Diploma Thesis

Microproccessors and Digital Systems Laboratory

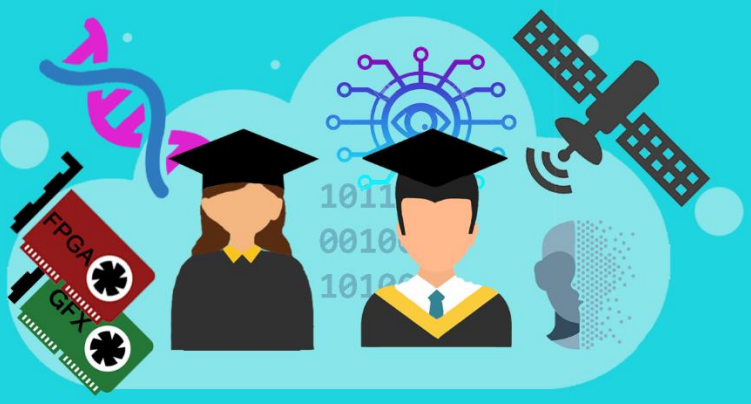

Machine learning based design space exploration for FPGA kernel acceleration

Designing hardware for FPGAs is a demanding and time-consuming process. Although algorithmic level methodologies, such as High-Level Synthesis (HLS), have provided abstraction layers compared to Register Transfer Level (RTL) (i.e., Verilog and VHDL), developers still need to familiarize with hardware concepts to efficiently accelerate computationally intensive parts of an application. Selecting manually the appropriate HLS directives is a challenging task even for human experts, mainly because of the large design decision space and its variation through different FPGAs. In this direction, Xilinx introduced Vitis, a framework that provides a unified OpenCL interface for programming edge (e.g., MPSoC ZCU102, ZCU104) and cloud (e.g., Alveo U50, U200) devices. Even though Vitis simplifies the designing process and lets developers focus on performance optimizations, it is not able to consider the architectural characteristics of the target FPGAs (available resources etc.). In this context, it is essential, for both academic and industrial purposes, to build tools that can automatically specify the appropriate HLS directives with respect to a specific architecture.

In this diploma thesis you will optimize, in terms of latency and consumed resources, Rodinia High-Level Synthesis benchmark suite using a mechanism that performs design space exploration (DSE). You will verify the inefficiency of exhaustive DSE and you will speed-up the process by using genetic algorithms and reinforcement learning techniques. This tool will be built upon Vitis HLS 2021.2, the state-of-the-art framework for designing kernels for Xilinx FGPAs.

PREREQUISITES:

Strong knowledge of **High Level Synthesis**, **C/C++**, **Python**

Desirable: Familiarity with **Machine Learning**, **Deep Learning**

RELATED MATERIAL:

- Cong, Jason, et al. "Understanding performance differences of FPGAs and GPUs." *2018 IEEE 26th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. IEEE, 2018.
- Cummins, Chris, et al. "End-to-end deep learning of optimization heuristics." *2017 26th International Conference on Parallel Architectures and Compilation Techniques (PACT)*. IEEE, 2017.
- Cummins, Chris, et al. "Programl: Graph-based deep learning for program optimization and analysis." *arXiv preprint arXiv:2003.10536* (2020).

- https://www.xilinx.com/support/documentation/sw_manuals/xilinx2020_2/ug1399-vitis-hls.pdf

CONTACT INFORMATION:

Aggelos Ferikoglou Ph.D, NTUA (aferikoglou@microlab.ntua.gr)

Dimosthenis Masouros Ph.D, NTUA (dmasouros@microlab.ntua.gr)

Professor Dimitrios Soudris, NTUA (dsoudris@microlab.ntua.gr)