REPTILES: Repeated Tiles of Sargantana, a RISC-V multicore based on OpenPiton

Noelia Oliete-Escuín[1]*, Arnau Bigas-Soldevila[1], Narcís Rodas[1], Albert Aguilera[1], Sajjad Ahmad[1], Jonathan Balkind[2] †, Xavier Carril[1], Max Doblas[1], Ivan Díaz[1], Roger Figueras[1], Alireza Foroodnia[1], Cesar Fuget[3] ‡, Ignacio Genovese[1], Raúl Gilabert[1], Abbas Haghi[1], Alexander Kropotov[1], Neiel Leyva[1], Oscar Lostes-Cazorla[1], Lorién López-Villellas[4] §, Davy Million[5] ¶, Alireza Monemi[1], Sérik Pérez[1], Juan Antonio Rodríguez[1], Víctor Soria-Pardos[1], Behzad Salami[1], Francesc Moll[1], Oscar Palomar[1], Miquel Moretó[1] and Lluc Alvarez[1]

April 21, 2025

Abstract

Chip industry continues advancing and expanding modern computing systems, resulting in more complex multicore processors. Conversely, academic projects face scalability challenges due to limited resources, highlighting the need for open-source frameworks that enable innovation and knowledge sharing. Recently, several open-source proposals have emerged, offering flexible and scalable designs, but fail to meet the performance demands of modern High-Performance Computing (HPC) applications. In this project, we present REPTILES, an open-source RISC-V multicore framework based on OpenPiton‖. REPTILES interconnects multiple Sargantana cores with the memory hierarchy of OpenPiton. Moreover, we present the new features incorporated in Sargantana and OpenPiton designs to improve the performance of HPC applications. We demonstrate that REPTILES presents suitable scalability, achieving a speedup of 3.1× on average with 4 cores. Additionally, we show that Sargantana's new features increase the performance of vector addition benchmark in a 9.3×.

Introduction

The chip industry continues developing and scaling modern systems, leading to increasingly complex multicore processors. However, academic projects struggle to achieve similar scalability due to limited resources and infrastructure. To overcome these challenges, the community needs open-source architecture frameworks where researchers can explore and design their ideas. Open-source frameworks provide flexibility and transparency, fostering innovation and collaboration.

In recent years, different proposals have emerged to provide open-source solutions. Nevertheless, with the growing demand for High-Performance Computing (HPC), current open-source projects struggle to meet the performance requirements of modern applications.

In this project, we present **REPTILES**, **REP**eated **TILE**s of **S**argantana, a RISC-V multicore based on OpenPiton. REPTILES is an open-source multi-core architecture that aims to provide an accessible HPC framework where researchers can develop, experiment with, and optimize HPC applications. Additionally, we detail the high-performance features included in the design.

We show that REPTILES presents strong scalability, achieving an average speedup of 3.1× with 4 cores. Furthermore, we demonstrate that Sargantana's new features enhance the performance of the vector addition benchmark by 9.3×.

System Overview and Features

The multicore system presented in this poster is based on OpenPiton and Sargantana, publicly available open-source designs based on Verilog and SystemVerilog. OpenPiton is a multi-core framework. The cores themselves are replicated Sargantana tiles, single-issue in-order RISC-V 64-bit processors.

We evaluate REPTILES in an FPGA prototype that includes 4 Sargantana cores, without SIMD unit, and a cache hierarchy composed of an in-house 16 KB L1 instruction cache, a 16 KB High-Performance L1 data cache, a 32 KB L1.5 cache, and a 64 KB shared L2 cache with 64 B cache blocks. The private cache levels are configured with 64 MSHRs and connected via 64-bit NoC buses to the L2 cache. The evaluated configuration is limited due to the resource limitations of the Alveo u55 FPGA. In addition, the FPGA prototype includes 16 GB of HBM main memory and Ethernet support. These features enable the use of a Software Development Vehicle (SDV) where a shared filesystem can be mounted, facilitating extensive benchmarking

*Corresponding author: Barcelona Supercomputing Center, e-mail: noelia.oliete@bsc.es
†University of California, Santa Barbara
‡University Grenoble Alpes, Inria
§University Zaragoza
¶University Grenoble Alpes, CEA
‖REPTILES open-source code: https://github.com/bsc-loca/openpiton

and live interactive demos.

OpenPiton Improvements and New Open-Sourced Features

The increasing need for HPC calls for open-source solutions. Initiatives like OpenPiton [1] have been developed to support scalable and customizable architectures. Nevertheless, these frameworks often face performance limitations that restrict their capability to handle compute-intensive workloads effectively.

OpenPiton operates in different architectures, such as SPARC v9, x86 and RISC-V architectures (RISC-V 32-bit and RISC-V 64-bit). The architecture of OpenPiton consists of a chipset and one or more tiles. The chipset includes modules that connect the tiles with the peripherals, such as the UART. Each tile integrates three NoC routers, the core, and the cache hierarchy. This hierarchy includes private L1 instruction and data caches, a private L1.5 cache, and a shared distributed L2 cache that implements a directory-based MESI coherence protocol. Although OpenPiton provides numerous advantages, it encounters performance constraints that limit its suitability for HPC domains.

Recent work proposes a set of improvements to the NoC and the memory hierarchy to upgrade OpenPiton to meet the HPC requirements [2]. In this work, we integrate these features into our design and introduce new ones. Specifically, the parametric NoC width from 64 bits up to 704 bits, the support for configurable cache block sizes (64, 32, and 16 bytes) for all the cache levels and the parametric number of MSHRs, associativity and the parallel SRAM access for the L1.5 and L2 caches. Additionally, we enhance our design with the connection of Sargantana with the High-Performance Data Cache (HPDcache) [3].

Sargantana Improvements and New Open-Sourced Features

Sargantana is a Linux-capable 64-bit RISC-V processor that implements the RV64G ISA and achieves a 1.26 GHz frequency using a 22nm technology node [4]. Since its open-source release, it has received several improvements, such as architecture support for more RISC-V extensions and general usability enhancements to the RTL simulation environment.

The most significant change has been the upgrade from the RISC-V Vector Extension (RVV) 0.7 version to 1.0. In [4], Sargantana only supported a small set of arithmetic vector instructions that could be used when manually vectorizing specific codes. Currently, the core supports most of the extension specifications, except for the LMUL>1 setting and vector FP instructions. It also implements renaming for vector configuration instructions (previously, they stalled the pipeline), leading to more remarkable performance in vectorized codes.

Other notable added extensions are Sdext for debugging support and Sscofpmf to enable reading the core performance counters in Linux via perf.

Regarding the RTL simulation environment, we added support for the saving and restoring feature using Verilator. This allows users to periodically create checkpoints during simulation that store the model of the design in an intermediate state. Later, the simulation can be resumed from that point without re-running, which can be very helpful for debugging.

System Evaluation

We evaluate the performance of REPTILES running the NAS Parallel Benchmarks with OpenMP over Linux in the FPGA system. Figure 1 shows the speedup achieved on each benchmark with 2, 3, and 4 threads with respect to 1 thread. When increasing the number of threads, we can observe suitable scalability for all the benchmarks and significant performance improvements. The configuration with 4 threads also achieves a speedup of 3.6× for the CG and EP benchmarks and an average of 3.1×.

Additionally, we analyze the performance of Sargantana with RVV extension executing a vector addition benchmark of 8-bit elements on a standalone basis. We observe that the Sargantana RVV version achieves a 9.3× speedup with respect to the scalar version.

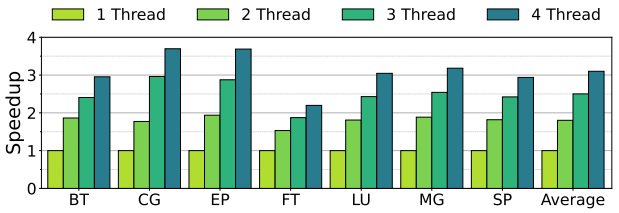

Figure 1: *Execution speedup of NAS benchmarks over 2, 3, and 4 threads with respect to 1 thread.*

References

[1] Jonathan Balkind et al. "OpenPiton: An Open Source Manycore Research Framework". In: *SIGARCH Comput. Archit. News* 44.2 (Mar. 2016), pp. 217–232. ISSN: 0163-5964. DOI: 10.1145/2980024.2872414. URL: https://doi.org/10.1145/2980024.2872414.

[2] Neiel Leyva et al. "OpenPiton4HPC: Optimizing OpenPiton Toward High-Performance Manycores". In: *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* 14.3 (2024), pp. 395–408. DOI: 10.1109/JETCAS.2024.3428929.

[3] César Fuguet. "HPDcache: Open-Source High-Performance L1 Data Cache for RISC-V Cores". In: *Proceedings of the 20th ACM International Conference on Computing Frontiers*. CF '23. Bologna, Italy: Association for Computing Machinery, 2023, pp. 377–378. ISBN: 9798400701405. DOI: 10.1145/3587135.3591413. URL: https://doi.org/10.1145/3587135.3591413.

[4] Victor Soria et al. "Sargantana: A 1 GHz+ in-order RISC-V processor with SIMD vector extensions in 22nm FD-SOI". In: *2022 25th Euromicro Conference on Digital System Design (DSD)*. IEEE. 2022, pp. 254–261.