

## Research context and motivation

**Computers are dead** † Aside from the undeniable yet expensive benefits brought by the advancements in the semiconductor manufacturing process, integrated computer systems **have not drastically changed** since their introduction in the 70s. The combined action of a shift in the programming paradigm towards **data-driven algorithms** and the ever-increasing **difficulties and costs** involved when trying to keep pace with Moore's Law exposed the limitations of traditional computer architectures:

- Researchers have pushed the state of the art in existing techniques (i.e., branch predictors, prefetching strategies) to the point that performance improvements hardly ever reach even **a few percentage points**.
- Traditional Von Neumann architectures **struggle to feed the hardware computing engines** with the massive amount of data required by modern computationally expensive workloads (did anyone mention machine learning yet?)

**Long live the computers** 🖥️ In this environment, the **open and extensible** RISC-V instruction set architecture (ISA) established itself as the perfect **architectural and microarchitectural playground**, paving the way to a so-called "Renaissance" in computer architecture towards the definition of next-generation **heterogeneous** computer systems.

## Addressed research questions/problems


**Open hardware for all!** 🌐 While the consolidation of RISC-V as the ISA of choice in both industry and academy is an excellent achievement for the affirmation of "open hardware", it comes with the side-effect of **fragmentation**:

- Several RISC-V implementations ⇒ **disperse** the scientific effort (PULP, among the others) ⇒ **steep learning curve** for new adopters
- Inconsistent software and OS support among frameworks (e.g., EEI).
- Several commonly used implementations are impractical to modify.

**Need for specialization** 🧑‍🔬 Most RISC-V implementations fit in the microcontroller range of products, however:

- Modern application software requires hardware support for **computationally-intensive workloads**, even on **edge devices** relying on limited power sources.
- In these systems, memory transfers often represent a bottleneck for performance and an unacceptable overhead for energy efficiency.

Both these aspects can be addressed using efficient and **specialized parallel accelerators**.



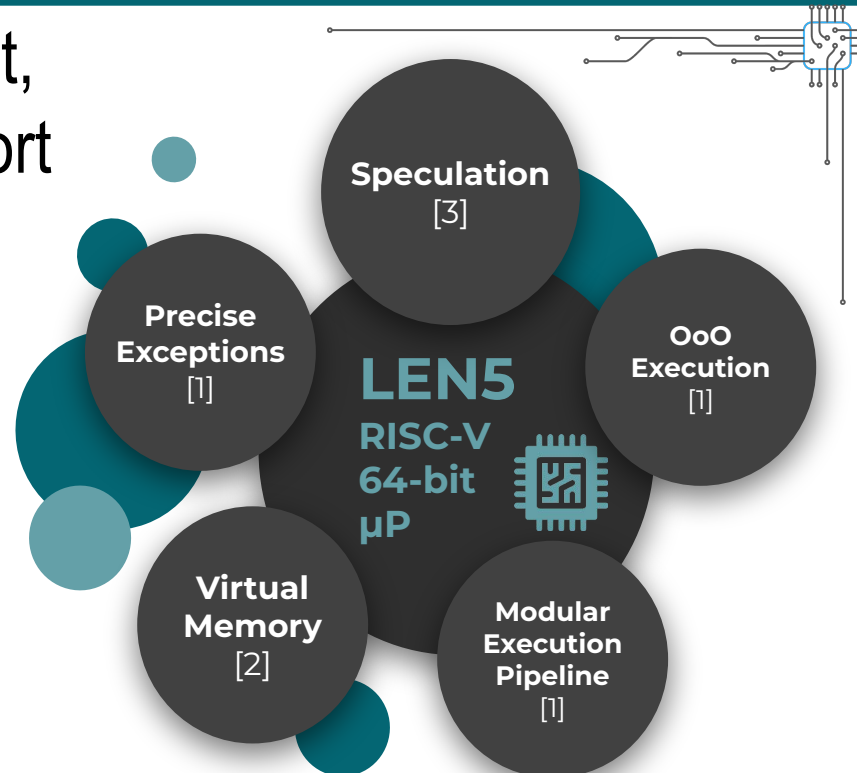
**Introducing LEN5** Starting from the assumptions above, the **LEN5** RISC-V microprocessor is being developed as a **modular** platform designed to be as simple as possible to configure, extend, and modify.

## Novel contributions

**LEN5** (Fig. 2) is a modular, highly-configurable, 64-bit, out-of-order, speculative RISC-V microprocessor with support for virtual memory, described in **SystemVerilog**.

**Features** 🧩

- RV64I ISA with "Zicsr" extension (M extension is coming)
- Small **prefetch queue** to compensate for memory latency
- Global **gshare branch predictor** and branch target buffer
- **Dynamically scheduled, out-of-order**, execution pipeline
- System-wide AXI-like (valid-ready) **handshaking** protocol
  - EU **latency-independent** execution pipeline
  - Easy to expand with **custom tightly-coupled EUs**
- Optional 2-level cache subsystem with hardware TLB and memory protection
- Optional tightly-coupled, edge-oriented vector supporting a subset of the V extension



[1] M. Caon, M. Martina, "Design of the Execution Pipeline for LEN5, an Out-of-order RISC-V Processor," 2019.  
 [2] M. Perotti, M. Martina, "Design of an OS Compliant Memory System for LEN5, a RISC-V Out of Order Processor," 2019.  
 [3] M. Andorno, M. Martina, "Design of the Frontend for LEN5, a RISC-V Out-of-Order Processor," 2019.

## Adopted methodologies




Fig. 1 LEN5 "spill cell".

**Modularity** 🧩 **"Spill cells"** (Fig. 1) used as output register:

- ✓ Break the **ready** signal back-propagation in a **transparent** way for upstream and downstream hardware
- ✓ They can be easily **bypassed** if they are not in the critical path
- ✓ **Reservation stations** for each EU, allowing for self-contained control

**Configurability** ⚙️

- ✓ Self-documented **configuration files**
  - Enable/disable features and extensions
  - Modify buffers size, pipeline depth, etc.
- ✓ make-based HW/SW build system

**Software** 💻

- ✓ "PULP-like" bootstrap code and bare-metal system calls.
- ✓ Automated QEMU debug via GDB.

**Hardware** 🛠️

- ✓ **SystemVerilog** RTL description
- ✓ Automated memory image creation (from SW) and simulation

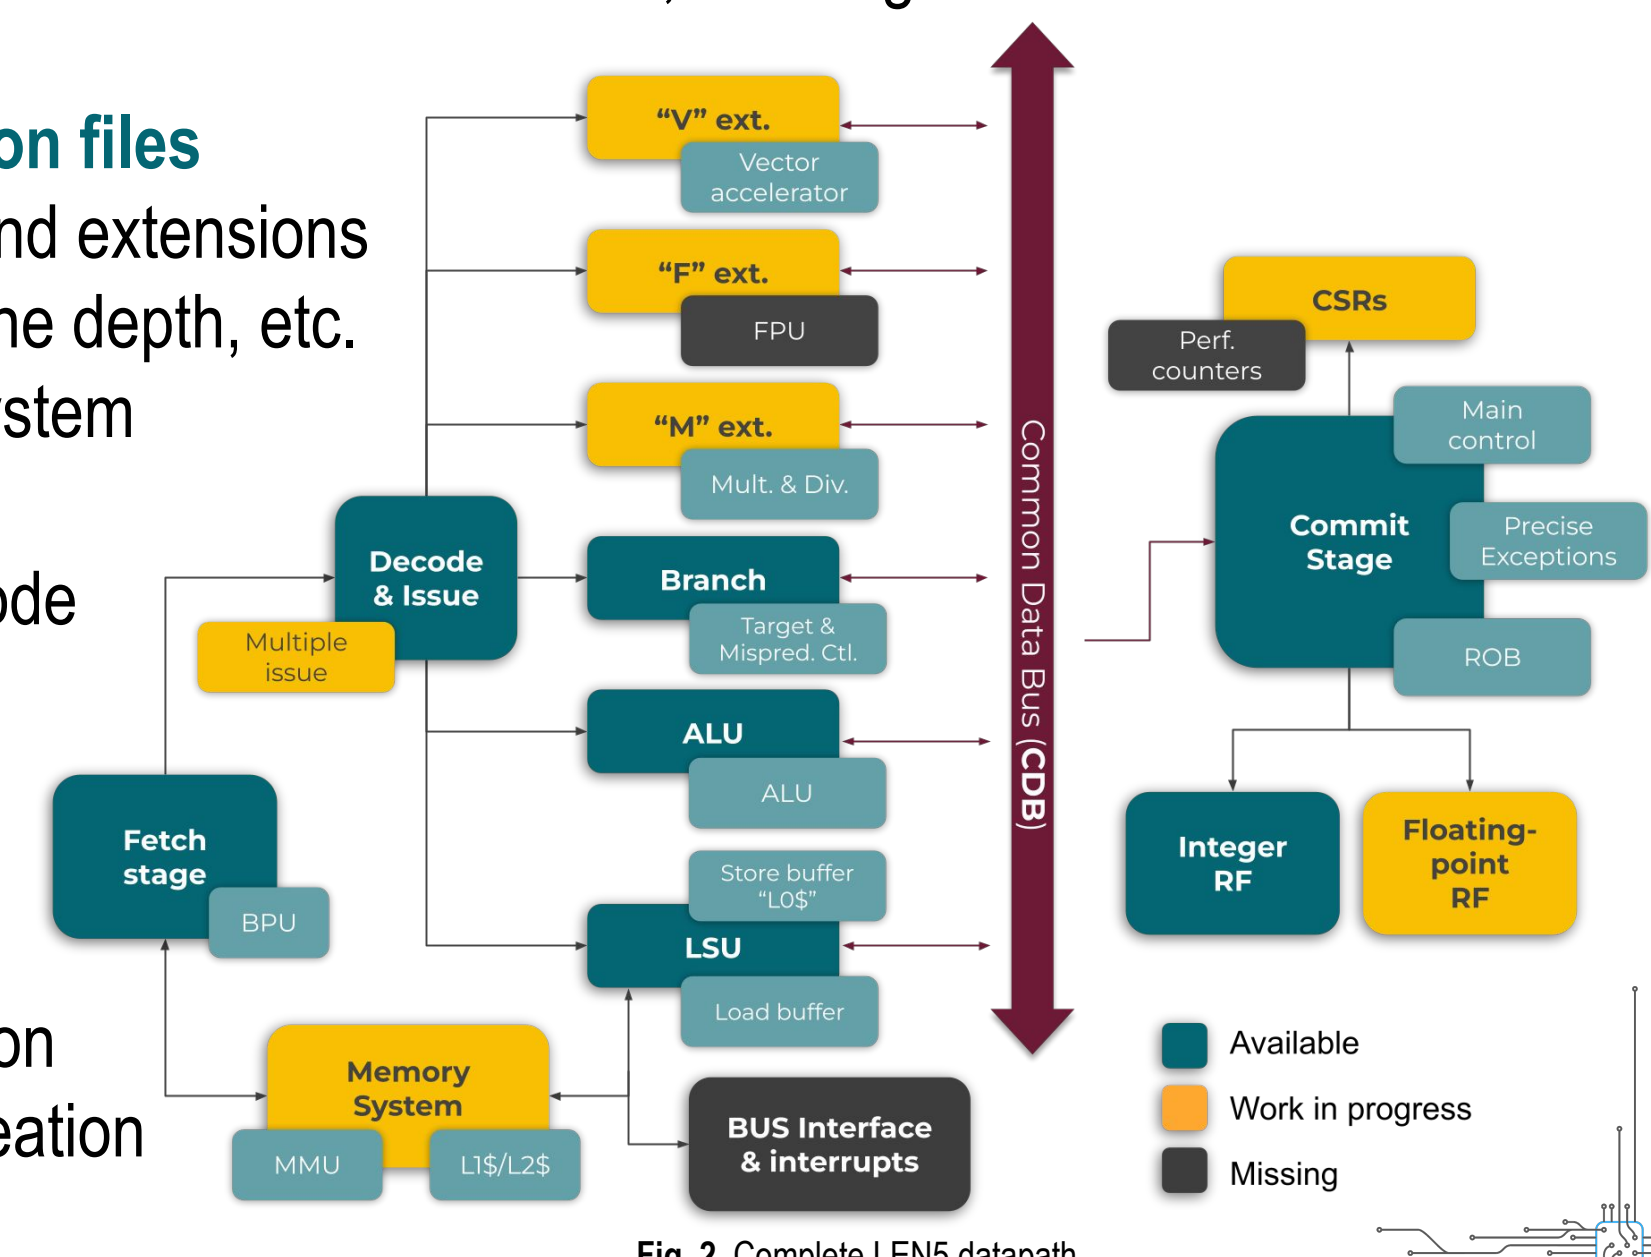


Fig. 2 Complete LEN5 datapath.

## Preliminary results

**Hello, world!** 🌍

A 280-instruction program printing a programmer's notorious first words is executed by **LEN5** with an **average IPC of 0.4**. Despite its simplicity, about 30% of the programs is represented by **isolated jump or branch instructions** that stress **LEN5's** branch prediction unit and force frequent flushes (7-cycle delay) of the execution pipeline.




Fig. 3 Preliminary synthesis results on UMC 65nm.

Better results shall be expected in proper benchmarks, that are underway.

**Synthesis** 🧩 Fig. 3 reports preliminary ASIC synthesis results obtained with four increasingly complex configurations of **LEN5**. The core is about 3 times larger than a simpler in-order core in its base version and can run at 700MHz.

## Future work

**Work in progress** 🧑‍🔬

The following milestones are under active development:

- Dual-issue, to properly take advantage of the out-of-order nature of **LEN5**
- Proper benchmarking to assess where **LEN5** places in the state of the art
- Edge-oriented tightly coupled vector accelerator

**Critical features** 🧩

Despite **LEN5** being in working conditions, some additional features are probably required to make it a suitable platform for advanced research projects:

- C (compressed) and A (atomic) instruction set extensions
- Software and hardware interrupt support, at least in machine-mode

**Optimization** ⚙️

The core components of **LEN5** would certainly benefit from an attentive analysis and optimization. This is probably a required step to compete with existing solutions.

## List of attended classes

**Hard skills:**

• 01UMNRV	Advanced deep Learning (didattica di eccellenza)	15/06/2021	6 CFU
• 01UKAIU	Advanced techniques for digital testing	04/07/2022	4 CFU
• 01SHORV	Nano & Quantum Computing	16/12/2021	8 CFU
• 02SFURV	Programmazione scientifica avanzata in matlab	25/05/2021	6 CFU
• 01DNHRV	System level low power techniques for IoT	15/07/2022	4 CFU
• 01TSBRV	Scienza dei dati applicata alle reti complesse	23/07/2021	4 CFU

**Soft skills:**

• 02LWHRV	Communication	29/03/2021	1 CFU
• 01PJMRV	Etica informatica	03/05/2021	4 CFU
• 01UNYRV	Personal branding	02/04/2021	1 CFU
• 01RISRV	Public speaking	29/03/2021	1 CFU
• 01SYBRV	Research integrity	14/04/2021	1 CFU
• 02RHORV	The new Internet Society: entering the black-box of digital innovation	16/06/2021	1 CFU
• 01UNXRV	Thinking out of the box	29/01/2021	1 CFU
• 01SWPRV	Time management	30/03/2021	1 CFU