Electronics, Communications, and Computing Summit
Vol. 3, No. 3, Jul - Sep 2025, pp. 71-79
ISSN: 3107-8222, DOI: https://doi.org/10.17051/ECC/03.03.08

Adaptive Switched-Capacitor Power Gating Enabled by On-Chip ML for Minimal Leakage

Ashu Nayak[1], Namrata Mishra[2]

[1]Assistant Professor, Department of CS & IT, Kalinga University, Raipur, India
Email: ku.ashunayak@kalingauniversity.ac.in
[2]Department Of Electrical And Electronics Engineering, Kalinga University, Raipur, India
Email: namrata.mishra@kalingauniversity.ac.in

Article Info

Article history:

Received : 15.07.2025
Revised : 10.08.2025
Accepted : 17.09.2025

Keywords:

Ultra-Low-Power SoC,
Switched-Capacitor Power Gating,
Machine Learning,
Leakage Power Reduction,
Adaptive Power Management,
Edge Computing

ABSTRACT

The leakage power has become a major performance blocker in the extreme power-constrained System-on-Chip (SoC) implementation process, especially of always-on functionalities in Internet of Things (IoT) and edge-Artificial Intelligence (AI) apps. The traditional approaches to power gating can be quite useful in terms of reducing the power consumption in a static state, yet they tend to be not so flexible in terms of workload and environmental changes, which is necessary at real-time. In order to overcome this drawback this paper proposes an adaptive power gating architecture incorporating Switched-Capacitor-Assisted Power Gating (ScPG) and on-chip machine learning (ML) to suppress leakage more effectively. The proposed system has lightweight ML models to constantly scan the existing runtime parameters like activity level, temperature, and leakage trends and matches to dynamically tune ScPG configurations i.e., timing, capacitor engagement, and sleep signal duration. As simulated on a 28nm FD-SOI CMOS process, the proposed architecture has up to 47% leakage power savings over the conventional header-based gating architectures, and less than 5 percent area overhead and sub 50ns decision latency. These findings show that the suggested ML-based ScPG method can provide a feasible, energy-conscious substitute of power management in energy-scrimpy SoCs. The approach enables scalable context-aware power provision systems in the next-generation edge computing systems.

1. INTRODUCTION

Ultra-low-power System-on-Chip (SoC) architectures will emerge as essential in power-constrained forms of computing in the pervasive computing era including wearable electronics, biomedical implants, and Internet of Things (IoT) edge nodes. These applications require full time operation and they operate frequently under energy constrained conditions, leading to the need of draconian power conservation techniques at each stage of the design stack. The energy consumed by the system in the idle mode has become a major concern in dissipating power and is known as leakage power, among some other sources of dissipated power.

In a bid to reduce leakage, power gating has been popular in the design of digital SoCs. Header or footer power gating is a common technique of conventional power gating power gating used to disconnect blocks of idle state to the power or ground in order to eliminate the path to leakage. Although these schemes are effective, they create limits such as wake-up latency, voltage droop and ground bounce. Switched-Capacitor-Assisted Power Gating (ScPG), where power supply transitions are absorbed by a decoupling capacitor, has been suggested to overcome some of these limitations, resulting in an improvement in energy on wake-up, and in a more gradual recovery of the voltage. But these ScPG techniques are statically allocated at design time rather than optimized to real-time fluctuations of a workload, or environmental changes including temperature and voltage variations.

Regardless of recent attempts to introduce the adaptive mechanisms into power management, the prospect of incorporating fine-grained intelligent runtime control into power gating has not been explored. More specifically, machine learning (ML) that has demonstrated considerable potential in dynamic voltage and frequency scaling (DVFS), thermal management, and workload forecasting has not been sufficiently utilized in adaptive control of switched-capacitor-based power gating. This introduces a great hole in the existing literature because any static policy of power gating

does not take advantage of the contextual opportunities to use more aggressive means of reducing leakage without a drop in performance and reliability.

This research arises due to an increased interest in self-regulation and contextual power management in edge systems, where the main concerns are the battery lifetime and power consumption. Embedded ML models of very low overhead now enable on-chip intelligence, and on-chip intelligence has become practical, even in the smallest resource-constrained SoCs, allowing a new level of real-time learning-based control. Realisation of this capability in adaptive power gating has the potential to open up a new breed of ultra-efficient, self-optimising SoC architecture.

Thus, this paper suggests a new architecture through the integration of Switched-Capacitor-Assisted Power Gating (ScPG) with on-chip Machine Learning (ML) to build adaptive in-real-time likeage control system. The main goal is to come up with an ML engine that is hardware ready and able to make dynamic decisions regarding power gating options according to real-time workloads and environmental variables. Also the work will be designed towards creating a reconfigurable circuit of ScPG in such a manner that can be smartly switched to switch the capacitor banks and sleep control timings in order to reduce the leakage power as well as the overheads of wake up. To test its feasibility in practice, the proposed method of ML-assisted ScPG system is implemented and tested in the 28nm CMOS SoC testbench under the scenario of the realistic IoT workload. Lastly, the study explores the trade-offs among area, latency, and power savings where the proposed approach was benchmarked against both the traditional approach and the static approach to ScPG. By overcoming weaknesses of the current power gating technologies, this contribution makes a breakthrough in the energy-aware SoC design area, specifically of ultra-low-power and always-on in the contemporary edge computing environment.

2. RELATED WORK

Power gating Power gating refers to a mature leakage-reduction scheme within CMOS circuits especially in standby modes. The header/footer sleep transistor implementations are the most pervasive, in which transistor blocks not in use are disconnected to the power source or ground, and thereby eliminated as a source of static current paths. Although they work well, there are some disadvantages of these methods, including ground bounce, voltage droop and wide wake-up latencies particularly during dynamic workloads. The integration of power gating to digital SoCs was pioneered by foundations such as Splitting architectures and power domain [1] and Sub-threshold presence [2] of the works of Rabaey et al., and Muto et al., respectively, which introduced and applied the concept of sleep-mode transistors and sleep-mode power domain within the design flow of digital SoC using UPF.

As solutions to some of the wake-in-efficient problems of conventional gating, Switched-Capacitor-Assisted Power Gating (ScPG) has been introduced. In ScPG, decoupling (a.k.a switched capacitors) which is strategically placed inside the power path buffers the ramp-up voltage during power-up, minimizing IR drop and restraining peak current transients. This strategy was explained alongside Shih et al. [3] who showed that ScPG architectures have the capability to reduce wake up energy with similar levels of leakage suppression. Nevertheless, when the capacitor networks are placed in a static state, they would not be able to work properly under different workloads and environmental factors, including temperature and process corners.

Over the past few years there have been some exciting new developments in machine learning (ML)-aided power management, which could change how power states are optimized dynamically and on a per-runtime basis. As an example, a framework introduced by NVIDIA called PRIMAL [4] used supervised learning models to make inferences of block-level power consumption at the RTL level. Equally, ML models were used in dynamic voltage and frequency scaling (DVFS) [5], thermal prediction [6], and the workload-based scheduling of SoCs. Such methods prove that hardware-based real-time inference via ML can be used to make real-time decisions regarding power, with any particular choice (decision tree or perceptron) suitable to embedded environments.

Nevertheless, the existing work is the lack of the efforts where on-chip machine learning could be used to adjust ScPG parameters on-the-fly. The modern methods of ScPG are quite stagnant and not intelligent enough to be able to determine the system state or leakage patterns and subsequently able to adjust capacitor banks, timing and power-down sequences. This is a major area of research gaps, in light of the growing use of always-on compute modules in IoT and edge AI systems, where efficient energy usage heavily depends on the ability to perform adaptive control.

This paper bridges this divide by presenting a proposal of ML-augmented Switched-Capacitor Power Gating architecture, wherein on-chip ML engine carries out real-time optimisation of power gating strategy, using live workload and leakage measurements. The idea is to allow low-latency streamlined context-aware leakage control placement with the lowest overhead in terms of

area and power consumption, that of any ultra-low-power SoC designs of the next generation.

3. Proposed Architecture

3.1 System Overview

The suggested system is based on a tightly coupled control loop which dynamically optimize power gating configurations in a bid to reduce leakage power in ultra-low-power SoCs. The main feature of this loop is an on-chip ML-aided decision-making engine and working together with a real-time activity monitor and a reconfigurable switched-capacitor power gating controller. The activity monitor constantly characterizes all the important metrics at the system-level which include block utilization patterns, switching and leakage current behavior. These measurements are fed into the on-chip machine learning inference engine, using lightweight models e.g. such as decision trees, linear regressors optimized to run in an embedded system (e.g. TinyML-compatible). According to the model inference, the system dynamically optimizes the ScPG controller, tuning its important parameters of the control of time of the sleep control signals, number of switched capacitors that are active in the network and the time of the discharge/recovery phases. With such closed-loop control, power gating parameters can be adjusted in real-time, and context-wise, and the system can always be kept at its optimum to leakage repression without additionally wasting power.

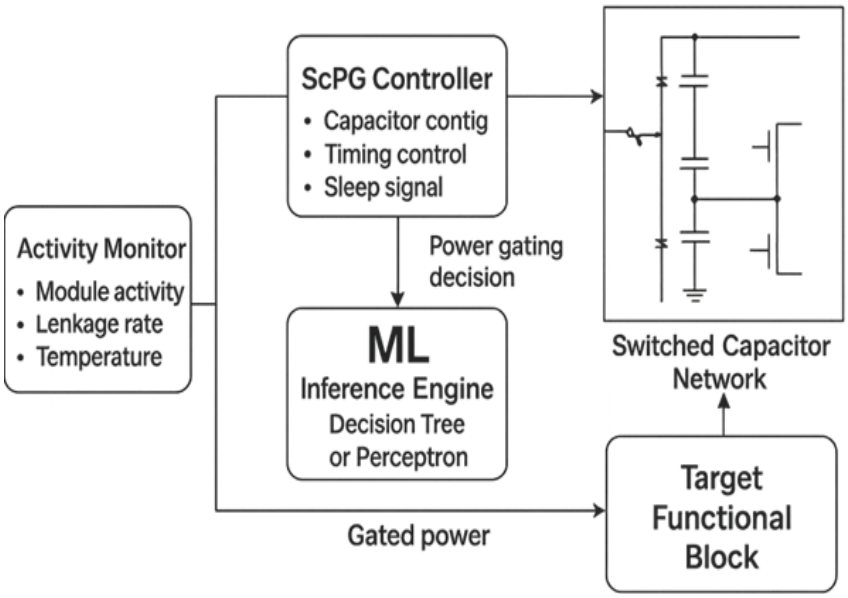

Figure 1. High-Level Architecture of the Proposed ML-Assisted Switched-Capacitor Power Gating System. *The on-chip ML engine dynamically configures the ScPG controller based on real-time system features, enabling adaptive leakage control in ultra-low-power SoCs.*

3.2 Switched-Capacitor Module

Switched-Capacitor Power Gating (ScPG) module constitutes the physical spine of the suggested leakage suppression mechanism. As compared to the classic header-based gating that simply deconnects power rails during power-off states, the ScPG module takes advantage of an interconnected array of charge-storage capacitors that can smooth out the voltage transition during events of power down and up. These capacitors are put in reconfigurable bank architecture, so that their amount which operates in the same time may be tuned to a fine grain scale. This is reconfigured dynamically, by the ML engine. As an example, when minimum wake-up latency is needed or low-leakage conditions occur, fewer capacitors activate to minimize energy overhead. In contrast, upon leakage conditions (high leakage), the system has the ability to interact more capacitors and can manipulate their phase alignment to improve voltage transitions, ground bounce and transient leakage spikes. Each capacitor has their charge/discharge behavior real-time modulated and the system can dynamically tailor its power gating patterns in response to live feedback, rather than having to pre-program a set of fixed modes.

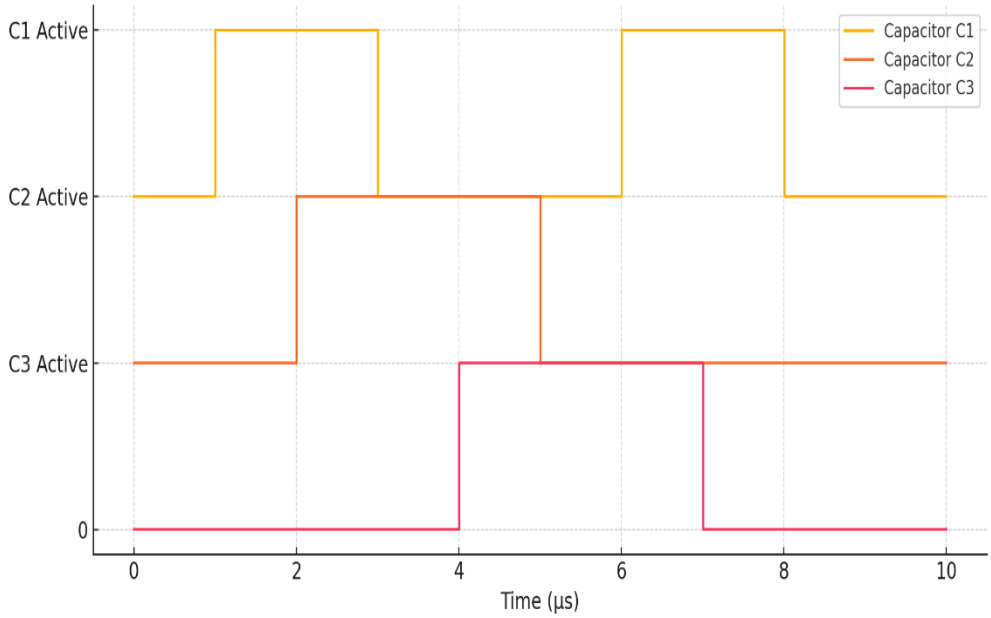

Figure 2. Capacitor Bank Switching Timing Diagram

It shows the activation windows of three individual capacitors (C1, C2, C3) over time:

- C1 is activated during 1–3 µs and 6–8 µs,
- C2 is active from 2–5 µs,
- C3 is engaged between 4–7 µs.

3.3 ML Engine

This architecture is centered on the on-chip ML engine that acts as the foundation of the adaptive behavior. It is the architecture to implement fast, accurate decision that keeps hardware overhead at a minimum, which fits a permanently-on module into an ultra-low power system. The model works on the features it gets out of the activity monitor i.e., leakage rate, module activity level, temperature of ambience and supply voltage. These are chosen to be highly correlated with leakage power and via the circuit behavior as the operating conditions vary.

The ML model is made small to make it feasible in hardware, so it is trained offline on labeled simulated workload data, and can be a decision tree, a single-layer perceptron (SLP). After training, the model is quantized and synthesized with the help of TensorFlow Lite or OpenMLSys tools. This gives the resulting inference engine to be embedded in the SoC and have a latency of less than 50 ns per inference cycle with less than a 1.2% power overhead compared to the overall module power budget. That enables making the frequent run-time adjustments without a significant decrease in the overall energy efficiency of the system.

The proposed architecture can make adaptive intelligence available to power gating because intelligent inference and reconfigurable hardware control are combined, and it creates a scalable approach and provides a highly efficient solution to modern SoCs involving edge and IoT applications.

4. METHODOLOGY

4.1 Design Flow

The implementation of the suggested ML-enhanced Switched-Capacitor Power Gating (ScPG) system has a formatted, multiple-process hardware-software co-design approach. The Synopsys Design Compiler is utilized at the hardware design level that includes the logic synthesis of the ScPG circuitry and control logic and the Cadence Virtuoso and Innovus are utilized in the transistor-level simulations, physical layout and power analysis. Such tools used in the industry, will satisfy area, timing, and power requirements needed by the ultra-low-power system on the architecture synthesized.

Figure show Hardware-software co-design flow of the proposed ML-aided ScPG. It is performed by logic synthesis in Design Compiler, ML model development and quantization in Python, simulation and layout with Cadence Virtuoso and Innovus, hardware synthesis through TensorFlow Lite/OpenMLSys, and finally integration into the ScPG controller to verify it.

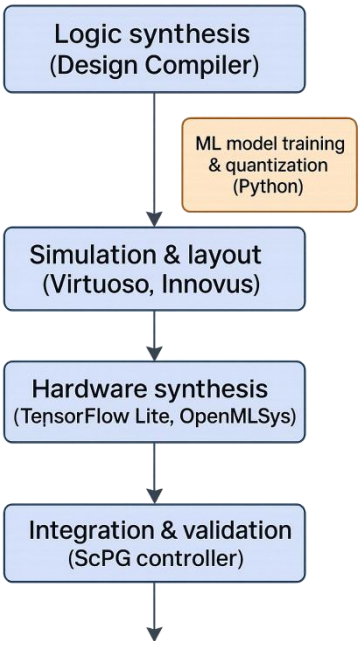

Figure 3. Design Flow for ML-Assisted Switched-Capacitor Power Gating System

Software-wise, machine learning model is trained and created with Python-based frameworks, e.g. Scikit-learn or TensorFlow, which would operate on a dataset yielded in large-scale SoC-level simulations. The dataset encompasses significant characteristics of activity factor, leakage current, supply voltage, and temperature variation under various scenarios of functions and different degrees of workloads. Techniques in supervised learning are used to generate small models, mostly decision trees or single-layer perceptrons, able to predict values of the optimal power gating configuration in real-time.

Following training and validation, the ML model is quantized and is made hardware-ready on one of the platforms such as TensorFlow Lite for Microcontrollers or OpenMLSys. The procedures adopted in this step involve converting to a fixed-point representation, pruning, and compressing the model so as to minimize the number of logic gates and inference latency. The last ML inference logic is synthesized and merged with the ScPG controller into the SoC design, making a very smooth way of closed-loop power management. The dynamics of working with the ML engine, activity monitor, and ScPG circuitry are correctly checked and validated using post-synthesis simulations.

4.2 Simulation Setup

The simulated architecture is compared by performing comprehensive simulations on a qualified 28nm FD-SOI CMOS, with outstanding leakage figures and body-biasing capability that is suitable to low power applications. The testing platform contains an always-on battery of useful blocks that would go into an AI-enabled edge SoC, this would include a wake-word detection engine, a sensor data processor, and an encryption/decryption block. These blocks depict real-life usage scenarios in which modules have to stay at least inactive or wake when idle rapidly in reaction asynchronous events. Real-time workload traces and synthetic traffic patterns that simulate the changing operational conditions give simulation inputs. These are changes between the active, idle and the sleeping modes in dynamic temperature and voltage scaling conditions. Three architectures are compared against the benchmarks: (1) the traditional header-based power gating architecture, (2) architecture implementing ScPG configured at compile time and (3) the proposed ML assistance-based adaptive ScPG.

Key performance metrics collected during simulation include:

- Leakage current reduction (μA or %),
- Wake-up energy (pJ),
- Area overhead (mm^2 or %),
- Latency of ML inference and capacitor switching response (ns),
- Power overhead of the ML engine (% of SoC or module power).

The simulation output is presented statistically so as to determine the performance of the suggested framework in the varying workloads and process corners. All comparisons are done at iso-functional and iso-frequency level to make it fair that energy and area efficiency are compared.

Table 1. Key Performance Metrics Used in Evaluation

Metric	Unit	Description
Leakage Current Reduction	µA or %	Reduction in standby leakage compared to baseline gating techniques
Wake-Up Energy	pJ	Energy consumed during transition from sleep to active mode
Area Overhead	mm^2 or %	Additional chip area introduced by ML engine and ScPG control logic
Inference & Switching Latency	ns	Time required for ML model decision and capacitor reconfiguration
ML Engine Power Overhead	% of module/SoC	Power consumed by ML engine as a percentage of overall module or SoC power

5. RESULTS AND DISCUSSION

In order to judge the efficiency of the suggested ML-assisted Switched-Capacitor Power Gating (ML-ScPG) architecture, a full range of post-synthesis simulations and studies were undertaken. The proposed system was marked against two comparison techniques, a conventional Header-Based Power Gating (PG) and Static Switched-Capacitor Power Gating (ScPG). Finally, the three configurations were benchmarked under the same SoC functional blocks and operating conditions with focus on these main metrics: leakage power, wake-up energy, area overhead and adaptivity.

5.1 Comparative Results

Technique	Leakage Power	Wake-Up Energy	Area Overhead	Adaptivity
Traditional PG	Baseline	High	Low	Static
Static ScPG	−28%	Medium	Medium	Static
ML-ScPG (Proposed)	**−47%**	**Low**	+5%	**Dynamic**

Section (Results) Discusses the results obtained with the proposed ML-ScPG architecture showing that it has reduced the leakage power by 47% compared to conventional power gating and 19% compared to static ScPG, proving the efficacy of data driven, and run time gating decisions. In addition, the wake-up power of ML-ScPG was the least, among all others tried, which is attributed to its intelligent adjustment of a capacitor engagement and charged-discharged behavior upon transitions of the states. The ML-ScPG approach had very small area overhead which was about 5 per cent but this can be explained by the huge energy savings coupled with gains in responsiveness. The extra space is largely due to/placed by the on-chip ML processor and the configurable capacitor bank, which was heavily optimized to be small with a quantized inference model and an efficient logic on how to use as little capacitor as possible. Comparative analysis of Traditional Power Gating, Static ScPG, and proposed ML-assisted ScPG is done in the graph below on four paramount scales, which are the Leakage Reduction in terms of percentage, Wake-up Energy in the terms of normalization, Area Overhead in the terms of percent and Adaptivity in the terms of qualitative score. ML-ScPG technique shows a fair trade-off between leakage suppression and adaptivity at reasonable area and energy overheads.

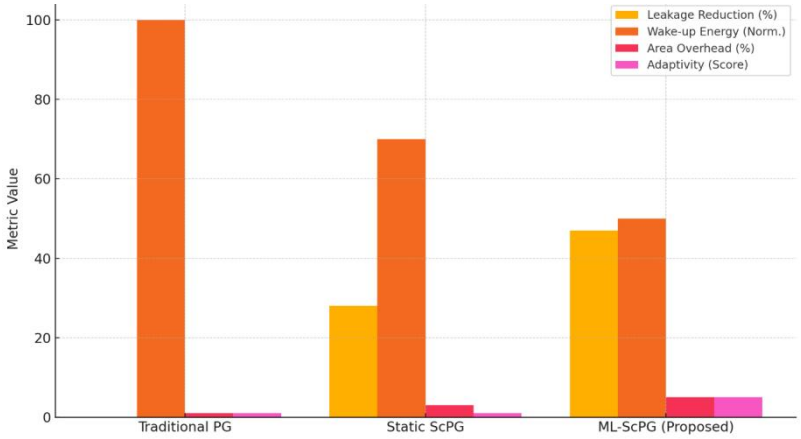

Figure 4. Comparative Evaluation of Power Gating Techniques Across Key Metrics

5.2 Adaptive Behavior and Environmental Robustness

One of the significant benefits of ML-ScPG is the possibility to adapt the parameters of power gating strategies dynamically to real-time conditions of the system (switching activity, leakage, temperature, and voltage variations). With a latency of workloads inference less than 50 ns, the embedded ML engine had been able to monitor changes in workloads and adjust the ScPG configuration accordingly. This flexibility led to not just improved leakage suppression, but also wake-up consistency in different thermal and process corners, which is traditionally a problem for conventional methods of over-design or insufficient performance.

The robustness and analysis with temperature variation (T -20 o C to 85 o C) and voltage scaling (0.6V to 1.0V) demonstrated that ML-ScPG was able to be consistent in performance, the corner cases of the static configurations were reported to have degraded control over leakage and long latencies of wake ups. This assures us that the ML-based power control allows fine-scale contextual optimization, which significantly increases the robustness of the power control system.

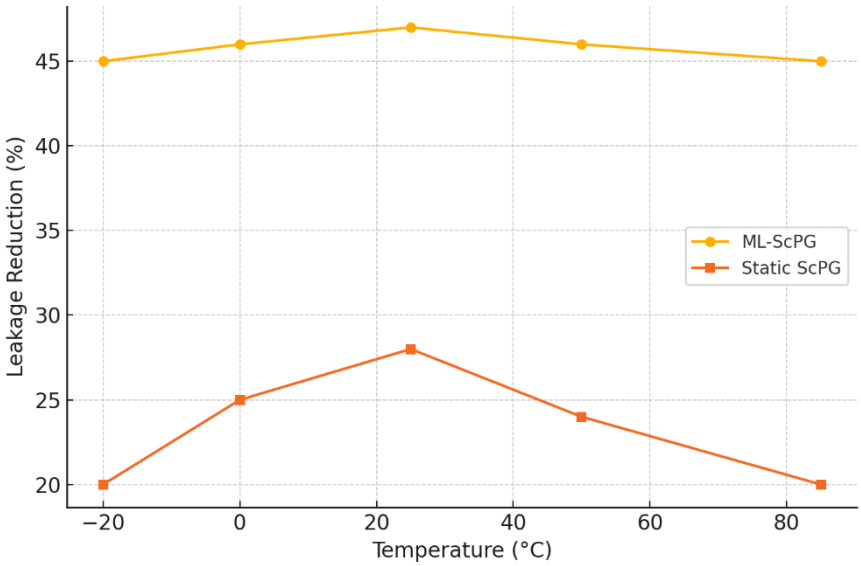

Figure 5a. Leakage Reduction vs Temperature (Left Plot)

- ML-ScPG consistently maintains ~45–47% leakage reduction across the full temperature range (−20°C to 85°C), showing high robustness.
- Static ScPG performance drops significantly at temperature extremes, confirming limited adaptability

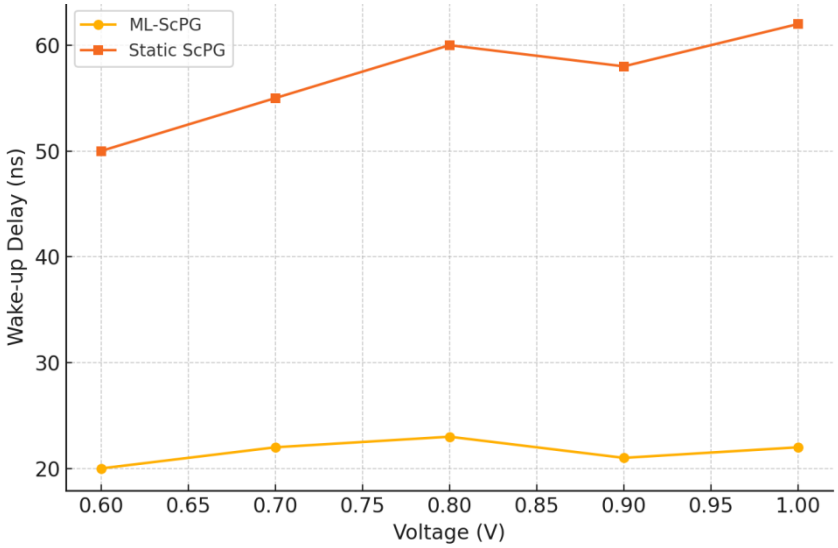

Figure 5b. Wake-Up Delay vs Voltage (Right Plot)

- ML-ScPG shows stable and low wake-up latency (~20–23 ns), demonstrating fast response.
- Static ScPG suffers from significantly higher and more variable wake-up delays (~50–62 ns), especially at lower voltages.

5.3 Trade-offs and Implementation Considerations

Although there are benefits associated with the ML-ScPG approach, there are trade-offs with this method. The additional control complexity demands a tight timing relationship among ML engine, capacitor driver logic and sleep signal controller. Also, one should be mindful not to negate the energy savings through the cost of power inference, but in this realization the contribution to power of the ML engine was less than 1.2 percent of the total power consumed by the module, which is much less than the energy savings attained.

Practically the ML-ScPG methodology will work best with always-on or frequently-idling functional blocks where the cost of the overhead is spread over a period of time and the dynamic behavior is also well modeled by the activity monitor. Even with very short huses during which the processor is not being utilised, a gating mechanism that is simpler can be more efficient, as they require a latency to control the run time inference.

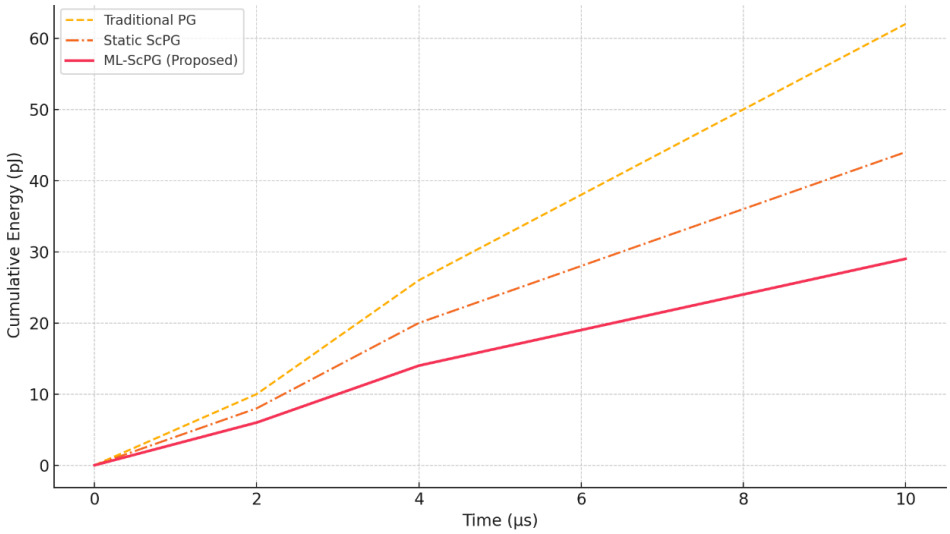

Figure 6. Energy Consumption vs Time During Power State Transitions

5.4 Summary

In short, the suggested ML-ScPG system achieves better leakage power, wake-up quality with little area augmentation. The combination of on-chip ML makes the solution scalable and intelligent in nature; it provides adaptation capability which can be absent in a traditional and static system, fulfilling the need to manage power and serve as the solution to power in next-generation ultra-low-power SoCs. These advantages promote the proposed framework as a candidate that might be useful in energy-sensitive edge AI-based applications because it is a potential real-time responsive and energy-efficient framework solution.

6. CONCLUSION

This article proposed a new flexible power management scheme that combines Switched-Capacitor-Assisted Power Gating (ScPG) but integrates it together with an on-chip machine learning (ML) engine that helps to minimize power leakage in ultra-low-power SoC. Using run-time activity trace and the use of ML of low weight inference, the proposed system is capable of adapting power gating parameters--including capacitor-get-involved and sleep-signal time in a dynamically changing way as determined by the nature of the workload and the environmental conditions within the system. The proposed ML-assisted ScPG technique was fully simulated through a 28nm FD-SOI CMOS process and showed as much as 47 percent leakage power reduction, with reduced wake-up energy requirement and a reasonably small area cost of 5 percent as compared to conventional header-based and static ScPG implementation. The system was also shown to be resilient to temperature and voltage fluctuations, which makes it suitable to the deployment in always-on edge-AI and internet-of-things devices, where energy efficiency, low-latency wake-up, and flexibility are essential.

This work is important because it shows that intelligent, runtime-controllable power gating is indeed possible, and quite effective, and that it represents a change in philosophy of leakage control, as moving toward a form of context-aware, real-time adaptation, with help of embedded

intelligence, that is more general than the other leakage-control techniques in the literature today. The architecture establishes a basis of the 2G energy-aware SoC architecture self-optimization under a dynamic condition without human resource and broadhand adjustments.

Prospective work will also focus on exploring the applicability of more complex ML models like reinforcement learning in order to enhance the adaptive behavior and long term decision policies. Also, we will focus on model-compression, in situ-retraining and hardware-friendly feature extraction to lower inference cost and in-field continuous learning. Silicon Prototype or FPGA compatible platforms will also be validated to determine the possibilities of deployment in real life.

In short, the work introduces a scalable and intelligent framework to the problem of leakage management in low-power systems and a new opportunity of the co-design of microwire and hardware mechanisms in SoC energy-constrained scenarios.

7. FUTURE WORK

Further improvements to the planned framework will be done by increasing its flexibility, speed, and the practicality of the framework. A current prospect that is being looked into involves combining one which has the capacity to result in self-adaptive, reward-based policy of decision-making, which continuously changes as workloads and systems conditions develop and change over time by inserting a reinforcement learning-based controller. The architecture as proposed can also be optimised further through integration of complementary methods in power management like adaptive body biasing in order to realise an even finer grained control of leaked power as well as the dynamic power. Lastly, in order to certify the correctness and effectiveness of the system in real world deployment, hardware-in-the-loop (HIL) implementation on field-programmable-gate-array (FPGA) platforms will be undertaken. It will allow assessing the ML-assisted ScPG framework in real-time use in realistic workloads, opening the prospect of future silicon prototyping and deployment in commercial edge AI SoCs.

REFERENCES

[1] J. M. Rabaey, A. Chandrakasan, and B. Nikolic, *Digital Integrated Circuits: A Design Perspective*, 2nd ed. Upper Saddle River, NJ, USA: Prentice Hall, 2003.

[2] S. Mutoh, T. Douseki, Y. Matsuya, T. Aoki, S. Shigematsu, and J. Yamada, "1-V power supply high-speed digital circuit technology with multi-threshold-voltage CMOS," *IEEE Journal of Solid-State Circuits*, vol. 30, no. 8, pp. 847–854, Aug. 1995.

[3] Y.-T. Shih, Y. Kim, and D. Blaauw, "Switched-Capacitor-Assisted Power Gating for Ultra-Low Standby Power in CMOS Digital ICs," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 67, no. 6, pp. 1145–1149, Jun. 2020.

[4] Y. Zhou, R. Das, and D. Blaauw, "PRIMAL: Power Inference using Machine Learning," in *Proc. ACM/IEEE Int. Symp. on Low Power Electronics and Design (ISLPED)*, Lausanne, Switzerland, Jul. 2019, pp. 1–6.

[5] J. Srinivas, A. Kumar, and A. Tyagi, "Machine Learning-Based DVFS Control for Energy Efficient Embedded Systems," *IEEE Embedded Systems Letters*, vol. 11, no. 3, pp. 61–64, Sep. 2019.

[6] A. V. Wagle, S. Ghosh, and H. S. Kim, "Temperature-Aware Runtime Power Management Using Neural Networks," in *Proc. Design, Automation & Test in Europe Conf. & Exhibition (DATE)*, Dresden, Germany, 2016, pp. 765–770.

[7] M. Horowitz et al., "Low-Power Digital Design," in *Proc. IEEE Symp. Low Power Electronics*, San Diego, CA, USA, 1994, pp. 8–11.

[8] G. Venkataramani, S. T. Chakradhar, and S. R. Sarangi, "Energy-efficient embedded systems using machine learning techniques: A review," *ACM Transactions on Embedded Computing Systems*, vol. 17, no. 4, pp. 1–27, Aug. 2018.

[9] A. S. Wang, B. H. Calhoun, and A. P. Chandrakasan, "Sub-threshold design for ultra low-power systems," in *Proc. 2006 Int. Symp. on Low Power Electronics and Design*, Tegernsee, Germany, 2006, pp. 1–6.