Journal of
Electronics Technology Exploration

Homepage: htttps://shmpublisher.com/index.php/joetex
p-ISSN 3025-3470 | e-ISSN 3026-1066

A Comparative Review of Performance and Energy Efficiency in ARM and RISC-V SoCs

Revan Muhammad Zaki[1], Satrio Budi Pratama[2], Muhammad Dzaky Althaf[3], Christina Jenifer Sibuea[4], Much Aziz Muslim[5]
[1, 2, 3, 4, 5]Department of Computer Science, Universitas Negeri Semarang, Indonesia

Article Info

Article history:

Received May 04, 2026
Revised May 20, 2026
Accepted June 29, 2026

Keywords:

RISC-V
ARM
Energy efficiency
Perfomance benchmarking
System-on-chip

ABSTRACT

The demand for energy-efficient computing systems is growing as workloads on embedded devices, edge computing, and scientific computing continue to expand. In this context, a comparison between the ARM and RISC-V architectures is important because both are widely used in System-on-Chip (SoC) designs but have different performance and energy efficiency characteristics. This article presents a comparative review of various benchmark studies on the Odroid XU4, Rock960, and Nezha D1 platforms using the NAS Benchmark, TensorFlow Lite Benchmark, and OpenFOAM. The summarized results show that RISC-V implementations generally have lower average power consumption but do not always yield a better performance-per-watt ratio. For some workloads, ARM platforms still demonstrate superior throughput and energy efficiency, while RISC-V performance is heavily influenced by the number of cores, accelerator support, and the maturity of the software ecosystem. Therefore, ARM remains more competitive for high-performance workloads, whereas RISC-V remains promising for power-efficient applications and continues to evolve.

This is an open access article under the CC BY-SA license.

Corresponding Author:

Revan Muhammad Zaki,
Department of Computer Science,
Universitas Negeri Semarang,
Sekaran, Gunungpati, Semarang, Indonesia.
Email: lsntl@ccu.edu.tw
https://doi.org/10.52465/joetex.v4i1.686

1. INTRODUCTION

The demand for high-performance computing (HPC) is growing as various applications such as weather forecasting, drug discovery, and scientific simulations continue to evolve. However, this increase in computing power is accompanied by a massive demand for power. Modern HPC systems require high processing resources, but at the same time pose serious challenges related to operational costs, thermal management, and environmental sustainability [1], [2]. Therefore, energy-efficient computing has become an increasingly important area of research, both in embedded systems and in high-performance systems.

In this context, system designers are increasingly utilizing system-on-chip (SoC) platforms as a practical approach to improving efficiency. SoCs integrate key components such as CPUs, GPUs, memory, and other controllers onto a single chip, making them well-suited for applications that require compact size, low power consumption, and adequate performance. At the architectural level, the instruction set architecture (ISA) plays a crucial role because it determines how software is translated into instructions that can be executed by the

hardware. Two ISA families currently receiving significant attention are ARM and RISC-V. ARM has long been the dominant architecture in mobile and embedded devices due to its mature ecosystem, extensive hardware support, and excellent performance-per-watt characteristics [3]. In contrast, RISC-V has emerged as an open ISA that offers design flexibility, ease of customization, and no licensing fees, making it attractive for both academic and industrial development [3].

Several recent studies indicate that the advantages of ARM and RISC-V depend heavily on the type of workload and the quality of their supporting ecosystems. In AI inference applications and edge-based systems, RISC-V platforms can demonstrate low power consumption, but performance gains are often influenced by accelerator support and software maturity [4]–[6]. In HPC-oriented studies, RISC-V systems still face limitations in per-core performance, memory behavior, and toolchain optimization, although multi-core designs show potential for workloads that emphasize throughput [7], [8]. Meanwhile, ARM-based platforms continue to demonstrate competitive performance and, in many cases, better energy efficiency when performance and power are considered together [9]. These findings indicate that comparing power consumption alone is insufficient to assess the suitability of an architecture.

However, the existing literature still leaves significant gaps. Many studies evaluate ARM and RISC-V separately, focus on only one type of benchmark, or emphasize either performance or power without placing both within the same analytical framework. Furthermore, comparisons of representative embedded SoCs remain fragmented, particularly when the research objective is to integrate synthetic benchmarks, AI inference, and real-world simulation workloads into a single discussion. Therefore, it remains necessary to clarify how ARM and RISC-V differ across various workload types and platform classes, as well as under what conditions each architecture is better suited for use.

To address this gap, this article compares three representative SoC platforms: the Odroid XU4 and Rock960 as ARM-based platforms, and the Nezha D1 as a RISC-V-based platform. These three platforms were selected because they represent practical devices in the embedded and edge-computing domains with distinct performance and power profiles. The comparison is based on benchmark studies using the NAS Parallel Benchmarks, TensorFlow Lite Benchmark, and OpenFOAM, which collectively cover scientific computing, AI inference, and real-world simulation workloads [10]. The purpose of this article is to present a more structured comparative review of the performance and energy efficiency of ARM and RISC-V SoCs, as well as to identify the architectural and system-level factors that influence their relative advantages.

2. METHOD

2.1 Hardware Platforms

This study uses three System-on-Chip (SoC) platforms representing two different architectures: ARM and RISC-V. These three platforms were selected to compare performance and energy efficiency across several different types of benchmarks. The general specifications of the platforms used are summarized in Table 1.

Table 1. Hardware platform specifications

Platform	Processor	Frequency	Architecture	RAM	Power Login	Operating System
Odroid XU4	Exynos5422 Octa-core	2 GHz dan 1.4 GHz	ARMv7l	2 GB LPDDR3	20W	Ubuntu 20.04.5 LTS
Rock960	Cortex-A72 Dual-core + Cortex-A53 Quad-core	1.8 GHz dan 1.4 GHz	AArch64	4 GB LPDDR3	24W	Ubuntu 20.04.4 LTS
Nezha D1	Allwinner D1 Single-core	1.0 GHz	RISC-V64	1 GB DDR3	10W	Debian GNU/Linux 11

These three platforms were used to represent different system characteristics in terms of processor architecture, memory capacity, and power consumption.

2.2 Benchmark Testing

To measure the performance and energy efficiency of each platform, this study used three types of benchmarks: the NAS Benchmark, the TensorFlow Lite Benchmark, and OpenFOAM. These three benchmarks were chosen because they represent three different types of workloads: scientific computing, artificial intelligence inference, and real-world simulations.

NAS Benchmark

The NAS Benchmark is used to measure the mathematical computing capabilities of each platform. This test focuses on floating-point operations commonly used in HPC systems [11]. The primary metric used is FLOPS (Floating-Point Operations per Second), which is the number of floating-point operations that can be processed in one second. The higher the FLOPS value, the greater the platform's computational capability.

TensorFlow Lite Benchmark

The TensorFlow Lite Benchmark is used to test the platform's ability to run artificial intelligence models on embedded systems. In this test, the MobileNet model was used for image recognition tasks. The primary parameter analyzed was inferences per second—that is, the number of predictions the system can generate in one second [12].

OpenFOAM

OpenFOAM is used as a real-world simulation to test the platform's ability to handle heavy computational loads. OpenFOAM is Computational Fluid Dynamics (CFD) software used to simulate fluid flow and heat transfer [13]. In this study, two main simulations were used, namely: Motorbike simulation and RotorDisk simulation. Both simulations were used to evaluate each platform's ability to handle non-synthetic workloads.

2.3 Measurement Parameters

The parameters measured in this study include:

1. Power Consumption (Watts)

Used to measure the average electrical power consumed during the test.

2. Execution Time

Used to measure the time required for the platform to complete the benchmark.

3. Energy Consumption (Joules)

Measures the total energy used during the testing process.

4. FLOPS

Used to measure floating-point computational capability per second.

5. Inferences per Second

Used to measure AI inference capability on the TensorFlow Lite Benchmark.

6. Temperature

Measures the device temperature during the testing process.

7. Performance per Watt

Used to measure energy efficiency based on the ratio of performance to power consumption.

2.4 Prevention of Thermal Throttling

To maintain the consistency of test results, this study uses a Python script to monitor the device temperature during the benchmark. The system will only run the next test once the device's temperature returns to its idle state. This method is used to prevent thermal throttling a condition in which the processor automatically reduces performance due to excessively high temperatures. As a result, test results are more stable and objective across platforms.

3. RESULT AND DISCUSSIONS

This study evaluates three System-on-Chip (SoC) platforms the Odroid XU4, Rock960, and Nezha D1 to compare the performance and energy efficiency of ARM and RISC-V implementations, as reported in a previous study [11]. The evaluation was conducted using the NAS Benchmark, the TensorFlow Lite (TFLite) Benchmark, and OpenFOAM as a real-world application representation. The findings of this study were also compared with two comparative studies: an evaluation of AI inference engines between the Sipeed. Maixduino (64-bit RISC-V) and Raspberry Pi 4B with a Coral USB accelerator (ARM) [12], and a study testing the 64-core RISC-V Pioneer Milk-V platform in the context of server-scale High Energy Physics (HEP) computing [13].

3.1 Power Consumption Analysis

Previous research reported that the Nezha D1 consistently had the lowest average power consumption compared to the other two platforms, namely the Rock960 and Odroid XU4, indicating that the RISC-V implementation on the Nezha D1 is capable of operating with relatively low power requirements [11]. Meanwhile, Rock960 recorded the highest power consumption in most test scenarios, while Odroid XU4 fell in the middle [11].

A similar pattern was observed in another study [12], in which the RISC-V-based Maixduino consumed only about 0.8–1.1 watts during inference, significantly less than the Raspberry Pi 4B (ARM), which drew about 4.6–4.9 watts. This difference in power consumption is more than fourfold, making the RISC-V platform

far more energy-efficient in absolute terms. Nevertheless, as also demonstrated in another study [13], RISC-V power consumption characteristics are not always uniform across all device classes. On the Pioneer Milk-V server platform, idle power consumption reaches approximately 80 Watts, comparable to a single-socket server. However, the dynamic power swing (the difference between idle and full-load) is relatively smaller, at around 60 Watts compared to approximately 85 Watts on a mid-range x86 server. This indicates that RISC-V's power characteristics are highly dependent on the scale of implementation and the device class used.
As emphasized in previous research [11], although low power consumption is often associated with better energy efficiency, this relationship is not always linear. Power consumption merely describes the amount of energy used per unit of time and does not reflect the platform's overall capability to complete a workload.

3.2 Performance Analysis

System performance was measured based on the time required to complete each benchmark [11]. The Rock960 delivered the best performance with the fastest execution times in most NAS Benchmark and TFLite Benchmark tests. In contrast, the Nezha D1 recorded the longest execution times in nearly all test scenarios. The limited number of processor cores and lower operating frequency (1.0 GHz) are suspected to be the primary factors causing the platform's low performance. Meanwhile, the Odroid XU4 demonstrated fairly competitive performance and, in some cases, was able to approach the results achieved by the Rock960 [11].
Another study provides a more nuanced perspective in the context of AI inference [12]. Maixduino (RISC-V) can process 16.8 to 127.3 frames per second depending on model depth, while the Raspberry Pi 4B (ARM) produces 11.9 to 86.9 frames per second. Interestingly, Maixduino actually outperforms the Raspberry Pi 4B in some model configurations with shorter inference latency on lightweight models. These findings indicate that in the domain of AI inference with dedicated neural accelerators (KPUs), the performance advantage does not always lie with ARM. However, when the Raspberry Pi 4B is paired with a Coral USB accelerator (ARM + TPU), its performance increases significantly to reach 250–833 frames per second, demonstrating the importance of hardware accelerator support in enhancing system performance [12].

In the context of more demanding HPC workloads, another study tested the 64-core RISC-V Pioneer Milk-V platform using the ROOT, DB12, and Geant4 benchmarks [13]. The results of the single-threaded ROOT benchmark show that RISC-V still lags significantly behind AMD and ARM. This is due to RISC-V's per-core performance remaining below that of both x86 and ARM architectures on numerical and statistical workloads [13].

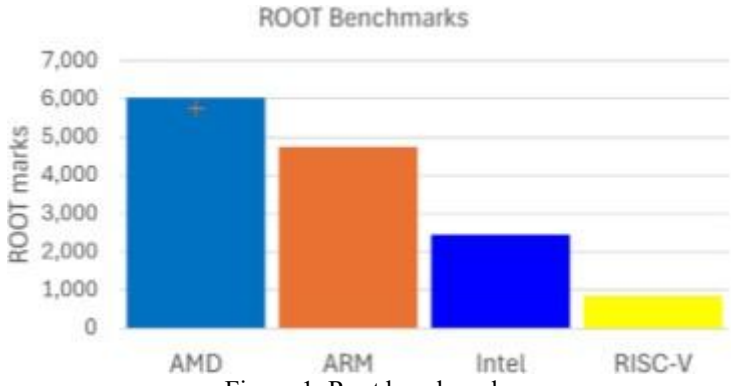

Figure 1. Root benchmarks

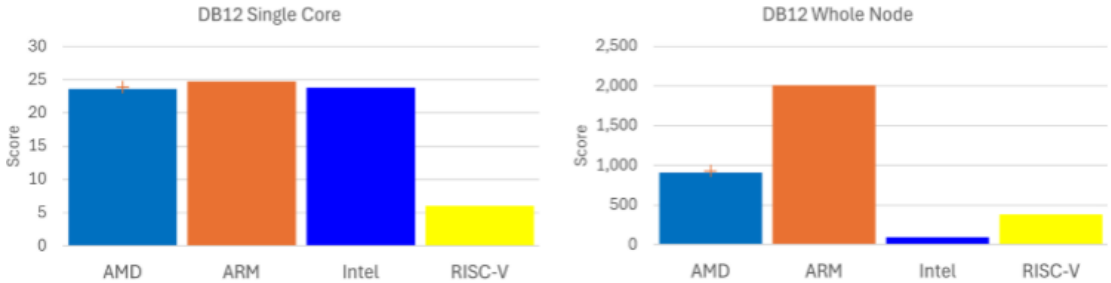

Figure 2. Comparison of DB12 single-core and whole-node scores across architectures

However, Figure 2 reveals a different dynamic in the DB12 benchmark. In single-core mode, a single RISC-V core does indeed perform significantly worse than an x86 or ARM core. However, thanks to the 64-core design of the Pioneer Milk-V, the RISC-V whole-node DB12 score is nearly four times higher than that

of a mid-range Intel i5 CPU, which has only four cores. This confirms that RISC-V's potential in high-throughput scenarios depends heavily on parallelism capabilities rather than just per-core speed.

3.3 Power Consumption Analysis

Calculating total power consumption based on a combination of power draw and execution time. Unlike the average power consumption results, the Nezha D1 actually showed the highest total energy consumption in most tests. This phenomenon occurs because the low power consumption cannot offset the long computation time required to complete the workload. Conversely, the Rock960 and Odroid XU4 have higher power consumption but can complete tasks in significantly less time, resulting in lower total energy usage [11].

The same dynamics were observed in a previous study [12]. Although Maixduino (RISC-V) consumes significantly less power, when measured in terms of energy efficiency (frames per second per Watt), the ARM platform with the Coral accelerator outperforms it dramatically, achieving energy efficiency of up to 175 frames per second per Watt on the lightweight model, compared to Maixduino's approximately 127 frames per second per Watt. This finding reaffirms that low power consumption does not always guarantee better energy efficiency if the resulting computational performance is insufficient [11].

3.4 Computational Throughput Analysis

In the NAS Benchmark, throughput is measured using FLOPS, while in the TFLite Benchmark it is measured based on the number of inferences per second. Previous research reported that the Rock960 consistently achieved the highest throughput values in both benchmarks, demonstrating its ability to handle computationally intensive workloads more effectively. The Odroid XU4 showed relatively competitive results, while the Nezha D1 produced the lowest throughput in nearly all tests [11].

In the context of HEP benchmarking, other studies have also reported that RISC-V lags behind AMD and ARM in terms of throughput on the events-per-second metric in Geant4 simulations [13]. Nevertheless, when measured per unit of energy (events per kWh), the gap between RISC-V and ARM narrows, indicating that the absolute throughput deficit does not fully represent RISC-V's actual energy efficiency [13].

3.5 Energy Efficiency Analysis

Previous studies calculated energy efficiency based on the performance-to-power consumption ratio: FLOPS per watt for the NAS Benchmark and inferences per second per watt for the TFLite Benchmark. The Odroid XU4 and Rock960 exhibit relatively competitive energy efficiency levels; in fact, in some tests, the Odroid XU4 demonstrated better efficiency despite its lower absolute performance. On the other hand, the Nezha D1 consistently yields the lowest energy efficiency values. Although its power consumption is the lowest, the low throughput results in a performance-to-power ratio that is not competitive [11].

Comparison with previous studies reinforces these findings. RISC-V energy efficiency depends heavily on the quality of the integrated accelerators and the maturity of its model compilation toolchain, not solely on the ISA architecture used [12]. Other research adds that the small dynamic power swing in RISC-V opens opportunities for optimization through power-gating techniques and workload-aware scheduling, so that RISC-V's energy efficiency potential at the server scale can still be significantly improved through software optimization [13].

3.6 Temperature Analysis

[11] Conducted temperature measurements to ensure that all testing proceeded without the influence of thermal throttling. The Nezha D1 maintained the lowest operating temperature across all benchmarks. The Rock960 exhibited higher temperatures but remained within safe operating limits, while the Odroid XU4 recorded the highest temperatures despite using an active cooling system. All devices operated below critical temperature thresholds, ensuring performance results reflect the devices' actual capabilities without distortion from thermal throttling. The lower thermal conditions on the Nezha D1 are also consistent with the findings in [12], where low-power RISC-V platforms tend to generate less heat, offering a tangible advantage for embedded and edge computing applications with limited cooling.

3.7 Validation Using OpenFOAM

To evaluate the relevance of the benchmark results to real-world applications, [11] used OpenFOAM (Computational Fluid Dynamics) with two main simulations, motorBike and rotorDisk. The results are consistent with the synthetic benchmarks: the Nezha D1 still has the lowest average power consumption but requires the longest execution time, the Rock960 demonstrates the best performance, and the Odroid XU4 falls in the middle. In terms of total energy consumption, the Nezha D1 again shows the highest value due to its longer computation time.

This validation approach using real-world workloads aligns with the methodology employed in other studies [13], which also utilized real scientific workloads (Geant4 ParFullCMS and AnaEx01). In both of these studies, the characteristic performance patterns observed across architectures in synthetic benchmarks were found to be consistent when tested on more complex real-world workloads. These findings reinforce the validity of benchmarking as a reliable predictor of performance [13].

3.8 General Discussion

Overall, the three studies reviewed confirm the same theme: ARM and RISC-V implementations have advantages in different aspects. [11] reports that the Nezha D1 (RISC-V) excels in power consumption and low operating temperatures, yet still lags in performance and energy efficiency, while the Rock960 (ARM) demonstrates the highest performance and the Odroid XU4 offers the most competitive performance-to-energy-efficiency balance.

There is a consistent common thread among the three studies: RISC-V, in its current implementations, excels in absolute power consumption, but has not yet been able to match ARM's energy efficiency when metrics are calculated comprehensively (performance per Watt). This gap varies depending on the context whether it is an embedded edge system [11], [12] versus an HPC server [13] and also depends on the availability of specialized accelerators and the maturity of the software ecosystem.

[13] Specifically notes that part of the RISC-V performance gap is not a permanent architectural barrier, but rather a result of immaturity, compiler toolchains, and a lack of architecture-specific optimizations—all of which are currently under active development. Therefore, as concluded in [11], platform selection must be tailored to application needs: for systems with power constraints, RISC-V can be an attractive option; however, for applications requiring high computational performance and optimal energy efficiency, the ARM platform still demonstrates a more significant advantage under current conditions.

4. CONCLUSION

Simulation results and analysis confirm that the Novelty Half-Bridge LLC Resonant Converter with Magnetizing Inductor and Hybrid Rectifier (NHB-LLCRC-MIHR) achieves significant improvements in efficiency, output voltage stability, and reductions in both switching and conduction losses compared to conventional converters. Operation at the resonant frequency (around 97.9 kHz) produces minimum impedance and maximum efficiency, reflecting the essential characteristics of LLC converters. The addition of *Lm* and a hybrid rectifier keeps the output voltage stable across a wide switching frequency range (50 kHz to 1 MHz). The resonant tank design and synchronous secondary MOSFET control support optimal energy transfer and minimize dead time during rectification. High DC conversion ratio efficiency persists due to low voltage drop in the hybrid rectifier and an optimized Lm, especially at high current levels. The proposed converter produces higher output voltage, lower voltage ripple, and a faster time to reach steady-state compared to the conventional converter. A larger Lm serves as a natural filter. Secondary MOSFETs deliver rapid and efficient rectification. The NHB-LLCRC-MIHR topology offers promising potential for precision power supplies, electric vehicle systems, and renewable energy sources requiring high efficiency and power stability. Recommended future work includes adaptive digital control for the hybrid rectifier and further transformer optimization to reduce physical size without sacrificing efficiency. Higher power testing and analysis of disturbance resilience in industrial environments remain essential. Experimental validation using a physical prototype and investigation of dynamic load effects on long-term performance remain important aspects for improvement.

REFERENCES

[1] World Health Organization, "Breast cancer." 2024.
[2] N. Harbeck and M. Gnant, "Breast cancer," *Lancet*, vol. 389, no. 10074, pp. 1134–1150, 2017.
[3] F. Bray *et al.*, "Global cancer statistics 2022: {GLOBOCAN} estimates of incidence and mortality worldwide for 36 cancers in 185 countries," *CA. Cancer J. Clin.*, vol. 74, no. 3, pp. 229–263, 2024.
[4] K. Kourou, T. P. Exarchos, K. P. Exarchos, M. V Karamouzis, and D. I. Fotiadis, "Machine learning applications in cancer prognosis and prediction," *Comput. Struct. Biotechnol. J.*, vol. 13, pp. 8–17, 2015.
[5] F. Pedregosa *et al.*, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.
[6] D. Delen, G. Walker, and A. Kadam, "Predicting breast cancer survivability: A comparison of three data mining methods," *Artif. Intell. Med.*, vol. 34, no. 2, pp. 113–127, 2005.
[7] M. D. Ganggayah, N. A. Taib, Y. C. Har, P. Lio, and S. K. Dhillon, "Predicting factors for survival of breast cancer patients using machine learning techniques," *BMC Med. Inform. Decis. Mak.*, vol. 19, p. 48, 2019.
[8] L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, pp. 5–32, 2001.
[9] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.*, vol. 13-17-Augu, pp. 785–794, 2016.
[10] H. He and E. A. Garcia, "Learning from imbalanced data," *IEEE Trans. Knowl. Data Eng.*, vol. 21, no. 9, pp. 1263–1284, 2009.
[11] N. V Chawla, K. W. Bowyer, and L. O. Hall, "SMOTE : Synthetic Minority Over-sampling Technique," vol. 16, pp. 321–357, 2002.

[12] A. Fernández, S. Garcia, F. Herrera, and N. V Chawla, "{SMOTE} for learning from imbalanced data: Progress and challenges, marking the 15-year anniversary," *J. Artif. Intell. Res.*, vol. 61, pp. 863–905, 2018.
[13] D. R. Cox, "The regression analysis of binary sequences," *J. R. Stat. Soc. Ser. B*, vol. 20, no. 2, pp. 215–242, 1958.
[14] T. M. Cover and P. E. Hart, "Nearest neighbor pattern classification," *IEEE Trans. Inf. Theory*, vol. 13, no. 1, pp. 21–27, 1967.
[15] T. Fawcett, "An introduction to ROC analysis," *Pattern Recognit. Lett.*, vol. 27, no. 8, pp. 861–874, 2006.