



PhD in Information Technology and Electrical Engineering

Università degli Studi di Napoli Federico II

PhD Student: Mirko Gagliardi

XXXI Cycle

Training and Research Activities Report – First Year

Tutor: Alessandro Cilardo

1. Information

- Mirko Gagliardi, master's degree in Computer Engineering in 2015 from the University of Naples Federico II.
- XXXI Cycle – ITEE
- Doctoral Fellowship funded by CeRICT
- Tutor: Alessandro Cilardo
- Collaboration with CeRICT in the context of European project MANGO. "MANGO: exploring Manycore Architectures for Next-GeneratiOn HPC systems" is a large-scale European project exploring heterogeneous manycore architectures for HPC, and innovative architectural mechanisms for High Performance, Power-efficiency and time-Predictability tomorrow HPC systems.

2. Study and Training activities

- MS Modules:
 - Advanced Computer Architecture and GPU Programming*, Alessandro Cilardo, 07/01/2016, 6 CFU;
 - Intelligenza Artificiale*, Flora Amato, 26/09/2016, 6 CFU;
- Ad Hoc Modules:
 - Communicating and Disseminating your Research Work*, Mo Mansouri, 16/03/2016, 3 CFU;
 - Scientific Writing*, Paolo Russo, 24/05/2016, 5 CFU;
 - Testing Automation*, Porfirio Tramontana, 12/01/2017, 3 CFU;
 - Le imprese e la ricerca*, Sforza Antonio, 28/02/2016, 4 CFU;
 - Interoperability, Semantic Technologies and Applications*, Flora Amato, 24/03/2017, 2 CFU;
- Seminars
 - Networks-On-Chip: Introduction And Advanced Topics*, Josè Flicht, 16/11/2015 – 18/11/2015, 1.8 CFU (1 extra credit for supplementary activities);
 - Adversarial Testing of Protocol Implementations*, Prof. Cristina Nita Rotaru, 23/09/16, 0.4 CFU;
 - Programmable Network conjugation*, Roberto Bifulco, 26/02/2016, 0.4 CFU;
 - Speech Technologies At Trinity College*, Loredana Cerrato, 02/03/2016, 0.2 CFU;
 - Challenging Real-Time Measurement Systems For Immersive Life-Size Augmented Environment*, Giovanni Caturano, 29/04/2016, 0.5 CFU;
 - Methodologies for embedded software validation*, Diego Tornese, 25/05/2016, 0.2 CFU;
 - Internet: la dimensione immateriale dell'esistenza*, Stefano Quintarelli, 19/05/2016, 0.4 CFU;
 - DDoS Detection in Cloud and Campus Networks*, Jill Jermin, 06/06/2016, 0.2 CFU;

- ix. *An Overview on Image Forensics with emphasis on physics-based scene verification*, Christian Riess, 18/05/2016, 0.2 CFU;
- x. *Half Day EMC Design and troubleshooting Course*, Arturo Mediano, 29/09/2016, 0.8 CFU;
- xi. *Modelling Critical Infrastructures (Cis)*, Peter Popov, 3-11-2016, 0.4 CFU;
- xii. *Security Operations in una Telco, esperienze e riflessioni dal campo*, Pietro Fabio Zamparelli, 11-11-2016, 0.4 CFU;
- xiii. *Minix 3: A reliable and Secure Operating System*, A. Tanenbaum, 30-11-2016, 0.4 CFU;
- xiv. *L' AIS ed il Sistema Nazionale per il monitoraggio marittimo*, Massimo Marazzo, 16-12-2016, 0.4 CFU;
- xv. *Video Coding / Transcoding in Heterogeneous HPC*, Mario Kovac, 7-12-2016, 0.4 CFU;
- xvi. *IBM Cognitive Computing: Challenges and Opportunities in Building an Artificial Intelligence Platform for Business*, Pietro Leo, 17-02-2017, 0.4 CFU;
- xvii. *Cognitive Computing and da Vinci robot: Research Proposal and Discussion*, Paolo Maresca, 17-02-2017, 0.2 CFU;
- xviii. *Smart Nanodevices for Theranostics*, Ilaria Rea, 24-02-2017, 0.3 CFU;
- xix. *Fuzzy Logic, Genetic Algorithms and their application to Next Generation Networks*, Leonard Bapoli, 14-03-2017, 0.8 CFU;
- xx. *From Mathematical formalization to Artificial Visual-Attention: Toward a Human-Like Robot Vision*, Kurosh Madani, 04-04-17, 0.4 CFU;
- xxi. *Living bots and alter ego*, Tamburrini Guglielmo, 04-04-2017, 0.4 CFU;
- xxii. *Dataflow SuperComputing for BigData*, Flora Amato, 12-04-2017, 0.6 CFU;

	Credits year 1								Credits year 2							
	Estimated	1	2	3	4	5	6	Summary	Estimated	1	2	3	4	5	6	Summary
Modules	20	0	6	3	5	0	6	20	10	0	7	2	0	0	0	9
Seminars	5	1.8	0.8	0.8	1.2	0	0.8	5.4	5	2	0.9	2.2	0	0	0	5.1
Research	35	7	4	6	5	8	5	35	45	6	4	6	10	9	10	45
	60	8.8	10.8	9.8	11.2	8.0	11.8	60.4	60	8	12	10	10	9	10	59

3. Research Activities

- Title: *Exploring Advanced Computer Architectures and Techniques to Improve Performance and Power Efficiency in Many Core Era*

Nowadays, **Exascale Computing** is constantly gaining importance in the Scientific Community and **High-Performance Computers** are now considered essential for some experiments in many sectors, such as digital signal processing applications, scientific intense calculation, computer graphics and even cinematic applications. However, most high-performance architectures mostly rely on general-purpose compute units such as CPUs and/or GPUs, which deliver adequate performance while ensuring programmability and application portability. Unfortunately, pure general-purpose hardware is affected by inherently limited power-efficiency (low GFLOPS-per-Watt) now considered as a primary metric.

In this context, current trends in HPC are increasingly moving towards **heterogeneous platforms**, systems made of different computational units, with specialized accelerators that complement general purpose CPUs, including digital signal processors (DSPs) or graphics processing units (GPUs), co-processors, and custom acceleration logic, enabling significant benefits in terms of both power and performance. Historically, HPC has never extensively relied on **FPGAs**, mostly because of the reduced support for floating-point arithmetic, as FPGAs are typically fixed-point oriented. Modern **Altera** and **Xilinx** FPGA families focus on floating-point operations reaching **10 TFLOPs** ideal peak performance. Furthermore, FPGAs are historically power efficient. Those are the essential reasons while very recent trends are putting more emphasis on the potential role of FPGAs, beyond very special-purpose acceleration. Architectural customization can play a key role in modern high-performance architectures, as it enables unprecedented levels of power-efficiency compared to CPUs/GPUs. In line with the above trends, academic and industrial research is focusing on **GPU-like** paradigms to introduce some form of programmability in FPGA design. In the last years, a few GPU-like projects have appeared, such as: Nyuzi, flexgrip, Maven VT, MIAOWGPU.

My research activity investigates the architectural requirements of emerging HPC applications and the most suitable power efficient and high performance solutions. My final goal is to model application-driven accelerator based on GPU-like paradigm that can best fit new emerging HPC application requirements. In this context, our research group founded the **Nu+ project**.

Nu+ project main focus is parametrization. Our aim is to build an application-driven architecture to achieve the best hardware/software configuration for any data-parallel kernel. Specialized data-parallel accelerators have been known to provide

efficiency than general-purpose processors for codes with significant amounts of regular data-level parallelism (DLP). However every parallel kernel has its own "ideal" configuration. The Nu+ architecture is a GPU-like open-source softcore. This HPC-oriented accelerator merges SIMT paradigm with a vector processor model. The core is RISC-based in-order pipeline. Its control unit is intentionally light. Our architecture masks memory and operation latencies deeply relying on hardware multithreading. Our aim is to keep the architecture control logic light and focus on a heavily parallelism and high data-parallel kernel oriented.

Nu+ project is part of a larger H2020 project named MANGO. The main MANGO objective is to define new-generation high-performance, power-efficient, deeply heterogeneous architectures.

4. Products

- Publications:
 - i. Cilaro Alessandro, Gagliardi Mirko. "*Customizable Heterogeneous Acceleration for Tomorrow's High-Performance Computing.*" High Performance Computing and Communications (HPCC), 2015 IEEE 7th International Symposium on Cyberspace Safety and Security (CSS), 2015 IEEE 12th International Conference on Embedded Software and Systems (ICSS), 2015 IEEE 17th International Conference on. IEEE, 2015.
 - ii. Cilaro Alessandro, Gagliardi Mirko, Donnarumma Ciro, "*A Configurable Shared Scratchpad Memory for GPU-like Processors.*" International Conference on P2P, Parallel, Grid, Cloud and Internet Computing. Springer International Publishing, 2016.
 - iii. Cilaro Alessandro, Gagliardi Mirko, Passaretti Daniele, "*NoC-Based Thread Synchronization in a Custom Manycore System.*" International Conference on P2P, Parallel, Grid, Cloud and Internet Computing. Springer International Publishing, 2017.