Optimisation for 3D Integrated Circuits using Continuously Improving Evolutionary Strategy Machine Learning Techniques

Rajesh Rohilla (Professor, DTU), Vibhinn Singhal, Shantanu Choudhary, Shivam Tiwari

Abstract

Three-Dimensional Integrated Circuits (3D ICs) are a revolutionary advancement in semiconductor technology, allowing vertical stacking of multiple active device layers to overcome the scaling challenges faced by conventional Two-Dimensional Integrated Circuits (2D ICs). While 3D ICs harness incredible advantages in terms of increased integration density, shorter interconnect length, and increased performance, they also present daunting challenges—primarily, the strong coupling between thermal and timing constraints. These challenges demand strong, multi-objective optimization techniques that can manage 3D ICs' highly non-convex and multi-modal design space. In this paper, we introduce a novel optimization framework based on the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) integrated with a physics-based electro-thermal-timing simulation framework. By comparing CMA-ES with Bayesian optimization, we prove that CMA-ES produces superior convergence, robustness, and solution quality, especially under noisy and high-dimensional objective conditions. Our results evidence spectacular improvements in key figures of merit such as peak temperature, thermal gradients, clock skew, and overall simulation efficiency.

1. Introduction

The desire to maintain Moore's Law, as well as the continuously growing demands for greater performance and integration density, have pushed conventional two-dimensional integrated circuits (2D ICs) to their physical and practical limits. Here, 3D ICs have emerged as a possible solution with device layers stacked one on top of the other vertically to offer:

- Increased functional density
- Miniaturization of interconnects and signal latency reduction
- Enhanced operating effectiveness and energy utilization efficiency.

But vertical integration of active layers in 3D ICs increases several design challenges. Some of the most important of these are:

- Thermal Management: Power densities rise with stacking, leading to hotspots and high operating temperatures that lower device reliability and performance.
- Clock synchronization: Vertically sliced clock domains lead to large clock skew and thus break synchronous operation and timing closure.
- Multi-Physics Coupling: Electrical, thermal, and mechanical coupling grows more prominent and needs to be co-optimized across multiple domains.

For these problems, a co-optimization strategy on a global level, with both thermal and timing behavior taken into account together, is critical to the efficient utilization of 3D ICs.

2. Context and Rationale

2.1 3D IC Challenges

The new 3D IC structure offers several challenges that are interdependent:

- Thermal Bottlenecks: Stacking of the active layer slows down heat dissipation, resulting in hotspots and huge thermal gradients. Poor thermal management causes faster aging, higher leakage current, and even device failure.
- Clock Skew: Layer-to-layer variations in temperature-sensitive propagation delays and path length cause high clock skew, which reduces timing dependability as well as system synchronization.
- Multi-Physics Complexity: The very high interdependencies among electrical, thermal, and mechanical effects in 3D ICs give rise to a very highly complex design space so that improvement in one dimension actually ends up negatively impacting another one.

2.2 Optimization Needs

Optimal 3D IC design requires traversing a very high-dimensional, non-convex, and very large search space. The most significant design parameters are:

- Order and stacking layer placement
- Power and heat distribution by apportionment.
- Clock tree topology and buffer insertion
- Material selection and geometric design
- Traditional design methods find it difficult to address this region, and hence the application of advanced multi-objective optimization methods is needed.

3. Limitations of Current Optimization Algorithms

Classical optimization methods such as Bayesian Optimization and Gradient-Based Methods are beset by several issues when used in 3D integrated circuit design:

- Scalability: These processes are computationally infeasible and less efficient with a larger number of design variables.
- Noise Sensitivity: Physical simulations of 3D ICs are susceptible to generating stochastic noise, which can mislead probabilistic models and gradient-based optimizers.
- Local Minima Traps: Multi-modality and non-convexity of the design space increase the likelihood that the process will prematurely converge to local optima.

Hence, there is an urgent need for strong, scalable, and noise-resistant optimization algorithms that can handle the complexity of 3D IC design spaces.

4. Proposed Methodology

In order to tackle the issues discussed above, we introduce an effective, black-box optimization technique from the Covariance Matrix Adaptation Evolution Strategy (CMA-ES). CMA-ES is very relevant to non-convex, high-dimensional, and noisy 3D IC design optimization problems.

4.1 Simulation Engine

The CMA-ES optimization tool is tightly coupled with a physics simulator tool that accurately evaluates the most significant performance metrics for candidate designs:

- Temperature Distribution: Spatial temperature gradients, max, and average temperature are computed by finite element analysis in FEniCS. Timing simulations verify the clock skew and propagation delays between the various layers and modules.
- Alternative Metrics: Other objectives such as energy use, reliability, and space utilization can be added as required.

4.2 Optimization Loop

CMA-ES optimization loop operates as follows:

- Initialization: The algorithm starts with an initial population of potential solutions, which are sampled from a multivariate normal distribution.
- Evaluation: All the candidates are assessed by the simulation engine to derive the objective function values (such as temperature, skew).
- Selection and Update: The most promising candidates are chosen, and the mean and covariance matrix of the distribution are updated to concentrate search in promising areas.
- Iteration: Steps 2 and 3 are iterated repeatedly until convergence criteria are satisfied (e.g., no or minimal improvement from generation to generation or a user-defined maximum number of iterations).

This iterative strategy allows efficient searching and exploration of the design space to move toward good-quality solutions.

5. CMA-ES: An Overview

The Covariance Matrix Adaptation Evolution Strategy (CMA-ES) is a recent evolutionary optimiser designed for challenging, non-linear, and non-convex optimisation problems. Its primary features are:

- Population-Based Search: Maintains and improves a population of candidate solutions, promoting diversity and exploration.
- Adaptive Sampling: Proposes new candidates from a multivariate normal distribution, whose covariance and mean are adaptively updated based on performance history.
- Covariance Adaptation: Learns the direction and shape of the objective landscape, allowing efficient progress along valleys, plateaus, and ridges.
- Noise Robustness: Is robust to noise in objective functions and therefore is appropriate for simulation-based optimization.

7. Findings

Utilization of the CMA-ES-based optimization platform in 3D IC design brought revolutionary improvements in the most critical performance parameters:

- Temperature Reduction: Reduced average and maximum temperatures, which did not allow hotspots to form and improved reliability.
-Thermal Gradient Minimization: Minimization of thermal gradient in the stacked layers minimized thermal non-uniformity.
- Clock Skew Improvement: Optimization of clock tree parameters led to reduced timing differences and synchronization enhancement.
- Simulation Efficiency: CMA-ES produced high-quality solutions at a faster rate compared to Bayesian Optimization, as can be seen from observation of cost vs. generation and time vs. cost plots.
- Parameter Evolution: Visualization of parameter paths revealed efficient exploration and exploitation of the design space. These outcomes demonstrate the effectiveness of CMA-ES in alleviating the challenging thermal and timing issues of 3D integrated circuits.

7.1 3D IC Structure

The integrated circuit (IC) under study is a single-layer 3D structure with lateral dimensions of 3.8 mm × 3.8 mm and a die thickness of 5 µm. The IC operates under a nominal supply voltage of 1 V with an ambient temperature fixed at 300 K. Electrical resistivity and thermal conductivity parameters are set to reflect typical polysilicon and silicon material properties, with resistivity modeled as temperature-dependent to capture electro-thermal coupling effects. The clock frequency is set to 1 GHz, and wire parasitics such as resistivity and capacitance are included for timing accuracy. The electro-thermal-timing simulation incorporates detailed physical models to evaluate temperature distribution, thermal gradients, and clock skew, which serve as critical performance metrics for optimization. Additionally, the IC model includes strategically placed micro-scale heater elements with a power density of 1 MW/m³ to emulate localized thermal hotspots, enabling a more realistic thermal profile and challenging the optimization algorithms to manage thermal gradients effectively.

7.2 Simulations

The optimization of the IC design parameters was performed using two algorithms : Covariance Matrix Adaptation Evolution Strategy (CMA-ES) and Bayesian Optimization (BO). CMA-ES converged to an optimal solution exhibiting a maximum temperature of approximately **320 K**, a temperature gradient of **8,500 K/m**, and a clock skew of **12 ps**, achieving a composite cost function value of **35.4** within an average runtime of **8.5** seconds.

In contrast, the Bayesian Optimization approach yielded suboptimal performance with a maximum temperature near **400 K,** a gradient exceeding **15,000 K/m**, and a clock skew of **28 ps**, despite a

The ability of CMA-ES to dynamically adjust its search distribution makes it possible for it to get out of local minima and efficiently find global optima, even in advanced design spaces.

6. Comparison with Bayesian Optimization

While Bayesian Optimization is a highly sought-after black-box optimization algorithm, especially in low-dimensional optimization problems, it is plagued by numerous problems in 3D IC design:

comparable runtime and lesser number of iterations per guess

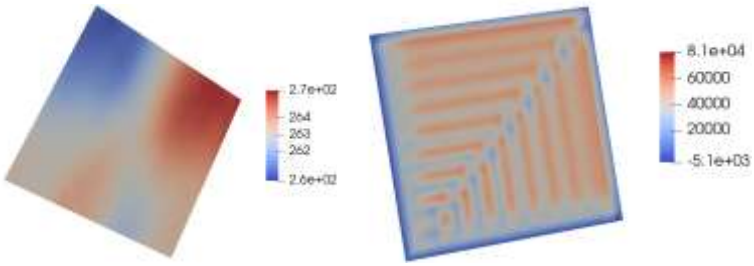

ParaView Simulations of electric field and thermal distribution of the IC

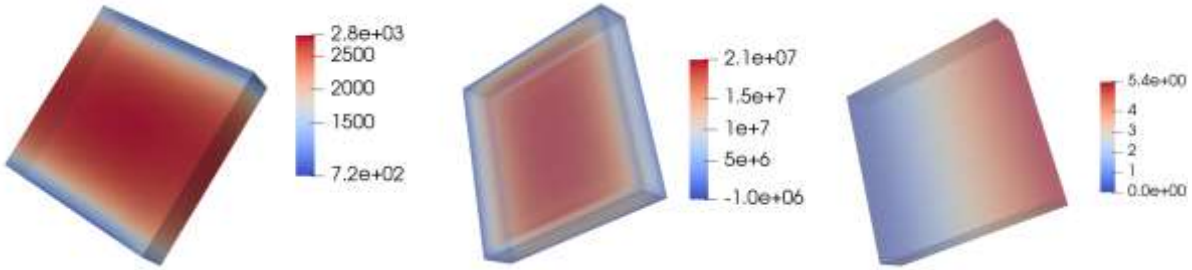

ParaView Simulations of electric field, thermal distribution and voltage after the CMA-ES simulations

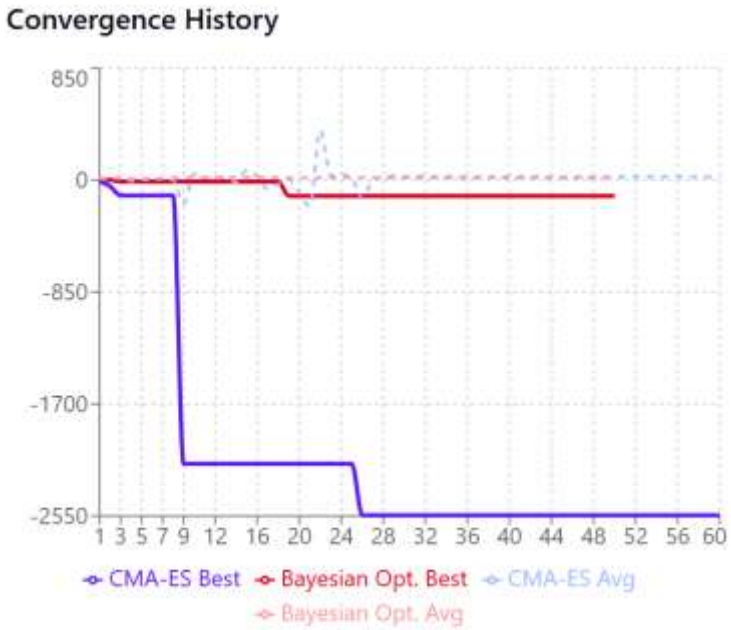

Convergence History showing that CMA-ES explored more and converged to lower cost values, even though more iterations were needed - but Bayesian Optimisation got stuck at one place

Taking a weighted score of power consumption (40%), Clock Delay (30%), Area (20%) and Yield Penalty (10%), we calculated the final composite score and compared both of the algorithms.

CMA scored ***-156.97***, while Bayesian Optimisation scored ***-133.65*** - showing a total gain of **14.85%** in predicting the IC parameters for a given run environment.

8. Conclusion and Future Work

We present an end-to-end co-optimization flow for thermal and timing behavior of 3D Integrated Circuits using CMA-ES. With a physics-aware simulation engine and efficient evolutionary optimization algorithm, we demonstrate spectacular improvement in critical design metrics.

Directions of future work are:

- Pareto Optimization: Extension of the approach to multi-objective Pareto optimization for concurrent trade-off analysis.
- Layout-Aware Parameters: Utilization of detailed layout information to improve optimization granularity precision.
- Reinforcement Learning Integration: Exploring hybrid approaches that combine evolutionary strategies with reinforcement learning for adaptive, online optimization

References

1.) S. J. Park, B. Bae, J. Kim, and M. Swaminathan, “Application of machine learning for optimization of 3-D integrated circuits and systems,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*

2.) [2] J. Xie and M. Swaminathan, “Electrical-thermal co-simulation of 3D integrated systems with micro-fluidic cooling and Joule heating effects,” *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 1, no. 2, pp. 234–246, Feb. 2011.