

A RISC-V Processor SoC With Integrated Power Management at Submicrosecond Timescales in 28 nm FD-SOI

Ben Keller, *Student Member, IEEE*, Martin Cochet, *Student Member, IEEE*, Brian Zimmer, *Member, IEEE*, Jaehwa Kwak, *Student Member, IEEE*, Alberto Puggelli, *Member, IEEE*, Yunsup Lee, *Student Member, IEEE*, Milovan Blagojević, *Member, IEEE*, Stevo Bailey, *Student Member, IEEE*, Pi-Feng Chiu, *Student Member, IEEE*, Palmer Dabbelt, Colin Schmidt, Elad Alon, *Senior Member, IEEE*, Krste Asanović, *Fellow, IEEE*, and Borivoje Nikolić, *Fellow, IEEE*

Abstract—This paper presents a RISC-V system-on-chip (SoC) with integrated voltage regulation, adaptive clocking, and power management implemented in a 28 nm fully depleted silicon-on-insulator process. A fully integrated simultaneous-switching capacitor DC-DC converter supplies an application core using a clock from a free-running adaptive clock generator, achieving high system conversion efficiency (82%–89%) and energy efficiency (41.8 double-precision GFLOPS/W) while delivering up to 231 mW of power. A second core serves as an integrated power-management unit that can measure system state and actuate changes to core voltage and frequency, allowing the implementation of a wide variety of power-management algorithms that can respond at submicrosecond timescales while adding just 2.0% area overhead. A voltage dithering program allows operation across a wide continuous voltage range (0.45 V–1 V), while an adaptive voltage-scaling algorithm reduces the energy consumption of a synthetic benchmark by 39.8% with negligible performance penalty,

demonstrating practical microsecond-scale power management for mobile SoCs.

Index Terms—Adaptive clock, energy-efficient processor, fine-grained adaptive voltage scaling (AVS), integrated voltage regulator, power management unit (PMU), RISC-V, voltage dithering.

I. INTRODUCTION

ENERGY efficiency is the key constraint in modern systems-on-chip (SoCs). Server-class chips are thermally limited and require better energy efficiency to improve performance, while the utility of mobile and IoT devices depends substantially on low energy consumption to prolong battery life. These constraints demand advanced power-management systems that can improve energy efficiency without sacrificing performance by adapting to the demands and constraints of particular workloads. However, many modern workloads demonstrate rapid changes in program behavior [1]–[3]. Mobile and IoT devices, often driven by an unpredictable sensor or a user input, have especially bursty workloads, with short periods of high compute demand followed by long periods of inactivity [4]. Accordingly, the most effective power-management systems must be able to respond at microsecond timescales, rapidly adjusting voltage and frequency to track changes in workload and maximize energy savings while minimizing performance loss.

One barrier to implementing fine-grained power management is the ability to quickly switch between different voltage and frequency operating modes. Most systems supply a small number of voltage domains via off-chip regulators. These off-chip components generally have slow transition times due to the large time constants of the off-chip passives, and the additional component count can substantially increase system cost. Integrated regulation can achieve much faster transition times because the passive elements are much smaller. However, area limitations and the lower quality of integrated inductors make it difficult to achieve high conversion efficiencies or power densities in fully integrated systems, negating any energy savings [5]–[9]. Similarly, most clock generation for digital circuits is done by large power-hungry phase-locked loop (PLL) circuits. These PLLs typically take many cycles to lock to a new frequency target. Thus, traditional regulation and clocking circuits severely limit the speed of voltage and frequency changes.

Manuscript received November 21, 2016; revised February 9, 2017; accepted March 21, 2017. Date of publication April 1, 2017; date of current version June 22, 2017. This paper was approved by Guest Editor Wim Dehaene. This work was supported in part by DARPA PERFECT under Award HR0011-12-2-0016, in part by GRC under Grant SRC/TxACE 1836.136, in part by Intel ARO, in part by AMD, in part by NVIDIA through the NVIDIA Fellowship, and in part by the NSF GRFP. (*Corresponding author: Ben Keller.*)

B. Keller, J. Kwak, S. Bailey, P.-F. Chiu, P. Dabbelt, C. Schmidt, 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: bkeller@eecs.berkeley.edu).

M. Cochet is with STMicroelectronics, 38920 Crolles, France, and with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA, and also with Aix-Marseille University & CNRS, IM2NP (UMR 7334), 13284 Marseille, France.

B. Zimmer was with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA. He is now with NVIDIA, Santa Clara, CA 95050 USA.

A. Puggelli was with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA. He is now with Lion Semiconductor Inc., San Francisco, CA 94107 USA.

Y. Lee was with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA. He is now with SiFive, Inc., San Francisco, CA 94107 USA.

M. Blagojević was with STMicroelectronics, 38920 Crolles, France and the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA. He is now with Intel Deutschland GmbH, Neubiberg 85579, Germany.

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.2017.2690859

TABLE I
COMPARISON TO PRIOR WORK

	This work	[10]	[11]	[41]	[42]	[43]
Technology	28nm FD-SOI	28nm FD-SOI	22nm FinFET	22nm SOI	20nm HKMG	65nm CMOS
Regulation Technique	Simultaneous-switching SC	Simultaneous-switching SC	Buck	LDO	-	Interleaved SC
Off-chip components	No	No	Yes	No	Yes	No
Conversion efficiency	82-89%	80-86%	90%	60-90%	-	60-81%
Conversion area overhead	36%	16%	-	-	-	17%
Energy efficiency (GFLOPS/W)	41.8 DP	26.2 DP	8 DP*	12 DP*	-	-
Transition time (regulators)	0.02 μ s	0.02 μ s	0.32 μ s*	300 μ s	-	2 μ s*
Response time (AVS)	<1 μ s	-	10 000 μ s*	-	\sim 100 μ s	30 μ s*

* Estimated

While integrated regulation has given some systems high conversion efficiencies, integrated power control logic is also required to achieve fine-grained power management. Some previous systems have lacked power-management capability [10]. Others have supported power management, but actuate voltage and frequency changes only in response to operating system requests at millisecond timescales [11], [12]. Hardware-actuated power management has been proposed in [3] and [13]–[19], but only a system implementation that includes fast efficient voltage regulation, lightweight clock generation, and dedicated power-management hardware can demonstrate the efficacy of the approach. Table I provides a summary of prior work.

This paper implements a processor SoC with integrated voltage regulation and power management. Simultaneous-switching switched-capacitor DC–DC (SC-DCDC) converters are tightly integrated with an adaptive clock generator to achieve high system conversion efficiencies and power densities suitable for a mobile-class processor [10]. The generated voltage and clock supply a RISC-V [20] application processor with a vector accelerator that loads the SC-DCDC converters with up to 231 mW of power demand. A separate fixed-voltage RISC-V power-management unit (PMU) processor can actuate changes to each of these systems in response to changes in workload at submicrosecond timescales, enabling fully autonomous power management that can reduce core energy consumption by up to 39.6% with minimal performance penalty. This paper expands upon the initial presentation in [21] with more detailed design descriptions of the PMU processor, the adaptive clock generator, and top-level physical design. Additional measurement results show the result of diverse power management algorithms, demonstrating the utility and flexibility of the SoC. Section II describes the system implementation in detail, and Section III details measured performance results of the system. Section IV describes the benefits of a variety of power-management algorithms running on the SoC, and Section V concludes this paper.

II. SoC DESIGN

Fig. 1 summarizes the components of the SoC. The system is partitioned into two voltage domains: 1) a variable-voltage

core domain that contains the application processor and 2) a fixed 1 V uncore domain that contains the power measurement and control blocks. Because the core can operate at varying frequencies and voltages, all digital communication between the core and the uncore uses level shifters and asynchronous queues.

The core voltage domain is supplied by an integrated SC-DCDC converter. The reconfigurable converter can operate in four modes to downconvert external 1 V and 1.8 V supplies. Each conversion mode represents a configuration of the flying capacitors that result in voltage conversion ratios of one-half or two-thirds of the external supply voltages, generating voltages averaging 1 V (bypass mode), 0.9 V (1.8 V 1/2 mode), 0.67 V (1 V 2/3 mode), and 0.5 V (1 V 1/2 mode). The SC-DCDC converter is partitioned into forty-eight $90 \times 90 \mu\text{m}$ unit cells and equipped with a lower-bound (hysteretic) controller. When the converter output voltage crosses a fixed external reference voltage, a toggle signal is generated and distributed to each unit cell in a clock tree balanced by the place-and-route tool. All unit cells switch simultaneously, and so the converter output voltage ripples significantly about the average value, improving the system conversion efficiency [10]. Because the generated voltage can ripple over a wide range, the core clock is produced by a free-running adaptive clock generator that modulates its frequency to track fine-grained voltage variations. Core SRAMs are implemented as custom 8T macros to enable low-voltage operation.

The frequency of the core clock and the SC-DCDC toggle signal are recorded by counters that can be read by the PMU processor, which resides in the uncore domain. The PMU is fully programmable and executes power-management programs that read on-chip counters and alter the SC-DCDC output voltage mode while the application core is running. In addition, the threshold voltage of the logic in the core voltage domain can be manipulated by an integrated back-bias generator that can supply up to 1.8 V of forward body bias (FBB) [22]. Measurement circuits, such as a programmable current mirror load and an integrated SC-DCDC waveform reconstruction unit, allow for the full characterization of key system parameters of the SoC [23]. The uncore voltage

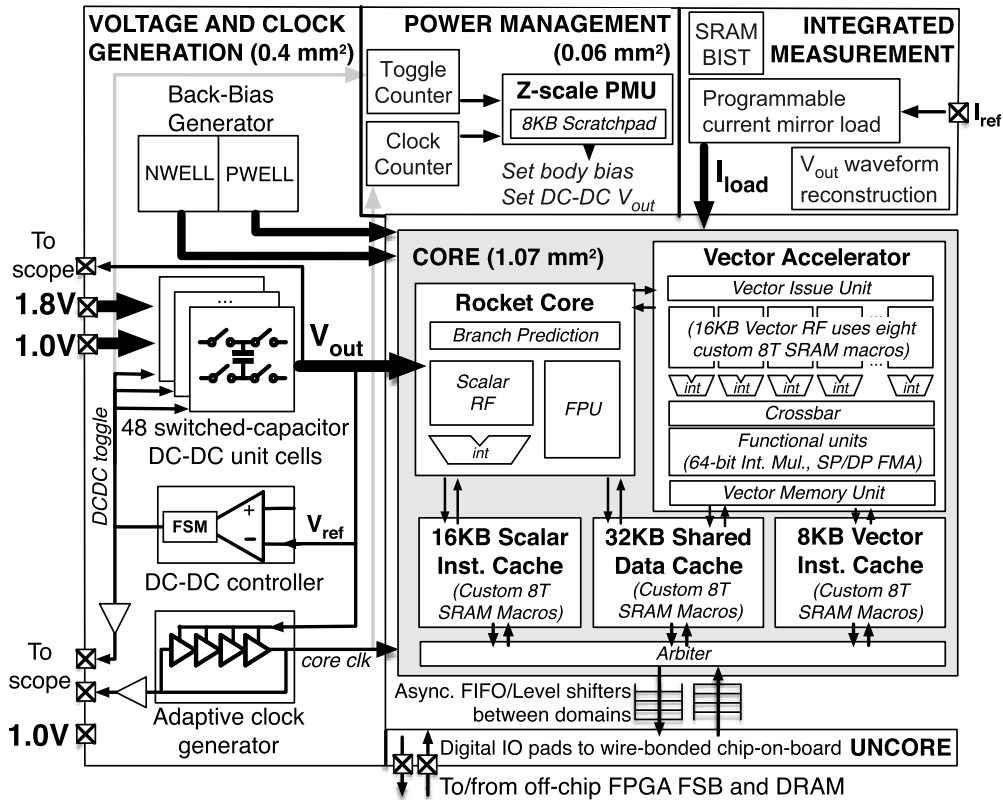


Fig. 1. Block diagram of the processor SoC.

domain also contains logic to serialize off-chip accesses for communication with the host FPGA system.

A. Compute Core

The core voltage domain contains an application processor based on the open-source Rocket processor [24] with a vector coprocessor similar to the implementations described in [10] and [25]. Rocket is a 64-bit in-order single-issue processor that implements the free and open RISC-V instruction set architecture (ISA) [20]. The scalar core implements a simple branch predictor, L1 caches, and a memory-management unit that supports page-based virtual memory and boots modern operating systems such as Linux.

A high-performance single-lane vector coprocessor accelerates compute-intensive workloads. The coprocessor implements a decoupled vector-fetch architecture that can achieve high compute efficiency via systolic dataflow [26]. The vector unit is tightly integrated with the Rocket core, which issues commands to the accelerator via a custom extension to the RISC-V ISA. The memory system, including the 32 KB L1 data cache, is shared between the scalar core and the vector accelerator, although the vector unit implements its own 8 KB instruction cache. Compared with the implementation in [10], the vector accelerator is improved by increasing floating-point throughput to two single-precision and double-precision fused multiply-add (FMA) units supporting up to four floating-point operations per cycle [27] and by adding separate long-latency functional units dedicated to the accelerator.

B. Adaptive Clock Generator

An adaptive clock generator supplies the clock to the core logic, improving system energy efficiency by tracking the voltage ripple produced by the simultaneous-switching SC-DCDC converters and adjusting the clock period on a cycle-by-cycle basis. Fig. 2 shows the schematic of the adaptive clock generator. The delay units are composed of four tunable delay banks supplied by the same rippling voltage as the core logic. Each bank uses a different cell for its delay element, and the banks can be tuned independently by changing the control settings of the multiplexers in each bank. The first bank uses a custom buffer cell designed to balance rise and fall times. The remaining banks each contain a library standard cell that frequently occurred in the core netlist produced by the synthesis tool. Each bank uses a different combination of pMOS/nMOS ratio and gate length, as these characteristics affect the voltage–frequency relationship of the cells, and the replica paths must be able to track similar cell variations in the critical paths of the core. This combination of multiple delay cells tracks the core critical path more accurately than a single standard cell [28]. The two delay units are identical, but can be independently tuned to adjust the duty cycle of the generated clock. This free-running adaptive clock generator was implemented alongside the clock generator from [10] that selects edges from a fixed reference for direct comparison.

The presence of insertion delay in the core clock tree negatively impacts the ability of the adaptive clock to accurately track the supply voltage. In the ideal case, the generated clock

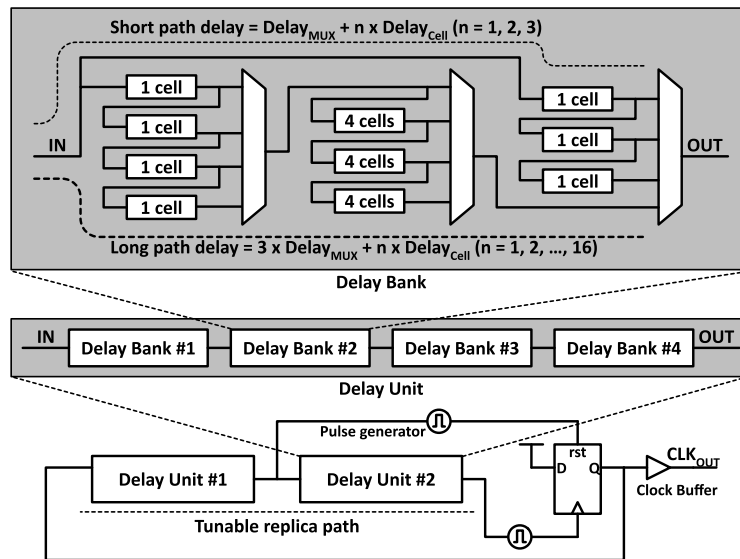


Fig. 2. Block diagram of the adaptive clock generator.

would instantaneously propagate to all clock sinks and the voltage of the core logic would be the same as that of the replica paths in the adaptive clock generator. In practice, the clock tree causes the clock edges to be received by the register sinks only after some delay. During this delay period, the core voltage further decreases, so the core operates more slowly than the equivalent replica path when the edges were generated. While the progression of the clock edges through the clock tree also slows down due to the decrease in voltage [29], this clock-data compensation effect does not fully guard against insertion delay effects because the voltage-dependent delay characteristics of the clock tree and logic critical paths may not be the same. Large insertion delays require additional timing margin that reduces system efficiency, so care was taken during physical design to minimize the insertion delay of the core clock tree.

C. Counters and Power Measurement

For power-management algorithms to be effective, the power consumption of the system must be known or estimated. Some power-management schemes implement complex models based on the activity of a number of key signals to estimate power consumption [12], [30], while others place a sense resistor between the supply and the load [31] or halt the normal operation of the voltage regulators to measure current draw [32]. In contrast, the SC-DCDC system in this paper enables simple, accurate, and noninvasive measurement of core power because power correlates directly with the switching frequency of the regulators. The SC-DCDC voltage regulators toggle when the supplied voltage reaches a known fixed reference voltage, and the capacitance of the system is fixed; therefore, each DC-DC toggle transfers a fixed amount of energy to the core. As shown in Fig. 3, a slow toggling frequency implies lower power consumption and an increased toggling frequency implies high power consumption. Therefore, we instrument the SC-DCDC toggle

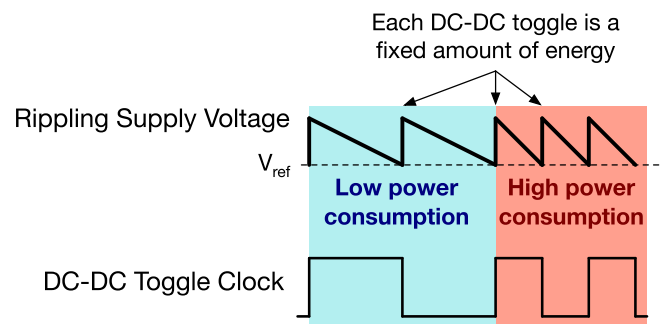


Fig. 3. Correlating the SC-DCDC toggle clock with core power consumption.

signal with a counter that compares its frequency with a fixed reference to measure power, while a second counter measures the core clock frequency.

D. Power Management Core

The PMU contains a 32-bit 3-stage single-issue in-order RISC-V processor based on Z-scale [33]. Implementing the PMU with a processor instead of fixed-function hardware enables experimentation with a wide range of power-management algorithms. The power-management core forgoes caches in favor of an 8 KB scratchpad memory, which is 128 bits wide, and mapped to the upper half of the 4 GB physical memory space (the lower half is reserved for memory traffic from the application processor). The PMU supports the RV32IM instruction extensions and is fully programmable via the RISC-V software toolchain. Power-management programs are written into the scratchpad during system boot. Fig. 4 shows the pipelines of the two processors in the SoC, and Table II compares their key features.

The three-stage design minimizes gate count while enabling sufficient performance to enable fine-grained power management. The first stage of the pipeline fetches an instruction out of a 128-bit instruction buffer (the value of the program counter is calculated in the previous cycle). The second stage

TABLE II
COMPARISON OF THE TWO PROCESSOR CORES

Feature	Rocket (excluding coprocessor)	Z-scale
Instruction Set	RISC-V RV64G	RISC-V RV32IM
Pipeline	Five-stage single-issue in-order	Three-stage single-issue in-order
Memories	32KB DCache, 16KB ICache	8KB unified scratchpad
Standard Cell Count	196K	18K
Area (Cells + Memories)	0.461 mm ²	0.053 mm ²




Fig. 4. Pipeline diagrams of the integrated processors.

of the pipeline decodes a RISC-V instruction, reads the register file, and executes an ALU instruction. Branches are resolved in the second stage, so the instruction in the fetch stage is flushed when a branch is taken. Writeback is isolated into the third stage of the pipeline, reducing the fanout delay on the write port of the register file. The result in the third stage is bypassed to the second stage, eliminating the need for stalls in some cases. The memory stage and the multiplication/division pipeline stages are also in the third pipeline stage. The multiply and divide units compute 1 bit of the operation each cycle, requiring 32 cycles to complete a multiply or divide. Loads and stores directly access the scratchpad. An extra pipeline register was added in the arbiter between instruction and data memory requests to eliminate a long critical path and allow the PMU to meet the setup time constraint of the uncore clock domain.

All system control registers that monitor and control the various subsystems of the chip, including the SC-DCDC converter voltage mode, are mapped into the PMU's control status register (CSR) space. Additionally, the application core and the PMU can communicate directly via inter-processor interrupts, and each has a register mapped directly into the CSR space of the other system, allowing arbitrary data to be communicated between the two cores. The PMU issues CSR reads to the SC-DCDC toggle counters and core clock counters to detect changes in workload. It can then issue CSR writes to actuate mode changes in a fast feedback loop as shown in Fig. 5.

E. Custom Low-Voltage SRAMs

Operation at low voltages is critical to achieving maximum energy efficiency, but on-chip SRAM typically limits the minimum operating voltage of the entire system as the small transistors in SRAM bitcells are especially vulnerable to process variation. All SRAM arrays in the core voltage domain use the same custom 4 KB 8T-based SRAM macro

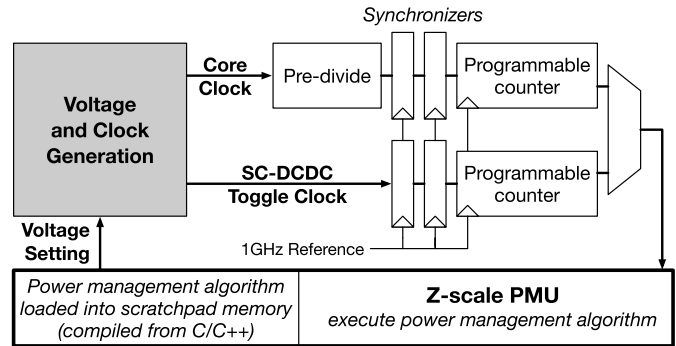


Fig. 5. Integrated power management system.




Fig. 6. Layout of the custom 8T SRAM macro.

shown in Fig. 6, which is logically organized as 512 entries of 72 bits (64 bits + 8 possible error-correcting code bits) and physically organized as two arrays of 128 rows \times 144 columns with two-to-one physical interleaving. Low-voltage operation is enabled by the 8T bitcell, where each transistor is larger than the equivalent high-density 6T bitcell, and by the fully depleted silicon-on-insulator (FD-SOI) process, which reduces the threshold voltage variation [34], [35]. While the arrays also implemented a negative-bitline write assist, the assist was not necessary to achieve minimum voltage operation.

F. Physical Design

The top-level physical design of the SoC was completed using a multiclock and multivoltage flow in a digital place-and-route tool. Fig. 7 shows the floorplan of the SoC. The design is partitioned into two voltage areas, with the core voltage area supplied by the SC-DCDC converters placed centrally. Numerous additional voltages and clocks are defined to supply both the core and the various analog and mixed-signal blocks that make up the SoC.

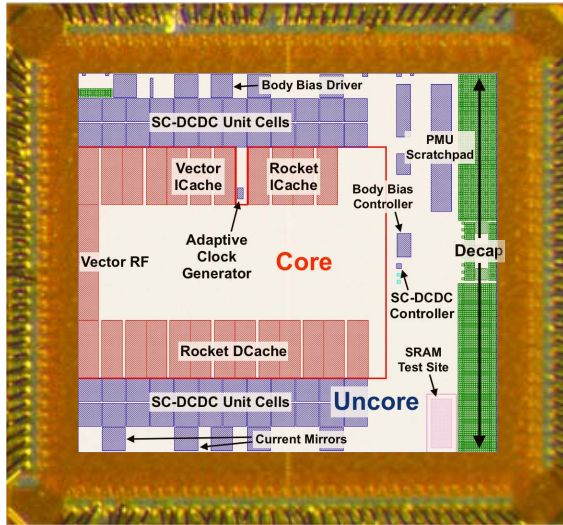


Fig. 7. Annotated die micrograph.

As described in Section II-B, reducing the insertion delay from the adaptive clock generator to the core clock sinks is critical to improving system efficiency. Accordingly, the clock generator itself was placed near the center of the core area, and a “peninsula” of the uncore voltage domain was extended to allow the routing of control signals from the block to the top level of the design hierarchy. The location of the core clock multiplexer, which allowed the selection of different core clock sources for test, was specified explicitly and placed near the center of the core area. These improvements combined to reduce core insertion delay by several hundred picoseconds.

Decoupling capacitance was added to improve the integrity of the 1.0 V and 1.8 V inputs to the SC-DCDC unit cells and help offset power delivery issues caused by the wire-bond packaging of the chip. A custom decoupling cell with MOS capacitors and a six-layer metal-oxide-metal 50 nm mesh add 539 pF of capacitance to the 1.8 V supply and 802 pF of capacitance to the 1.0 V supply.

III. EXPERIMENTAL RESULTS

A prototype system was designed and implemented [21] in 28 nm ultrathin body and BOX FD-SOI technology [36]. Fig. 7 shows the die micrograph, and Table III summarizes the key features.

A. Measurement Setup

The test setup is similar to that presented in [10]. The die is wire bonded to a small test board that includes decoupling capacitance and pinouts for measurement, including measurement pins dedicated to the generated adaptive clock and the dc-dc toggle clock. This test board connects to a Zedboard that includes an FPGA and a network-accessible ARM core. The serialized 16-bit digital interface of the test chip is connected to programmable logic in the FPGA, which can route memory traffic to local DRAM and emulate system calls. Software running on the ARM core acts as a host controller for the tethered test chip processors.

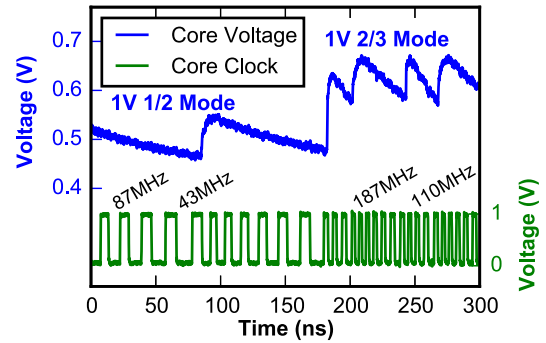


Fig. 8. Measured waveforms of the core voltage and adaptive clock during a mode transition.

TABLE III
CHIP SUMMARY

Technology	28 nm FD-SOI
Die Area	1665 $\mu\text{m} \times$ 1818 μm (3.03 mm ²)
Core Area	895 $\mu\text{m} \times$ 1193 μm (1.07 mm ²)
Standard Cells	568K
Converter Area	48 \times 90 $\mu\text{m} \times$ 90 μm (0.39 mm ²)
Core Voltage	0.48 V to 1 V (bypass mode)
Core Power	1.2 mW to 231 mW (bypass mode)
Core Frequency	20 MHz to 797 MHz (bypass mode)
Peak Energy Efficiency	41.8 GFLOPS/W (1/2 1V mode)
Conversion Efficiency	82-89%
AVS Transition Time	<1 μs
Peak AVS Energy Savings	39.8%

To bring up the system, external supplies are adjusted and the uncore brought out of reset. Next, the clocking system is initialized, including the adaptive clock generator, followed by all other IP blocks, including the SC-DCDC converters, which can be initialized to any operating mode. Finally, programs are loaded into memory and the cores are brought out of reset. The PMU begins program execution first, followed by the application core.

B. Integrated Voltage Regulation

Fig. 8 shows the core voltage and core clock during a transition between voltage modes. The core clock adapts automatically to the voltage ripples within each mode and to the large voltage transition between modes.

Table IV shows the measured system conversion efficiency of the SoC operating under each SC-DCDC mode. Because the instantaneous voltage and current consumption of the SC-DCDC output voltage cannot be easily measured, system conversion efficiency is calculated by comparing the energy cost of a long-running computation with the energy cost under a fixed supply voltage and clock that takes the same time to complete [10]. In this calculation, 100% efficiency represents a lossless regulator supplying the core as it operates at the maximum frequency achievable at that voltage. The efficiency calculation therefore accounts for both regulator inefficiency

TABLE IV
SYSTEM CONVERSION EFFICIENCY

DC-DC Mode	Efficiency (Adaptive Clock)	Efficiency (Fixed Clock)
1/2 1.8V	88.7%	76.6%
2/3 1.0V	85.0%	75.4%
1/2 1.0V	81.6%	61.9%




Fig. 9. Measurement of core performance running a tuned matrix-multiply benchmark with the SC-DCDC in bypass mode. The number in each cell is the core energy efficiency in double-precision GFLOPS/W.

and the energy cost of operating at a frequency less than the maximum possible.

The adaptive clock is tuned at each voltage setting by sweeping the settings of its replica delay path and choosing the fastest setting that still results in correct core functionality. The adaptive clocking system provides a large improvement in system conversion efficiency because the core is able to operate at a higher average frequency as the supply voltage ripples, reducing the amount of energy required to complete the same amount of work. When supplied by the SC-DCDCs, the processor achieves a peak energy efficiency of 41.8 double-precision GFLOPS/W running an FMA microbenchmark on the vector coprocessor in 1/2 1 V mode. The processor is able to boot Linux and run user programs while powered by the rippling supply voltage and adaptive clock.

C. System Energy Efficiency

Fig. 9 shows the processor functionality across wide voltage and frequency ranges. The SC-DCDC converter is placed into bypass mode for characterization, allowing the measurement of processor performance under fixed voltage and frequency. The best energy efficiency in bypass mode of 54.0 double-precision GFLOPS/W is achieved at 500 mV and 40 MHz.

Fig. 10 shows the best frequency achievable at each operating point and the total energy consumed by a fixed-duration matrix-multiply benchmark at that operating point. The application of FBB increases performance but results in higher leakage power.

D. Clock Generation

Fig. 11 compares the generated frequency of the free-running adaptive clock generator with the design from [10]. Both designs reliably generate a core clock that can supply the core area while complex software such as Linux is executed on the application core. The two designs track voltage similarly at the slowest delay settings, but the tracking varies

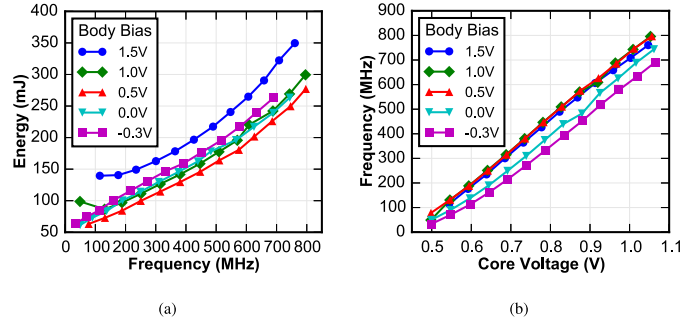


Fig. 10. Measured core frequency and energy running a matrix-multiply benchmark on the vector coprocessor. Positive body bias values denote FBB. (a) Effect of body bias on total energy. (b) Effect of body bias on the voltage-frequency curve.

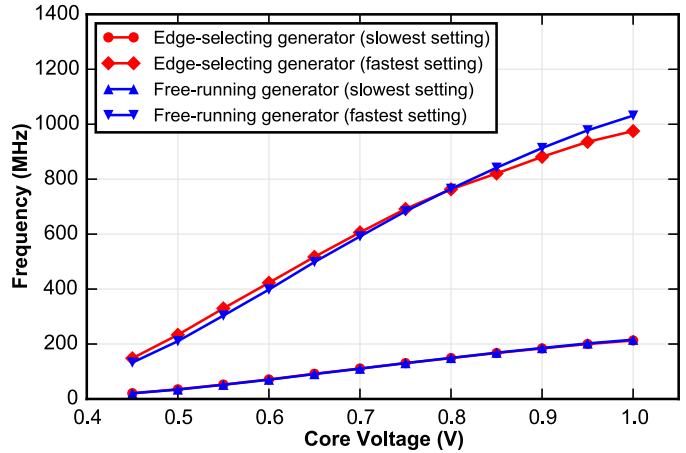


Fig. 11. Measured performance of the free-running adaptive clock generator compared with the edge-selecting clock generator from [10].

at the fastest setting because the free-running design oscillates entirely in the variable-voltage domain, while part of the timing loop in the edge-selecting design is level-shifted to the fixed 1 V supply. The free-running generator achieves the same functionality as the adaptive clocking scheme from [10], while eliminating the need for the distribution of a fixed reference and reducing design area.

Fig. 12 compares the voltage-dependent delay characteristics of the four delay banks, normalized to the delay of the custom buffer cells in bank 3. Each bank’s result was measured by recording the generated clock’s frequency after selecting the maximum delay through that bank and the minimum delay through the remaining banks. The cells with small pMOS/nMOS ratios and larger gate lengths have larger delays at lower voltages. The wide variation in voltage-dependent delays between the different delay banks (up to 18% at 0.45 V) validates the need for multiple different standard cells to achieve accurate critical path tracking.

E. Power Measurement Counter

Fig. 13 shows the measured core voltage and SC-DCDC toggle clock. Fig. 14 compares the core power measured by the bench equipment with the SC-DCDC toggle frequency measured by the integrated counter described in Section IV. The correlation is monotonic and approximately linear for each




Fig. 12. Measurement of voltage-dependent delay tracking of the delay banks in the adaptive clock generator. The frequencies are normalized to those of Delay Bank #3 to emphasize the differing voltage-dependent delay characteristics of the four banks.

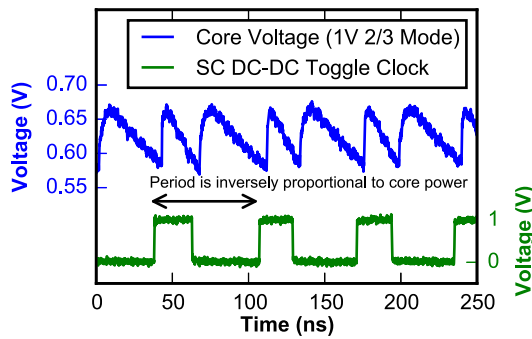


Fig. 13. Measured waveforms of the core voltage and SC-DCDC toggle clock.




Fig. 14. Measured correlation of core power with the SC-DCDC switching rate.

of the three conversion modes, confirming the practicality of using this toggle frequency to estimate core power and system load.

IV. POWER MANAGEMENT

The programmable PMU allows the implementation of a wide variety of power management algorithms to improve energy efficiency. Several different experiments demonstrate the flexibility of the system in implementing common energy-saving techniques.

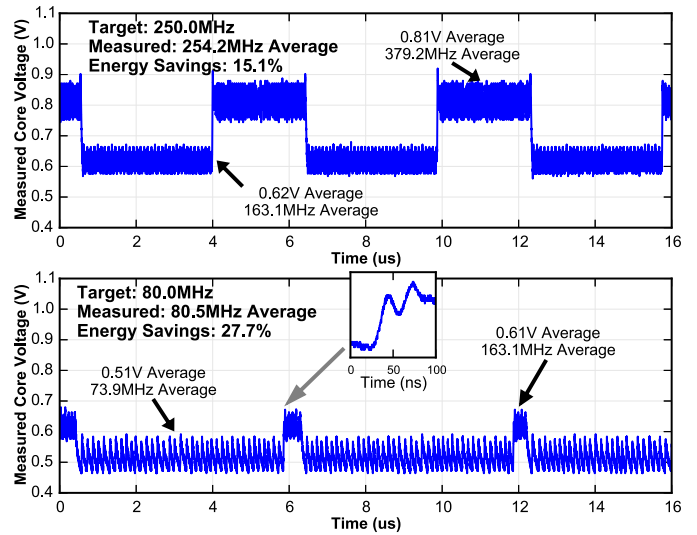


Fig. 15. Core voltage measured as the PMU runs a frequency-locking algorithm.

A. Frequency Locking

Operating systems or users may request a particular operating frequency target to match a particular performance demand. However, systems with SC-DCDC regulators operate most efficiently only at the maximum frequency of each discrete operating voltage mode. Nonetheless, the prototype system is able to lock to a particular frequency target via hopping, in which the PMU automatically adjusts the voltage and frequency of the system to achieve an *average* frequency that is indistinguishable from the desired frequency at longer timescales. Despite the small number of discrete voltage settings, an arbitrary effective clock frequency can be achieved by rapidly switching between two voltage configurations. To implement this algorithm, the PMU first calibrates the system by using the core clock counter to measure the average operating frequency at each voltage mode. Then a target frequency is provided to the PMU, which polls the core clock counter and dithers the voltage setting to achieve the target frequency in aggregate.

The results of the experiment are shown in Fig. 15. Without dithering, the processor would need to operate only in the higher mode to guarantee that the performance target is met, which would incur up to 40% more energy than the dithered approach. The choice of hopping frequency presents a tradeoff between increased fidelity to the target effective frequency and the more frequent occurrence of transition overheads, which can increase energy consumption. In this paper, the energy cost associated with transitions between voltage modes is small because the processor continues to operate as the clock frequency adjusts during the mode transition. In high-to-low mode transitions, no charge is wasted, but in some low-to-high transitions, the flying capacitance is charged to 1 V, consuming wasted energy

$$E_{loss} = \frac{1}{2} C_{fly} \Delta V^2. \quad (1)$$

In the worst-case transition from 1 V 2/3 mode to 1.8 V 1/2 mode, the flying capacitance is charged from roughly




Fig. 16. Measured energy consumption of the system while the PMU executes a voltage dithering algorithm. The blue curve was measured with the SC-DCDC converters in bypass mode, showing the energy consumption of the system with no voltage conversion losses. The dithering program allows the core to operate at a wide range of effective operating points to trade off performance with energy efficiency.

0.33 V, resulting in an E_{loss} of roughly 1 nJ. At a typical core operating power of 50 mW, this loss is equivalent to the energy consumed by just 20 ns of normal operation. Accordingly, a hopping frequency of approximately 6 μ s was chosen. This frequency is still quite fast, but is slow enough that transition energy costs can be neglected.

B. Voltage Dithering

Many voltage generation schemes can efficiently generate only a few discrete output modes [37]. Nonetheless, a wide range of effective operating voltages can be achieved through voltage dithering, in which the voltage is rapidly switched between two operating modes to achieve an average voltage between them [38], [39]. Theoretical analysis has shown that voltage dithering between just a few modes can achieve near-optimal efficiency [40]. Fig. 16 shows the measured results of a PMU program for voltage dithering. The program rapidly switches between neighboring voltages at a ratio specified at runtime. Each measured point represents the completion of a fixed-length matrix-multiply benchmark. The dithering algorithm enables a wide operating range for the core, bounded only by the lowest and highest voltage settings of the SC-DCDC converter.

Fig. 17 shows the conversion efficiency of the system at each operating point as calculated by the method described in [10]. Two different dithering programs were run on the PMU. The first program simply switches the voltage mode setting of the SC-DCDC converters, without changing any other system settings; the delay settings of the replica paths in the adaptive clock generator were tuned to the best setting that could function across both operating modes. The results are shown by the red points in Fig. 17. Because the best setting of the replica paths changes according to operating mode, the conversion efficiencies of this approach are less than optimal for part of the dithering range. The second program switches both the voltage mode and the delay settings of the replica paths according to a pre-characterization of the best adaptive clock setting associated with each voltage mode. This program is able to speed the generated clock at the higher

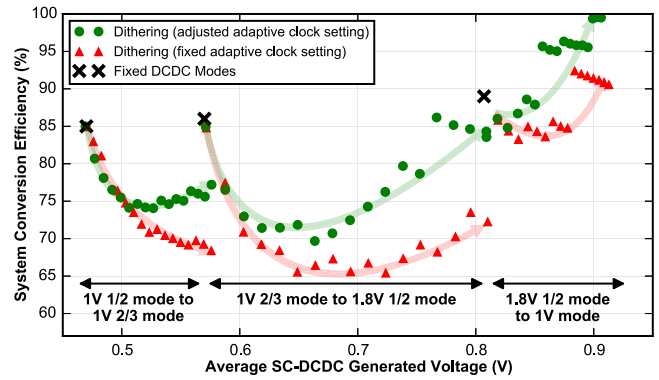


Fig. 17. Measured conversion efficiency of the SC-DCDC regulators and adaptive clocking system while the PMU runs the dithering program. The system conversion efficiency is measured by calculating the energy overhead of completing the same fixed workload in the same amount of time as the system in bypass mode. The red points show the measured results from a PMU program in which the delay settings of the adaptive clock are not adjusted upon voltage mode transitions, while the green points show the results of an improved PMU program in which the adaptive clock settings are adjusted after each voltage transition.

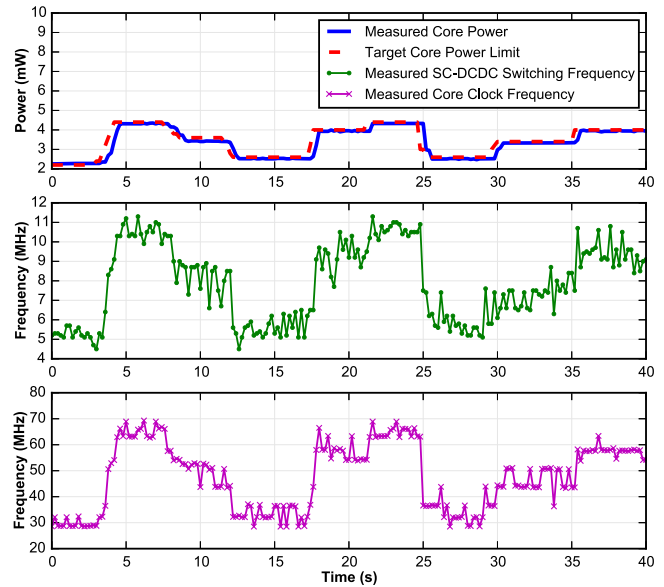


Fig. 18. Measured results as the PMU adaptively scales frequency within a user-specified power budget.

voltage settings, leading to higher conversion efficiencies. In all, the efficiencies of the second program range from 70% to 100%, depending on the voltage mode and dithering ratio. Conversion efficiencies while dithering are not as high as the efficiencies of the fixed operating modes in Table IV because the external voltage reference used by the comparator cannot be tuned for a particular SC-DCDC mode.

C. Power Envelope Tracking

A system-level absolute power limit is a common power management constraint. Fig. 18 shows the results of a power management algorithm executed on the PMU that maximizes core performance within a user-specified power budget. The




Fig. 19. Measured core voltage as the PMU executes AVS in response to variable workload.

power management program polls an externally writable control register that stores the absolute power limit for the program. The PMU core then monitors the SC-DCDC toggle counters to continuously estimate core power using the quadratic model described in [23] with pre-characterized coefficients. If the estimated core power is above the specified limit, core frequency is decreased, and if it is below the limit, core frequency is increased. In this way, the best possible performance is automatically obtained while the user-specified power budget is respected.

D. Fine-Grained Adaptive Voltage Scaling

The PMU can also use the integrated counters to coordinate fine-grained adaptive voltage scaling (AVS) on-chip without any explicit guidance from the programs running on the processor. Extremely fine-grained AVS algorithms can exploit much shorter periods of inactivity than software-based AVS algorithms. In the proposed algorithm, core power is used as a marker of program phase. When core power is higher, the core is likely to be executing a compute-intensive program region, and a high voltage is necessary to maximize performance. When the core power is lower, the core is likely to be waiting idle for off-chip communication in a memory-bound program region, and energy can be saved with minimal performance impact by reducing the voltage. For this experiment, the core runs a synthetic benchmark that alternates between the compute-intensive and idle phases at a timescale of tens of microseconds.

Fig. 19 shows the core voltage measured during the execution of the benchmark and the AVS power-management algorithm. The algorithm switches the core voltage between the 1.8 V 1/2 mode and the 1 V 2/3 mode, actuated by core power estimates determined by continuously polling the SC-DCDC toggle counter. When the core voltage is high and the toggle rate drops below a threshold, this corresponds to an idle program period, so the PMU reduces the core voltage to save energy. When the core voltage is low and the toggle

rate exceeds a threshold, the workload has increased and the PMU increases the core voltage. The system is able to detect changes in workload in less than 1 μ s and adjust the core voltage in response. Without integrated voltage regulators and power management, the system would not be able to respond within the timescales of the workload variation. The results of the power-management algorithm are therefore compared against continuous operation in the higher voltage mode, which would otherwise be required to meet the same performance target. The power-management algorithm reduces the energy consumed by 39.8%, and the fast response incurs negligible ($<0.2\%$) performance penalty compared with this baseline, demonstrating the efficacy of fine-grained AVS at improving energy efficiency with fast workload tracking.

V. CONCLUSION

This processor SoC couples an energy-efficient RISC-V core and vector accelerator with a power-management processor, integrated voltage regulators, and an adaptive clock generator, allowing for improvements in system energy efficiency through the use of power-management algorithms running in microsecond-scale feedback loops entirely on-die. The four discrete output voltages of the integrated SC-DCDC regulator can be dithered to produce a wide continuous range of effective core voltages and frequencies. An AVS algorithm can use integrated counters to track changes in program phase and respond at submicrosecond timescales, allowing for substantial energy savings with minimal performance penalty. Taken together, these technologies represent a powerful tool for the energy-efficient design of future mobile SoCs.

ACKNOWLEDGMENT

The authors would like to thank A. Cathelin, P. Flatresse, S. Clerc, F. Abouzeid, P. Roche, J.-L. Autran, J. Bachrach, G. Choi, B. Richards, J. Dunn, all the students, faculty, staff, and sponsors of the Berkeley Wireless Research Center and the ASPIRE Lab for their contributions, and STMicroelectronics for the fabrication donation.

REFERENCES

- [1] W. Kim, M. S. Gupta, G.-Y. Wei, and D. Brooks, "System level analysis of fast, per-core DVFS using on-chip switching regulators," in *Proc. IEEE 14th Int. Symp. High Perform. Comput. Archit.*, Salt Lake City, UT, USA, Feb. 2008, pp. 123–134.
- [2] K. Rangan, G. Wei, and D. Brooks, "Thread motion: Fine-grained power management for multi-core systems," in *Proc. Int. Symp. Comput. Archit.*, Jun. 2009, pp. 302–313.
- [3] S. Eyerman and L. Eeckhout, "Fine-grained DVFS using on-chip regulators," *ACM Trans. Archit. Code Optim. (TACO)*, vol. 8, no. 1, pp. 1–24, Apr. 2011.
- [4] A. Shye, B. Scholbrock, and G. Memik, "Into the wild: Studying real user activity patterns to guide power optimizations for mobile architectures," in *Proc. IEEE/ACM Int. Symp. Microarchitecture*, Dec. 2009, pp. 168–178.
- [5] 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 Modeling Power Electron.*, Jun. 2010, pp. 1–7.
- [6] 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. 167–168.
- [7] S. Clerc *et al.*, "8.4 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 *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2015, pp. 150–151.

- [8] 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 *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2015, pp. 154–155.
- [9] M. Turnquist *et al.*, "Fully integrated DC-DC converter and a 0.4 V 32-bit CPU with timing-error prevention supplied from a prototype 1.55V Li-ion battery," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2015, pp. 320–321.
- [10] B. Zimmer *et al.*, "A RISC-V vector processor with simultaneous-switching switched-capacitor DC-DC converters in 28 nm FDSOI," *IEEE J. Solid-State Circuits*, vol. 51, no. 4, pp. 930–942, Apr. 2016.
- [11] N. Kurd *et al.*, "Haswell: A Family of IA 22 nm processors," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2014, pp. 112–113.
- [12] A. Naveh, D. Rajwan, A. Ananthkrishnan, and E. Weissmann, "Power management architecture of the 2nd generation Intel core microarchitecture, formerly codenamed Sandy Bridge," in *Proc. IEEE Hot Chips Symp.*, Aug. 2011, pp. 1–33.
- [13] D. N. 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.
- [14] D. Marculescu, "On the use of microarchitecture-driven dynamic voltage scaling," in *Proc. Workshop Complexity-Effective Design*, Jun. 2000, pp. 1–6.
- [15] H. Li, C. Y. Cher, K. Roy, and T. N. Vijaykumar, "Combined circuit and architectural level variable supply-voltage scaling for low power," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 13, no. 5, pp. 564–575, May 2005.
- [16] P. Juang, Q. Wu, L.-S. Peh, M. Martonosi, and D. W. Clark, "Coordinated, distributed, formal energy management of chip multiprocessors," in *Proc. Int. Symp. Low Power Electron. Design (ISLPED)*, Aug. 2005, pp. 127–130.
- [17] E. Beigne *et al.*, "An asynchronous power aware and adaptive NoC based circuit," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1167–1177, Apr. 2009.
- [18] V. Gupta and M. Singh, "Energy conservation in asynchronous systems using self-adaptive fine-grain voltage scaling," in *Proc. Int. Green Comput. Conf.*, Jun. 2013, pp. 1–8.
- [19] W. Godycki, C. Torng, I. Bukreyev, A. Apsel, and C. Batten, "Enabling realistic fine-grain voltage scaling with reconfigurable power distribution networks," in *Proc. 47th Annu. IEEE/ACM Int. Symp. Microarchitecture (MICRO)*, Dec. 2014, pp. 381–393.
- [20] *The RISC-V ISA*, accessed on Nov. 14, 2016. [Online]. Available: <http://riscv.org>
- [21] B. Keller *et al.*, "Sub-microsecond adaptive voltage scaling in a 28 nm FD-SOI processor SoC," in *Proc. IEEE Eur. Solid-State Circuits Conf.*, Sep. 2016, pp. 269–272.
- [22] M. Blagojević, M. Cochet, B. Keller, P. Flatresse, A. Vladimirescu, and B. Nikolić, "A fast, flexible, positive and negative adaptive body-bias generator in 28 nm FDSOI," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2016, pp. 61–62.
- [23] M. Cochet *et al.*, "On-chip supply power measurement and waveform reconstruction in a 28 nm FD-SOI processor SoC," in *Proc. IEEE Asian Solid-State Circuits Conf.*, Nov. 2016, pp. 125–128.
- [24] K. Asanović *et al.*, "The rocket chip generator," Dept. Elect. Eng. Comput. Sci., Univ. California, Berkeley, CA, USA, Tech. Rep. UCB/ECS-2016-17, Apr. 2016.
- [25] 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.
- [26] Y. Lee, "Decoupled vector-fetch architecture with a scalarizing compiler," Ph.D. dissertation, Dept. Elect. Eng. Comput. Sci., Univ. California, Berkeley, CA, USA, May 2016.
- [27] *IEEE Standard for Floating-Point Arithmetic*, IEEE Std. 754, 2008.
- [28] J. Kwak and B. Nikolić, "A self-adjustable clock generator with wide dynamic range in 28 nm FDSOI," *IEEE J. Solid-State Circuits*, vol. 51, no. 10, pp. 2368–2379, Oct. 2016.
- [29] D. Jiao, J. Gu, and C. H. Kim, "Circuit design and modeling techniques for enhancing the clock-data compensation effect under resonant supply noise," *IEEE J. Solid-State Circuits*, vol. 45, no. 10, pp. 2130–2141, Oct. 2010.
- [30] D. C. Snowdon, S. M. Petters, and G. Heiser, "Accurate on-line prediction of processor and memory energy usage under voltage scaling," in *Proc. IEEE Int. Conf. Embedded Softw.*, Sep. 2007, pp. 84–93.
- [31] R. McGowen *et al.*, "Power and temperature control on a 90-nm Itanium family processor," *IEEE J. Solid-State Circuits*, vol. 41, no. 1, pp. 229–237, Jan. 2006.
- [32] Y. Sinangil *et al.*, "A self-aware processor SoC using energy monitors integrated into power converters for self-adaptation," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2014, pp. 139–140.
- [33] Y. Lee, "Z-scale: Tiny 32-bit RISC-V systems," presented at the OpenRISC Conf., Geneva, Switzerland, Oct. 2015.
- [34] N. Planes *et al.*, "28 nm FDSOI technology platform for high-speed low-voltage digital applications," in *Proc. Symp. VLSI Technol. (VLSIT)*, Jun. 2012, pp. 133–134.
- [35] B. Zimmer, O. Thomas, S. O. Toh, T. Vincent, K. Asanović, and B. Nikolić, "Joint impact of random variations and RTN on dynamic writeability in 28 nm bulk and FDSOI SRAM," in *Proc. IEEE Eur. Solid-State Device Res. Conf.*, Sep. 2014, pp. 98–101.
- [36] P. Flatresse *et al.*, "Ultra-wide body-bias range LDPC decoder in 28 nm UTBB FDSOI technology," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers (ISSCC)*, Feb. 2013, pp. 424–425.
- [37] 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.
- [38] S. Lee and T. Sakurai, "Run-time voltage hopping for low-power real-time systems," in *Proc. Design Autom. Conf.*, Jun. 2000, pp. 806–809.
- [39] B. H. Calhoun and A. P. Chandrakasan, "Ultra-dynamic voltage scaling (UDVS) using sub-threshold operation and local voltage dithering," *IEEE J. Solid-State Circuits*, vol. 41, no. 1, pp. 238–245, Jan. 2006.
- [40] V. Gutnik and A. P. Chandrakasan, "Embedded power supply for low-power DSP," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 5, no. 4, pp. 425–435, Dec. 1997.
- [41] Z. Toprak-Deniz *et al.*, "Distributed system of digitally controlled microregulators enabling per-core DVFS for the POWER8 microprocessor," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2014, pp. 98–99.
- [42] V. Krishnaswamy *et al.*, "Fine-grained adaptive power management of the SPARC M7 processor," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2015, pp. 74–76.
- [43] D. Bol *et al.*, "SleepWalker: A 25-MHz 0.4-V sub-mm² 7- μ W/MHz microcontroller in 65-nm LP/GP CMOS for low-carbon wireless sensor nodes," *IEEE J. Solid-State Circuits*, vol. 48, no. 1, pp. 20–32, Jan. 2013.



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

He has held research internship positions at NVIDIA, Santa Clara, CA, USA and the National Institute of Standards and Technology, Gaithersburg, MD, USA. His current research interests include energy-efficient microprocessor design, fine-grained adaptive voltage scaling, and innovative digital hardware design paradigms.



Martin Cochet (S'14) was born in Saint-Etienne, France, in 1991. He received the Dipl.Ing. degree from École Polytechnique, Palaiseau, France, in 2012, the M.S. degree from the Imperial College, London, U.K., in 2013, and the Ph.D. degree in electronic engineering from Aix-Marseille University, Marseille, France, and STMicroelectronics, Crolles, France, in 2016.

In 2009, he was an Officer Cadet with the French Air Force, Avord, France. In 2012, he was a Research Intern with the University of California, Santa Barbara, CA, USA. In 2015, he was a Visiting Scholar at the University of California, Berkeley, CA, USA. His current research interests include circuit design for digital system-on-chip power management for energy efficiency and high speed communications.



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 now with the Circuits Research Group, NVIDIA Corporation, Santa Clara, CA, USA. His current research interests include soft error resilience and energy-efficient digital design, with an emphasis on low-voltage SRAM design and variation tolerance.



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

From 2006 to 2009, he was a Staff Researcher with GCT Research, Seoul, South Korea, designing the digital controller of wireless communication systems. He was an Intern with Intel Corporation, Hillsboro, OR, USA, in 2011, and Advanced Micro


Devices, Inc., Sunnyvale, CA, USA, in 2013. His current research interests include energy-efficient microprocessor design, self-adjustable clock generators, and advanced synchronization circuits.



Alberto Puggelli (S'09–M'17) received the B.Sc. degree from the Politecnico di Milano, Milan, Italy, in 2006, two M.Sc. degrees in electrical engineering (*summa cum laude*) from the Politecnico di Torino, Turin, Italy, in 2008, and 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 an Intern Analog Designer with ST-Ericsson, Plan-les-Ouates, Switzerland, in 2009, and with Texas Instruments Inc., Dallas, TX, USA, from 2011 to 2012. He is currently the Director of Technology with Lion Semiconductor Inc., San Francisco, CA, USA. He has authored or co-authored over 20 publications in IEEE/ACM conference proceedings and journals. He holds eight U.S. patents. His current research interests include the design of hybrid dc-dc voltage regulators.


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



Yunsup Lee (S'09) received the B.S. degree in computer science and electrical engineering from the Korea Advanced Institute of Science and Technology, Daejeon, South Korea, in 2005, the M.S. degree in computer science from the University of California, Berkeley, CA, USA, in 2011, and the Ph.D. degree in computer science from the University of California, Berkeley, CA, USA, in 2016.

He started the RISC-V instruction set architecture project with A. Waterman, D. Patterson, and K. Asanović in 2010, and co-founded SiFive, Inc.,

San Francisco, CA, USA, where he serves as the Chief Technology Officer. His current research interests include the design of energy-efficient computing systems.



Milovan Blagojević (M'12) was born in Kraljevo, Serbia, in 1987. He received the B.S. and M.S. degrees in electrical engineering and computer science from the School of Electrical Engineering, University of Belgrade, Belgrade, Serbia, in 2010 and 2012, respectively, and he defended the Ph.D. thesis from the Institute Supérieur d'Electronique de Paris, Paris, France, in 2015.


He was an Intern with the HW/SW Co-Design Group, Institute for Informatics, Erlangen, Germany, in 2010, and with Intel, Belgrade, in 2011, where he was involved in imaging algorithms. From 2012 to 2015, he was a Visiting Researcher at STMicroelectronics and the University of California, Berkeley, CA, USA, where he was involved in FD-SOI technology and energy-efficiency optimization in modern digital System-on-Chips (SoC). He is currently an SoC Architect with the iCDG Department, Intel Deutschland GmbH, Munich, Germany, where he is working on a next-generation baseband modem.



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 from the University of California, Berkeley, CA, USA, in 2014, where he is currently pursuing the Ph.D. degree.

He was a Summer Research Intern with the Jet Propulsion Laboratory, Pasadena, CA, USA, in 2014, and NVIDIA Corporation, Santa Clara, CA, USA, in 2015. His current research interests include

digital integrated circuit design methodologies, digital signal processing and algorithms, and machine learning.



Pi-Feng Chiu (S'10) received the B.S. and 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 with 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, and Western Digital, Irvine,

CA, USA, in 2014 and 2016, respectively. Her current research interests include resilient memory circuit design and emerging nonvolatile memories.



Palmer Dabbelt received the B.S. degree in engineering from The University of Illinois, Urbana, CA, USA, in 2011. He is currently pursuing the Ph.D. degree with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, USA.

His current research interests include energy-efficient microprocessor design.



Colin Schmidt was born in Olney, MD, USA, in 1991. He received the B.S. degree in electrical and computer engineering and computer science from Cornell University, Ithaca, NY, USA, in 2013. He is currently pursuing the Ph.D. degree with the University of California, Berkeley, CA, USA.

His current research interests include data-parallel architectures and their compilers.



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

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

Prof. Asanović is an ACM Distinguished Scientist.



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.

He was with Lion Semiconductor, Wilocity, Cadence, Xilinx, Oracle, Intel, AMD, Rambus, Hewlett Packard, and IBM Research, where he was involved in digital, analog, and mixed-signal integrated circuits for computing, test and measurement, and high-speed communications. He has been an Associate Professor of Electrical Engineering and

Computer Sciences and the Co-Director of the Berkeley Wireless Research Center, University of California, Berkeley, CA, USA, since 2007. His current research interests include energy-efficient integrated systems, including the circuit, device, communications, and optimization techniques used to design them.

Dr. Alon was a recipient of the IBM Faculty Award in 2008, the 2009 Hellman Family Faculty Fund Award, and the 2010 UC Berkeley Electrical Engineering Outstanding Teaching Award. He has co-authored papers that received the 2010 ISSCC Jack Raper Award for Outstanding Technology Directions Paper, the 2011 Symposium on VLSI Circuits Best Student Paper Award, and the 2012 and 2013 Custom Integrated Circuits Conference Best Student Paper Award.



Borivoje Nikolić (S'93–M'99–SM'05–F'17) received the Dipl.Ing. and M.Sc. degrees in electrical engineering from the University of Belgrade, Beograd, Serbia, in 1992 and 1994, respectively, and the Ph.D. degree 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 co-authored the book

Digital Integrated Circuits: A Design Perspective (2nd ed., Prentice-Hall, 2003). His current research interests include digital, analog and RF integrated circuit design, and very-large-scale integration (VLSI) implementation of communications and signal processing systems.

Dr. Nikolić was a recipient of the NSF CAREER Award in 2003, the College of Engineering Best Doctoral Dissertation Prize, and the Anil K. Jain Prize for the Best Doctoral Dissertation in Electrical and Computer Engineering at University of the California at Davis in 1999, and the City of Belgrade Award for the Best Diploma Thesis in 1992. He was also a recipient of the Best Paper Awards at the IEEE International Solid-State Circuits Conference, the Symposium on VLSI Circuits, the IEEE International SOI Conference, the European Solid-State Device Research Conference, the European Solid-State Circuits Conference, the S3S Conference, and the ACM/IEEE International Symposium of Low-Power Electronics. From 2014 to 2015, he was a Distinguished Lecturer of the IEEE Solid-State Circuits Society.