

A Sub- μ W Energy-Performance-Aware IoT SoC With a Triple-Mode Power Management Unit for System Performance Scaling, Fast DVFS, and Energy Minimization

Xinjian Liu¹, Graduate Student Member, IEEE, Sumanth Kamineni¹, Member, IEEE, Jacob Breiholz¹, Benton H. Calhoun¹, Fellow, IEEE, and Shuo Li¹, Member, IEEE

Abstract—This article presents an ultra-low-power (ULP) Internet-of-Things (IoT) system-on-chip (SoC) using a triple-mode power management unit (PMU) to achieve self-adaptive power-performance scaling and energy-minimized operation. The proposed PMU comprises three modes: energy-aware (EA) mode, performance-aware (PA) mode, and minimum energy point (MEP) tracking mode. By controlling a microprocessor with the three modes, the SoC can adaptively scale its frequency and supply voltage based on either the input energy availability or the task priority. To achieve robust and rapid mode transitions, the SoC adopts fast dynamic voltage and frequency scaling (DVFS) and fast load transient response (FLTR) through asynchronous control. For energy-minimized operation, a sub-nW constant-energy-cycle (CEC) algorithm keeps the microprocessor operating at the MEP with a 0.026-mm² area overhead. In addition, the on-chip integration of a bias generator (BG), clock (CLK), and power-on-reset block empowers the SoC to be a fully self-contained system. Fabricated in 65-nm CMOS, measurement results show that the SoC has a minimum power consumption of 194.3 nW at 180 Hz. The proposed PMU achieves 5.2-nW quiescent power and 92.6% peak efficiency while maintaining >80% efficiency from 190 nW to 3 mW. The MEP tracking (MEPT) circuits achieve <2.3% energy per cycle error and <18 mV voltage tracking error. The measured quiescent power of the MEPT circuits in the idle mode is 379 pW, which only accounts for 0.19% of the total system power. Measurements of the triple-mode transitions show that this SoC is well suited for resource-constrained IoT applications.

Index Terms—Buck converter, energy aware, fast dynamic voltage and frequency scaling (DVFS), high efficiency, Internet of Things (IoT), minimum energy point (MEP) tracking, performance aware, performance scaling, power management unit

Manuscript received 10 July 2023; revised 7 November 2023; accepted 27 December 2023. This article was approved by Associate Editor Taekwang Jang. This work was supported in part by the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy (EERE) under Award DE-EE0008225 and in part by the NSF Nanosystems Engineering Research Center (NERC) Advanced Self-Powered Systems of Integrated Sensors and Technologies (ASSIST) Center under Grant EEC-1160483. This article was presented in part at the IEEE International Solid-State Circuits Conference, February 2022 [DOI: 10.1109/ISSCC42614.2022.9731758]. (Corresponding author: Shuo Li.)

Xinjian Liu, Sumanth Kamineni, Jacob Breiholz, and Benton H. Calhoun are with the Department of Electrical and Computer Engineering, University of Virginia, Charlottesville, VA 22903 USA (e-mail: x15sp@virginia.edu).

Shuo Li is with the Department of Electrical Engineering, Yale University, New Haven, CT 06511 USA (e-mail: shuol.li@yale.edu).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/JSSC.2024.3350449>.

Digital Object Identifier 10.1109/JSSC.2024.3350449

(PMU), sub-nW quiescent power, system-on-chip (SoC), wide dynamic range.

I. INTRODUCTION

MINIATURIZATION of modern Internet-of-Things (IoT) devices mandates aggressive power reduction of systems-on-chip (SoCs) to extend the lifespan of devices that rely on limited energy sources. However, for many real-world applications, data sensing, processing, and wireless transferring necessitate fast response to meet the specific performance constraints, especially for multi-task applications, as shown in Fig. 1. This results in a huge challenge for circuit design that needs to tradeoff between energy consumption and the required performance to truly enable (EN) ultra-low-power (ULP) applications, such as wearable electronics, implantable health care, smart home, and security.

Thus far, a wide variety of ULP SoCs and techniques have been reported to optimize the energy and performance tradeoffs. Duty cycling is commonly used to save average power [1], [2], [3], [4]. However, for multi-task applications, frequent turning on/off the loads may lead to degraded performance due to the settling time and extra energy cost during the start-up phase. Another technique that has been widely used is dynamic voltage and frequency scaling (DVFS) [5], [6], [7], [8], [9], [10]. By scaling the operating frequency and supply voltage of a digital component (e.g., a microprocessor) with fast DVFS tracking techniques [11], [12], [13], [14], the power consumption and performance can be balanced. However, as the power of the circuits keeps scaling down with the supply voltage into the subthreshold region (supply voltage less than threshold voltage), the DVFS technique is sub-optimal from an energy efficiency point of view. In deep subthreshold, the energy consumed per cycle or operation (E_{PC}/E_{PO}) no longer decreases with the supply voltage but instead rises due to the leakage energy that integrates over a longer operation cycle. This opposite trend generates the minimum energy point (MEP). Therefore, keeping digital circuits operating at the MEP is necessary to maximize the amount of work that can be completed on a fixed energy budget [16], [17], [18].

Therefore, for ULP IoT applications, SoCs need to have ultra-low quiescent power, a highly efficient power management unit (PMU) for a longer system lifetime, performance

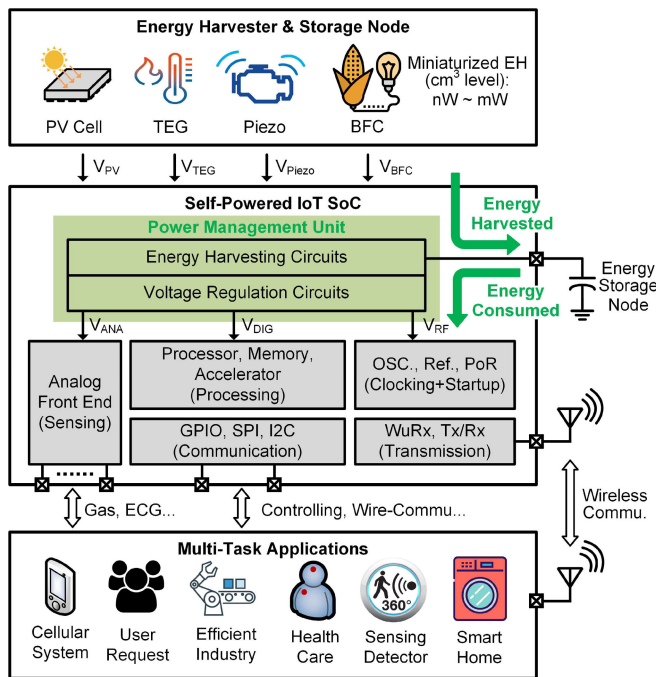


Fig. 1. Typical top-level block diagram of self-powered IoT SoCs for multi-task applications.

scaling based on both the input energy and task priority for an energy and performance tradeoff, fast DVFS for rapid mode transitions, and energy minimization to ensure energy-optimized operation. However, most prior arts that use DVFS to determine the system operating point are only based on the performance requirements [5], [7], [9], input energy conditions [10], or manual/programmable control [6], [8], [19]. None of the previous SoCs and PMUs have fully coordinated the input energy condition, energy efficiency of the loads, and performance requirements simultaneously at a nanowatt scale of power consumption. Also, for MEP tracking (MEPT), none of the previous work achieved sub-microwatt power and high energy delivery efficiency (>90%) [16], [17], [18].

Therefore, to address these challenges, we present a ULP SoC with energy and performance managed by a triple-mode PMU that integrates energy–performance-aware, performance scaling, event-driven fast DVFS, fast load transient response (FLTR), and sub-nanowatt MEPT [19]. The highlighted contributions of our work over the prior art are listed as follows.

- 1) The SoC integrates both PMU and microprocessor in the control loop with three operation modes, including energy-aware (EA) mode, performance-aware (PA) mode, and MEPT mode. The proposed triple mode together with fast DVFS coordinates the power, energy, and performance tradeoffs at nanowatt-scale system power consumption.
- 2) A constant-energy-cycle (CEC) MEPT algorithm and circuits are proposed, achieving a low tracking error of <18 mV (E_{PC} tracking error of <2.3%). Furthermore, compared to previous methods, the proposed algorithm results in a significant power overhead reduction of over two orders of magnitude.
- 3) Compared with prior art, the PMU achieves the highest peak efficiency, 5.2-nW quiescent power, and

>10⁴ dynamic power range with >80% efficiency (>100× larger than prior art). The SoC also achieves the lowest power of 194.3 nW at 180 Hz with all functions, including clock (CLK), bias generator (BG), and power-on-reset for a fully deployable IoT solution.

In this article, we expand the work from [20] to further illustrate the design consideration, tradeoffs, and circuit implementations for energy-constrained multi-task IoT applications. The rest of this article is organized as follows. Section II describes the system architecture, system behaviors, and mode control (MC). Section III explains the proposed MEPT algorithm with its design considerations and accuracy analysis. Section IV demonstrates the circuit implementations followed by the measurement results and comparison with the state-of-the-art SoCs in Section V. Finally, Section VI concludes this article.

II. SYSTEM ARCHITECTURE AND MC

A. Triple-Mode System-Level Power Management

To prolong the system lifetime, it is important for the power consumption of the SoC to align with the available input power and avoid depleting the energy storage node. Therefore, an EA mode is required to detect the input energy condition and adaptively scales the power of the load components based on the input energy availability. Usually, the output of the PMU drives the components on the SoC to manage the power. Therefore, V_{OUT_PMU} needs to scale with the input voltage level ($V_{STORAGE}$) in the EA mode, as shown in Fig. 2. On the other hand, for multi-task applications, to meet the performance requirements, a PA mode is required to scale the SoC speed and power in response to event priority. Once the energy availability drops under a certain threshold, the SoC needs to work in subthreshold to reduce power consumption and utilize the remaining energy in an energy-efficient manner for a maximized system lifetime. This necessitates an MEPT mode to track the MEP in the subthreshold region. Fig. 2 shows how the proposed SoC switches among the three modes to adaptively scale its power and performance based on both the input conditions and event priority levels.

B. Analysis of DVFS and MEPT Technologies

As discussed in Section I, for multi-task applications, duty cycling may not be beneficial due to the additional overhead in power and speed caused by frequent on-and-off operations. In contrast, the DVFS technique provides more flexible adaptivity, which is suitable for always-on operations. For multi-task IoT applications, our design uses both DVFS and MEPT methods. When V_{DD} is much larger than the threshold voltage, V_{th} , digital loads operate at a high frequency where the dynamic power dominates and the leakage power is negligible. The total power consumption can then be calculated by the following equation:

$$\begin{aligned}
 P_{Total} &\approx P_{DYN} = C_{EFF} V_{DD}^2 f_{REQ} \\
 &= C_{EFF} V_{DD}^2 \frac{I_{DS,SAT}}{K L_{DP} C_{OUT} V_{DD}} \propto V_{DD} (V_{DD} - V_{th})^2 \quad (1)
 \end{aligned}$$

where C_{EFF} , K , L_{DP} , C_{OUT} , and $I_{DS,SAT}$ stand for the average effective switched capacitance of the entire circuit, a delay




Fig. 2. Timing waveforms for the triple-mode transition based on input energy conditions and performance requirements.

fitting parameter, depth of the circuit, the output capacitance of a characteristic inverter, and the saturation current of a CMOS transistor, respectively. Based on (1), by scaling down the voltage and frequency of the digital loads, the power consumption can be reduced significantly.

Furthermore, fast DVFS is desired to allow the SoC to quickly move to the new operating point for faster system responses. DVFS can also effectively improve the energy efficiency, which can be quantified as energy per cycle, E_{PC} , as represented by the following equation:

$$E_{PC} = E_{DYN} = C_{EFF} V_{DD}^2. \quad (2)$$

Based on (2), the energy needed for each operation cycle reduces as the supply voltage decreases. However, this conclusion no longer holds when the supply voltage approaches the subthreshold region. In the subthreshold region, the energy consumption per cycle from leakage, E_{LEAK} , increases exponentially as V_{DD} decreases, which generates an MEP, as shown in Fig. 3 and the following equations [15]:

$$E_{LEAK} = W_{EFF} K C_{OUT} L_{DP} V_{DD}^2 e^{-V_{DD}/nV_{th}} \quad (3)$$

$$E_{Total} = V_{DD}^2 (C_{EFF} + W_{EFF} K C_{OUT} L_{DP} V_{DD} e^{-V_{DD}/nV_{th}}) \quad (4)$$

where W_{EFF} and n present the total effective width that contributes to leakage current and the subthreshold slope, respectively. Hence, although the DVFS technique remains effective in optimizing power consumption and energy efficiency when V_{DD} is significantly larger than V_{th} , its impact diminishes as the circuit approaches the subthreshold region. This necessitates low-power (LP) MEPT circuits to achieve optimized energy efficiency for subthreshold operation.

C. Proposed SoC Architecture and MC

This work proposes a triple-mode PMU-processor-in-loop architecture to flexibly optimize the energy and performance requirements. The architecture of the proposed IoT SoC and PMU is illustrated in Fig. 4. The SoC includes a triple-mode PMU, a microprocessor with its memory, and I/O peripherals. It operates in three modes: EA, PA, and MEPT. The priority of these modes is $EA > PA > MEPT$. The selection of the modes is automated based on both the input and load information. Depending on the chosen mode, the corresponding signal (SEL_{EA} , SEL_M , or SEL_{PA}) is selected to control the reference voltage, V_{REF} . Fig. 5 shows the flowchart of the

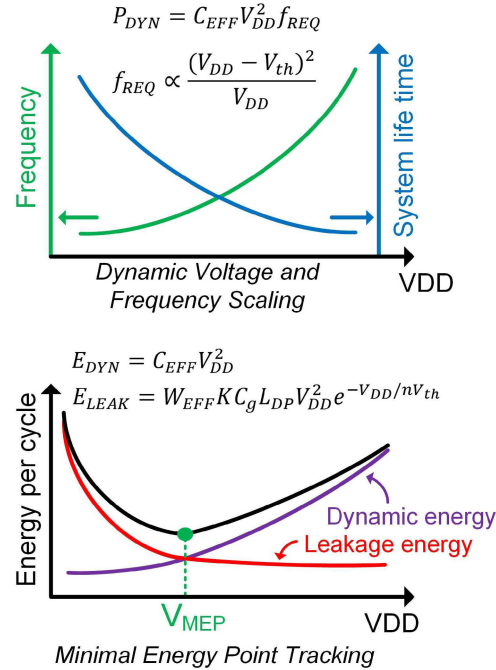


Fig. 3. DVFS and MEPT methods for balancing energy and performance.

MC algorithm. In default, the SoC works in the EA mode, where the DVFS operating point is proportional to the input voltage level. When prioritized tasks arise, the processor changes the SEL_{PA} value based on the pre-programmed lookup table (LUT). If SEL_{PA} is larger than SEL , which indicates that the performance requirement is not met, the system goes into the PA mode and sets $SEL = SEL_{PA}$, ensuring that the DVFS operating point aligns with the task priority. Whenever SEL_{PA} changes to a larger value, the EN_{PA} signal (in Fig. 4) goes from 0 to 1, triggering the fast DVFS function through the asynchronous control [14]. Since the DVFS control circuit is only triggered at the rising edge of the PMU clock, the EN_{PA} signal generates an extra pulse on the clock line to allow the PMU to promptly respond and track the new DVFS operating point. This overcomes the potential delay caused by the low PMU clock frequency due to pulse frequency modulation (PFM) control in low output power cases. Once the input energy voltage level is lower than a programmable threshold, the system transitions into the MEPT mode. In this mode, the SEL is set to SEL_M and EN_M is set to 1.

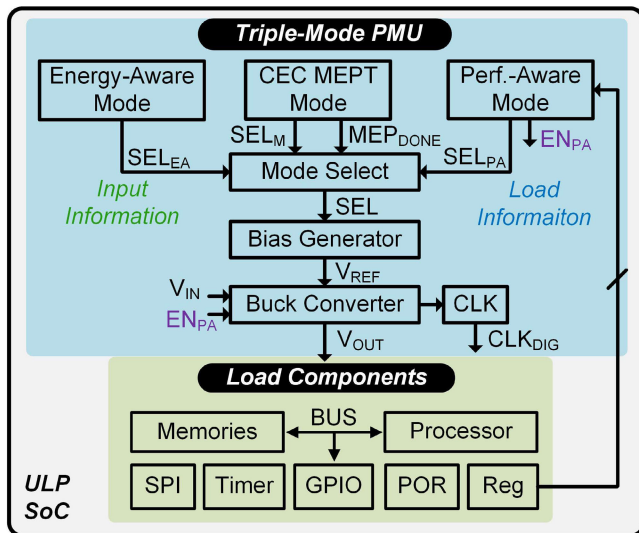


Fig. 4. Proposed SoC architecture with the triple-mode PMU.

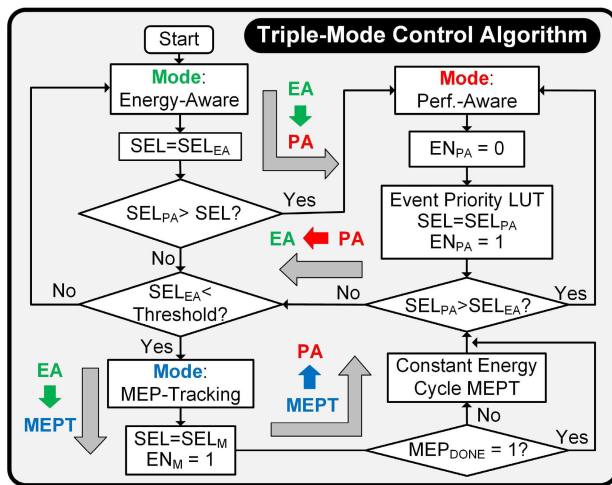


Fig. 5. Triple-MC algorithm.

A hill-climbing-based MEPT algorithm is executed to track the MEP until the MEPT is completed (indicated by $MEP_{DONE} = 1$), and the system remains at the MEP. If the $SEL_{PA} > SEL_{EA}$ or the input energy voltage level starts to rise above the threshold again, the PMU exits the MEPT mode and transitions back to the PA or EA mode. Therefore, by selectively activating the three SEL signals, each mode can work independently without complex mode transitions. This flexibility enables the system to adaptively switch among modes, optimizing the power–performance tradeoff in a dynamic and autonomous manner.

III. MEPT AND ANALYSIS

A. Conventional MEPT Method

The conventional methods of MEPT are demonstrated in Fig. 6. The sample-and-hold method [16] frequently samples the voltage of the output capacitance, C_L . By quantifying the voltage droop at $C_L(V_2 - V_1)$, the energy consumed by the load for every N cycle can be digitized. Since

E_{PC} is proportional to $V_1(V_2 - V_1)$, which is derived from $\Delta E = 1/(2C\Delta V^2)$, the MEPT can be achieved by a hill-climbing algorithm. However, this method requires a high-frequency clock along with a continuous-time comparator to quantify the voltage drop ($V_2 - V_1$), which is not suitable for sub-microwatt applications. The second technique in [17] achieves MEPT by regulating the dynamic-leakage power ratio to a process, voltage and temperature (PVT)-independent constant value through V_{DD} and body bias searching. By providing different clock frequencies to the load component and comparing the operating frequency of the dc–dc converter, a dynamic-leakage power ratio can be indirectly calculated for MEPT. However, it requires deeply depleted channel CMOS technology to enable body biasing. Rahman et al. [18] propose an MEPT scheme with performance regulation. With a fixed operating frequency for the switched capacitor (SC)-based converter and constant input voltage, E_{PC} can be computed. However, a 30-MHz fixed clock is needed, leading to high power consumption. Therefore, for ULP applications, we propose a CEC MEPT method to achieve accurate MEPT with lower power and area overhead, compared with the prior art mentioned above.

B. Proposed CEC MEPT

The proposed CEC MEPT enables the PMU to deliver near-constant energy at each power delivery cycle to the load side and record the load clock cycles through digital counters. E_{PC} can be indirectly calculated by the outputs of the counters. By comparing the counter outputs at adjacent two V_{OUT} 's, the circuit can approach the MEP and lock the operating point once it finishes. Fig. 7 shows the architecture of the MEPT circuits. To match the critical path of the microprocessor regardless of the voltage ripples, a tunable replica oscillator (TR-OSC) with the unified-clock-and-power architecture [22] is implemented to automatically scale the load frequency with the supply voltage. The TR-OSC drives both the load circuits and the asynchronous counters. Two asynchronous counters are implemented to record the load clock cycles with low power overhead.

Fig. 8(a) and (b) shows the algorithm and timing waveform of the proposed CEC MEPT, respectively. After the PMU goes into the MEPT mode, $SEL = SEL_M = 111$, which indicates that the DVFS operating point is controlled by the MEPT block and set to the highest voltage value, 580 mV. After a fixed number (16) of power delivery cycles, V_{OUT} stabilizes, and the MEPT process starts. For the first power delivery cycle, the first asynchronous counter, Counter_H, is enabled to count the load clock. Once V_{OUT} drops back to V_{REF} , Counter_H stops counting, and V_{REF} decreases one step (20 mV) to a lower voltage level. Then, V_{OUT} reaches the new V_{REF} , and Counter_L is enabled to count the load clock, followed by a comparison of the two counters to decide the tracking direction. If Counter_L has a smaller output, which indicates that E_{PC} is higher at the lower V_{OUT} , the MEP is missed. Then, the circuit jumps back to the previous V_{OUT} point and exits.

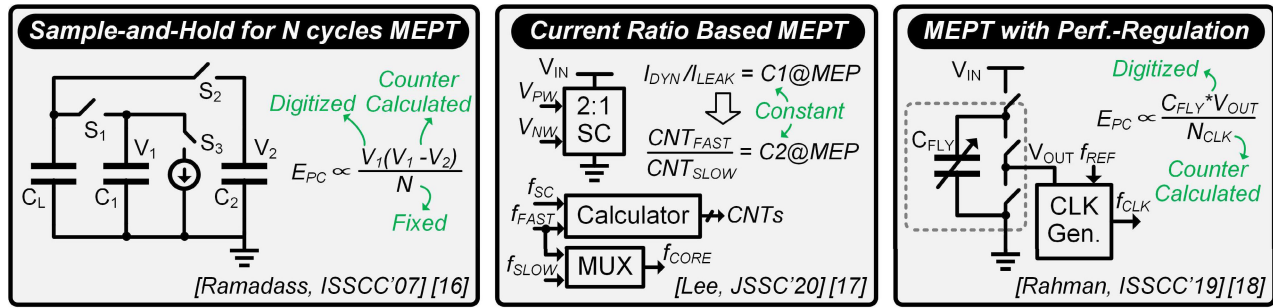


Fig. 6. Conventional MEPT algorithms [16], [17], [18].




Fig. 7. Architecture of the proposed CEC MEPT.

The energy delivered from V_{IN} to V_{OUT} per power delivery cycle can be calculated by the following equation:

$$\begin{aligned}
 E_{CYCLE} &= I_{AVE} V_{IN} T_{HS} \eta_{PS-EFFI} \\
 &= \frac{(V_{IN} - V_{OUT})}{2L} T_{HS} V_{IN} T_{HS} \eta_{PS-EFFI} \\
 &= \frac{T_{HS}^2 (V_{IN} - V_{OUT}) V_{IN}}{2L} \eta_{PS-EFFI} \quad (5)
 \end{aligned}$$

where T_{HS} and L present the width of the HS on time and inductance, respectively. Since N , L , and V_{IN} are constant, and for two adjacent V_{OUT} 's where the step is 20 mV (V_{OUT} ranges 0.4–0.58 V), V_{OUT} and $\eta_{PS-EFFI}$ can be approximately regarded as a constant. Therefore, the energy delivered to the load side can be a constant value for two adjacent V_{OUT} 's. For N cycle power delivery ($N = 1$ in this work), if the counter's output is M_{COUNT} , the load component E_{PC} can be calculated by the following equation:

$$E_{PC} = \frac{N T_{HS}^2 (V_{IN} - V_{OUT}) V_{IN}}{2L M_{COUNT}} \eta_{PS-EFFI}. \quad (6)$$

Therefore, by comparing the counter values for two adjacent V_{OUT} , the MEPT can be achieved. The selection of a 20-mV step size involves a tradeoff between the accuracy of the MEPT and the power/area overhead. A larger step size reduces the number of DVFS operating points in a fixed searching range, resulting in a reduction in power and area overhead, as fewer voltage references are needed. However, this may come at the cost of lower resolution in the MEPT. In our design, we considered the load MEP curve and chose a 20-mV step size to ensure that the energy difference between two consecutive steps closely approaches the limit that the MEPT circuit can distinguish accurately. Further reducing the

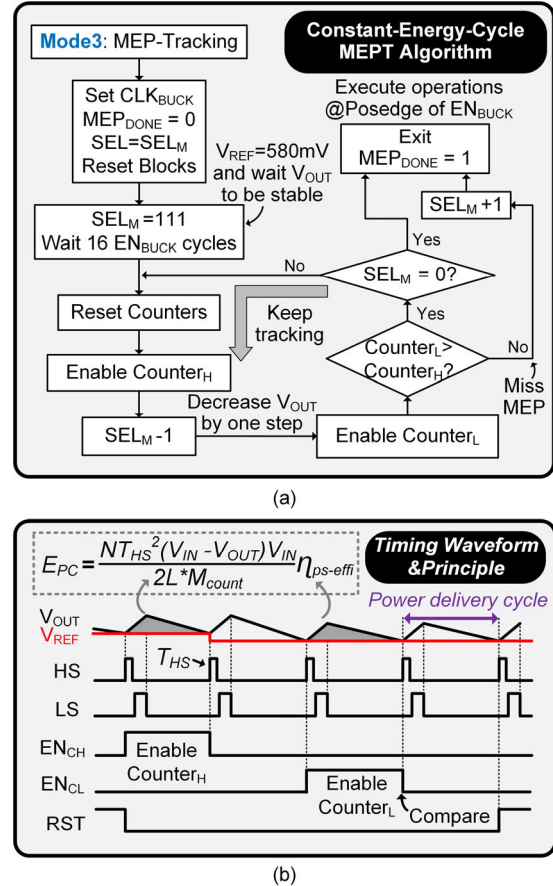


Fig. 8. (a) Flowchart of the proposed CEC MEPT algorithm. (b) Timing waveform of the proposed CEC MEPT.

step size would not yield improvements in resolution, as it would reach a point where the tracking error introduced by the MEPT circuits starts to impose limitations on accuracy. The tracking error introduced by the MEPT will be explored in Section III-C.

C. Accuracy Analysis for the CEC MEPT

The tracking errors of the proposed CEC MEPT are mainly from assuming that the energy delivered to the load is constant for each power delivery cycle. This assumption includes a few approximations that lead to MEPT inaccuracy. In this section, we categorize and quantify the errors associated with this assumption, as shown in Fig. 9.




Fig. 9. Categories and mechanisms of the MEPT errors.

- 1) ΔE_{STEP} : The energy difference (%) between two adjacent power delivery cycles caused by the V_{OUT} changes.
- 2) $\Delta E_{PS-EFFI}$: The energy difference (%) between two adjacent power delivery cycles caused by $\eta_{PS-EFFI}$.
- 3) ΔE_{NOISE} and ΔE_{SAMPLE} : The energy differences (%) between two adjacent power delivery cycles caused by the inaccurate V_{REF} and V_{OUT} comparison due to the comparator noise and low sampling frequency.

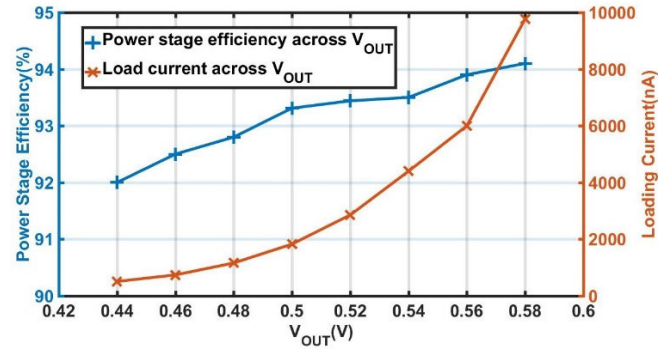
For ΔE_{STEP} , according to (5), the term $(V_{IN} - V_{OUT})$ in the numerator changes while tracking the MEP. Therefore, the energy delivered to the load side is different when V_{OUT} is at V_{OUT1} and V_{OUT2} . The difference can be calculated by the following equation:

$$\Delta E_{STEP} = \left| \frac{(V_{IN} - V_{OUT1}) - (V_{IN} - V_{OUT2})}{(V_{IN} - V_{OUT1})} \right| = \frac{20\text{mV}}{V_{IN} - V_{OUT1}}. \quad (7)$$

Therefore, the maximum error happens when $V_{IN} - V_{OUT1}$ is at a minimum ($V_{IN} = 1.5\text{ V}$ and $V_{OUT1} = 580\text{ mV}$ in this design). The maximum error $\Delta E_{STEP-MAX}$ is then 2.2%, which means that the delivered energy at the current step is 2.2% higher than the energy delivered at the previous step.

For $\Delta E_{PS-EFFI}$, when V_{OUT} decreases, the voltage stress on the inductor goes higher, which leads to degraded power delivery efficiency. The power delivery efficiency is simulated across V_{OUT} with 1.5-V V_{IN} demonstrating a 0.2%–0.6% efficiency drop for two adjacent V_{OUT} 's, as shown in Fig. 10, where the digital core and TR-OSC are the loading components. In other words, at 1.5-V V_{IN} , the amount of energy delivered to the load side is smaller due to the decreased efficiency, which counteracts the effects from ΔE_{STEP} . Thus, the error from ΔE_{STEP} and $\Delta E_{PS-EFFI}$ together is less than 2.2%.

In terms of ΔE_{NOISE} , due to the noises from the comparator and reference voltage, the power delivery happens late or earlier than the time when V_{OUT} reaches V_{REF} . The strong-arm-based dynamic comparator along with the reference voltage is simulated with transient noise analysis [23], [24]. With a 1.5-V V_{IN} and 50-kHz clock, the results show a 120–140- μV rms equivalent input noise at 1σ and 0.4–0.58 V V_{OUT} . Therefore, the energy difference due to the noise can be calculated based on the energy stored on the

Fig. 10. Simulated power stage efficiency across V_{OUT} with 1.5-V V_{IN} when loading the digital core (RISC-V and SRAM) and TR-OSC.

load capacitor, as illustrated by the following equation:

$$\Delta E_{NOISE} = \frac{C_{OUT}V_{OUT}^2 - C_{OUT}(V_{OUT} - 0.14)^2}{2E_{CYCLE}} = C_{OUT} \frac{0.28 * V_{OUT} - 0.0196}{2E_{CYCLE}}. \quad (8)$$

Therefore, we can get the maximum ΔE_{NOISE} at 0.58 V V_{OUT} : $\Delta E_{NOISE-MAX} = 81.4C_{OUT}/E_{CYCLE}$. C_{OUT} is controllable and has a tradeoff versus ripple voltage and tracking speed [14]. According to (5), when T_{HS} is 0.8 μs and inductance is 22 μH in this design, the energy delivered to the load per cycle (E_{CYCLE}) is 20 nJ with 1.5-V V_{IN} and 0.58-V V_{OUT} . To achieve a low ΔE_{NOISE} , for example <2%, C_{OUT} needs to be <4.9 μF .

In this design, we have chosen 4.7 μF as C_{OUT} , which leads to a maximum 2% ΔE_{NOISE} .

For ΔE_{SAMPLE} , the error is introduced by the delayed regulation when the comparator clock frequency is too low, V_{OUT} may drop below V_{REF} before the power delivery cycle happens. Like ΔE_{NOISE} , the power delivery happens at a voltage deviated from V_{REF} . At the highest tracking voltage, where the load current is assumed at around 10 μA (according to the power of the digital core), V_{OUT} decreases with a slope of 2.1 V/s based on the equation $\Delta V/\Delta T = I_{LOAD}/C_{OUT}$ when C_{OUT} is 4.7 μF . To achieve $\Delta E_{SAMPLE} < 2\%$, the voltage droop should be <140 μV . The sampling frequency needs to be larger than 15 kHz. In our design, once the circuit goes into MEPT mode, the clock frequency is automatically set to around 50 kHz to achieve ΔE_{SAMPLE} that is <0.1%, which is negligible. Considering the error caused by ΔE_{STEP} , $\Delta E_{PS-EFFI}$, and ΔE_{NOISE} , the maximum overall error can be calculated by

$$\Delta E_{TOTAL-MAX} = (1 - \Delta E_{PS-EFFI-MIN})(1 + \Delta E_{STEP-MAX}) \times (1 + \Delta E_{NOISE-MAX}) - 1 < 3.6\%.$$

Therefore, the energy delivery to the load side at two adjacent V_{OUT} values differs only by 3.6% at the worst case, showing the accuracy of our tracking scheme.

IV. CIRCUIT IMPLEMENTATION

A. System Implementation

Fig. 11 shows the system block diagram of the proposed IoT SoC. It includes a digital core, a clock and reset generator

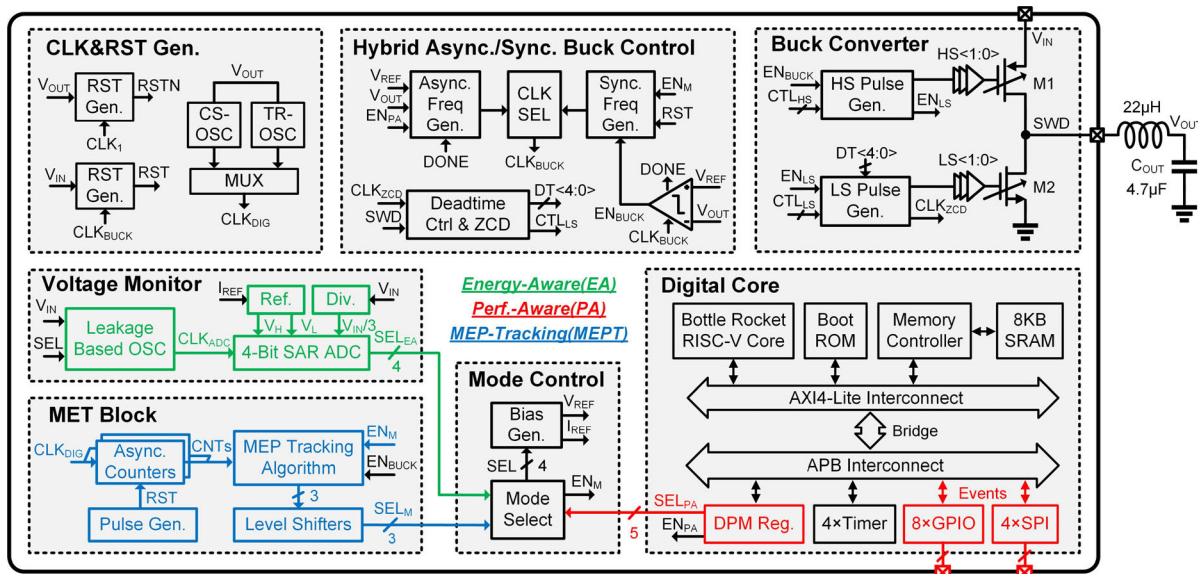


Fig. 11. System block diagram of the ULP IoT SoC with the proposed triple-mode PMU.

block, a buck converter with hybrid control scheme [14], a voltage monitor (VM), an MEPT block, and an MC block. The digital core consists of a 32-bit reduced instruction set computer five (RISC-V) processor, a boot ROM, a memory controller, an 8-kB SRAM, and peripherals. For energy awareness, a 4-bit asynchronous SAR analog to digital converter (ADC) in the VM is clocked by a ULP low-frequency leakage current-based current-starving OSC (CS-OSC) to monitor the input voltage level. For performance awareness, the micro-processor keeps sampling the I/O interfaces and maps the task priority to SEL_{PA4} , which is the last four bits of SEL_{PA} . The most significant bit (MSB) of SEL_{PA} is the comparison results of SEL_{EA} and SEL_{PA4} , indicating the need to switch to the PA mode. The buck converter includes a length-tunable power stage, two pulse generators, and a hybrid async./sync. control scheme [14] for fast DVFS and FLTR. The MEPT block is digitally implemented with a hill-climbing algorithm. To achieve low quiescent power, the entire PMU uses a custom 2.5-V I/O device standard-cell library except for the MEPT block that uses 1.2-V core devices. This allows the MEPT block to be powered by V_{OUT} rails for lower dynamic power and a smaller area overhead. Therefore, with all these functions and techniques integrated on-chip, the SoC can flexibly and adaptively scale its power and performance based on both input conditions and performance requirements.

B. Digital Core

The digital core features an 8-kB SRAM macro and a 32-bit Bottle-Rocket RISC-V microcontroller class processor core with a basic three-stage pipeline that implements the RV32IMC instruction set [25]. Besides the RISC-V core, there are existing ULP digital cores that can push the floor power down to the nanowatt level, including the ARM-Cortex M series [9], [17], [18] and MSP430 series [10], [26]. However, they both are not open source and require a license fee. The MSP430 series is from Texas Instrument, resulting in

a more limited selection of development tools and libraries. In contrast, the RISC-V core has multiple different versions to choose from for different applications for free with an expanding ecosystem, which is the main reason we choose this type of core.

The RISC-V interfaces to a boot-ROM, peripherals, and an embedded memory controller with the SRAM, through a custom Acorn RISC Machine (ARM) Advanced Microcontroller Bus Architecture (AMBA) eXtensible Interface 4 (AXI4)-Lite bus and an ARM AMBA Advanced Peripheral Bus (APB). The peripherals include an 8-bit general-purpose input-output (GPIO) interface, four serial peripheral interface (SPI) masters, four timers, and a configuration block that contains memory-mapped registers for communication with on-chip blocks. Moreover, a joint test action group (JTAG) debug module is included as part of the Bottle-Rocket package. The memory controller enables a modular memory interface supporting most bus protocols, including the ARM Cortex and the RISC-V processors. The memory-mapped 8-kB SRAM in the system serves as both the instruction and data memory. The memory is custom-designed for self-powered systems with a high- V_{TH} 6T bit-cell to operate at sub-threshold voltages between 0.4 and 1.2 V, achieving nanowatt-level power consumption. Designing an SRAM bit-cell to operate at a sub-threshold voltage that satisfies the requirements, such as read/write stability, leakage, power consumption, operation voltage, operating frequency, and density, is a multi-dimensional design-space exploration process. The bit-cell type, bit-cell size, device type, assist techniques, and micro-architecture are all important SRAM knobs. We implement our design using a bit-cell design tool [27] that automates the bit-cell generation process for a given user specification. This auto-generation flow decides the bit-cell design knobs by performing the multi-dimensional design-space exploration and replacing the human engineer hours with machine computing time. The SRAM includes four 2-kB banks, each with two sub-banks. Each sub-bank shares the same peripheral circuitry, offering speed, area, and leakage

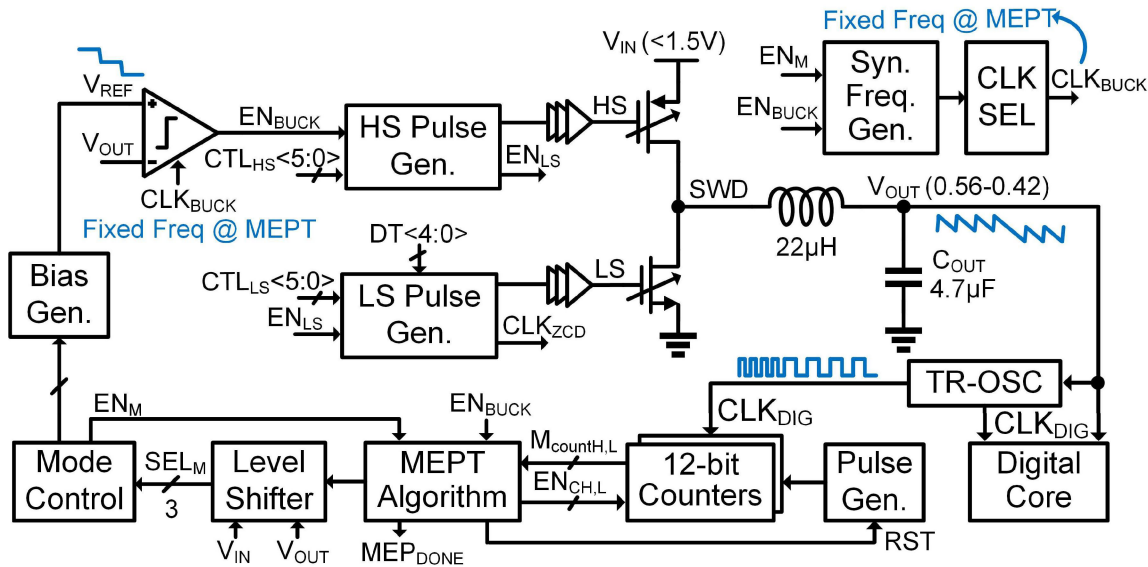


Fig. 12. Detailed schematic of the proposed asynchronous CEC MEPT.

improvements. Read and write assist circuitry is also included in the SRAM to further enhance the robustness. Except for the SRAM, the rest of the blocks of the digital core are synthesized using a standard automated place and route flow.

C. MEPT Circuits

Fig. 12 shows the circuit implementation of the proposed CEC MEPT. The circuit reuses the existing PMU circuits and signals. Therefore, it only needs two 12-b asynchronous counters, a pulse generator, an MEPT algorithm control block, and three level shifters. The bit width of the counter needs to be large enough to guarantee that the counter will not overflow during counting. This can be calculated based on the energy delivered from V_{IN} to V_{OUT} per power delivery cycle and the energy per cycle of the load components. For example, in our design, the energy per power delivery is typically 10–40 nJ between 0.42 and 0.55 V and the energy per cycle of the digital core is approximately 32 pJ at 0.55 V. Therefore, the number of load cycles that is needed to consume the 40-nJ energy is 1250. To achieve low power and area overhead with high efficiency, three key techniques are utilized.

- 1) The EN_{BUCK} and CLK_{DIG} signals are reused as inputs for the MEPT algorithm control block and asynchronous counters, respectively. By leveraging the existing EN_{BUCK} signal to indicate a new power delivery cycle for the buck converter and utilizing the CLK_{DIG} signal as the load clock cycles, there is no need for additional clocks or signals, thereby saving power.
- 2) The whole function is achieved through a digital hill-climbing algorithm-based feedback loop. This avoids using high-power analog components and allows the function to be implemented with V_{OUT} as supply voltage, resulting in power reduction without compromising performance.
- 3) As discussed in Section III-C, the buck converter clock frequency is regulated at a >50 -kHz range

to minimize the tracking error from low sampling frequency. Therefore, with all these technologies, the MEPT block achieves a high tracking accuracy with only a 0.026 -mm² area overhead and sub-nW power consumption.

D. Sub-nW Buck Converter and Hybrid Control

The digital buck converter consists of power stage and pulse generators along with its hybrid synchronous and asynchronous control. The comparator compares V_{REF} and V_{OUT} and controls the power delivery cycles by turning on/off the power stage transistors. Adaptive deadtime and zero current detector (ZCD) detection are implemented together to avoid extra conduction loss and reverse current for higher power efficiency. The synchronous control loop regulates the output voltage and adaptively changes the clock frequency using a PFM scheme to provide high efficiency during light load conditions. The asynchronous loop has the functions of: 1) detecting the voltage droop or DVFS requests and 2) generating an asynchronous pulse to over-clock the comparator and regulate the output voltage, achieving microsecond-level fast tracking. By utilizing those techniques [14], the PMU achieves fast DVFS and FLTR with sub-nanowatt power overhead.

E. Sub-Nanowatt Asynchronous ADC

A conventional SAR ADC uses a binary search procedure to sense and digitize the input voltage, which often ranges from V_{IN} to ground (GND). For an N -bit ADC, the smallest analog increment corresponding to a 1-LSB change (voltage resolution) is only $V_{IN}/2^N$. As the input voltage increases, the number of output bits of the ADC must increase to maintain the same voltage resolution, which leads to increased area and power overhead. In this application, it is unnecessary to cover the whole input voltage range from V_{IN} to GND since the battery would be almost depleted after the voltage drops 50% from the nominal voltage [28]. In this design, as shown in

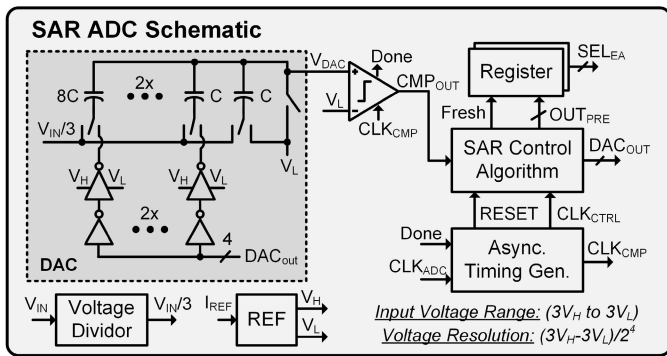


Fig. 13. Detailed schematic of the asynchronous SAR ADC.




Fig. 14. Schematic of the nW BG.

Fig. 13, the input voltage range that needs to be digitized is only 1.5–2.5 V. To efficiently sense this input voltage range, we implement an ADC with two references: V_H and V_L . With the voltage divider, it achieves an input voltage resolution of $(3V_H - 3V_L)/2^d$ with a smaller number of output bits, compared with conventional schemes with a full reference range.

F. Sub-Nanowatt BG and Clocks

The BG uses a beta multiplier current reference with a 28-M Ω on-chip resistor to generate supply-independent voltage references, as shown in Fig. 14. It generates 23 references, of which 15 are allocated for DVFS with a voltage step of around 40 mV ranging from 0.5 to 1.1 V, and eight are designated for MEPT with a voltage step of around 20 mV ranging from 0.42 to 0.56 V. To save power, each current mirror branch generates two voltage references instead of one. Those references are selected by a multiplexer (MUX) to generate V_{REF} . Each of the references has a 6- or 8-pF on-chip decoupling capacitor. The clock generator of the core consists of two ring OSCs, which includes a TR-OSC and a tunable CS-OSC, as shown in Fig. 15. The TR-OSC contains different delay stages [29]. Each stage includes different types of digital logic cells (inverters, NOR gates, and resistors) with independent delay tunability. By experimentally tuning the delay of each stage in the TR-OSC, the critical path of the digital core can be emulated by the TR-OSC for MEPT. The CS-OSC provides a larger operating frequency range for the SoC and enables low-frequency operation.

V. EXPERIMENTAL RESULTS

The SoC is fabricated in a bulk planar 65-nm LP CMOS process. The die photograph is shown in Fig. 17 with a

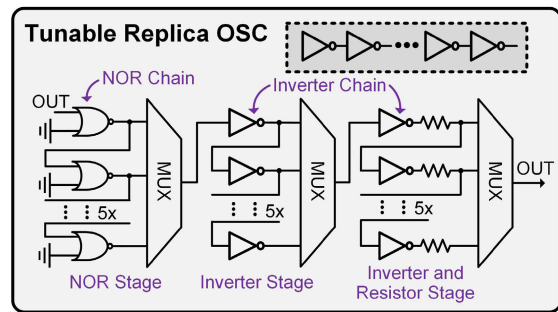


Fig. 15. Schematic of the TR-OSC.

die area of 1.56×1.95 mm. The chip is tested with a QFN100 package, an LPS5030-223MRC inductor, a 10- μ F input capacitor, and a 4.7- μ F load capacitor.

A. System Operation

Fig. 16(a) shows the flowchart of the C code that is loaded into the digital core for DVFS control and Fig. 16(b) shows the setup of programming the digital core. Two methods are available for this design to program the digital core, which both are verified experimentally. The first method is to utilize gcc-toolchain [30] to convert the C program into an executable and linkable format (elf) file, followed by using GDB and Open on-chip debugger (OCD) [31] to directly load the program into the SRAM by communicating with the JTAG port over a J-link probe [32]. Another method is to program the digital core over the non-volatile memory (NVM). Once the digital core boots up, it automatically loads programs over SPI. By loading the desired program into the NVM by an IO-3200 pattern generator and logic analyzer (PGLA), the digital core can load and execute the program stored in the NVM. As shown in Fig. 16(a), the loaded program lets the digital core periodically scan the GPIO port and decide the time of enabling the PA mode by comparing SEL_{PA} and SEL_{EA} . Once the program is executed, the SoC adaptively switches among the three modes based on both the input and output conditions. Fig. 16(c) shows the measured triple-mode transition waveform. A Keysight B2902A sourcemeter generates a triangle V_{IN} voltage, while three GPIO signals, which represent different priority levels, are controlled by a PGLA. At the beginning, when V_{IN} increases, the PMU is in the EA mode so that V_{OUT} also increases with V_{IN} . In the EA mode, if the ADC output changes, the PMU enables the asynchronous loop to quickly track the new reference voltage. Whenever an event occurs (GPIO signal goes to 1), the digital core maps the event address with its LUT and changes SEL_{PA} . If SEL_{PA} is larger than the SEL, the system goes into the PA mode and the digital core changes the EN_{PA} signals (shown in Fig. 11). This change enables the asynchronous loop of the PMU to quickly regulate V_{OUT} to the new reference voltage, as discussed in Section II-C. Therefore, when V_{IN} is relatively low, the system has a higher chance to move into the PA mode when an event occurs since the SEL is likely smaller than SEL_{PA} , indicating that the performance requirements are not met. After V_{IN} goes below 1.5 V, the MEPT block is enabled to start tracking to

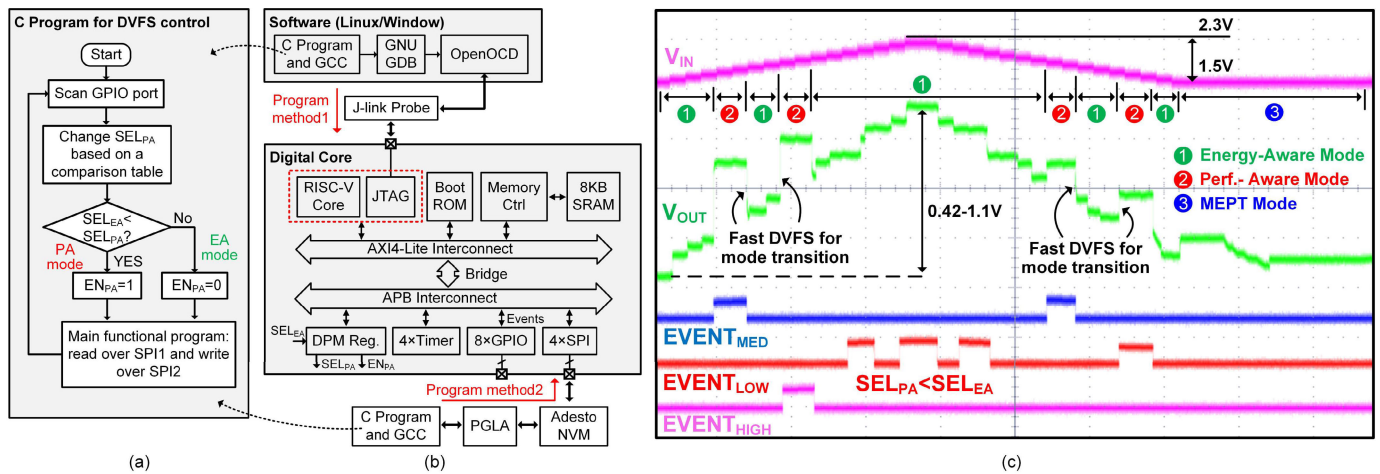


Fig. 16. (a) Flowchart of the C code loaded into digital core for DVFS control. (b) Testing setups and two methods of programming the digital core. (c) Measured triple-mode transition among the EA mode with V_{IN} changing between 1.5 and 2.3 V, the PA mode with triple prioritized events, and the MEPT mode after V_{IN} drops below 1.5 V.




Fig. 17. Chip micrograph of the IoT SoC.

keep the system operating at the MEP until V_{IN} charges up again or a prioritized event occurs. Thanks to the TR-OSC, the clock frequency can automatically scale with the voltage of the digital core to ensure its functionality during DVFS transitions. With the async./sync. control scheme validated by our previous work in [14], the PMU achieves 8.32-mV/ μ s up-tracking and 4.64-mV/ μ s down-tracking speed. When the load current changes from 45 nA to around 1 mA within 100 ns, the voltage droop is 56 mV and the settling time is 183 μ s.

B. PMU Efficiency

Fig. 18(a) shows the measured power efficiency of the triple-mode PMU across output power. The efficiency is calculated by using the load power at V_{OUT} divided by the power measured at V_{IN} . A sourcemeter is used at V_{OUT} to provide load current. The MEPT block is powered by V_{OUT} , so its power is measured separately and added into the input power (P_{PMU}). The results show that the PMU achieves a 92.6% peak efficiency and maintains an efficiency >80% from

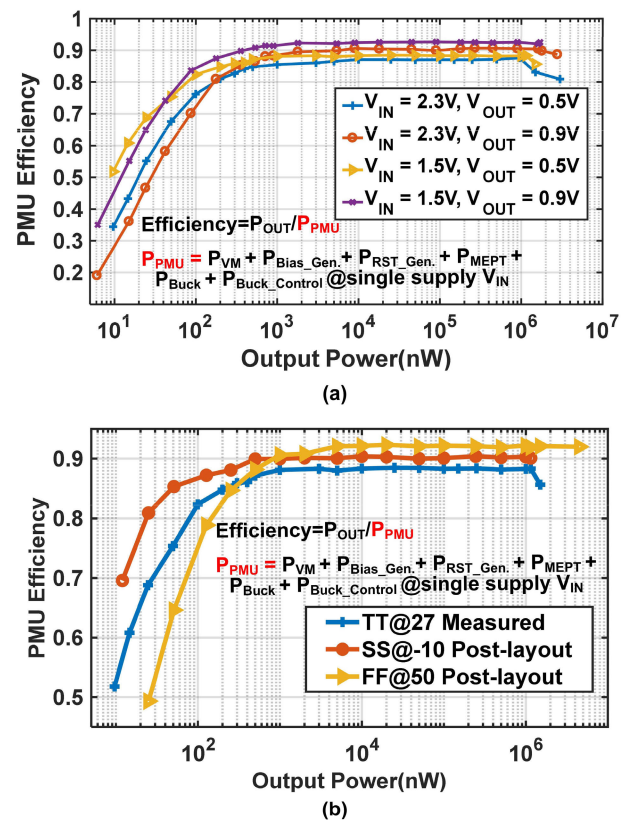


Fig. 18. (a) Measured power efficiency of the triple-mode PMU across output power at different input and output voltages. (b) Measured and simulated power efficiency of the triple-mode PMU across output power at 1.5-V V_{IN} and 0.5-V V_{OUT} across different corners and temperatures.

190 nW to 3 mW, providing over four orders of magnitude of the load power range. Fig. 18(b) shows the measured and simulated (post-layout) power efficiency of the triple-mode PMU across output power at 1.5-V V_{IN} and 0.5-V V_{OUT} across different corners and temperatures. Thanks to the PFM control, the measured switching frequency of the PMU automatically scales between 21 Hz and 163 kHz according to the load current.

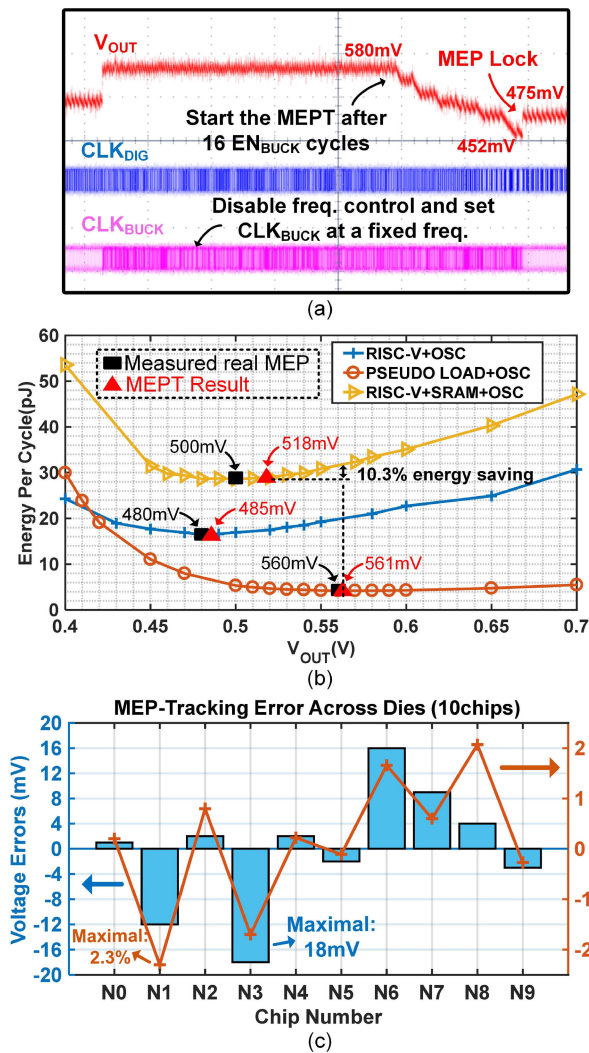


Fig. 19. (a) Measured MEPT waveform, (b) measured MEPT accuracy for different loads, and (c) measured MEPT accuracy in terms of voltage error and E_{PC} across ten dies.




Fig. 20. Setup for efficiency and quiescent power measurement.

C. MEPT Accuracy

The measured MEPT waveform is shown in Fig. 19(a). After the SoC goes into the MEPT mode, V_{OUT} is regulated at 580 mV, while the clock of the buck converter is fixed at a >50 -kHz frequency. After V_{OUT} is stable, the MEPT process starts. Fig. 19(b) shows the MEPT accuracy across load variation. The pseudo loads are composed of multi-threshold [high threshold (HVT)/standard threshold (SVT)/low threshold (LVT)] power-gateable digital counters with tunable

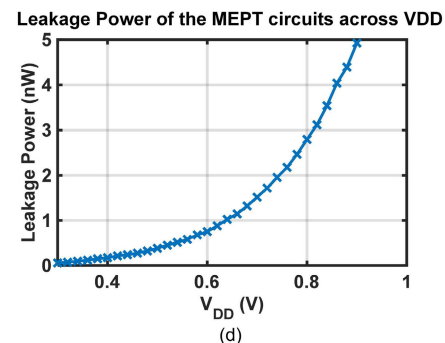
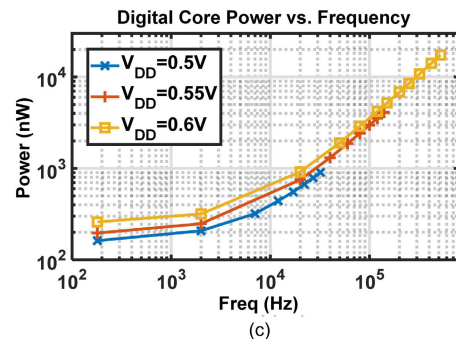
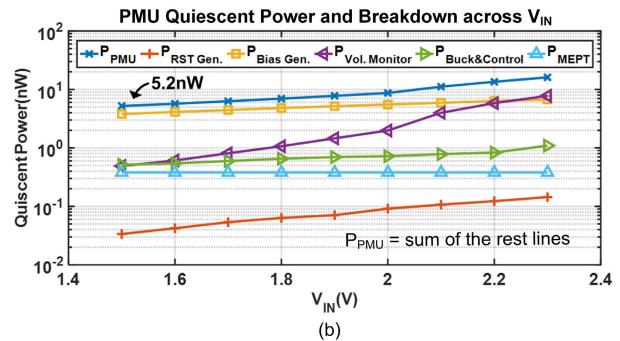



Fig. 21. (a) Measured power breakdown of the IoT SoC and PMU and (b) quiescent PMU power breakdown of the PMU across V_{IN} . (c) Power consumption of digital core across frequency at 0.5, 0.55, and 0.6 V. (d) Measured leakage power across supply voltage for the MEPT circuits.

load capacitors. Therefore, the leakage and dynamic power of the pseudo loads are controllable. Three load combinations, including RISC-V with TR-OSC, pseudo loads with TR-OSC, and digital core (RISC-V + SRAM) with TR-OSC, are tested to verify the MEPT accuracy. The MEPT result is the mode value of MEPT results over 20 iterations and it achieves <18 -mV error compared with the real (measured) MEP. Ten chips are tested with the pseudo load and TR-OSC as loading components, and the results show that the maximum voltage error is 18 mV and the maximum E_{PC} error is 2.3%, as shown in Fig. 19(c). This allows the load circuits to operate close to

TABLE I

COMPARISON OF THE PMU-ENABLED IoT SoC WITH STATE-OF-THE-ART

	[16] ISSCC'07	[18] ISSCC'19	[17] JSSC'20	This Work
Technology	65nm	55nm CCD	65nm	65nm
Processor and SRAM	N/A	Cortex-M0+8KB	Cortex-M3+512B	RISC-V+8KB
Regulated Voltage (V)	0.25-0.7	0.48-0.75	0.35-0.58	0.4-1.1
Operating Frequency	N/R	100KHz-6MHz	1.1Hz-38MHz	180Hz-5.7MHz
PMU Architecture	Sync. Buck	Cascade SC	C _{FLY} Tuned SC	Async./Sync. Buck
Power Management Techniques	MEPT	MEPT	MEPT+Perf. Aware	MEPT+Perf. Aware +Energy Aware
MEPT voltage error or E _{PC} error	N/R	<4.7%	≤5mV	<20mV/<2.3%
MEPT Power Consumption	N/R	84nW*	2μW	412pW
MEPT Area	0.05mm ²	N/R	0.043mm ²	0.026mm ²
Fast DVFS	No	No	No	Yes
Fast load Response	No	No	No	Yes
Dynamic Load Range with Efficiency > 80%	1μW-100μW (100)	N/R	N/R	190nW-3mW (1.57x10 ⁴)
PMU Peak Efficiency (%)	86	N/R	82*	92.6
System Power Consumption	1.23μW-116.2μW*	>110nW*	>2.4μW*	194.3nW-598.6μW
Components included in Minimum System Power	PMU+FIR	Cortex-M0+SRAM +CLK+TIMER	N/R	RISCV+SRAM+IOs+PMU+TIMERS+ROM+CLK

*Calculated/Observed from waveforms

the real MEP. Also, compared with using a fixed operating voltage, the MEPT can bring an energy saving up to 10.3% when the load component changes.

D. System Power Breakdown and Analysis

Fig. 20 shows the measurement setup for the PMU efficiency and quiescent power measurement. A Keithley 2401 sourcemeter is configured as a load at the V_{OUT} rail. The quiescent power, including leakage, is measured by a Keithley 6430 sub-Femto sourcemeter with high accuracy. Fig. 21(a) shows the system and PMU power breakdown. The SoC has a minimum system power consumption of 194.3 nW at 180 Hz clocked by the CS-OSC. In addition, the proposed PMU achieves 5.2-nW quiescent power. The loss from the control circuits, including DVFS, is less than 10% by reusing the existing signals and circuits [14] and the major power loss comes from the BG due to the 23 voltage references that are required for DVFS and MEPT. The power overhead of the MEPT circuits in the idle state, with a V_{OUT} of 0.5 V, accounts for only 0.19% of the total system power. Fig. 21(b) shows the quiescent power breakdown of the PMU across V_{IN} . The total quiescent power of the PMU achieves 5.2 nW at 1.5 V, with the BG component contributing the most to power consumption. Fig. 21(c) shows the power consumption of the digital core, which includes SRAM and RISC-V processor across frequency and supply voltage. The total power of the digital core and TR-OSC is 904 nW when the core operates at 32 kHz with a supply voltage of 0.5 V. The maximal operation frequency of the digital core is 32, 140, and 520 kHz, at 0.5, 0.55, and 0.6 V, respectively. The TR-OSC is manually tuned to align with the maximal frequencies with a margin added for reliable operations. The measured leakage power of the MEPT circuit across supply voltage is illustrated in Fig. 21(d), showing that the MEPT only consumes 379 pW at 0.5 V.

E. Comparison to State-of-the-Art Works

Table I compares the proposed PMU-enabled SoC with state-of-the-art works, which have not previously targeted the nanowatt-level power range. Our PMU maintains a high efficiency over a load range that is $>100\times$ than the prior art and achieves the highest peak efficiency. Thanks to the hybrid buck control scheme, this PMU also features fast DVFS and FLTR which previous works do not support. Besides, the proposed triple-mode power management allows the SoC to coordinate both the input and load conditions to achieve a flexible tradeoff between energy and performance. The CEC MEPT circuit achieves $<2.3\%$ E_{PC} error with $>100\times$ power overhead reduction and the lowest area overhead of 0.026 mm². Finally, the SoC and PMU achieve MEPT for energy minimization, performance regulation, and available input energy awareness while simultaneously allowing these techniques to be applied to ULP, nanowatt-scale SoCs.

VI. CONCLUSION

To achieve energy optimization while simultaneously regulating the performance of the SoC for sub-microwatt IoT applications, we present a 194-nW SoC with a triple-mode PMU that achieves available energy adaptability, performance scaling, and MEPT. By utilizing a PMU-processor-in-loop control architecture, the SoC can self-scale its performance–power based on both the input available energy and performance requirements. Fast DVFS and FLTR for mode transition are achieved through a sub-nanowatt hybrid asynchronous and synchronous control for the buck converter. With the proposed CEC MEPT, the SoC can track the most efficient operating point for energy minimization with <18 -mV voltage error and $<2.3\%$ E_{PC} error. Thanks to the digital implementation and reuse of existing signals, the MEPT circuits achieve the lowest 0.026-mm² area overhead and 412-pW active power overhead, which only accounts for 0.19% of the total system power and achieves $>100\times$ power reduction compared with prior arts. With all those techniques mentioned above and ULP design for SAR ADC, BG, and OSCs, the IoT SoC and PMU achieve a minimum 194- and 5.2-nW quiescent power, respectively. All these results and features make this SoC well-suited for ULP IoT applications.

REFERENCES

- [1] A. Dissanayake, H. L. Bishop, S. M. Bowers, and B. H. Calhoun, "A 2.4 GHz-91.5 dBm sensitivity within-packet duty-cycled wake-up receiver," *IEEE J. Solid-State Circuits*, vol. 57, no. 3, pp. 917–931, Mar. 2022.
- [2] Y.-S. Noh, J.-I. Seo, H.-S. Kim, and S.-G. Lee, "A reconfigurable DC–DC converter for maximum thermoelectric energy harvesting in a battery-powered duty-cycling wireless sensor node," *IEEE J. Solid-State Circuits*, vol. 57, no. 9, pp. 2719–2730, Sep. 2022.
- [3] S. S. Amin and P. P. Mercier, "MISIMO: A multi-input single-inductor multi-output energy harvester employing event-driven MPPT control to achieve 89% peak efficiency and a 60,000x dynamic range in 28 nm FDSOI," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2018, pp. 144–146.
- [4] D. S. Truesdell, S. Li, and B. H. Calhoun, "A 0.5-V 560-kHz 18.8-fJ/cycle on-chip oscillator with 96.1-ppm/°C steady-state stability using a duty-cycled digital frequency-locked loop," *IEEE J. Solid-State Circuits*, vol. 56, no. 4, pp. 1241–1253, Apr. 2021.

- [5] T. D. Burd, T. A. Pering, A. J. Stratakos, and R. W. Brodersen, "A dynamic voltage scaled microprocessor system," *IEEE J. Solid-State Circuits*, vol. 35, no. 11, pp. 1571–1580, Nov. 2000.
- [6] J. Myers et al., "A 12.4 pJ/cycle sub-threshold, 16 pJ/cycle near-threshold ARM Cortex-M0+ MCU with autonomous SRPG/DVFS and temperature tracking clocks," in *Proc. Symp. VLSI Circuits*, Jun. 2017, pp. C332–C333.
- [7] P. A. Meinerzhagen et al., "An energy-efficient graphics processor in 14-nm tri-gate CMOS featuring integrated voltage regulators for fine-grain DVFS, retentive sleep, and V_{MIN} optimization," *IEEE J. Solid-State Circuits*, vol. 54, no. 1, pp. 144–157, Jan. 2019.
- [8] D. S. Truesdell, J. Breiholz, S. Kamineni, N. Liu, A. Magyar, and B. H. Calhoun, "A 6–140-nW 11 Hz–8.2-kHz DVFS RISC-V microprocessor using scalable dynamic leakage-suppression logic," *IEEE Solid-State Circuits Lett.*, vol. 2, no. 8, pp. 57–60, Aug. 2019.
- [9] P. Prabhat et al., "27.2 M0N0: A performance-regulated 0.8-to-38 MHz DVFS ARM cortex-M33 SIMD MCU with 10 nW sleep power," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2020, pp. 422–424.
- [10] L. Lin, S. Jain, and M. Alioto, "Integrated power management for battery-indifferent systems with ultra-wide adaptation down to nW," *IEEE J. Solid-State Circuits*, vol. 55, no. 4, pp. 967–976, Apr. 2020.
- [11] X. Liu, C. Huang, and P. K. T. Mok, "A high-frequency three-level buck converter with real-time calibration and wide output range for fast-DVFS," *IEEE J. Solid-State Circuits*, vol. 53, no. 2, pp. 582–595, Feb. 2018.
- [12] J.-G. Kang, M.-G. Jeong, J. Park, and C. Yoo, "A 10 MHz time-domain-controlled current-mode buck converter with 8.5% to 93% switching duty cycle," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2018, pp. 424–426.
- [13] S. Pan and P. K. T. Mok, "A 10-MHz hysteretic-controlled buck converter with single on/off reference tracking using turning-point prediction for DVFS application," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 66, no. 11, pp. 4502–4515, Nov. 2019.
- [14] X. Liu, B. H. Calhoun, and S. Li, "A sub-nW 93% peak efficiency buck converter with wide dynamic range, fast DVFS, and asynchronous load-transient control," *IEEE J. Solid-State Circuits*, vol. 57, no. 7, pp. 2054–2067, Jul. 2022.
- [15] B. H. Calhoun, A. Wang, and A. Chandrakasan, "Modeling and sizing for minimum energy operation in subthreshold circuits," *IEEE J. Solid-State Circuits*, vol. 40, no. 9, pp. 1778–1786, Sep. 2005.
- [16] Y. K. Ramadass and A. P. Chandrakasan, "Minimum energy tracking loop with embedded DC–DC converter enabling ultra-low-voltage operation down to 250 mV in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 43, no. 1, pp. 256–265, Jan. 2008.
- [17] J. Lee et al., "A self-tuning IoT processor using leakage-ratio measurement for energy-optimal operation," *IEEE J. Solid-State Circuits*, vol. 55, no. 1, pp. 87–97, Jan. 2020.
- [18] F. U. Rahman, R. Pamula, A. Boora, X. Sun, and V. Sathe, "19.1 computationally enabled total energy minimization under performance requirements for a voltage-regulated 0.38-to-0.58 V microprocessor in 65 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2019, pp. 312–314.
- [19] W. Lim, I. Lee, D. Sylvester, and D. Blaauw, "8.2 batteryless sub-nW Cortex-M0+ processor with dynamic leakage-suppression logic," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2015, pp. 1–3.
- [20] X. Liu, S. Kamineni, J. Breiholz, B. H. Calhoun, and S. Li, "A 194 nW energy-performance-aware IoT SoC employing a 5.2 nW 92.6% peak efficiency power management unit for system performance scaling, fast DVFS and energy minimization," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2022, pp. 1–3.
- [21] X. Wu et al., "A 20-pW discontinuous switched-capacitor energy harvester for smart sensor applications," *IEEE J. Solid-State Circuits*, vol. 52, no. 4, pp. 972–984, Apr. 2017.
- [22] F. U. Rahman et al., "A unified clock and switched-capacitor-based power delivery architecture for variation tolerance in low-voltage SoC domains," *IEEE J. Solid-State Circuits*, vol. 54, no. 4, pp. 1173–1184, Apr. 2019.
- [23] T. Rabuske and J. Fernandes, "Noise-aware simulation-based sizing and optimization of clocked comparators," *Anal. Integr. Circuits Signal Process.*, vol. 81, no. 3, pp. 723–728, Dec. 2014.
- [24] S. Art, "Keeping things quiet: A new methodology for dynamic comparator noise analysis," *EE J. Chalk Talk Ser.*, Dec. 2016. Accessed: Nov. 2022. [Online]. Available: https://www.cadence.com/content/dam/cadencewww/global/en_US/videos/tools/custom_ic_analog_rf_design/NoiseAnalysisposting201612Chalk%20Talk.pdf
- [25] K. Asanović. (Apr. 2016). "The rocket chip generator." EECS Dept., Univ. California, Berkeley, CA, USA. Tech. Rep. UCB/EECS-2016-17. [Online]. Available: <https://www2.eecs.berkeley.edu/Pubs/TechRpts/2016/EECS-2016-17.pdf>
- [26] L. Lin, S. Jain, and M. Alioto, "Sub-nW microcontroller with dual-mode logic and self-startup for battery-indifferent sensor nodes," *IEEE J. Solid-State Circuits*, vol. 56, no. 5, pp. 1618–1629, May 2021.
- [27] S. Kamineni, "Design methodologies and frameworks for autonomous synthesis of system on chip (SoC) components," M.S. thesis, Dept. Elect. Comput. Eng., Univ. Virginia, Charlottesville, Virginia, 2023. Available: [Online]. Available: https://libraetd.lib.virginia.edu/public_view/zc77sr31x
- [28] *CR2032 Coin Battery Datasheet*. Accessed: Nov. 2022. [Online]. Available: <https://data.energizer.com/pdfs/cr2032.pdf>
- [29] J. Tschanz, K. Bowman, S. Walstra, M. Agostinelli, T. Karnik, and V. De, "Tunable replica circuits and adaptive voltage-frequency techniques for dynamic voltage, temperature, and aging variation tolerance," in *Proc. Symp. VLSI Circuits*, Jun. 2009, pp. 112–113.
- [30] *Pre-Built RISC-V GCC Toolchain Binaries*. Accessed: Nov. 2023. [Online]. Available: <https://www.sifive.com/software>
- [31] *RISC-V OpenOCD Official Release*. Accessed: Nov. 2023. [Online]. Available: <https://github.com/riscv/riscv-openocd>
- [32] *SEGGER J-Links Probe*. Accessed: Nov. 2023. [Online]. Available: <https://www.segger.com/products/debug-probes/j-link>



Xinjian Liu (Graduate Student Member, IEEE) received the B.Eng. degree in microelectronics from Fudan University, Shanghai, China, in 2019. He is currently pursuing the Ph.D. degree in electrical engineering with the University of Virginia (UVA), Charlottesville, VA, USA.

He joined UVA in July 2019. In the summer of 2022, he was a Silicon Design Intern with Everactive, Charlottesville. His research interests include low-power dc–dc converters, power management unit, and Internet-of-Things (IoT) system-on-chip design.

design.

Mr. Liu was a Winner of the 2019–2020 IEEE SCS International Student Circuits Video Contest Award and a recipient of the 2023–2024 IEEE SCS Predoctoral Achievement Award. He was also a Winner of the 2023 Link Lab Distinguished Research Award from UVA. He serves as a reviewer for IEEE JOURNAL OF SOLID-STATE CIRCUITS.



Sumanth Kamineni (Member, IEEE) received the B.Tech. degree from Sri Venkateswara University, Tirupathi, India, in 2012, the M.Tech. degree from VIT University, Vellore, India, in 2015, and the Ph.D. degree from the University of Virginia, Charlottesville, VA, USA, in 2023.

He previously worked as a CAD Engineer at Microchip Technology, Chennai, India, from 2015 to 2017. Currently, he works as a Senior Circuit Design Engineer of SRAM at Nvidia, Santa Clara, USA. His research interests include memory

design, machine learning-based circuit design and automation, autonomous SoC synthesis, electronic design automation (EDA), STD cell-compatible unit cell design for circuit synthesis, and open-source EDA tools.



Jacob Breiholz received the B.S. and Ph.D. degrees in electrical engineering from the University of Virginia, Charlottesville, VA, USA, in 2015 and 2021, respectively.

From 2021 to 2023, he was a Senior SoC Design Engineer at Everactive, Inc., Charlottesville, where he focused on digital design and implementation for ultra-low power, energy-harvesting wireless systems-on-chip (SoCs). Currently, he is a Staff Physical Design Engineer at ASIC North, Williston, VT, USA. His research interests include low-power

digital circuit design and self-powered SoC design for Internet-of-Things (IoT) applications.



Benton H. Calhoun (Fellow, IEEE) received the B.S. degree in electrical engineering from the University of Virginia, Charlottesville, VA, USA, in 2000, and the M.S. degree and the Ph.D. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2002 and 2006, respectively.

In January 2006, he joined the Department of Electrical and Computer Engineering, University of Virginia, where he is currently the Alice M. and Guy A. Wilson Professor of Electrical and Computer Engineering. He is the Campus Director and the Technical Thrust Leader of the NSF Nanosystems Engineering Research Center (ERC) for Advanced Self-Powered Systems of Integrated Sensors and Technologies (ASSIST). He co-founded and is the Co-CTO at Everactive, Inc., Charlottesville, which is selling self-powered, energy-harvesting wireless sensing solutions in the industrial Internet-of-Things (IoT) market. His research has emphasized energy-efficient and sub-threshold circuit design for self-powered, battery-less wireless sensing systems. Starting from fundamental advances in sub-threshold circuits, he has expanded his work to include complete self-powered nodes for IoT and body-worn applications. He has coauthored *Sub-Threshold Design for Ultra-Low-Power Systems* (Springer, 2006) and authored *Design Principles for Digital CMOS Integrated Circuit Design* (NTS Press, 2012). He has over 260 peer-reviewed publications and 30 issued U.S. patents that contribute to the field of energy-efficient circuits and systems for self-powered and energy-constrained applications. His research interests include self-powered wireless sensors for the IoT, battery-less systems, body area sensor networks, low-power digital circuit design, system-on-chip architecture and circuits for energy-constrained applications, system-driven embedded hardware/software design, wakeup receivers, energy-harvesting-power management units, sub-threshold digital circuits, sub-threshold SRAM, energy-efficient communication, power harvesting and delivery circuits, low-power mixed-signal design, and medical applications for low-energy electronics.



Shuo Li (Member, IEEE) received the B.Eng. degree in microelectronics from the University of Electronic Science and Technology of China, Chengdu, China, in 2013, the M.S. degree in microelectronics from Fudan University, Shanghai, China, in 2016, and the Ph.D. degree in electrical engineering from the University of Virginia, Charlottesville, VA, USA, in 2021.

Prior to joining Yale University, New Haven, CT, USA, he was a Post-Doctoral Fellow with the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Champaign, IL, USA, from 2021 to 2023. He is currently a Post-Doctoral Associate of electrical engineering at Yale University. His research interests include analog/digital/mixed-signal integrated circuits and systems for intelligent Internet-of-Things (IoT) edge, in-memory/neuromorphic computing for edge artificial intelligence (AI), energy harvesting and power management units, and systems-on-chip for ultra-low-power IoT applications.

Dr. Li was a recipient of the IEEE International Symposium on Circuits and Systems (ISCAS) Best Paper Award in 2020 and a Winner of the IEEE SSCS 2019-2020 International Student Circuit Contest. He also serves as a reviewer for IEEE JOURNAL OF SOLID-STATE CIRCUITS, IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS I: REGULAR PAPERS, IEEE TRANSACTIONS ON VERY LARGE-SCALE INTEGRATION SYSTEMS, and IEEE TRANSACTIONS ON BIOMEDICAL CIRCUITS AND SYSTEMS.