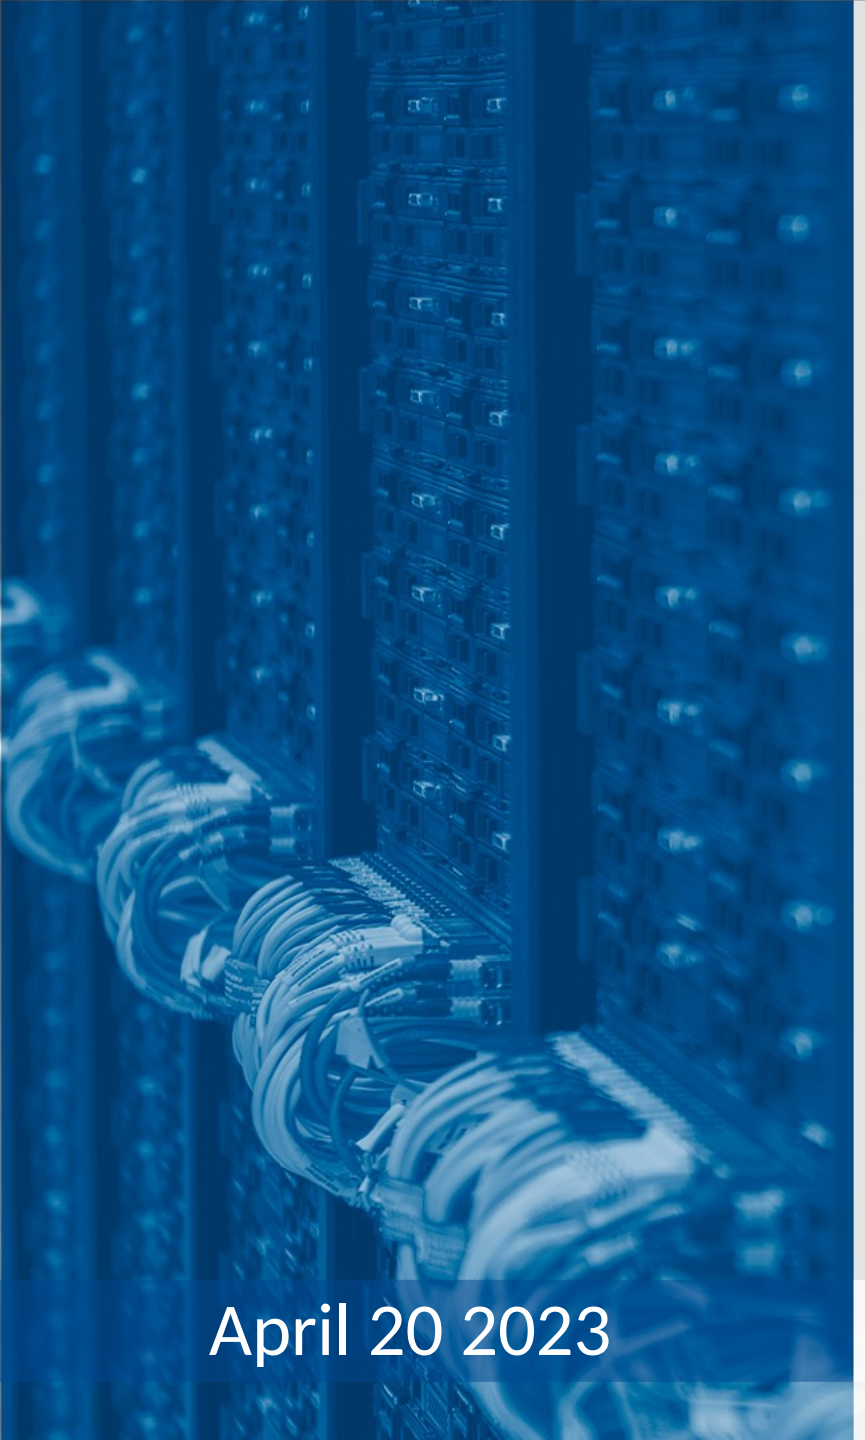

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

DRAC: Designing RISC-V-based Accelerators for Next Generation Computers

Miquel Moretó

UPC and BSC

April 20 2023

Universidad Complutense de Madrid (UCM), Madrid (Spain)

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

BSC-CNS objectives

Supercomputing services
to Spanish and EU researchers

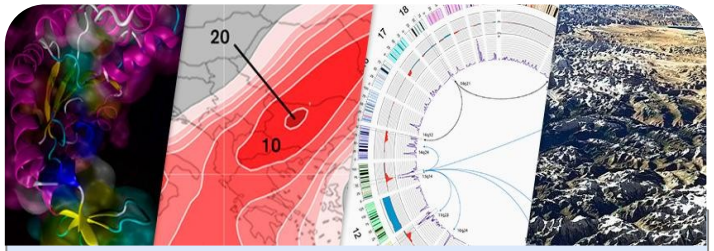

R&D in Computer, Life, Earth and
Engineering Sciences

PhD programme, technology
transfer, public engagement

BSC-CNS is a consortium that includes

Spanish Government	60%	
Catalan Government	30%	Generalitat de Catalunya Departament de Recerca i Universitats
Univ. Politècnica de Catalunya (UPC)	10%	UNIVERSITAT POLITÈCNICA DE CATALUNYA BARCELONATECH

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

MareNostrum 4

Total peak performance:	**13,9** Pflops
General Purpose Cluster:	11.15 Pflops
MN4 CTE-Power:	1.57 Pflops
MN4 CTE-ARM:	0.65 Pflops
MN4 CTE-AMD:	0.52 Pflops

PRACE

Access: prace-ri.eu/hpc_acces

RES
RED ESPAÑOLA DE SUPERCOMPUTACIÓN

Access: bsc.es/res-intranet

BSC Barcelona Supercomputing Center
Centro Nacional de Supercomputación

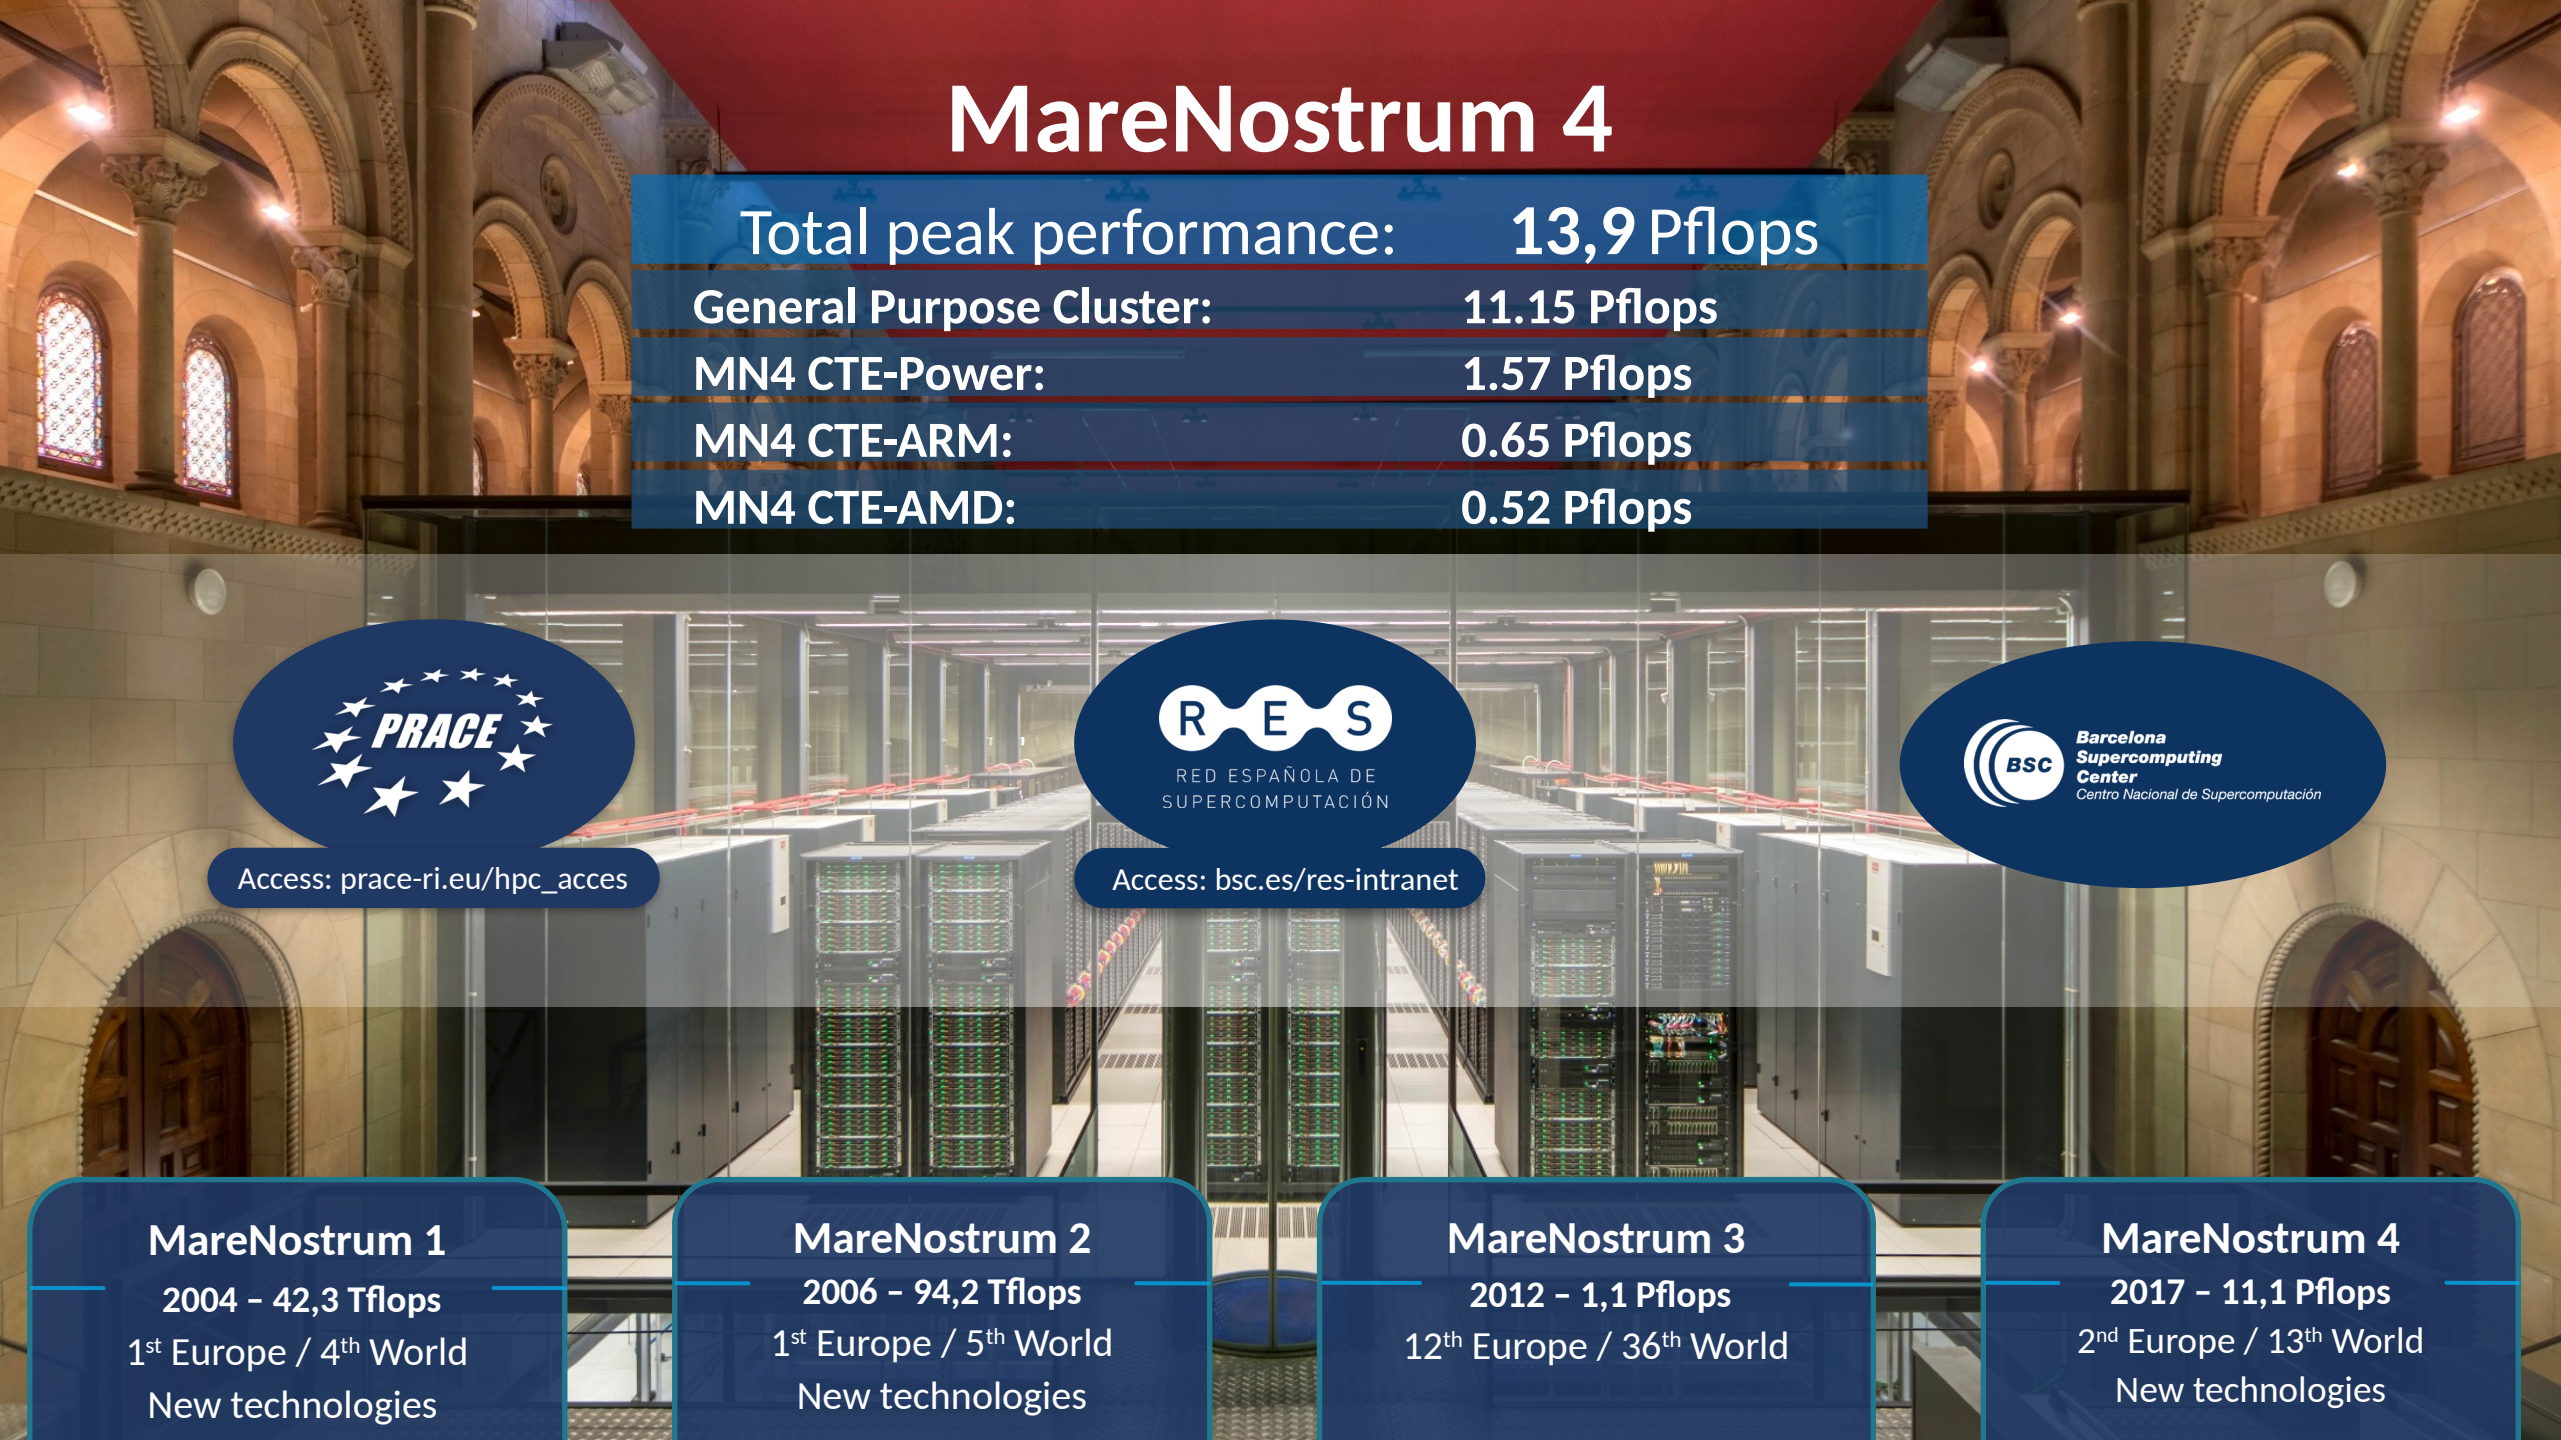

MareNostrum 1
2004 – 42,3 Tflops
1st Europe / 4th World
New technologies

MareNostrum 2
2006 – 94,2 Tflops
1st Europe / 5th World
New technologies

MareNostrum 3
2012 – 1,1 Pflops
12th Europe / 36th World

MareNostrum 4
2017 – 11,1 Pflops
2nd Europe / 13th World
New technologies

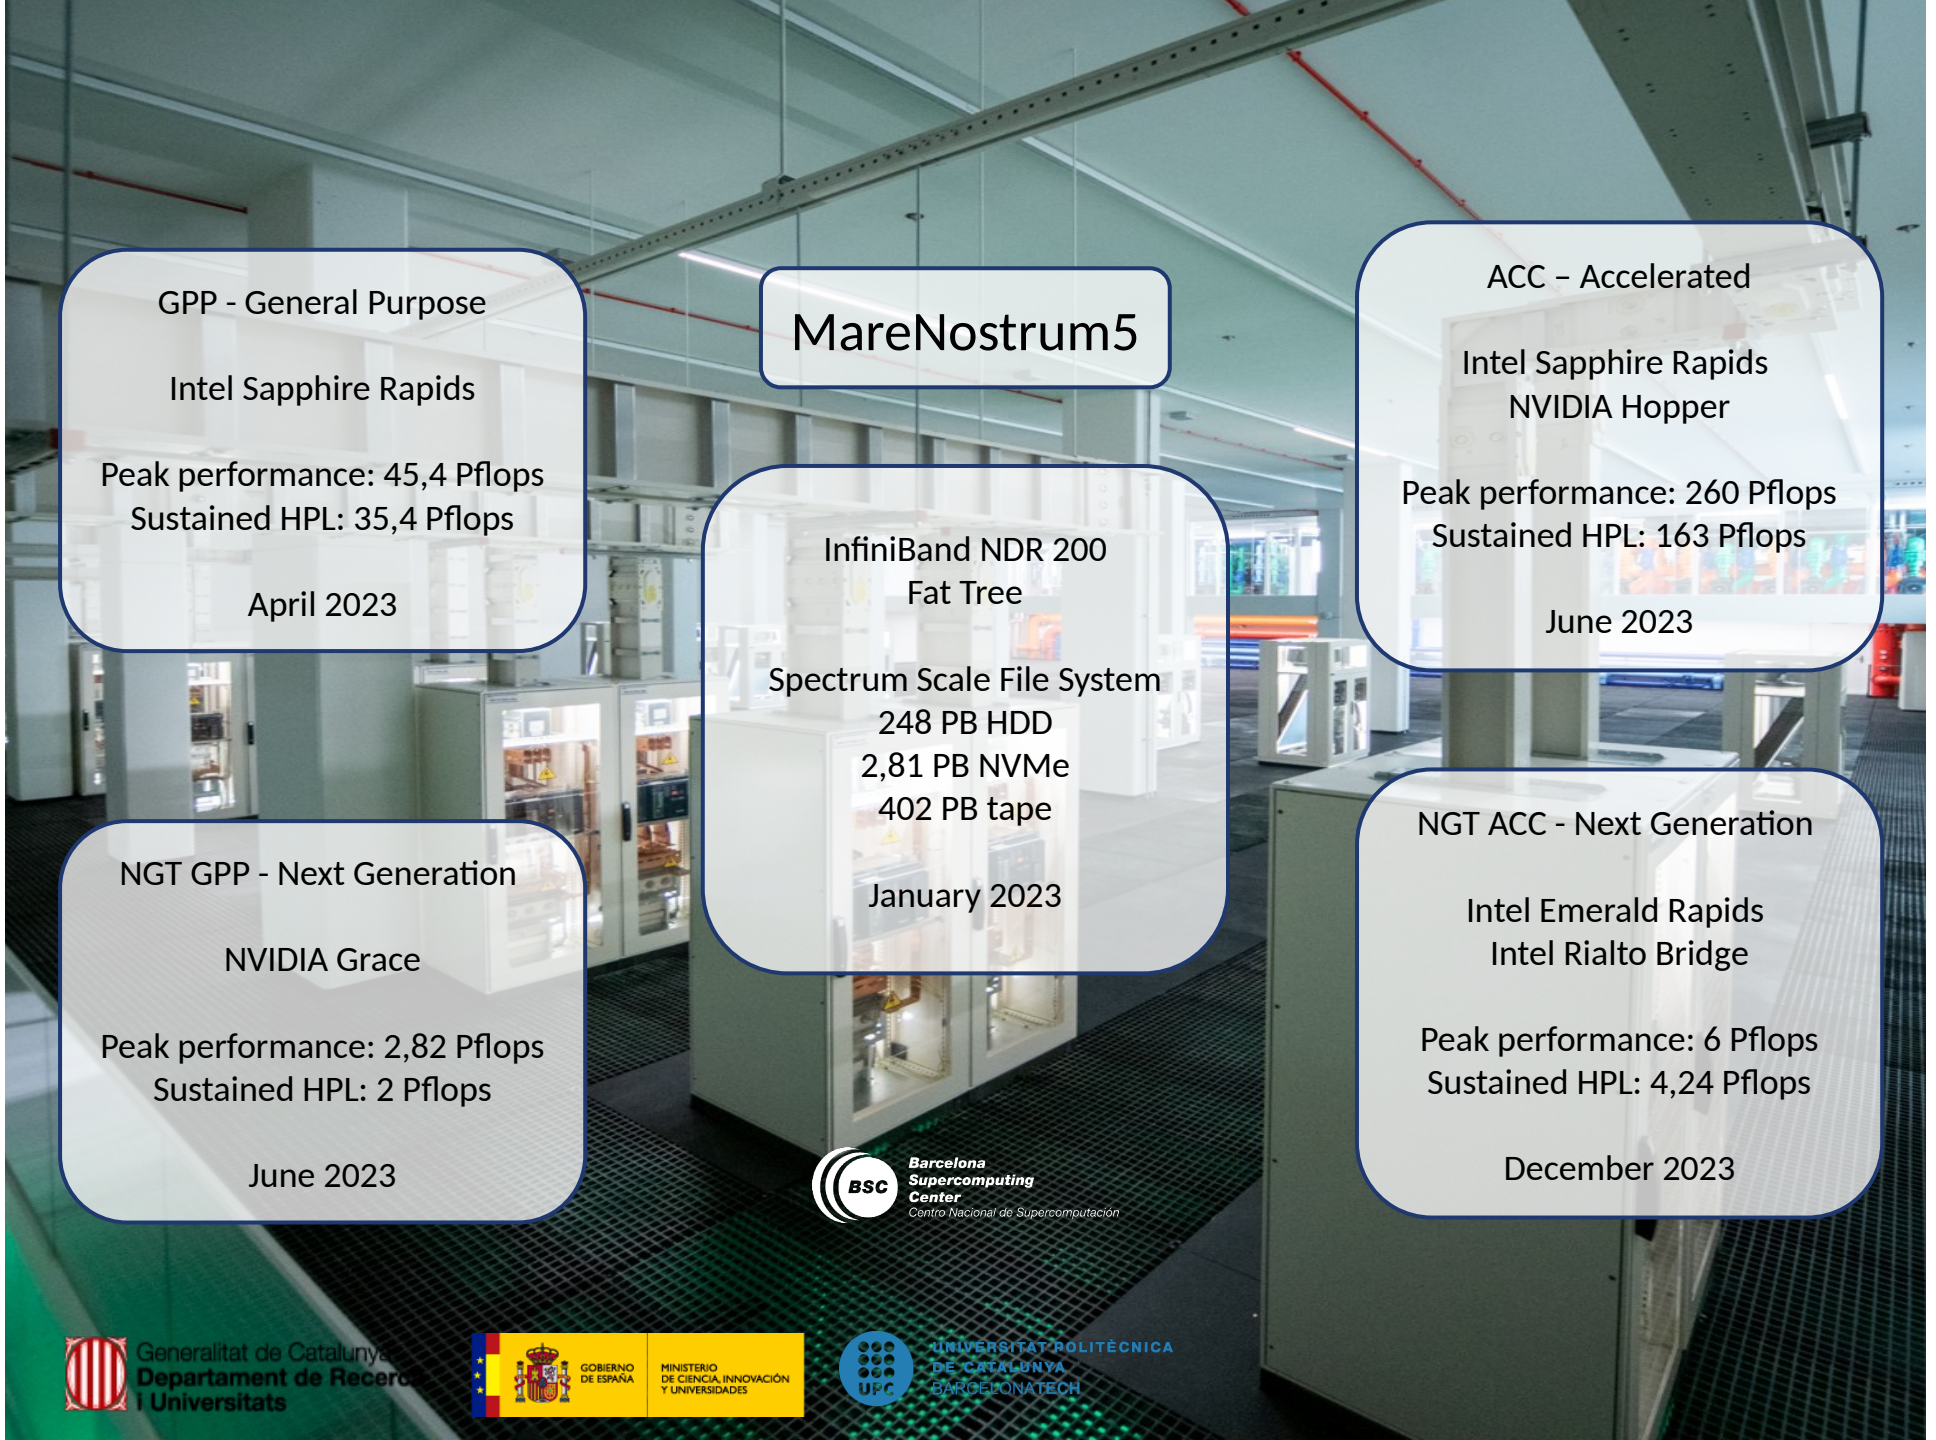
MareNostrum5
GPP - General Purpose
Intel Sapphire Rapids
Peak performance: 45,4 Pflops
Sustained HPL: 35,4 Pflops
April 2023
InfiniBand NDR 200
Fat Tree
Spectrum Scale File System
248 PB HDD
2,81 PB NVMe
402 PB tape
January 2023
ACC – Accelerated
Intel Sapphire Rapids
NVIDIA Hopper
Peak performance: 260 Pflops
Sustained HPL: 163 Pflops
June 2023
NGT GPP - Next Generation
NVIDIA Grace
Peak performance: 2,82 Pflops
Sustained HPL: 2 Pflops
June 2023
NGT ACC - Next Generation
Intel Emerald Rapids
Intel Rialto Bridge
Peak performance: 6 Pflops
Sustained HPL: 4,24 Pflops
December 2023
BSC
Barcelona Supercomputing Center
Centro Nacional de Supercomputación
Generalitat de Catalunya
Departament de Recerca i Universitats
GOBIERNO DE ESPAÑA
MINISTERIO DE CIENCIA, INNOVACIÓN Y UNIVERSIDADES
UPC
UNIVERSITAT POLITÈCNICA DE CATALUNYA
BARCELONATECH

BSC Staff Evolution

BSC Staff evolution 2005 - 2022, plus forecast 2023 & 2024

Data at 30th June 2022
(Including collaborators)

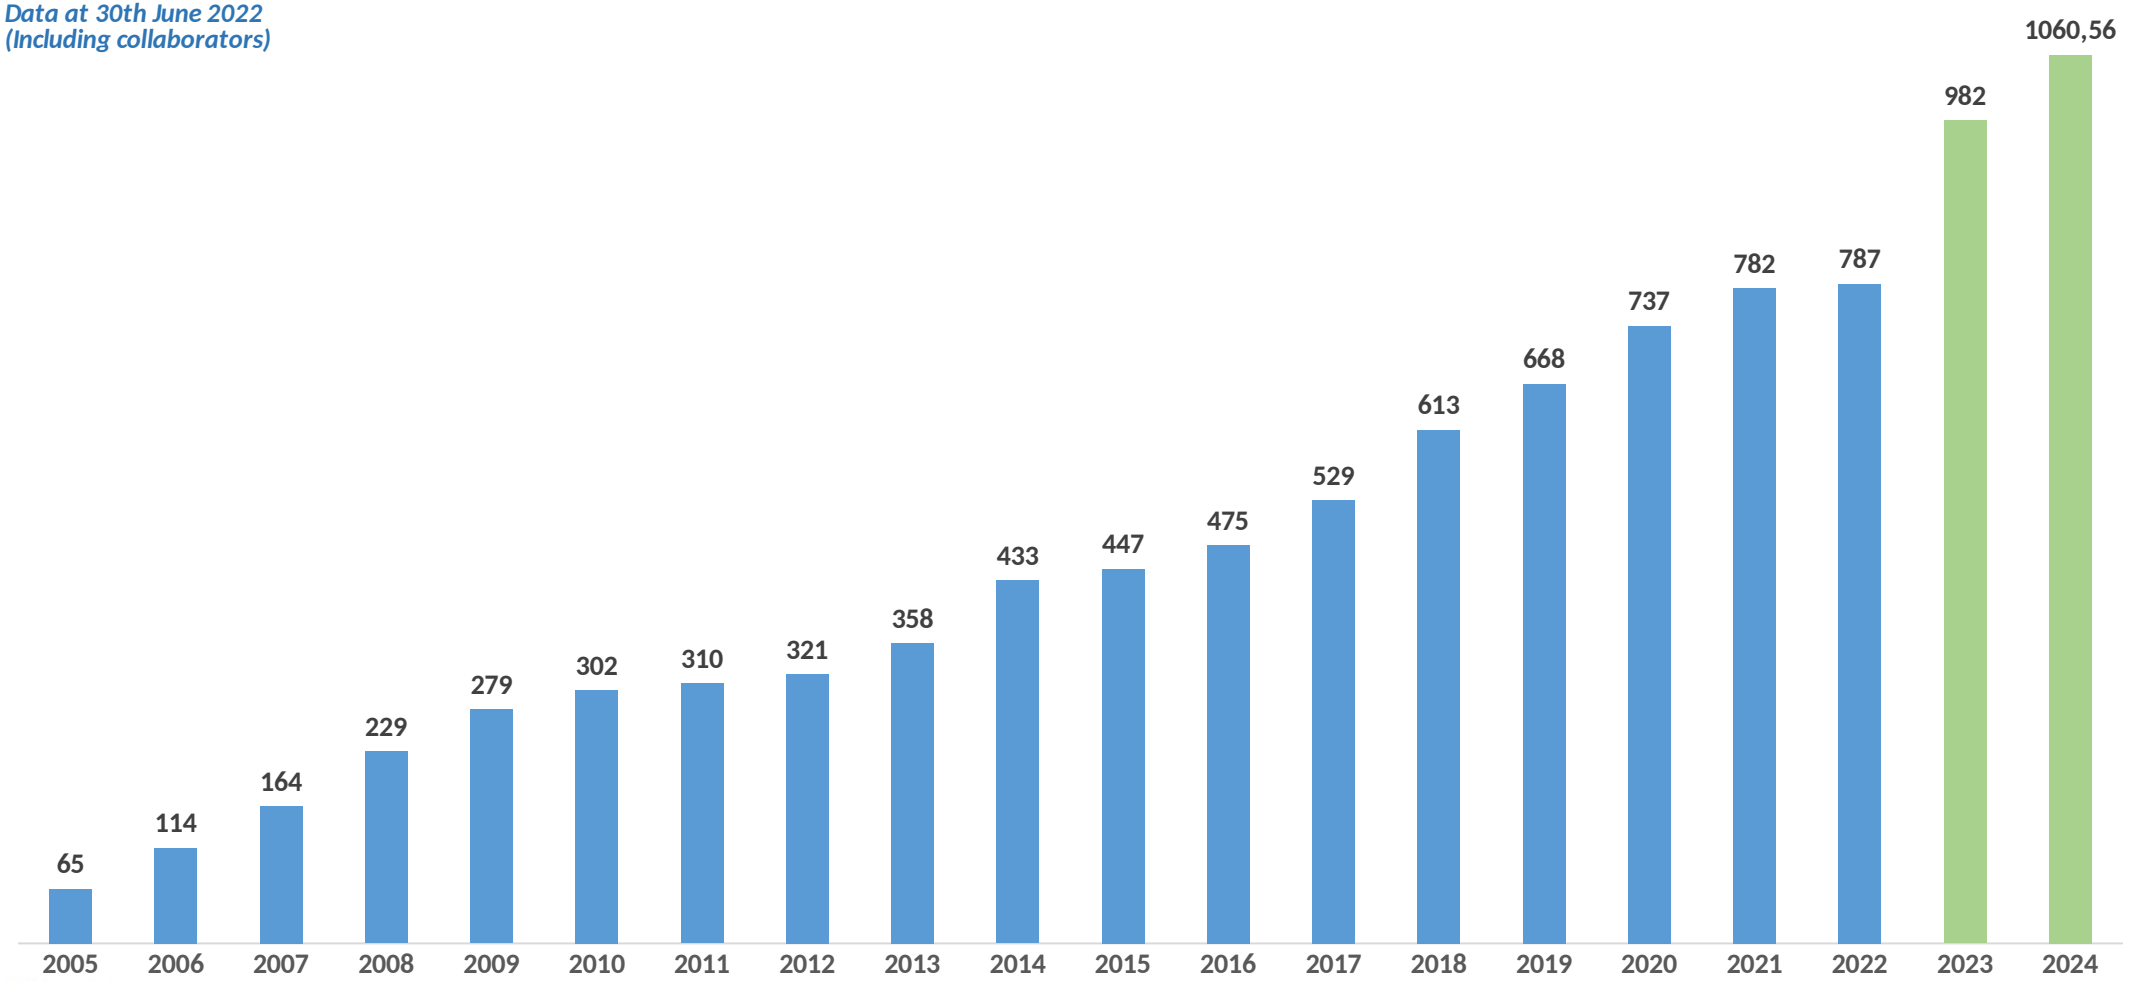

Barcelona Supercomputing Center
Centro Nacional de Supercomputación


Barcelona
Supercomputing
Center
Centro Nacional de Supercomputación

Collaborations with Global IT Industry

Mission of BSC Scientific Departments

To influence the way machines are built, programmed and used: programming models, performance tools, Big Data, Artificial Intelligence , computer architecture, energy efficiency

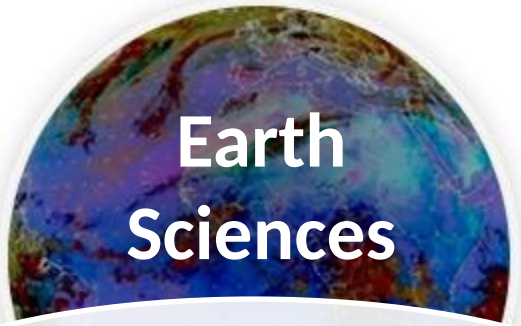

To develop and implement global and regional state-of-the-art models for short-term air quality forecast and long-term climate applications

To understand living organisms by means of theoretical and computational methods (molecular modeling, genomics, proteomics)

To develop scientific and engineering software to efficiently exploit super-computing capabilities (biomedical, geophysics, atmospheric, energy, social and economic simulations)

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

The RISC-V Revolution!

Today´s Technology Trends

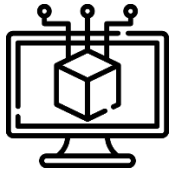

Massive penetration of Open Source Software

- 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

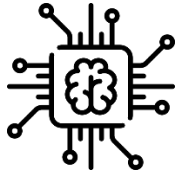

New Open Source Hardware Momentum from IoT and the Edge to HPC

- 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?**

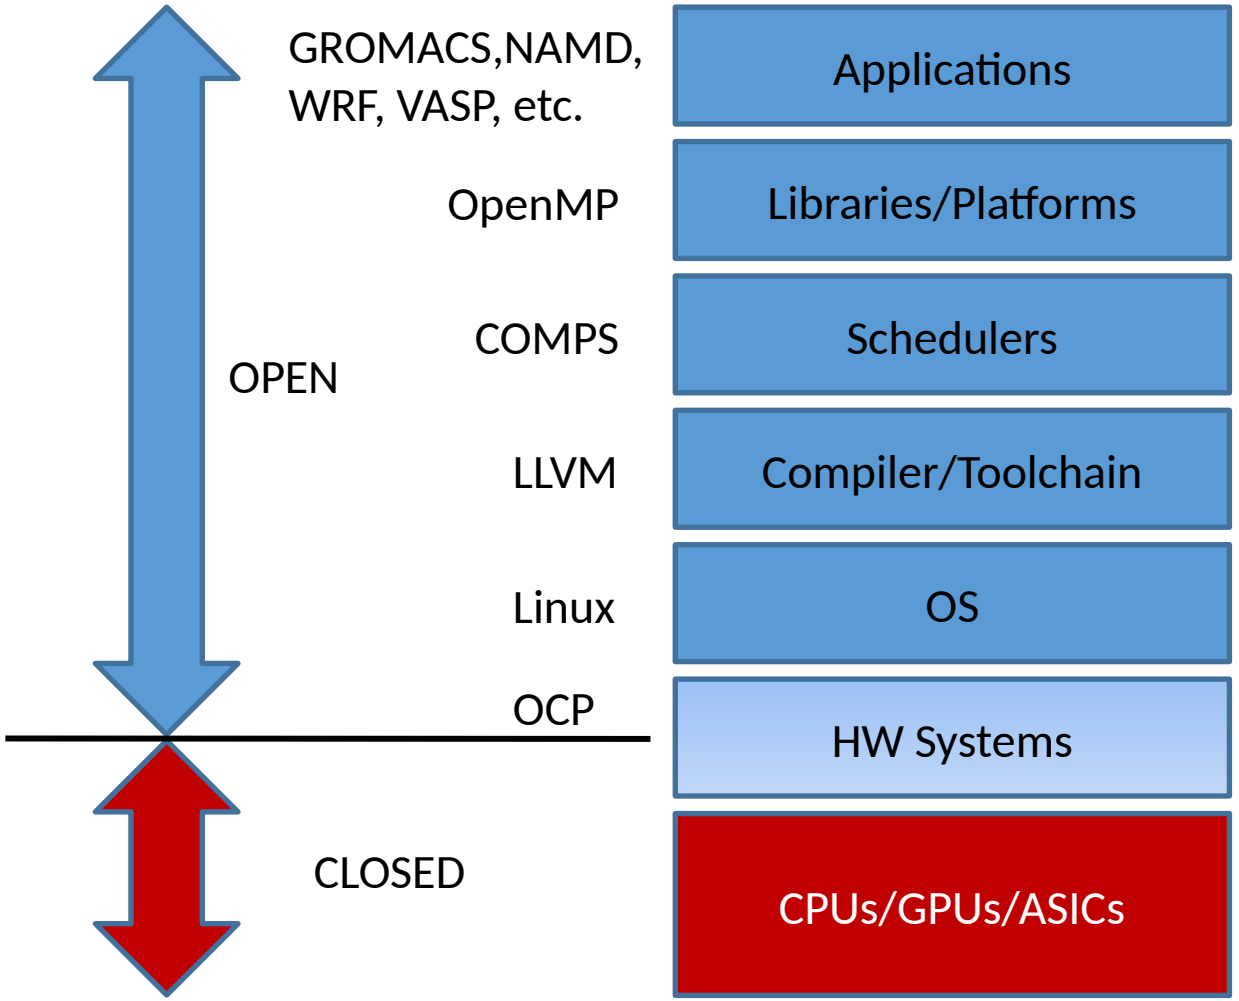

HPC Tomorrow

- 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.**

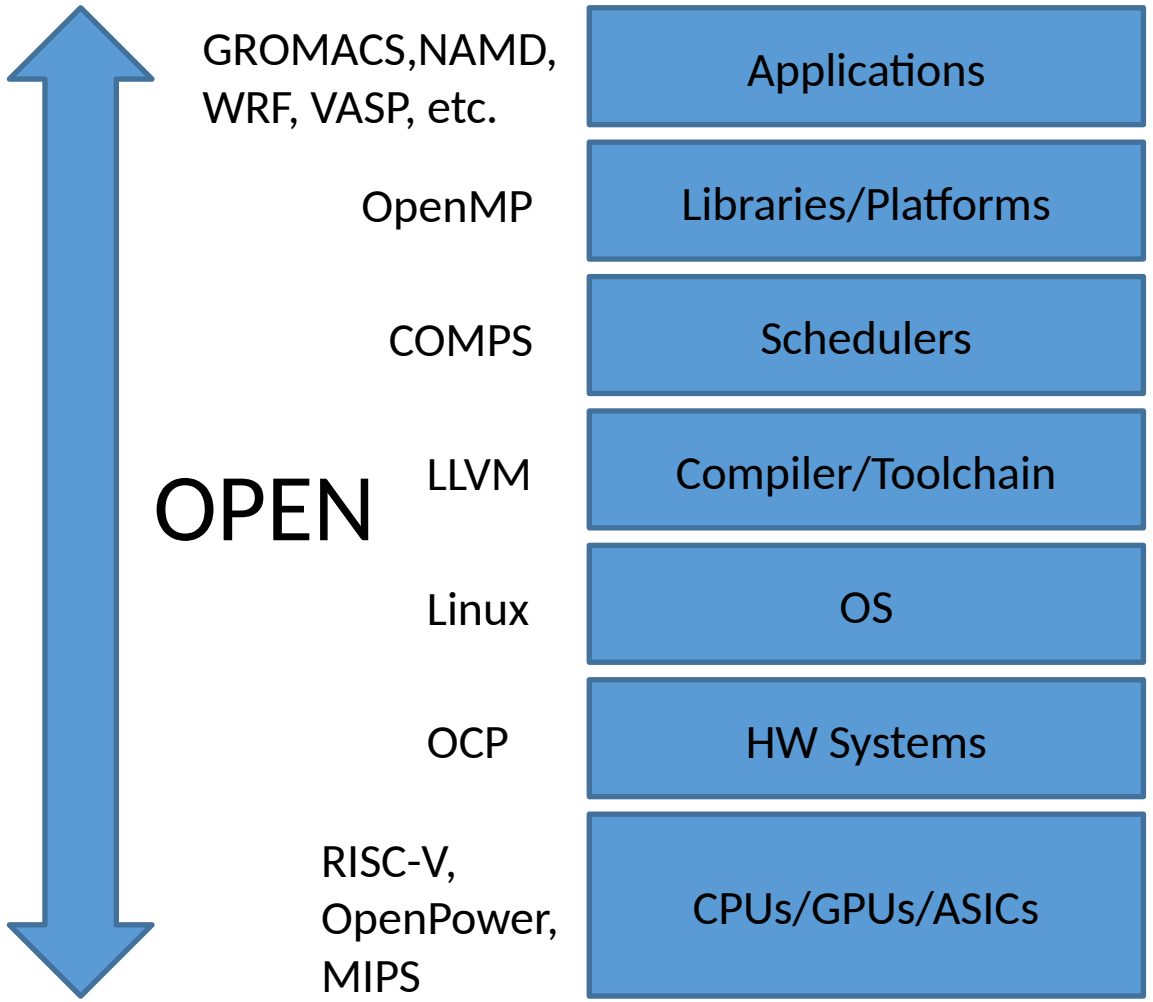

The European Processor Initiative (EPI)

The European Processor Initiative

The European Processor Initiative (EPI) under the SGA1 of the Framework Partnership Agreement (FPA: 800928), to design and implement a roadmap for a new family of low-power European processors for extreme scale computing, high-performance Big-Data and a range of emerging applications.

- **History:** Remember MontBlanc? BSC leads the RISC-V HPC accelerator development
- **Consortium (SGA1):**

28 partners from 10 European countries to Coordinate: Bull SAS (France)

- **Budget:** €80M (100% funded)
- **Duration:** 36 months (01/12/2018-31/12/2021)
- **5 Streams** (4 Technical and 1 Management/Exploitation/C&D)

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Key contact point
Jesús Labarta jesus.labarta@bsc.es

EPI MAIN OBJECTIVE

To develop European microprocessor and accelerator technology

- Strengthen competitiveness of EU industry and science

SiPearl

Rhea

Arm-based general purpose CPU

EPAC

RISC-V based Accelerators

BSC, SemiDynamics, EXTOLL, FORTH, ETHZ, UniBo, UniZG, Chalmers, CEA, E4, Menta, ZPT, …

OVERALL TECHNOLOGY ROADMAP

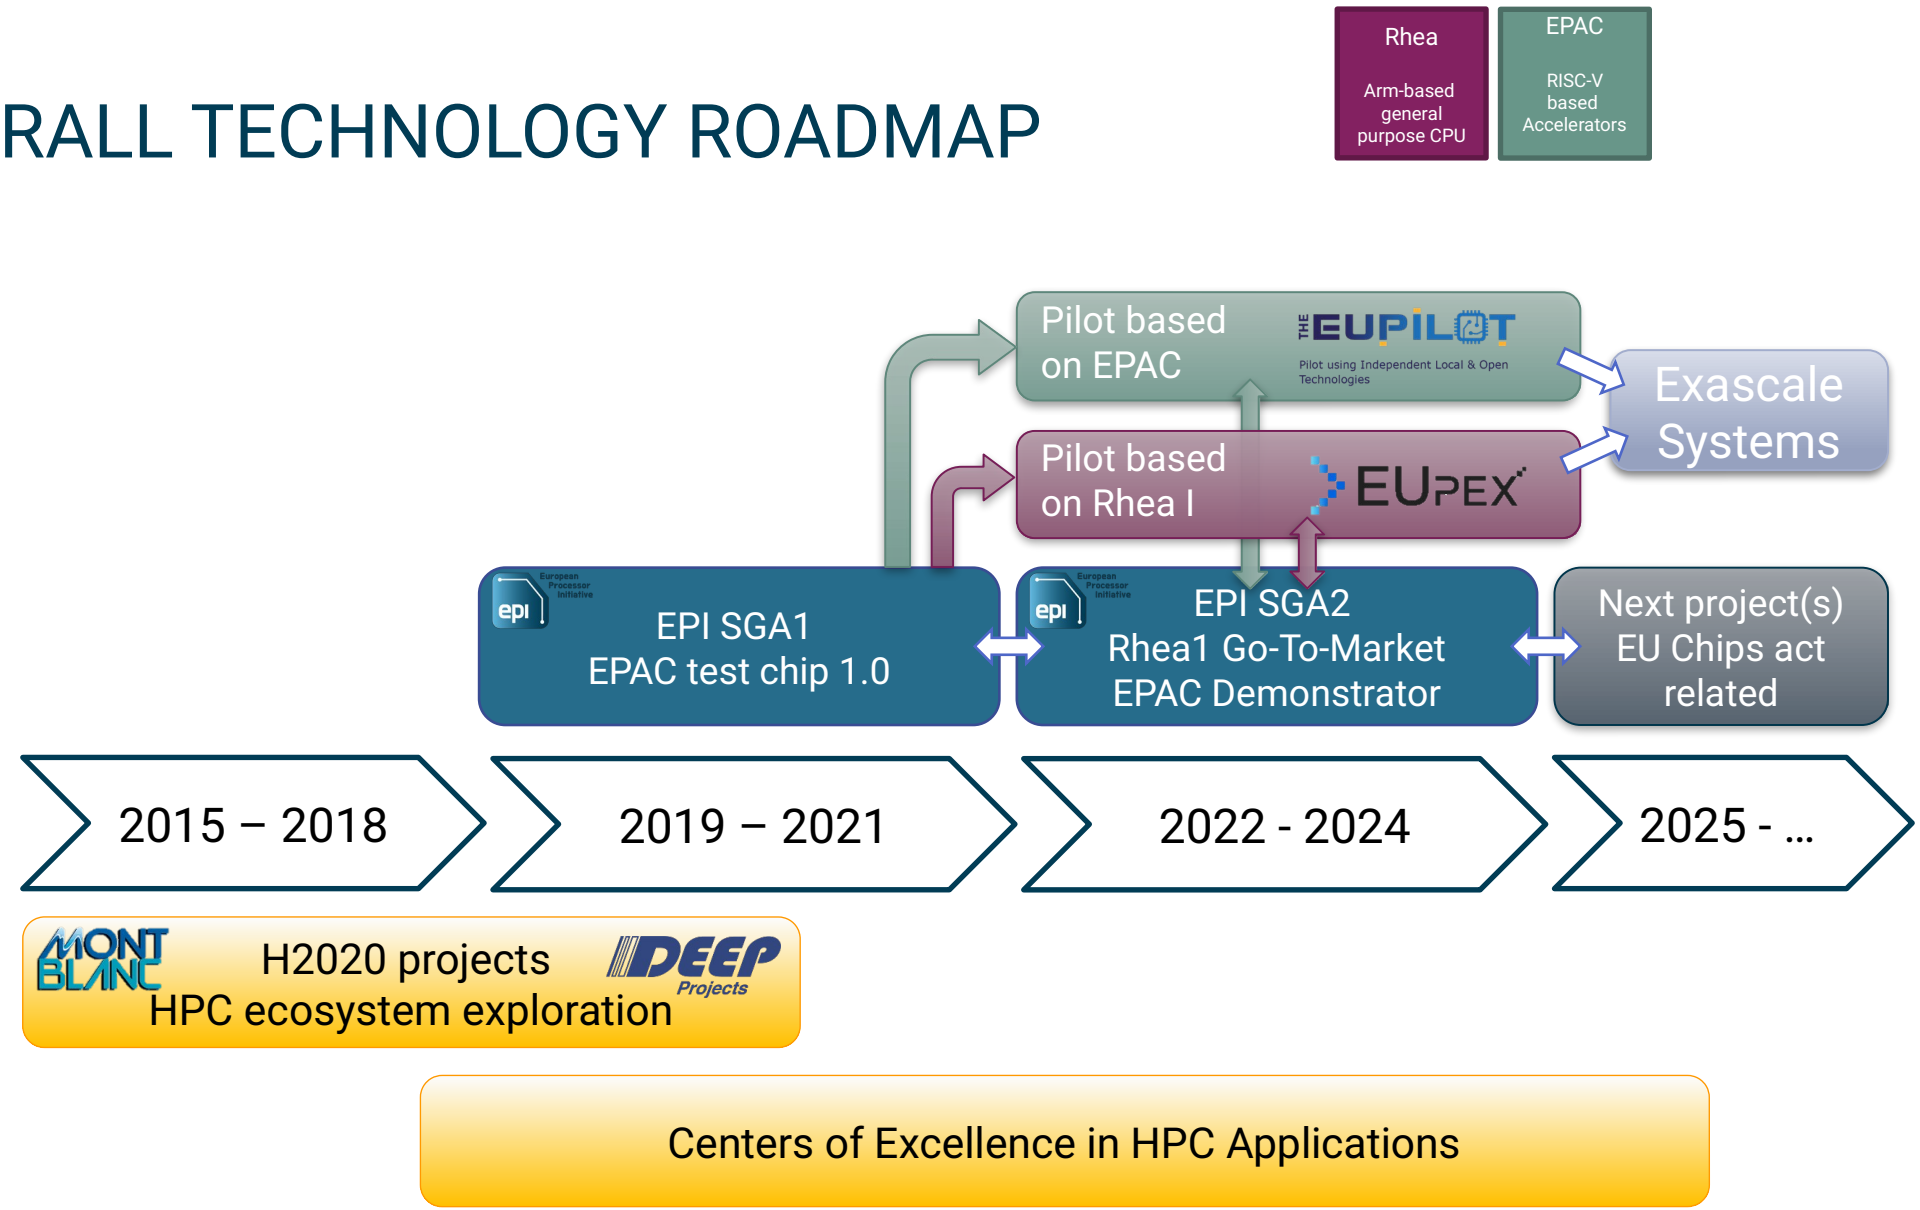

EPAC: A RISC-V ACCELERATOR

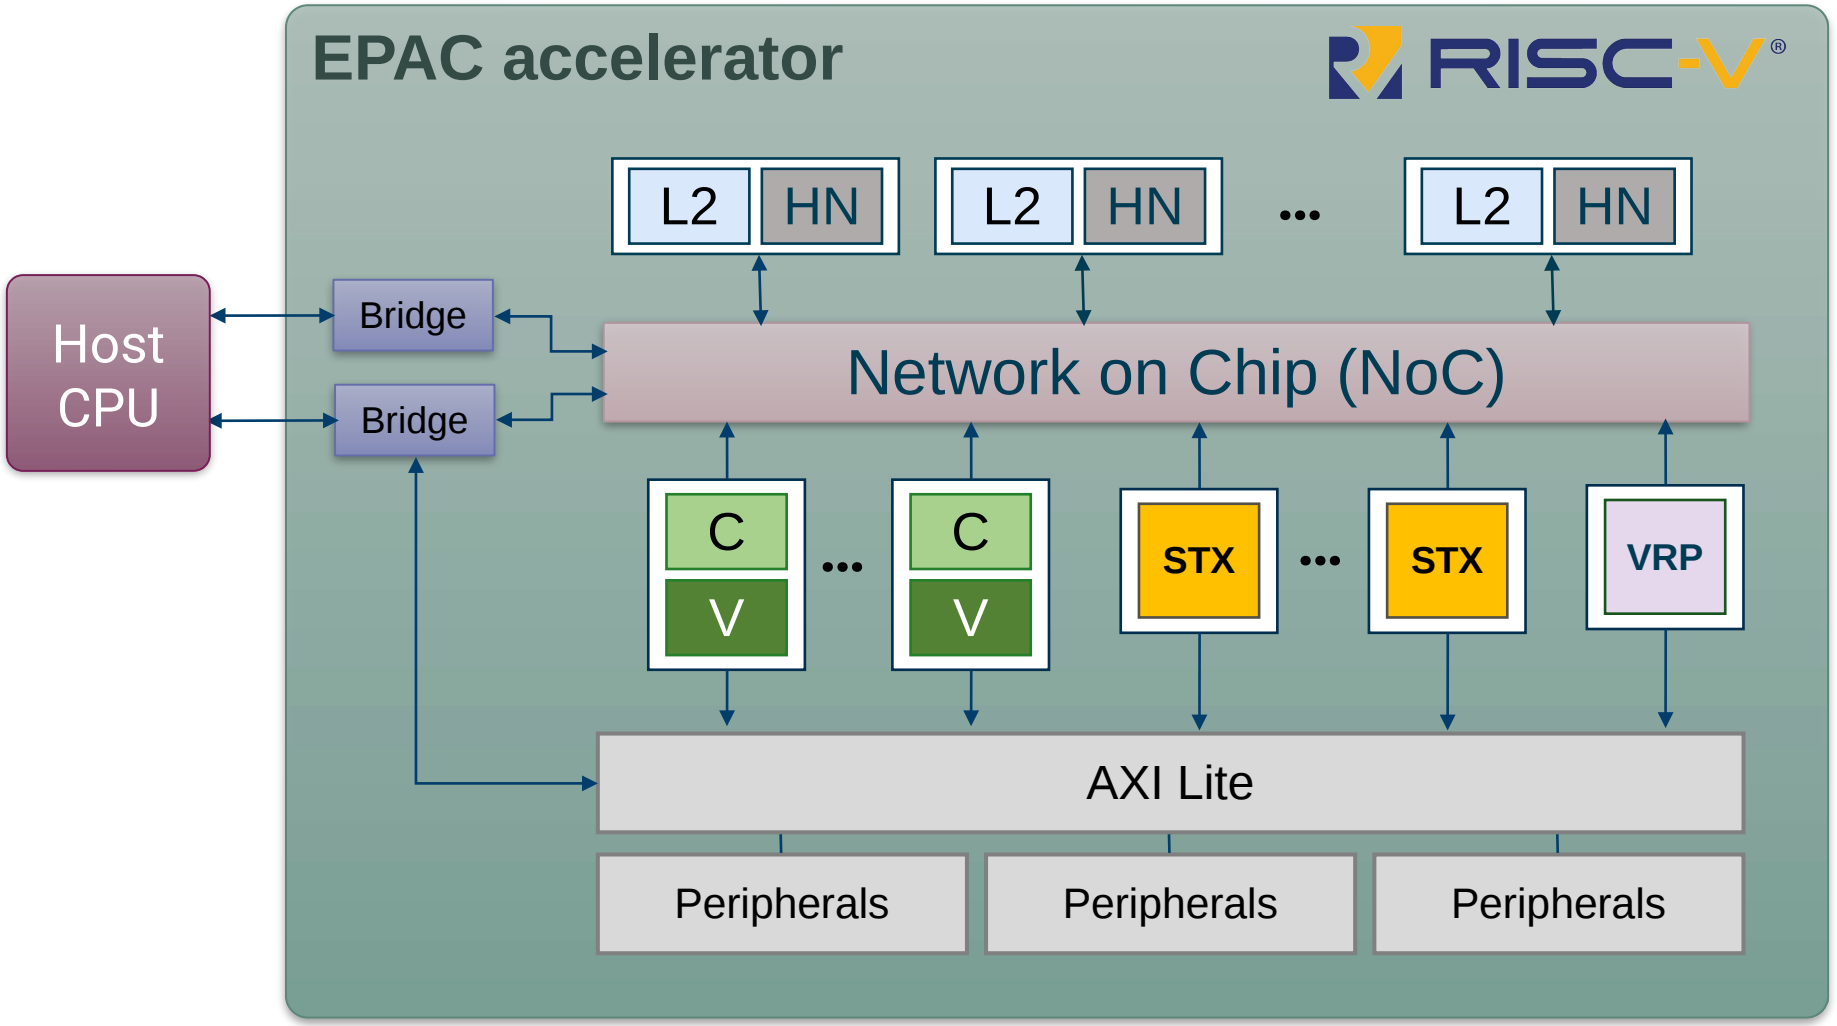

RISC-V core and VPU

F. Minervini, O. Palomar. RISC-V Summit 2021
"*Vitruvius: An Area-Efficient Decoupled Vector Accelerator for High Performance Computing*"

RISC-V core: Avispado

- 2-way in-order core
- Full HW-support for unaligned accesses
- Cache: L1I$ =16KB, L1D$ = 32KB

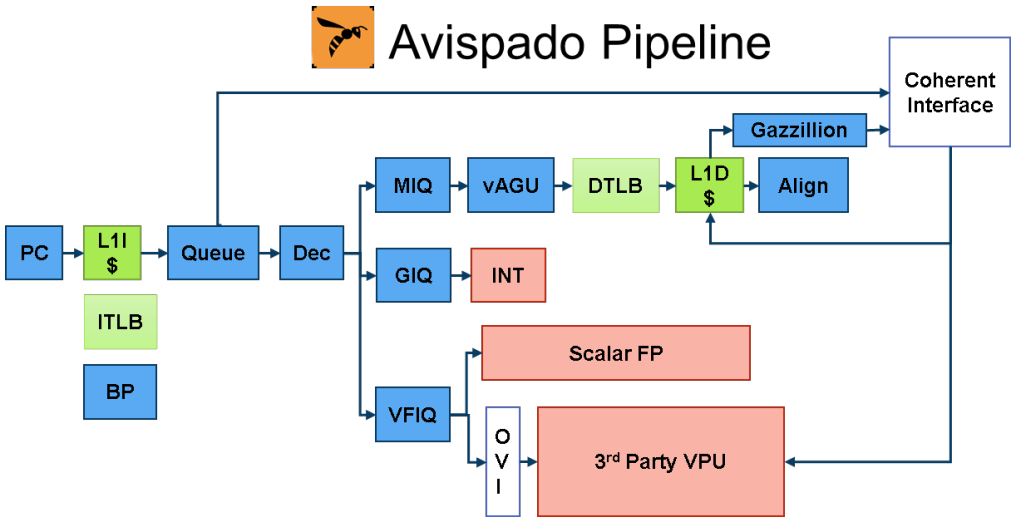

Key contact point
Roger Espasa (SemiDynamics)

VPU

- Long vectors: 256 DP elements
  - #Functional Units (FUs) << Vector Length (VL)
  - 1 vector instruction can take several cycles
- 8 Lanes per core
  - FMA/lane: 2 DP Flop/cycle
- 40 physical registers, some out of order

Architecture	**Vector register size** (1 cell = 1 double element)									
Intel AVX512	**D1**	**D2**	**D3**	**D4**	**D5**	**D6**	**D7**	**D8**		
Arm Neon	**D1**	**D2**								
A64FX	**D1**	**D2**	**D3**	**D4**	**D5**	**D6**	**D7**	**D8**		
NEC Aurora SX	**D1**	**D2**	**D3**	**D4**	**D5**	**D6**	**D7**	**D8**	**…**	**D256**
EPAC VPU	**D1**	**D2**	**D3**	**D4**	**D5**	**D6**	**D7**	**D8**	**…**	**D256**

Key contact point
Adrián Cristal adrian.cristal@bsc.es

EPAC programming environment

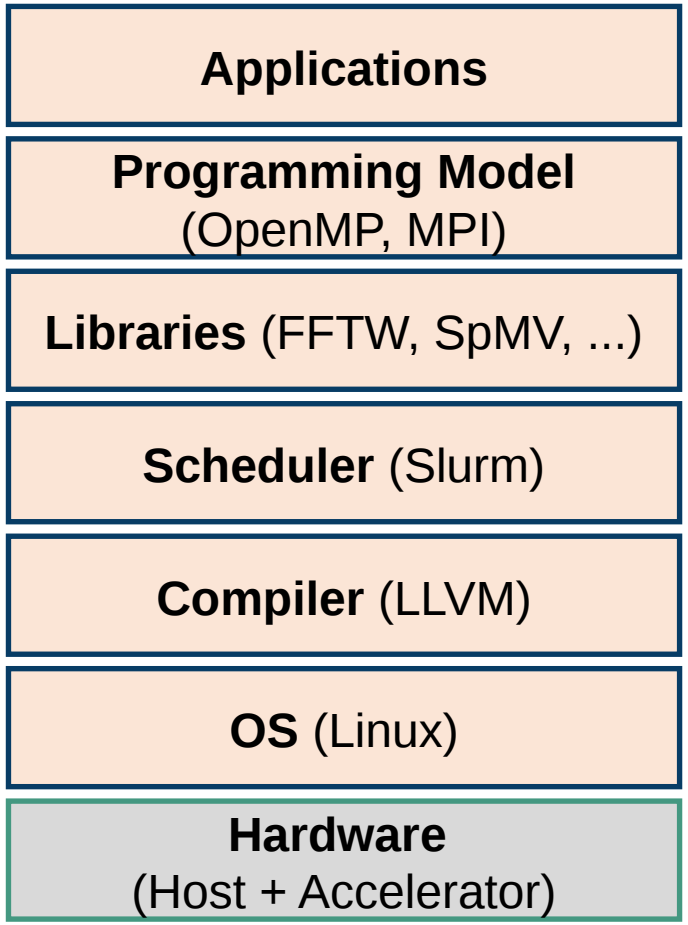

- Offload of processes from host to accelerator
- **Interoperability MPI + OpenMP**
  - Accelerator can run MPI and OpenPM processes
- **Task-based models**
  - Taskify MPI calls
  - Single mechanism
    - Concurrency
    - Locality and data management
- **Long vectors** (256 elements, 8 lanes per core)
  - Decouple front-end from back-end
  - Convey access pattern semantics to the architecture
  - Vector length agnostic (VLA) programming and architecture

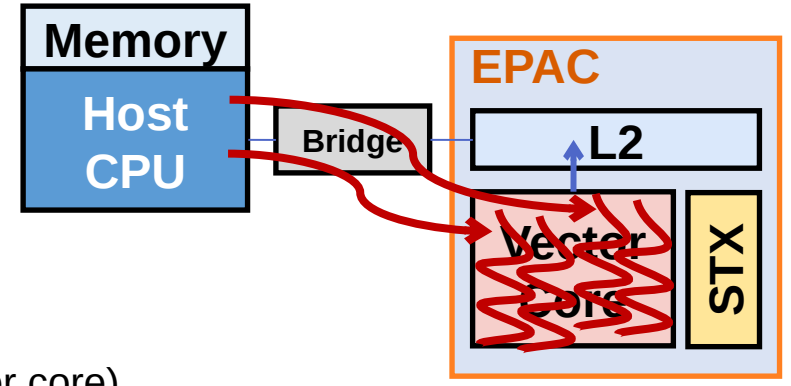

How to use the V-extensions?

- **Assembler:** always a valid option but not the most pleasant
- **C/C++ builtins (intrinsics)**
  - Low-level mapping to instructions
  - Allows embedding it into an existing C/C++ codebase
  - Allows relatively quick experimentation
- **#pragma omp simd** (aka *"Semi automated vectorization"*)
  - Relies on vectorization capabilities of the compiler
    - Usually works but gets complicated if the code calls functions
  - Also usable in Fortran
- **Autovectorization**: let compiler work for you ;-)

Interested in compiling your code in RVV?
Roger Ferrer roger.ferrer@bsc.es

RISC-V platforms (SDV)

Commercial and FPGA-based

- Hardware/Software infrastructure for Continuous Integration and RTL check
  - HiFive commercial hardware (scalar)
  - EPAC RTL (1core) on FPGA

- Platform to:
  - Demonstrate a full HPC software stack
    - Linux, compiler, libraries, job scheduler, MPI
  - Test latest RTL with complex codes
    - Advanced performance analysis tools
    - Accurate timing available

RISC-V Commercial (only scalar)

EPAC FPGA-based implementation (vector support)

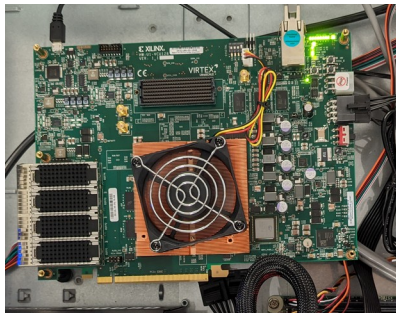

Full SDV

Interested in testing your code on EPAC?
Filippo Mantovani filippo.mantovani@bsc.es

EPAC 1.0 Test Chip (Tapeout Q2 2021)

- Chip fabrication Q2 2021
- Global Foundries 22nm (GF22FDX)
- Final Top level chip floorplan
- Total area:
  - 5943 X 4593 um^2
  - (27.297 mm^2)

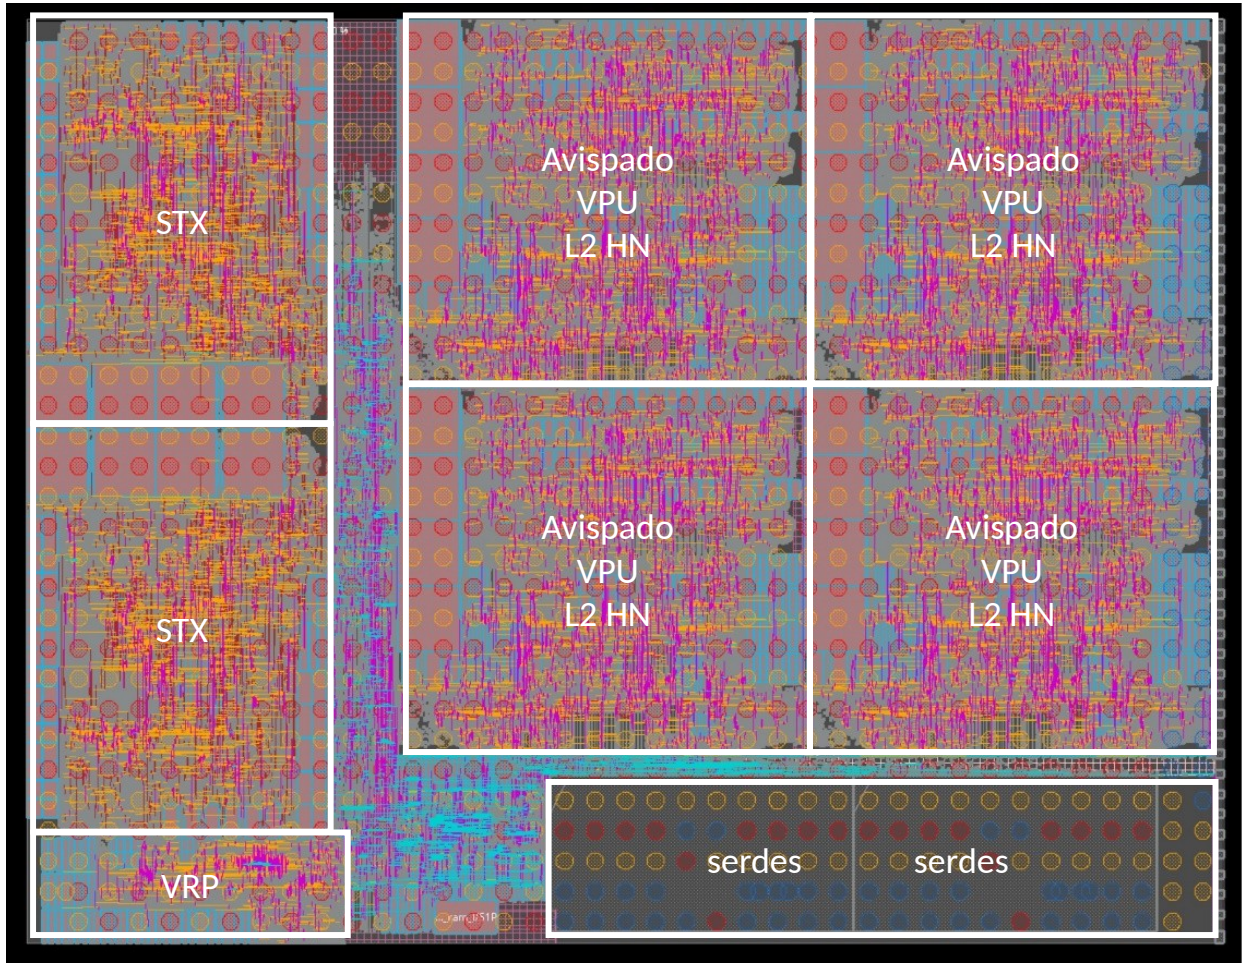

EPAC 1.5 Test Chip (Tapeout Q4 2022)

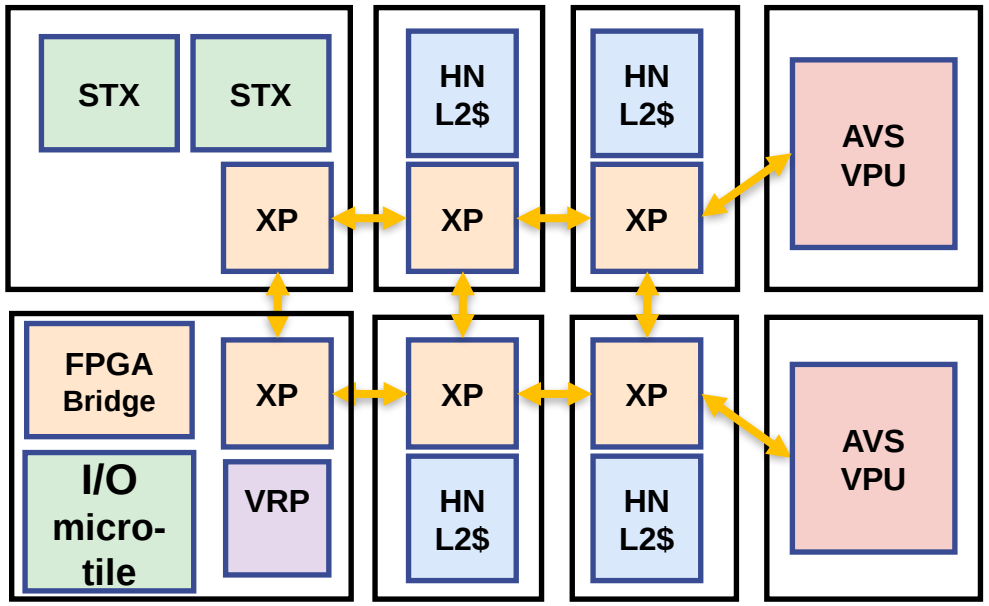

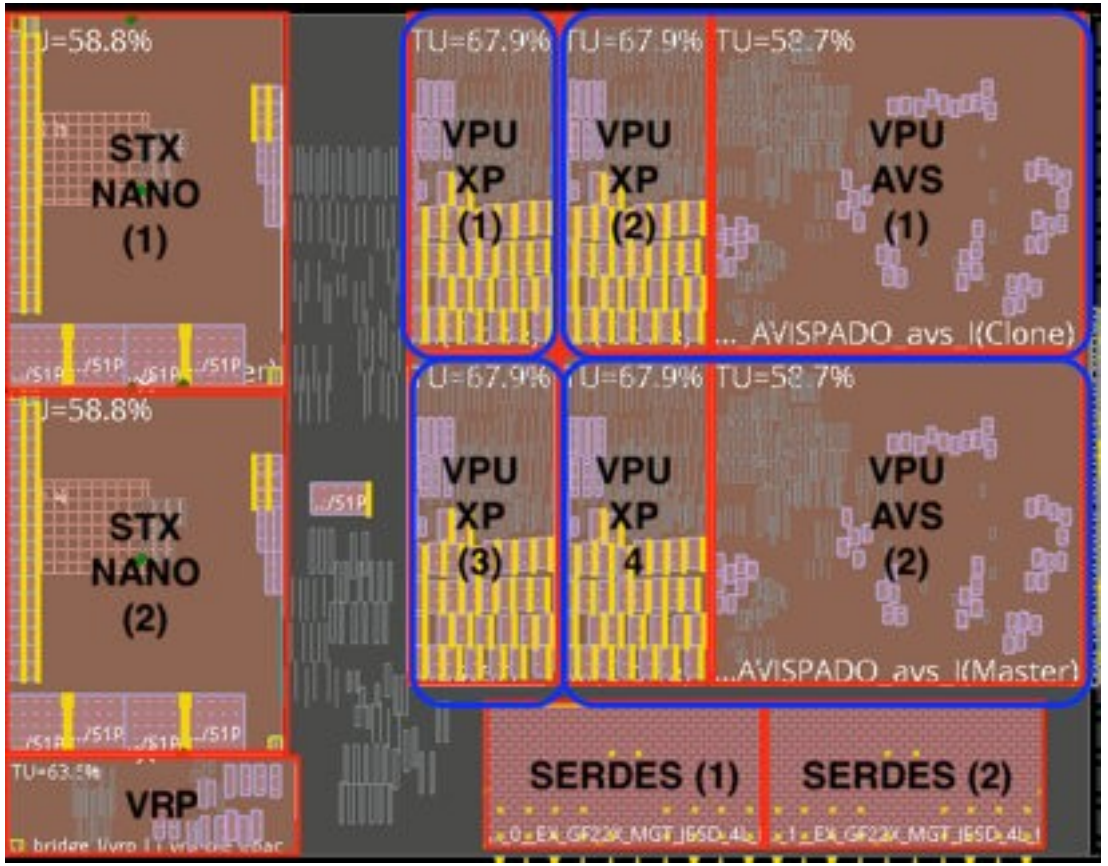

DRAC Overview

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

Generalitat de Catalunya Departament de Recerca i Universitats

Unió Europea Fons Europeu de Desenvolupament Regional

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.

DRAC Project

Designing RISC-V-based Accelerators for next generation Computers

- **DRAC**: **D**esigning **R**ISCV-based **A**ccelerators for next generation **C**omputers
  - 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

RIS3CAT
Research and Innovation Strategies for Smart Specialisation a CATALUNYA

UNIVERSITAT POLITÈCNICA DE CATALUNYA BARCELONATECH
Facultat d'Informàtica de Barcelona

UNIVERSITAT DE BARCELONA

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

UAB Universitat Autònoma de Barcelona

UNIVERSITAT ROVIRA I VIRGILI

Generalitat de Catalunya
Departament de Recerca i Universitats

Unió Europea
Fons Europeu de Desenvolupament Regional

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.

DRAC General Objectives

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 local ecosystem for custom hardware design and fabrication

6 Transfer DRAC technology to local and international companies

7 Transfer DRAC technology to the university system using educational kits based on DRAC designs

Generalitat de Catalunya
Departament d'Empresa i Coneixement
Secretaria d'Universitats i Recerca

Unió Europea
Fons Europeu
de Desenvolupament Regional

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.

Lagarto RISC-V Tapeouts

Designing RISC-V-based Accelerators for next generation Computers

2019 2020 2021 2022 2023

Lagarto (May'19)
- Lagarto Hun 5-stage in-order
- 150MHz (external)
- TSMC 65nm
- 2.5mm2

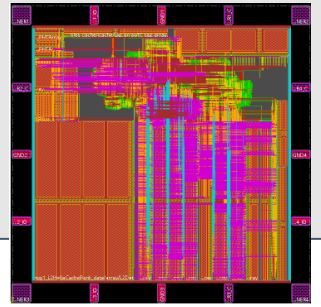

DVINO (Apr'21)
- Lagarto Hun in-order
- VPU
- PLL 600 MHz
- SDRAM mem cont
- HyperRAM
- VGA
- ADC
- TSMC 65nm
- 8mm2

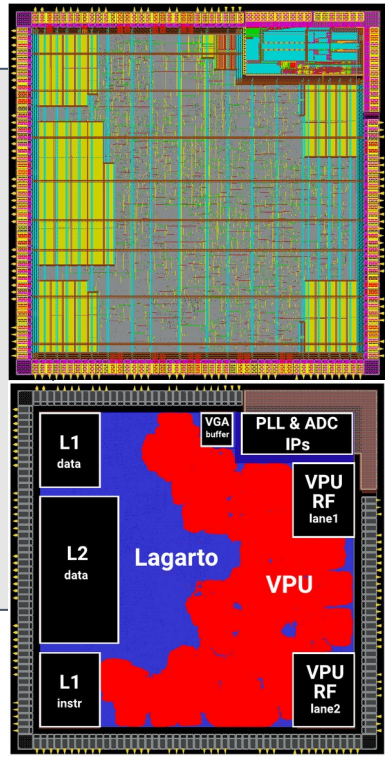

Sargantana (Feb'22)
- Sargantana 7-stage in-order
- PLL 1.2 GHz
- Custom extensions
- SDRAM
- Prototype analog IPs: SerDes 8GHz
- GF 22nm
- Area: 2.9mm2

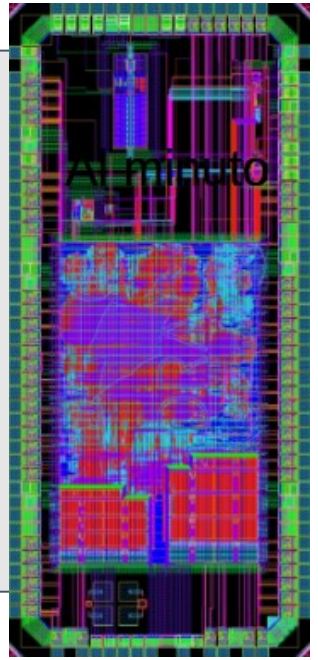

Kameleon (Dec'22)
- Lagarto Ka 11-stage ooo
- PLL 1.2 GHz
- Automotive Accel
- Crypto Accel
- Genomic Accel
- PICOS Accel
- SerDes 8GHz
- GF 22nm
- Area: 9mm2

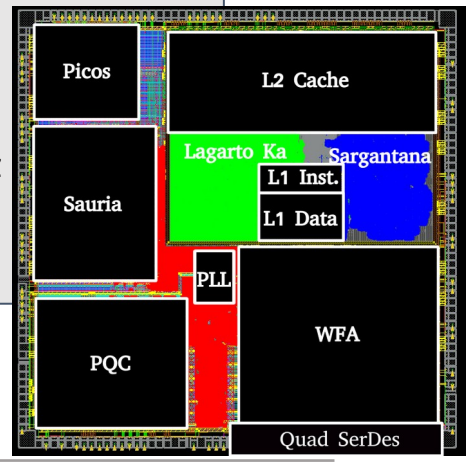

Generalitat de Catalunya
Departament de Recerca i Universitats

Unió Europea
Fons Europeu de Desenvolupament Regional

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.

Lagarto: First RISCV Tapeout (May 2019)

- 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 2.5mm2
- Fabrication and bringup
  - Submitted in May 2019
  - Samples received in Sep 2019
  - Bringup with custom PCB in Oct 2019
  - Linux boot in Dec 2019

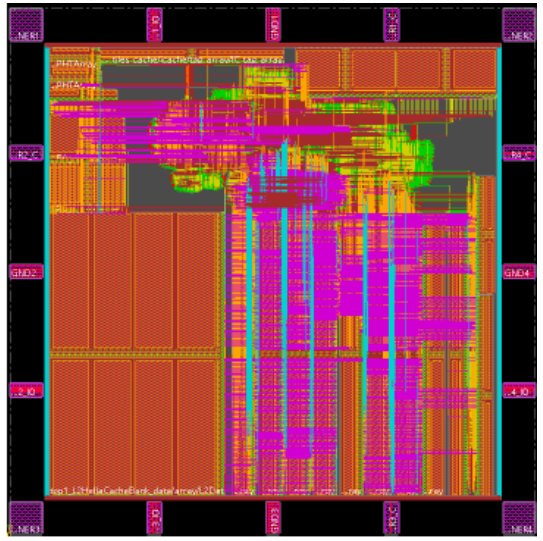

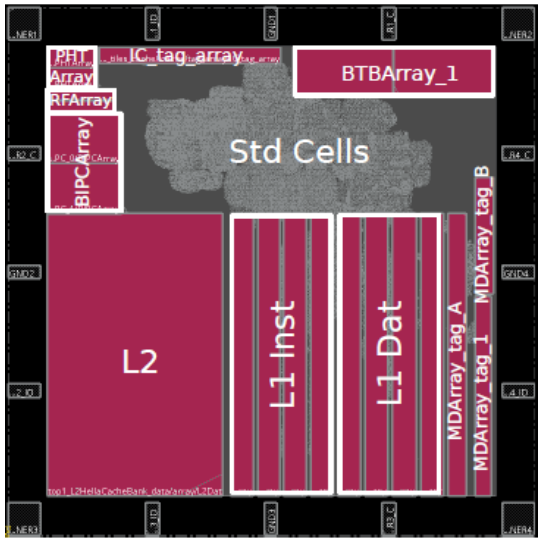

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
UNIVERSITAT POLITÈCNICA DE CATALUNYA BARCELONATECH
Centro Nacional de Microelectrónica IMB
Generalitat de Catalunya Departament de Recerca i Universitats
Unió Europea Fons Europeu de Desenvolupament Regional
The DRAC project with -file 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.

From Architecture to an Actual System

RTL design and verification → IC design → PCB design and assembly → System characterization

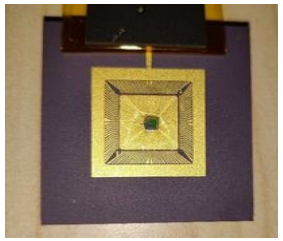

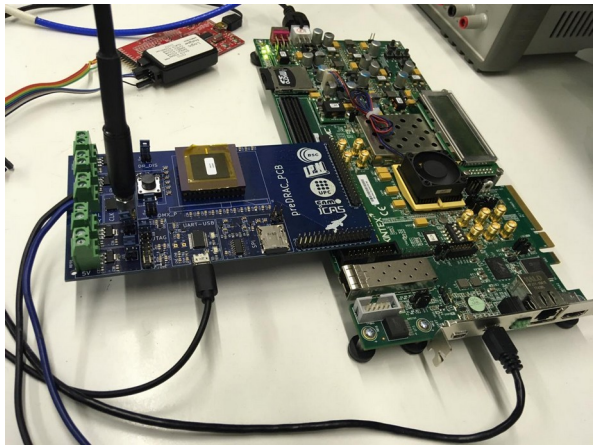

Technology information

IC design flow
Physical verification
Prepare for manufacturing

Generalitat de Catalunya
Departament de Recerca i Universitats

Unió Europea
Fons Europeu de Desenvolupament Regional

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.

DVINO: 2nd Lagarto Tapeout (Apr 2021)

- 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: 8.6mm2

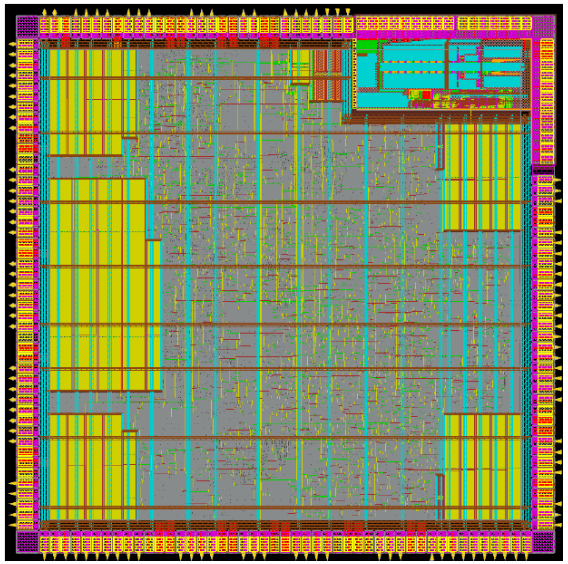

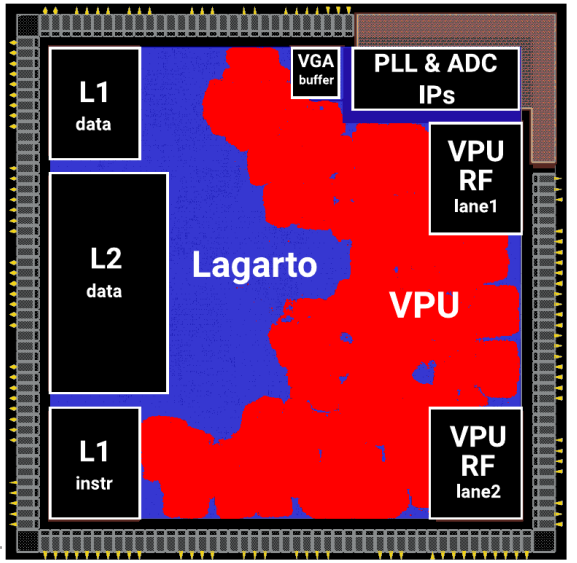

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

UNIVERSITAT POLITÈCNICA DE CATALUNYA
BARCELONATECH

Centro Nacional de Microelectrónica IMB

Generalitat de Catalunya
Departament de Recerca i Universitats

Unió Europea
Fons Europeu de Desenvolupament Regional

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.

Sargantana: 3rd Lagarto Tapeout (Feb 2022)

- 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: 2.9mm2

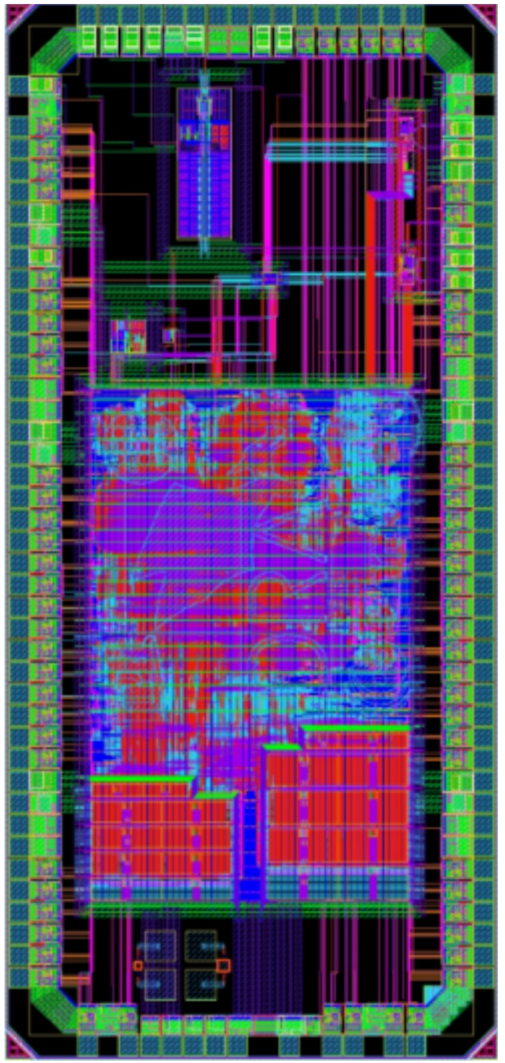

Generalitat de Catalunya Departament de Recerca i Universitats

Unió Europea Fons Europeu de Desenvolupament Regional

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.

Kameleon: 4th Lagarto Tapeout (Dec 2022)

Designing RISC-V-based Accelerators for next generation Computers

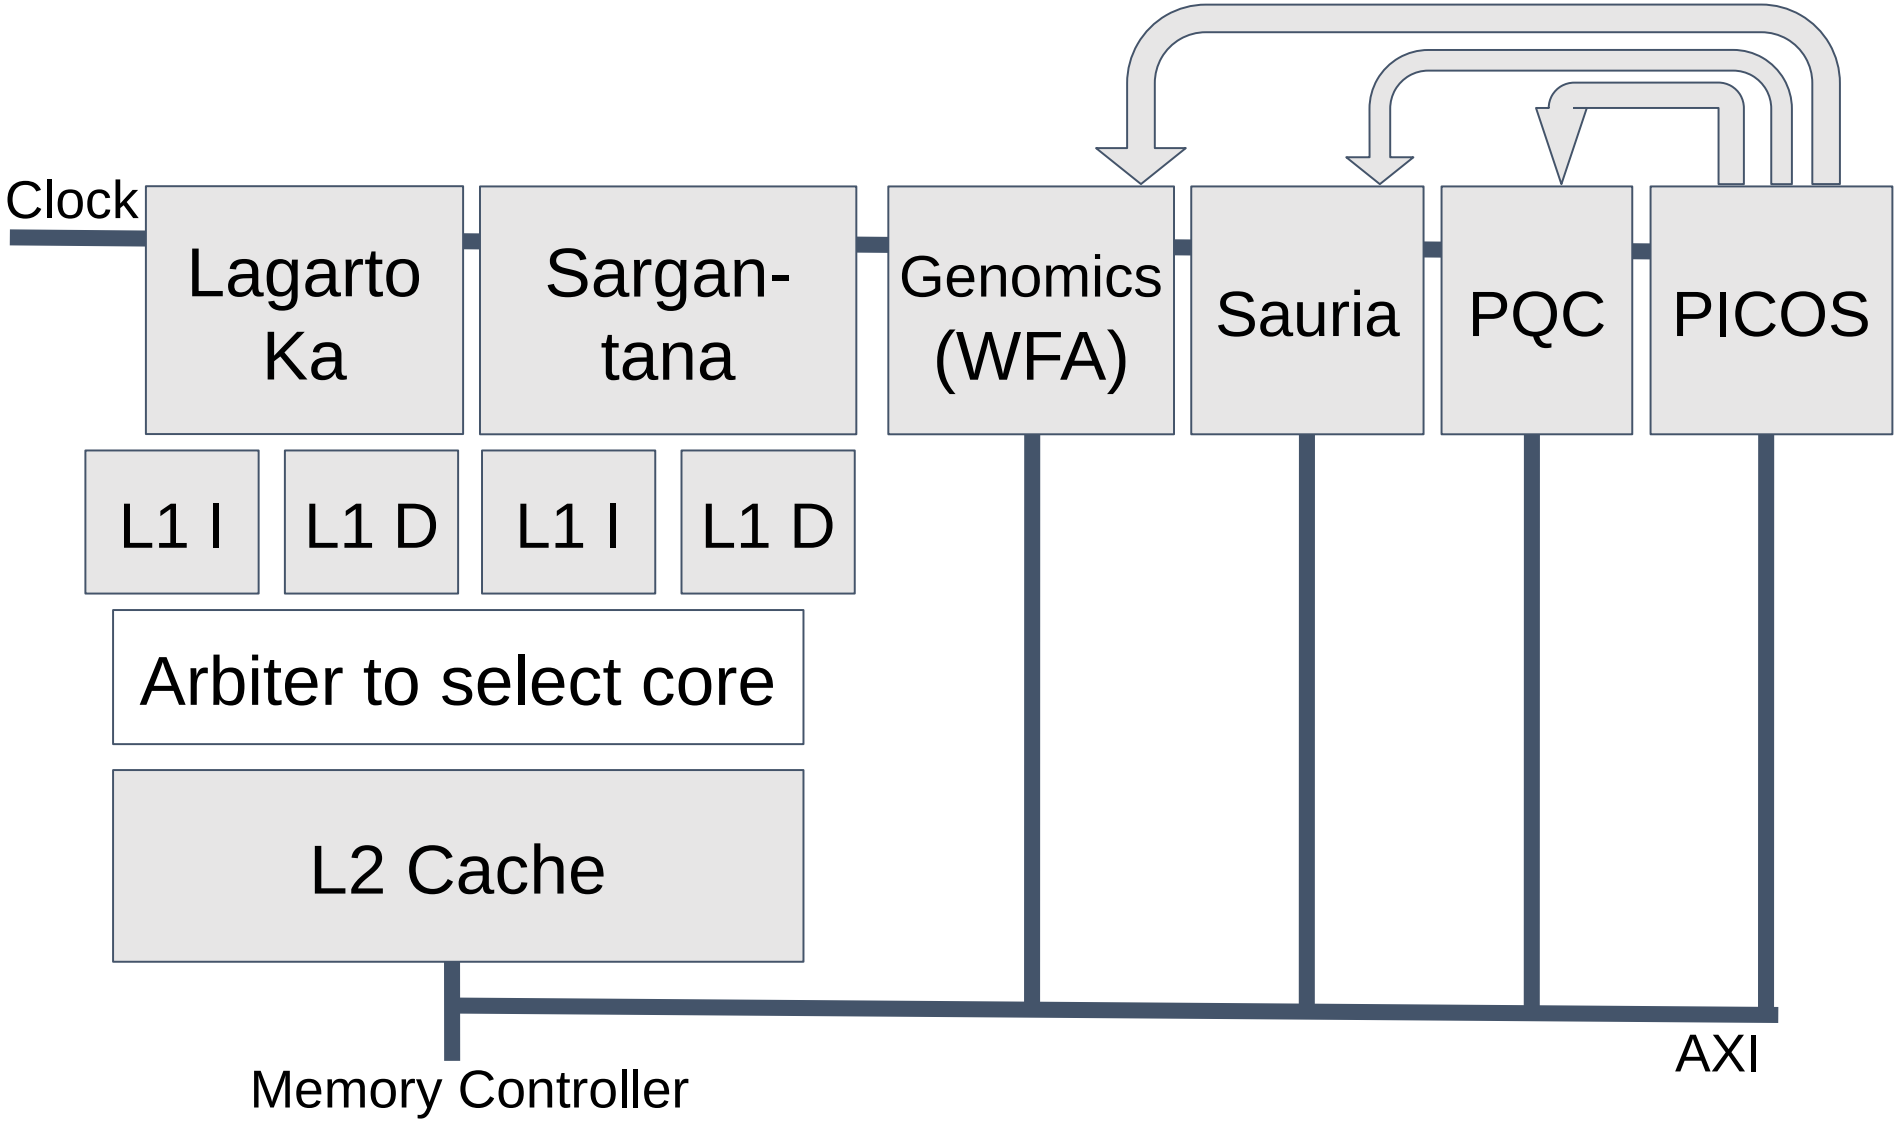

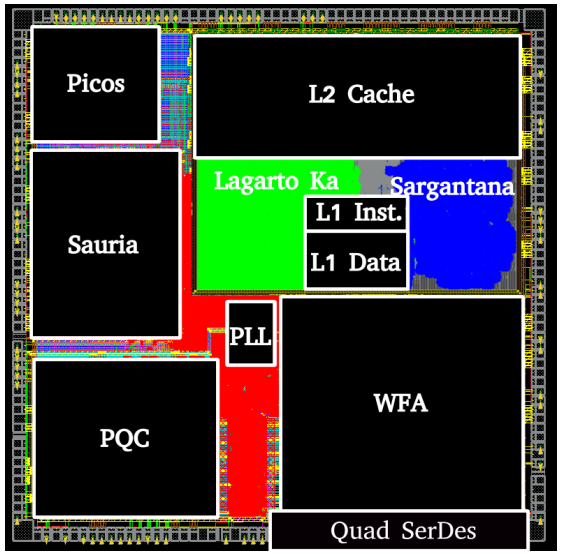

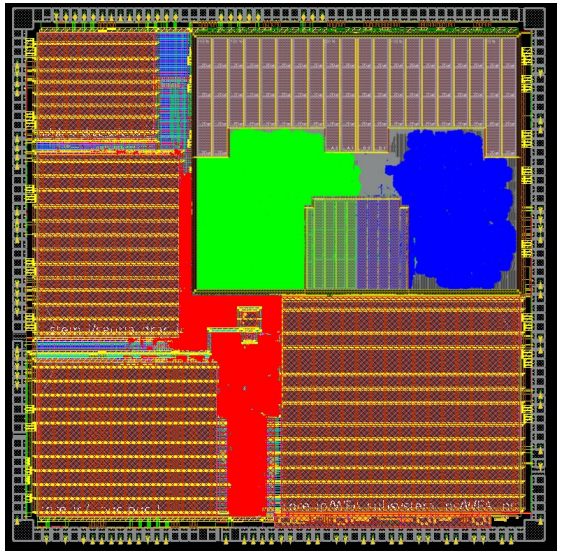

Generalitat de Catalunya
Departament de Recerca i Universitats

Unió Europea
Fons Europeu de Desenvolupament Regional

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.

Lagarto Ka Out-of-Order Core

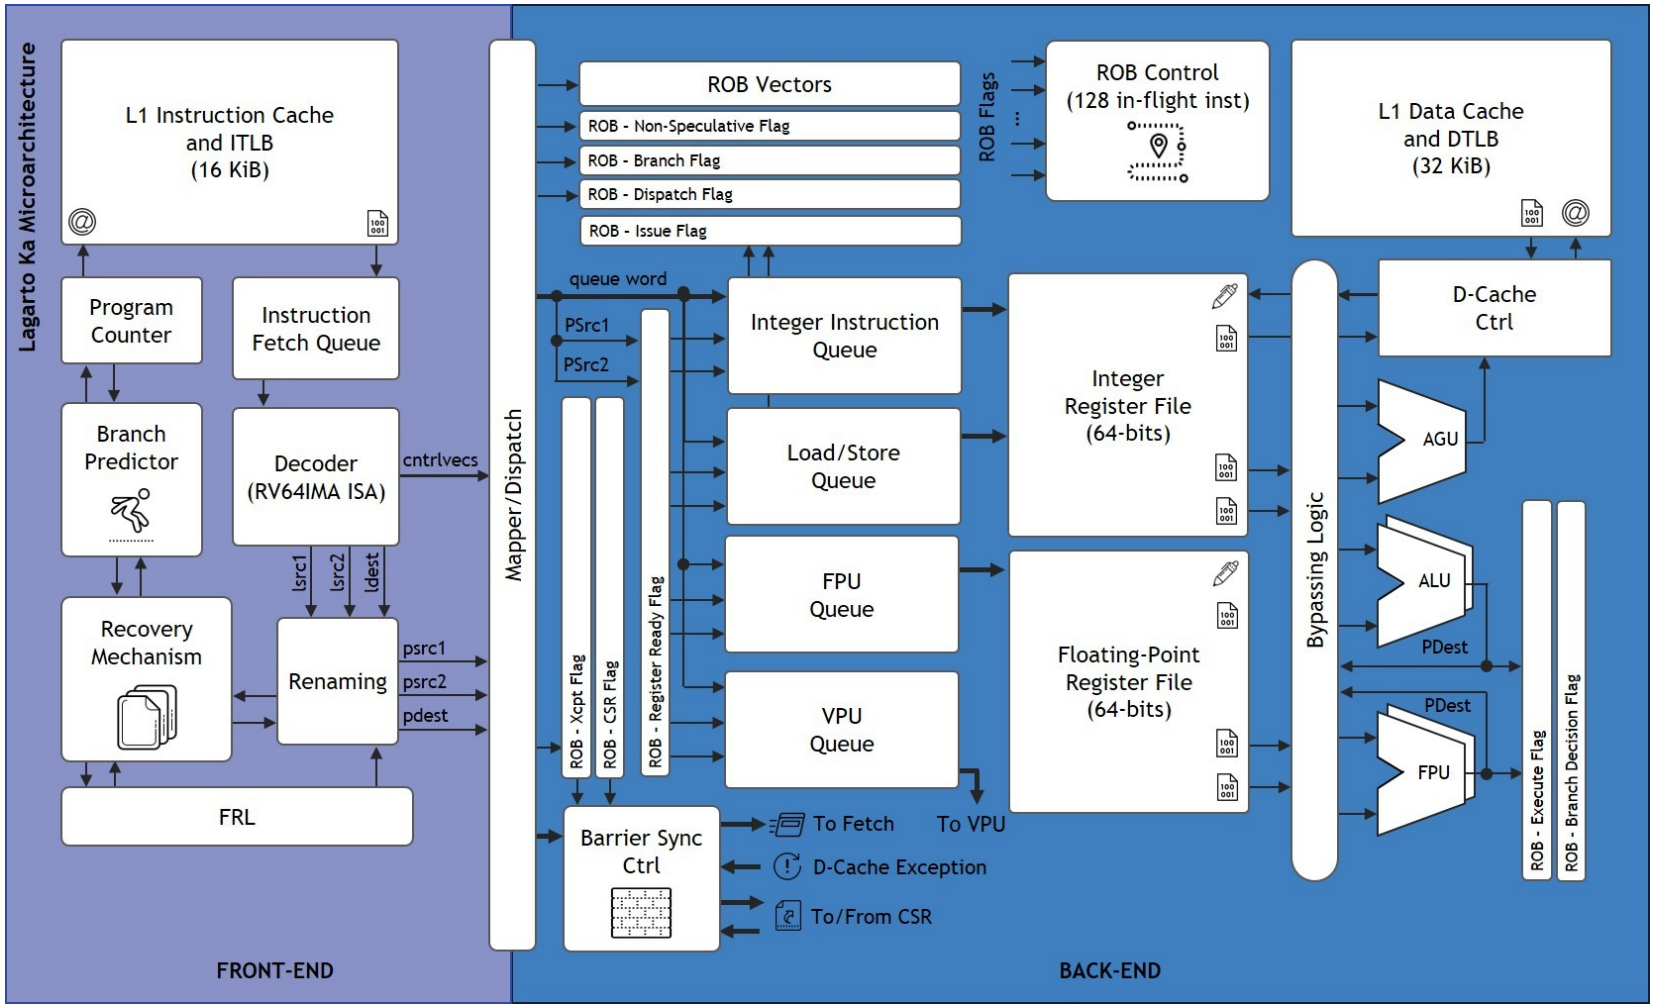

Current Features

- 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)

Generalitat de Catalunya
Departament de Recerca i Universitats
Unió Europea
Fons Europeu de Desenvolupament Regional
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.

SAURIA Physical Design

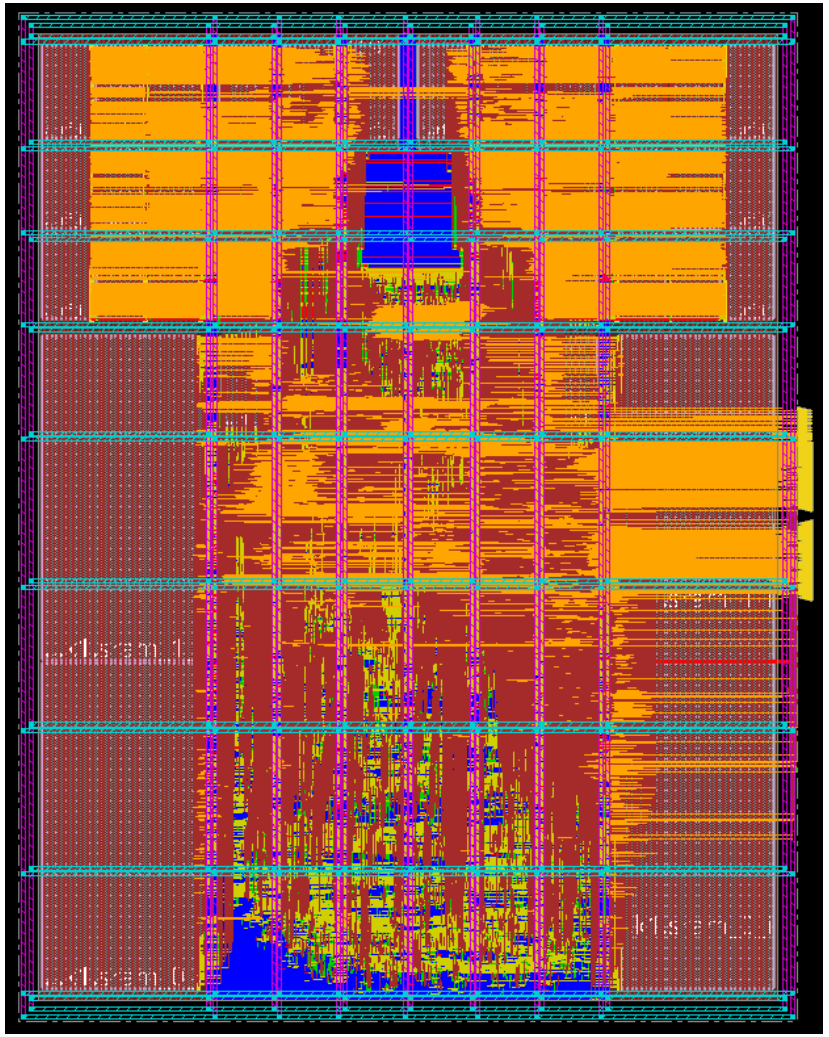

Specs:

8x16 Array: 128 Processing Elements (PE)
Approximate logic in PE multipliers & adders
1.00 x 0.95 mm = 0.95 mm2
128 GFLOP/s @ 500MHz

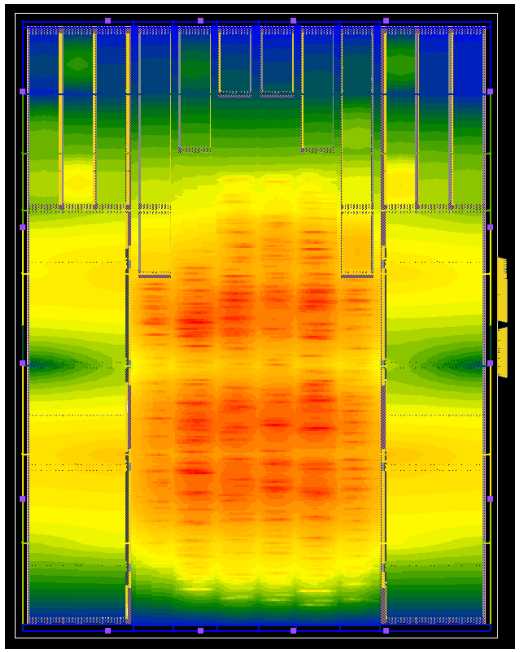

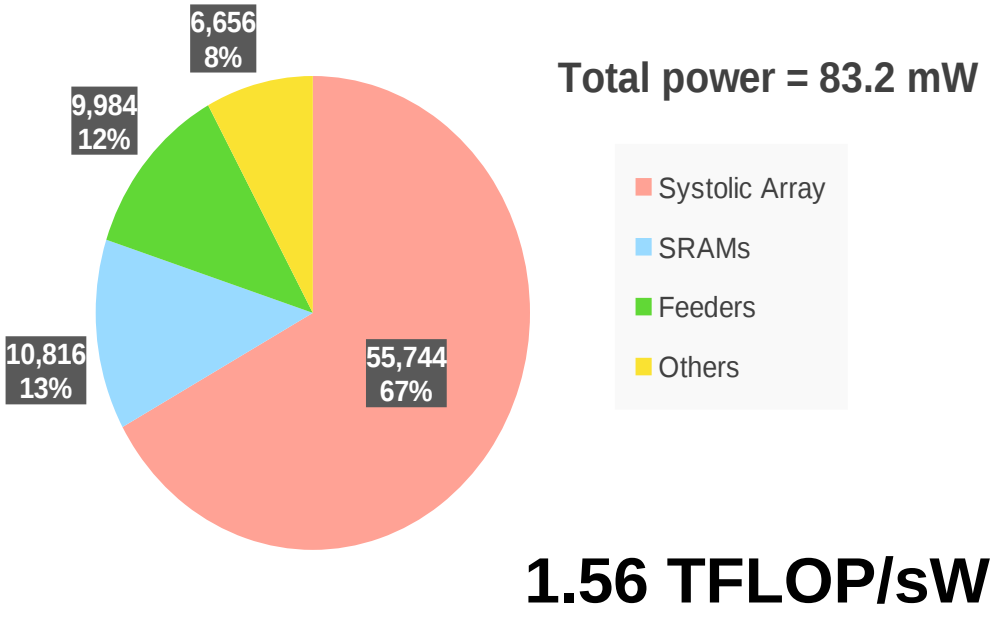

1.56 TFLOP/sW

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

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
- **Monolythic 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: 1.6mm2 in Global Foundries 22nm
  - Performance speedups: 515x with 10K reads, 10% error

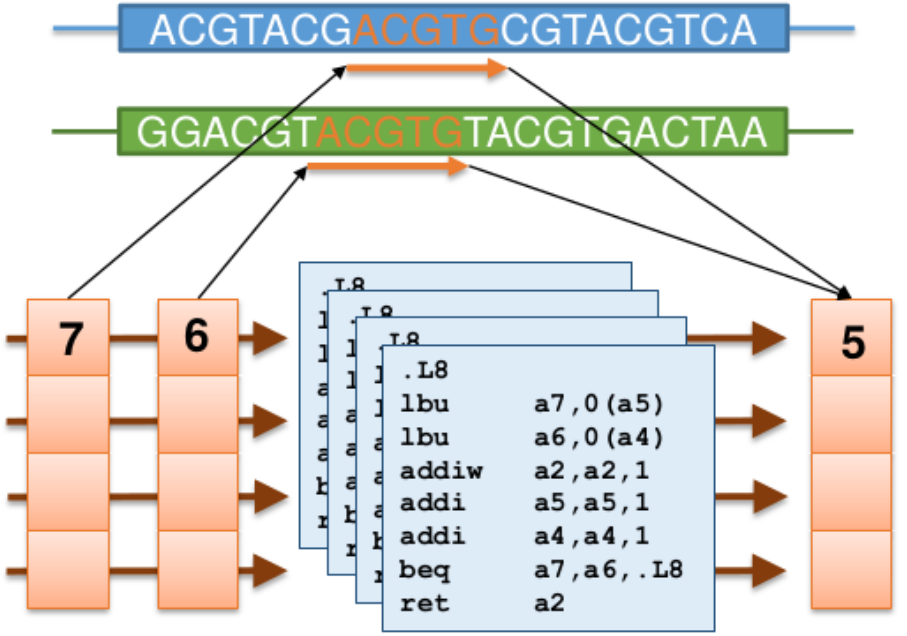

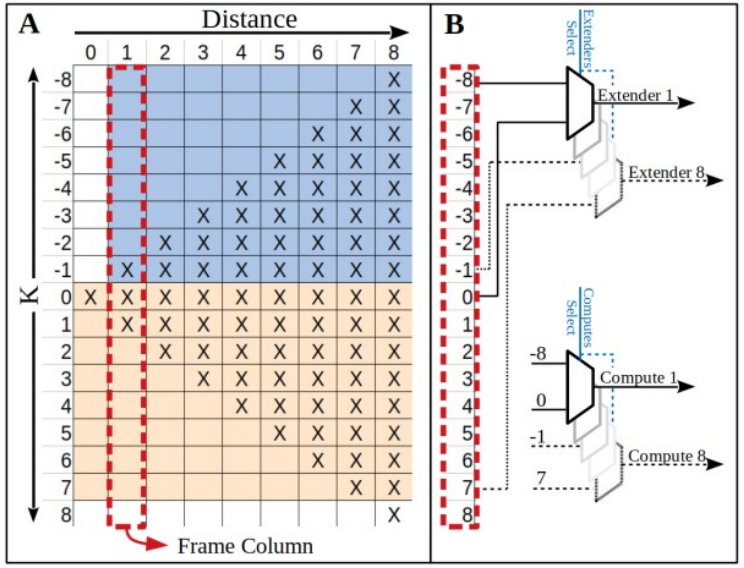

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

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]
  - Monolythic CME accelerator based on HLS
  - Integration of the CME accelerator in Lagarto SoC via AXI interface

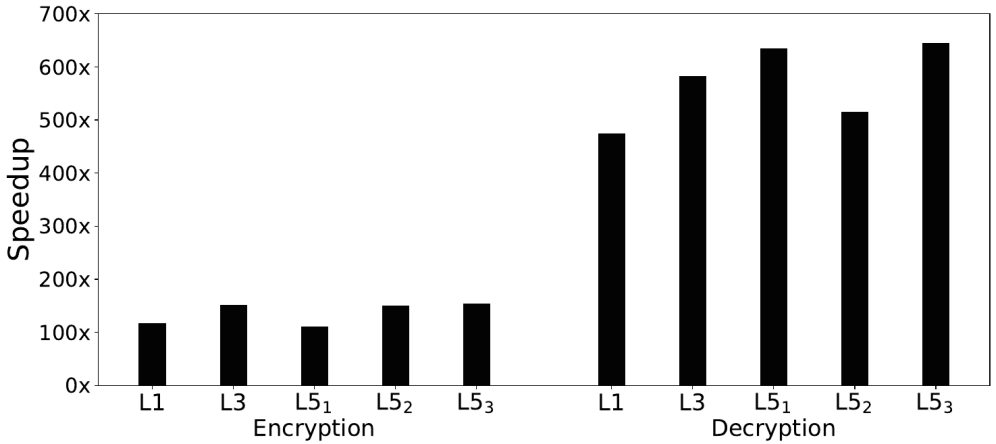

- Accelerate other KEM and digital signature (DS) schemes:
  - **NRTU KEM** and **Crystals-Kyber KEM** / **Crystals-Dilithium DS**
  - Both algorithms rely on very different operations and will require different acceleration techniques.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

[FPL'21] V. Kostalampros et al., HLS-Based HW/SW Co-Design of the Post-Quantum Classic McEliece Cryptosystem. FPL 2021.

Building a RISC-V Ecosystem in Barcelona

Designing RISC-V-based Accelerators for next generation Computers

Lagarto Team (Sept 2019)

DRAC KoM (Feb 2020)

DRAC F2F (Jul 2022)

Generalitat de Catalunya
Departament de Recerca i Universitats

Unió Europea
Fons Europeu de Desenvolupament Regional

The DRAC project
within the framewo

Building a RISC-V Ecosystem in Barcelona

Designing RISC-V-based Accelerators for next generation Computers

DRAC Final Workshop (December 2022)

Generalitat de Catalunya
Departament de Recerca i Universitats

Unió Europea
Fons Europeu de Desenvolupament Regional

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.

Lagarto Roadmap

First Steps Towards a Lagarto Multicore

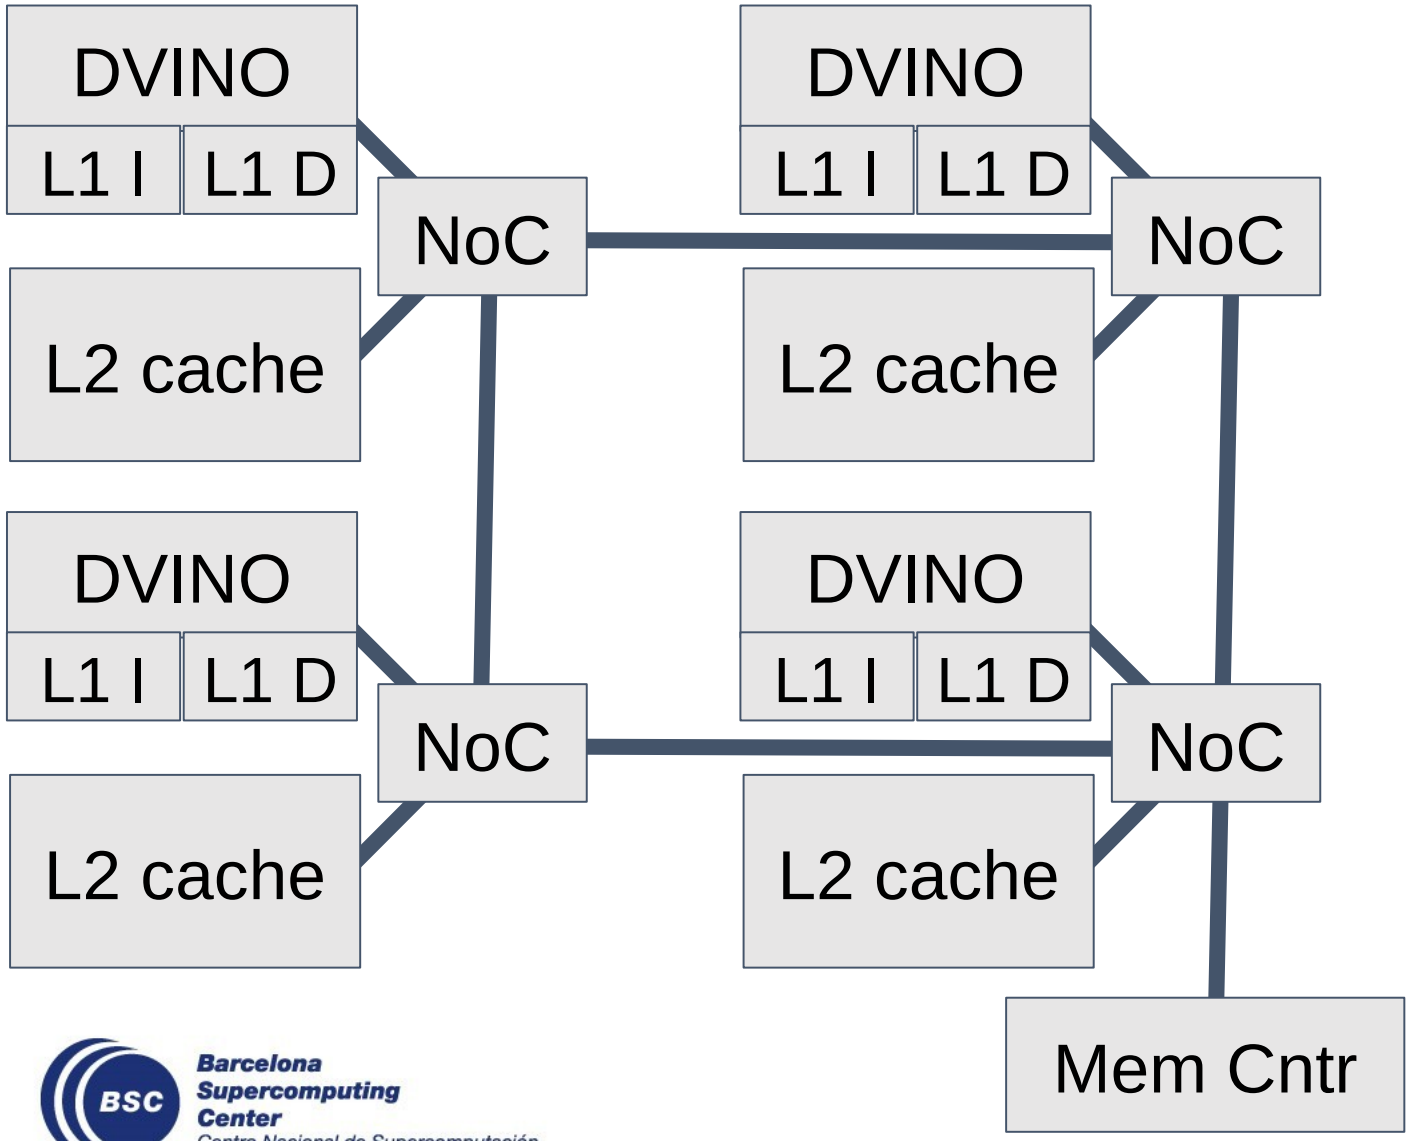

- Multicore design based on:
  - DVINO processor
  - RISC-V ISA support: I, M, A, F, D, C, V
  - 4-lane VPU
  - OpenPiton 2-level cache hierarchy (priv. L1, shared L2)
- Current status
  - Linux boot (openSBI)
  - RTL simulation of parallel applications (up to 64 cores)
  - Multiple memory controllers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

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

European Processor Initiative

MEEP
MareNostrum Experimental Exascale Platform

Towards a RISC-V Heterogeneous Manycore

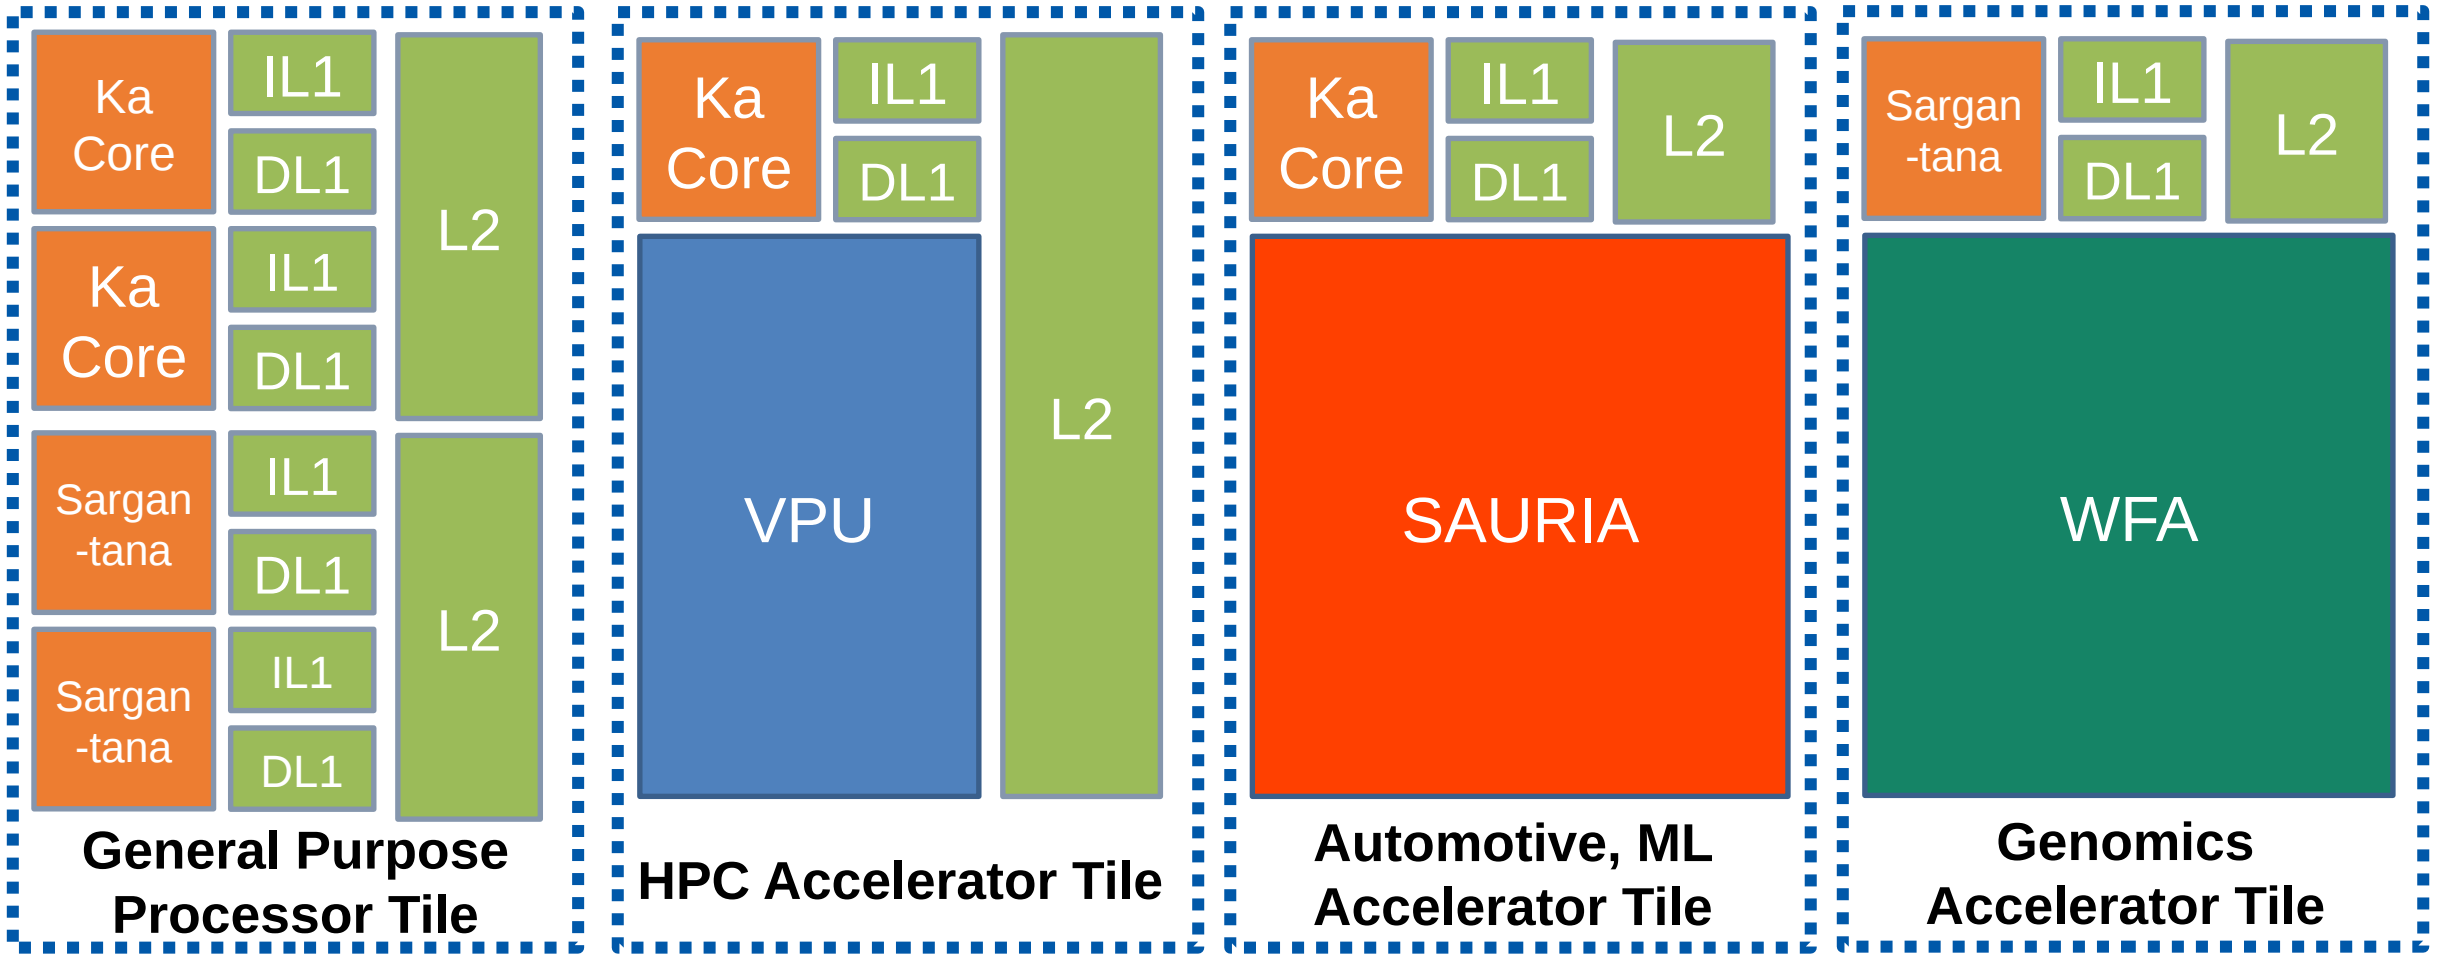

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Other Domain-Specific Accelerator Tiles (sparse, security, safety, etc.)

Chips Act, PERTE Chip, Intel and more!

EUROPEAN CHIPS ACT

The European Chips Act will ensure that the EU strengthens its semiconductors ecosystem, increases its resilience, as well as ensure supply and reduce external dependencies.

1. Strengthen Europe's research and technology leadership towards smaller and faster chips

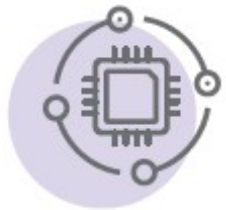

2. Build and reinforce capacity to innovate in the design, manufacturing and packaging of advanced chips

3. Put in place a framework to increase production capacity to 20% of the global market by 2030

4. Address the skills shortage, attract new talent and support the emergence of a skilled workforce

5. Develop an in-depth understanding of the global semiconductor supply chains

The Chips Act should result in additional public and private investments of more than **€15 billion**.

These investments will complement:

- **existing programmes** and actions in research & innovation in semiconductors (Horizon Europe, Digital Europe programme)
- **announced support** by Member States.

In total, **more than €43 billion of policy-driven investment will support the Chips Act until 2030**, which will be broadly matched by long-term private investment.

3. EXECUTION

ACTIONS & BUDGET (I)

	M€
COMPONENT I. BOLSTERING SCIENTIFIC CAPACITY	**1.165**
ACTION 1: Development of R&D&i on cutting-edge and alternative architecture microprocessors	475
ACTION 2: Development of R&D&i on integrated photonics	150
ACTION 3: Development of R&D&i on quantum chip development	40
ACTION 4: Budget line for the IPCEI on Microelectronics and Communication Technologies	500
COMPONENT II. DESIGN STRATEGY	**1.330**
ACTION 5: Creation of cutting-edge alternative architecture microprocessor fabless companies	950
ACTION 6: Creation of pilot lines	300
ACTION 7: Creation of a network for education, training and skills-building in relation to semiconductors	80

3. EXECUTION

ACTIONS & BUDGET (II)

	M€
COMPONENT III. CONSTRUCTION OF FABRICATION PLANTS IN SPAIN	9.350
ACTION 8: Creating fabrication capacity at sizes below 5 nm	7.250
ACTION 9: Creating fabrication capacity at sizes above 5 nm	2.100
COMPONENT IV. STIMULATING THE ICT MANUFACTURING INDUSTRY IN SPAIN	400
ACTION 10: ICT manufacturing industry incentive scheme	200
ACTION 11: Creation of a chips fund	200
GOVERNANCE	5
Special Commissioner for the Microelectronic and Semiconductors Project	5
TOTAL PUBLIC INVESTMENT	12.250

Intel Labs Barcelona are back!

- New joint Intel – BSC Laboratory to design HPC processors based on RISC-V technology
- Funding: 400M$ in the next 10 years. Headcount: ~200 (estimated)

Other companies will also come to Spain!

Path to Zettascale

Developing European Hardware/Software Technology

Full Stack Open Source HPC Ecosystem

Build Full System based on RISC-V: MN6 and many others

European & Global collaboration

Intel and BSC: Continuing to collaborate into the Zettascale era

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

intel

Research & Product

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

BSC is building a New Lab!
100+ Job Opportunities

Research Engineer/Researcher

R&D for Zettascale and beyond: Applications to Accelerators

If your experience and/or motivation include any of the disciplines and skills below, check out our QR:

Hardware (Processor architecture, micro-architecture, accelerators, memory hierarchy, memory controllers, HBM, DRAM, non-volatile memory, RTL design, VHDL, verilog, SystemC, System verilog, Synopsys, Cadence, Mentor Graphics, synthesis, place and route, timing closure, packaging, PCB design, verification, validation, CI, post-silicon debug, DFT, gate-level simulation,...)

Software (programming models, MPI, compilers (LLVM), SYCL, OneAPI, Tensorflow, PyTorch, Apache Spark, CI/CD, operating systems, managed runtimes, OpenMP, task-based programming models, containers, security, fault-tolerance, virtualization, C/C++, Tcl, Python, Perl/Csh,...)

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Reference: 176_23_CS_Z_R0-4-RE1-4

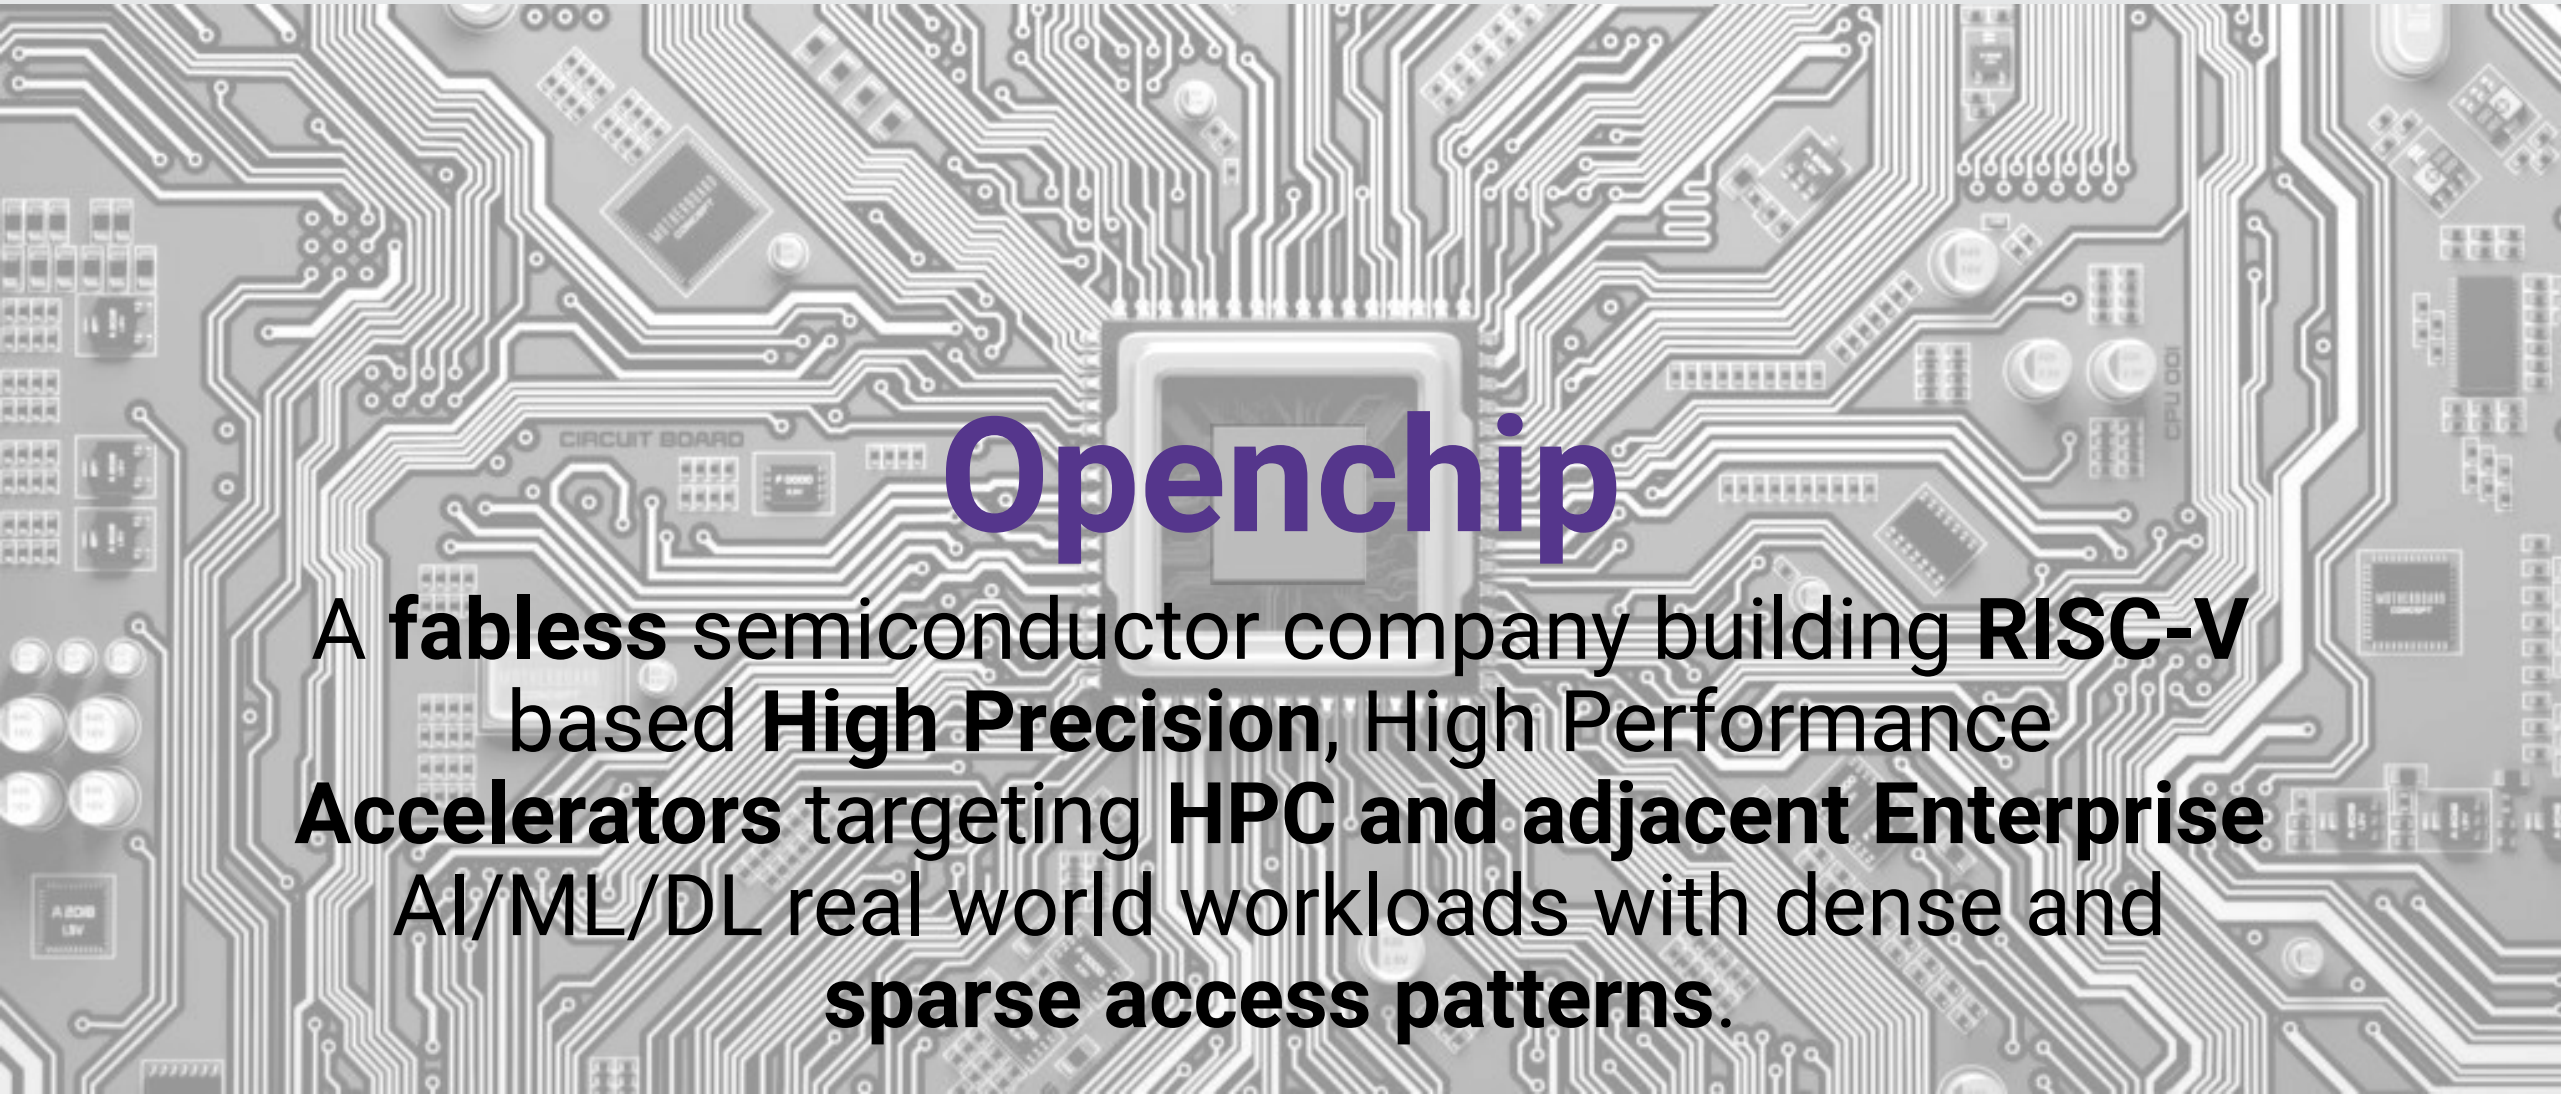

Openchip

A **fabless** semiconductor company building **RISC-V** based **High Precision**, High Performance **Accelerators** targeting **HPC and adjacent Enterprise** AI/ML/DL real world workloads with dense and **sparse access patterns**.

Conclusions

- 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 become the European domestic solution
- We need your help! Join us to contribute to the RISC-V open movement:
  - Contribute to open source community with in-house designs and tools
  - Contribute to European RISC-V econsystem and projects
  - Master/PhD student and research engineer positions are available

Generalitat de Catalunya
Departament de Recerca i Universitats

Unió Europea
Fons Europeu de Desenvolupament Regional

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.

RISC-V Summit Europe 2023

- Join us in Barcelona between June 5 and 9 2023!!!

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Early registration deadline: Apr 30 2023

Designing RISC-V-based Accelerators for next generation Computers

Thank you!

Generalitat de Catalunya
Departament d'Empresa i Coneixement
Secretaria d'Universitats i Recerca

Unió Europea
Fons Europeu
de Desenvolupament Regional

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.