Ph.D. in Information Technology
Thesis Defense

October 12th, 2026
at 9:30 am
Room PT1 – building 20A

Giovanni GOZZI – XXXVIII Cycle

Empowering high-level synthesis methodologies with an openmp runtime

Supervisor: Prof. Fabrizio Ferrandi

Abstract:

This thesis investigates the High-Level Synthesis (HLS) of parallel hardware accelerators with a particular focus on OpenMP. It adopts a compiler-first methodology that directly maps OpenMP runtime primitives into hardware components, leveraging a modern compilation flow based on Clang. This approach ensures robustness across applications and provides a systematic foundation for supporting future OpenMP extensions.
All generated accelerators share a unified architectural template composed of a sequential region and one or more parallel regions. The sequential region is responsible for initialization and serial operation, while the parallel regions carry out concurrent computations. Depending on how parallelism is expressed in the input code, two complementary synthesis flows are provided: one targeting loop-based parallel regions and another designed for task-based parallelism. The former instantiates multiple processing elements to exploit iteration-level concurrency, whereas the latter builds a dependency-aware distributed controller capable of dynamically scheduling and synchronizing independent tasks. Within this architecture, memory parallelism plays a central role: to overcome bandwidth and latency bottlenecks, the architecture supports multiple memory channels coupled with flexible arbitration mechanisms. Both centralized arbiters and Network-on-Chip (NoC) interconnects are supported, enabling a trade-off between area utilization and memory-request throughput depending on workload characteristics. Experimental evaluation demonstrates improvements in throughput across loop-based and task-based benchmarks, with moderate area and frequency overhead. Ultimately, this thesis presents a methodology that exploits compiler-level optimizations and at the same time hides hardware implementation details, and can be used to implement an HLS tool capable of efficiently translating complex parallel software into hardware.

PhD Committee

Prof. Gianluca Palermo, Politecnico di Milano

Prof. Jose Cano Reyes, University of Glasgow

Prof.ssa Stefania Perri, Università della Calabria