

# Meet Monte Cimone: Exploring RISC-V High Performance Compute Clusters

Andrea Bartolini<sup>1</sup>, Federico Ficarelli<sup>1,2</sup>, Emanuele Parisi<sup>1</sup>, Francesco Beneventi<sup>1</sup>, Francesco Barchi<sup>1</sup>, Daniele Gregori<sup>3</sup>,

Fabrizio Magugliani<sup>3</sup>, Marco Cicala<sup>3</sup>, Cosimo Gianfreda<sup>3</sup>, Daniele Cesarini<sup>2</sup>, Andrea Acquaviva<sup>1</sup>, Luca Benini<sup>1</sup>

<sup>1</sup>University of Bologna, Italy <sup>2</sup>CINECA, Italy <sup>3</sup>E4 Company S.p.A., Italy

**Abstract**—The new open and royalty-free RISC-V ISA is attracting interest across the whole computing continuum, from microcontrollers to supercomputers. High-performance RISC-V processors and accelerators have been announced, but RISC-V-based HPC systems will need a holistic co-design effort, spanning memory, storage hierarchy interconnects and full software stack. In this paper, we describe Monte Cimone, a fully-operational multi-blade computer prototype and hardware-software test-bed based on U740, a double-precision capable multi-core, 64-bit RISC-V SoC. Monte Cimone does not aim to achieve strong floating-point performance, but it was built with the purpose of “priming the pipe” and exploring the challenges of integrating a multi-node RISC-V cluster capable of providing an HPC production stack including interconnect, storage and power monitoring infrastructure on RISC-V hardware. We present the results of our hardware/software integration effort, which demonstrate a remarkable level of software and hardware readiness and maturity - showing that the first generation of RISC-V HPC machines may not be so far in the future.

## I. INTRODUCTION

The strategic role of High Performance Computing (HPC) systems is widely acknowledged in many fields, from weather forecasting to drug design. With the pervasive digitalization of our society, high performance computers fuel the most disruptive mega-trends, from the deployment of artificial intelligence (AI) at scale (e.g. for training large machine learning models) to industrial internet-of-things (IoT) applications (e.g. for creating and maintaining digital twins).

The key challenge in designing HPC systems today and in the foreseeable future is increasing compute efficiency, to meet the rapidly growing performance demand (10x every four years) within a constant or modestly increasing power budget, while facing the slow-down of Moore’s Law. To exacerbate the efficiency challenge, while integrated circuits technology is still delivering device density increases (albeit as a slower pace), power consumption does not scale down at the same rate. Hence power density grows and it is increasingly difficult to meet thermal design power specifications without compromising performance. Disruptive technologies, such as quantum or optical computing may bring long-term relief in some specific application areas, but there is no silver bullet in sight.

To tackle the efficiency issue, academia and industry are aggressively pursuing architectural innovation and co-design strategies to develop HPC systems that mitigate the efficiency limitations of programmable architectures through various forms of specialization and domain-specific adaptation. Instruction Set Architectures (ISAs) have to evolve rapidly to sustain architectural

evolution and domain adaptation, and the advent of the RISC-V open, royalty-free and extensible ISA has been a major step toward accelerating innovation in this area.

Currently, high-performance 64bit (RV64) RISC-V processors and accelerator chips are being designed, promising prototypes are demonstrated in numerous publications [3] and products are announced at a fast cadence [1], [2]. It is thus reasonable to expect that high-performance chips based on RISC-V will be available as production silicon within the next couple of years. However, building a HPC system requires significantly more than just high-performance chips. Many think that the RISC-V software stack and system platform are extremely immature, and will need several additional years of development effort before full applications could be run, benchmarked and optimized on a RISC-V-based HPC system. Our goal is dispel this overly conservative notion.

The main contribution of this work is to present Monte Cimone, the first physical prototype and test-bed of a complete RISC-V (RV64) compute cluster, integrating not only all the key hardware elements besides processors, namely main memory, non-volatile storage and interconnect, but also a complete software environment for HPC, as well as a full-featured system monitoring infrastructure. Further, we demonstrate that it is possible to run full fledged, mainstream HPC applications on Monte Cimone today.

To the best of our knowledge, this is the first RISC-V cluster which is fully operational and supports a baseline HPC software stack, proving the maturity of the RISC-V ISA and the first generation of commercially available RISC-V SoCs. We also evaluated the support for Infiniband network adapters which are recognised by the system, but are not yet capable to support RDMA communication.

Future work will focus on improving the software stack to achieve higher memory utilisation (i), to implement dynamic power and thermal management (ii), overcome the limitation in the Infiniband support (iv), extend Monte Cimone with PCIe RISC-V based accelerators (v).

## REFERENCES

- [1] RISC-V Exchange: Cores & SoCs. <https://riscv.org/exchange/cores-socs>.
- [2] Alexander Dörflinger, Mark Albers, and et al. A comparative survey of open-source application-class risc-v processor implementations. In *Proceedings of the 18th ACM International Conference on Computing Frontiers*, CF ’21, page 12–20, 2021.
- [3] John L. Hennessy and David A. Patterson. A new golden age for computer architecture. *Communications of the ACM*, 62(2):48–60, 2019.