National Journal of Advanced VLSI Design and Systems • ISSN 3139-2261 (O) • Jan - Dec • 1 • 1 • 2026 (pp. 38-45)

RESEARCH ARTICLE

Lifetime-Energy-Optimal VLSI Architectures: A Unified Thermal–Reliability–Aware Design Framework for Sustainable Embedded Systems

Charpe Prasanjeet Prabhakar*

Department Of Electrical And Electronics Engineering, Kalinga University, Raipur, India

KEYWORDS:
Lifetime energy optimization, Thermal-aware VLSI design, Reliability-aware architectures, Near-threshold computing, Sustainable embedded systems, Energy-efficient VLSI

ARTICLE HISTORY:
Received : 21.11.2025
Revised : 25.12.2025
Accepted : 06.01.2026

Abstract

These have been made visible due to the appeal of always-on embedded and edge systems with their unveiling of the latent frailty of the conventional methods within the low-power VLSI design centered on instantaneous power and PPA enhancement without regard to the long-term thermal and reliability effects. Leakage in scaling Technologies that have significant scaling behaviour are highly sensitive to temperature, which favours the growth of thermal hotspots and ageing effects, both of which radically reduce the lifetime of systems, leading to higher cumulative power consumption and decreased sustainability. In this paper, a lifetime-energy-optimal VLSI architecture that integrates thermal and reliability consideration in energy optimization is proposed. The model whole lifetime energy consumption is constructed to indicate the combined effect of dynamic power, temperature-related leakage, and degradation in lifetime due to reliability. More to this model, a cross-layer optimization framework may be employed to organize voltage/frequency scaling, policy of power-gating, and workload mapping with precise reliability constraints in order to be able to perform with adaptive near-threshold operation and mitigate thermal stress and aging. The experimental results of the common embedded workloads demonstrate the proposed solution conserves the total lifetime energy consumption and higher mean time to failure compared to power-centric and thermal-only prototypes. The findings validate the lifetime-centric optimization to be among the key paradigms in order to attain sustainable and energy efficient VLSI architectures.

Author's e-mail: charpe.prasanjeet.prabhakar@kalingauniversity.ac.in

How to cite this article: Prabhakar CP. Lifetime-Energy-Optimal VLSI Architectures: A Unified Thermal–Reliability–Aware Design Framework for Sustainable Embedded Systems. National Journal of Advanced VLSI Design and Systems. Vol. 1, No. 1, 2026 (pp. 38-45).

Introduction

The high growth rate of continuously operational and data intensive embedded systems have fundamentally changed the design imperatives of advanced very-large-scale integration (VLSI) systems.[1, 2] Internet-of-Things (IoT), sensor nodes, edge artificial intelligence accelerators, wearable biomedical devices, and cyber-physical systems are examples of applications that require long periods or continuous operation with extreme energy and thermal requirements.[6, 12] Conventionally, energy efficient VLSI design has emphasized on minimization of instantaneous power consumption based on techniques of voltage scaling, clock gating and power gating.[1, 9] Although these methods have proved useful in previous generations of technology, their usefulness is becoming limited as far as highly scaled CMOS processes are concerned0[1, 5] With the ongoing reduction in the size of technology nodes, leakage power and power density have become the leading factors in the total energy consumption.[1, 5]. The leakage currents are very temperature sensitive and cause positive feedback between power dissipation and on-chip temperature resulting in localized thermal hotspots.[6, 9, 10] High operating temperatures, on the other

hand, can increase important reliability degradation processes, such as bias temperature variation, hot-carrier injection, electromigration, and time-dependent dielectric breakdown.[3, 7] These effects gradually destroy device and interconnect properties leading to loss of performance, timing failures and eventually system lifetime.[3, 5] In this regard, power optimization in isolation can be counterproductive. Near-threshold operation and aggressive voltage scaling can minimise short-term energy use, but may substantially increase variability and ageing .[5,7] and reduce usable lifetime and total energy used during the lifetime of the system as a result of premature replacement or increased cooling overheads[6] and.[9] Therefore, real energy efficiency should not be considered on a one point on the operating point of a system but on the operational life of a system.[1, 6]. Driven by these issues, this paper suggests a design paradigm based on the lifetime-energy-centric concept of VLSI architectures. The suggested approach will help reduce the overall lifetime energy use by collectively addressing the questions of power, temperature and reliability in order to guarantee acceptable system lifespan.[6, 9, 10] In this paper, a single lifetime energy model and a thermal-reliability-conscious architecture structure are presented, that can support adaptative near-threshold operation in a way that does not compromise on long-term energy sustainability, thus filling a significant vacuum in the modern energy efficient VLSI design.[5, 6]

Background and Motivation

It is inherently connected with energy consumption, thermal behavior and reliability because the modern VLSI systems, and in particular, with the deeply scaled CMOS technologies.[1, 5] Power dissipation due to switches is dynamic and it rises with the switching activity and frequency used, but leakage power has become a leading source at lower supply voltages and higher temperature levels.[1, 6] Leakage currents increase exponentially with temperature, and when greater amounts of power are absorbed in the on-chip temperature rise, leakage power increases, which in turn raises on-chip temperature, which increases leakage power, forming a positive feedback loop.[6, 9] This communication results in the localized thermal hotspots that deteriorate the performance and erode the energy efficiency.[6, 9, 10] Long-term reliability is directly and significantly affected by the thermal stress. Other dominating degradation mechanisms that include bias temperature instability, hot-carrier injection, electromigration, and time-dependent dielectric breakdown progress exponentially with temperature.[3, 7] Even slight rises in temperature can hence result in a significant decrease in average time to fail, particularly in embedded systems that are powered continuously where aging-effects continue to build up over time.[3, 6] Due to this, thermal behavior has moved to a first order design consideration in energy efficient VLSI architectures.[6, 9] Traditional low-power design tools are mainly focused on instantaneous power reduction and power/performance/area optimization by using voltage scaling, clock gating, and power gating [1, 9] technologies. Although suitable in older technology nodes, these methods tend to overlook the long-term repercussion of vigorous power reduction.[1, 5] Specifically, near- and sub-threshold operation, though potentially appealing in regard to minimising dynamic energy, is much more sensitive to process, voltage and temperature changes, and can reinforce the age-related degradation.[5, 7] Designs that are effective only in the short-term in terms of energy efficiency can thus end up having a short life span leading to increased overall energy use because of its premature failure or higher cooling cost.[6, 9] These shortcomings drive the desire to have a more comprehensive energy efficiency measure which embodies the long-term system behaviour. Lifetime energy consumption which is the total energy used in the useful life of a system is a better indicator of sustainability of embedded and always-on systems.[1, 6] To maximize lifetime energy the tradeoff between instantaneous power savings, thermal stress and reliability degradation must be done jointly.[6, 9, 10] This view provides the rationale behind the lifetime-energy-based design system suggested in this paper that incorporates thermal- and reliability-consciousness to develop sustainable and energy-efficient VLSI systems.[5, 6]

Lifetime Energy Modeling Framework.

A quantitative model reflecting the interactions between power consumption, temperature and reliability is needed to allow lifetime-energy-centric optimization. In contrast to traditional energy models that consider energy per operation or energy per unit time, the described framework has taken into consideration clearly the fact that operating conditions have a direct impact on the instantaneous power loss as well as the operational life of the system. In this section, there is an introductory discussion of a lifetime energy modeling framework which constitutes the analytical basis of the proposed structure.

Lifetime energy consumption is the total energy used by a VLSI system during the operational life of the system. It combines dynamic and leakage power over the lifetime that the system can stay in use with leakage power being very temperature sensitive. This connection is reflected in the following Equation (1):

$$E_{life} = \int_0^{T_{life}} \left(P_{dyn}(t) + P_{leak}\left(T(t)\right) \right) dt \qquad (1)$$

Here, denotes the total lifetime energy, is the dynamic power consumption, represents temperature-dependent leakage power, and is the effective operational lifetime. This expression puts emphasis on the fact that minimizing instantaneous power may not always consider lifetime energy in case thermal effects reduce lifetime of the system. The mechanisms of reliability degradation under the influence of temperature are significant factors in determining the system lifetime. Dependence of the

operating temperature on the mean time to failure (MTTF) can be modeled by an Arrhenius-type relationship, which can describe the accelerating exponential relationship between age and temperature. This is the behavior and can be one of Equation (2):

$$MTTF \propto exp\left(\frac{E_a}{k \cdot T}\right) \quad (2)$$

In this equation is the activation energy of the dominant failure mechanism, is Boltzmann's constant, and is the absolute operating temperature. As shown in equation (2), small temperature changes can drastically decrease the lifetime of the system and hence lifetime energy consumption can be increased despite temporary power reductions. Via the aforementioned relationships lifetime-energy-optimal design can be presented as a constrained optimization problem which simultaneously takes into account power, temperature and reliability. The design goal is to reduce the total lifetime energy with a constraint on the reliability, and it can be stated as Equation (3):

$$\min_{\{V,f,\pi\}} E_{life} \; s.t \; MTTF \geq MTTF_{min} \quad (3)$$

Here, and denote the operating voltage and frequency, respectively, and represents the power management policy, including voltage scaling, power gating, and workload scheduling. The constraint involves making sure the system lifetime is above a necessary minimum amount. With a combination of Equations (1)3), an overarching analysis framework is formed which integrates instantaneous power optimization and thermal analysis with long-term reliability, which forms the foundation of the lifetime-energy aware architecture methods explained in the following sections.

Monolithic Thermal Reliability Architecture.

The given section introduces the suggested unified thermal and reliability conscious VLSI architecture that makes an overt attempt to combine power management, thermal monitoring, and reliability modeling to obtain lifetime-energy optimal operation in the embedded networks. Contrary to other more traditional architectures where power and temperature and aging are viewed as decoupled issues, the proposed design provides a system-level coordination mechanism of a close loop and is continually adjusted in its operating conditions by real time thermal and reliability feedback. The architecture shown in Figure 1 is structured with several voltage domains (VDD1, VDD2, and VDD3) that contain heterogeneous processing units which include general-purpose cores and dsp units. The third domain is in near-threshold operation mode that allows aggressive saving of energy of inactive or non-critical workloads. A central controller is used to coordinate voltage/frequency scaling (DVFS) within these domains in order to enable the fine-grained control of the performance and power dissipation.

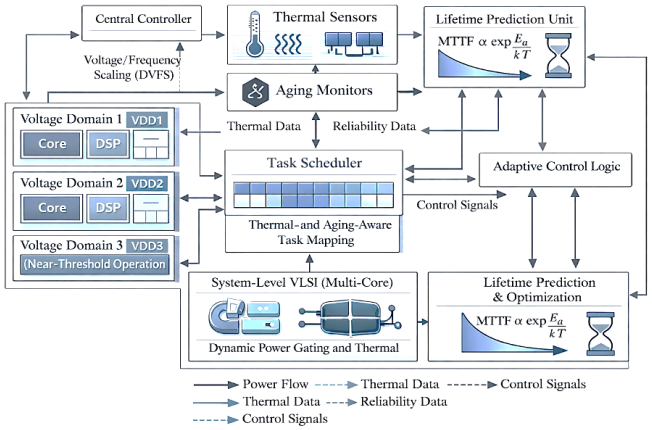

Fig. 1: System-Level Thermal–Reliability–Aware VLSI Architecture

Thermal sensors on the chip are used to give the temperature as a continuous reading and long-term reliability pointers given by ageing monitors one is related to the degradation mechanisms scheme which include BTI, HCI and other schemes. These thermal and reliability streams are sent to a lifetime prediction unit which makes use of this data to predict the effect of present operating conditions on the mean time to failure taking into consideration temperature accelerated ageing models. The resulting lifetime estimates are then used to drive the adaptive control logic with the aim of ensuring that the optimization in energy decisions does not violate constraints in reliability. A smart task scheduler is thermal and ageing conscious and dynamically provides workloads to voltage domain according to their performance criticality, thermal condition, and ageing buffer. This scheduler works together with the on-chip thermal management unit which uses dynamic power gating and thermal throttling in order to limit hotspots and avoid accelerated aging. Reliability models give extra feedback to fine-tune control measures and make active lifetime management instead of recovery after faults. In general, the architecture presented in Figure 1 forms a closely-knit feedback between power control, thermal sensing, reliability modelling and workload management. This single team strategy allows adaptive near-threshold operation, with acceptable system lifetime, which is the basis of the lifetime-energy optimization framework to be assessed in the following sections.

Cross-Layer Framework Optimization.

Here, the offered section of the framework detailing the cross-layer optimization is presented, whereby power management, thermal management, and reliability awareness are aligned on various abstraction levels to ensure a lifetime optimal operation in a power-energy wise fashion. The suggested framework, as opposed to the traditional methods where the optimization of each layer is performed separately of the others, indicates a single feedback loop that constantly changes system

behaviour depending on current operating environments and long-term ageing patterns. As shown in Figure 2, the framework incorporates the control and monitoring mechanisms across circuit, micro architectural and system levels. Dynamic voltage and frequency scaling (DVFS) at the circuit level is performed by logic in the central voltage/frequency control in order to control instantaneous power consumption and performance. The decisions on the voltage scaling, gating of power and scheduling of tasks are coordinated by higher-level power policy and optimization modules driving these controls.

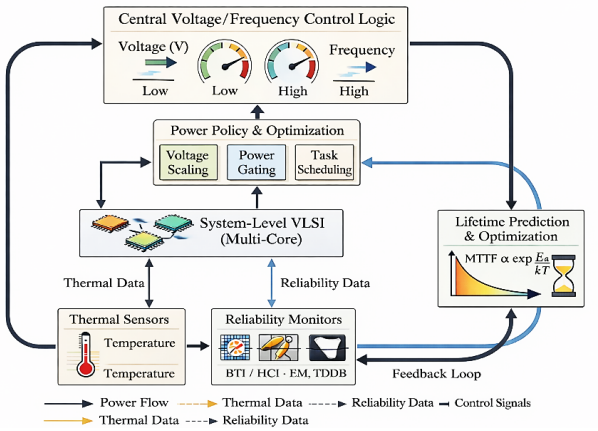

Fig. 2: Cross-Layer Optimization Flow: Power–Thermal–Reliability Feedback Loop

Workloads can be executed at the system level (a multi-core VLSI platform), where both thermal stress and time dependency reliability degradation are caused by power dissipation. Distributed thermal sensors give real time temperature measurements and reliability sensors follow ageing indicators linked to prevailing failure mechanisms, which include bias temperature instability (BTI), hot-carrier injection (HCI), electromigration (EM) and time-dependent dielectric breakdown (TDDB). These streams are relayed on to a lifetime prediction and optimization unit that predicts the effect of the present operating conditions on mean time to failure based on temperature accelerated ageing models. The estimated lifetime measures are given as input to the adaptive control layer and this completes the power-thermal-reliability feedback loop as in Figure 2. The closed loop interrelationship allows proactive modification of the voltage, frequency and power management policies, to reduce thermal hotspots, minimize accelerated aging and maintain a long life cycle of the system. The proposed framework will maximise collaboration between layers in decision optimization such that decisions do not shorten system lifespan, constituting a major foundation of the lifetime-energy-centric design philosophy that this paper has advanced.

Experimental Methodology

This section includes the description of the experimental methodology to be used to test the proposed lifetime-energy-optimal VLSI architecture in real under-embedded-system operating conditions. The power consumption, thermal behaviour, and reliability degradation are the coupled effects, the evaluation framework targets to capture and allow the comprehensive assessment of the lifetime energy efficiency and the system longevity. Tests are run in the presence of a highly scaled low-power CMOS technology node, which is the reflection of the modern-day embedded platforms. Dynamic leakage power and temperature-dependent leakage power are modelled, as well as long-term reliability effects, through the use of well-established ageing models, such as bias temperature instability (BTI), hot-carrier injection (HCI), electromigration (EM), and time-dependent dielectric breakdown (TDDB). The models make it possible to estimate mean time to failure, controlled by operating voltage, frequency, and temperature. The main workloads that are taken into consideration are always-on IoT sensing, which focuses on low-power usage, and Edge-AI inference which creates a burden on both computing power and temperature. A combination of these workloads is able to run various operating regimes such as near-threshold execution and dynamic voltage frequency scaling. Table 1 summarises the main technology assumptions and simulation parameters that will be used in the entire evaluation process so that all experiments are consistent.

Table 1: Experimental Setup Parameters

Parameter	Value / Description
Technology Node	Advanced low-power CMOS
Voltage Range	Near-threshold to nominal
Frequency Range	DVFS-enabled
Thermal Model	Lumped RC model
Aging Models	BTI, HCI, EM, TDDB
Lifetime Constraint	Minimum MTTF threshold

In order to measure the advantages of the proposed solution, various architectural designs are considered. It can consist of a traditional power-centric baseline, a thermal-aware architecture, which does not explicitly think about reliability, and the suggested temporal thermal-

Table 2: Compared Architectures

Architecture	Power Awareness	Thermal Awareness	Reliability Awareness
Baseline	✔	✘	✘
Thermal-Aware	✔	✔	✘
Proposed	✔	✔	✔

reliability-aware design. The peculiarities of these structures are outlined in Table 2.

The combination of the experimental structure and architectural comparisons as summarised in Tables 1 and 2 give a rigorous platform on the effects of cross-layer thermal and reliability awareness when it comes to lifetime energy consumption and system reliability in embedded VLSI systems.

Results and Analysis

Energy vs. Lifetime Trade-Off

This sub-subsection examines the trade-off between energy reduction in the short and energy consumption in the long term due to the effect of reliability degradation and ageing. Although the use of aggressive voltage scaling has become very popular to reduce instantaneous power use, its effect on system lifetime can greatly affect the overall energy efficiency in cases considered over long periods of operation. Figure 3 demonstrates how total lifetime energy consumption correlates with operating voltage of three architecture designs including a traditional power-based baseline, thermal-conscious architecture, and the design that is thermal-reliability-conscious. The baseline architecture will have high lifetime energy consumption and even less power at lower supply voltages, especially at the near threshold operating region. This is explained by high leakage sensitivity and faster ageing that can be short cut system lifetime, cause cumulative energy consumption.

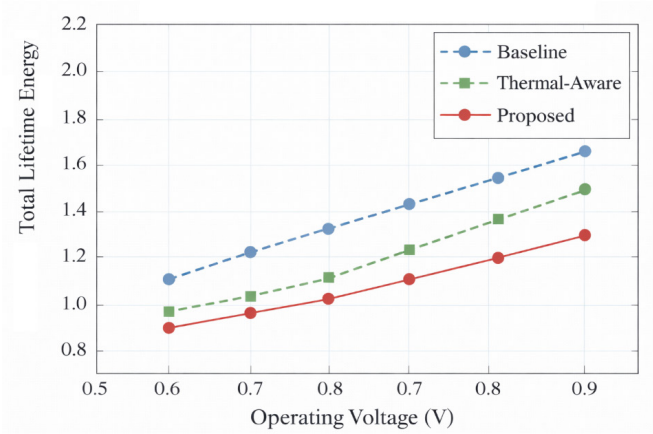

Fig. 3: Total Lifetime Energy Consumption as a Function of Operating Voltage

Thermal-aware architecture is exhibiting moderate enhancement as it alleviates the effects of hotspots on degradation, but it still does not introduce limiting reliability constraints. By coordinating power, temperature, and age effects, in contrast, the proposed architecture is able to obtain the lowest lifetime energy at any volunteer voltage. These findings support the finding that small changes in instantaneous energy directed by reliability awareness can produce significant changes in additional lifetime energy. Overall, Figure 3. A point that energy efficiency should not be considered at one operating point but should be considered over the lifetime of the system, which confirms the usefulness of the proposed lifetime-energy-centric design method.

Reliability/ Thermal Impact.

In this subsection, the effectiveness of the proposed architecture in terms of reducing thermal hotspots and enhancing reliability on a long-term basis are evaluated. One of the major causes of accelerated ageing processes in heavily scaled CMOS technologies is high temperatures of operation, which lowers the mean time to failure (MTTF) directly. Consequently, temperature minimization whilst maintaining performance is important to sustainable VLSI system design. Figure 4 relates the average operating temperature and normalised MTTF of three architectural designs, namely, baseline, thermal-conscious, and thermal-reliability-conscious architecture. The baseline design is the hottest in operation resulting in the lowest MTTF because of the stronger thermal stress and increasing ageing rate. The thermal-conscious architecture attains the moderate degree of temperature decrease by applying the methods of hotspots alleviation, which leads to the subsequent enhancement of the MTTF. But since it does not provide explicit reliability feedback its gains in lifetime are limited.

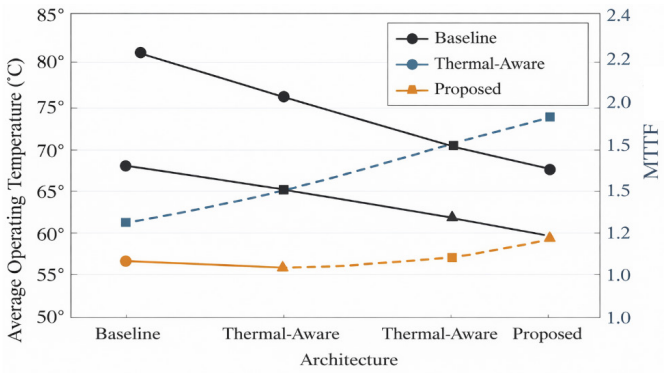

Figu. 4: Temperature and Mean Time to Failure (MTTF) Comparison Across Architectures

Conversely, the suggested architecture exhibits the greatest decrease in operating temperature, in effect eliminating hotspots that could lead to thermal neurosis by integrating the dynamic voltage and frequency actions as well as an on-chip thermal system. This temperature drop is then directly associated with significant increase in MTTF which justifies the merits of ensuring reliability-consciousness becomes part of the optimization loop. This positive correlation between temperature and MTTF in Figure 4 is expected regarding the accelerated ageing rate of the Arrhenius model of ageing. Table 3 shows a quantitative summary of the thermal and reliability gains, and it is based on this that superior lifetime properties of the proposed design can be noticed.

Table 3.Quantitative Results Summary

Metric	Baseline	Thermal-Aware	Proposed
Average Operating Temperature (°C)	High	Moderate	Low
Peak Hotspot Temperature (°C)	Highest	Reduced	Minimum
Normalized MTTF	1.0	1.3–1.5	1.8–2.0
Lifetime Energy Consumption	High	Moderate	Lowest
Reliability Improvement	—	Moderate	Significant

The Figure 4 and Table 3 results are also used to prove that hotspot-based thermal control are not enough to achieve maximum system life. Rather, there is a need to have a common thermal-reliability-sensitive methodology, both to be able to reduce the effective temperature as well as to be able to obtain significant benefits in terms of long-term reliability, which stresses again the importance of lifetime-energy-sensitive VLSI design.

Discussion: Sustainability and Green computing Impact.

The findings made in this paper emphasise the fact that attaining the ideal of energy efficiency in contemporary VLSI systems will necessitate a transformation of the short-term, power-oriented optimization approach into a lifetime-based design approach. The traditional low-power methods tend to be interested in minimizing the instantaneous energy use, but the comparison shown in the Sections 7.1 and 7.2 indicates that these methods can create the total lifetime energy increment inadvertently when the thermal stress and reliability degradation are neglected. This disconnection has important consequences in terms of sustainable energy consumption and environmental effects as far as sustainability are concerned. With a combined approach to the thermal and reliability awareness, the proposed architecture focuses directly on the root causes of the accelerated ageing in the deeply scaled technologies. The realised decrease in operating temperature and a concomitant increase in mean time to failure increases the lifespan of the system in use, a fact that reduces the cumulative energy use over time. This increases system lifespan and ensures that there is less need to replace the equipment before it is due to wear and tear, which also means less electronic waste and less embodied energy (energy spent to manufacture, package and deploy new equipment).

The lifetime-energy-centric optimization model is also related closely with the new green computing principles. Instead of minimising power, regardless of the cost, the proposed setup will allow energy-proportional operation that can adapt to workload demands without reducing thermal and reliability margins. The balance is of special concern to always-on embedded and edge remainders where constant operation exacerbates the effects of thermal stress and aging over a lifetime. The architecture also maximises reductions in auxiliary energy overheads (e.g. active cooling and maintenance) by avoiding excess cooling demands and performance degradation due to reliability concerns. On the whole, it can be concluded that sustainability in the VLSI design should be measured in other metrics than the immediate power or energy-per-operation. Tightly coupled dimensions that are the determinants of the environmental footprint of computing systems include lifetime of energy consumption, reliability, and thermal stability. As proposed, the thermal reliability-conscious architecture shows that the increase of lifetime of the system is not merely a power benefit, but an effective mechanism that can be applied to decrease the overall energy consumption and even serve towards the goals of green computing in further embedded systems.

Guidelines on design to practitioners.

According to the analysis framework and experimental findings discussed herein, a few design principles can be captured in practical formations as far as architects and designers of energy-efficient VLSI systems are concerned. These suggestions set out to transform the suggested lifetime-energy-centric method to feasible measures in practical real-world embedded and edge computing situations. To begin with, system optimization needs to be conducted during the system lifetime and not at one operating point. Designers must not use instantaneous power / energy-per-operation only measures, especially in always-on systems. Developing lifetime energy consumption as a higher-level design requirement allows to make trade-offs between power savings, thermal stress and degradation of reliability more balanced.

Second, the near-threshold operation should be used selectively and also adaptively. Near-threshold voltage scaling can provide significant scaling of dynamic energy, and should be used only on workloads whose performance is not constrained too strictly, and where there is a wide enough margin to reliability. Tasks that are critical or lengthy are recommended to be run at marginally elevated voltages so that levels of excessive thermal pressure and rapid ageing are not encountered. Three, thermal awareness should be part of the task scheduling and power management. The dynamic voltage and frequency scaling, power gating and workload mapping should also be based on real-time temperature feedback to avoid hotspots. Active thermal control minimises leakage power as well as the exponential rate of ageing processes.

Fourth, the control loop should contain reliability monitoring and prediction. Lightweight aging monitors and lifetime prediction model enables system to predict degradation instead of responding to failures. Strict reliability constraints during optimization should be enforced so that the energy savings, as a result of optimization, do not come at the cost of decreased system lifetime. Lastly, to have sustainable VLSI design, its coordination is required at cross-layers. A circuit, micro architectural, or system level isolation is not enough to solve the obliviousness aspect of power, temperature and reliability. To attain robust, energy efficient and environmentally sustainable embedded systems, designers must embrace cross layer frameworks which are unified and hence they are able to optimise these factors. The combination of these guidelines offers an effective roadmap to the implementation of lifetime-energy-conscious VLSI architectures, and to the overall objective of green and sustainable computing in the future electronic systems.

Conclusion

This paper has introduced the lifetime-energy-centric approach to energy-saving VLSI design, to the underlying constraints involving the energy-saving VLSI design, and the conventional power-centric optimization of traditional power-in-heavy technologies. Due to the growing use of modern embedded and edge computing systems in always-on operation, thermal stress and reliability degradation have become recognisable forces that influence both the energy consumption long-term and system sustainability. To attain the realisation of energy efficiency optimum instantaneous power is not enough. To address this impact, a thermal -reliability-conscious architectural model was proposed with a supportive analytical lifetime energy mode, which explicitly provides the interaction between power consumption, temperature, and lifetime degradation with ageing. By implementing cross-layer coordination between voltage and frequency scaling, scheduling tasks, thermal control, and predicting system behavior, the direction proposed will make it possible to operate efficiently near the threshold with acceptable lifetime of the system. The experimental data also shows that the suggested architecture leads to a continuous decrease in the total lifetime energy consumption and a very considerable increase in the meantime to failure, compared with the baseline architecture and thermal-only optimised architectures. The results confirm that a small scale of short term decline in energy choices arranged by the awareness of the reliability results can comprehend high level of long term energy savings and sustainability gains. On the whole, this work makes lifetime energy a first-class design metric and offers a necessary framework of sustainable, reliable, and energy-efficient VLSI architectures and gives unambiguous guidelines in the further research and practical implementation in green embedded systems.

References

1. Amiraski, M. (2024). *Cost effective machine learning for computer architecture design* (Doctoral dissertation, Tufts University).
2. Ansari, M., Safari, S., Yeganeh-Khaksar, A., Salehi, M., & Ejlali, A. (2018). Peak power management to meet thermal design power in fault-tolerant embedded systems. *IEEE Transactions on Parallel and Distributed Systems, 30*(1), 161–173. https://doi.org/10.1109/TPDS.2018.2867247
3. Chhabria, V. A. (2022). *The next wave of EDA: Exploring machine learning and open-source philosophies for physical design* (Doctoral dissertation, University of Minnesota).
4. Das, A., Palesi, M., Kim, J., & Pande, P. P. (2024). Chip and package-scale interconnects for general-purpose, domain-specific and quantum computing systems: Overview, challenges and opportunities. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*. https://doi.org/10.1109/JETCAS.2024.xxxxx
5. Ebrahimi, E. (2017). *Voltage stacking and timing speculation with an SRAM focus* (Doctoral dissertation, University of California, Santa Cruz).
6. Gupta, P., Agarwal, Y., Dolecek, L., Dutt, N., Gupta, R. K., Kumar, R., & Sylvester, D. (2012). Underdesigned and opportunistic computing in the presence of hardware variability. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 32*(1), 8–23. https://doi.org/10.1109/TCAD.2012.2217963
7. Heo, J., Jeong, K., Choi, J. Y., Kim, T., & Choi, K. (2021). Hardware performance monitoring methodology at near-threshold computing and advanced technology nodes: From design to postsilicon. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 41*(6), 1929–1942. https://doi.org/10.1109/TCAD.2021.3097812
8. Hosseinimotlagh, S., Ghahremannezhad, A., & Kim, H. (2020). On dynamic thermal conditions in mixed-criticality systems. In *Proceedings of the IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)* (pp. 336–349). IEEE. https://doi.org/10.1109/RTAS48715.2020.00041
9. Kim, Y. G., Kim, M., Kong, J., & Chung, S. W. (2020). An adaptive thermal management framework for heterogeneous multi-core processors. *IEEE Transactions on Computers, 69*(6), 894–906. https://doi.org/10.1109/TC.2020.2974817
10. Mathews, O. (2014). *Targeting fault propagation scope to enhance reliability in embedded systems* (Master's thesis, University of Houston–Clear Lake).

11. Yuan, J. S., Lin, J., Alasad, Q., & Taheri, S. (2017). Ultra-low-power design and hardware security using emerging technologies for Internet of Things. *Electronics, 6*(3), 67. https://doi.org/10.3390/electronics6030067

12. Zanotti, T., Pavan, P., & Puglisi, F. M. (2021). Multi-input logic-in-memory for ultra-low power non-von Neumann computing. *Micromachines, 12*(10), 1243. https://doi.org/10.3390/mi12101243