Last Update: Jan 2024

Alessandro Capotondi

Fixed Term Research Fellow (RTD-A)
Department of Physics, Informatics and Mathematics (FIM)
University of Modena and Reggio Emilia

Email: alessandro.capotondi@unimore.it
ORCID: https://orcid.org/0000-0001-8705-0761

EDUCATION

Ph.D, Electrical, Electronic and Information Engineering. (2016)
Alma Mater Studiorum Università di Bologna
Dissertation: *Programming Models and Tools for Many-Core Platforms*
Dissertation Advisors: Luca Benini, Prof., Andrea Marongiu, Prof.
Abstract:
Parallel programming models are key to effectively harness the computational power of heterogeneous manycore SoC. This thesis presents a set of techniques and HW/SW extensions that enable performance improvements and that simplify programmability for heterogeneous many-core platforms. The thesis contributions cover vertically the entire software stack for many-core platforms, from hardware abstraction layers running on top of bare-metal, to programming models; from hardware extensions for efficient parallelism support to middleware that enables optimized resource management within manycore platforms.

Master, Computer Engineering, Alma Mater Studiorum Università di Bologna. (2012)
Thesis: *Sviluppo e Ottimizzazione di Runtime Layer per OpenMP su Piattaforma Embedded Many-Core (Desing and Optimization of OpenMP Runtime for Embedded Many-Core Platform)*
Thesis Advisor: Luca Benini, Prof., Andrea Marongiu, Prof.

Bachelor, Computer Science, Università degli Studi di Urbino. (2008)

TEACHING EXPERIENCE

Teaching Assistant, Università di Modena e Reggio Emilia, (2020-present)
Courses: *High Performance Computing (Master Degree)*
Heterogeneous Parallel Architectures Laboratory
- OpenMP Programming
- CUDA Programming
- Xilinx HLS HW Desing

Teaching Assistant, Alma Mater Studiorum Università di Bologna, (2015-2019)
Courses: *Hardware-Software Design of Embedded System M*
Heterogeneous Parallel Architectures
- OpenMP Programming

Teaching Assistant, Alma Mater Studiorum Università di Bologna. (2019)
Courses: *Fondamenti di Informatica T (Informatics Fundamentals)*
- C Programming

RESEARCH EXPERIENCE

Postdoctoral Fellow, Department of Physics, Informatics and Mathematics (FIM)
University of Modena and Reggio Emilia. (2019-current)
Research activities cover heterogeneous programming models for autonomous vehicles, such as Drones and Car; the activities aim to create more easy-to-use, productive, and predictable abstractions to effectively use hardware accelerator mapped on FPGA SoC. The work is based on HERO (https://pulp-platform.org/hero.html), a RISCV-V heterogeneous many-core accelerator developed in collaboration with ETH Zurich and the University of Bologna. Another active branch of research is focusing on efficient inference techniques for ultra-low-power edge computing, such as low-bitwidth quantization, TinyML, and AutoML.

Postdoctoral Fellow, Department of Electrical, Electronic and Information Engineering (DEI)
Alma Mater Studiorum Università di Bologna. (2016-2019)
Actively worked on several research topics such as Parallel Programming Model design for many-core accelerators; Hardware/Software exploration of many-core extensions for efficient work parallelization; Programmability improvements for FPGA-based CNN accelerators; Definition of new directives and offloading paradigms in OpenMP for novel many-core heterogeneous architectures; Study of efficient workload dispatch on NUMA affected many-core platforms.

Visiting Researcher, Integrated Systems Laboratory, ETH Zurich. (2016)
Co-creator, compiler and OpenMP designer of Open Heterogeneous Research Platform (HERO). HERO combines a PULP-based open-source parallel manycore accelerator implemented on FPGA with a hard ARM Cortex-A multicore host processor running full-stack Linux. HERO offers a complete hardware and software platform which advances the state of the art of transparent accelerator programming using the OpenMP v4.5 Accelerator Model. The programmer writes a single application source file for the host and uses OpenMP directives for parallelization and accelerator offloading.

GRANTS AND FELLOWSHIPS
Research Grant Winner: H2020 ECSEL COMP4DRONES
Research Grant Winner: H2020 CLASS
Research Grant Winner: POR-FESR 2014-2020 Open-Next
Research Grant Winner: EU ERC Multitherman
Research Grant Winner: EU FP7 Virtical Project

PROFESSIONAL MEMBERSHIPS
Institute of Electrical and Electronics Engineers (IEEE) Member since 2015.

PUBLICATIONS
Google Scholar: 266 citations, H-index: 8

Book and Thesis
[b1] *Alessandro Capotondi:* ***Programming models and tools for many-core platforms.*** *University of Bologna, Italy, 2016*

Journal and Transaction Articles
[j8] *Miguel de Prado, Manuele Rusci, Alessandro Capotondi, Romain Donze, Luca Benini, Nuria Pazos:* ***Robustifying the Deployment of tinyML Models for Autonomous Mini-Vehicles.*** *Sensors 21(4): 1339 (2021)*

[j7] *Paolo Meloni, Daniela Loi, Gianfranco Deriu, Marco Carreras, Francesco Conti, Alessandro Capotondi, Davide Rossi:* ***Exploring NEURAghe: A Customizable Template for APSoC-Based CNN Inference at the Edge.*** *IEEE Embed. Syst. Lett. 12(2): 62-65 (2020)*

[j6] *Alessandro Capotondi, Manuele Rusci, Marco Fariselli, Luca Benini:* ***CMix-NN: Mixed Low-Precision CNN Library for Memory-Constrained Edge Devices.*** *IEEE Trans. Circuits Syst. II Express Briefs 67-II(5): 871-875 (2020)*

[j5] *Alessandro Capotondi, Andrea Marongiu, Luca Benini:* ***Runtime Support for Multiple Offload-Based Programming Models on Clustered Manycore Accelerators.*** *IEEE Trans. Emerging Topics Comput. 6(3): 330-342 (2018)*

[j4] *Igor Loi, Alessandro Capotondi, Davide Rossi, Andrea Marongiu, Luca Benini:*
The Quest for Energy-Efficient I$ Design in Ultra-Low-Power Clustered Many-Cores. *IEEE Trans. Multi-Scale Computing Systems 4(2): 99-112 (2018)*

[j3] *Paolo Meloni, Alessandro Capotondi, Gianfranco Deriu, Michele Brian, Francesco Conti, Davide Rossi, Luigi Raffo, Luca Benini:* ***NEURAghe: Exploiting CPU-FPGA Synergies for Efficient and Flexible CNN Inference Acceleration on Zynq SoCs.*** *TRETS 11(3): 18:1-18:24 (2018)*

[j2] *Andrea Marongiu, Alessandro Capotondi, Luca Benini:*
Controlling NUMA effects in embedded manycore applications with lightweight nested parallelism support. *Parallel Computing 59: 24-42 (2016)*

[j1] *Andrea Marongiu, Alessandro Capotondi, Giuseppe Tagliavini, Luca Benini:*
Simplifying Many-Core-Based Heterogeneous SoC Programming With Offload Directives. *IEEE Trans. Industrial Informatics 11(4): 957-967 (2015)*

Conference and Workshop Papers

[c14] *Andreas Kurth, Koen Wolters, Björn Forsberg, Alessandro Capotondi, Andrea Marongiu, Tobias Grosser, Luca Benini:* ***Mixed-data-model heterogeneous compilation and OpenMP offloading.*** *CC 2020: 119-131*

[c13] *Micaela Verucchi, Gianluca Brilli, Davide Sapienza, Mattia Verasani, Marco Arena, Francesco Gatti, Alessandro Capotondi, Roberto Cavicchioli, Marko Bertogna, Marco Solieri:* ***A Systematic Assessment of Embedded Neural Networks for Object Detection.*** *ETFA 2020: 937-944*

[c12] *Manuele Rusci, Alessandro Capotondi, Luca Benini:* ***Memory-Driven Mixed Low Precision Quantization for Enabling Deep Network Inference on Microcontrollers.*** *MLSys 2020*

[c11] *Manuele Rusci, Marco Fariselli, Alessandro Capotondi, Luca Benini:* ***Leveraging Automated Mixed-Low-Precision Quantization for Tiny Edge Microcontrollers.*** *IoT Streams/ITEM@PKDD/ECML 2020: 296-308*

[c10] *Leonardo Ravaglia, Manuele Rusci, Alessandro Capotondi, Francesco Conti, Lorenzo Pellegrini, Vincenzo Lomonaco, Davide Maltoni, Luca Benini:* ***Memory-Latency-Accuracy Trade-Offs for Continual Learning on a RISC-V Extreme-Edge Node.*** *SiPS 2020: 1-6*

[c9] *Andreas Kurth, Alessandro Capotondi, Pirmin Vogel, Luca Benini, Andrea Marongiu:* ***HERO: an open-source research platform for HW/SW exploration of heterogeneous manycore systems.*** *ANDARE@PACT 2018: 5:1-5:6*

[c8] *Manuele Rusci, Alessandro Capotondi, Francesco Conti, Luca Benini:* ***Quantized NNs as the definitive solution for inference on low-power ARM MCUs?: work-in-progress.*** *CODES+ISSS 2018: 12*

[c7] *Alessandro Capotondi, Andrea Marongiu:* ***Enabling zero-copy OpenMP offloading on the PULP many-core accelerator.*** *SCOPES 2017: 68-71*

[c6] *Alessandro Capotondi, Andrea Marongiu:* ***On the effectiveness of OpenMP teams for cluster-based many-core accelerators.*** *HPCS 2016: 667-674*

[c5] *Alessandro Capotondi, Germain Haugou, Andrea Marongiu, Luca Benini:* ***Runtime Support for Multiple Offload-Based Programming Models on Embedded Manycore Accelerators.*** *COSMIC@CGO 2015: 4:1-4:10*

[c4] *Davide Rossi, Francesco Conti, Andrea Marongiu, Antonio Pullini, Igor Loi, Michael Gautschi, Giuseppe Tagliavini, Alessandro Capotondi, Philippe Flatresse, Luca Benini:* ***PULP: A parallel ultra low power platform for next generation IoT applications.*** *Hot Chips Symposium 2015: 1-39*

[c3] *Alessandro Capotondi, Andrea Marongiu, Luca Benini:* ***Enabling Scalable and Fine-Grained Nested Parallelism on Embedded Many-cores.*** *MCSoC 2015: 297-304*

[c2] *Marco Balboni, Marta Ortín-Obón, Alessandro Capotondi, Hervé Tatenguem Fankem, Alberto Ghiribaldi, Luca Ramini, Víctor Viñals, Andrea Marongiu, Davide Bertozzi:* ***Augmenting manycore programmable accelerators with photonic interconnect technology for the high-end embedded computing domain.*** *NOCS 2014: 72-79*

[c1] *Andrea Marongiu, Alessandro Capotondi, Giuseppe Tagliavini, Luca Benini:* ***Improving the programmability of STHORM-based heterogeneous systems with offload-enabled OpenMP.*** *MES 2013: 1-8*

Informal Publications (Non-Peer Reviewed)

[i6] *Miguel de Prado, Romain Donze, Alessandro Capotondi, Manuele Rusci, Serge Monnerat, Luca Benini, Nuria Pazos:* ***Robust navigation with tinyML for autonomous mini-vehicles.*** *CoRR abs/2007.00302 (2020)*

[i5] *Leonardo Ravaglia, Manuele Rusci, Alessandro Capotondi, Francesco Conti, Lorenzo Pellegrini, Vincenzo Lomonaco, Davide Maltoni, Luca Benini:* ***Memory-Latency-Accuracy Trade-offs for Continual Learning on a RISC-V Extreme-Edge Node.*** *CoRR abs/2007.13631 (2020)*

[i4] *Manuele Rusci, Marco Fariselli, Alessandro Capotondi, Luca Benini:* ***Leveraging Automated Mixed-Low-Precision Quantization for tiny edge microcontrollers.*** *CoRR abs/2008.05124 (2020)*

[i3] *Manuele Rusci, Alessandro Capotondi, Luca Benini:* ***Memory-Driven Mixed Low Precision Quantization For Enabling Deep Network Inference On Microcontrollers.*** *https://arxiv.org/abs/1905.13082 (2019)*

[i2] *Paolo Meloni, Alessandro Capotondi, Gianfranco Deriu, Michele Brian, Francesco Conti, Davide Rossi, Luigi Raffo, Luca Benini:* ***NEURAghe: Exploiting CPU-FPGA Synergies for Efficient and Flexible CNN Inference Acceleration on Zynq SoCs.*** https://arxiv.org/abs/1712.00994 *(2017)*

[i1] *Andreas Kurth, Pirmin Vogel, Alessandro Capotondi, Andrea Marongiu, Luca Benini:*
HERO: Heterogeneous Embedded Research Platform for Exploring RISC-V Manycore Accelerators on FPGA. https://arxiv.org/abs/1712.06497 *(2017)*