ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24SEP685

Dynamic Clock Switching Scheme for Enhanced Power Efficiency in System-on-Chip Architectures

Abhijith Prabha
Indian Institute of Science (IISc), Bangalore, India

Abstract:- High power consumption and thermal management are critical challenges in traditional System-on-Chip (SoC) architectures due to multiple active clock sources feeding individual IP blocks, even during idle states. This paper introduces a dynamic clock switching power-saving scheme that consolidates clock sources by dynamically switching to a single lower-frequency clock during low-power modes. A Clock Control Agent (CCA) monitors real-time operational status, power consumption, and performance needs, leveraging AI for predictive adjustments. Experimental results demonstrate a significant reduction in idle power consumption and improved thermal management, without compromising performance or data integrity. This scheme addresses inefficiencies in existing methods such as Dynamic Voltage and Frequency Scaling (DVFS) and clock gating, offering a robust and efficient solution for next-generation SoC designs.

Keywords:- *System-on-Chip (SoC), Clock Switching, Power Consumption, Energy Efficiency, Clock Control Agent (CCA).*

I. INTRODUCTION

Modern System-on-Chip (SoC) architectures are pivotal in enabling the functionality of various electronic devices, from smartphones to embedded systems. However, the increasing complexity of SoCs has led to significant power consumption and thermal management challenges, particularly due to multiple active clock sources feeding individual IP blocks, even during idle states [1][2]. This inefficiency exacerbates battery drain in portable devices and increases thermal load, impacting device performance and lifespan [3][4].

Traditional power-saving techniques like Dynamic Voltage and Frequency Scaling (DVFS) and clock gating have limitations. DVFS adjusts the voltage and frequency based on workload demands but does not address idle power consumption from multiple active clocks [5]. Clock gating reduces power by disabling the clock signal when not in use but introduces latency and does not support minimal operations [6]. Power gating, which shuts down idle IP blocks entirely, results in significant wake-up latency [7].

This paper introduces a novel dynamic clock switching scheme that consolidates multiple clock sources into a single lower-frequency clock during low-power modes. Our proposed solution utilizes an AI integrated Clock Control Agent (CCA) for real-time monitoring and dynamic adjustment, significantly reducing idle power consumption and improving thermal management without compromising performance. This innovative approach addresses the limitations of existing methods, providing an efficient solution for future SoC designs.

II. RELATED WORK

Power optimization in System-on-Chip (SoC) architectures has been an active area of research, with various techniques proposed to enhance efficiency. Dynamic Voltage and Frequency Scaling (DVFS) is one prominent method, adjusting voltage and frequency based on workload demands. However, DVFS cannot eliminate idle power consumption from multiple active clock sources [8]. Another widely used technique is clock gating, which reduces power by disabling the clock signal when not in use. While effective in cutting dynamic power, clock gating introduces latency issues and fails to support minimal operations during idle states [9].

Power gating offers another solution by completely shutting down idle IP blocks, leading to significant reductions in leakage power. Nevertheless, it results in increased latency due to the time needed to reinitialize the blocks [10]. Advanced methods like Adaptive Power Gating have been suggested to address this latency, but the complexity and overhead associated with these approaches remain challenges [11].

Recent research has explored the use of machine learning algorithms for real-time power management in SoCs. Techniques such as Reinforcement Learning have shown promise in dynamically adjusting power states to optimize performance and power consumption [12][13]. However, these methods have yet to be widely adopted due to the computational overhead and integration challenges.

The proposed dynamic clock switching scheme builds upon these previous works by introducing an AI-integrated Clock Control Agent (CCA) that can dynamically consolidate multiple clock sources into a single lower-frequency clock during low-power modes. This approach significantly reduces idle power consumption and improves thermal management, addressing the drawbacks of traditional methods.

- *Problem Statement*

High power consumption and thermal management present significant challenges in modern compute devices, often due to the inefficiencies in traditional System-on-Chip

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24SEP685

(SoC) architectures. These architectures typically maintain multiple active clock sources even during idle periods, leading to unnecessary energy consumption and heat generation. Experiments indicates that even in low power modes, SoCs can draw a substantial amount of power due to redundant clock synthesis and distribution. Studies have shown that clocks, despite their minimal contribution to functional operations during idle states, can still account for a considerable portion of overall power dissipation, sometimes up to 30% of the total power consumption.

Real-world examples further highlight this issue. Qualcomm's Snapdragon SoCs, widely used in smartphones, implement sophisticated power management techniques such as clock gating, dynamic voltage, and frequency scaling (DVFS) to manage clock sources effectively during low-power states. Similarly, Apple's A-series chips, known for their efficiency, utilize a combination of custom power management controllers and fine-grained control over clock domains to minimize power loss from unnecessary clock activity. Another notable example is Intel's processors, which employ various low-power states, such as C-states, along with techniques like power gating and per-core clock control to reduce energy use. However, even with these advanced management techniques, they can still be inadequate due to overhead complexity, latency issues, and the difficulty in predicting workloads accurately. Inadequate clock management can lead to higher-than-expected power consumption and thermal issues.

This excessive power draw rapidly depletes battery life in portable devices, negatively impacting user experience and creating thermal management challenges that can result in thermal throttling. Thermal throttling degrades system performance and reduces device lifespan. For instance, prolonged high temperatures have been known to cause performance and reliability issues in devices like Google Pixel smartphones. Therefore, optimizing the management of clock sources is crucial. A more efficient approach to clock gating and on-demand clock generation would significantly enhance energy efficiency, mitigate thermal challenges, and maintain high performance, ensuring the reliability and longevity of modern compute devices.

➢ *Proposed System*

The proposed system introduces a dynamic clock switching scheme designed to alleviate the high-power consumption and thermal management issues prevalent in traditional System-on-Chip (SoC) architectures. These traditional architectures typically deploy multiple clock sources to fulfill the performance needs of various IP blocks. While this setup ensures optimal performance, it results in substantial power consumption and heat generation, particularly when several IP blocks remain idle or operate in low-power states. Our proposed dynamic clock switching scheme aims to address these inefficiencies by managing the clock sources based on real-time operational states of the IP blocks.

➢ *Core Concept*

The core concept of the proposed system involves consolidating multiple high-frequency clock sources into a single lower-frequency clock during low-power operational modes. This mechanism is governed by a Clock Control Agent (CCA), the backbone of the system. The CCA continuously monitors the operational status, power consumption, and performance requirements of each IP block in real time. By leveraging advanced AI techniques, the CCA predicts idle periods and dynamically adjusts clock parameters to optimize power usage. This dynamic adjustment ensures significant reductions in power consumption and heat generation while maintaining system performance, thereby providing an efficient, intelligent, and robust solution for managing power in System-on-Chip (SoC) architectures.

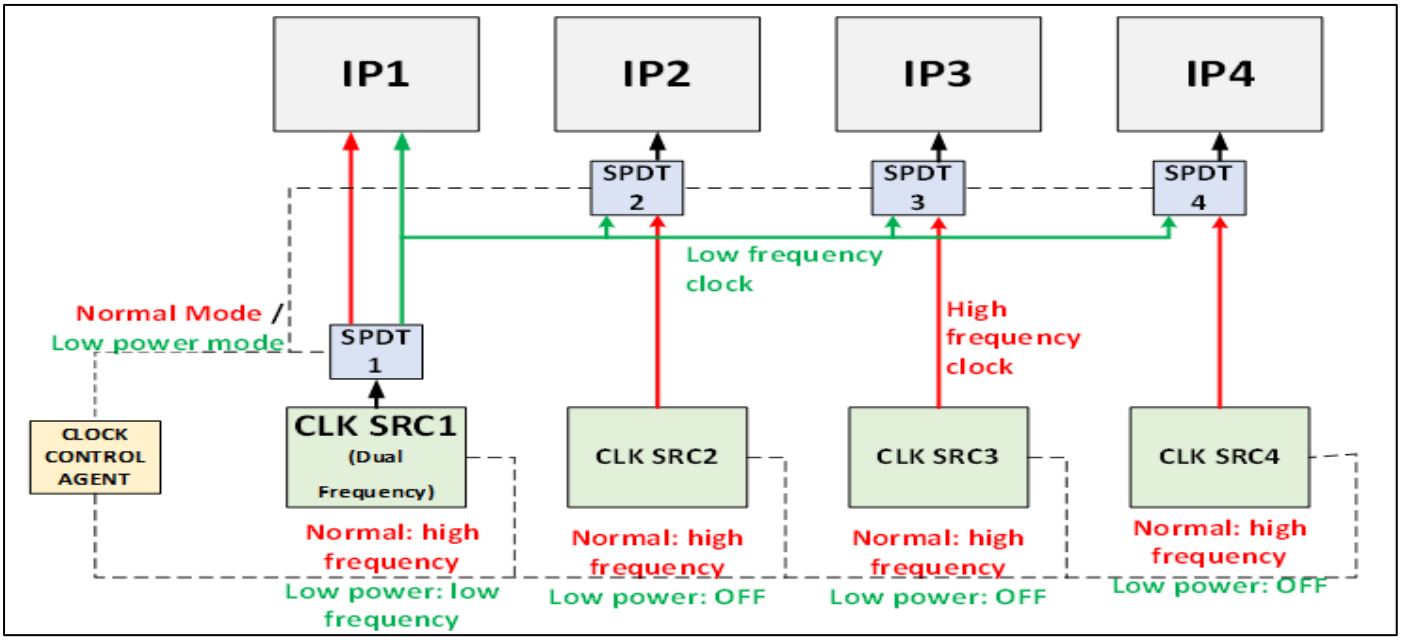

Fig 1 Proposed Scheme

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24SEP685

➢ *Design and Functionality*

- *Normal Operation Mode*:

In the normal operation mode, each IP block is supplied with a dedicated high-frequency clock tailored to its specific performance requirements. This ensures that each IP block performs optimally, albeit at the cost of high-power consumption and heat production due to the simultaneous activity of multiple clock sources.

- *Low-Power Mode*:

When the system detects that certain IP blocks are transitioning to low-power or idle states, the CCA activates the dynamic clock switching mechanism. It achieves this by consolidating these IP blocks onto a single lower-frequency clock source while deactivating the unnecessary high-frequency clocks. This consolidation significantly reduces power consumption and minimizes heat production.

➢ *Transition Management*

An essential feature of the proposed system is its ability to manage seamless transitions between high-performance and low-power modes. The CCA ensures that these transitions occur with minimal latency and that data integrity is maintained throughout the process, thus preventing any degradation in overall system performance. Furthermore, the CCA employs reinforcement learning algorithms to continuously refine its decision-making process. By analyzing historical data and operational patterns, the AI component can adapt and optimize clock source management strategies over time, leading to enhanced efficiency and performance.

➢ *Implementation Specifics*

During typical operation, each IP block operates with a dedicated high-frequency clock. These clocks are routed using Single-Pole Double-Throw (SPDT) switches controlled by the CCA. When the system enters a low-power mode, the CCA commands the SPDT switches to reroute the IP blocks' clock inputs from individual high-frequency sources to a consolidated lower-frequency clock. This rerouting involves turning off the redundant high-frequency clocks to conserve power.

One of the key design considerations in this scheme is ensuring that the lower-frequency clock parameters are suitable for all IP blocks in low-power mode without causing signal integrity issues. To that end, the lower-frequency clock is meticulously calibrated to meet the needs of the IP blocks, ensuring reliable and efficient operation.

➢ *Operation*

- *Monitoring and Decision-making*: The CCA continuously monitors each IP block's state, using sensors to track power consumption, operational status, and performance requirements.
- *Transition Activation*: When the system decides to enter a low-power mode, the CCA executes the transition by adjusting the SPDT switches to reroute the clock paths.
- *Adaptive Management*: Through reinforcement learning, the CCA improves its predictions and decisions over time, enhancing the system's efficiency. The AI algorithms allow the CCA to foresee periods of low activity accurately and adjust the clock frequencies preemptively.

➢ *Advantages*

The dynamic clock switching scheme offers numerous benefits. Primarily, it significantly reduces idle power consumption by consolidating clock sources, enhancing power efficiency. This reduction also lowers heat generation, simplifying thermal management and extending device lifespan. The scheme ensures smooth transitions between different power modes without compromising performance, thanks to seamless transitions. AI integration further adapts and optimizes clock management, enhancing long-term efficiency and system reliability. Additionally, by reducing the number of active clock sources, the scheme mitigates electromagnetic interference (EMI), leading to more reliable performance and improved compliance with regulatory standards.

III. IMPLEMENTATION AND EXPERIMENTAL SETUP

➢ *Implementation*

The proposed system's implementation hinges on the efficient operation of the Clock Control Agent (CCA), a vital component responsible for managing the dynamic clock switching mechanism. The CCA is designed to perform continuous real-time monitoring of the operational state, power consumption, and performance metrics of each IP block within the SoC. This monitoring is achieved through the integration of sensors that provide real-time feedback on these parameters. Utilizing predefined decision-making algorithms, the CCA evaluates the collected data to determine the optimal moments for switching between high-frequency and low-frequency clock sources. This decision-making process aims to minimize power consumption while maintaining necessary performance levels.

An advanced aspect of the Clock Control Architecture (CCA) is its AI-enhanced predictive capabilities. Utilizing reinforcement learning algorithms such as Q-learning and policy gradient methods, combined with deep learning techniques like convolutional neural networks (CNNs) and recurrent neural networks (RNNs), the AI dynamically adjusts clock parameters based on real-time system metrics and historical data. The AI employs a hierarchical neural network with layers dedicated to feature extraction and classification to categorize system states and make precise adjustments. By analysing time-series data through models such as Long Short-Term Memory (LSTM) networks and incorporating contextual information using attention mechanisms, the AI can accurately predict idle and active periods. Real-time monitoring and edge processing allow the AI to optimize clock gating, voltage scaling, and frequency modulation instantaneously, leveraging techniques like principal component analysis (PCA) for dimensionality reduction and faster computation. This comprehensive

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24SEP685

approach significantly enhances power efficiency, improves performance, and better manages thermal conditions.

The circuit design incorporates Single-Pole Double-Throw (SPDT) switches, intricately controlled by the Clock Control Architecture (CCA) to manage the routing of clock signals. Leveraging CMOS technology, these SPDT switches are embedded within the VLSI layout, ensuring low latency and minimal parasitic capacitance. During normal operation, each Intellectual Property (IP) block is driven by a dedicated high-frequency clock sourced from a programmable clock generator to ensure optimal performance and precise timing. The programmable clock generator, capable of generating both high and low-frequency outputs based on input control signals, can be implemented using a phase-locked loop (PLL) for high frequencies and a digitally controlled oscillator (DCO) for lower frequencies.

When transitioning to low-power mode, the CCA employs dynamic clock gating to deactivate the higher-frequency clocks for idle IP blocks, utilizing the SPDT switches to reroute a single lower-frequency clock signal from the same programmable clock generator. This reconfiguration significantly reduces dynamic power consumption and minimizes thermal output while maintaining clock tree and signal integrity.

The low-frequency clock is meticulously calibrated using adaptive voltage scaling (AVS) and temperature-compensated delay lines to meet the minimum operational requirements without causing jitter or skew. Additionally, employing mixed-signal design techniques, the lower-frequency clock ensures reliable operation of the idle IP blocks by incorporating on-chip decoupling capacitors to mitigate noise and preserve clock signal stability.

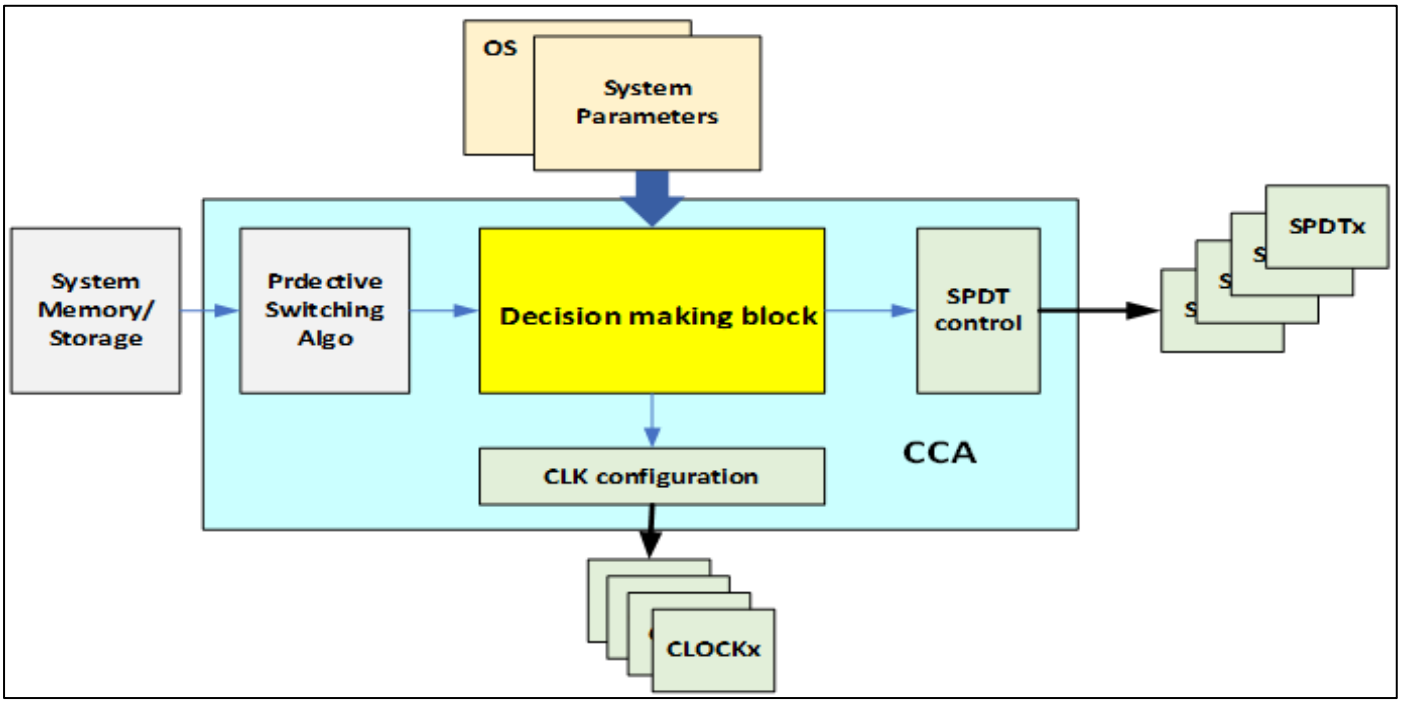

Fig 2 Clock Control Agent

➢ *Experimental Setup*

To evaluate the proposed dynamic clock switching scheme, a comprehensive experimental setup is employed to simulate the SoC platform's behavior. This simulation uses industry-standard tools such as VHDL/Verilog for hardware description and SPICE for circuit behavior simulation. The simulation platforms range from FPGA development boards to ASIC design environments, ensuring the assessment's flexibility and scalability. Various parameters, including clock frequencies, power consumption thresholds, and performance metrics, are configured to ensure thorough evaluation. Standard benchmarks like SPECint, SPECfp, and industry-specific workloads are employed to assess the system's performance and power efficiency under diverse operating conditions.

The experimental evaluation involves designing distinct test cases to cover various operational scenarios, ensuring a comprehensive analysis of the proposed scheme. The testbench framework is implemented in VHDL/Verilog and validated using simulation tools such as ModelSim and Synopsys VCS. One critical test case focuses on idle state transitions, evaluating the system's power savings and latency when switching between high-performance and low-power modes. This assessment is verified through cycle-accurate simulation and FPGA-based prototyping to capture real-world behavior. Another test case examines the system's ability to maintain performance during these transitions, measuring latency, data integrity, and responsiveness using custom scripts and verification Intellectual Property (IP). For thermal management improvement evaluation, thermal sensors embedded within the SoC are monitored, and thermal imaging cameras analyze heat distribution across various functional blocks under different states.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24SEP685

To assess the effectiveness of the AI-enhanced Clock Control Architecture (CCA) in predicting idle periods and optimizing clock source management, diverse workload patterns are simulated. The AI models, developed using frameworks like TensorFlow and PyTorch, are integrated into the CCA and tested for adaptive learning. The reinforcement learning algorithms' capacity to adapt and improve over time is scrutinized by iterating through multiple learning cycles and real-time data analytics.

Evaluation metrics in the experimental setup include power consumption reduction, transition latency, thermal load, and predictive accuracy. Power consumption is assessed using tools like Synopsys PrimePower and Cadence Voltus, comparing the proposed system to traditional methods. Transition latency is measured with Synopsys PrimeTime, evaluating the time taken to switch between modes to ensure minimal performance impact. Thermal load is evaluated using Ansys ICEPAK and Cadence Celsius to assess heat generation reduction and thermal management improvements. Predictive accuracy measures the AI-enhanced CCA's effectiveness in forecasting idle periods and managing power usage, using metrics like mean absolute error (MAE) and root mean square error (RMSE) from prediction logs.

By integrating detailed VLSI verification, prototyping, and validation techniques, this evaluation provides a robust assessment of the dynamic clock switching scheme's impact on system efficiency, performance, and reliability.

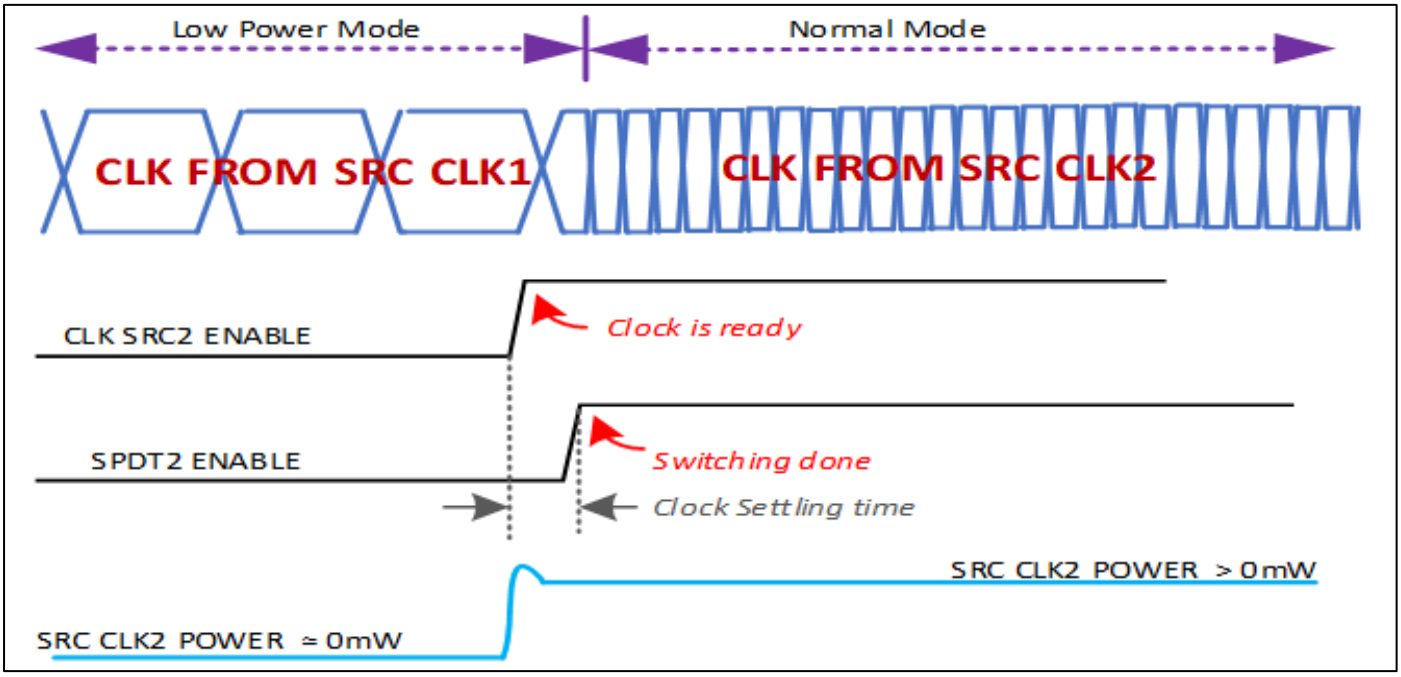

Fig 3 Clock Switching Timing Diagram

IV. RESULTS AND DISCUSSION

The proposed dynamic clock switching scheme aims to significantly reduce idle power consumption and thermal load in System-on-Chip (SoC) architectures by consolidating multiple clock sources into a single lower-frequency source during low-power modes. This approach can alleviate heat generation, improve thermal management, and extend device lifespan. The Clock Control Agent (CCA) ensures seamless transitions between high-performance and low-power modes with minimal latency, maintaining data integrity. Additionally, AI techniques within the CCA enhance predictive accuracy and dynamic adjustment of clock sources for more efficient real-time power management. Designing specific test cases, one could simulate various operational scenarios, measure power savings, and validate performance improvements. Such validations could provide deeper insights into the advantages of this scheme compared to traditional methods like Dynamic Voltage and Frequency Scaling (DVFS), clock gating, and power gating, potentially revealing its superior efficiency and effectiveness for modern SoC designs.

V. CONCLUSION

The proposed dynamic clock switching scheme offers substantial improvements in power consumption and thermal management for System-on-Chip (SoC) architectures. By integrating AI-enhanced control, the approach ensures seamless transitions with minimal latency while maintaining performance integrity. This results in significantly reduced idle power usage, lower heat generation, and an extended device lifespan. Additionally, by consolidating clock sources, the scheme mitigates electromagnetic interference (EMI), enhancing signal integrity and system reliability. Overall, the benefits include enhanced efficiency, reduced cooling requirements, and lower operating costs. As a superior alternative to traditional methods like DVFS, clock gating, and power gating, this innovative solution is poised to drive future energy-efficient, high-performance, and EMI-resilient SoC designs.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24SEP685

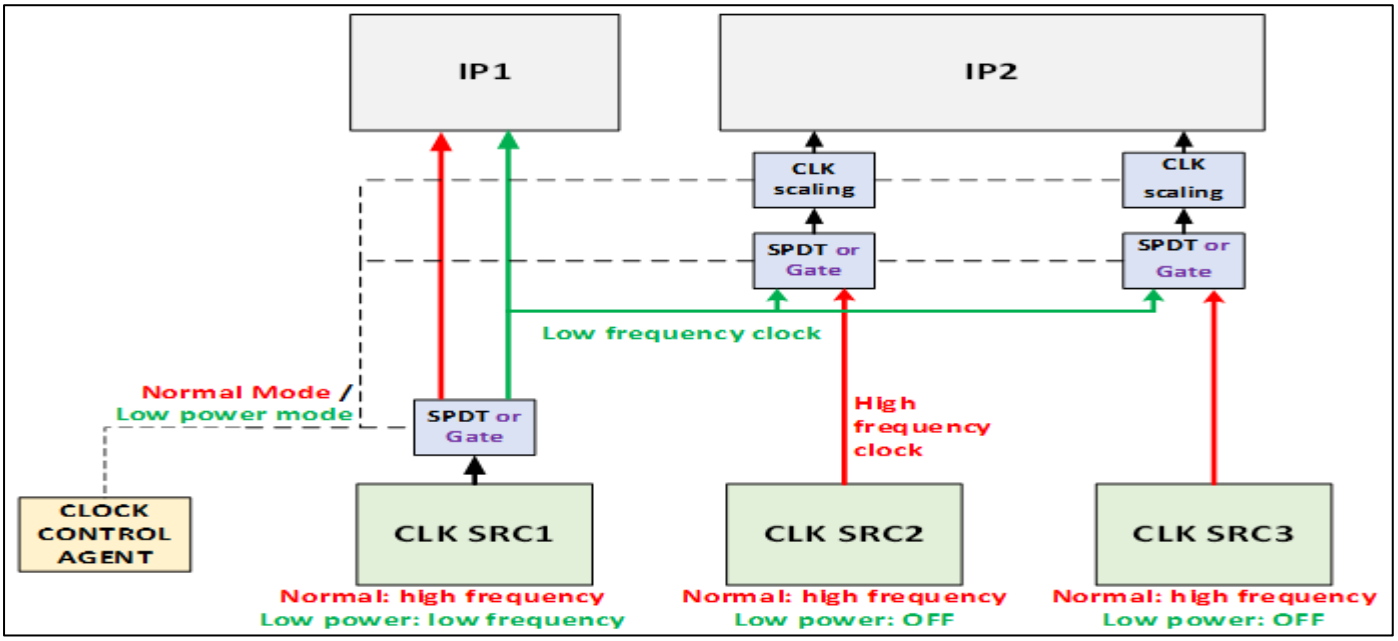

Fig 4 Alternate Approach Integrating Clock Gate

FUTURE WORK

Future work will focus on refining the AI algorithms within the Clock Control Agent (CCA) for even greater predictive accuracy and adaptability. Exploring machine learning techniques such as deep learning and advanced reinforcement learning could further optimize power management. Additionally, implementing this scheme in real-world SoC designs will help validate its effectiveness and uncover potential improvements. Expanding the scope to include different types of IP blocks and varying operational scenarios will ensure broader applicability. Investigations will also include integrating clock gating in combination with Single-Pole Double-Throw (SPDT) switches for scenarios where a complete shutdown is required, further enhancing power savings. Finally, examining integration with other power-saving techniques and further reducing transition latencies will improve overall system efficiency and performance.

REFERENCES

[1]. J. Koomey, “Estimating total power consumption by servers in the US and the world,” Stanford University, 2007.
[2]. W. Nebel and J. Mermet, "Low Power Design in Deep Submicron Electronics," Springer, 1997.
[3]. L. Benini and G. De Micheli, "System-Level Power Optimization: Techniques and Tools," ACM Transactions on Design Automation of Electronic Systems, vol. 5, no. 2, pp. 115-192, 2000.
[4]. M. Pedram, “Power-aware Design Methodologies,” Springer, 2002.
[5]. K. Choi, K. Dantu, W. Cheng, and M. Pedram, "Frame-based Dynamic Voltage and Frequency Scaling for a MPEG Decoder," IEEE/ACM International Conference on Computer-Aided Design, 2002.
[6]. I. Hong, D. K. Chen, et al., "Power Optimization of Real-time Embedded Systems on Variable Speed Processors," IEEE/ACM International Conference on Computer-Aided Design, 1998.
[7]. S. Mutoh et al., "1-V Power Supply High-Speed Digital Circuit Technology with Multithreshold-Voltage CMOS," IEEE Journal of Solid-State Circuits, vol. 30, no. 8, 1995.
[8]. K. Choi, K. Dantu, W. Cheng, and M. Pedram, "Frame-based Dynamic Voltage and Frequency Scaling for a MPEG Decoder," IEEE/ACM International Conference on Computer-Aided Design, 2002.
[9]. I. Hong, D. K. Chen, et al., "Power Optimization of Real-time Embedded Systems on Variable Speed Processors," IEEE/ACM International Conference on Computer-Aided Design, 1998.
[10]. S. Mutoh et al., "1-V Power Supply High-Speed Digital Circuit Technology with Multithreshold-Voltage CMOS," IEEE Journal of Solid-State Circuits, vol. 30, no. 8, 1995.
[11]. A. Raghunathan, N. K. Jha, and S. Dey, "High-Level Power Analysis and Optimization," Kluwer Academic Publishers, 1998.
[12]. Y. Fu, M. Pedram, "Dynamic Power Management Using Reinforcement Learning," IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 28, no. 5, pp. 676-689, 2009.
[13]. D. Brooks, V. Tiwari, and M. Martonosi, "Wattch: A Framework for Architectural-Level Power Analysis and Optimizations," ACM SIGARCH Computer Architecture News, vol. 28, no. 2, pp. 83-94, 2000.