UNIVERSITÀ DEGLI STUDI DI NAPOLI FEDERICO II

itee PhD information technology electrical engineering

DIETI UNINA

Vincenzo Maisto

Innovative Computing Architectures for Green Computing

Tutor: Alessandro Cilardo

Cycle: XXXVII

Year: Second

itee PhD information technology electrical engineering

Ministero dell'Università e della Ricerca

PON RICERCA E INNOVAZIONE 2014 - 2020

My background

- MSc degree: Computer Engineering at UNINA
- Research group/laboratory: SECLab
- PhD start date: 1st January 2022
- Scholarship type: MUR PON
- Partner company under DM 1061: A3cube Inc.

Research field of interest

- **Green Computing Architectures**
- **Advanced Hardware/Software Co-design** on heterogeneous MPSoC classes:

Research field of interest

- **Green Computing Architectures**
- **Advanced Hardware/Software Co-design** on heterogeneous MPSoC classes:
  - Application-class SoC design

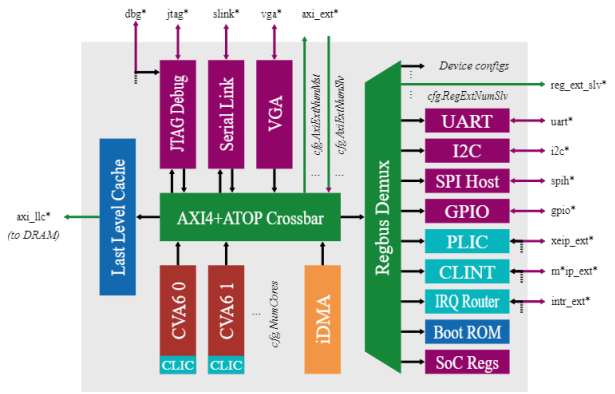

Research field of interest

- **Green Computing Architectures**
- **Advanced Hardware/Software Co-design** on heterogeneous MPSoC classes:
  - Application-class SoC design
  - Edge MPSoC platforms

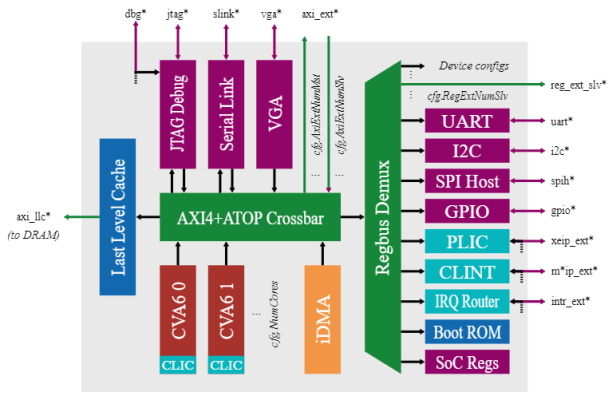

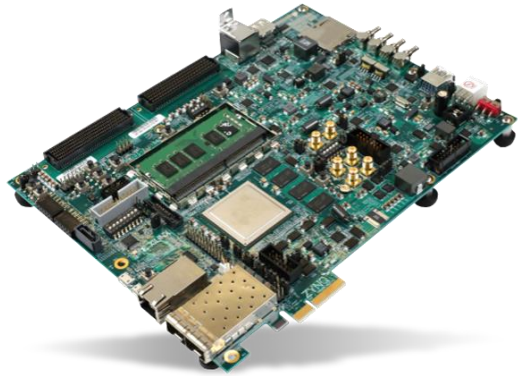

Research field of interest

- **Green Computing Architectures**
- **Advanced Hardware/Software Co-design** on heterogeneous MPSoC classes:
  - Application-class SoC design
  - Edge MPSoC platforms
  - Server-class Accelerator Cards

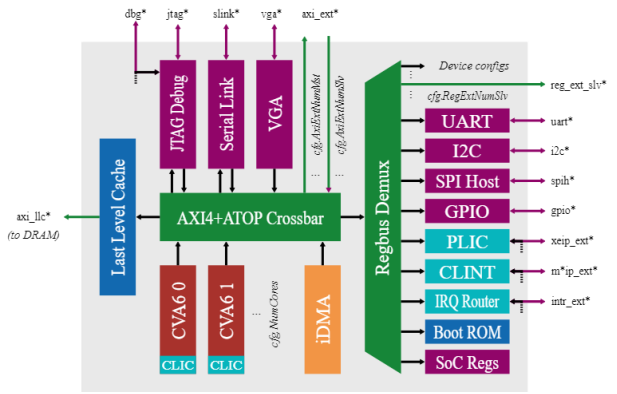

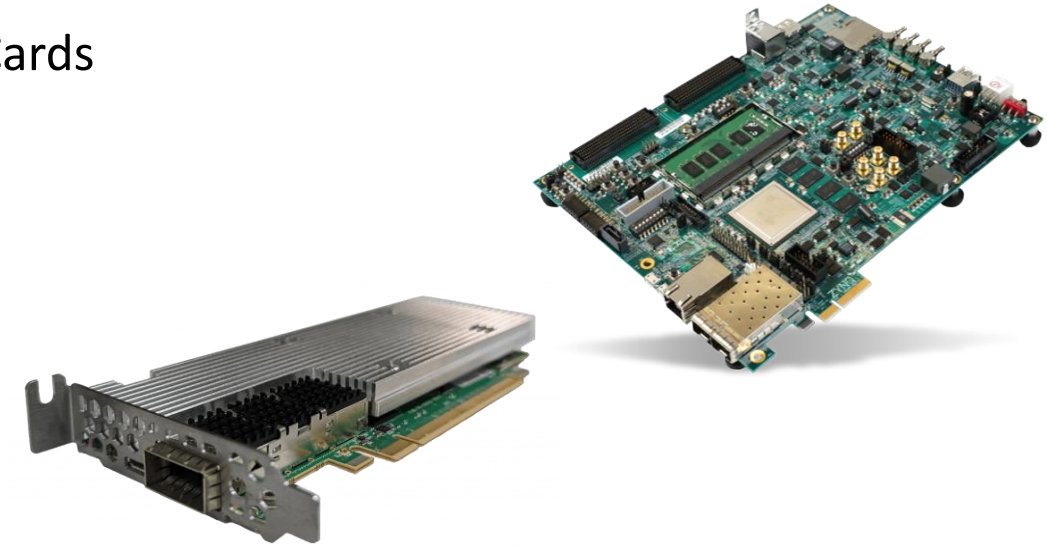

Research field of interest

- **Green Computing Architectures**
- **Advanced Hardware/Software Co-design** on heterogeneous MPSoC classes:
  - Application-class SoC design
  - Edge MPSoC platforms
  - Server-class Accelerator Cards
- Objective: **Heterogeneous, Scalable and Energy Efficient Architectures for HPC**

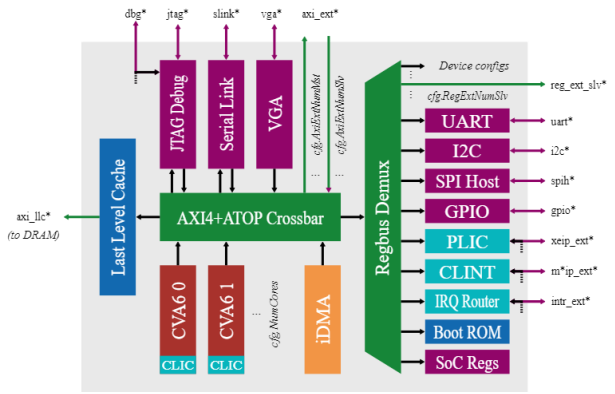

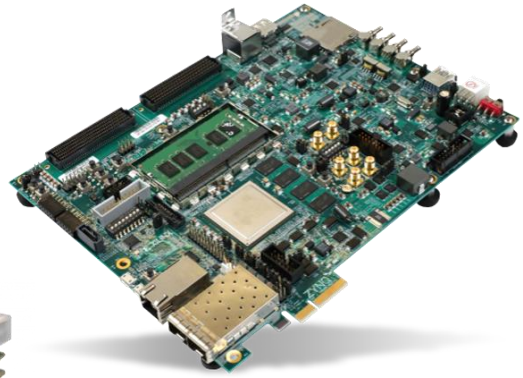

Summary of study activities

- Ad hoc PhD courses / schools:
  - Statistical data analysis for science and engineering research, ITEE, UNINA;
- Six months visiting period at ETH Zurich
  - Hosted by the PULP group at IIS (Integrated Systems Lab)

ETH zürich

- Conferences / events attended:
  - Participation to the “10 Years of PULP” workshop in Lugano

Summary of activities per credits/months

- Reported credits to ITEE:

Year	Tot Year	Courses	Seminars	Research
1	64.4	29	6.4	29
2	56.25	4	3.25	49
3	0	0	0	0
Total	**120.85**	**33**	**9.25**	**78**

- Reported months w.r.t. DM 1061:

Months	Reported	Expected
In department	13.5	**24**
In industry	4.5	**6**
Abroad	6	**6**
Total	**24**	**36**

Research Activity: Green Computing Architectures

Problem:

1. **Acceleration** of industrial and scientific computing workloads
2. **Technologically heterogeneous** computing architectures
3. **Energy efficiency and power consumption**
4. **Scalable** solutions for datacenter scales

Objectives:

1. **Scalability and energy efficiency**
2. **Heterogeneous HPC architectures**

Research Activity: Overview

Methodology:

1. **State-of-the-art** of the acceleration methodologies and cutting-edge technologies
   - AI co-processors
   - High Level Synthesis
   - Dynamic Partial Reconfiguration
   - Distributed file systems
   - Cluster computing
   - High Bandwidth Memories
   - Vector Extensions
   - On-chip interconnection
2. **Evaluation and analysis** of innovative hardware computing platforms and software stacks
   - Xilinx ZCU102, Intel Arria10, Intel Agilex, Xilinx Virtex, Xilinx Alveo …
   - Vitis-AI, Intel OPAE, OpenCL, HLS, Intel OFS, SYCL/DPC++, oneAPI, …
3. **Advanced hardware/software co-design** of innovative architectures

Software platform design  **Hardware platform design**

Research Methodology:
Advanced Hardware/Software Co-design

Software drivers and libraries

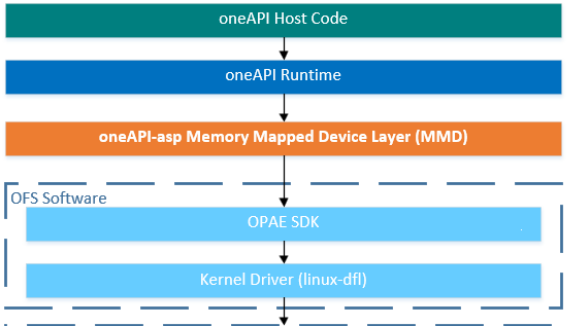

Hardware accelerator design

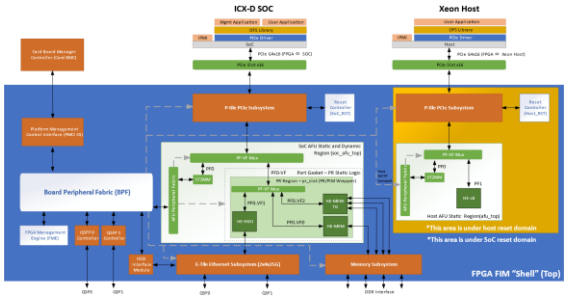

Research Methodology:
Advanced Hardware/Software Co-design

Software drivers and libraries

Hardware platform and interfaces

Hardware accelerator design

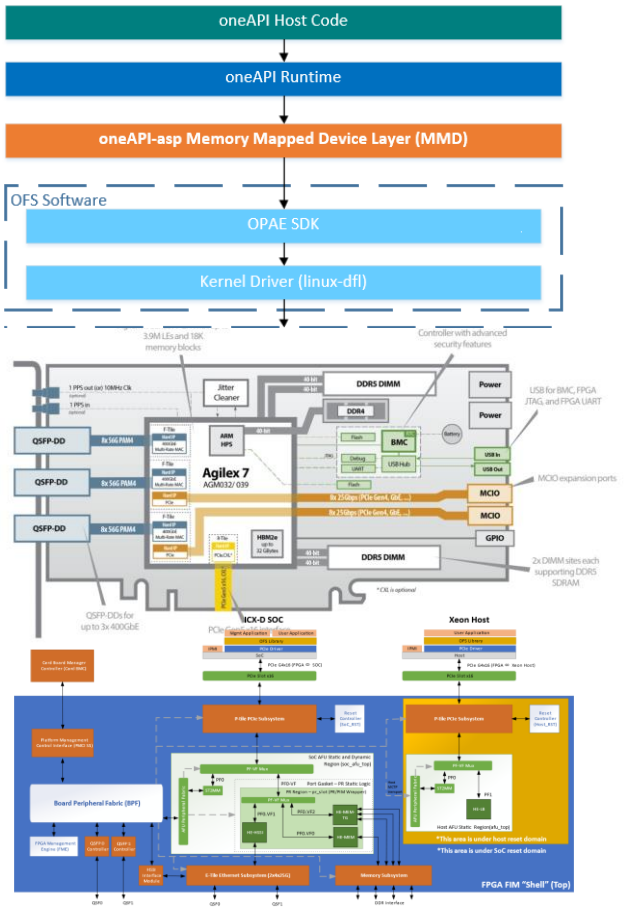

Research Methodology:
Advanced Hardware/Software Co-design

HPC middlewares

Software drivers and libraries

Hardware platform and interfaces

Hardware accelerator design

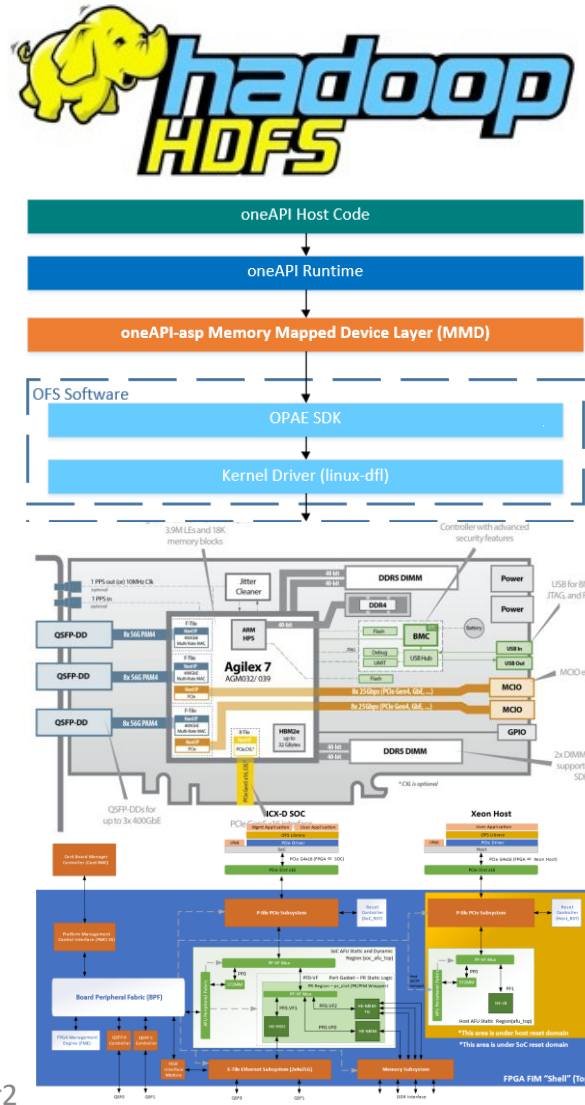

Products

[P1]	Cilardo, A., Maisto, V., Mazzocca, N., & Rocco di Torrepadula, F. (2023). An approach to the systematic characterization of multitask accelerated CNN inference in edge MPSoCs. *ACM Transactions on Embedded Computing Systems*. DOI: https://doi.org/10.1145/3611015
[P2]	Cilardo, A., Maisto, V., Mazzocca, N., Rocco di Torrepadula, F. Knowledge Distillation for EdgeAI: A Systematic Evaluation of the Energy Efficiency and Accuracy Trade-off. *[in preparation for IEEE Transactions on Sustainable Computing]*
[P3]	GitHub PRs for Ara: • [Draft] Refactoring hw source code • [Draft] Extend sw build flow for Linux environment • [Draft] Bug fixes and vstart CSR support • [Draft] Introduce virtual memory support in Ara GitHub PRs for CVA6, CVA6-SDK and Cheshire *[in preparation]*
[P4]	Maisto, V., Perotti, M., Cilardo, A., Benini, L. *Virtual Memory Support for RISC-V Vector Extension: A Quantitative Evaluation of Runtime Performance and Energy Efficiency.* *[in preparation]*
[P5]	Maisto, V., Cilardo, A., Billi, E. *Datacenter-scale Acceleration of Distributed Filesystems: Erasure Codes FPGA Offloading with SYCL High-Level Design and PCIe SR-IOV.* *[in preparation]*

Plans for Next Year

- *First year:*
  - *Transversal technological study of several MPSoC classes*
  - ***Edge-class** platforms*
- *Second year:*
  - *1st part: **Server-class** platforms*
  - *2nd part: Energy-efficient application-class **MPSoC design***

Plans for Next Year

- *First year:*
  - *Transversal technological study of several MPSoC classes*
  - ***Edge-class** platforms*
- *Second year:*
  - *1st part: **Server-class** platforms*
  - *2nd part: Energy-efficient application-class **MPSoC design***
- Third year:
  - Finalize the work on the three tracks

Plans for Next Year

- *First year:*
  - *Transversal technological study of several MPSoC classes*
  - ***Edge-class** platforms*
- *Second year:*
  - *1st part: **Server-class** platforms*
  - *2nd part: Energy-efficient application-class **MPSoC design***
- Third year:
  - Finalize the work on the three tracks
  - Structured proposal for the new era of **energy efficient heterogeneous and scalable HPC**

Thank you for the attention