Ph.D. in Information Technology
Thesis Defense

April 27th, 2026
at 3:00 pm
Room 1B – building 20A

Davide GALLI – XXXVIII Cycle

Deep Learning Against Obfuscation: Boosting Side-Channel Analysis Across Segmentation, Attack, and Defense

Supervisor: Prof. Davide Zoni

Abstract:

Side-channel attacks exploit unintended information leakage emitted by cryptographic devices to extract sensitive data. Such attacks represent a critical threat to the security of modern computing platforms, especially in the context of pervasive IoT devices that continuously process sensitive data. This thesis addresses the challenges of evaluating and enhancing the security of digital circuits against side-channel attacks, focusing on realistic scenarios and reducing the costs and complexity of the analysis. To achieve this goal, it presents an open-source framework that uses deep learning to resolve critical problems at every stage of the side-channel workflow, from trace acquisition to the development of lightweight countermeasures.

Isolating and aligning cryptographic operations within side-channel traces is crucial for successful attacks. This task becomes particularly challenging in real-world scenarios where noise and countermeasures can obscure the signals of interest. A primary contribution of this thesis is Hound, a deep-learning pipeline for advancing cryptographic operations segmentation in traces highly affected by hiding countermeasures, eliminating the need for trigger infrastructure. The novel Chameleon dataset is introduced publicly to provide a realistic benchmark for evaluating segmentation and attack methodologies. Once segmented and isolated, side-channel traces can be analyzed to recover secret information. Classical attacks often struggle against traces protected by hiding techniques, which introduce significant temporal misalignment. This work explores attack strategies against dynamic frequency scaling countermeasures, introduces the DFS_DESYNCH dataset for benchmarking methodologies on highly desynchronized traces, and proposes Owl, a novel deep learning-assisted attack to tackle these challenges. Owl leverages deep learning to preprocess and realign traces, enabling effective key recovery even in the presence of significant temporal misalignment. On the defense side, the thesis investigates the impact of run-time variability, such as dynamic voltage and frequency scaling, on side-channel leakage and attack

resistance. It proposes Rabbit, a hardware module for fine-grained and random frequency changes on FPGAs, and demonstrates its effectiveness through extensive experimental campaigns. Results identify frequency scaling as the most robust hiding technique among those inspected, while Rabbit demonstrates high resistance to the tested advanced attacks without compromising system performance.

All the experimental campaigns leverage JARVIS, a novel open-source hardware-software framework for side-channel research on FPGA-based IoT-class systems. JARVIS provides a complete environment for configuring, monitoring, and analyzing cryptographic operations, enabling researchers to deploy and assess countermeasures with high observability and control. Overall, this thesis delivers a comprehensive workflow for side-channel analysis in realistic environments, combining open-source platforms, novel datasets, advanced segmentation and attack techniques, and innovative countermeasures. The contributions support the development of secure digital systems and provide valuable resources for the research community, advancing the state of the art in side-channel analysis

Gabriele MONTANARO – XXXVIII Cycle

Optimization of FPGA-Based Heterogeneous Multi-Core Systems-on-Chip for Data Centers

Supervisor: Prof. Davide Zoni

Abstract:
In recent years, the demand for computational power--especially from artificial intelligence applications--has grown at an unprecedented pace. To address this need, massive investments have been made in the construction of data centers and high-performance computing systems. However, traditional architectures such as CPUs and GPUs have shown clear limitations: their versatility comes at the cost of low efficiency in domain-specific contexts, both in terms of performance and energy consumption. This has led to the adoption of highly specialized hardware accelerators, capable of improving computational performance while simultaneously reducing energy costs.

Among the available solutions, FPGAs stand out thanks to their ability to implement architectures tailored to specific operations, combined with the flexibility of rapid reprogramming, which enables much shorter development cycles compared to ASICs. Nevertheless, their cost and non-negligible energy consumption make optimized usage essential. Furthermore, the ever-increasing size of modern FPGAs allows them to host multiple applications simultaneously, but also dramatically enlarges the design space of possible configurations. Exhaustively testing all configurations becomes infeasible, which has motivated the development of design space exploration approaches to accelerate the identification of optimal solutions.

In this context, the research presented in this thesis introduces a set of architectural and methodological techniques to facilitate design space exploration in systems with multiple hardware accelerators instantiated on a single FPGA. The work is structured into two main directions.

The first concerns the development of a prototyping platform for heterogeneous systems, enhanced with features designed to: (i) expand the design space by balancing energy consumption, resource utilization, and computational performance; (ii) collect execution data from the explored configurations; and (iii) accelerate the exploration process through the use of FPGAs' dynamic partial reconfiguration. Experimental results demonstrate that these techniques reduce the design space exploration time by up to 29×.

The second research direction focuses on optimization methodologies, pursued along two complementary levels: at the application level, by identifying which sections should run in software and which in hardware; and at the system level, by determining the best allocation of accelerators to maximize performance while minimizing area usage. In the latter case, the use of machine learning techniques makes the methodology architecture-agnostic while drastically reducing the number of configurations that need to be evaluated. Results show that the optimal configuration can be identified by exploring less than 0.1% of the overall design space.

PhD Committee

Prof. Michele Carminati, Politecnico di Milano

Prof. Stjepan Picek, Radboud University

Prof. Dimitrios Soudris, National Technical University of Athens