

Modeling FPGA-Based Systems via Few-Shot Learning

Gagandeep Singh, *TUe, Netherlands & ETH Zürich*

Dionysios Diamantopolous, *IBM Research Europe*

Juan Gómez-Luna, *ETH Zürich*

Sander Stuijk, *TUe, Netherlands*

Onur Mutlu, *ETH Zürich*

Henk Corporaal, *TUe, Netherlands*

Contact: g.singh@tue.nl

Machine-learning-based models have recently gained traction as a way to overcome the slow downstream implementation process of FPGAs by building models that provide fast and accurate performance predictions. However, these models suffer from two main limitations: (1) a model trained for a specific environment cannot predict for a new, unknown environment; (2) training requires large amounts of data (features extracted from FPGA synthesis and implementation reports), which is cost-inefficient because of the time-consuming FPGA design cycle. In various systems (e.g., cloud systems), where getting access to platforms is typically costly, error-prone, and sometimes infeasible, collecting enough data is even more difficult. Our research aims to answer the following question: for an FPGA-based system, can we leverage and transfer our ML-based performance models trained on a low-end local system to a new, unknown, high-end FPGA-based system, thereby avoiding the aforementioned two main limitations of traditional ML-based approaches? To this end, we propose a transfer-learning-based approach for FPGA-based systems that adapts an existing ML-based model to a new, unknown environment to provide fast and accurate performance and resource utilization predictions.

Keywords: Machine learning; FPGA; modeling; performance prediction

DOI: <https://doi.org/10.1145/3431920.3439460>