



Floorplan-Informed Performance/Power/Area Predictions in ASIC HLS with GNNs

Motivation

High-level synthesis (HLS) enables rapid development of application-specific integrated circuits (ASICs) by translating behavioral descriptions into register-transfer-level implementations. However, the final circuit depends on how operations are mapped to available hardware components and how those components are placed and connected on the chip. Hence, achieving an efficient design requires exploring a large space of architectural choices, including loop unrolling, pipelining, resource sharing, and memory organization, each of which affects **Performance, Power, and Area (PPA)**. Evaluating these alternatives through logic synthesis and physical implementation is computationally expensive, while early HLS estimates may fail to capture downstream optimizations and technology-dependent effects.

Therefore, accurate, low-cost PPA prediction can help designers compare many alternatives quickly and reserve expensive implementation runs for the most promising designs. A novel prediction approach used in the recent state-of-the-art works employs graph neural networks (GNNs), exploiting their ability of learning relationships between a design's computational structure, synthesis choices, and implementation quality, enabling faster identification of Pareto-optimal architectures.

Prior works have exploited a hierarchical approach from source to post-place-and-route level [1], gaining deeper knowledge of inner-loop dataflow via latency profiling, but they are solely based on FPGA-based HLS. On the ASIC HLS path, even in the latest works [2], they are based on optimizations made on FPGA implementations, trying to leverage knowledge transfer techniques to enhance the prediction quality. This thesis investigates **how direct characterization of ASIC**

implementations can provide the predictor with information about the target technology and the physical effects that influence the final circuit.

Thesis Targets




Figure 1: Abstraction of the Thesis Flow

The target of this thesis, presented abstractly in Figure 1, is to employ GNN methodologies based on the hierarchical approach, tested to be effective in the FPGA field [1], providing a floorplan-informed QoR metrics predictor starting from a high-level description. The predictor will combine program structure with information collected from ASIC implementations and floorplan conditions, such as the available chip space and its shape. Across the flow, we will use the CIRCT framework [3], allowing us to present several abstractions of the workload evaluated, starting from high-level descriptions (e.g., kernel level, operational level, gate level, etc.).

The steps that will be followed in this thesis will include:

- **Offline microkernel profiling:** Based on microkernel benchmarks, we will measure critical metrics for several abstraction layers of the hardware implementation (e.g., bit growth at the gate level, delay/area/power information at post-synthesis level, etc.). For this reason, RTL description will be provided by the CIRCT framework, and the open-source framework of OpenROAD[4] will be used for post-synthesis and post-implementation analysis, respectively, providing the necessary metrics. The collected information will enrich the program graphs used to train the GNN predictor. Of course, this benchmarking is not parametrized only by the nature of the workloads, but also **by the technology corner** used to be synthesized (nm, nominal voltage, operating temperature, etc.)
- **Online prediction and validation:** For a new high-level described kernel, a structured graph will be augmented with the chosen target and floorplan conditions. The trained GNN will predict PPA without requiring an implementation of every alternative. Promising designs will then be implemented to evaluate whether the approach helps find good designs with less exploration time (orders of magnitude less time 😊).

In conclusion, through this work the student will gain practical experience in novel hardware compilation frameworks, open source ASIC implementation tools, and graph neural networks, exploring how machine learning can support real chip-design decisions.

RELATED MATERIAL:

[1] M. Gao, J. Zhao, Z. Lin, and M. Guo, "Hierarchical Source-to-Post-Route QoR Prediction in High-Level Synthesis with GNNs," Design, Automation & Test in Europe Conference (DATE), 2024.

[Paper](#) · [Code](#)

[2] S. Zheng, Z. Ding, C. Wang, B. Yu, M. Wong, Y. Sun, and J. Cong, "Enhancing ASIC HLS Design Optimization via Knowledge Transfer from FPGA HLS Models," IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 2026.

[Paper](#)

[3] CIRCT Project, "Circuit IR Compilers and Tools," official documentation.

[Project](#) · [HLS in CIRCT](#) · [Verilog Generation](#)

[4] The OpenROAD Project, “OpenROAD-flow-scripts: RTL-to-GDS Flow,” official documentation.

[Documentation](#) · [Tutorial](#) · [Code](#)

CONTACT INFORMATION:

- Georgios Alexandris, Ph.D. Student, NTUA: (galexandris@microlab.ntua.gr)
- Elena Vouvali, M.Eng., Collaborator, NTUA: (elenavouvali@gmail.com)
- Panagiotis Chaidos, Ph.D. Student, NTUA: (pchaidos@microlab.ntua.gr)
- Sotirios Xydis, Assistant Professor, NTUA: (sxydis@microlab.ntua.gr)