoffs due to different gate types or different relative contribution of gate or wire delay. Second, the insertion delay of the clock tree means that the replica and critical path see different voltages, but the clock tree itself will compensate to diminish this effect [34]. Therefore, after many simulated cycles there is a distribution of extra FO4 stages that could be computed by the core before the margined adaptive clock edge arrives. The average of the distribution corresponds to the overhead of the adaptive clock, and the numerical simulation predicts an average cycle time increase of 7%. The losses due to nonideal clocking are already included in the system efficiency measurement, so this prediction serves as an estimate of the relative contribution of nonideal clocking to total losses.

E. Effect of V_{ref} on Efficiency

As discussed in Section III-D, the choice of V_{ref} sets the size of the output voltage ripple, and the load current automatically modulates the switching frequency f_{sw} . For the ideal case shown in Fig. 1, simultaneous-switching converters have essentially zero losses from $P_{C_{fly}}$, but in reality, a simultaneous-switching converter will still charge-share with the intrinsic capacitance of the load. Fig. 20 analytically compares the efficiency as a function of switching frequency for three 1 V 2:1 mode converters: a conventional interleaved converter, the proposed simultaneous-switching converter, and a hypothetical simultaneous-switching converter with no load capacitance. Because the interleaved converter has more charge-sharing losses, it incurs high losses for large ripple sizes at low switching frequencies, and therefore has a higher optimal switching frequency. A simultaneous-switching converter that closely matches the implemented system, with an output load capacitance equal to the converter capacitance, has charge-sharing




Fig. 21. Effect of the lower bound reference voltage on the average output voltage, average power, and maximum processor frequency for the 1 V 2/3 mode.




Fig. 22. Effect of lower bound reference voltage on efficiency for the 1 V 2/3 mode.




Fig. 23. Die micrograph.

losses with the output load that cause an approximate 5% efficiency loss versus an ideal simultaneous-switching converter. No explicit decoupling capacitance was added to the core in order to minimize charge-sharing losses.

Charge sharing also causes the average output voltage to fall below the ideal divided output voltage for each converter type. Measurements confirm that charge sharing with the processor's intrinsic capacitance causes the average output voltage to change for different V_{ref} choices, as shown in Fig. 21. While an optimal V_{ref} maximizes efficiency, suboptimal V_{ref} points could be chosen to achieve finer-grain control of the average output voltage (and therefore average performance) than switching between the fixed conversion topologies. As the load current changes, the optimal V_{ref} will also change.

Fig. 22 shows measured system efficiency for different V_{ref} (and therefore different ripple size and switching frequency). The peak efficiency occurs at a point where the sum of charge sharing losses (proportional to ripple size) and switching losses (proportional to switching frequency and therefore inversely

proportional to ripple size) is minimized. Additionally, larger ripples increase the range in voltages seen by the adaptive clock, reducing tracking between the replica path and the core's critical path, and further decreasing efficiency for lower V_{ref} and larger ripple. In this implementation, V_{ref} is chosen for maximal long-term average efficiency, but in future work, V_{ref} can be automatically changed to match expected load conditions.

V. CONCLUSION

The combination of the RISC-V architecture, low-voltage SRAM, and wide operating range DVFS enabled by on-chip voltage conversion and adaptive clocking achieves 26.2 GFLOPS/W with the 1 V 1/2 DC-DC configuration when computing double-precision matrix-multiplication using the vector accelerator. A simultaneous-switching SC DC-DC built with MOS capacitors and a centralized lower-bound controller reconfigures to provide four output voltages between 0.45 and 1 V, and achieves high converter efficiency by avoiding charge sharing. An adaptive clock translates high converter efficiency to high system efficiency by maximizing clock frequency for the voltage waveform to the core. Measurement results show that the system achieves 80%–86% system efficiency, with losses attributed to traditional converter switching losses, charge-sharing with the intrinsic capacitance of the core, and imperfect clock tracking. The simultaneous-switching approach described in this paper provides a low cost and high efficiency DVFS solution for low-power mobile devices.

ACKNOWLEDGMENT

The authors would like to thank T. Burd, J. Dunn, O. Thomas, and A. Vladimirescu for their contributions.

REFERENCES


- [1] T. Burd, T. Pering, A. Stratakos, and R. 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. Gupta, G.-Y. Wei, and D. Brooks, "System level analysis of fast, per-core DVFS using on-chip switching regulators," in *Proc. IEEE Int. Symp. High Perform. Comput. Archit.*, Feb. 2008, pp. 123–134.
- [3] D. Truong *et al.*, "A 167-processor computational platform in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1130–1144, Apr. 2009.
- [4] S. Digne *et al.*, "Within-die variation-aware dynamic-voltage-frequency-scaling with optimal core allocation and thread hopping for the 80-core teraFLOPS processor," *IEEE J. Solid-State Circuits*, vol. 46, no. 1, pp. 184–193, Jan. 2011.
- [5] Z. Toprak-Deniz *et al.*, "Distributed system of digitally controlled microregulators enabling per-core DVFS for the POWER8 microprocessor," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2014, pp. 98–99.
- [6] N. Kurd *et al.*, "Haswell: A family of IA 22 nm processors," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2014, pp. 112–113.
- [7] E. Burton *et al.*, "FIVR: Fully integrated voltage regulators on 4th generation Intel Core SoCs," in *Proc. IEEE Appl. Power Electron. Conf. Expo.*, Mar. 2014, pp. 432–439.
- [8] A. Nalamalpu *et al.*, "Broadwell: A family of IA 14 nm processors," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2015, pp. C314–C315.
- [9] M. Seeman, V. Ng, H.-P. Le, M. John, E. Alon, and S. Sanders, "A comparative analysis of switched-capacitor and inductor-based DC-DC conversion technologies," in *Proc. IEEE Workshop Control Model. Power Electron.*, Jun. 2010, pp. 1–7.
- [10] H. Krishnamurthy *et al.*, "A 500 MHz, 68% efficient, fully on-die digitally controlled buck voltage regulator on 22 nm tri-gate CMOS," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2014, pp. 1–2.

- [11] E. Alon and M. Horowitz, "Integrated regulation for energy-efficient digital circuits," *IEEE J. Solid-State Circuits*, vol. 43, no. 8, pp. 1795–1807, Aug. 2008.
- [12] T. 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 *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2014, pp. 90–91.
- [13] R. Jain *et al.*, "A 0.45–1 V fully integrated reconfigurable switched capacitor step-down DC–DC converter with high density MIM capacitor in 22 nm tri-gate CMOS," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2013, pp. C174–C175.
- [14] R. Jain, S. Kim, V. Vaidya, K. Ravichandran, J. Tschanz, and V. De, "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.
- [15] H.-P. Le, S. 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.
- [16] S. Clerc *et al.*, "A 0.33 V/–40 C process/temperature closed-loop compensation SoC embedding all-digital clock multiplier and DC–DC converter exploiting FDSOI 28 nm back-gate biasing," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2015, pp. 1–3.
- [17] S. Kim *et al.*, "Enabling wide autonomous DVFS in a 22 nm graphics execution core using a digitally controlled hybrid LDO/switched-capacitor VR with fast droop mitigation," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2015, pp. 1–3.
- [18] M. Turnquist *et al.*, "Fully integrated DC–DC converter and a 0.4V 32-bit CPU with timing-error prevention supplied from a prototype 1.55 V Li-ion battery," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2015, pp. C320–C321.
- [19] A. Waterman, Y. Lee, D. A. Patterson, and K. Asanović, "The RISC-V instruction set manual—Volume I: Base user-level ISA," EECS Dept., Univ. California, Berkeley, CA, USA, Tech. Rep. UCB/EECS-2011-62, May 2011 [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2011/EECS-2011-62.html>
- [20] R. Jevtic *et al.*, "Per-core DVFS with switched-capacitor converters for energy efficiency in manycore processors," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 23, no. 4, pp. 723–730, Apr. 2015.
- [21] M. D. Seeman, "A design methodology for switched-capacitor DC–DC converters," Ph.D. dissertation, EECS Dept., Univ. California, Berkeley, CA, USA, May 2009 [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-78.html>
- [22] S. Lee and T. Sakurai, "Run-time voltage hopping for low-power real-time systems," in *Proc. 37th Annu. Des. Autom. Conf.*, 2000, pp. 806–809.
- [23] Y. Lee *et al.*, "A 45 nm 1.3 GHz 16.7 double-precision GFLOPS/W RISC-V processor with vector accelerators," in *Proc. IEEE Eur. Solid-State Circuits Conf.*, Sep. 2014, pp. 199–202.
- [24] Y. Lee *et al.*, "Exploring the tradeoffs between programmability and efficiency in data-parallel accelerators," *ACM Trans. Comput. Syst.*, vol. 31, no. 3, pp. 6:1–6:38, Aug. 2013.
- [25] C. Batten, R. Krashinsky, S. Gerding, and K. Asanović, "Cache refill/access decoupling for vector machines," in *Proc. Int. Symp. Microarchit.*, Dec. 2004, pp. 331–342.
- [26] R. M. Russell, "The CRAY-1 computer system," *Commun. ACM*, vol. 21, no. 1, pp. 63–72, Jan. 1978.
- [27] H.-P. Le, J. Crossley, S. 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 *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2013, pp. 372–373.
- [28] M. Seeman and S. Sanders, "Analysis and optimization of switched-capacitor DC–DC converters," in *Proc. IEEE Workshops Comput. Power Electron.*, Jul. 2006, pp. 216–224.
- [29] T. Kobayashi, K. Nogami, T. Shirotori, Y. Fujimoto, and O. Watanabe, "A current-mode latch sense amplifier and a static power saving input buffer for low-power architecture," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 1992, pp. 28–29.
- [30] J. Kwak and B. Nikolic, "A 550–2260 MHz self-adjustable clock generator in 28 nm FDSOI," in *Proc. IEEE Asian Solid-State Circuits Conf.*, Nov. 2015, pp. 1–4.
- [31] B. Zimmer *et al.*, "A RISC-V vector processor with tightly-integrated switched-capacitor DC–DC converters in 28 nm FDSOI," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2015, pp. C316–C317.
- [32] N. Planes *et al.*, "28 nm FDSOI technology platform for high-speed low-voltage digital applications," in *Proc. IEEE Symp. VLSI Technol.*, 2012, pp. 133–134.
- [33] D. Jacquet *et al.*, "A 3 GHz dual core processor ARM Cortex-A9 in 28 nm UTBB FD-SOI CMOS with ultra-wide voltage range and energy efficiency optimization," *IEEE J. Solid-State Circuits*, vol. 49, no. 4, pp. 812–826, Apr. 2014.
- [34] K. Bowman, C. Tokunaga, T. Karnik, V. De, and J. Tschanz, "A 22 nm all-digital dynamically adaptive clock distribution for supply voltage droop tolerance," *IEEE J. Solid-State Circuits*, vol. 48, no. 4, pp. 907–916, Apr. 2013.
- [35] J. Kwong, Y. K. Ramadass, N. Verma, and A. P. Chandrakasan, "A 65 nm sub-microcontroller with integrated SRAM and switched capacitor DC–DC converter," *IEEE J. Solid-State Circuits*, vol. 44, no. 1, pp. 115–126, Jan. 2009.



Brian Zimmer (S'09–M'15) received the B.S. degree in electrical engineering from the University of California, Davis, CA, USA, in 2010, and the M.S. and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, CA, USA, in 2012 and 2015, respectively.

He is currently with the Circuits Research Group, Nvidia Corporation, Santa Clara, CA, USA. His research interests include energy-efficient digital design, with an emphasis on low-voltage SRAM design and variation tolerance.



Yunsup Lee (S'09) received the B.S. degree in both computer science and electrical engineering from Korea Advanced Institute of Science and Technology, Daejeon, Korea, in 2005, and the M.S. degree in computer science from the University of California, Berkeley, CA, USA, in 2011, where he is currently pursuing the Ph.D. degree.

He is the Chief Technology Officer with SiFive, Inc., San Francisco, CA, USA. His research interests include design of energy-efficient data-parallel accelerators and free and open instruction sets such

as RISC-V.

Mr. Lee was the recipient of the NVIDIA Graduate Fellowship from 2012 to 2014.



Alberto Puggelli (S'09) received the B.Sc. and two M.Sc. degrees in electrical engineering (*summa cum laude*) from the Politecnico di Milano, Milan, Italy, and the Politecnico di Torino, Turin, Italy, in 2006 and 2008, respectively. He received the M.Sc. degree in computer science and the Ph.D. degree in electrical engineering and computer science from the University of California, Berkeley, CA, USA, in 2013 and 2014, respectively.

He was with ST-Ericsson in 2009 and with Texas Instruments in 2011 and 2012, as an Intern Analog

Designer. He is currently the Director of Technology with Lion Semiconductor Inc., San Francisco, CA, USA. His research interests include the design of hybrid DC–DC voltage regulators.

Dr. Puggelli was the recipient of two Gold Medal Awards for the Best Student from the Politecnico di Milano. He received the AEIT Fellowship Isabella Sassi Bonadonna in 2010.



Jaehwa Kwak (S'10) received the B.S. and M.S. degrees in electrical engineering from Seoul National University, Seoul, Korea, in 2004 and 2006, respectively. He is currently pursuing the Ph.D. degree in electrical engineering at the University of California, Berkeley, CA, USA.

From 2006 to 2009, he was a Staff Researcher with GCT Research, Seoul, Korea, and worked on designing the digital controller of the wireless communication system. During his Ph.D. course, he was a Design Intern with Intel Corporation, Hillsboro,

OR, USA, in Summer 2011, and with the Advanced Micro Devices, Inc., Sunnyvale, CA, USA, in Summer 2013. His research interests include energy efficient microprocessor design, including self-adjustable clock generators, and advanced synchronization circuits.



Ruzica Jevtić (A'11–M'13) was born in Belgrade, Serbia, in 1981. She received the B.S. degree in electrical engineering from the University of Belgrade, Belgrade, Serbia, in 2004, and the Ph.D. degree with European Ph.D. mention in electrical engineering from the Technical University of Madrid, Madrid, Spain, in 2009.

She was a Postdoctoral Fellow with the University of California, Berkeley, CA, USA, from 2011 to 2013, where she was engaged in low power circuit design for energy efficient microprocessors. She was a Researcher with the University of Carlos III, Leganes, Spain, from 2013 to 2014. She is currently an Assistant Professor with the University Antonio de Nebrija, Madrid, Spain. Her research interests include fully integrated switched capacitor DC-DC design and soft error circuit resiliency.

Dr. Jevtić was the recipient of the Marie Curie International Outgoing Fellowship and Marie Curie Industry-Academia Partnerships and Pathways Fellowship.



Ben Keller (S'12) received the B.S. degree in engineering from Harvey Mudd College, Claremont, CA, USA, in 2010, and the M.S. degree in electrical engineering from the University of California, Berkeley, CA, USA, in 2015. Since 2012, he has been pursuing the Ph.D. degree at the Department of Electrical Engineering and Computer Sciences at the University of California, Berkeley.

He was a Research Intern with the NVIDIA, Santa Clara, CA, USA, and the National Institute of Standards and Technology, Gaithersburg, MD,

USA. His research interests include energy-efficient microprocessor design, fine-grained DVFS, and innovative digital hardware design paradigms.



Steven (Stevo) Bailey (S'11) was born in Richmond, VA, USA, in 1989. He received the B.S. degree in engineering science and the B.A. degree in physics from the University of Virginia, Charlottesville, VA, USA, in 2012, and the M.S. degree in electrical engineering from the University of California, Berkeley, CA, USA, in 2014. He is currently pursuing the Ph.D. degree at the University of California, Berkeley.

He was a Summer Research Intern with the Jet Propulsion Laboratory, Pasadena, CA, USA, in 2014, and with Nvidia Corporation, Santa Clara, CA, USA,

in 2015. His research interests include resilient digital integrated circuit design, digital signal processing and algorithms, and machine learning.



Milovan Blagojević (M'12) received the B.Sc. and M.Sc. degrees in electrical engineering from the University of Belgrade, Belgrade, Serbia, in 2010 and 2012, respectively. He defended a CIFRE Ph.D. thesis realized in cooperation among three institutions, the Berkeley Wireless Research Center, Berkeley, CA, USA, STMicroelectronics, Crolles, France, and Institut supérieur d'électronique de Paris, Paris, France, in December 2015.

He joined Intel Mobile Communications Group, Munich, Germany, in January 2016. His research

interests include energy-efficiency and power management of modern nanoscale digital VLSI systems, ultra low power architectural and circuit solutions for Internet of Things system on chips, and hardware and software solutions for 2-D and 3-D imaging systems.



Pi-Feng Chiu (S'10) received the B.S. and the M.S. degrees in electronic engineering from National Tsing Hua University, Hsinchu, Taiwan, in 2009 and 2010, respectively. She is currently pursuing the Ph.D. degree in electrical engineering at the University of California, Berkeley, CA, USA.

From 2010 to 2012, she was with the Industrial Technology Research Institute, Hsinchu, Taiwan. She joined the Berkeley Wireless Research Center, Berkeley, CA, USA, in 2012. She was an Intern with Samsung Electronics, CA, USA, in 2014. Her

research interests include resilient memory circuit design and emerging non-volatile memories.



Hanh-Phuc Le (S'07–M'13) received the B.S., M.S., and Ph.D. degrees in electrical engineering from Hanoi University of Science and Technology, Hanoi, Vietnam, KAIST, Daejeon, Korea, and the University of California (U.C.), Berkeley, CA, USA, in 2004, 2006, and 2013, respectively.

He recently joined the Department of Electrical and Computer and Energy Engineering, University of Colorado, Boulder, CO, USA, as an Assistant Professor. During his Ph.D. in electrical engineering at U.C. Berkeley in 2012, he cofounded Lion Semiconductor, San Francisco, CA, USA, where he served as the Chief Technology Officer. He held R&D positions at Oracle, Intel, Rambus, JDA Tech in Korea and the Vietnam Academy of Science and Technology, Ho Chi Minh, Vietnam. He has authored/coauthored 1 book chapter and over 30 journal and conference papers in the area of integrated power electronics and energy-efficient systems. He is also the inventor of nine U.S. patents (five granted and four pending). His research interests include miniaturized/on-die power conversions, smart power ICs and integrated systems for mobile/wearable applications.


Dr. Le was the recipient of the 2012–2013 IEEE Solid-State Circuits Society Predoctoral Achievement Award and the 2013 Sevin Rosen Funds Award for Innovation at U.C. Berkeley.



Po-Hung Chen (S'10–M'12) received the B.S. degree in electrical engineering from National Sun Yat-Sen University, Kaohsiung, Taiwan, the M.S. degree in electronics engineering from National Chiao-Tung University, Hsinchu, Taiwan, and the Ph.D. degree in electrical engineering from the University of Tokyo, Tokyo, Japan, in 2005, 2007, and 2012, respectively.

In 2011, he was a Visiting Scholar with the University of California, Berkeley, CA, USA, where he conducted research in power management cir-

cuits. Since 2012, he has been an Assistant Professor with the Department of Electronics Engineering, National Chiao-Tung University. His research interests include power management IC for energy harvesting, fully integrated power management ICs, wireless power transmission, and low-voltage low-power CMOS analog circuits.



Nicholas Sutardja (S'12) received the B.S. degree in electrical engineering and computer science and the B.A. degree in applied mathematics from the University of California, Berkeley, CA, USA, in 2012, where he is currently pursuing the Ph.D. degree in electrical engineering.

He worked on high-speed wireline receivers with Altera, San Jose, CA, USA, and sensors for pulse oximetry with ADI, Thief River Falls, MN, USA, in the summer of 2011 and 2014, respectively. His

research interests include mixed signal ICs, energy-efficient high-speed link systems, analog design methodologies, biomedical devices, and sensors.



Rimantas Avizienis was born in Los Angeles, CA, USA, on January 14, 1977. He received the B.S. and M.S. degrees in computer science from the University of California, Berkeley, CA, USA, in 1999 and 2011, respectively, where he is currently pursuing the Ph.D. degree in computer science.

From 1999 to 2008, he was with the Center for New Music and Audio Technology, University of California, Berkeley. His research interests include low-power VLSI design and software-defined radio.



Andrew Waterman received the B.S.E. degree in electrical and computer engineering from Duke University, Durham, NC, USA, in 2008, and the M.S. and Ph.D. degrees in computer science from the University of California, Berkeley, CA, USA, in 2011 and 2016, respectively.

He is Chief Engineer with SiFive, Inc., San Francisco, CA, USA.



Brian Richards (M'09) received the B.S. degree in electrical engineering from California Institute of Technology, Pasadena, CA, USA, in 1983, and the M.S. degree in electrical engineering and computer science from the University of California, Berkeley, CA, USA, in 1986.

In 1986, he joined the Research Staff with the University of California, Berkeley, where he worked on large-scale digital system design projects including the Infopad portable wireless multimedia terminal. A founding member of the Berkeley Wireless

Research Center, he is continuing the development and support of several ASIC and FPGA system design CAD tool flows.



Philippe Flatresse was born in Brest, France, in 1970. He received the Ph.D. degree in microelectronics from the Institut National Polytechnique de Grenoble, Grenoble, France, and CEA LETI, Grenoble, France, in 1998.

In 2000, he joined STMicroelectronics Central R&D, Crolles, France, to deploy the SOI digital design within the company. Thanks to this work, he has pioneered the partially and fully depleted SOI technologies and demonstrated their key advantages for low-power high-performance digital applications.

As a Design Architect, his research interests include exploration, development, and implementation of ultra-low power platforms able to work in an energy-efficient way on an ultra-wide range of operating points targeting high-growth application areas such as e-health, Internet of Things, and wearable human-computer interfaces.



Elad Alon (S'02–M'06–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.

In January 2007, he was an Associate Professor of electrical engineering and computer sciences as well as a Co-Director of the Berkeley Wireless Research Center (BWRC), University of California, Berkeley, CA, USA. He has held consulting, visiting, or advisory positions at Lion Semiconductor, Wilocity, Cadence, Xilinx, Oracle, Intel, AMD,

Rambus, Hewlett Packard, and IBM Research, where he worked on digital, analog, and mixed-signal integrated circuits for computing, test and measurement, and high-speed communications. His research interests include energy-efficient integrated systems, including the circuit, device, communications, and optimization techniques used to design them.

Dr. Alon was the recipient of the IBM Faculty Award in 2008, the 2009 Hellman Family Faculty Fund Award as well as the 2010 UC Berkeley Electrical Engineering Outstanding Teaching Award. He has coauthored 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 and 2013 Custom Integrated Circuits Conference Best Student Paper Awards.



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.

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

Dr. Asanović is an ACM Distinguished Scientist.



Borivoje Nikolić (S'93–M'99–SM'05) 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 in electrical and computer engineering from the University of California, Davis, CA, USA, in 1999.

In 1999, he joined the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, USA, where he is currently a National Semiconductor Distinguished Professor of Engineering. He is coauthor of the book *Digital Integrated Circuits: A Design Perspective*, 2nd ed. (Prentice-Hall, 2003). His research interests include digital, analog, and RF integrated circuit design and VLSI implementation of communications and signal processing systems.

Dr. Nikolić was a Distinguished Lecturer of the IEEE Solid-State Circuits Society in 2014–2015. He was the 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 at the University of California, Davis, in 1999, as well as the City of Belgrade Award for the Best Diploma Thesis in 1992. For work with his students and colleagues, he received the Best Paper Awards at the IEEE International Solid-State Circuits Conference, Symposium on VLSI Circuits, the IEEE International SOI Conference, the European Solid-State Device Research Conference, the European Solid-State Circuits Conference, the S3S Conference, and the ACM/IEEE International Symposium on Low-Power Electronics.