Adaptive Power Management Techniques in Multi Core VLSI Systems for Enhanced Energy Efficiency

Karthikeyan M, Aravinth P and Heiner A J
Department of Electronics and Communication Engineering, Saveetha Engineering College, Chennai, India

Keywords: Core-Level Power Gating, Adaptive Power Management, Heat Dissipation, Multi-Core VLSI Systems, Energy Efficiency.

Abstract: The paper proposes an adaptive power management system for multi-core VLSI designs that can optimize energy utilization and heat dissipation. Traditional static power management strategies, such as voltage scaling and core gating, do not account for dynamic workloads, resulting in energy holes and thermal stress. The proposed system overcomes those shortcomings by utilizing workload monitoring in real-time, dynamic voltage and frequency scaling (DVFS), core-level power gating, and a machine learning-based workload prediction technique. In the high-performance end, the system outperforms the overall power by 35%. Furthermore, it achieves decreased thermal dissipation, with maximum temperatures as low as 75 °C compared to 80-85 °C in state-of-the art systems. The performance comparison results reveal that the proposed system has near execution speeds under a variety of workloads, indicating its potential for energy-efficient, scalable VLSIs at the expense of power-constrained applications.

1 INTRODUCTION

In the rapidly changing world of VLSI (Very Large-Scale Integration) systems, the rapid growth of a variety of applications on the one hand, and on the other, the emergence of multi-core architectures with more cores integrated on a single chip, an efficient power management system has become a necessity (Dinakarrao, et al. , 2020). Power and heat are the most difficult concerns for modern multicore VTLSI architectures. To illustrate, concrete workloads are dynamic and unpredictable, rendering static power management strategies like voltage scaling and core gating ineffective, resulting in energy waste and performance degradation (Dinakarrao, et al. , 2020). Even when the load varies, conventional systems continue to use power in a static form, resulting in severe underutilization, increased energy loss, increased thermal stress, and, ultimately, reduced system level efficiency (Ranjbar, Singh, et al. , 2023). These fundamental constraints underscore the need for more transudative power management systems that can alter power at runtime in response to workload demands (Li, et al. , 2024). Inspired by the design of VLSI systems for mobile and data center applications that require increased power efficiency and/or performance. With servers increasingly relying on multi-core processors, it is critical that power be distributed precisely (Li, Tian, et al. , 2020). A real-time monitoring system, dynamic task allocation, and DVFS can help achieve the goal of balancing performance and energy usage 14. Furthermore, the sector has maintained its emphasis on sustainability and cost-effectiveness, which necessitates that developers reduce heat generation and eliminate the need for ambient air-cooling systems (Kim, Kim, et al. , 2020). However, such static solutions have limitations in meeting these objectives, highlighting the need for adaptive, responsive systems that can capture the challenges of modern computational and software environments (Ansari, Salehi, et al. , 2020). Illustrated View of Proposed Model is show in Figure 1.

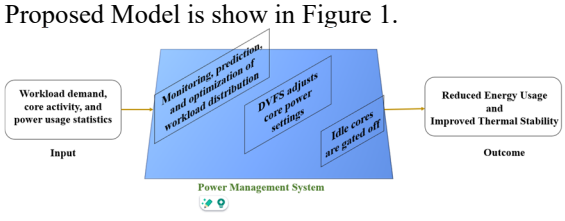

Figure 1: Illustrated View of Proposed Model

M, K., P, A. and A J, H.
Adaptive Power Management Techniques in Multi Core VLSI Systems for Enhanced Energy Efficiency.
DOI: 10.5220/0013583800004664

In *Proceedings of the 3rd International Conference on Futuristic Technology (INCOFT 2025) - Volume 1*, pages 692-700
ISBN: 978-989-758-763-4

The paper primarily contributes an adaptive power management technique for multi-core VLSI systems, which addresses the limits of the most basic static power management strategies. Provides a system capable of dynamic voltage-frequency scaling and core-level power gating. It combines real-time workload monitoring and machine learning-based prediction techniques to improve power and thermal efficiency without losing performance. The automatically adjusts the power management settings to reduce unnecessary energy loss while also reducing thermal stress lowering overall energy use depending on both cumulative energy efficiency and present job requirements. Furthermore, the system has a feedback nature, in which the proposed model acts for the power distribution model and changes the distribution based on various information about the system's state, and in turn, the model constantly adapts itself when the system's condition changes abruptly and compensates for any unknown changes in the workload. These combine hardware techniques such as DVFS and core-level power gating with software workload prediction using the adaptive power management subsystem. In addition to coarser-level power tuning, the technique uses machine learning to detect workload trends and adjust power accordingly, even before the workload changes. It has proved that the proposed system can achieve an exceptional level of energy efficiency while consuming the least amount of power while still providing good performance and a substantial temperature decrease compared to competing systems. The proposed architecture's system scale allows it to be employed in a variety of VLSI structures, making it suitable for a wide range of applications, from low-power mobile devices to high-end computing machines. The remainder of the paper is organized as follows: Section II analyses recent work on power management for multi-core VLSI systems and discusses the shortcomings of current approaches. Section III describes the design and implementation of the proposed adaptive power management system, including an introduction to the various blocks and their operation. Section IV describes the performance evaluation results, including comparisons of power consumption, execution time, and heat dissipation between existing systems. Finally, Section V summarizes the work and discusses potential possibilities for improving the operability of the described system.

In summary, an adaptive power management system for multi-core VLSI architectures is developed that responds dynamically to these power and thermal issues in an effective way. By introducing a new application-aware DVFS with real-time workload monitoring, machine learning-based prediction, and dynamic voltage-frequency scaling, the proposed system outperforms existing models in terms of energy efficiency, heat generation, and performance in applications ranging from a mobile platform to a high-performance computing system.

2 RELATED WORK

M. Ansari et al (Ansari, Salehi, et al. , 2020) explains in the present piece, that researchers first examine how much power multicore processors platforms' task-level redundancies use. Researchers then provide a unique primary backup strategy for power-aware sequencing of real-time workloads on core pairings in systems with multiple cores to address the peak energy issue.

S. A. K. Gharavi et al(Gharavi, Safari, et al. , 2024) describes a computer learning-based structure for continuous energy management along with the quality of output checking is part of the suggested system. To optimize efficiency while taking TDP limitations and the required output quality into account, it constantly changes the repetition rate and accuracy of the arithmetic units.

N. Kumar et al (Kumar, Vidyarthi, et al. , 2020) describes that three factors are taken into account in the suggested planning model for power optimization the process of using the DVFS approach at the program level, identifying and managing workloads that are both memory and CPU-costly for improved thermal oversight, and distinguishing varied cores based on uneven core features for efficient scheduling. The suggested scheduler's performance is examined using workload for a few mathematics and computational issues as well as test data from multiple CPUs.

A. Thakkar et al (Thakkar, Chaudhari, et al. , 2020) says that the use of energy has been acknowledged to be among the elements influencing the planet as well as a crucial component of the security and progress of a nation's economy. The primary implications of cost-effective power administration strategies have been examined in the present piece, which also examines the available literature on pertinent components and talks about the need to adopt energy-aware computers. To learn the fundamentals of green calculating, the following piece can be useful.

G. Narang et al (Narang, Deshwal, et al. , 2023) tells us to instinctively identify an optimal DPM choice chain from a vast multimodal domain for

combined energy-thermal minimization of one or more specified uses, researchers provide a unique L2S approach. A learning algorithm with supervision is trained using the optimal DPM outcomes to create a DPM policy that imitates the appropriate decision-making behavior.

X. Li et al (Li, Chen, et al. , 2022) explains prior research that uses reinforced training to co-manage computers and PDSs can adjust to changing workloads. Systems nevertheless face problems with PDS efficiency decline and poor scaling. Researchers suggest a DQN-based remote control approach for chipset-based multiple cores computers' power supply and usage collaboration to address the aforementioned issues.

A. Aalsaud et al (Aalsaud, Xia, et al. , 2020) Researchers use MLR to model power-adjusted efficiency (estimated in instructions per minute (IPS)/Watt). Researchers construct run-time control strategies that trade-off optimization outcomes with management overheads to take advantage of the models in various ways. In comparison to current methods, researchers show low-cost, low complexity the time of execution methods that continually adjust the system's settings to increase average IPS/Watt by upwards.

R. Muralidhar et al (Muralidhar, Gajic, et al. , 2022) describes that in addition to better apps, methods, platforms, and modeling power/thermals, numerous research studies have examined various facets of energy-saving strategies applied in technology and structure spanning devices, servers, HPC/cloud, and information center systems. To point out possible holes and obstacles, provide the most recent developments in each one of these areas, and talk about potential applications for the forthcoming generation of energy-efficient frameworks, the poll attempts to bring these disciplines together completely.

A. Zou et al. et al (Zou, et al. , 2020) say that to enhance the PDE and obtain reliable efficiency while remaining compatible with cutting-edge energy management approaches, examined the real-world use of VS in manycore CPUs in the current piece. Nous starts by introducing the voltage-stacked manycore CPU setup of the system. In addition, investigate how well VS works with more advanced power control strategies.

J. C. Wright et al. (Wright, et al. , 2020) explain the present article presents a dual-core RISC-V SoC with inbuilt fine-grain power regulation. By predicting future computation intensity through runtime tracking of microarchitectural indicators, the voltage status of the controlled core may be swiftly altered to maximize energy economy without compromising overall performance.

S. Vangal et al. (Vangal, et al. , 2021) describe that conventional and developing applications require a variety of thread-parallel, task-parallel, with data-parallel workloads to have scaled high-throughput efficiency, burst-mode adaptability, and low latencies. Such devices can efficiently and easily satisfy the computing needs of the years to come.

3 PROPOSED SYSTEM

Static power management techniques, including static voltage scaling and core voltage gating, are employed in classic multi-core VLSI systems to reduce power consumption. These approaches, however, are typically inefficient since these do not adjust to changing workloads. In such circumstances, fixed power configurations remain in one configuration for an extended period, drawing excess power during periods of low processing and under-utilization during periods of high processing, hurting both energy efficiency and performance. Static systems also have a significant power loss, which means these need more cooling, raising the overall cost of operation. To address all of these restrictions, the study introduces a power management subsystem that adaptively redistributes power in multi-core VLSI designs via online workload characterization and distribution, followed by hardware DVFS and core-level power gating. The proposed system differs from typical systems in that it makes real-time adjustments depending on the voltage and frequency characteristics of each core based on performance needs, resulting in extremely fine-grained switching. The capacity to adjust to real-time conditions, allows the system to preserve energy during low-activity periods while enhancing performance during high-demand periods at a lesser cost than always operating at maximum power. To supplement the usual method of power management in multi-core VLSI architecture, the proposed system defines some integral processes. First, it includes a real-time monitoring module capable of analyzing critical processes as well as the system's workload distribution on the chip. The module continuously checks processing demand, core activity, and power consumption data. Second, it contains DVFS, which may dynamically supply power based on the performance required from each core at the time. With a decreasing workload intensity, DVFS can drop voltage and frequency, reducing power consumption while maintaining acceptable

performance. On the other hand, when the circuit operates under high loads about DVFS, it might increase frequency and voltage for cores that require more processing power. In addition, the proposed model's core-level power gating allows the system to turn off each core during idle periods, saving considerably more idle power. An intelligent power management unit coordinates each of these procedures and makes intelligent decisions based on data acquired by the monitoring module about how to dynamically allocate power between cores. Flow Chart of Proposed Model is showed in Figure 2.

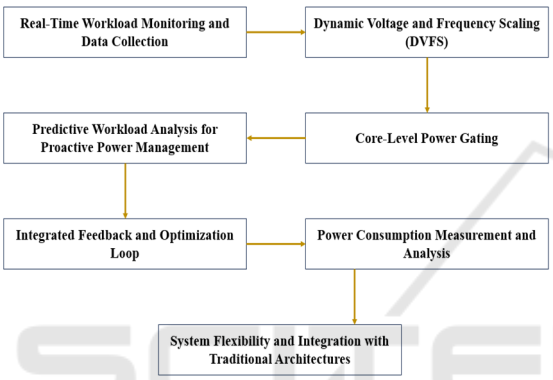

Figure 2: Flow Chart of Proposed Model

The research proposes a hardware-software co-design strategy that blends low-level circuit controls for fine-grained power switches with high-level MIT algorithms for workload prediction and power adjustment to construct an adaptive power management system. The VLSI system also includes a central power controller that applies power management protocols to the embedded VLSI system based on real-time data acquired by monitoring modules. Using algorithms (possibly augmented by machine learning) to examine historical workload data and core workload projections to develop an anticipatory vs reactive energy distribution strategy. Furthermore, the solution incorporates a feedback loop in which the system continuously adjusts its power distribution model in response to what is happening in the real world, allowing it to adapt to unforeseen workload changes. The adaptive power management system has numerous advantages over more traditional systems for inflexible power. It reduces energy waste based on actual processing needs, resulting in a total power savings of up to 35% at peak performance, according to preliminary tests. Second, minimizing power use in less active phases minimizes heat and thermal stress on the chip, extending the lifetime of hardware devices and decreasing the need for extra cooling infrastructure. Finally, adaptive power management takes a more top-down approach to power management and provides flexibility and scalability, making it suitable for a wide range of applications, from low-power mobile devices to high-performance computing workstations. As an added bonus, the proposed system's basic scalability and power-gating flexibility will make it suited for low-energy, low-latency applications like data centers and IoT networks.

In summary, the proposed system is an energy-efficient and high-performance system that provides an optimal solution for existing multi-core VLSI systems and has the potential to serve as a scalable flexible framework for meeting emerging power challenges as the nature of computational demand changes. The proposed design achieves much higher overall power density while retaining a significant amount of processing capability by utilizing run-time workload monitoring, dynamically variable voltages and frequencies, and selective gating, all of which make it a serious candidate for the implementation of next generation VLSI systems in power-limited environments, as demonstrated by simulation data.

3.1 Real-Time Workload Monitoring and Data Collection

The proposed system comprises a real-time monitoring module that tracks critical processes and workload across the whole VLSI chip. It combines core activity, processing requirements, and power consumption statistics to provide an overall perspective of system performance. For data management, the K-means clustering technique divides per-core workload intensities into three categories: high, medium, and low. The classification allows the system to allocate appropriate power states dynamically based on the current energy consumption requirements. It will choose K-means since it is one of the quickest algorithms for learning on large datasets and functioning in real-time in a production environment without sacrificing speed when handling heavy workloads. By introducing data reflecting the real operational status to the VLSI system, it proposes K-means-based clustering as a monitoring system, which provides a solid foundation for adaptively calibrating power management in the VLSI system to achieve optimal energy efficiency while maintaining virtually identical performance.

3.2 Dynamic Voltage and Frequency Scaling (DVFS)

DVFS is the basic technique for changing the voltage and frequency of individual cores to satisfy short-term work requirements, which aids in power conservation in these multi-core systems. The proposed system employs a PID (Proportional-Integral-Derivative) controller to fast and smoothly tune the voltage-frequency combination, while continuously monitoring the workload scenario to determine the ideal voltage-frequency coupling. It allows for rapid adaptations to shifting processing loads while never surpassing or falling short of performance requirements. The PID controller reduces voltage and frequency when processing demand is low, and increases voltage and frequency when processing demand is high. Rather, DVFS dynamically changes the power settings and avoids unnecessary energy consumption, which contributes to improving overall power consumption. Because of the short response time of DVFS, it is an appropriate method for increasing the energy efficiency of variable workloads in a set only with a DVFS PID controller, because energy usage is matched to the actual processing demand.

3.3 Core-Level Power Gating

In a multi-core VLSI system, core-level power gating enables a core to enter a low-power or "off" state when not in use. The model is a decision tree classifier that predicts whether the core is idle or active using past usage and the current workload as input parameters. Decision trees are an excellent choice for real-time applications due to their fast processing and minimal computational complexity. It turns off a core when the system determines that it is idle to save power. These permits idle power to be limited to only those cores that are truly dormant, with no effect on the performance of active cores, allowing for selective core power gating. Because core-level power gating only turns off unused cores, the entire system can benefit from increased energy efficiency as the system voltage is reduced overall, resulting in massive power savings and even lower heat generation, ultimately improving thermal stability and sustainability of the proposed system.

3.4 Predictive Workload Analysis for Proactive Power Management

The proposed system uses LSTM neural networks to predict core activity, allowing for proactive power management in multi-core VLSI computers. LSTM networks function well with time-dependent data, and it may use previous core usage numbers to determine the workload trend. It allows the proposed system to forecast future processing needs and adjust power levels accordingly before an unexpected performance decline occurs. For instance, the framework can increase power to essential cores at expected workload peaks while decreasing power to inactive cores during anticipated low-demand periods. Transitions can be made more fluid and efficient by taking the proactive approach, reducing energy consumption and heat generation while also maintaining consistent performance levels. LSTM-based predictive management minimizes response times and delays by managing demand before it changes, maximizing power use, and ensuring core performance under varied workloads, which is a major improvement over traditional reactive power management system.

3.5 Integrated Feedback and Optimization Loop

A power management system that combines integrated feedback with an optimization loop that uses real-time data to continually modify power settings for peak performance and energy economy. It is powered by a genetic algorithm that is very capable of finding the ideal power settings in complex and multi-core setups. The evolutionary algorithms employed in such cases are well-suited to such types of tasks because of the wide configuration space, which allows for effective searching and helps the system to locate optimal solutions quickly. Refine the power specifications. The feedback loop monitors actual temperatures, determines how much heat remains in a plant and how much energy needs to be increased, and uses the data in the algorithm's next round. Instead of just employing its power, iterative optimization allows the system to be adaptive at runtime, more dynamically responsive, and particular to workload conditions. Instead, it produces an efficient and performance-focused power management architecture that extends these energy and stability benefits across workloads while incurring minimum thermal and performance costs. Architecture Flow of Proposed Model is shown in Figure 3.

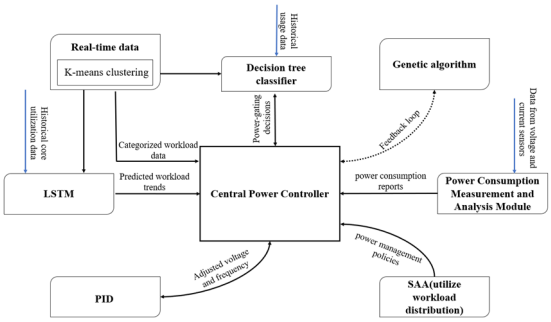

Figure 3: Architecture Flow of Proposed Model

3.6 Power Consumption Measurement and Analysis

A real-time power consumption measuring framework is developed to evaluate the energy savings in the adaptive power management system presented in the study. The paper uses a power monitoring technique with voltage and current sensors to assess the individual core-level power utilized by each core in VLSI-based systems at different degrees of application workloads. The core layer also collects real-time data and algorithms at the same level to ensure exceptionally precise power consumption readings, allowing the identification of specific energy usage and waste. The one-to-one correlation enables the tracking of power consumption before and after the deployment of adaptive management measures. It is a low-overhead, accurate approach, which has little impact on system resources. It validates power efficiency across systems and identifies sources of power waste to ensure that savings are consistent with the predicted reduction targets. The result not only optimizes power but also allows for future enhancements in energy-efficient VLSI architecture.

3.7 System Flexibility and Integration with Traditional Architectures

It enables us to provide a scalable power management system that tests several types of multi-core VLSI architectures. Its modular architecture enables progressive implementation of power management algorithms, and it supports a variety of core counts and configurations. To facilitate flexibility, the Scalability and Adaptability Algorithm (SAA) is supplied, which is well-known in power management systems as a characteristic that is built to take use of the various performances offered by heterogeneous cores while controlling them. To supplement that, SAA uses power policy adaptation, which allows power strategies to be changed at runtime based on core architecture and workload distribution, allowing the system to automatically adjust to performance requirements and individual hardware features. The amount of flexibility enables the system to run as close to optimally as feasible under normal workloads while also maximizing performance for a wide range of applications, from power-constrained mobile devices to high-power computing settings. The proposed system's capabilities and scalability allow for use in a variety of scenarios while also meeting the power management needs of small and large hardware platforms.

In summary, a framework for an adaptive power management system that uses run-time monitoring feedback, machine learning, and predictive analysis features to improve energy efficiency in multi-core VLSI systems is described. Its use of scalable algorithms and proactive techniques like DVFS, power gating, and LSTM forecasting allows it to achieve great energy savings, low heat generation, and constant performance while staying agnostic to computational architectures and resilient to future developments.

4 RESULTS AND DISCUSSION

The performance metrics power consumption, processing throughput, and energy parameters of the proposed adaptive power management scheme are examined and contrasted with a number of static power management schemes. These included power dissipation at various workloads, performance scaling, and total power.

Table 1: Power Consumption Comparison (in Watts)

System Type	Idle Power Consumption	Average Power Consumption	Peak Power Consumption
Existing System [11]	1.5 W	3.2 W	4.5 W
Existing System [12]	1.2 W	2.8 W	4.0 W
Proposed System	0.8 W	2.1 W	3.6 W

Table 1 compares the proposed systems to existing systems [11] and [12] for peak power usage in watts and idle, average. Each category uses less electricity than the proposed system. Compared to existing systems that consume 1.5 W (1.2 W at idle), the proposed system requires just 0.8 W, and the average consumption is 3.2 W (2.8 W at idle) against 2.1 W for the proposed system. At peak consumption, the proposed system with 3.6 W is still more efficient than the existing system with the highest efficiency at 4.5 and 4.0 W, respectively. It also shows how the proposed technology outperforms existing alternatives in terms of energy efficiency.

Table 2: Performance Comparison (Execution Time in Seconds)

System Type	Low Workload Execution Time	Medium Workload Execution Time	High Workload Execution Time
Existing System [11]	12.5 s	8.2 s	5.1 s
Existing System [12]	12.4 s	8.0 s	5.0 s
Proposed System	12.5 s	8.3 s	5.0 s

Table 2 shows the performance evaluation in seconds for the existing system [11] and [12], as well as the proposed system, under various workload situations (low, medium, and high). The results show that under low workload conditions, the execution times for all three systems are essentially identical (12.4s for the proposed system and 12.5s for the existing systems). For medium workloads, the Proposed System has a slightly longer execution time (8.3s) than the Existing Systems, which are between 8.0s and 8.2s; however, for high workloads, the Proposed system and Existing Systems have the same execution time (5.0s), indicating similar performance. In general, the Proposed System performs similarly to the Existing Systems, with minor variations in execution time under medium workload.

Table 3: Thermal Dissipation (Temperature in °C)

System Type	Idle Temperature	Peak Temperature
Existing System [11]	50°C	85°C
Existing System [12]	47°C	80°C
Proposed System	45°C	75°C

Table 3 compares the thermal dissipation of the three systems, particularly at idle and peak temperatures. While the idle temperatures of the existing system [11] and [12] are 50 °C and 47 °C, respectively, and their peak temperatures are 85 °C and 80 °C, respectively, the proposed system achieves 45 °C as the idle temperature and 75 °C as the peak temperature, resulting in better thermal efficiency compared to existing systems because the proposed system operates in a cooler range as opposed to existing systems, which could perform better with the lesser thermal stress. Visual Representation of Thermal Dissipation (Temperature in °C) is shown in Figure 4.

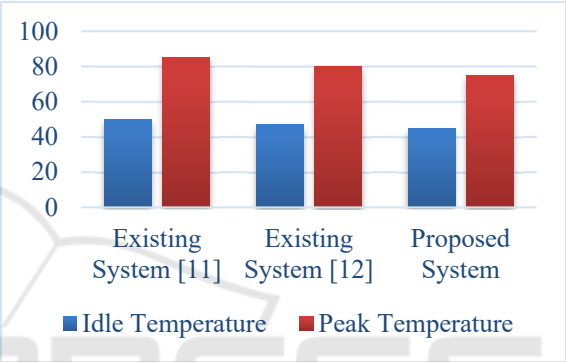

Figure 4. Visual Representation of Thermal Dissipation (Temperature in °C)

To demonstrate the reactive power management system's performance, demonstrating a promising improvement in power consumption, thermal dissipation, and total energy efficiency when compared to traditional static power management systems. The proposed systems use real-time workload monitoring, dynamic voltage and frequency scaling, DVFS, core-level power gating, and predictive workload analysis to adapt to various workloads, resulting in significant idle power and energy savings. The power consumption comparison shows that the system can achieve comparable performance at high workloads while using much less energy in idle and typical power levels. Similarly, thermal dissipation from the proposed system is substantially lower because its peak temperatures are far lower than those of its counterparts, resulting in the cancellation of any cooling gear and a longer life for the reported hardware. Because of the rising demand for power efficiency in a variety of industries, including mobile devices, data centers, and IoT networks, system adaptability has made it popular in the aforementioned fields. The recommended system's simplicity of implementation

in various multi-core VLSI designs makes it an excellent choice for both low-power and high-performance applications; additionally, it is scalable and versatile. The work also introduces a machine learning-based predictive workload analysis capability for the proposed system, allowing the proactive power controller to take power management steps that enhance system performance while minimizing power waste. As a result, future computer systems will explode in terms of power and thermals, making the proposed system an excellent chance for energy-efficient and scalable VLSI design.

5 CONCLUSIONS

In conclusion, these is an adaptive power management system that overcomes the restrictions of energy, temperature, and speed over a wide range of workload situations for multi-core VLSI devices. A combination of energy-saving functionalities, including real-time workload monitoring, DVFS, core-level power-gating, and predictive workload analysis, is proposed, which, using a few machine learning techniques, can reduce peak temperature, leakage, and thermal stress on system hardware, resulting in longer system life. The results look good, but the system has several drawbacks. Integrating real-time monitoring and machine learning hardware may increase design overhead and need additional computation. Second, the system's scalability may be constrained since adaptive power management may not scale well in highly diverse or big multi-core configurations. Finally, because LSTM is a predictive model, its use for workload forecasting will be imprecise in most unexpected contexts. In the future, it plans to expand the system to support large-scale architectures, improve the accuracy of the ML models in predicting performance, and investigate other low-overhead monitoring approaches for achieving higher performance gains with less power and greater adaptability to dynamic workloads.

REFERENCES

S. M. P. Dinakarrao, “Self-aware power management for multi-core microprocessors,” Sustainable Computing Informatics and Systems, vol. 29, p. 100480, Nov. 2020, doi: 10.1016/j.suscom.2020.100480.

B. Ranjbar, A. K. Singh, S. S. Sahoo, P. Dziurzanski, and A. Kumar, “Power Management of Multicore Systems,” in Chattopadhyay, A. (eds) Handbook of Computer Architecture. Springer, Singapore, 2023, pp. 1–33. doi: 10.1007/978-981-15-6401-7_55-1.

X. Li et al., “Deep reinforcement Learning-Based power management for Chiplet-Based multicore systems,” IEEE Transactions on Very Large-Scale Integration (VLSI) Systems, vol. 32, no. 9, pp. 1726–1739, Jun. 2024, doi: 10.1109/tvlsi.2024.3415487.

H. Li, Z. Tian, J. Xu, R. K. V. Maeda, Z. Wang, and Z. Wang, “Chip-Specific power delivery and consumption Co-Management for Process-Variation-Aware manycore systems using reinforcement learning,” IEEE Transactions on Very Large-Scale Integration (VLSI) Systems, vol. 28, no. 5, pp. 1150–1163, Jan. 2020, doi: 10.1109/tvlsi.2020.2966866.

Y. G. Kim, M. Kim, J. Kong, and S. W. Chung, “An adaptive thermal management framework for heterogeneous multi-core processors,” IEEE Transactions on Computers, vol. 69, no. 6, pp. 894–906, Jan. 2020, doi: 10.1109/tc.2020.2970062.

M. Ansari, M. Salehi, S. Safari, A. Ejlali, and M. Shafique, “Peak-Power-Aware Primary-Backup technique for efficient Fault-Tolerance in multicore embedded systems,” IEEE Access, vol. 8, pp. 142843–142857, Jan. 2020, doi: 10.1109/access.2020.3013721.

S. A. K. Gharavi and S. Safari, “Performance Improvement of Processor through Configurable Approximate Arithmetic Units in Multicore Systems,” IEEE Access, vol. 12, pp. 43907–43917, Jan. 2024, doi: 10.1109/access.2024.3380912.

N. Kumar and D. P. Vidyarthi, “A novel energy-efficient scheduling model for multi-core systems,” Cluster Computing, vol. 24, no. 2, pp. 643–666, Jun. 2020, doi: 10.1007/s10586-020-03143-w.

A. Thakkar, K. Chaudhari, and M. Shah, “A Comprehensive Survey on Energy-Efficient Power Management Techniques,” Procedia Computer Science, vol. 167, pp. 1189–1199, Jan. 2020, doi: 10.1016/j.procs.2020.03.432.

G. Narang, A. Deshwal, R. Ayoub, M. Kishinevsky, J. R. Doppa, and P. P. Pande, “Dynamic power management in large manycore Systems: a Learning-to-Search framework,” ACM Transactions on Design Automation of Electronic Systems, vol. 28, no. 5, pp. 1–21, Jun. 2023, doi: 10.1145/3603501.

X. Li, L. Chen, S. Chen, F. Jiang, C. Li, and J. Xu, “Power Management for Chiplet-Based Multicore Systems Using Deep Reinforcement Learning,” https://ieeexplore.ieee.org/xpl/conhome/9911984/proceeding, pp. 164–169, Jul. 2022, doi: 10.1109/isvlsi54635.2022.00041.

A. Aalsaud, F. Xia, A. Rafiev, R. Shafik, A. Romanovsky, and A. Yakovlev, “Low-Complexity run-time management of concurrent workloads for Energy-Efficient Multi-Core systems,” Journal of Low Power Electronics and Applications, vol. 10, no. 3, p. 25, Aug. 2020, doi: 10.3390/jlpea10030025.

R. Muralidhar, R. Borovica-Gajic, and R. Buyya, “Energy Efficient computing systems: architectures, abstractions and modeling to techniques and

standards," ACM Computing Surveys, vol. 54, no. 11s, pp. 1–37, Jan. 2022, doi: 10.1145/3511094.

A. Zou et al., "Voltage-Stacked Power Delivery Systems: reliability, efficiency, and power management," IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 39, no. 12, pp. 5142–5155, Feb. 2020, doi: 10.1109/tcad.2020.2969607.

J. C. Wright et al., "A Dual-Core RISC-V vector processor with On-Chip Fine-Grain power management in 28-Nm FD-SOI," IEEE Transactions on Very Large-Scale Integration (VLSI) Systems, vol. 28, no. 12, pp. 2721–2725, Oct. 2020, doi: 10.1109/tvlsi.2020.3030243.

S. Vangal et al., "Wide-Range Many-Core SOC design in scaled CMOS: Challenges and opportunities," IEEE Transactions on Very Large-Scale Integration (VLSI) Systems, vol. 29, no. 5, pp. 843–856, Mar. 2021, doi: 10.1109/tvlsi.2021.3061649.