



*Designing RISC-V-based Accelerators  
for next generation Computers*

# DRAC Designing RISC-V-based Accelerators for next generation Computers

Miquel Moretó (BSC and UPC)

European RISC-V HPC Roadmap Workshop, Toulouse, January 16 2023



***The RISC-V Revolution!***

# Today's Technology Trends




Massive penetration of Open Source Software

- IoT (Arduino),
- Mobile (Android),
- Enterprise (Linux),
- HPC (Linux, OpenMP, etc.)



Moore's Law + Power = Specialization (HW/SW Co-Design)

- More cost effective
- More performant
- Less Power

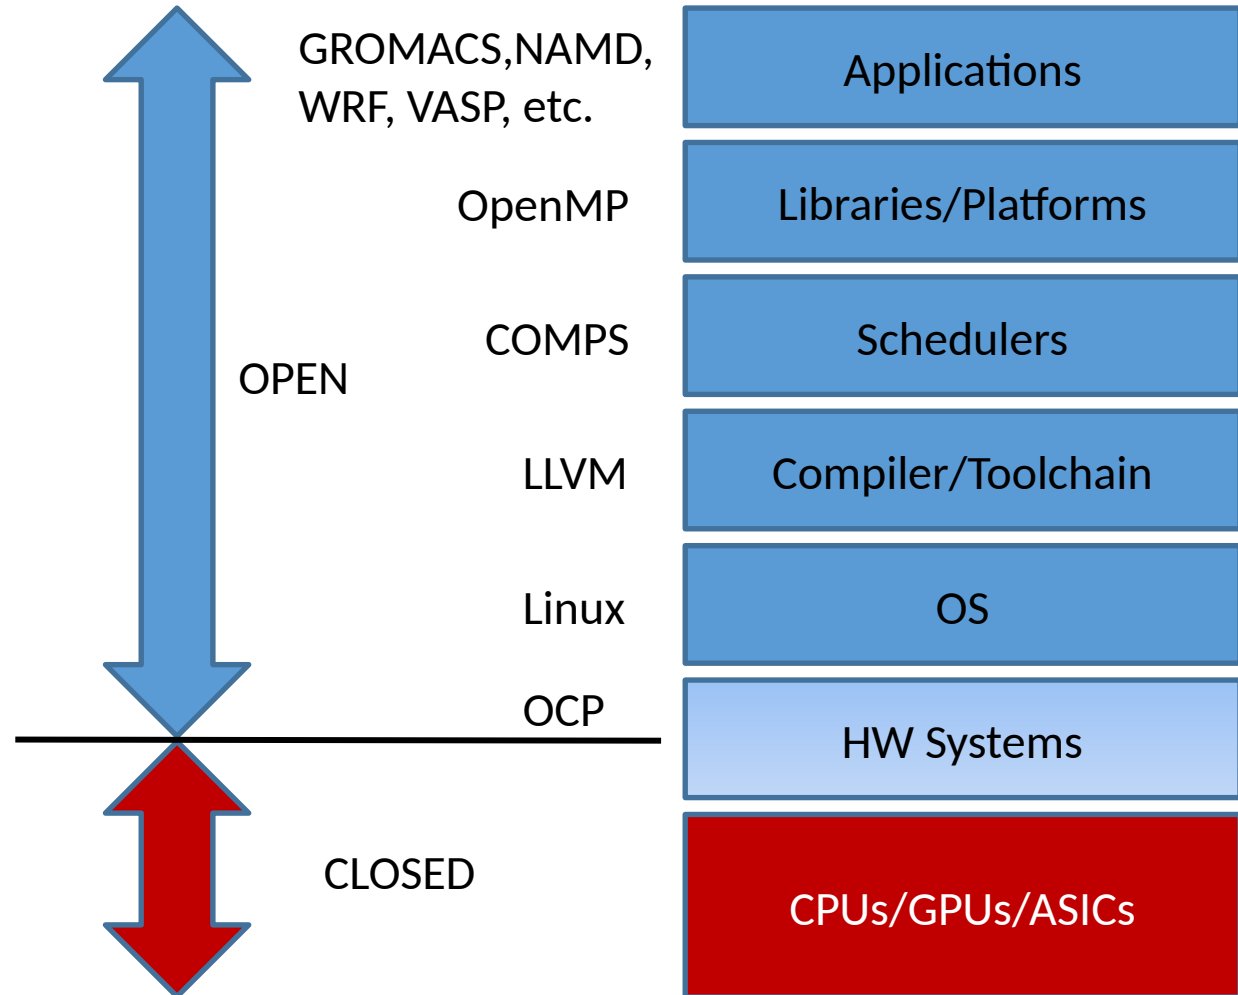


New Open Source Hardware Momentum from IoT and the Edge to HPC

- RISC-V
- OpenPOWER

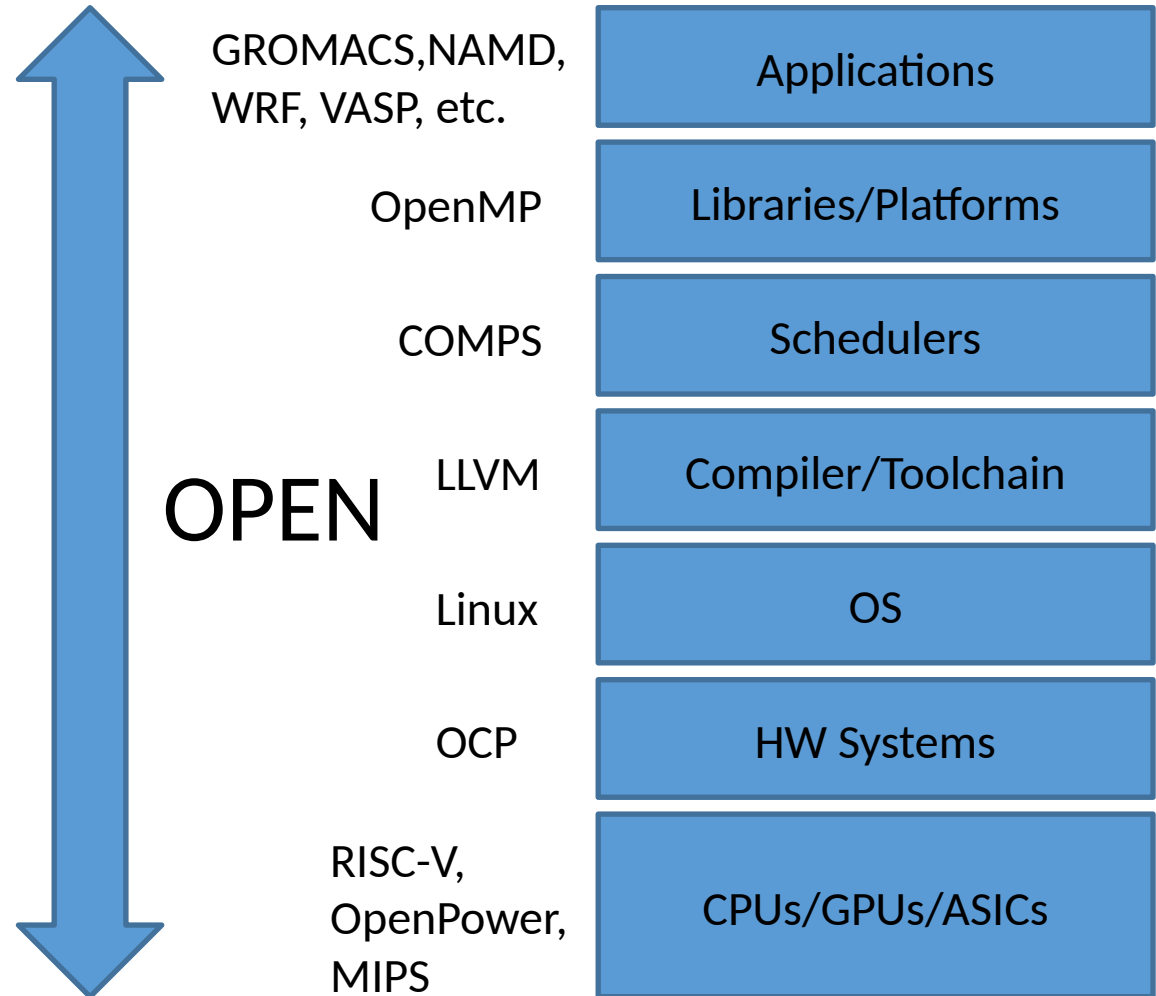
# HPC Today

- Europe has led the way in defining a common open HPC software ecosystem
- Linux is the de facto standard OS despite proprietary alternatives
- Software landscape from Cloud to IoT already enjoys the benefit of open source
- Open source provides:
  - A common platform, specification and interface
  - Accelerates building new functionality by leveraging existing components
  - Lowers the entry barrier for others to contribute new components
  - Crowd-sources solutions for small and larger problems
- **What about Hardware and in particular, the CPU?**



# HPC Tomorrow

- Europe can lead the way to a completely **open SW/HW stack for the world**
- RISC-V provides the open source hardware alternative to dominating proprietary non-EU solutions
- Europe can achieve complete technology independence with these foundational building blocks
- Currently at the same early stage in HW as we were with SW when Linux was adopted many years ago
- **RISC-V can unify, focus, and build a new microelectronics industry in Europe.**





# ***DRAC Overview***

# It all started in Mexico...

- CIC-IPN Lagarto I design in Mexico (2012-2017)
  - 5-stage single issue in-order pipeline, MIPS-based microcontroller
  - FPGA implementation capable of booting Linux
- BSC and CIC-IPN Lagarto Initiative (2018 onwards)
  - MIPS to RISC-V
  - FPGA to ASIC
- European Processor Initiative (EPI) (Dec 2018 - Dec 2023)
  - Flagship project (80M€ Phase 1; 70M€ Phase 2; 26 partners)
  - BSC leads the European RISC-V Vector Accelerator (EPAC)
- July 2018: Come to my office...
  - RIS3CAT “Emerging Technologies” call in Nov 2018



European Processor Initiative

- **DRAC: Designing RISCV-based Accelerators for next generation Computers**

- Consortium: BSC (coord.), UPC, UAB, UB, URV
- Dates: June 2019 – June 2023
- Budget: 4M€ (50% co-funded by Generalitat)
- Alignment with the European Processor Initiative (EPI) project:
  - Focus on RISC-V-based accelerator developed in Barcelona
  - Promote RISC-V in the CS degrees in Catalan universities
  - Build IC design teams capable of taping out DRAC technology: RTL design, verification and physical design



- 1 Design of an out-of-order general purpose RISC-V **processor**
- 2 Design of **accelerators** and required hardware support to have secure processors that incorporate post-quantum cryptographic schemes and virtualization techniques
- 3 Design of accelerators for genomics data analytics
- 4 Design of efficient and low power processors for autonomous navigation applications
- 5 Building a Catalan ecosystem for custom hardware design and fabrication
- 6 Transfer DRAC technology to local and international companies
- 7 Transfer DRAC technology to the Catalan university system using educational kits based on DRAC designs


2019

2020


2021

2022



2023




**Lagarto (May'19)**  
 - Lagarto Hun 5-stage in-order  
 - 150MHz (external)  
 - TSMC 65nm  
 - 2.5mm<sup>2</sup>



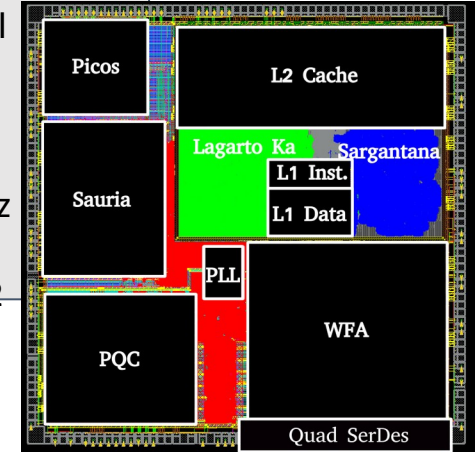
**DVINO (Apr'21)**  
 - Lagarto Hun in-order  
 - VPU  
 - PLL 600 MHz  
 - SDRAM mem cont  
 - HyperRAM  
 - VGA  
 - ADC  
 - TSMC 65nm  
 - 8mm<sup>2</sup>



**Sargantana (Feb'22)**  
 - Sargantana 7-stage in-order  
 - PLL 1.2 GHz  
 - Custom extensions  
 - SDRAM  
 - Prototype analog IPs: SerDes 8GHz  
 - GF 22nm  
 - Area: 2.9mm<sup>2</sup>



**Kameleon (Dec'22)**  
 - Lagarto Ka 11-stage ooo  
 - PLL 1.2 GHz  
 - Automotive Accel  
 - Crypto Accel  
 - Genomic Accel  
 - PICOS Accel  
 - SerDes 8GHz  
 - GF 22nm  
 - Area: 9mm<sup>2</sup>

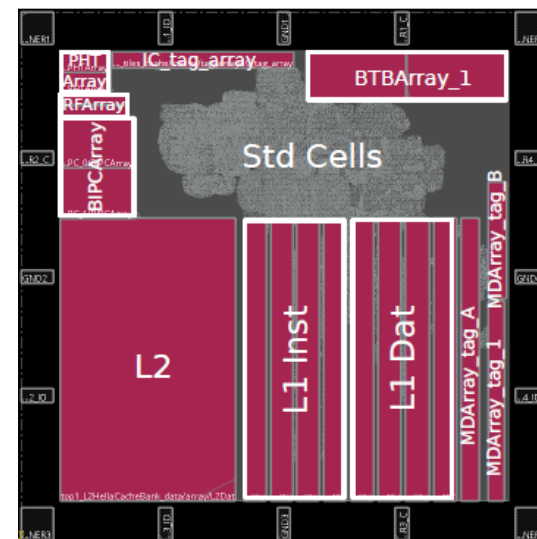



- Target design:



- Lagarto Hun in-order scalar core, 5 stages, single issue, RV64IMA
- 16KB L1 caches, 64KB L2 cache, TLB
- Memory controller on the FPGA side via packetizer
- Debug ring via JTAG
- Target technology: TSMC 65nm, area fits in 2.5mm<sup>2</sup>

- Fabrication and bringup


- Submitted in May 2019
- Samples received in Sep 2019
- Bringup with custom PCB in Oct 2019
- Linux boot in Dec 2019


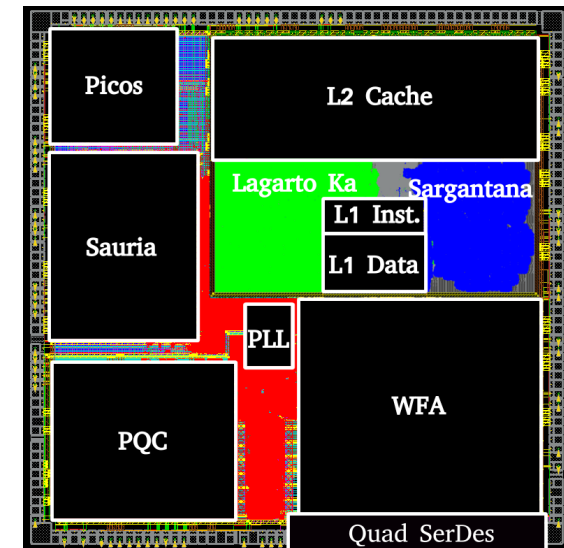
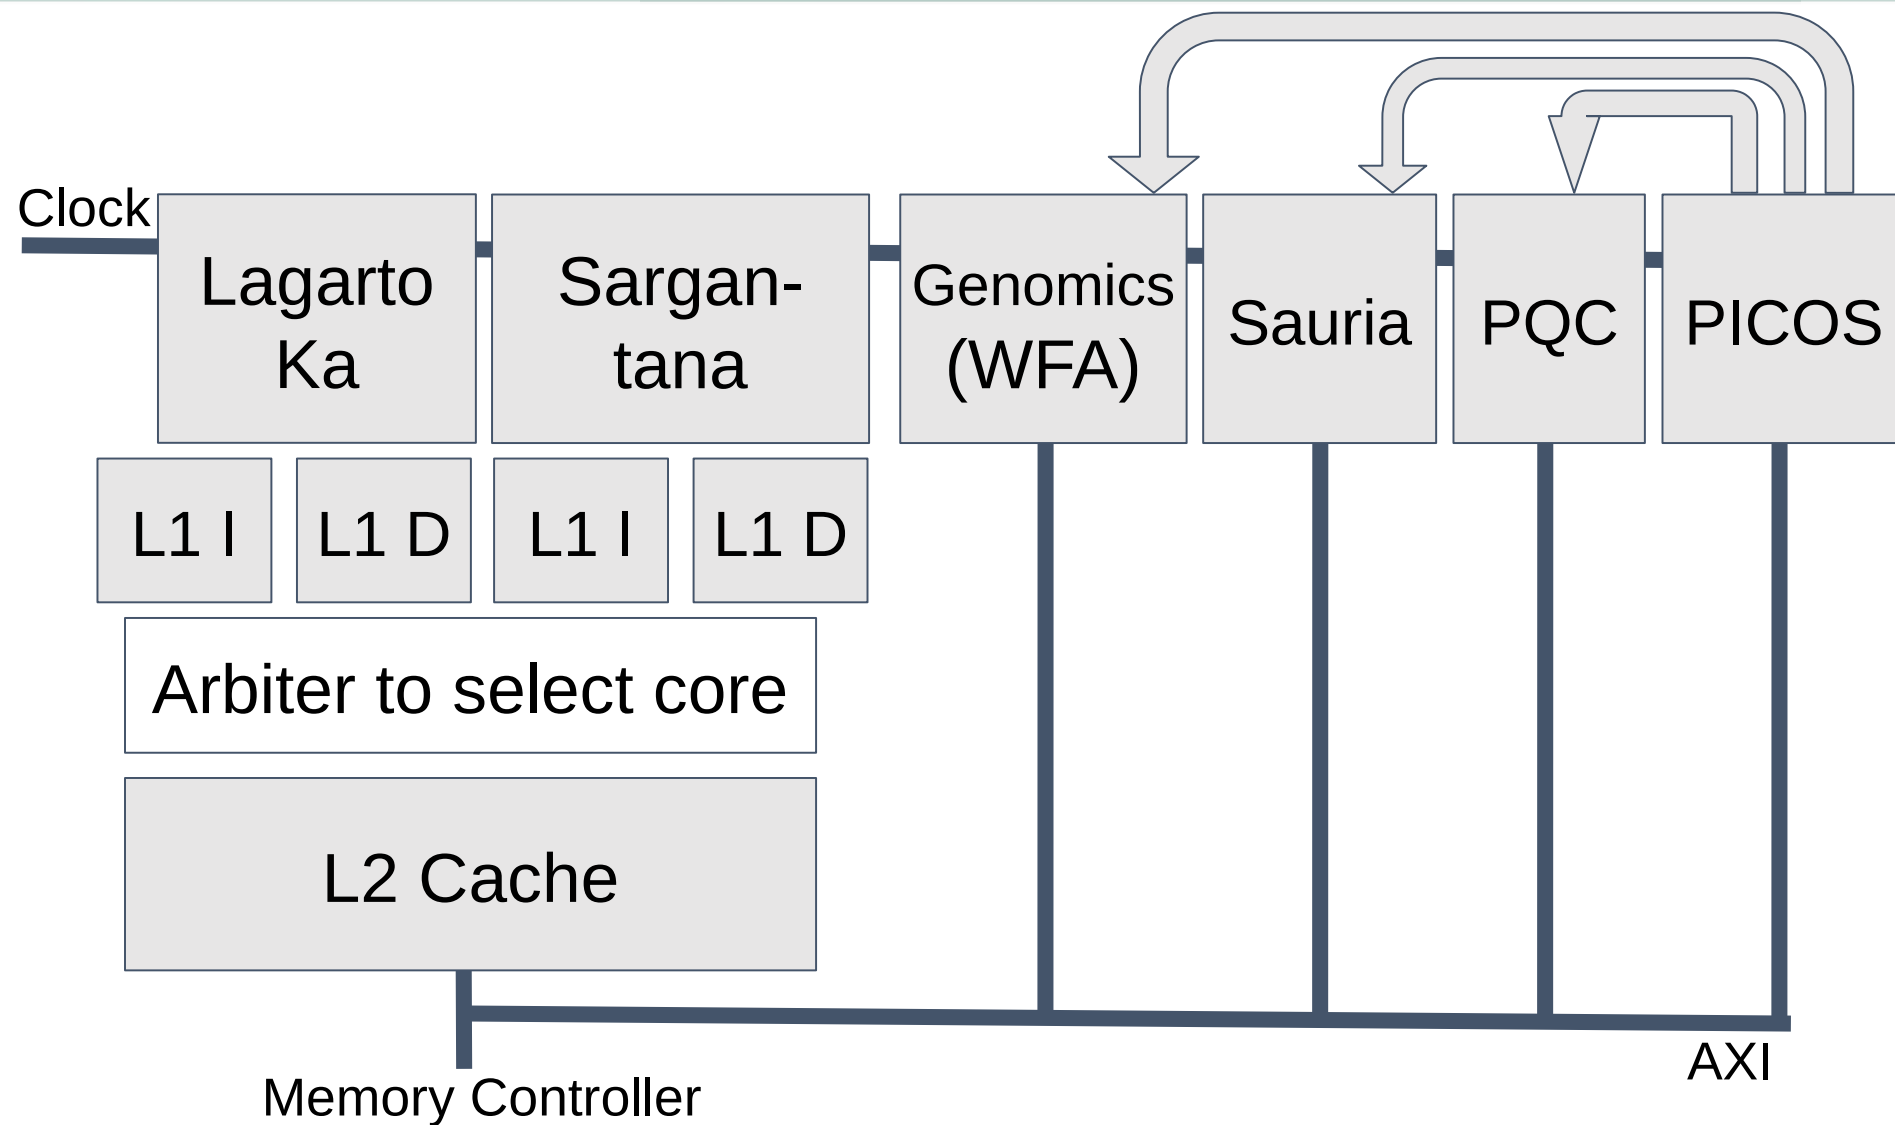


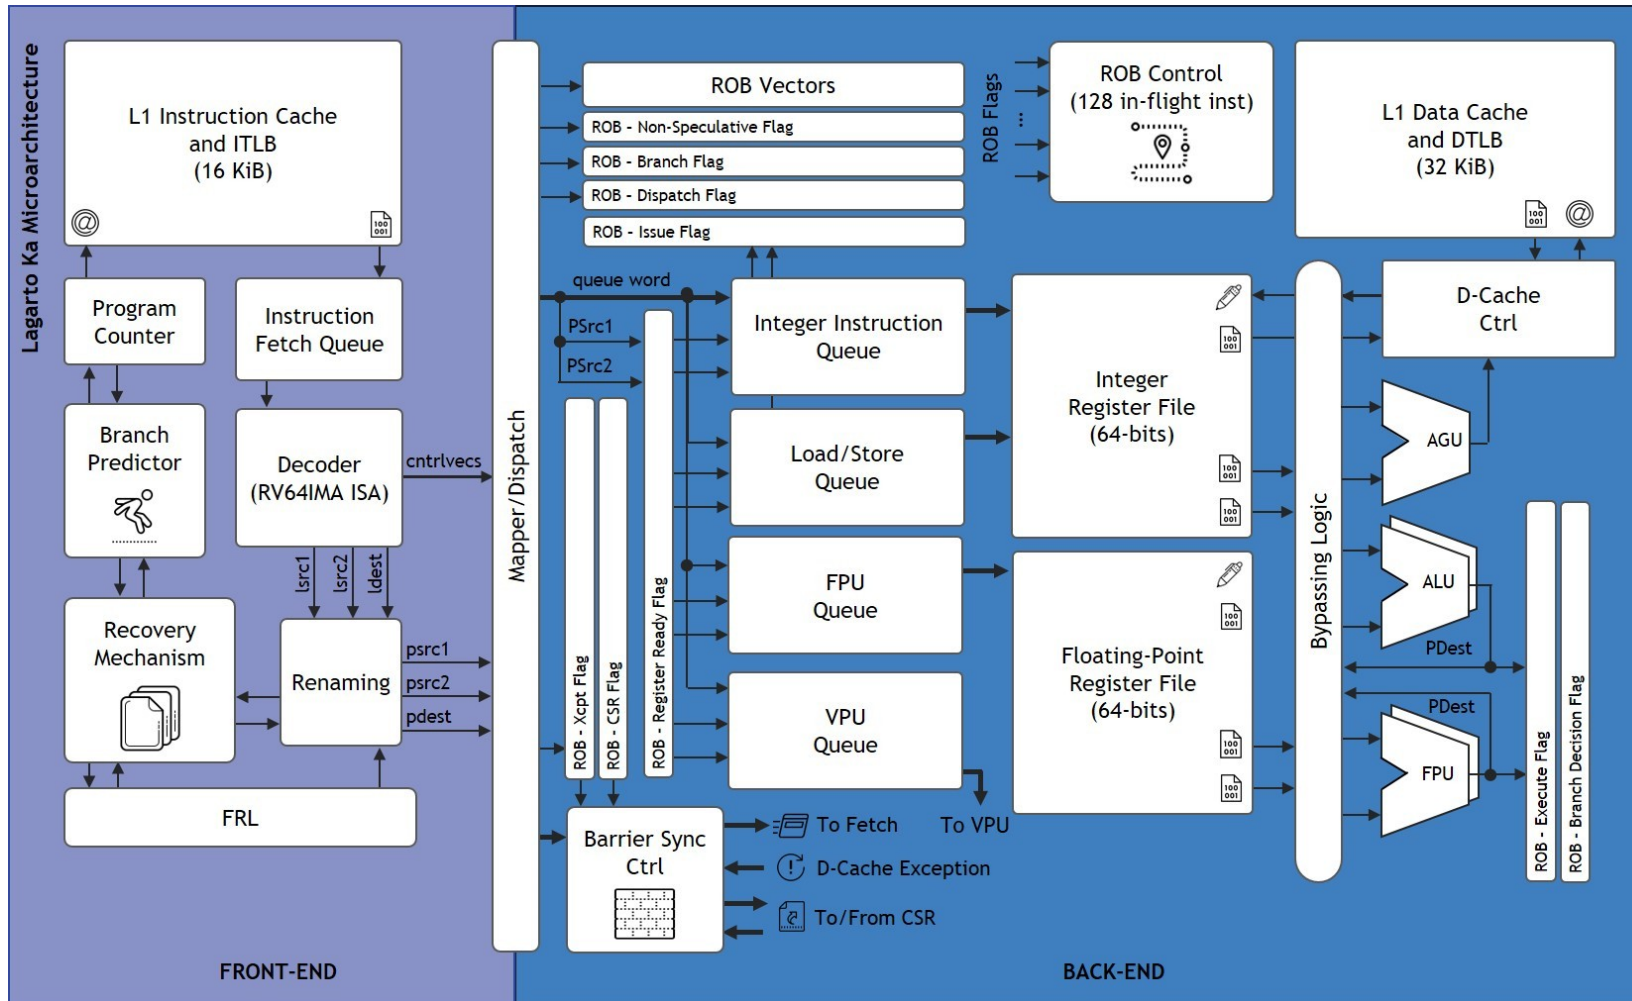
- DRAC Vector IN-Order (DVINO) processor details:
  - Lagarto Hun scalar pipeline, 5-stage, in-order, RV64IMA
  - Hydra 2-lane (VPU-DRAC-1.0), 4096-bit vector length
  - Internal PLL. DVINO can run at 600, 400, 300 and 200MHz
  - In-house L1 instruction cache and PMU
  - L1 data and L2 caches from lowRISC 0.2
  - Multiple contr: JTAG, UART, SPI, VGA, SDRAM and Hyperram.
  - In-house JTAG-based debug-ring
  - Technology node: TSMC 65nm (Europractice)
  - Area: 8.6mm<sup>2</sup>



- Sargantana in-order processor details:
  - Lagarto Hun pipeline, 7-stage, in-order, RV64IMAFD (RV64G)
  - Support for floating point operations (single and double precision)
  - Integer SIMD VPU, 128-bit vector length, custom instructions
  - Internal PLL. Sargantana can run above 1.1GHz
  - In-house L1 instruction cache and PMU
  - L1 data and L2 caches from lowRISC 0.2
  - Multiple controllers: JTAG, UART, SPI, SerDes, and Hyperram
  - In-house JTAG-based debug-ring
  - Technology node: GF 22nm (Europractice)
  - Area: 2.9mm<sup>2</sup>








## Current Features

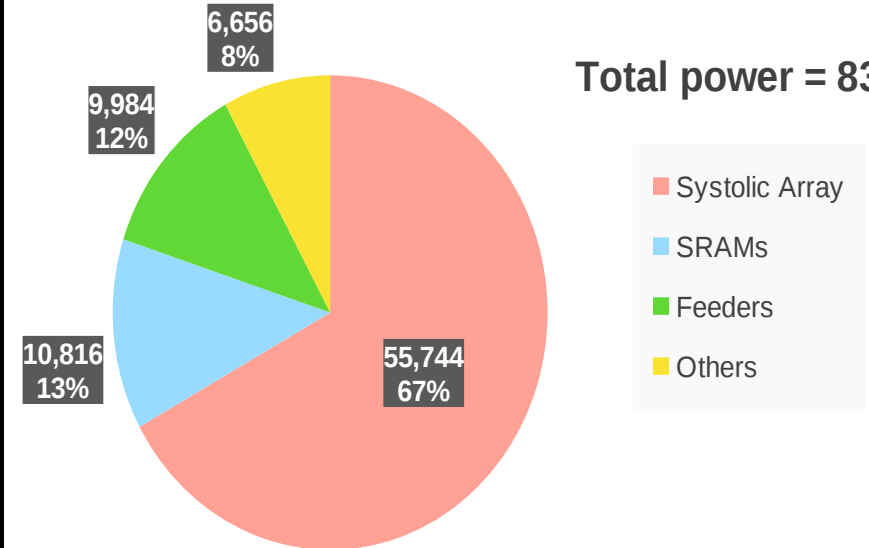

- 2-way 64-bit out-of-order architecture
- RV64IMA ISA
- 11-stage pipeline implementation
- Parameterized branch predictor:
  - › BTB 16-128 entries
  - › BHT 16-128 entries
  - › RAS 2-8 entries
- ROB 128-entries
- Low-power Integer queue (out-of-order issue)
- In-order Load/Store Queue
- Hit-under-miss support
- Configurable, L1 caches
  - › 16 KiB L1 I-cache (Typical)
  - › 32 KiB L1 D-cache (Typical)

# SAURIA Physical Design



## Specs:

8x16 Array: 128 Processing Elements (PE)  
Approximate logic in PE multipliers & adders  
1.00 x 0.95 mm = 0.95 mm<sup>2</sup>  
128 GFLOP/s @ 500MHz

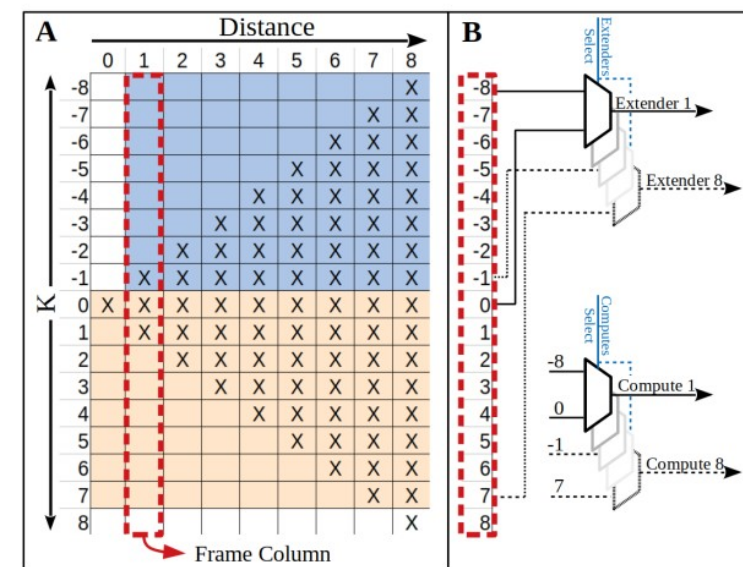
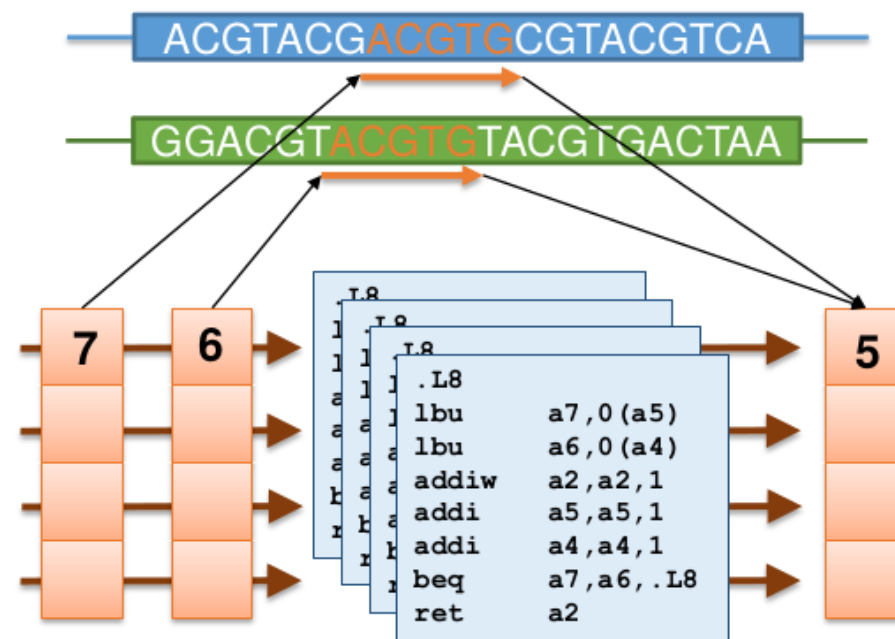


Total power = 83.2 mW

**1.56 TFLOP/sW**

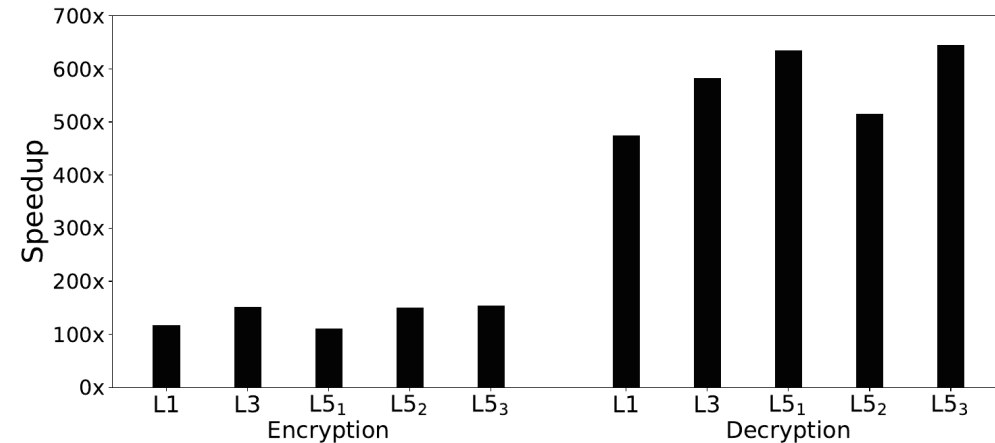
# Genome Alignment Acceleration (WFA)

- Pairwise alignment/mapping DNA/RNA sequences with the novel Wavefront Alignment (WFA) algorithm
- Target: provide specific hardware support for the most time-consuming operations of the algorithm.
- **ISA extensions:**
  - [vmax\_vv] Vectorial maximum.
  - [vmax3inc\_vv] Vectorial “3-way” maximum fused with increment.
  - [vcnt] Scalar “count consecutive matches”
  - [vcnt\_vv] Vectorial “count consecutive matches”.
- Use narrow-integers 16-bit or 8-bit integers
- **Monolithic accelerator** integrated with Lagarto SoC
  - Single aligner (due to area limitations) capable of 64 ops/cycle
  - Current Place and Route (PnR) results: 1.1GHz (typical corner)
  - Area: 1.6mm<sup>2</sup> in Global Foundries 22nm
  - Performance speedups: 515x with 10K reads, 10% error



# PQC Acceleration

- Main Goal: **RISC-V acceleration** of different PQC schemes
- **Classic McEliece (CME) KEM acceleration**:
  - HW/SW co-design implementation of CME KEM @ Zynq Ultrascale [FPL'21]
  - Monolithic CME accelerator based on HLS
  - Integration of the CME accelerator in Lagarto SoC via AXI interface
- Accelerate other KEM and digital signature (DS) schemes:
  - **NRTU KEM** and **Crystals-Kyber KEM / Crystals-Dilithium DS**
  - Both algorithms rely on very different operations and will require different acceleration techniques.






# Building a RISC-V Ecosystem in Barcelona

*Designing RISC-V-based Accelerators for next generation Computers*

DRAC KoM (Feb 2020)



Lagarto Team (Sept 2019)

DRAC F2F (Jul 2022)



DRAC Final Workshop (December 2022)

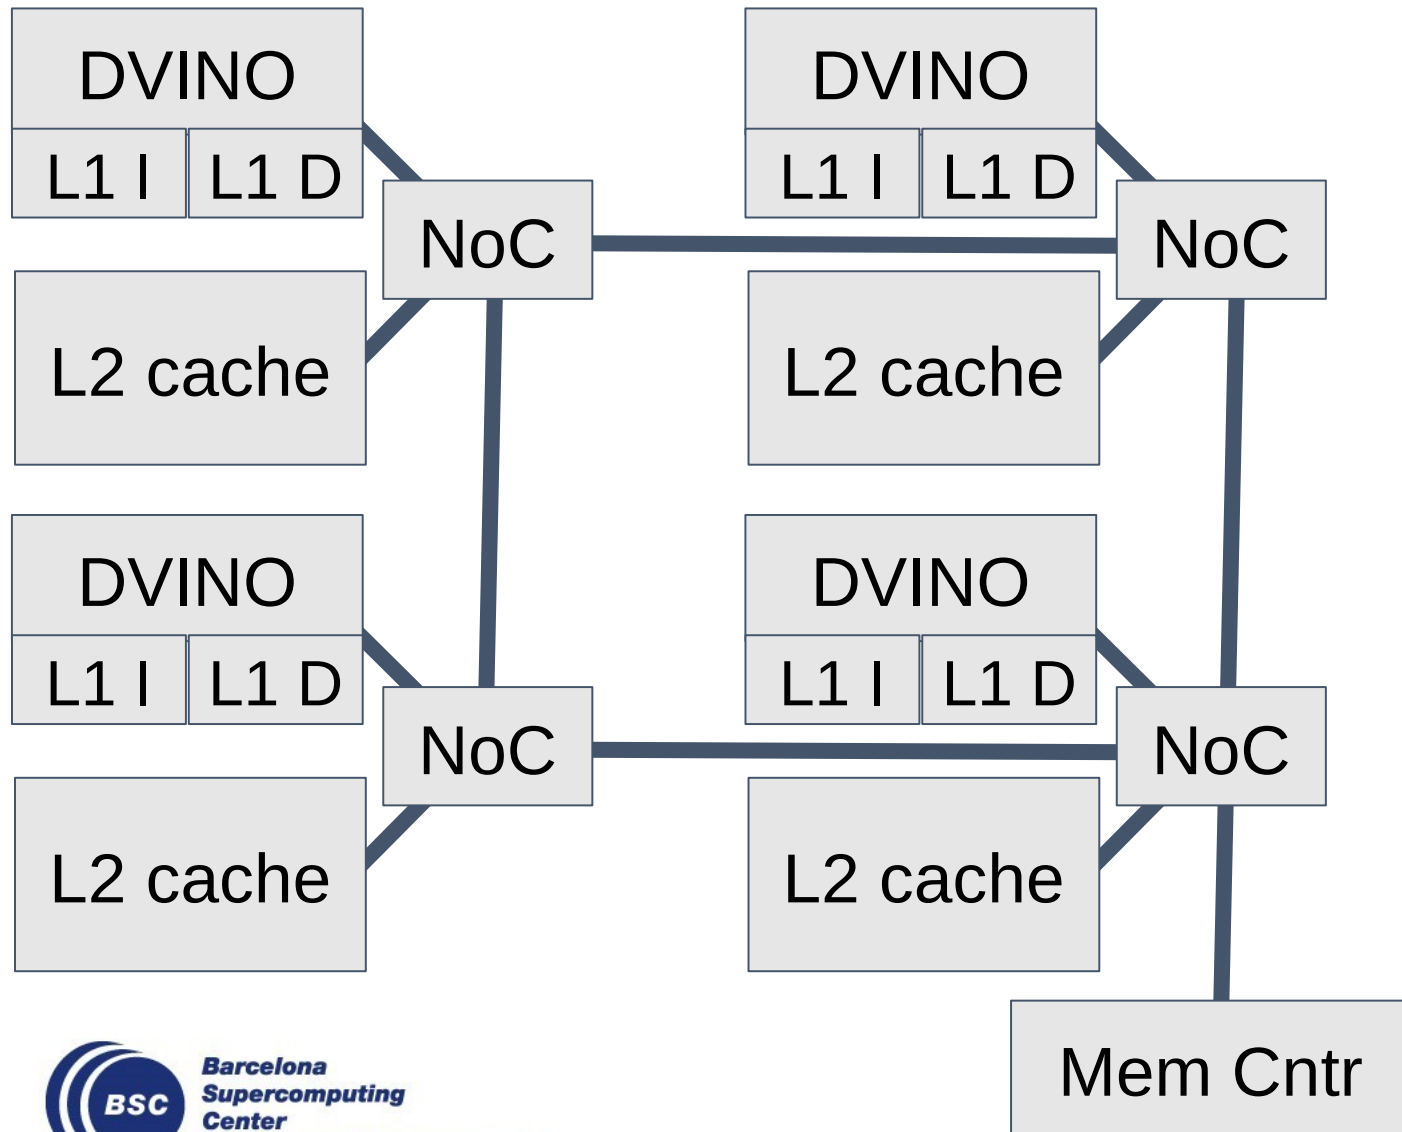
- When DRAC started in 2019: minimal presence of RISC-V in the Catalan ecosystem
  - Few courses in the MIRI Master at UPC
- Current situation:
  - UAB redesigned the CS degree to incorporate RISC-V as the main ISA
  - UPC incorporated RISC-V in the computer engineering specialization of the CS degree and several masters (MEE, MEI, MIRI)
  - UB also introduced RISC-V in CS and EE degrees
  - URV will incorporate RISC-V in the main computer architecture courses in 2023
  - Many other universities are working towards the same direction in the context of the Spanish Red RISC-V network (UCM, UMU, USC, UIB, etc.)





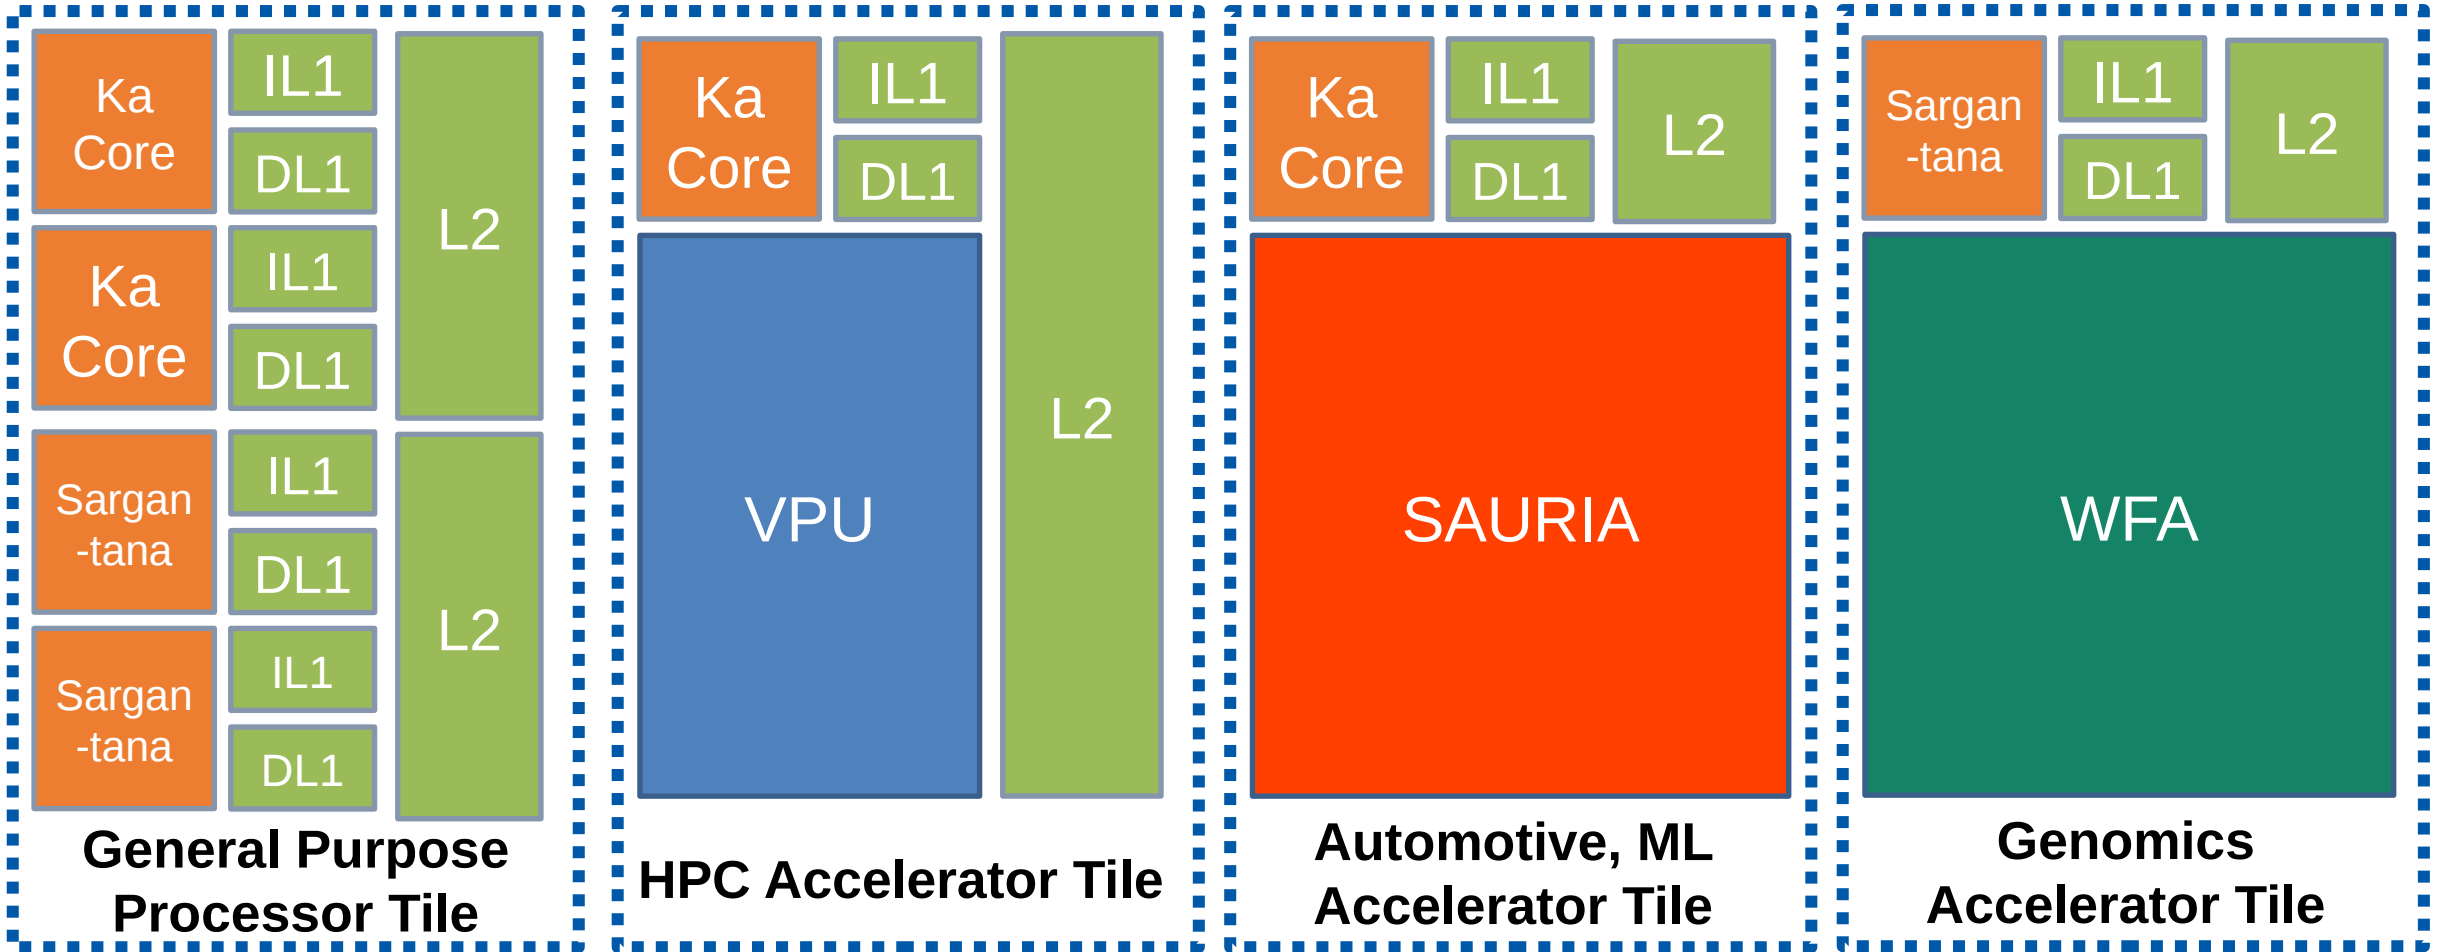
# *Lagarto Roadmap*

# First Steps Towards a Lagarto Multicore



- Multicore design based on:
  - DVINO processor
  - RISC-V ISA support: I, M, A, F, D, C, V
  - 4-lane VPU
  - OpenPiton 2-level cache hierarchy (priv. L1, shared L2)
- Current status
  - Linux boot (openSBI)
  - RTL simulation of parallel applications (up to 64 cores)
  - Multiple memory controllers
  - FPGA-ready

# Towards a RISC-V Heterogeneous Manycore



- Open source hardware design opportunities and challenges!!
  - Many open source RTL designs available, including cores, SoCs and accelerators
  - Design toolflow partially open (simulators, testing, FPGA emulation), but still some pieces are missing (verification, Place&Route, bringup)
  - Technology-related IP is completely closed
  - SW ecosystem still under development
  - Ideal for teaching, research and startups!!
  - Potential to be an alternative to dominating proprietary non-EU solutions
- DRAC Project:
  - Contribute to open source community with in-house designs and tools
  - Design, verify and fabricate RISC-V-based processors and accelerators
  - Contribute to European RISC-V ecosystem and projects



*Designing RISC-V-based Accelerators for next generation Computers*

# Thank you!



The DRAC project with -file number 001-P-001723- has been 50% co-financed with € 2,000,000.00 by the European Union Regional Development Fund within the framework of the ERDF Operational Program of Catalonia 2014-2020, with the support of Generalitat of Catalonia. Copyright 2020 © All Rights Reserved.