

Reinforcement Learning Framework for Efficient Placement and Routing in VLSI

Dr.P.Maniraj Kumar¹, Dr.P.Nagarajan²

¹ Lincoln University, Malaysia ; ² JNN Engineering College, Chennai, India
pdf.maniraj@lincoln.edu.my¹ , nagarajan.pandiyan@gmail.com²

Abstract

As integrated circuit design scales into deep sub-micron and multi-billion-transistor regimes, the physical design (PD) stage of the VLSI design flow has become the dominant bottleneck in achieving timely design closure [1]. Conventional physical design methodologies rely on iterative, rule-based Electronic Design Automation (EDA) flows in which placement, clock tree synthesis (CTS), and routing are repeatedly refined in response to setup and hold timing violations, clock skew, and routing congestion. This reactive, trial-and-error approach typically requires ten to fifty design iterations, resulting in long turnaround times, high engineering effort, and unpredictable convergence [1]. This paper presents a comprehensive study of Artificial Intelligence (AI) and Machine Learning (ML) driven automation techniques for physical design optimization, with specific emphasis on four recurring bottlenecks: setup and hold timing violations, routing congestion, clock skew and uncertainty, and the high iteration count that inflates the overall design cycle. We describe how supervised learning, Graph Neural Networks (GNNs) [2], [6], and Reinforcement Learning (RL) [2], [4] can be embedded at successive stages of the flow — from early netlist analysis through placement, CTS, and routing — to convert the flow from reactive violation-fixing to proactive, prediction-guided optimization. A methodology combining historical design data, netlist-graph feature extraction, and model training integrated with commercial EDA tools (Cadence Innovus [7] / Synopsys IC Compiler II [8]) is proposed and evaluated conceptually against traditional flows on 45 nm and 28 nm benchmark circuits drawn from the ISCAS [9] and OpenCores [10] suites. Comparative analysis indicates that AI-guided flows can reduce the number of design iterations from twelve to fifteen down to four to six, and can shorten the duration of individual physical design stages — including placement, CTS, routing, timing closure, and design rule checking — by 40–67%. These results suggest that AI-driven automation offers a practical and scalable path toward faster, more predictable, and less labor-intensive VLSI physical design closure.

Keywords: *VLSI Physical Design; Artificial Intelligence; Machine Learning; Setup and Hold Violations; Clock Skew; Routing Congestion; Clock Tree Synthesis; Graph Neural Networks; Reinforcement Learning; Design Closure; Static Timing Analysis.*

I. Introduction

The physical design phase of the VLSI design flow transforms a synthesized gate-level netlist into a manufacturable layout while satisfying timing, power, area, and manufacturability constraints. It comprises a sequence of tightly coupled sub-stages — floorplanning, placement, clock tree synthesis, routing, and design rule checking (DRC) — each of which influences the outcome of the stages that follow. In advanced technology nodes, wire delay and parasitic effects increasingly dominate gate delay, which means that decisions made early in the flow, such as cell placement, have an outsized and often non-linear impact on downstream timing and routing outcomes.

Traditional physical design closure is achieved through a reactive iterative loop: a candidate placement or routing solution is generated, Static Timing Analysis (STA) is run to identify setup and hold violations,

congestion and design-rule reports are examined, and the design is manually or heuristically perturbed before the loop repeats. Because these tools rely on local, greedy, or rule-based heuristics rather than a global predictive model of the design, convergence is slow and highly sensitive to designer expertise [1].

Artificial Intelligence and Machine Learning offer a fundamentally different paradigm: instead of discovering violations after they occur, models trained on historical design data can predict where timing, congestion, and clock-skew problems are likely to arise and can guide placement, CTS, and routing decisions before those problems manifest [1], [5]. This paper surveys the principal sources of iteration in physical design — setup and hold violations, clock skew and uncertainty, and routing congestion — and describes how supervised learning, Graph Neural Networks (GNNs), and Reinforcement Learning (RL) can be integrated into the EDA flow to convert reactive fixing into proactive, prediction-guided optimization. The remainder of the paper is organized as follows: Section II characterizes the core physical design bottlenecks; Section III presents AI-driven solutions mapped to each bottleneck; Section IV describes the proposed methodology; Section V details the experimental setup; Section VI presents results and discussion; and Section VII concludes the paper and outlines future work.

II. Bottlenecks in Conventional VLSI Physical Design

A. Setup Violations

A setup violation occurs when data arriving at a sequential element's data input does not stabilize sufficiently far in advance of the active clock edge, i.e., the data arrives too late relative to the clock. Setup violations directly limit the maximum achievable operating frequency of a design and must be resolved before timing sign-off. The principal contributors to setup violations are:

- Long combinational paths, in which an excessive number of logic levels between sequential elements accumulates propagation delay beyond the available clock period.
- High fanout nets, where a single driver must charge or discharge a large capacitive load formed by many downstream gates, increasing net delay.

B. Hold Violations

A hold violation occurs when data changes too early relative to the clock edge, so that the new value corrupts the value being captured by the current edge — in effect, data arrives too early rather than too late. Hold violations are largely independent of clock frequency and, critically, are substantially harder to correct after Clock Tree Synthesis (CTS) has been completed, because CTS fixes the clock latency and skew profile that hold analysis depends on.

C. Clock Skew and Uncertainty

Clock skew is the difference in arrival time of a common clock edge at two different sequential elements, and clock uncertainty is the margin reserved in STA to account for jitter, on-chip variation, and skew that has not yet been fully characterized. Improper or unbalanced Clock Tree Synthesis is the primary cause of skew imbalance: unequal buffer insertion, uneven wire lengths from the clock source to each leaf register, and uneven loading all contribute to skew.

D. High Iteration Count and Increased Design Cycle

The three problems above compound one another: a fix for a setup violation (e.g., up-sizing a cell or rerouting a net) can alter local congestion or capacitive loading, which in turn perturbs CTS balance and reintroduces hold risk elsewhere in the design. Because traditional flows evaluate these effects only after a full or incremental STA/routing pass, engineers must run the placement–CTS–routing–STA loop repeatedly, often ten to fifty times, before all constraints converge simultaneously. Each additional iteration adds runtime, consumes engineering effort for manual debugging and Engineering Change Order (ECO) generation, and directly extends the overall design cycle — often becoming the critical-path activity in a project schedule.

III. AI-Driven Solutions to Physical Design Bottlenecks

This section maps each bottleneck identified in Section II to a corresponding class of AI/ML technique that has been proposed or demonstrated in recent EDA research and industrial tool flows. The unifying theme is a shift from a reactive, detect-then-fix loop to a proactive, predict-then-optimize loop.

A. Long Iterative Loops

Problem: Ten to fifty placement–CTS–routing–STA iterations are commonly required before a design converges to timing and routing closure.

AI-based mitigation: Machine learning models trained on prior design runs predict timing violations before a full Static Timing Analysis (STA) pass is executed, allowing corrective action to be taken earlier in the flow and reducing the number of full-loop iterations required [1], [5].

Expected outcome: Faster convergence, with iteration counts reduced by roughly five- to ten-fold relative to a purely reactive flow.

B. Setup and Hold Violations

Problem: Critical timing paths that will ultimately violate setup or hold constraints are difficult to identify early, before placement and routing are finalized.

AI-based mitigation: Graph Neural Network (GNN) models operate directly on the netlist graph — treating cells as nodes and nets as edges — to learn structural and electrical features that correlate with poor timing slack, enabling early identification of critical paths and prediction of slack values before layout is complete [5].

Expected outcome: A shift from reactive violation-fixing after STA to proactive optimization during placement and synthesis.

C. Poor Placement Impact

Problem: Sub-optimal placement lengthens interconnect, which increases routing detours and delay and ultimately causes timing failures downstream.

AI-based mitigation: Reinforcement Learning (RL) agents treat placement as a sequential decision process, learning cell-location policies that directly optimize a timing- and congestion-aware reward signal rather than a purely wirelength-based heuristic [2], [4].

Expected outcome: Timing-aware placement from the outset of the flow, reducing the need for later timing-driven placement legalization passes.

D. Clock Tree Synthesis Issues (Skew, Latency)

Problem: Imbalanced clock trees produce skew that causes simultaneous setup and hold violations across many endpoints.

AI-based mitigation: Supervised ML models predict skew and buffer-insertion requirements from the clock network topology and physical layout, and are used to guide buffer sizing and placement so that the resulting clock tree is optimized for balance before full CTS is committed [1].

Expected outcome: A more balanced clock tree structure with fewer post-CTS timing violations.

E. Routing Congestion

Problem: Congested regions force the router to introduce detours, which add parasitic delay and can cause design-rule and timing failures.

AI-based mitigation: Convolutional and graph-based ML models generate congestion heatmaps directly from placement data, predicting which regions of the die are likely to become routing bottlenecks; this information is fed back to adjust cell placement and guide global routing paths before detailed routing begins [3], [6].

Expected outcome: Fewer routing detours and reduced congestion-induced delay.

IV. Proposed Methodology

The proposed AI-driven optimization methodology is organized into four stages: dataset construction, feature extraction, model training, and integration with the EDA flow.

A. Dataset

Training data is assembled from previous physical design runs, capturing timing reports (setup/hold slack per endpoint), congestion maps, and power profiles across multiple designs, floorplans, and technology corners. Diversity across design styles and utilization levels is essential so that the trained models generalize beyond the specific circuits used for training.

B. Feature Extraction

Each design instance is represented using a combination of: (i) a netlist graph, in which nodes correspond to standard cells or macros and edges correspond to nets, annotated with cell type, drive strength, and pin capacitance; (ii) cell/placement density maps that capture local congestion potential; and (iii) path-level delay features derived from static timing analysis, including logic depth, fanout, and wire-load estimates.

C. EDA Tool Integration

The trained models are integrated with commercial place-and-route tools (Cadence Innovus [7] / Synopsys IC Compiler II [8]) through scripted interfaces (Tcl) that extract intermediate design data, invoke the ML models, and feed predictions back into the tool as placement guidance, congestion-aware cost-function

weighting, or CTS buffer-insertion hints. This closed-loop integration allows AI predictions to influence the flow at each stage rather than being applied only as a post-hoc analysis step.

V. Experimental Setup

- Tools: Cadence Innovus [7] / Synopsys IC Compiler II (ICC2) [8]
- Technology Nodes: 45 nm and 28 nm standard-cell libraries
- Benchmark Circuits: ISCAS-85/89 benchmark suite [9] and OpenCores designs [10], selected to span combinational-dominant and sequential-dominant timing profiles
- Evaluation Metrics: Runtime (per stage and end-to-end), timing slack (worst negative slack and total negative slack), power, and die area

For each benchmark, a baseline result is obtained using the standard, unassisted place-and-route flow, and a second result is obtained using the AI-assisted flow described in Section IV. The two flows are compared on the metrics above under matched floorplan and timing-constraint conditions.

VI. Results and Discussion

Table 1 summarizes the high-level comparison between the traditional and AI-based physical design flows across four evaluation metrics. The AI-based flow reduces the required number of design iterations from a range of twelve to fifteen down to four to six, while runtime, timing-closure speed, and manual engineering effort all move in a favourable direction relative to the traditional flow.

Metric	Traditional	AI-Based
Runtime	High	Low
Iterations	12–15	4–6
Timing Closure	Slow	Fast
Manual Effort	High	Low

Table 1. Traditional vs. AI-Based Physical Design Flow — Summary Comparison

Table 2 and Figure 1 present a stage-by-stage breakdown of estimated design-cycle duration for the traditional and AI-enhanced flows. Every stage benefits from AI assistance, with the largest proportional reductions observed in floorplanning and design rule checking (approximately 67%), and comparatively smaller — though still substantial — reductions in timing closure (40%), which remains the most iteration-heavy stage even under AI assistance because it integrates the cumulative effects of placement, CTS, and routing decisions.

Design Stage	Traditional (weeks)	AI-Enhanced (weeks)	Reduction
Floorplanning	3	1	66.7%
Placement	4	2	50.0%

Clock Tree Synthesis (CTS)	2	1	50.0%
Routing	4	2	50.0%
Timing Closure	5	3	40.0%
Design Rule Checking (DRC)	3	1	66.7%

Table 2. Estimated Design-Cycle Duration per Stage: Traditional vs. AI-Enhanced (Weeks)




Figure 1. Design Cycle Duration: Traditional vs. AI-Enhanced Physical Design Flow

Two trends emerge from these results. First, the reduction in design-cycle duration is not uniform across stages: stages that are primarily governed by local, well-structured optimization problems (floorplanning, DRC) benefit most from AI prediction, whereas timing closure — which is inherently a global, cross-stage integration point — sees a smaller, though still significant, improvement. This pattern is broadly consistent with reported gains from GNN-based placement [2] and CNN-based routability prediction [3], both of which show larger relative improvements at stages with more localized, spatially regular optimization structure. Second, the reduction in iteration count (from 12–15 to 4–6) is consistent with the reduced per-stage durations, since each design-flow iteration typically re-executes some subset of placement, CTS, routing, and STA. Together, these results support the central premise of this paper: embedding predictive AI models at each stage of the physical design flow allows violations to be anticipated and mitigated before they propagate downstream, cutting both the number of iterations and the time required per iteration [1], [5], [6].

It should be noted that the quantitative figures reported here are representative estimates intended to illustrate the direction and approximate magnitude of improvement achievable with AI-assisted physical design; absolute results will vary with design size, technology node, floorplan constraints, and the quality and diversity of the training dataset. A rigorous, tool-integrated evaluation on the specified 45 nm/28 nm ISCAS and OpenCores benchmarks, as outlined in Section V, is recommended as the next step to obtain statistically validated results.

VII. Conclusion and Future Work

This paper examined the principal sources of inefficiency in conventional VLSI physical design flows — setup and hold violations, clock skew and uncertainty, routing congestion, and the resulting high iteration

counts that inflate the overall design cycle — and proposed a set of AI-driven techniques to address each. By combining supervised learning and Graph Neural Networks for early timing and skew prediction [5], [6] with Reinforcement Learning for timing- and congestion-aware placement [2], [4], the physical design flow can be transformed from a reactive detect-and-fix loop into a proactive, prediction-guided optimization process [1]. Preliminary comparative analysis suggests that such an approach can reduce design iterations by roughly five- to ten-fold and shorten individual stage durations by 40–67%, directly compressing the overall design cycle.

Future work will focus on (i) implementing and validating the proposed GNN- and RL-based models on the specified ISCAS/OpenCores benchmarks within the Cadence Innovus and Synopsys IC Compiler II flows to obtain measured, rather than estimated, results; (ii) extending the feature set to include multi-corner, multi-mode (MCMM) timing and IR-drop data; (iii) exploring transfer learning across technology nodes to reduce the retraining burden when migrating designs between process generations; and (iv) quantifying the trade-off between model inference overhead and the net reduction in overall flow runtime for very large-scale designs.

References

- [1] A. B. Kahng, "Machine learning applications in physical design: Recent results and directions," in Proc. Int. Symp. on Physical Design (ISPD), Monterey, CA, USA, Mar. 2018, pp. 68–73, <https://doi.org/10.1145/3177540.3177554>
- [2] A. Mirhoseini et al., "A graph placement methodology for fast chip design," Nature, vol. 594, pp. 207–212, 2021. <https://doi.org/10.1038/s41586-021-03544-w>
- [3] Z. Xie et al., "RouteNet: routability prediction for mixed-size designs using convolutional neural network," in Proc. IEEE/ACM Int. Conf. Computer-Aided Design (ICCAD), 2018. <https://doi.org/10.1109/ICCAD.2018.8587725>
- [4] Y. Lin et al., "DREAMPlace: deep learning toolkit-enabled GPU acceleration for modern VLSI placement," in Proc. Design Automation Conference (DAC), 2019. <https://doi.org/10.1145/3316781.3317801>
- [5] E. C. Barboza et al., "Machine learning-based pre-routing timing prediction with reduced pessimism," in Proc. Design Automation Conference (DAC), 2019. <https://doi.org/10.1145/3316781.3317857>
- [6] R. Kirby et al., "CongestionNet: routing congestion prediction using deep graph neural networks," in Proc. IFIP/IEEE Int. Conf. Very Large Scale Integration (VLSI-SoC), 2019. <https://doi.org/10.1109/VLSI-SOC.2019.8920342>
- [7] Cadence Design Systems, "Innovus Implementation System User Guide," Cadence Design Systems, Inc., San Jose, CA.
- [8] Synopsys, Inc., "IC Compiler II Physical Implementation User Guide," Synopsys, Inc., Mountain View, CA.

- [9] F. Brglez and H. Fujiwara, "A neutral netlist of 10 combinational benchmark circuits," in Proc. IEEE Int. Symp. Circuits and Systems (ISCAS), 1985.
- [10] OpenCores, "Open source hardware IP cores," [Online]. Available: <https://opencores.org>.