MANTECH Publications

Journal of Research in VLSI Design, Tools and Technology
Volume 10, Issue 1, January-April, 2025

Enhancing Power Efficiency in Very Large Scale Integration (VLSI) Circuits through Advanced Low-Power Design Techniques for Modern Electronic Systems

Anuja R. Deshmukh
Assistant Professor
Department of Electronics and Telecommunication Engineering
Nitte Mahalinga Adyantaya Memorial Institute of Technology
Email id: *anujardeshmukh@rediffmail.com*

Abstract

The growing demand for portable and high-performance electronic devices has placed significant emphasis on reducing power consumption in Very Large Scale Integration (VLSI) circuits. As transistor sizes shrink and device density increases, power efficiency becomes a primary concern in the design of modern integrated circuits. This paper explores the crucial low-power design techniques implemented in VLSI circuits to enhance power efficiency without compromising performance. It presents a comprehensive discussion on sources of power dissipation, the importance of low-power design in current and future technology, and the various techniques at the circuit, logic, and architectural levels used to mitigate power consumption. The study also addresses the key challenges and potential future directions in low-power VLSI design.

Special attention is given to the integration of machine learning in dynamic power management and the role of 3D ICs in reducing interconnect power losses. The findings of this paper aim to assist designers and researchers in developing sustainable, energy-efficient electronics for the next generation of devices.

Keywords: *Low-power design, VLSI, power optimization, CMOS technology, power dissipation, circuit-level techniques, energy-efficient design, leakage power.*

INTRODUCTION

With the rapid development of portable computing and communication devices, power efficiency has emerged as a primary constraint in Very Large Scale Integration (VLSI) circuit design. As modern systems increasingly demand higher functionality and performance within smaller, power-constrained environments—such as smartphones, wearable devices, and IoT modules—energy-aware design has become critical. Not only must designers consider processing speed and chip area, but they must also account for battery longevity, heat dissipation, and system reliability, which are all directly impacted by power consumption.

As technology nodes continue to scale down to sub-10nm and beyond, transistors have become smaller, faster, and more densely packed. However, this density brings about greater challenges, especially due to increased leakage currents, higher switching activity, and complex interconnect structures. These issues compound to make power a critical design bottleneck, not just a secondary consideration.

The evolution of VLSI technology from micrometer to nanometer scale has facilitated the integration of millions to billions of transistors on a single chip, enabling immense computational power.

However, this evolution has also introduced significant power-related problems, such as increased chip temperature, reduced signal integrity, electromigration risks, and diminished battery life in mobile applications. As a result, today's VLSI designers are compelled to explore innovative low-power design techniques that operate at all levels—device, circuit, architecture, and system—to ensure sustainable performance.

LITERATURE REVIEW

Over the years, numerous researchers and engineers have proposed methods to address the growing issue of power dissipation in VLSI circuits. In the early stages of CMOS development, the focus was predominantly on reducing dynamic power, as it was the primary contributor. Strategies like minimizing capacitance, reducing supply voltage, and optimizing logic transitions were widely adopted. Switching activity was a central point of concern, and early models revolved around reducing redundant transitions through logic restructuring and operand isolation.

As the technology scaled down beyond the sub-100nm region, the contribution of static (leakage) power increased significantly, sometimes accounting for over 40% of total power. This shift in power dynamics led to the development of power gating, multi-threshold CMOS (MTCMOS), and dual-Vdd techniques, which were instrumental in controlling leakage without compromising performance drastically.

Many researchers have also explored architectural and system-level techniques, such as clock gating, which disables the clock signal to idle modules, and dynamic voltage and frequency scaling (DVFS), which adjusts power consumption based on workload. These methods have proven particularly useful in microprocessors, where idle times are frequent.

In recent years, the focus has expanded to include emerging paradigms like approximate computing, which deliberately allows for minor inaccuracies to save power in error-tolerant applications such as image processing and machine learning. Moreover, machine learning-based power prediction models have been introduced to dynamically manage energy consumption at runtime. These advancements highlight the growing interest in intelligent, adaptive, and context-aware low-power VLSI systems that go beyond traditional techniques.

POWER DISSIPATION SOURCES IN VLSI CIRCUITS

Power dissipation in VLSI circuits primarily arises from three major sources: dynamic power, static (leakage) power, and short-circuit power. Understanding these sources is essential for selecting the appropriate low-power design techniques.

Dynamic Power Consumption

This is the most significant component of power in traditional CMOS circuits. It results from the charging and discharging of load capacitances during logic state transitions (0 to 1 or 1 to 0). The dynamic power PdynP_{dyn}Pdyn is mathematically expressed as:

Pdyn=αCLVdd2fP_{dyn} = \alpha C_L V_{dd}^2 fPdyn=αCLVdd2f

Where:

- α\alphaα = Switching activity factor
- CLC_LCL = Load capacitance
- VddV_{dd}Vdd = Supply voltage
- fff = Operating frequency

The square dependency on supply voltage makes voltage scaling an effective technique for dynamic power reduction. However, reducing voltage impacts speed, requiring careful trade-offs. Excessive clock frequencies and unnecessary transitions also increase dynamic power, emphasizing the need for efficient clock management and signal transition minimization.

Static (Leakage) Power Consumption

With technology nodes scaling down to below 45nm, leakage power has become a significant concern. It occurs even when the circuit is idle and no switching is taking place. Leakage arises from several mechanisms:

- **Subthreshold leakage:** Due to carriers tunneling through the transistor channel when Vgs < Vth.
- **Gate oxide tunneling:** Caused by electrons leaking through ultra-thin gate dielectrics.
- **Junction leakage:** From reverse-biased pn-junctions at the drain and source.

Leakage power is directly affected by threshold voltage (Vth) and device temperature. As devices operate faster with lower Vth, they become more leaky. Techniques like high-Vth transistors, transistor stacking, and power gating are widely used to manage static power.

Short-Circuit Power

This type of power dissipation occurs during logic transitions when both PMOS and NMOS transistors conduct simultaneously for a brief moment, creating a direct path from Vdd to ground. Though smaller in magnitude than dynamic and leakage power, it becomes significant in high-frequency designs with fast rise/fall times. The amount of short-circuit power depends on:

- Input transition time
- Output load
- Supply voltage

Short-circuit power can be reduced by optimizing the signal slope, proper transistor sizing, and buffer insertion techniques.

Table 1: Comparison of Power Dissipation Types in VLSI Circuits

Power Type	Cause	Formula / Dependence	Dominance (Technology Node)
Dynamic Power	Charging/discharging of capacitances	$P = \alpha CV^2f$	High in older technologies (>100nm)
Leakage Power	Subthreshold leakage, gate tunneling	Exponentially dependent on Vth and Temp	Dominant in modern nodes (<65nm)
Short-Circuit Power	Both NMOS & PMOS ON during transitions	Depends on input transition time & Vdd	Moderate

SWITCHING ACTIVITY REDUCTION TECHNIQUES

Clock Gating- Clock gating is one of the most widely used dynamic power reduction techniques. It disables the clock signal to specific circuit modules when they are inactive, thus preventing unnecessary transitions. Since the clock network contributes significantly to switching activity due to its continuous toggling, gating it when not required drastically reduces power consumption. Integrated clock gating (ICG) cells are commonly used in design flows to automate this technique at the RTL or gate level. Clock gating is often combined with power-aware synthesis tools to optimize gating conditions further.

Operand Isolation- Operand isolation involves using latches or AND gates to block the propagation of input signals to functional units that are not currently being used. This prevents unneeded internal switching in complex blocks like multipliers or ALUs. It is particularly useful in datapath designs, where only a subset of functional units is active at any given time. By masking signal transitions, operand isolation can minimize glitching and spurious activity in idle blocks, offering both power and performance improvements.

Bus Encoding- Bus encoding schemes aim to reduce the number of bit transitions that occur on interconnect buses. Since every 0-to-1 or 1-to-0 transition consumes energy, encoding the data to reduce hamming distance between successive values can lead to power savings. Common methods include Gray coding, where only one bit changes at a time, and Bus-Invert coding, which inverts the data if the number of switching bits exceeds a threshold. These techniques are especially effective in address and data buses of microcontrollers and processors, where high-frequency toggling is frequent.

SUPPLY VOLTAGE REDUCTION TECHNIQUES

Dynamic Voltage Scaling (DVS)- DVS adjusts the supply voltage dynamically based on workload demands. When full performance is not needed, the voltage and clock frequency are reduced, thereby lowering dynamic power quadratically. DVS is highly effective in systems like mobile processors, where performance requirements vary significantly depending on the task. Implementing DVS requires voltage regulators and control logic to ensure seamless transitions between voltage levels. However, care must be taken to avoid timing violations and performance degradation when voltage is lowered excessively.

Multi-Vdd Design- In multi-Vdd architecture, the chip is partitioned into voltage islands, each operating at an optimal supply voltage based on its performance-criticality. For example, a high-speed processor core may run at 1.0V, while peripheral blocks like timers or UARTs may operate at 0.8V or lower. Level shifters are used at the interfaces between different voltage domains to maintain signal integrity. Though it introduces design complexity and area overhead, multi-Vdd provides a balance between energy efficiency and performance by localizing power optimization.

THRESHOLD VOLTAGE MANAGEMENT TECHNIQUES

Multi-Threshold CMOS (MTCMOS)

MTCMOS is a well-established method for reducing subthreshold leakage power. It utilizes low-Vth transistors in critical timing paths to maintain speed, while high-Vth transistors are used in non-critical paths or for power-gating purposes. Sleep transistors (either PMOS or NMOS) are often introduced between power rails and the logic block to control leakage when the circuit is idle. MTCMOS is especially beneficial in standby-heavy applications like memory blocks and embedded systems, where leakage power can dominate.

Variable Threshold CMOS (VTCMOS)

VTCMOS dynamically modifies the threshold voltage of a transistor using body biasing techniques. Forward body biasing (FBB) reduces Vth for performance gain, while reverse body biasing (RBB) increases Vth to reduce leakage during idle states. This dynamic control offers flexibility in balancing power and delay, especially in adaptive systems where the operating condition changes over time. VTCMOS requires triple-well or twin-well CMOS

processes for independent body control, which adds to fabrication complexity but significantly boosts energy efficiency.

LEAKAGE POWER REDUCTION TECHNIQUES

Power Gating- Power gating effectively cuts off the power supply to idle or standby circuit regions using high-Vth sleep transistors. When a block is not in use, it is isolated from the Vdd or ground rail, eliminating leakage almost entirely. Upon wake-up, the power is re-applied, and the block resumes operation. This technique is heavily used in microcontrollers, mobile SoCs, and memory arrays, where idle periods are frequent. The key challenge lies in minimizing wake-up latency and inrush current, which can cause performance hiccups or voltage droop.

Transistor Stacking- The transistor stacking effect exploits the fact that stacked OFF transistors leak less current than a single OFF transistor. When multiple transistors are connected in series and turned off, the intermediate node voltages cause a reduction in Vgs across individual transistors, thereby reducing subthreshold leakage. Though stacking increases area and delay slightly, it is a low-cost leakage mitigation technique suitable for low-speed or background logic. It's also used in combination with high-Vth devices or sleep modes for greater efficiency.

Table 2: Comparison of Leakage Reduction Techniques

Technique	Method	Area Overhead	Performance Impact	Leakage Reduction
Power Gating	Disconnect power supply to blocks	High	Low during active	Very High
MTCMOS	Use of high-Vth transistors	Moderate	Minimal	High
Transistor Stacking	Series connection of transistors	Low	Can be significant	Moderate

Description: This table compares popular leakage reduction methods in terms of area cost, performance hit, and effectiveness in reducing leakage.

ARCHITECTURAL LEVEL POWER OPTIMIZATION

Parallelism and Pipelining- Parallelism and pipelining are fundamental architectural strategies for improving throughput and energy efficiency. Parallelism involves executing multiple instructions or processes simultaneously, reducing total execution time. On the other hand, pipelining breaks down the execution process into stages, allowing multiple instructions to be processed in an overlapping manner. These techniques not only enhance performance but also allow the system to operate at lower clock frequencies and supply voltages without degrading throughput, thereby reducing dynamic power consumption. Modern multicore processors and DSP architectures leverage both strategies to meet power-performance trade-offs efficiently.

Resource Sharing- Resource sharing is a cost-effective approach to power reduction, especially in application-specific integrated circuits (ASICs) and reconfigurable hardware like FPGAs. By allowing multiple operations or tasks to use a common arithmetic or logic unit, designers can minimize redundant hardware and lower the total capacitance switching during operation. This reduces both dynamic and leakage power while also saving area. Proper control logic ensures that resources are allocated efficiently without causing timing bottlenecks or performance degradation.

Voltage Islands- Voltage islands allow independent voltage domains within a single chip, so each region can operate at the minimum required voltage level for its functionality. Performance-critical blocks such as CPU cores or memory controllers can use higher voltages, while non-critical blocks like I/O controllers or timers operate at lower voltages. This strategy enables granular power optimization, especially in System-on-Chip (SoC) designs. However, implementing voltage islands requires additional components such as level shifters, isolation cells, and domain-specific power management units, which can increase design complexity and verification efforts.

DESIGN AUTOMATION AND EDA SUPPORT

Power-Aware Synthesis Tools- Modern EDA tools now incorporate power-aware design methodologies throughout the digital design flow—from RTL synthesis to physical design. These tools allow designers to specify power constraints, such as power budgets, clock gating opportunities, and multi-voltage domains, early in the process. Using tools like Synopsys

Design Compiler, Cadence Genus, or Siemens Calypto, engineers can generate optimized netlists with clock gating, power gating, and other low-power structures automatically inserted. This ensures that power optimization is not just a post-layout step but an integrated part of the design cycle.

RTL Power Estimation and Optimization

Early-stage power estimation at the Register Transfer Level (RTL) allows designers to make architecture-level decisions before moving to costly gate-level or physical implementations. RTL power estimation uses models based on switching activity, capacitance, and logic structure to predict power consumption. By incorporating power analysis into early design stages, inefficient logic paths, high-activity nets, and redundant computations can be identified and corrected early. This proactive approach improves design productivity and reduces overall design time and iterations.

CHALLENGES IN LOW-POWER VLSI DESIGN

Increased Design Complexity Low-power techniques often require the introduction of additional control logic, clock gating circuitry, voltage islands, and sleep transistors, which increases both the area and complexity of the design. This complicates not only circuit-level design but also functional verification, power-aware simulation, and timing analysis. Managing corner cases and interactions between power domains adds further burden during debugging and integration.

Trade-offs Between Power, Performance, and Area (PPA)

Optimizing for low power almost always involves a trade-off. For example, reducing supply voltage slows down the circuit, while increasing transistor stacking or using high-Vth devices can degrade timing performance. Similarly, implementing multiple voltage domains or power gating structures adds area overhead. Finding the right balance between these competing design goals—performance, power, and area—requires careful architectural exploration and algorithmic decision-making.

Process Variations As transistor sizes shrink, manufacturing process variations—such as threshold voltage fluctuation, channel length variation, and oxide thickness inconsistencies—can significantly affect power consumption, delay, and yield. These variations make power

estimation and optimization unpredictable, often requiring the use of adaptive body biasing, on-chip sensors, and robust timing margins to ensure reliable operation across all process corners.

Thermal Management Despite advances in power reduction, high transistor density and increased switching activity in modern chips generate substantial heat. If not properly managed, thermal issues can lead to reliability degradation, reduced performance, or even permanent device failure. Hence, low-power designs must also be thermal-aware, using techniques such as thermal-aware floor planning, dynamic thermal management, and heat-spreading materials during packaging.

SCOPE FOR FUTURE RESEARCH

Machine Learning-Based Power Management AI and machine learning techniques are increasingly being adopted for predictive power modeling and adaptive power control. These models can learn from runtime data to optimize clock frequency, supply voltage, and resource allocation in real time. Reinforcement learning algorithms, for instance, are being explored to implement energy-efficient dynamic voltage scaling in real-world applications. This paves the way for self-optimizing circuits that adapt to environmental changes and workload patterns.

Approximate Computing For applications where perfect accuracy is not mandatory—such as multimedia processing, sensor data fusion, and AI inference—approximate computing enables aggressive power savings. By relaxing the correctness of computations, it becomes possible to reduce logic complexity, transition activity, and memory accesses. This emerging paradigm is particularly useful in edge AI devices, where battery life and computation speed are more critical than absolute precision.

3D Integration and Emerging Devices- 3D ICs involve vertical stacking of logic and memory dies, reducing wire length and improving power efficiency by lowering interconnect delays and capacitive loading. Emerging devices like FinFETs, Gate-All-Around (GAA) FETs, and Tunnel FETs also offer improved energy efficiency due to their steep subthreshold slopes and reduced leakage. Research in these areas promises to overcome the limitations of planar CMOS, pushing the boundaries of low-power nanoelectronics.

Energy Harvesting and Self-Powered Systems- In the context of IoT and wearable electronics, energy harvesting systems use ambient energy sources like solar, RF, piezoelectric, or thermal gradients to power circuits. These ultra-low-power systems demand specialized design techniques that ensure operation even under fluctuating power availability. Circuit blocks must be capable of voltage scaling, energy-aware scheduling, and power-aware data retention, making this a fertile area for innovation.

CONCLUSION

The increasing demand for portable, high-performance, and energy-efficient devices has placed low-power VLSI design at the forefront of modern electronics. As device geometries continue to shrink and performance expectations rise, power optimization at every design level—from transistors to system architecture—has become a necessity rather than an option.

This paper has outlined critical techniques for reducing dynamic and static power, including switching activity minimization, supply and threshold voltage management, architectural innovations, and EDA-supported automation. Despite these advances, challenges such as design complexity, process variations, and thermal issues persist. The scope for future innovation lies in embracing machine learning, approximate computing, 3D integration, and self-powered circuit designs.

With continuous research and interdisciplinary collaboration, the future of low-power VLSI holds promising solutions for sustainable, high-performance computing.

REFERENCES

1. Chandrakasan, A. P., Sheng, S., & Brodersen, R. W. (1992). Low-power CMOS digital design. IEEE Journal of Solid-State Circuits, 27(4), 473–484.
2. Rabaey, J. M., Chandrakasan, A., & Nikolic, B. (2003). Digital integrated circuits: A design perspective (2nd ed.). Prentice Hall.
3. Roy, K., Mukhopadhyay, S., & Mahmoodi-Meimand, H. (2003). Leakage current mechanisms and leakage reduction techniques in deep-submicrometer CMOS circuits. Proceedings of the IEEE, 91(2), 305–327.

4. Kim, C. H., & Roy, K. (2002). Dynamic Vt SRAM: A leakage tolerant cache memory for low voltage microprocessors. Proceedings of the International Symposium on Low Power Electronics and Design (ISLPED), 251–254.
5. Mutoh, S., et al. (1995). 1-V power supply high-speed digital circuit technology with multithreshold-voltage CMOS. IEEE Journal of Solid-State Circuits, 30(8), 847–854.
6. Hwang, C. H., & Wu, Y. (2015). Design of low power systems using multi-voltage techniques. Microelectronics Journal, 46(10), 1031–1040.
7. Mohan ram, K., & Touba, N. A. (2003). Cost-effective approach for reducing power consumption in scan testing. Proceedings of the VLSI Test Symposium (VTS), 160–165.
8. Powell, M. D., et al. (2000). Gated-Vdd: A circuit technique to reduce leakage in deep-submicron cache memories. Proceedings of ISLPED, 90–95.
9. Narendra, S., & Chandrakasan, A. (2006). Leakage in nanometer CMOS technologies. Springer.
10. Wang, A., Calhoun, B. H., & Chandrakasan, A. (2006). Sub-threshold design for ultra-low-power systems. Springer.
11. Liu, W., et al. (2010). A survey of clock gating techniques in power-aware processors. ACM Computing Surveys, 42(3), 1–36.
12. Magen, N., et al. (2004). Interconnect-power dissipation in a microprocessor. Proceedings of ISLPED, 263–266.
13. Ye, Y., Borkar, S., & Roy, K. (1999). A technique for standby leakage reduction in CMOS circuits. Proceedings of the IEEE International Conference on Computer Design, 82–87.
14. Hanchate, N., & Ranganathan, N. (2004). A new technique for leakage reduction in CMOS circuits using self-controlled stacked transistors. IEEE Transactions on VLSI Systems, 13(5), 503–513.
15. Coskun, A. K., Ayala, J., & Atienza, D. (2013). Dynamic thermal management in 3D multicore architectures. IEEE Design & Test, 30(3), 76–84.
16. Zhang, Y., et al. (2016). Machine learning-based power management for embedded systems. ACM Transactions on Embedded Computing Systems, 15(4), 1–26.
17. Han, J., & Orshansky, M. (2013). Approximate computing: An emerging paradigm for energy-efficient design. European Test Symposium, 1–6.