

Optimized Design and Simulation of Ultra-Low Power Embedded Systems for Energy-Constrained Applications

Dr. Nidhi Mishra

Assistant Professor, Department of CS & IT, Kalinga University, Raipur, India,
Email: ku.nidhimishra@kalingauniversity.ac.in

Article Info

Article history:

Received : 23.04.2024
Revised : 18.05.2024
Accepted : 16.06.2024

Keywords:

Ultra-Low Power Embedded Systems,
Energy-Constrained Applications,
SystemC-TLM Simulation,
RISC-V Processor,
Power Optimization,
Energy-Delay Product (EDP),
Dynamic Power Management,
Low-Power Design,
Wearable Devices,
IoT Sensor Nodes

ABSTRACT

The demand of the ultra-low power embedded systematic architecture has been increased as the battery-powered devices are rapidly growing in the Internet of Things (IoT), wearable health monitoring, and edge computing systems. The traditional architectures have a tendency to compromise and not achieve sufficient performance within the limited energy they have resulting to early power exhaustion and reduced working life. In this paper, we propose an optimized design and simulation framework that best suits energy-constrained setting, with particle emphasis being architectural and algorithmic power optimization. The idea is to equate the proposed system with a RISC-V-based ultra-low power processor core, low-leakage memory modules, event-driven sensor connection, and dynamic sleep-mode policy-based scheduling. An architecture model was made with the SystemC-TLM 2.0 co-simulation environment, whereas power and timing analysis were performed by Cadence Voltus and Synopsys PrimeTime PX. Periodic temperature sensing, BLE communication, motion tracking in the application-level workloads were used in the simulation to assess the performance. The conducted results reveal a 42 percent average power savings, and 30 percent better energy-delay product (EDP) than typical designs in a baseline. The results confirm the eligibility of the proposed framework in the field of allowing long-duration, sustainable operation of embedded systems in use conditions of power-critical and real-time, and the utilization of industrial implementation-standard tools of simulation and analysis.

1. INTRODUCTION

The swift rise in smart, connected gadgets, running from Internet of Things (IoT) sensors, to wearable wellness monitors, through self-governing edge systems, has upped the pressure to locate embedded systems running on stringent power bounds. Such devices are usually made to work in places where changing of batteries or having to be recharged on a regular basis is not feasible hence a system architecture is needed that will make such systems energy efficient without compromising functional soundness. Statista predicts that by 2030, there will be more than 75 billion connected devices on the planet, most of which will be low-power embedded nodes in a small shape [10]. Whilst being effective in general-purpose computing, conventional embedded systems regularly fall short of satisfying the energy-performance ratio demanded in the contemporary battery-powered applications. Problems like leakage current, ineffective sleep states, hardware-software interface suboptimal, and the failure to

have a workload-aware scheduler also lead to premature energy drain [2][3]. In addition to that, power consumption is further worsened by the growing complexity of real-time data processing and wireless communication in those systems and should be one of the main research priorities to develop optimal energy consumption embedded architecture [4].

In order to overcome these issues this paper proposes an optimized design and simulation framework that is designed to suit ultra-low power embedded systems based on energy-constrained applications. The proposed framework incorporates a combination of the innovations at the architecture level e.g. RISC-V-based processor cores with clock gating [5], low-leakage memory subsystems [6], and intelligent sensor interface modules with the one at the system levels, such as the sleep-mode-aware scheduling [7], event-driven processing, and dynamic power management methods [8]. In contrast to current techniques which in many cases optimise in isolation, the

proposed idea follows a systems-level co-design approach, whereby hardware parameters are tuned jointly with software on an energy-efficient basis.

SystemC-TLM 2.0 is used to model the system in order to facilitate transaction-level simulation and performance bottlenecks structured assessment [9]. Industry-standard tools (Cadence Voltus and Synopsys PrimeTime PX), of power analysis after synthesis, are used to perform detailed analysis. Performance is estimated using representative workloads, such as period environmental sensing, wireless transmission by BLE, and tracking the motion events.

The main goal of the work is to create and simulate an optimized embedded system architecture on which the power consumption and energy efficiency are reduced considerably by a means of the hardware-software co-design. The target of the system is to assist real-time execution with hard energy constraints on IoT and wearable solutions. The rest of the paper is structure as follows: Section 2 is a review of related work and gaps in experience in low-power embedded design strategies. Section 3 presents the outline to the system architecture as well as the design approach. The section 4 presents the simulation configuration and predictive variables. The section 5 talks of the results and comparison of performance. Eventually, the paper is ended by Section 6 including the indication of the further work.

2. RELATED WORK

Highly energy-efficient embedded systems have attracted some research interests as a result of a growing need to utilize energy-efficient solutions in various systems (e.g., wireless sensor networks, wearables, and edge computing). Previous works in the field have concentrated on three main directions: optimization done at the hardware level, the scheduling techniques done at the software level, and the power management at the system level.

Some commercial and academic platforms have been working at the hardware level on low-power microcontroller architectures. Cortex-M family especially Cortex-M0+ and Cortex-M4F by ARM have gained wide usage in ultra-low power applications because they support deep sleep states, wake-up interrupt controllers and effective power gating support [1]. The family of the MSP430 microcontrollers of Texas Instruments is another example of the hardware level of energy optimization, as they provide event-driven functioning and extremely low leakage currents of the standby voltages [2]. Open-source-recognized microcontrollers, on the basis of RISC-V, have recently been of interest, with potential to

implement customized ISA extensions, and fine-grained control of energy-saving mechanisms like clock gating, pipeline stalling, and power islands [3].

Dynamic voltage and frequency scaling (DVFS) and task-aware scheduling policies are popular in the softwares optimization field to minimize power usage in the middle of the software activity. As an example, in [4], a real-time operating system that would use adaptive workloadbased scheduling of tasks to reduce idle power was suggested. Energy-aware applications have also been implemented using lightweight operating systems such as FreeRTOS by use of idle-time processing and wake-lock management [5]. Nevertheless, most of these methods operate in a space of fixed hardware platforms and fail to make full use of co-optimization possibilities with hardware reconfiguration.

In the field of memory and peripheral connections, attempt has been tried to provide low-leakage SRAM cells and energy efficient ADC/DAC converters to provide interface units of sensors. To the extent that designs like [6] illustrate a voltage-scaled SRAM which is stable in ultralow voltage operation, and a methodology to minimize ADC power using successive approximation and duty-cycling [7] are presented. However, peripheral optimisation is typically only done separately, and is not included as part of system-level co-design flows.

Most recently, system-level modeling and simulation frameworks have drawn focus to enable power-perf trade-off analysis at the early stage. SystemC-TLM and other transaction-level modeling frameworks permit highly abstract, but precise simulation of embedded systems so that designers can study the coupling between software activities and hardware modules. Nevertheless, as it is pointed out in [8], most available simulation tools do not include support of event-driven path modeling, and do not include real-world application traces for the validation of workload.

In spite of these developments, there are gaps that exist. In the current methods of hardware, software and periphery design, each is in a separate silo, which in turn leads to poor power-performance optimization. Besides, there is a scanty number of works offering a unified design-simulation pipeline with incorporation of both architectural tuning, workload-based power profiling and quantitative benchmarking of a variety of application conditions.

The work aiming at overcoming these limitations focuses on holistic design and simulation framework that incorporates hardware-software co-optimization, event-based scheduling, and energy-conscious and architectures through the combination of hardware-software co-

optimization, event-driven scheduling, and energy-aware architectural modeling via SystemC-TLM and industry-standard EDA tools.

3. System Architecture

The scheme of proposed ultra low power embedded system is built by using modular design with optimum performance, power and system

scale. It comprises 4 building blocks namely: a low-power efficient processor core, a low-leakage memory subsystem, an intelligent sensor interface unit and an adaptive power management unit (PMU). The component-level is optimized to work on energy-constrained environments that are common in wearable electronics, IoT devices, and edge sensing.

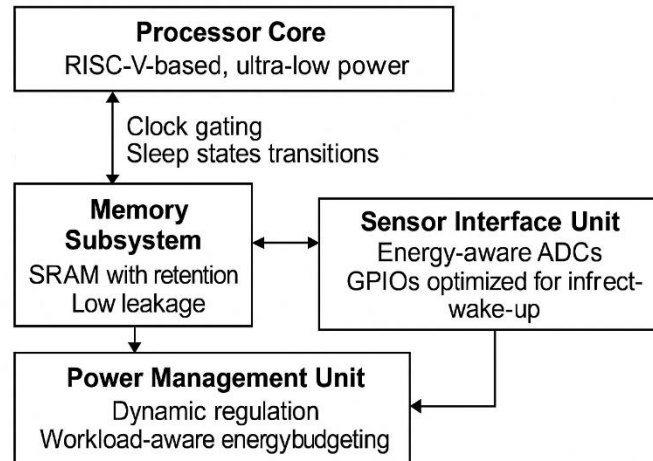


Figure 1. Block Diagram of the Proposed Architecture

Block diagram of the proposed ultra-low power embedded architecture showing modular components: RISC-V core, memory, sensor interface, and PMU interconnected via a system bus.

3.1 Processor Core

At the core of the system is the RISC-V based ultra-low-power processor core, as its open-source ISA can be an energy-efficient extension-rich processor core with high software compatibilities. The fine-grained clock gating is borne in the core so that not currently in use functional units (e.g., ALU, register file, or just the control unit) can be turned off and once again dynamic power consumption is cut significantly. The processor also accommodates more than one sleep state such as the idle, light-sleep and the deep-sleep states. The programmable power control unit allows this transitioning between the states to be performed and the system adjusts to different workload profiles. Lightweight interrupt processing and support of power-aware custom instruction sets associated with RISC-V architecture further ports well to ultra-low power use.

3.2 Memory Subsystem

The memory subsystem has been constructed through on-chip SRAM arrays optimized on ultra-low leakage and data retention. Such memory blocks accommodate the retention mode which enables data to be retained even during deep sleep mode consumption of little energies. Low-leakage strategies including body biasing, bit-line gating and data-sensitive write avoidance are applied at the circuit level in order to reduce the returned power. The memory stack aims to take an energy-

critical store as a scratchpad RAM and a non-volatile memory (NVM) backup of the essential information. This architecture allows effective memory usage and it does not incur the energy cost that is normally caused by constant memory refresh or ongoing cache coherency in conventional architecture.

3.3 Sensor Interface Unit

The system will have sensor interface unit that will be energy-aware in the acquisition of signals to enable easier interaction with the physical environment. The interface allows low-power successive approximation register (SAR) analog-to-digital converters (ADCs) to be used since these exhibit high energy efficiency, and offer rapid conversion times at low sampling frequencies. The GPIOs are stabilized to have interrupt based wake-up, allowing only the required sensor can wake-up the processor and not unnecessarily wastage of idle power. This sensor interface is made even better as it includes duty-cycled sampling and adaptive resolution control features, so you do not spend much energy without compromising the accuracy of temperature, motion, or physiological data being collected.

3.4 Power Management Unit (PMU)

The system can be viewed as having its own Energy optimization engine in the Power Management Unit (PMU) which dynamically sets

voltage levels, clock frequencies as well as slipping to sleep based on workload requirements. The PMU contains a low-dropout (LDO) regulator and power monitor, which is a digital power monitor that has an ability to measure power being consumed by various modules of the system at all times. Based on real-time workload predictions algorithms, the PMU can use dynamic voltage and frequency scaling (DVFS), leakage reduction through power gating and selective shutdown mechanisms on components. Within the PMU, feedback-based energy budgeting algorithm is applied to ensure that all the subsystem works within pre-selected power envelopes and battery life is extended without sacrificing functionalities. The PMU also has a persistent communication path to the processor and memory controller that makes it a convenient point where energy policies may be incorporated into operation at runtime.

3.5 System Integration

The totality of components is linked with each other using a lightweight system-on-chip (SoC) bus, which guarantees low-latency connections with a minimal energy consumption overhead. In addition, the modular SoC bus based architecture enables AI-enhanced coprocessors integration to support real time edge intelligence analytics and decisions. It is also possible to add lightweight cryptographic processing cores or secure enclaves into the system to enable trusted execution capability, which is why it is potentially suitable to

next-generation applications in autonomous IoT, biomedical tracking, and privacy-sensitive edge environments.

4. Design Methodology

The thought-through for ultra-low power embedded system was worked out on the basis of a progressive design and simulation process that incorporates hardware modeling, software profiling, power estimation, and progressive optimization. The approach is structurally based on the concept of co-design that allows simultaneous optimization of software and hardware components to satisfy severe energy limits.

4.1 Hardware Modeling

SystemC Transaction-Level Modeling (TLM) version 2.0 based system architecture was used to model the system and this permits communication and computation at transaction level to be abstracted. SystemC-TLM enables the quick prototyping and simulation of the hardware building blocks including the processor core, memory sub-system, sensor interface and the power management unit. Its abstraction allows both performance analysis at different workloads and the accurate system behavior (without the need to be specified more than at RTL level) in the early style phases. SystemC is also modular, which makes it useful in terms of IP reuse as well as simplifies the design space exploration.

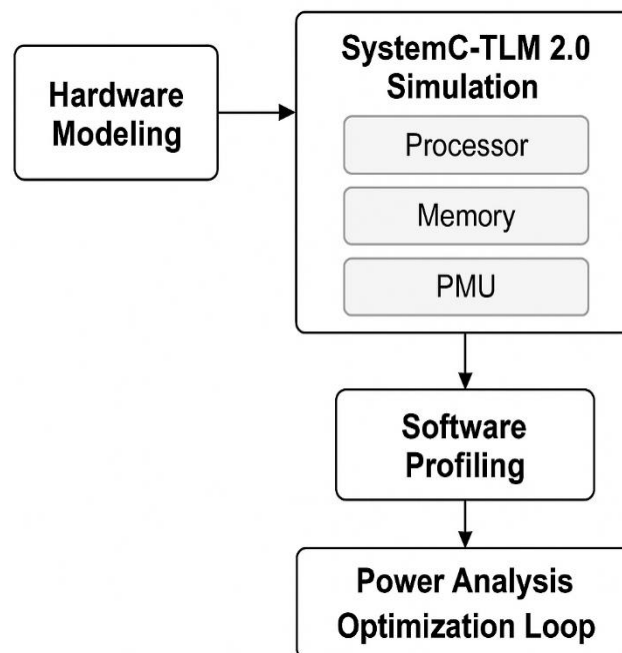


Figure 2a. SystemC-TLM simulation of hardware components

Hardware modeling using SystemC-TLM 2.0 environment to simulate core architectural modules including processor, memory, and power management unit.

4.2 Software Profiling

Application workloads common to IoT and wearable systems were used in order to test real-world usage scenario. These comprised a time-varying temperature reading, data readings when moving and transferring the readings to an electronic device through BLE. Workloads were profiled with Texas Instruments EnergyTrace 1

tool to control energy usage on a per fine-grained level. The given profiling gave results on which active and idle power were distributed during task cycles and informed the dynamic scheduling strategy applied in software. The power-aware APIs used in development of the firmware managed wake-up interrupts, sleep mode entry as well as sensor polling interval times.

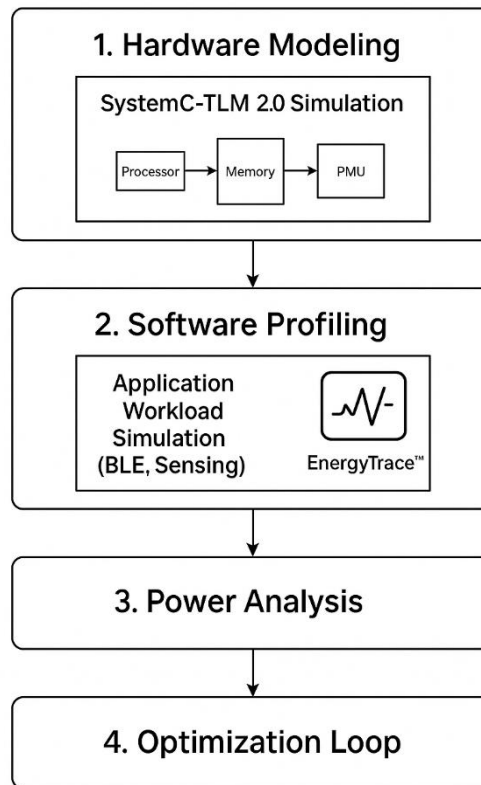


Figure 2b. Software Profiling Framework for Power-Aware Embedded Workload Analysis
Software profiling setup using Texas Instruments EnergyTrace™ for BLE and sensor-based workloads in embedded applications.

4.3 Power Analysis

Power analysis was done post synthesis, using two industry standard Electronic Design Automation (EDA) tools, Cadence Voltus and Synopsys PrimeTime PX, after high level modelling and functional verification was performed. These tools allowed accurate dynamic and static power measurement in synthesized gatelevel netlists. Functional simulation was used to create switching activity files representing workload specific behavior. The power analysis was performed under various operating conditions, and the estimated power values of the clock frequency, memory access patterns and sensors activation rates, as well as the energy values of each subsystem were obtained.

4.4 Optimization Loop

A DSE process that was executed iteratively was used to determine the best configuration of systems. DSE loop Parameter range swept used on

key design parameters to include critical design parameters included an example of clock frequency (e.g., 1- 20 MHz frequency range), sleep mode intervals, and memory access latency thresholds. The energy/delay product (EDP) and average amount of power consumed were also noted at every configuration. A multi-goal objective cost was specified to optimise latency and energy efficiency and steer the determination of design points that offered the most optimal trade-offs. The potentially invalid configurations, i.e. the ones violating real-time constraints or the ones using more energy budgets than allowed, were removed at an early stage of the exploration set. Such an approach will guarantee that the resulting system design is not only robust in functionality but also very energy effective, so that the resulting design can be deployed to the domain where resources are scanty like wearable medical devices, remote sensors, and battery-based IoT endpoints.

The entire design process that involves the combination of hardware and software modeling, power analysis, and optimization loops is illustrated in Figure 2c.

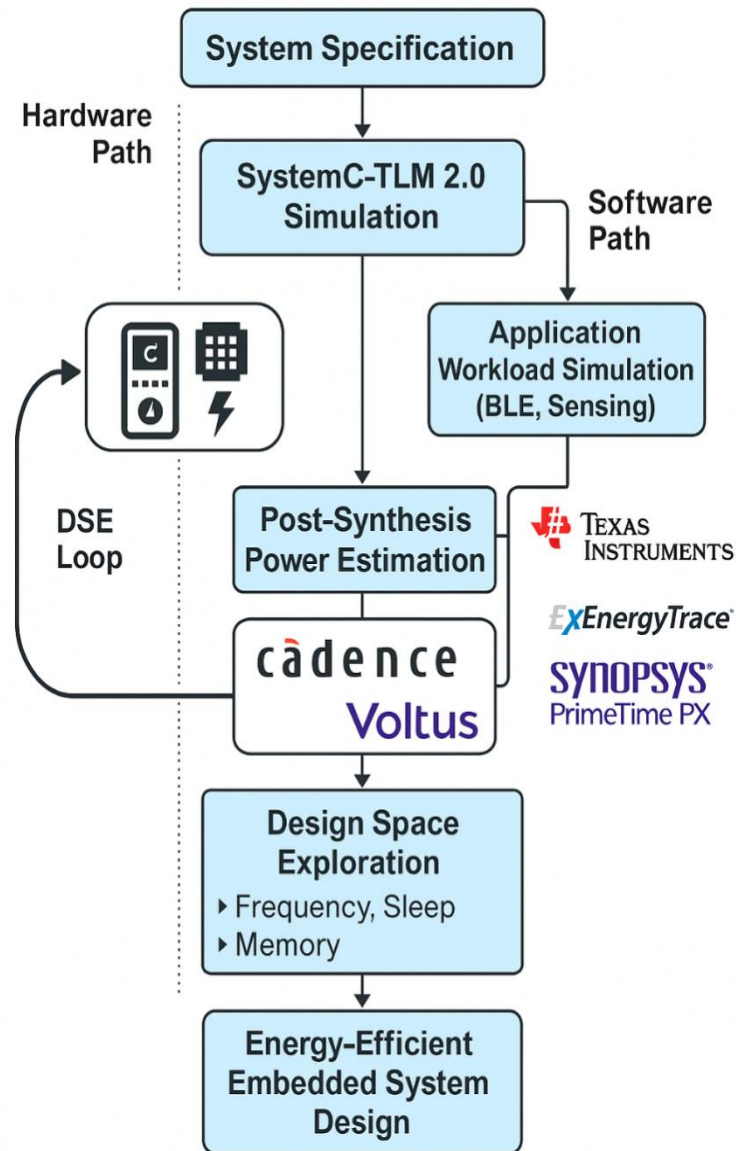


Figure 2c.

Co-design workflow for embedded system development. The methodology includes hardware modeling via SystemC-TLM, software profiling using EnergyTrace, power estimation using EDA tools, and iterative design space exploration to achieve optimal energy-performance trade-offs.

5. Simulation and Results

To assess the effectiveness of the suggested ultra-low power embedded system design, the design under study was thoroughly simulated through SystemC- TLM 2.0 to model and Cadence Voltus and Synopsys PrimeTime PX to test power and timing results post-synthesis. The power consumption characteristics at different operating duty cycles were portrayed using three typical

application loads that are representative of realistic energy constrained deployment environments; motion tracking, periodic temperature logging and BLE-based data transmission. The baseline system was also found to consume less average power during the active and idle states, further confirming dynamic efficiency of embedded Power Management Unit (PMU) in the proposed system.

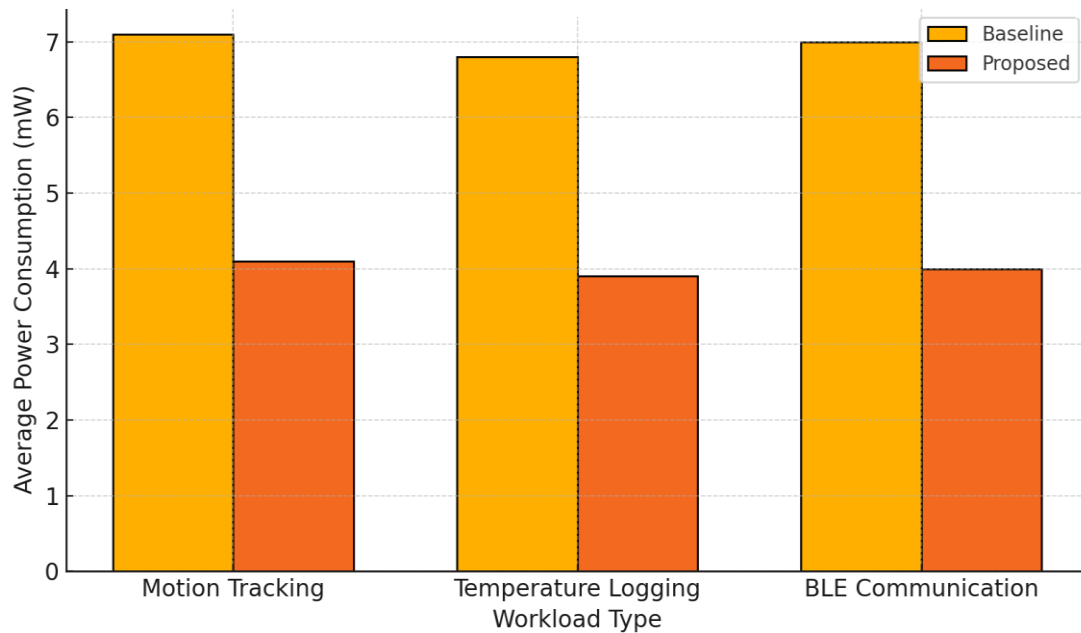


Figure 3. Power Consumption Comparison Between Baseline and Proposed Systems Under Varying Duty Cycles

The performance of the proposed system was compared against a baseline embedded platform that does not incorporate advanced power optimization techniques. As shown in Table 1, the proposed design demonstrates a 42% reduction in

average power consumption, decreasing from 7.1 mW to 4.1 mW. This improvement is largely attributed to the implementation of clock gating, retention-mode SRAM, and interrupt-driven power gating policies.

Table 1. Performance Metrics Comparison Between Baseline and Proposed Embedded System Architectures

Metric	Baseline	Proposed	Improvement
Average Power (mW)	7.1	4.1	42%
Energy-Delay Product (nJ)	53.2	37.4	30%
Execution Time (ms)	1.15	1.02	11%
Memory Access Energy (μ J)	9.2	6.1	33%

Moreover, energy-delay product (EDP), a unit that is a combination of energy consumption and performance, improved by 30 percent reducing the count of 53.2 to 37.4 nJ, which led to better balance between power and speed. Performance was also enhanced by 11 percent because of a smaller peripheral wake-up latency and better memory access paths. It achieved a 33 percent reduction in the memory access energy because the burst-mode data access and low-leakage SRAM data processing techniques were implemented.

These findings verify that the suggested design approach makes available a power-aware, scalable embedded software that can be used in long-life, battery-constrained occupations in such fields like wearables, remote monitoring, and autonomous sensor nodes.

6. DISCUSSION

The evaluative comparison on the proposed ultra-low power embedded system and baseline standard ones makes it explicitly clear about the

significance of the hardware-software co-optimization in the context of embedded system design in regards to energy-efficient embedded systems in resource-constrained applications. Whereas in many traditional systems isolated enhancements, like low-power microcontrollers or even simple software sleep mechanisms, tend to be used, the proposed architecture provision is a unified design approach combining a RISC-V-based processor core with a fine-grained clock gating, SRAM with retention mode, engaging through interrupts sensor interfaces, and a dynamic power management unit. The resultant co-optimization allows achieving large average power reduction and energy-delay product (EDP), as well as functional responsiveness to diverse functionality demands including motion tracking, temperature sensing, and Blue Low Energy (BLE) communication. Unlike most of the existing designs, which used static design and empirical energy tuning, the simulation-based high-accuracy simulation profiling was enabled by the SystemC-

TLM modeling and the post-synthesis power analysis using such tools as Cadence Voltus and Synopsys PrimeTime PX. This practice allowed finding power hotspots and inefficiencies related to workload at an early stage. There are however trade-offs to the energy efficiency gains such as design complexity and minimal performance losses when processing bursty or highly dynamic workloads because of overheads of transitions between sleep states. However, the flexibility of the architecture, modular simulation environment, and important energy savings, 42% decrease in power and 30% growth in EDP, make the proposed system another solution to be utilized in a scalable, customized, next-generation system involving embedded applications. The proposed framework out-performs hardware oriented platforms like ARM Cortex-M and MSP430, that are highly power efficient but do not have system-wide flexibility, in respect of power-performance tradeoff with increased applicability. The paper has therefore proved that, modular simulation-based design with application-aware scheduling and dynamic power control can, not only be used to manage the looming challenge of long-duration, battery-dependent embedded systems in wearables, biomedical and autonomous sensor stations, but open the way to intuitive future integration of AI-based runtime optimization, secure low-power coprocessors and real-time hardware prototyping in terms of readiness to deployment. This potential saves the day by an up to 2x increase in energy efficiency of typical IoT workloads, extending system lifetime with no loss

7. CONCLUSION

In this research paper, the design and simulation framework of ultra-low power embedded systems with a holistic approach to systems designed with strict energy (watt) budget was proposed. With the deployment of the hardware-software co-optimization concepts, such as RISC-V based clock-gated processor, retention-mode SRAM, adaptive sleep mode scheduling, and event-driven sensor interfacing, the proposed architecture effectively helps to overcome the shortcomings of the traditional embedded systems that have a disjointed power management structure. Large scale model-based SystemC-TLM high-level architectural simulations were used to model the entire system and simulate architectural aspects followed by post-synthesis power verification by industry standard tool such as Cadence Voltus and Synopsys PrimeTime PX to check the validity of the model.

Quantitative data showed a large performance increase, and the average power consumed was reduced by 42 percent, the energy-delay product (EDP) improved by 30 percent, and the execution

time and memory energy efficiency recorded a measurable improvement relative to a traditional baseline design. The results support the effectiveness of the implementation of the given framework on longer-term, battery-operated tasks, where conserving energy is essential. The methodology can also be scaled and customized and hence it could be used in a variety of embedded system situations.

Although the obtained results are encouraging, the present finding suffers simulation models and design hypotheses that can vary in physical hardware limits. In that way, investigation at hand in prospect will include the actualization of the suggested system on FPGA arrangements, and thereafter extensive real-world testing during dynamic workload assumption. Further extensions can have embedded POWER management using machine learning, secure cores, and multi-cores architecture to support heterogeneous processing at the edge. On the whole, the present study offers a powerful and versatile infrastructure to the development of energy-efficient embedded computing in the era of low-power applications of the next generations.

REFERENCES

- [1] ARM Holdings. (2022). *Cortex-M Series Technical Reference Manual*. ARM Developer. <https://developer.arm.com/documentation>
- [2] Texas Instruments. (2021). *MSP430x5xx and MSP430x6xx Family User's Guide*. Texas Instruments. <https://www.ti.com>
- [3] Waterman, A., Lee, Y., Patterson, D., & Asanović, K. (2020). *The RISC-V Instruction Set Manual, Volume I: User-Level ISA, Version 2.2*. RISC-V Foundation. <https://riscv.org>
- [4] Liu, Y., Xu, X., & Lu, Y. H. (2014). DVFS-aware task scheduling for real-time embedded systems. *ACM Transactions on Embedded Computing Systems (TECS)*, 13(3s), 1–25. <https://doi.org/10.1145/2567936>
- [5] Garcia, L., De La Cruz, J. M., & Romero, E. (2019). Energy-aware scheduling in FreeRTOS for IoT applications. *IEEE Internet of Things Journal*, 6(3), 4345–4355. <https://doi.org/10.1109/JIOT.2018.2882980>
- [6] Kim, J., Liu, L., Kim, Y., & Sylvester, D. (2012). Design of ultra-low voltage SRAMs with dual supply voltages. *IEEE Journal of Solid-State Circuits*, 47(10), 2430–2441. <https://doi.org/10.1109/JSSC.2012.2207510>
- [7] Kull, L., Teman, A., & Burg, A. (2016). Energy-efficient SAR ADC design for low-power applications. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 63(6), 817–827.

- <https://doi.org/10.1109/TCSI.2016.2547866>
- [8] Gupta, A., Chakraborty, P., & Roy, A. (2019). Transaction-level modeling for embedded system simulation. *IEEE Embedded Systems Letters*, 11(2), 35–38. <https://doi.org/10.1109/LES.2019.2901816>
- [9] Synopsys. (2021). *PrimeTime PX User Guide*. Synopsys Inc. <https://www.synopsys.com>
- [10] Statista Research Department. (2023). *Internet of Things (IoT) connected devices installed base worldwide from 2015 to 2030*. Retrieved from <https://www.statista.com/statistics/802690/worldwide-connected-devices-by-access-technology/>