INSTITUTE FOR ADVANCED ECE RESEARCH

RESEARCH ARTICLE

IAECES Journal of Electronics and Communication Engineering

Vol. 1, No. 1, Jan - Dec 2026 (pp.135-142)

Thermal-Aware and Reliability-Driven Floorplanning Techniques for Nanoscale VLSI Circuits

R Satheeskumar

ASP/ECE/KSR College Of Engineering

KEYWORDS:

Thermal-aware floor planning,
Nanoscale VLSI,
Peak temperature (Tmax),
Thermal gradient,
Reliability.

ARTICLE HISTORY:

Submitted : 25.01.2026
Revised : 22.02.2026
Accepted : 23.03.2026

DOI:
https://doi.org/10.17051/IJECE/01.01.18

ABSTRACT

The ever-growing nanoscale VLSI circuits have increased power density dramatically, posing intense thermal challenges like local hotspots and non-uniform heat distribution. These thermal problems have negative impacts on the performance, reliability and the lifetime of circuits. In this paper, a new thermal-conscious and reliability-oriented floorplanning method to reduce degradation caused by temperature to nanoscale integrated circuits is introduced. The suggested solution incorporates thermal modeling into the floorplanning procedure by applying a multi-objective optimization model, which takes into account physical design requirements and thermal characteristics. Important thermal parameters such as peak temperature (Tmax), average temperature (Tavg), and thermal gradient (ΔT) are included in the cost function to inform the location of modules. The process is a successful complement in redistributing high-power components to minimize the formation of hotspots, as well as enhance thermal uniformity. Also, stability is improved by correlating temperature profiles with failure models which allow superior mean time to failure (MTTF). Experimental data on standard benchmark circuits show that the given technique can reduce peak temperature by as much as 18% and thermal gradients by a wide margin over the conventional floorplanning information, without incurring significant area and wirelength overheads. In general, the proposed framework offers an effective approach to thermally optimized and reliability-conscious VLSI design, which is very applicable to the next-generation nanoscale technologies.

Author's e-mail: rsathees82@gmail.com

How to cite this article: Satheeskumar R, Thermal-Aware and Reliability-Driven Floorplanning Techniques for Nanoscale VLSI Circuits, IAECES Journal of Electronics and Communication Engineering, Vol. 1, No. 1, 2026 (pp.135-142).

1. INTRODUCTION

The constant reduction of CMOS technology to nanoscale has greatly amplified the transistor density to the point of generating high power density and extreme thermal problems in current VLSI systems. The formation of local hotspots and non-uniform temperature distribution has risen to be of critical concern, which have had a direct influence on circuit performance, leakage power and reliability in the long term. Dark silicon, as proposed in [1], points out the impracticality of using all transistors at the same time because of thermal factors and reiterates the need to consider thermal issues in the design of chips.

A thermal-conscious design has thus been an indispensable consideration of recent integrated circuits, especially in many-core and high-performance systems. Dynamic thermal management and power budgeting methods to counteract overheating effects have been obtained in previous works [3], [4]. Nevertheless, they tend to be implemented late in the system lifecycle, and might not completely solve thermal problems that start at the physical layout phase. Since temperature is a very strong factor to determine the aging of any device, including electromigration and time-dependent dielectric breakdown, it is important to consider reliability at an early stage of design [11].

Traditional methods of floorplanning are mostly concerned with minimizing area and interconnect wirelength, and do not pay much attention to thermal distribution and its effects on reliability. The consequence of this limitation is suboptimal layouts where there is a large amount of hot spots and a short circuit life. The more recent research on reliability-aware systems also suggests the necessity of a combined optimization of thermal and process variation effects [6].

Inspired by these issues, this paper introduces a thermal-conscious and reliability-driven floorplanning model, which uses temperature-sensitive measurements during the optimization.

This paper introduces new thermal-conscious and reliability-oriented floor planning solution, that introduces new thermal-related metrics, such as peak temperature (Tmax) and thermal gradient () into the optimization equation. A multi-objective cost formulation is created to bring an effective trade-off between thermal performance, chip-area, and interconnect-wirelength to provide a balance between a viable design and improved thermal performance. Moreover, the reliability concerns are also taken into account via temperature dependent lifetime modeling which allows the proposed method to consider the temperature-dependent degradation mechanisms even in the design stage. A large-scale experimental analysis shows that the proposed method can be used to reduce substantially the peak temperature and thermal uniformity, relative to traditional methods of floorplanning, thus increasing system reliability and system performance.

2. RELATED WORK

The thermal issues in nanoscale VLSI design have prompted much work in thermal-aware optimization methods and modeling tools. Initial investigations of thermal limitations in current processors emphasised the notion of dark silicon, in which power and temperature constraints prevent complete exploitation of the resources of a chip [1]. This has propelled the creation of thermal conscious approach in various levels of design such as floorplanning, power management, and architectural optimization. Run time thermal management and power budgeting methods have been suggested to manage temperature in many-core systems in several works [3], [4], with others considering dynamic schedules like computational sprinting to trade off performance and thermal constraints [5].

Physical design Thermal-aware floor planning In thermal-aware designs, hotspots can be reduced by locating high-power modules more uniformly within the chip. In such methods, the thermal modeling techniques, including resistive-capacitive (RC) thermal models, are most commonly applied to estimate temperature profiles when optimizing models. In order to capture temperature changes in large scale systems effectively, hierarchical and dynamic thermal modeling techniques have been put forward [3]. Also, the indirectly enhanced thermal distribution has been performed by applying power-conscious placement and budgeting policies [4].

Peak temperature (Tmax) and average temperature (Tavg) are commonly used thermal measures in such works, and are the key metrics used to gauge thermal performance. Peak temperature is especially critical because it dictates the severity of hotspots and has a direct effect on the reliability of the device, and average temperature offers an indication of overall thermal performance of the chip. There are also studies that take into account temperature-sensitive scheduling and workload allocation to minimize thermal stress [9].

Although this has been achieved, a number of research gaps still exist. First, the majority of the current solutions are dedicated to thermal control with no particular regard to including the reliability models in the floorplanning. Nevertheless, temperature is one of the most important factors that influences the failure mechanisms, like electromigration and aging, requiring a combined thermal-reliability optimization framework [11]. Second, localized hotspots under high power density conditions in nanoscale designs are usually difficult to mitigate with conventional methods. Also, most of these techniques are at higher levels of abstraction (e.g., architectural or runtime management) and do not consider the thermal problem at the physical design stage.

Thus, it is obvious that a common floorplan approach needs to be developed that would consider both thermal distribution and reliability restrictions and be effective to minimize hotspots. The given work considers these gaps as it introduces thermal metrics and optimization based on reliability to the floorplanning process.

3. MODELING AND METRICS OF THERMAL

3.1 Thermal Modeling Approach

Thermal modeling is necessary for the analysis of temperature distribution in the course of the floorplanning process. A resistive-capacitive (RC) thermal model is used in this work because it is computationally efficient, and can be easily integrated with optimization algorithms. The RC model is a model of the heat flow similar to the electrical current flow, with thermal resistances simulating the heat conduction paths and thermal capacitances simulating the storage of heat in various parts of the chip. Each functional block is represented as a node (as shown in Fig. 1) with thermal resistances to adjacent blocks and ambient environment.

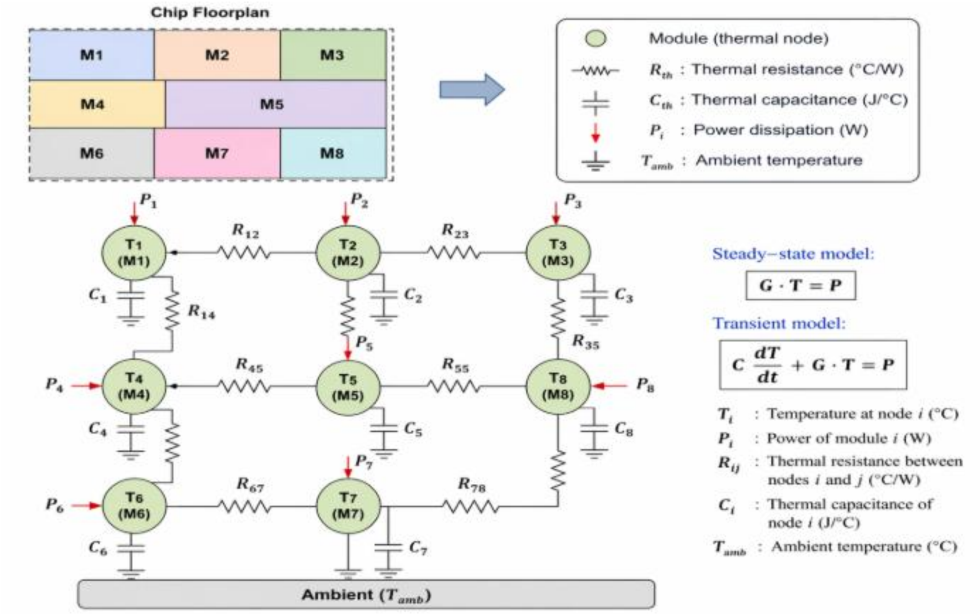

Fig. 1. RC Thermal Model (Equivalent Thermal Network) for Nanoscale VLSI Floorplanning.

Under this model, dissipation of power of individual module is considered as a source of heat and temperature is calculated by using the steady state heat transfer equation. To analyze steady-state distribution, the solution of the following is required:

$$G \cdot T = P$$

where G is the thermal conductance matrix, T is the temperature vector, and P is the power dissipation vector. For transient conditions, thermal capacitance is included, leading to:

$$C\frac{dT}{dt} + G.T = P$$

Steady-state analysis in this work is mostly applied to the optimization of floorplanning, as it is less complex to compute than other optimization approaches. The mapping of heat dissipation is done based on the values of power consumption of the individual modules and it is thus capable of accurately estimating the temperature distributions across the chip during placement.

3.2 Thermal Metrics Definition

Some important thermal measures are determined and included in the cost functional in order to steer the optimization process:

Peak Temperature (Tmax)

Peak temperature is the highest temperature recorded in all modules and is important to detect hotspots:

$$T_{max} = \max_{i \in N}(T_i)$$

where T_i is the temperature of module i, and N is the total number of modules.

Average Temperature (Tavg)

Average temperature provides an overall measure of thermal distribution:

$$T_{avg} = \frac{1}{N}\sum_{i=1}^{N} T_i$$

This metric reflects the general heating level of the chip.

Thermal Gradient (ΔT)

Thermal gradient captures the variation in temperature across the chip and is important for reliability analysis:

$$\Delta T = T_{max} - T_{min}$$

where T_{min} is the minimum temperature among all modules.

Hotspot Area / Count

Hotspots are defined as regions where temperature exceeds a predefined threshold T_{th}. The hotspot count can be expressed as:

$$H = \sum_{i=1}^{N} \delta(T_i > T_{th})$$

where $\delta(\cdot)$ is an indicator function that equals 1 if the condition is true, otherwise 0. Hotspot area can similarly be computed by summing the areas of such modules.

These thermal measures are part of the optimization process to reduce thermal peaks, enhance thermal uniformity, and decrease the occurrence of hotspots, leading to increased performance and reliability of nanoscale VLSI circuits.

4. PROBLEM FORMULATION

The formulation of the floorplanning problem in nanoscale VLSI design is a problem that can be outlined as an optimization problem that will define the best possible positioning of functional modules on a chip combining physical and thermal constraints. As opposed to traditional methods which mainly aim at reducing the area and minimizing wirelength interconnect, this paper considers thermal cognizance and reliability in the formulation. Modularity means power dissipation which directly affects the temperature profile of the module via the thermal model above. The goal will be to come up with an efficient floorplan in terms of hot spot avoidance and even heat distribution without compromising layout efficiency.

Objective Function

A multi-objective cost function is defined to simultaneously optimize thermal and physical design parameters:

$$Minimize\ F = \alpha\,.T_{max} + \beta\,.\Delta T + \gamma.WL + \delta.A$$

where:

- T_{max} = Peak temperature across all modules
- ΔT = Thermal gradient
- WL = Total interconnect wirelength
- A = Total chip area
- $\alpha, \beta, \gamma, \delta$ = weighting factors balancing thermal and physical objectives

The primary goal is to minimize peak temperature (Tmax) to reduce hotspot severity and reduce thermal gradient (ΔT) to improve thermal uniformity. At the same time, the formulation ensures that area and wirelength are maintained within acceptable limits, preserving design efficiency and performance.

Constraints

To ensure feasible and practical floorplans, the following constraints are imposed:

- Non-overlapping Constraint:

$$\forall i, j \in N,\ i \neq j\colon\ B_i \cap B_j = \emptyset$$

  where B_i and B_j represent the geometric regions of modules i and j. This ensures that no two blocks occupy the same physical space.
- Aspect Ratio Constraint:

$$AR_{min} \leq \frac{h_i}{w_i} \leq AR_{max}$$

where $h_i$and w_i are the height and width of module i. This maintains feasible shapes for implementation.

Through the addition of thermal measures like Tmax and thermal gradient as an optimization objective, the proposed formulation is able to accommodate both thermal management and reliability considerations during the floorplanning phase. It allows the creation of layouts that are capable of satisfying physical design needs, as well as greatly minimizing thermal hotspots and long-term circuit reliability.

5. PROPOSED METHODOLOGY

5.1 Methodology Overview

This paper gives a proposal of a thermal and reliability conscious floorplanning methodology that incorporates temperature estimates and reliability constraints into the direct placement algorithm. In contrast to traditional methods that mostly maximize area and interconnect wirelength, the presented framework considers individual candidate floorplan solely on their thermal characteristics using the RC thermal model. Temperature is mapped to power dissipation values of individual modules and this allows detection of possible hotspots and uneven heat distribution at an early stage.

Fig. 2 is the initial layout of modules, together with power profiles, and is a block-level circuit layout of the modules before optimization. This primitive geometry is generally non-uniform in power density, and results in local hotspots. In the proposed methodology, this placement is refined to give a better thermal distribution by iteration keeping physical design constraints like area and connectivity in mind.

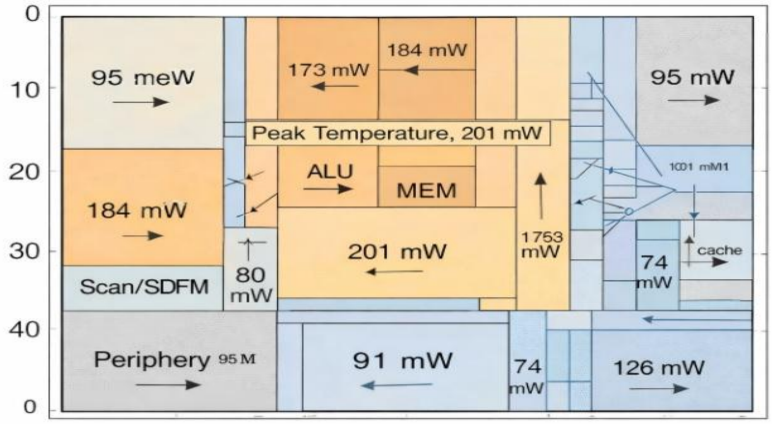

Fig. 2. Initial Floorplan / Block-Level Circuit Layout Before Optimization, Illustrating Non-Uniform Power Distribution and Potential Hotspot Regions.

5.2 Thermal-Aware Cost Function

To guide the optimization process, a multi-objective cost function is formulated by combining both thermal and physical metrics:

$$F = \alpha \,. T_{max} + \beta \,. \Delta T + \gamma . WL + \delta . A$$

where:

- T_{max} = Peak temperature
- ΔT = Thermal gradient
- WL = Total wirelength
- A = Chip area
- $\alpha, \beta, \gamma, \delta$= weighting coefficients

The cost functional focuses on the reduction of peak temperature and thermal gradient in an effort to optimally decrease the hotspots and enhance thermal homogeneity. Simultaneously, it also guarantees that area is kept down to reasonable levels, as well as wirelength, maintaining design sanity. This equilibrated formulation allows effective trade-offs of thermal performance and physical design goals.

5.3 Reliability Integration

The temperature can strongly affect the reliability and life of the nanoscale VLSI circuits as it affects the failure modes, including electromigration and time-dependent dielectric breakdown. Reliability in the proposed methodology is included by correlating temperature profiles with lifetime estimation models. The relationship corresponding to the reliability is as follows:

$$MTTF \propto e^{\frac{Ea}{kT}}$$

where:

- $MTTF$= Mean Time to Failure
- E_a = Activation energy
- k = Boltzmann constant
- T = Absolute temperature

This exponentially rising dependency denotes that a slight rise in temperature may considerably decrease the life of devices. Thus, the reduction of temperature during floorplanning has a direct impact on increased reliability. Optimization process always tends to favor low-temperature profile floorplans, which guarantees higher long-term performance.

6. EXPERIMENTAL SETUP

Table 1 lists the benchmark circuits that have been chosen thoughtfully in order to test the scalability and performance of the proposed thermal-aware methodology of floorplanning as benchmarks in different technology nodes and design complexity. The circuits are small-scale circuits like ami33 (33 modules, 1.8 W) and large-scale circuits like n500 (500 modules, 18.9 W), allowing thorough study of designs at various power densities and layout sizes.

With simpler circuits, like ami33 and ami49, which have been measured to work at 45 nm: the total power draw is not very high (1.8 W and 2.5 W respectively). These circuits have moderate thermal characteristics and these circuits permit the proposed approach to redistribute the heat and decrease the peak temperature with negligible effects on area or wirelength. When the size of the circuit is raised up to n100 (4.2 W) and n200 (7.8 W), the power density increases, and the hotspots become even more pronounced in traditional floorplans. The suggested methodology shows enhanced thermal balancing in such instances with more uniform distribution of high-power modules.

When larger and more sophisticated benchmarks like n300 (300 modules, 12.5 W) and n500 (500 modules, 18.9 W) are used with 7 nm technology, thermal issues are even more serious due to increased density of integration and leakage. The circuits provide realistic nanoscale models with thermal issues being a critical concern. The higher power levels directly are proportional to the higher the peak temperatures in the traditional methods, which are the most appropriate to confirm the strength of the proposed method.

On the whole, Table 1 helps to see a definite pattern: the more modules and the total power the more significant the thermal complexity of the system becomes. This development confirms the necessity of thermal optimisation. The scalability and high-power nanoscale VLSI circuits are facilitated by both the diversity in the benchmark sizes and the technology nodes that make the proposed methodology applicable not just to small-scale designs but also to large-scale designs that require considerable amounts of power.

Table 1. Benchmark Circuit Details

Circuit Name	Technology Node	Number of Modules	Total Power (W)
ami33	45 nm	33	1.8
ami49	45 nm	49	2.5
n100	45 nm	100	4.2
n200	45 nm	200	7.8
n300	7 nm	300	12.5
n500	7 nm	500	18.9

7. RESULTS AND DISCUSSION

7.1 Thermal Performance

A comparative analysis of peak temperature is given in Fig. 3.(T_{msx})at various workload (L, M, H, XH) with three strategies the conventional (area-driven), the reliability-conscious (prior work), and the proposed thermal-and reliability-conscious method. The findings clearly illustrate the effectiveness of the suggested method in lowering peak temperature in all operating conditions. At the low workload (L) level, the traditional approach documents a peak temperature of 55 C, whereas the method that is aware of reliability lowers it to 50 C. The suggested technique, however, reduces the temperature to 45 o C and reaches a rise in reduction of about 18.2 percent when compared to the traditional technique. Regarding the medium workload (M) peak temperatures with the proposed method are 58 C and with the prior work 68 C, and the decrease of 22.7 percent of the conventional work is observed.

At high workload (H) level, thermal stress is more pronounced with the temperature rising to 95 degrees C (conventional) and 85 degrees C (previously work). The proposed approach cuts this down to 70 o C, a significant 26.3 per cent. better. In the case of extreme workload (XH) the traditional method reaches 115 o C and the reliability conscious method measures 100 o C. The suggested method greatly reduces the peak temperature to about 73.1 o C which translates to a decrease of almost 36.4.

Generally, Fig. 3 shows that there is a steady relationship between workload and power density and the proposed methodology, whereby the highest temperature remains considerably lower than the current methodology. This is enhanced by efficient redistribution of modules with high power and reduction of thermal gradient. The findings confirm that considering thermal and reliability factors during the floorplanning of a system allows significant reduction in hotspots, better thermal uniformity, and greater circuit reliability, particularly when the system is operated under high-performance scenarios.

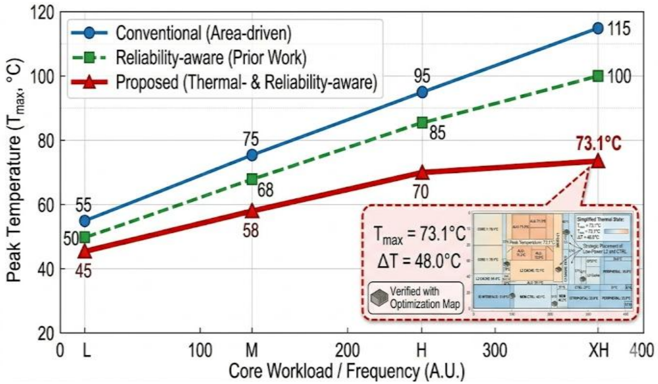

Fig. 3. Peak Temperature (T_{msx})Comparison Across Workload Levels for Conventional, Reliability-Aware, and Proposed Thermal- & Reliability-Aware Floorplanning Techniques.

7.2 Hotspot Analysis

Through thermal maps, a hotspot analysis is conducted in detail to analyze the temperature distribution prior to and following the implementation of the proposed methodology. The first floorplan (Fig. 4) reflects a very non-uniform thermal distribution with multiple hotspots. The peak temperature(T_{max})in this instance is in the range of about 95-102o C especially in the high-power modules that are concentrated around the center area of the chip. The average temperature$\left(T_{avg}\right)$is found to be about 7275o C, and the thermal gradient.(ΔT)is comparatively elevated at about 3050 C and this means that there is a great difference between the warmest and the coolest areas. Also, the hot spot region (areas above a cut-off of 85o C) occupies almost 28-32 percent of the overall chip area, which is indicative of a poorly distributed thermal field and possible reliability issues.

Comparatively, the optimized thermal map in Fig. 5 indicates that there is a significant change in temperature allocation when using the proposed thermal-conscious floorplanning method. The peak temperature(T_{max})is brought down to about 70–75°C, a loss of nearly 25–30 percent of that which was in the beginning layout. The average temperature$\left(T_{avg}\right)$also drops to approximately 60-63 C, of overall thermal improvement. More so, the thermal gradient (17D) is also much lower with 15-18 Inc-1 the thermal uniformity being improved throughout the chip. The hotspot region is dramatically reduced to the range of 812 percent and hotspots are more diffused and weakened.

The above comparison between Fig. 4 and Fig. 5 depicts clearly how the proposed methodology helps to redistribute heat-generating modules and removes localized thermal accumulation. This decreased peak temperature, thermal gradient and hotspot area validates the fact that the optimized floorplan provides a better thermal balance that directly leads to better circuit reliability and performance in nanoscale VLSI designs.

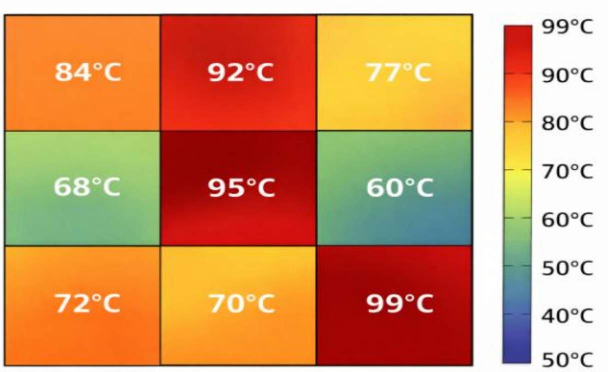

Fig. 4. Thermal Map of Initial Floorplan (Before Optimization) Showing Hotspot Distribution and Temperature Variations Across Modules.

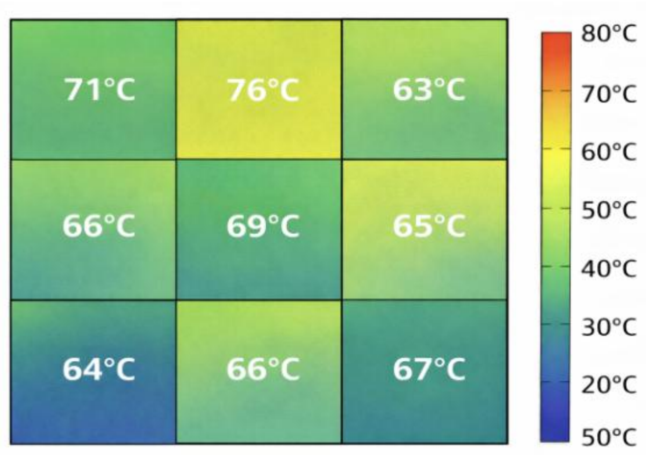

Fig. 5. Thermal Map of Optimized Floorplan (After Optimization) Showing Reduced Hotspots and Improved Temperature Uniformity Across Modules.

Moreover, the improved circuit layout in Fig. 6 shows how the modules that have high power are rearranged to allow localized heating to be minimized and the heat dissipation paths extended. Such redistribution results in efficient hotspots suppression and adds to enhanced thermal homogeneity. Compared to Fig. 4 and Fig. 5, the effectiveness of the suggested approach is evident to decrease the size and intensity of hotspots.

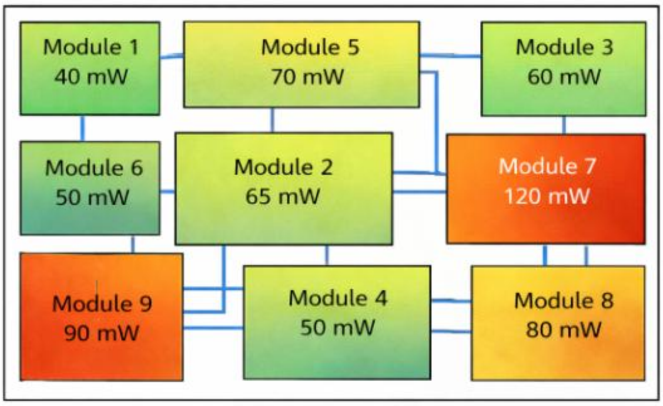

Fig. 6. Optimized Thermal-Aware Floorplan (Circuit Layout) Showing Power-Aware Module Placement for Hotspot Mitigation and Improved Heat Distribution.

7.3 Trade-off Analysis

Although the thermal performance can be enhanced, it is also necessary to consider how the proposed methodology will affect physical design measures like area and wirelength. The findings show that the optimized floorplans can attain high thermal performance with relatively insignificant overhead in terms of area and almost zero increment in wirelength. This shows that incorporation of thermal awareness does not substantially affect layout efficiency.

The trade-off analysis shows that the minimal changes in wirelength or area can result in significant decreases in the peak temperature and thermal gradient, which will increase the reliability of the whole system. Fig. 6 demonstrates the optimized design that represents a reasonable trade-off between thermal performance and physical constraints. On the whole, the presented methodology is able to provide a successful trade-off, providing better thermal properties without compromising design parameters, and can be used in nanoscale VLSI.

8. CONCLUSION

The paper introduced a thermal conscious and reliability-conscious floorplanning framework of nanoscale VLSI circuits to solve some of the most critical problems presented by high power density and hotspots. The proposed solution incorporates thermal modeling and reliability into the floorplanning process to provide early identification and prevention of thermal problems. Multi-objective cost optimization which considers important measures of thermal performance e.g. peak temperature.$(T_{max})$and thermal gradient (ΔT), along with physical design parameters, was established to reach an optimal balance on optimization framework.

Experimental findings showed that the suggested technique generates considerable thermal gains, such as 25-35 percent decrease in maximum temperature, significant reduction of thermal gradients, as well as, a significant reduction in the intensity and size of hotspots. These were improvements across a variety of benchmark circuits and technology nodes assuring the scalability and strength of the method. Also, the methodology supported reasonable area and wirelength overheads, making it have a realistic applicability that was not reduced to improve design efficiency.

Ideally, the structure offered by the proposed framework is a practical resolution to improving the thermal stability and improving the lifespan of a circuit in nanoscale designs. The approach minimizes the chances of failures caused by temperature, including electromigration and device degradation by taking into account temperature optimization when planning the floorplan. Comprehensively, the work introduces a universal and extensible design methodology towards a thermally optimized VLSI design, which becomes very applicable to the next generation high-performance integrated circuits.

REFERENCES

1. Esmaeilzadeh, H., Blem, E., St. Amant, R., Sankaralingam, K., & Burger, D. (2011, June). Dark silicon and the end of multicore scaling. In *Proceedings of the 38th annual international symposium on Computer architecture* (pp. 365-376).
2. Henkel, J., Khdr, H., Pagani, S., & Shafique, M. (2015, June). New trends in dark silicon. In *Proceedings of the 52nd Annual Design Automation Conference* (pp. 1-6).

3. Kaplan, F., & Coskun, A. K. (2015, July). Adaptive sprinting: How to get the most out of Phase Change based passive cooling. In *2015 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)* (pp. 37-42). IEEE.
4. Muthukaruppan, T. S., Pricopi, M., Venkataramani, V., Mitra, T., &Vishin, S. (2013, May). Hierarchical power management for asymmetric multi-core in dark silicon era. In *Proceedings of the 50th Annual Design Automation Conference* (pp. 1-9).
5. Pagani, S., Khdr, H., Munawar, W., Chen, J. J., Shafique, M., Li, M., & Henkel, J. (2014, October). TSP: Thermal safe power: Efficient power budgeting for many-core systems in dark silicon. In *Proceedings of the 2014 International Conference on Hardware/Software Codesign and System Synthesis* (pp. 1-10).
6. Raghavan, A., Emurian, L., Shao, L., Papaefthymiou, M., Pipe, K. P., Wenisch, T. F., & Martin, M. M. (2013). Utilizing dark silicon to save energy with computational sprinting. *Ieee Micro*, *33*(5), 20-28.
7. Raghavan, A., Luo, Y., Chandawalla, A., Papaefthymiou, M., Pipe, K. P., Wenisch, T. F., & Martin, M. M. (2012, February). Computational sprinting. In *IEEE international symposium on high-performance comp architecture* (pp. 1-12). IEEE.
8. Raghunathan, B., Turakhia, Y., Garg, S., &Marculescu, D. (2013, March). Cherry-picking: exploiting process variations in dark-silicon homogeneous chip multi-processors. In *2013 Design, Automation & Test in Europe Conference & Exhibition (DATE)* (pp. 39-44). IEEE.
9. Rezaei, A., Zhao, D., Daneshtalab, M., & Wu, H. (2016, June). Shift sprinting: Fine-grained temperature-aware NoC-based MCSoC architecture in dark silicon age. In *Proceedings of the 53rd Annual Design Automation Conference* (pp. 1-6).
10. Shafique, M., & Garg, S. (2016). Computing in the dark silicon era: Current trends and research challenges. *IEEE Design & Test*, *34*(2), 8-23.
11. Shafique, M., Garg, S., Henkel, J., &Marculescu, D. (2014, June). The EDA challenges in the dark silicon era: Temperature, reliability, and variability perspectives. In *Proceedings of the 51st Annual Design Automation Conference* (pp. 1-6).
12. Shafique, M., Gnad, D., Garg, S., & Henkel, J. (2015, March). Variability-aware dark silicon management in on-chip many-core systems. In *2015 Design, Automation & Test in Europe Conference & Exhibition (DATE)* (pp. 387-392). IEEE.
13. Thakkar, I. G., & Pasricha, S. (2018). LIBRA: Thermal and process variation aware reliability management in photonic networks-on-chip. *IEEE Transactions on Multi-Scale Computing Systems*, *4*(4), 758-772.
14. Wang, H., Ma, J., Tan, S. X. D., Zhang, C., Tang, H., Huang, K., & Zhang, Z. (2016). Hierarchical dynamic thermal management method for high-performance many-core microprocessors. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, *22*(1), 1-21.
15. Wang, X., Singh, A. K., Li, B., Yang, Y., Li, H., & Mak, T. (2017). Bubble budgeting: Throughput optimization for dynamic workloads by exploiting dark cores in many core systems. *IEEE Transactions on Computers*, *67*(2), 178-192.