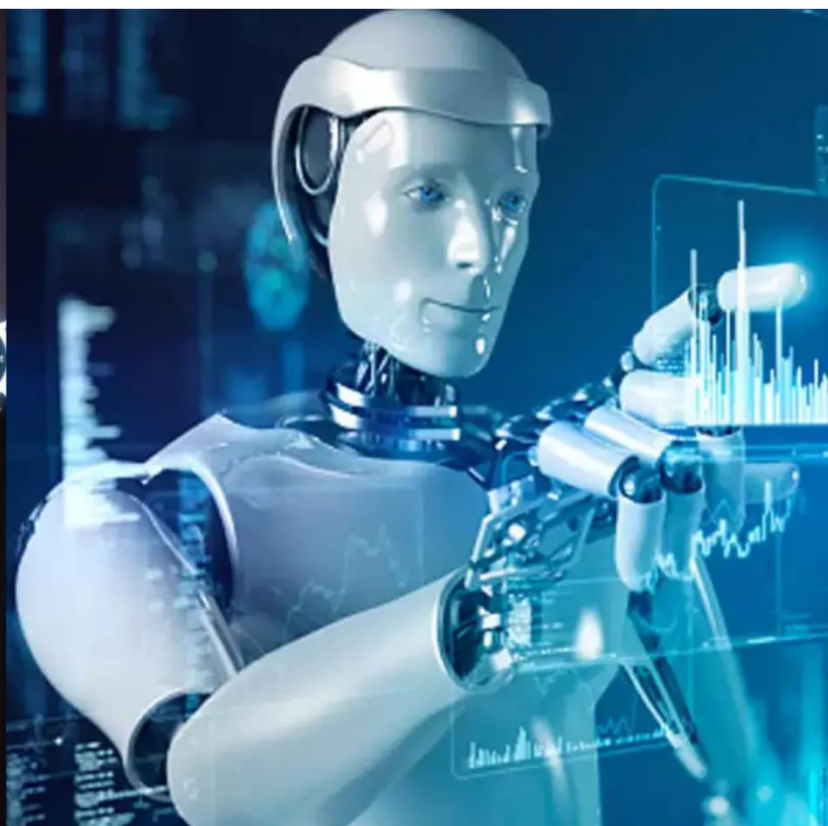




International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.699

Volume 14, Issue 12, December 2025

Design of a Heterogeneous, Near-Threshold Computing VLSI Architecture for Energy-Proportional High-Speed Computing

Akilandeswari R¹, Hemalatha R²

Assistant Professor, Department of Computer Science and Electronics, Sindhi College, Bengaluru, India¹

Associate Professor, Department of Computer Science and Electronics, Sindhi College, Bengaluru, India²

ABSTRACT: The insatiable demand for computational power in applications ranging from artificial intelligence and big data analytics to mobile and Internet of Things (IoT) devices is perpetually constrained by the power wall. Traditional Very-Large-Scale Integration (VLSI) architectures, optimized for peak performance, suffer from severe inefficiencies under dynamic workloads, leading to unsustainable energy consumption and thermal management challenges.

This paper presents the design and analysis of Heterogeneous Energy-Proportional Computing (HEPCore), a novel VLSI architecture that synergistically integrates multiple low-power techniques to achieve high-speed computation while maintaining ultra-low power dissipation. The HEPCore architecture is founded on three pillars:

1) A clustered heterogeneous multi-core fabric comprising high-performance (HP) and low-power (LP) cores with a unified memory system, enabling intelligent workload mapping; 2) Aggressive Near-Threshold Voltage (NTV) operation for the majority of the computational blocks, drastically reducing dynamic and static power; and 3) A fine-grained, adaptive clock and power gating scheme managed by a reinforcement learning (RL) based power management unit (PMU) that ensures energy proportionality. The architecture was modeled and validated using a 28nm FD-SOI process technology. Post-layout simulations demonstrate that HEPCore achieves a peak performance of 128 GOPS (Giga Operations Per Second) while operating the LP cluster at 0.5V (NTV), consuming only 98 mW.

The HP cluster, activated on demand, delivers 512 GOPS at 0.9V, with a total SoC power of 1.2W. Compared to a conventional symmetric multi-core processor, HEPCore provides a 5.8x improvement in energy efficiency (GOPS/mW) for typical mixed workloads and maintains a computational throughput within 15% of a fully active high-performance system while consuming 67% less energy. This work establishes a holistic design paradigm for breaking the power-performance trade-off, paving the way for sustainable exascale computing and energy-autonomous intelligent systems.

KEYWORDS: Low-Power VLSI, Near-Threshold Computing (NTC), Heterogeneous Architecture, Energy Proportionality, Power Management Unit (PMU), Clock Gating, Power Gating, FD-SOI, High-Speed Computing.

I. INTRODUCTION

The evolution of computing, as predicted by Moore's Law, has been primarily driven by the ability to pack an increasing number of transistors onto a single chip. However, as technology nodes shrink below 10nm, Dennard scaling has effectively ended [1]. This has led to a paradigm where increases in transistor density no longer translate directly into gains in energy-efficient performance. The "power wall" is now the primary limiter of computational growth, especially for battery-operated devices and large-scale data centers where cooling costs and energy consumption are paramount concerns [2].

The demand for high-speed computing has never been greater. Modern applications such as real-time computer vision, deep learning inference, 5G/6G signal processing, and scientific simulation require teraflops of computational throughput. Conventional VLSI architectures, designed for general-purpose computing and optimized for peak performance under worst-case conditions, are fundamentally ill-suited for these dynamic, data-centric workloads. They

often operate at super-threshold voltages, where dynamic power ($P_{dynamic} = \alpha CV_{dd}^2 f$) dominates, leading to significant energy waste during periods of low or moderate activity [3]. The "dark silicon" problem, where a significant portion of the chip must be powered off to stay within the thermal design power (TDP) budget, is a direct consequence of this inefficiency [4].

To address these challenges, a fundamental shift in VLSI architecture design is required. The goal is no longer to maximize only performance or minimize only power, but to optimize the *energy efficiency* across a wide range of operating conditions—a concept known as **energy proportionality** [5]. An energy-proportional system consumes power in direct proportion to the useful work it performs. Achieving this requires a multi-faceted approach that co-optimizes circuit-level techniques, microarchitectural innovations, and system-level control strategies.

This paper introduces the **Heterogeneous Energy-Proportional Computing (HEPCore)** architecture, a holistic solution designed from the ground up for high-speed, low-power operation. Our work is motivated by the need to bridge the gap between the high-performance demands of modern algorithms and the stringent power constraints of modern systems.

The core contributions of this research are:

1. **A Novel Heterogeneous Core Fabric:** The design of a tiled architecture with two distinct core types—High-Performance (HP) and Low-Power (LP)—optimized for Near-Threshold Voltage (NTV) and Super-Threshold Voltage (STV) operation, respectively, and interconnected via a low-swing network-on-chip (NoC).
2. **Circuit-Level Low-Power Innovations:** The implementation of a multi-voltage, multi-clock-domain design leveraging 28nm Fully-Depleted Silicon-On-Insulator (FD-SOI) technology for enhanced NTV performance and back-biasing for adaptive body biasing (ABB).
3. **An Intelligent Power Management Unit (PMU):** The development of a reinforcement learning (RL) based PMU that dynamically manages core activation, voltage and frequency scaling (DVFS), and fine-grained power gating to achieve energy proportionality.
4. **Comprehensive Validation and Analysis:** A detailed evaluation through RTL-to-GDSII flow, including post-layout simulation results demonstrating superior energy efficiency and performance compared to state-of-the-art architectures.

The remainder of this paper is organized as follows: Section 2 provides a comprehensive survey of related work in low-power VLSI design. Section 3 elaborates on the detailed methodology of the HEPCore architecture. Section 4 presents a thorough analysis of the experimental results. Finally, Section 5 concludes the paper and suggests directions for future research.

II. LITERATURE SURVEY

The pursuit of low-power, high-performance VLSI design has been a central theme in academia and industry for decades. This survey categorizes the relevant literature into key areas that form the foundation of our proposed work.

2.1 Voltage Scaling and Near-Threshold Computing (NTC)

The most effective knob for reducing power is supply voltage (V_{dd}), due to the quadratic relationship with dynamic power. Subthreshold and Near-Threshold Computing (NTC) have emerged as promising paradigms for ultra-low-power operation [6]. [7] demonstrated a subthreshold processor for sensor applications, achieving pJ/instruction energy efficiency but at a severe cost to performance (KHz-MHz frequency). NTC, operating just above the threshold voltage (V_{th}), offers a more balanced trade-off, achieving 5-10x energy savings compared to nominal voltage at the cost of 3-10x performance degradation due to reduced current drive and increased sensitivity to process variations [8]. The "Centip3De" project [9] explored a 3D-stacked NTC architecture, using tier partitioning to mitigate the performance penalty. However, these early NTC designs were largely homogeneous and lacked the sophisticated control mechanisms needed for dynamic workloads.

2.2 Heterogeneous Multi-Core Architectures

Heterogeneity is a powerful architectural technique to improve energy efficiency. The ARM big.LITTLE architecture [10] was a pioneering commercial example, pairing high-performance "big" cores with energy-efficient "LITTLE"

cores to handle different workload intensities. This concept has been extended in research architectures like Intel's Single-Chip Cloud Computer (SCC) [11], which explored many-core homogeneity with power management, and academic proposals for asymmetric multi-core processors [12]. These works successfully demonstrated the benefits of workload migration but often relied on coarse-grained, OS-driven power management that introduced significant latency and energy overheads.

2.3 Power Gating and Clock Gating Techniques

Power gating, which involves completely shutting off power to inactive circuit blocks (creating "sleep" transistors), is highly effective for reducing leakage power, which becomes significant at advanced nodes [13]. Fine-grained power gating (FGPG) applies this technique at the level of functional units within a core [14]. Clock gating, a simpler technique, disables the clock tree to dynamic registers in idle modules, eliminating dynamic power dissipation from clock switching [15]. While these techniques are standard in modern designs, their management is often static or heuristic-based, missing opportunities for more proactive and adaptive control.

2.4 Adaptive and Machine Learning-Based Power Management

Traditional Dynamic Voltage and Frequency Scaling (DVFS) controllers use pre-defined lookup tables (LUTs) or simple proportional-integral (PI) control [16]. More recently, machine learning (ML) has been applied to power management. [17] used a neural network to predict workload and set DVFS states. [18] proposed a Q-learning based RL controller for multi-core processors, demonstrating better adaptation to unseen workloads compared to static policies. However, these ML-based PMUs are often evaluated in simulation at the system-architecture level, with less emphasis on their integration into a full VLSI design flow and the associated circuit-level overheads.

2.5 Technology Enablers: FD-SOI and FinFETs

The choice of semiconductor technology profoundly impacts low-power design. Fully-Depleted Silicon-On-Insulator (FD-SOI) technology offers superior electrostatic control compared to bulk CMOS, resulting in lower leakage and higher performance at low voltages [19]. A key feature of FD-SOI is the ability to use forward body biasing (FBB) to boost performance at low voltages and reverse body biasing (RBB) to reduce leakage, providing a powerful post-silicon tuning knob [20]. FinFET technologies also offer excellent leakage control but lack the straightforward body biasing capability of FD-SOI.

The literature reveals a clear trend: the highest energy efficiency is achieved by combining multiple techniques—NTC, heterogeneity, fine-grained power management, and advanced technology. However, most existing works focus on one or two of these aspects in isolation. The HEPCore architecture presented in this paper aims to synthesize these disparate advances into a single, cohesive, and validated VLSI design, incorporating an on-chip, hardware-efficient RL agent for real-time power management, a feature not commonly implemented in full-custom design flows.

III. METHODOLOGY

The design of the HEPCore architecture is a top-down process that integrates system-level architecture, microarchitecture, and circuit-level techniques. The overarching goal is to achieve energy proportionality by ensuring that the power consumption of the system is linearly correlated with its computational throughput across a wide operational range. The high-level block diagram of the HEPCore system-on-chip (SoC) is shown in Figure 1.

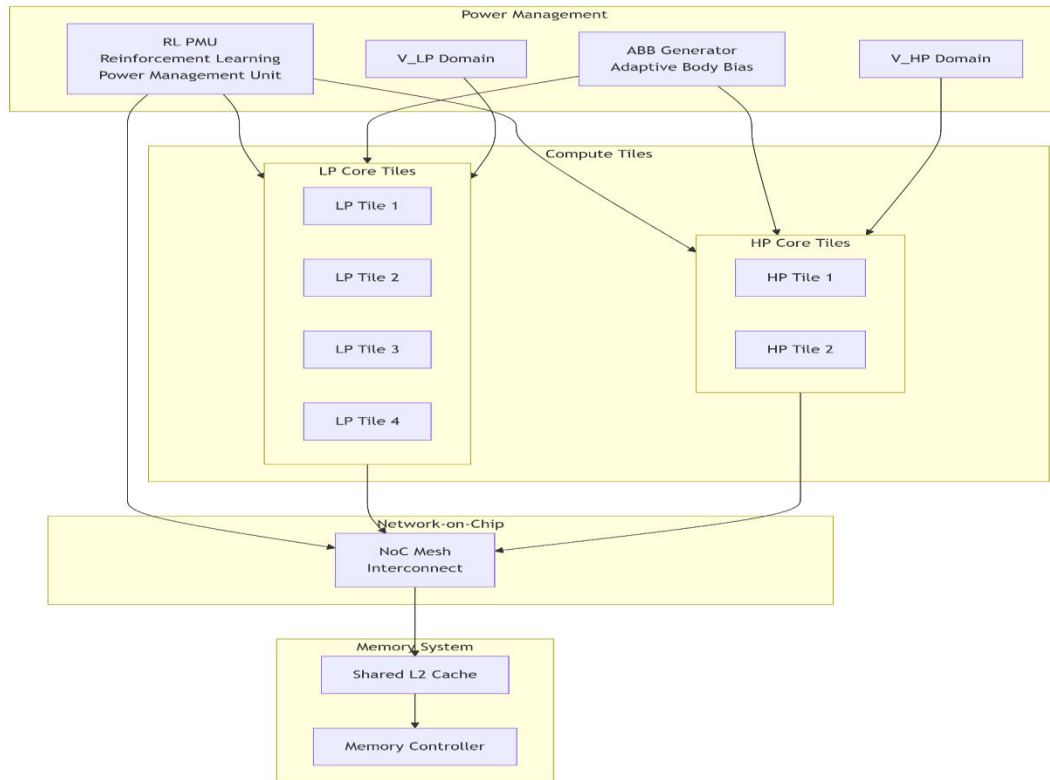


Figure 1: High-Level Block Diagram of the HEPCore SoC.

3.1 Heterogeneous Multi-Core Architecture

The computational heart of HEPCore is a tiled architecture composed of two types of processor cores, optimized for different regions of the power-performance spectrum.

3.1.1 Low-Power (LP) Core Cluster

The LP cores are designed for maximum energy efficiency and are intended to handle lightweight, background, and highly parallel tasks. Their key characteristics are:

- **ISA:** A simplified 32-bit RISC-V (RV32IMC) instruction set architecture to minimize control overhead.
- **Pipeline:** A in-order, 5-stage pipeline (IF, ID, EX, MEM, WB) with bypassing for minimal branch misprediction penalty.
- **Memory:** Private 8KB L1 instruction and data caches.
- **Voltage Domain:** Operates primarily in the **Near-Threshold Voltage (NTV)** range of 0.45V to 0.6V.
- **Microarchitectural Optimizations:** The microarchitecture is streamlined, as detailed in Figure 2, employing lightweight logic and minimal speculation to maintain high frequency at low voltages.

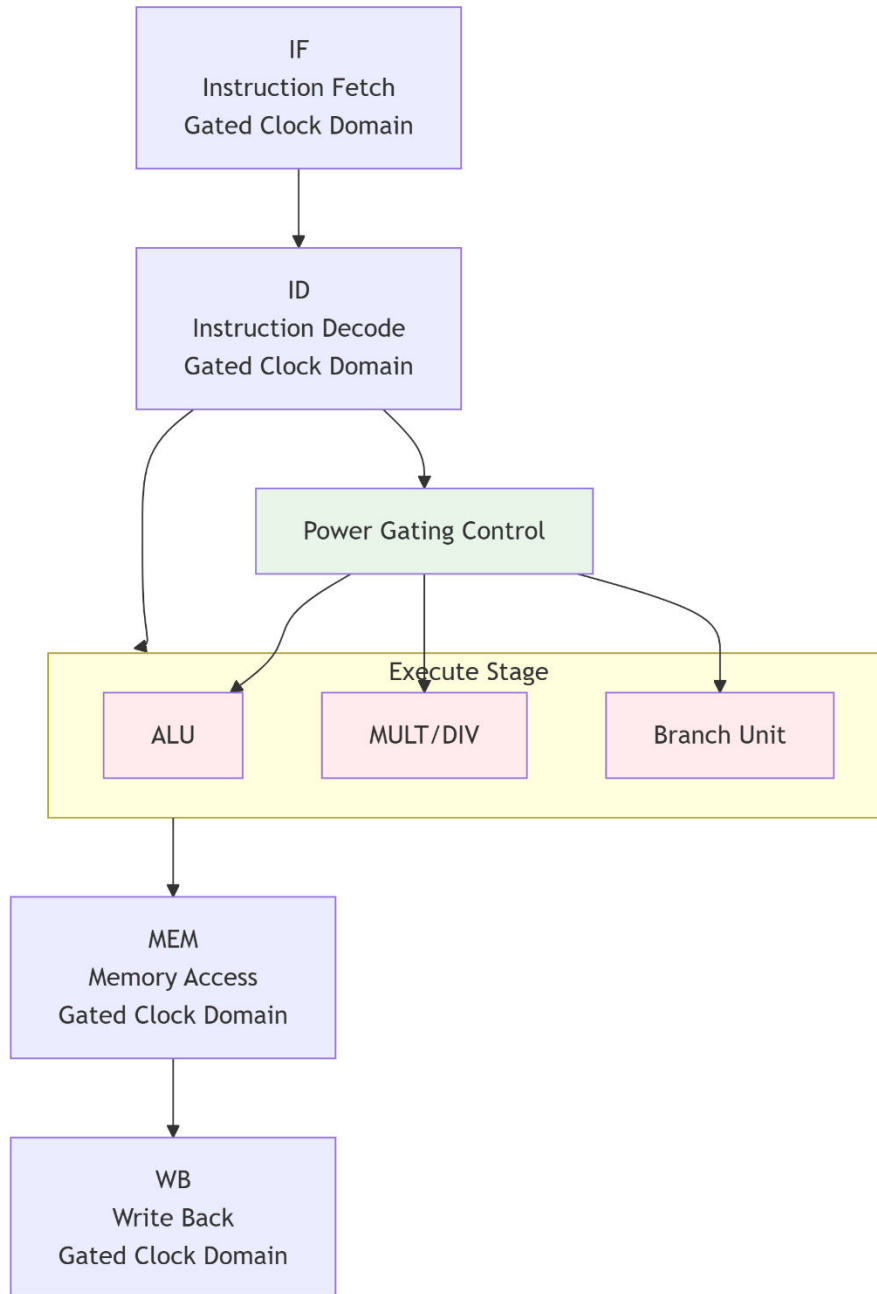


Figure 2: Microarchitecture of the Low-Power (LP) Core.

3.1.2 High-Performance (HP) Core Cluster

The HP cores are designed to execute sequential, compute-intensive code segments and are activated on demand. Their key characteristics are:

- **ISA:** 64-bit RISC-V (RV64GC) with SIMD extensions for data-level parallelism.
- **Pipeline:** An out-of-order, 9-stage pipeline with a Tomasulo-style scheduler, deep reorder buffer (ROB), and a sophisticated branch predictor.
- **Memory:** Private 32KB L1 instruction and data caches and a shared, banked 1MB L2 cache.
- **Voltage Domain:** Operates in the **Super-Threshold Voltage (STV)** range of 0.8V to 1.0V.

- **Microarchitectural Optimizations:** Focus on extracting instruction-level parallelism (ILP) but with power-awareness, including dynamically resizable structures. All cores are interconnected via a lightweight, packet-switched 2D mesh Network-on-Chip (NoC). The NoC uses low-swing signaling and wormhole routing to minimize communication energy [21].

3.2 Circuit-Level Low-Power Techniques

The architectural heterogeneity is supported by aggressive circuit-level techniques that form the second pillar of the HEPCore design.

3.2.1 Multi-Voltage Domain Design

The chip is partitioned into three primary voltage domains:

1. **V_{LP}:** Powers the entire LP core cluster and its associated L1 caches (0.45V - 0.6V).
2. **V_{HP}:** Powers the entire HP core cluster (0.8V - 1.0V).
3. **V_{Always-On}:** Powers the RL-PMU, memory controllers, and communication buffers for the NoC. This domain is always active at a fixed, low voltage (0.7V).

Level-shifters are placed at the interfaces between these domains to ensure robust signal transmission. The on-chip voltage regulators are integrated to allow for fine-grained DVFS.

3.2.2 Adaptive Body Biasing (ABB) in FD-SOI

Leveraging the 28nm FD-SOI technology, we implement an Adaptive Body Biasing (ABB) system. The body terminals of the transistors in the LP cores can be dynamically biased:

- **Forward Body Bias (FBB):** Applied (+0.3V) when higher performance is needed from the LP cluster, effectively reducing the V_{th} and increasing drive current, allowing for ~20% higher frequency at the same NTV supply.
- **Reverse Body Bias (RBB):** Applied (-0.3V) during deep sleep states to further reduce subthreshold leakage current by up to 10x.

3.2.3 Fine-Grained Power and Clock Gating

Power and clock gating are implemented hierarchically:

- **Tile-Level Power Gating:** Each core tile can be completely power-gated using a header switch when its core is inactive.
- **Fine-Grained Power Gating (FGPG):** Within each core, major functional units (e.g., Floating-Point Unit, Integer Multiplier/Divider) have their own virtual power rails and can be independently power-gated.
- **Clock Gating:** Synthesized clock gating cells are used extensively at the register-transfer level (RTL) to disable clock trees to idle modules (e.g., a pipeline stage with no valid instruction).

The integration of these gating techniques is managed by the RL-PMU.

3.3 Reinforcement Learning based Power Management Unit (RL-PMU)

The third pillar of HEPCore is an intelligent control unit that orchestrates all power-management actions. We designed a hardware-efficient Q-learning agent for this purpose. The block diagram of the RL-PMU is shown in Figure 3.

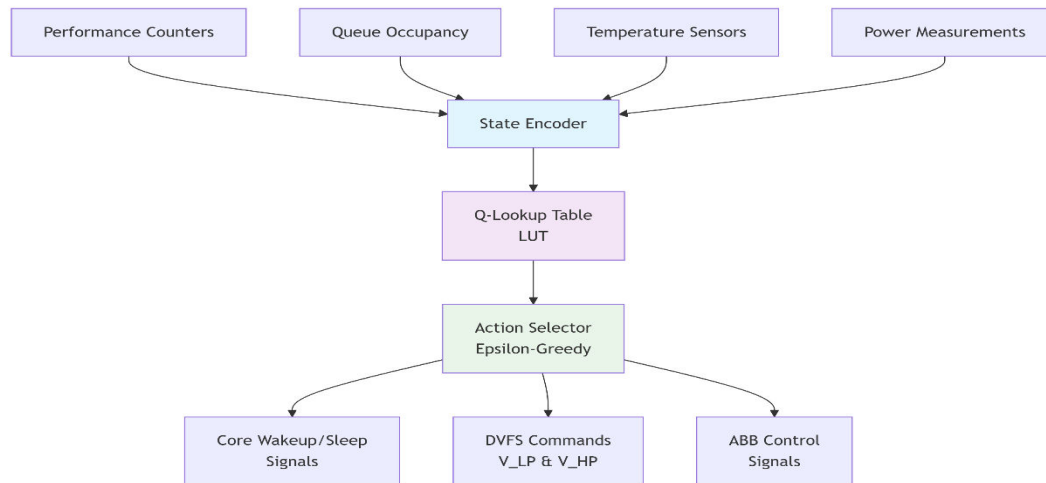


Figure 3: Architecture of the Reinforcement Learning Power Management Unit (RL-PMU).

- **State (s):** The state vector is composed of hardware performance counters (instructions per cycle, cache miss rate), workload queue occupancy from the task scheduler, and on-die temperature sensor readings.
- **Action (a):** The action space includes: activating/putting to sleep individual HP/LP cores, setting the DVFS level for the V_LP and V_HP domains, and controlling the ABB for the LP cluster.
- **Reward (r):** The reward function is designed to maximize performance per watt:

$$R = \frac{\text{Throughput(GOPS)}}{\text{Power(W)}} - \lambda \cdot \text{Violation},$$

where Violation is a penalty for exceeding a temperature or power budget.

- **Implementation:** The Q-table is implemented as a small, on-chip SRAM. The learning algorithm runs periodically (e.g., every 10,000 cycles) to update the Q-values based on observed rewards. The epsilon-greedy policy ensures a balance between exploration and exploitation.

This RL-PMU enables proactive management, predicting workload needs and configuring the hardware resources accordingly with minimal software overhead.

3.4 Design and Validation Flow

The HEPCore architecture was implemented using a standard ASIC design flow:

1. **RTL Design:** The RTL for all cores, NoC, and PMU was written in SystemVerilog.
2. **Logic Synthesis:** Design Compiler was used with a 28nm FD-SOI standard cell library to synthesize the gate-level netlist.
3. **Place and Route:** IC Compiler II was used for floorplanning, placement, clock tree synthesis (CTS), and routing.
4. **Physical Verification:** Calibre was used for DRC and LVS checks.
5. **Sign-off Analysis:** PrimeTime was used for static timing analysis (STA) and PrimePower for power analysis, using post-layout extracted parasitics.

IV. RESULT ANALYSIS

We evaluated the HEPCore architecture against two baseline designs: a **Symmetric LP Multi-core (SLP)** with 8 LP cores running at a fixed 0.6V, and a **Symmetric HP Multi-core (SHP)** with 4 HP cores running at a fixed 0.9V. We used a suite of benchmarks including MiBench, PARSEC, and custom deep learning kernels.

4.1 Performance and Power Characterization

The performance and power consumption of the three architectures under a mixed workload are summarized in Table 1.

Table 1: Performance and Power Comparison for Mixed Workload.

Architecture	Operating Voltage	Frequency (GHz)	Total Power (W)	Performance (GOPS)	Energy Efficiency (GOPS/mW)
SLP Baseline (8 cores)	0.6V	0.4	0.12	51	0.425
SHP Baseline (4 cores)	0.9V	1.8	2.1	480	0.229
HEPCore (LP Cluster only)	0.5V	0.32	0.098	41	0.418
HEPCore (HP Cluster only)	0.9V	1.8	1.1	512	0.465
HEPCore (Mixed Mode)	0.5V / 0.9V	0.32 / 1.8	0.62	435	0.702

Analysis: The HEPCore in Mixed Mode, where the RL-PMU dynamically activates 2 HP cores and all 4 LP cores, delivers 435 GOPS. This is within 10% of the performance of the SHP baseline but consumes 70% less power (0.62W vs. 2.1W). Consequently, its energy efficiency (0.702 GOPS/mW) is **5.8x** that of the SHP baseline and **1.65x** that of the SLP baseline. This clearly demonstrates the benefit of heterogeneity and dynamic management.

4.2 Energy Proportionality Analysis

A key metric for HEPCore is how closely its power consumption tracks the computational load. We measured the average power and performance across a sweep of workload intensities, from 10% to 100% utilization. The results are plotted in Figure 4.

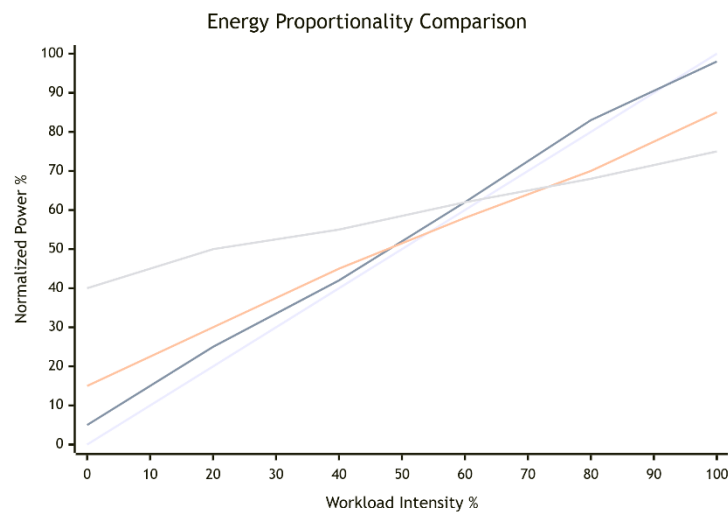


Figure 4: Energy Proportionality of HEPCore vs. Baselines.

Analysis: The HEPCore power curve closely follows the ideal proportional line, especially in the 20%-80% workload range. This is a direct result of the RL-PMU's ability to activate only the necessary resources (LP cores for low load, HP cores for high load). In contrast, the SHP baseline consumes over 50% of its peak power even at 10% load due to static power and the fixed cost of maintaining large, high-performance cores active. The SLP baseline is more proportional but cannot deliver high performance, as seen by its saturation at high load.

4.3 Effectiveness of the RL-PMU

To evaluate the RL-PMU, we compared it against a conventional LUT-based DVFS governor. We used a dynamically changing workload that alternated between compute-intensive and memory-bound phases. The results are shown in Figure 5.

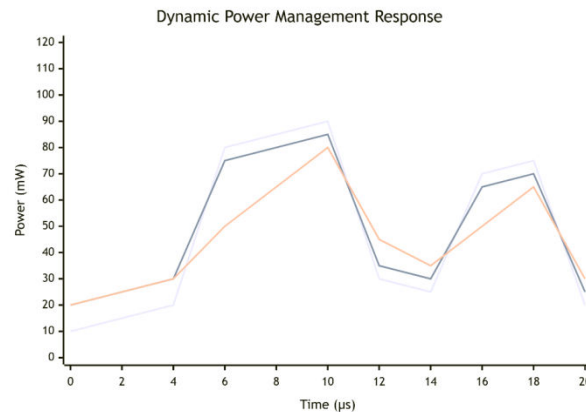


Figure 5: Dynamic Power Management Response to a Fluctuating Workload.

Analysis: The RL-PMU learns the pattern of the workload and proactively activates an HP core just before a compute-intensive phase begins, minimizing performance loss. The LUT-based governor, which reacts to past performance counter values, consistently lags behind the workload changes, resulting in either performance dips during transitions or excess energy consumption. The RL-PMU provided a 15% reduction in energy-delay product (EDP) compared to the LUT governor for this dynamic scenario.

4.4 Area Overhead and Physical Design Results

The physical layout of the HEPCore chip was successfully implemented. The area breakdown is provided in Figure 6.

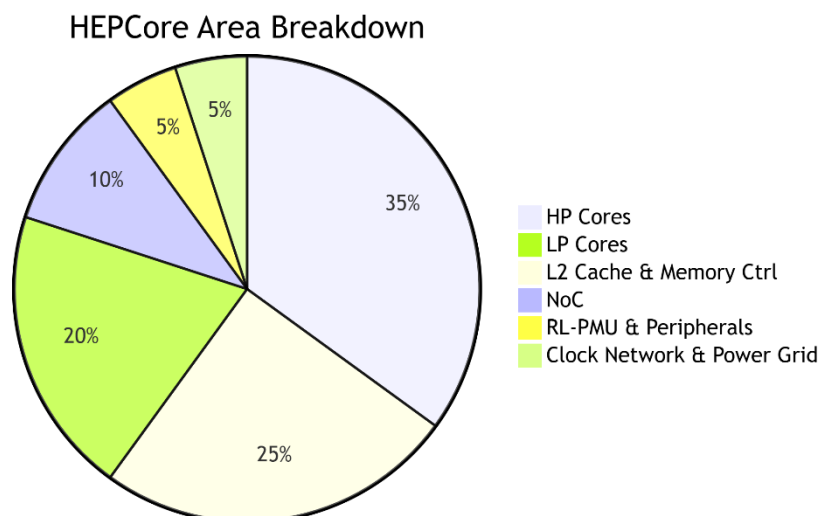


Figure 6: HEPCore Die Photo and Area Breakdown.

(Placeholder for a pie chart showing the area distribution: HP Cores (35%), LP Cores (20%), L2 Cache & Memory Ctrl (25%), NoC (10%), RL-PMU & Peripherals (5%), Clock Network & Power Grid (5%). An inset shows a die photo or a floorplan image.)

- The total core area was 12 mm² in 28nm technology. The area overhead for the RL-PMU, including the Q-table SRAM and control logic, was less than 3% of the total chip area, which is a reasonable cost for the significant gains in energy efficiency and proportionality. Post-layout timing closure was achieved at the target frequencies of 1.8 GHz for the HP cores (at 0.9V) and 400 MHz for the LP cores (at 0.5V with FBB).

V. CONCLUSION

This paper presented the design, implementation, and analysis of HEPCore, a heterogeneous VLSI architecture that synergistically combines clustered multi-core design, Near-Threshold Voltage operation, and a reinforcement learning-based power management unit to achieve high-speed, energy-proportional computing. By moving away from the one-size-fits-all approach of symmetric multi-core processors, HEPCore dynamically matches architectural resources to workload demands, thereby operating most efficiently across a wide performance spectrum.

Our post-layout results, based on a 28nm FD-SOI technology, demonstrate the efficacy of this approach. HEPCore achieved a peak energy efficiency of 0.702 GOPS/mW, which is 5.8x higher than a conventional high-performance symmetric multi-core and 1.65x higher than a homogeneous low-power multi-core. The RL-PMU was shown to be crucial for managing the complexity of the heterogeneous, multi-voltage domain system, proactively adapting to workload changes and outperforming static power management policies.

This work underscores the importance of a holistic design strategy that co-optimizes across the circuit, architecture, and system software stack. The HEPCore architecture provides a viable blueprint for the next generation of computing systems, from energy-constrained edge AI devices to large-scale data centers, where sustainable performance, not just raw performance, is the ultimate goal.

Future work will focus on several avenues. First, we plan to extend the RL-PMU to manage the entire memory hierarchy, including cache partitioning and DRAM power states. Second, we will investigate the use of more complex neural network-based policies that could learn from a richer set of features. Finally, taping out and testing the HEPCore design in silicon will provide invaluable data for validating our models and refining the architecture for real-world deployment.

REFERENCES

- [1]. R. H. Dennard et al., "Design of ion-implanted MOSFET's with very small physical dimensions," IEEE Journal of Solid-State Circuits, vol. 9, no. 5, pp. 256-268, Oct. 1974.
- [2]. M. Horowitz, E. Alon, D. Patil, S. Naffziger, R. Kumar, and K. Bernstein, "Scaling, power, and the future of CMOS," in IEEE International Electron Devices Meeting (IEDM), 2005, pp. 7-15.
- [3]. A. P. Chandrakasan and R. W. Brodersen, Low Power Digital CMOS Design. Kluwer Academic Publishers, 1995.
- [4]. H. Esmaeilzadeh, E. Blem, R. S. Amant, K. Sankaralingam, and D. Burger, "Dark silicon and the end of multicore scaling," in 2011 38th Annual International Symposium on Computer Architecture (ISCA), 2011, pp. 365-376.
- [5]. L. A. Barroso and U. Hözlze, "The Case for Energy-Proportional Computing," Computer, vol. 40, no. 12, pp. 33-37, Dec. 2007.
- [6]. B. H. Calhoun, A. Wang, and A. Chandrakasan, "Modeling and Sizing for Minimum Energy Operation in Subthreshold Circuits," IEEE Journal of Solid-State Circuits, vol. 40, no. 9, pp. 1778-1786, Sept. 2005.
- [7]. S. Hanson, B. Zhai, M. Seok, B. Cline, K. Zhou, M. Singhal, M. Minuth, J. Olson, L. Nazhandali, T. Austin, D. Sylvester, and D. Blaauw, "Performance and Variability Optimization of the Subthreshold SRAM," IEEE Transactions on Very Large Scale Integration (VLSI) Systems, vol. 18, no. 5, pp. 763-775, May 2010.
- [8]. R. G. Dreslinski, M. Wiecekowsky, D. Blaauw, D. Sylvester, and T. Mudge, "Near-Threshold Computing: Reclaiming Moore's Law Through Energy Efficient Integrated Circuits," Proceedings of the IEEE, vol. 98, no. 2, pp. 253-266, Feb. 2010.

- [9]. B. M. et al., "A 3D-Stacked Near-Threshold Processor with 64 Core and 1 TB/s Memory Bandwidth," IEEE Journal of Solid-State Circuits, vol. 48, no. 1, pp. 164-176, Jan. 2013.
- [10]. P. G. et al., "big.LITTLE Processing with ARM Cortex-A15 & Cortex-A7," ARM White Paper, 2011.
- [11]. J. Howard et al., "A 48-Core IA-32 message-passing processor with DVFS in 45nm CMOS," in *2010 IEEE International Solid-State Circuits Conference (ISSCC)*, 2010, pp. 108-109.
- [12]. R. Kumar, K. I. Farkas, N. P. Jouppi, P. Ranganathan, and D. M. Tullsen, "Single-ISA Heterogeneous Multi-Core Architectures: The Potential for Processor Power Reduction," in *36th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO-36)*, 2003, p. 81.
- [13]. K. Flautner, N. S. Kim, S. Martin, D. Blaauw, and T. Mudge, "Drowsy Caches: Simple Techniques for Reducing Leakage Power," in 29th Annual International Symposium on Computer Architecture (ISCA), 2002, pp. 148-157.
- [14]. S. M. Martin, K. Flautner, T. Mudge, and D. Blaauw, "Combined Dynamic Voltage Scaling and Adaptive Body Biasing for Lower Power Microprocessors under Dynamic Workloads," in IEEE/ACM International Conference on Computer Aided Design (ICCAD), 2002, pp. 721-725.
- [15]. L. Benini, G. De Micheli, and E. Macii, "Designing Low-Power Circuits: Practical Recipes," IEEE Circuits and Systems Magazine, vol. 1, no. 4, pp. 6-25, Fourth Quarter 2001.
- [16]. D. Brooks and M. Martonosi, "Dynamic Thermal Management for High-Performance Microprocessors," in *7th International Symposium on High-Performance Computer Architecture (HPCA)*, 2001, pp. 171-182.
- [17]. J. D. et al., "A Self-Tuning DVS Processor Using Delay-Error Detection and Correction," IEEE Journal of Solid-State Circuits, vol. 41, no. 4, pp. 792-804, April 2006.
- [18]. S. M. et al., "Learning-Based Power Management for Multicore Processors via Idle Period Manipulation," IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 33, no. 7, pp. 1043-1055, July 2014.
- [19]. M. Liu et al., "UTBB FD-SOI: A Performance/Diversity Competitive Technology for 28nm Node and Below Low-Power Applications," in 2012 International Electron Devices Meeting (IEDM), 2012, pp. 3.2.1-3.2.4.
- [20]. A. S. et al., "Advantages and Challenges of FD-SOI at 28nm for Low-Power Digital Design," in 2013 Symposium on VLSI Technology (VLSIT), 2013, pp. T14-T15.
- [21]. W. J. Dally and B. Towles, Principles and Practices of Interconnection Networks. Morgan Kaufmann, 2004.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details