

ETH zürich



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

RISC-V is the Future

The Open Platform for the Space Renaissance

Luca Benini

lbenini@iis.ee.ethz.ch, luca.Benini@unibo.it

PULP Platform

Open Source Hardware, the way it should be!



pulp-platform.org

@pulp_platform

[company/pulp-platform](https://www.linkedin.com/company/pulp-platform)

[youtube.com/pulp_platform](https://www.youtube.com/pulp_platform)




The Space “Renaissance”



Cumulative number of objects launched into space

This includes satellites, probes, landers, crewed spacecrafts, and space station flight elements launched into Earth orbit or beyond.



Data source: United Nations Office for Outer Space Affairs (2025)

OurWorldinData.org/space-exploration-satellites | CC BY

Note: When an object is launched by a country on behalf of another one, it is attributed to the latter.

Our World in Data

Multiple growing markets

- NTN: multiorbit constellations (vLeo, MEO GEO...)
- Earth Observation (EO) and Geospatial Data
- Space-Based Defense and Security
- ...and more (e.g. data-centers in space, space manufacturing, etc.)

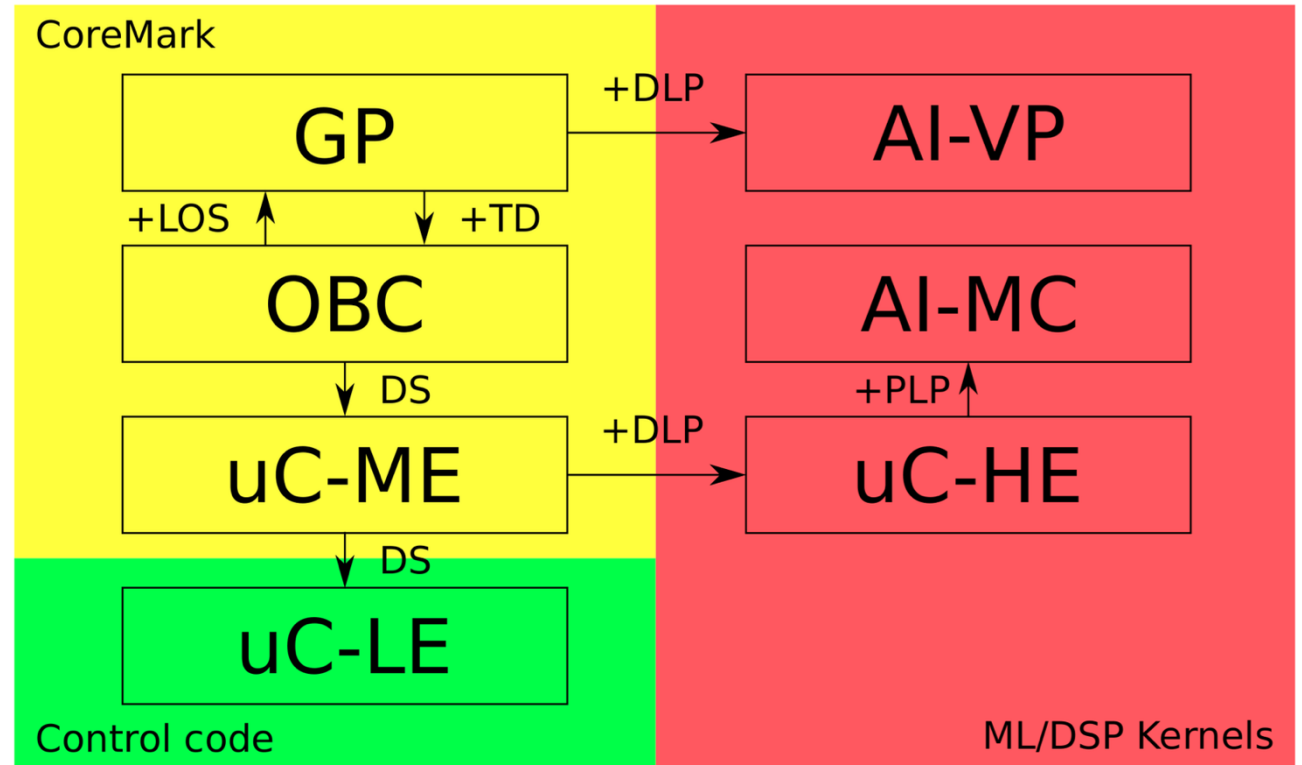
AI is key in all applications

- Federated learning and network/spectrum management
- Anomaly detection
- Constellation management, debris and collision avoidance

Space Missions Increasingly Rely on On-Board Processing



- **Not only On Board Computer**
 - Needed for real-time tasks
- **Different levels of microcontrollers**
 - Low/Mid/High end cores
 - For data acquisition/processing tasks
- **General Purpose processors**
 - To orchestrate complex systems
 - Running Linux-like Oses
 - More relaxed real-time tasks
- **And more and more ML/DSP cores**
 - For compute intensive AI applications



G. Furano, S. Di Mascio, A. Menicucci and C. Monteleone, "A European Roadmap to Leverage RISC-V in Space Applications," 2022 IEEE Aerospace Conference (AERO), Big Sky, MT, USA, 2022, pp. 1-7, doi: 10.1109/AERO53065.2022.9843361

Large variation of processor solutions are needed

But how is RISC-V going to help?



- **RISC-V THE open ISA**
 - Originally developed at UC Berkeley
 - Managed by RISC-V international since 2015 (4.5K members, global!)
- **Simple Base ISA (RV32 / RV64 / RV128)**
 - Extensions to cover many aspects (vector, matrix..)
- **Open development**
 - Technical working groups where members discuss and propose new extensions
 - Public review and comments, ratified by the Board of Directors
 - Eases sovereignty and export control concerns
- **Allows processors to be designed and extended easily**
 - While allowing a common SW infrastructure to be built around it.

The ISA is open, implementations can be open or proprietary

And for AI?

It's the software → **flexibility** key for fast evolution!

Need an **open standard** to counter a monopoly



RISC-V: The Free and Open RISC Instruction Set Architecture

Meta



tenstorrent

Key aspect of RISC-V: space for ISA Extensions



- RISC-V has Reserved opcodes for standard extensions
- Rest of opcodes free for custom implementations
- Custom extensions can be standardized
 - Standard extensions will be frozen/not change in the future

inst[4:2]	000	001	010	011	100	101	110	111
inst[6:5]								(> 32b)
00	LOAD	LOAD-FP	<i>custom-0</i>	MISC-MEM	OP-IMM	AUIPC	OP-IMM-32	48b
01	STORE	STORE-FP	<i>custom-1</i>	AMO	OP	LUI	OP-32	64b
10	MADD	MSUB	NMSUB	NMADD	OP-FP	<i>reserved</i>	<i>custom-2/rv128</i>	48b
11	BRANCH	JALR	<i>reserved</i>	JAL	SYSTEM	<i>reserved</i>	<i>custom-3/rv128</i>	≥ 80b

Extensibility is fundamental in the RISC-V ISA!

Extensions at work: Achieving ~100% dotp Unit Utilization



8-bit Convolution

- HW Loop
- LD/ST with post increment
- 8-bit SIMD sdotp
- 8-bit sdotp + LD

N

```

RV32IMC
addi a0,a0,1
addi t1,t1,1
addi t3,t3,1
addi t4,t4,1
lbu a7,-1(a0)
lbu a6,-1(t4)
lbu a5,-1(t3)
lbu t5,-1(t1)
mul s1,a7,a6
mul a7,a7,a5
add s0,s0,s1
mul a6,a6,t5
add t0,t0,a7
mul a5,a5,t5
add t2,t2,a6
add t6,t6,a5
bne s5,a0,1c000bc
    
```

RV32IMCXpulp

N/4

```

lp.setup
p.lw w1, 4(a0!)
p.lw w2, 4(a1!)
p.lw x1, 4(a2!)
p.lw x2, 4(a3!)
pv.sdotsp.b s1, w1, x1
pv.sdotsp.b s2, w1, x2
pv.sdotsp.b s3, w2, x1
pv.sdotsp.b s4, w2, x2
end
    
```

can we remove?

Yes! dotp+ld

N/4

```

Init NN-RF (outside of the loop)
lp.setup
pv.nnsdotup.h s0,ax1,9
pv.nnsdotsp.b s1, aw2, 0
pv.nnsdotsp.b s2, aw4, 2
pv.nnsdotsp.b s3, aw3, 4
pv.nnsdotsp.b s4, ax1, 14
end
    
```

9x less instructions than RV32IMC

14.5x less instructions at an affordable area cost (50%)

RISC-V Enables Domain Specific Architectures (DSAs)



Multiple Scales of acceleration

Extensions to processor cores

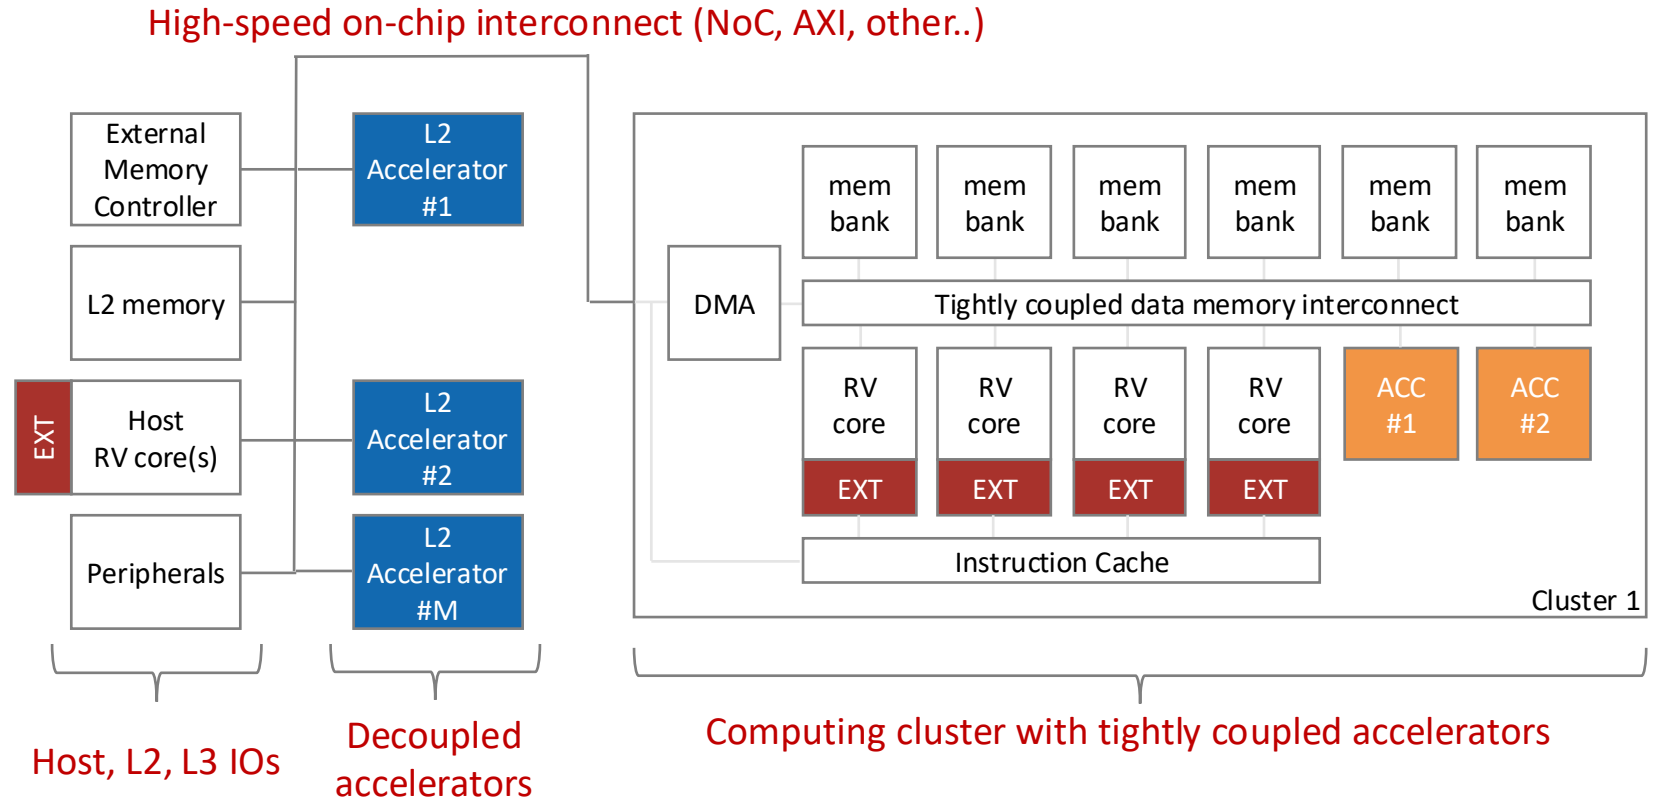
- Explore new extensions
- Efficient implementations

Shared-memory Accelerators

- Domain specific
- Local memory

Multiple Decoupled Accelerators

- Communication
- Synchronization



RISC-V is a key enabler → max agility, enabling SW build-up, without vendor lock-in



Specialization for AI in perspective

Kraken: Using 22FDX tech, NT@0.6V, High utilization, minimal IO & overhead

Energy-Efficient RISC-V Core → **20pJ (8bit)**



ISA-based 10-20x → **1pJ (4bit)**



XPULP



Configurable DP 10-20x → **100fJ (4bit)**



RBE



Highly specialized DP 100x → **1fJ (ternary)**



CUTIE, SNN

PULP has developed an Open Toolbox for DSAs



RISC-V Cores and Vector Units

RI5CY <i>CV32E</i>	Zero R <i>lbex</i>	Snitch	Spatz	Ariane <i>CVA6</i>	ARA
RV32	RV32	RV32	RVV	RV64	RVV


Peripherals

JTAG	SPI
UART	I2S
DMA	GPIO

Interconnects


LIC	HCI
APB	FlooNoC
AXI4	

Platforms



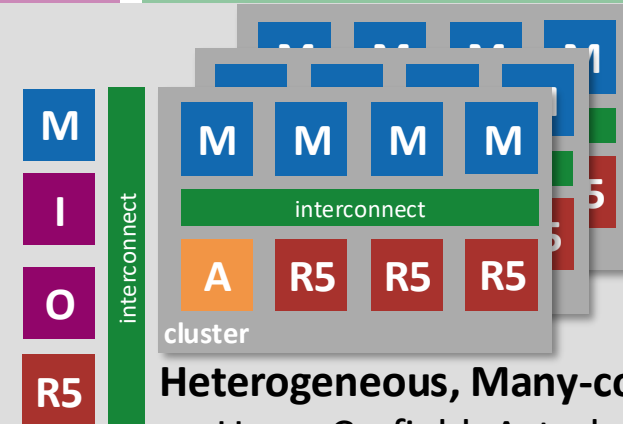
Single core

- PULPino, PULPissimo
- Cheshire



Multi-core

- OpenPULP
- ControlPULP



Heterogeneous, Many-core

- Hero, Carfield, Astral
- Occamy, Mempool

IOT

<https://github.com/pulp-platform>

HPC

Accelerators and ISA extensions

XpulpNN, XpulpTNN	ITA (Transformers)	RBE, NEUREKA (QNNs)	FFT (DSP)	REDMULE (FP-Tensor)
----------------------	-----------------------	------------------------	--------------	------------------------

DSA for Space: Reliable, Safe and Secure (Efficient, Scalable)



Reliable, safe and secure architectures

- Supporting mixed-criticality systems, trade-off between performance and reliability
- Better/faster virtualization support: vCLIC, cache partitioning
- Efficient implementations of RISC-V extensions:
Zicfiss: Control-Flow Integrity Shadow Stack, **Zicfilp**: Landing Pads

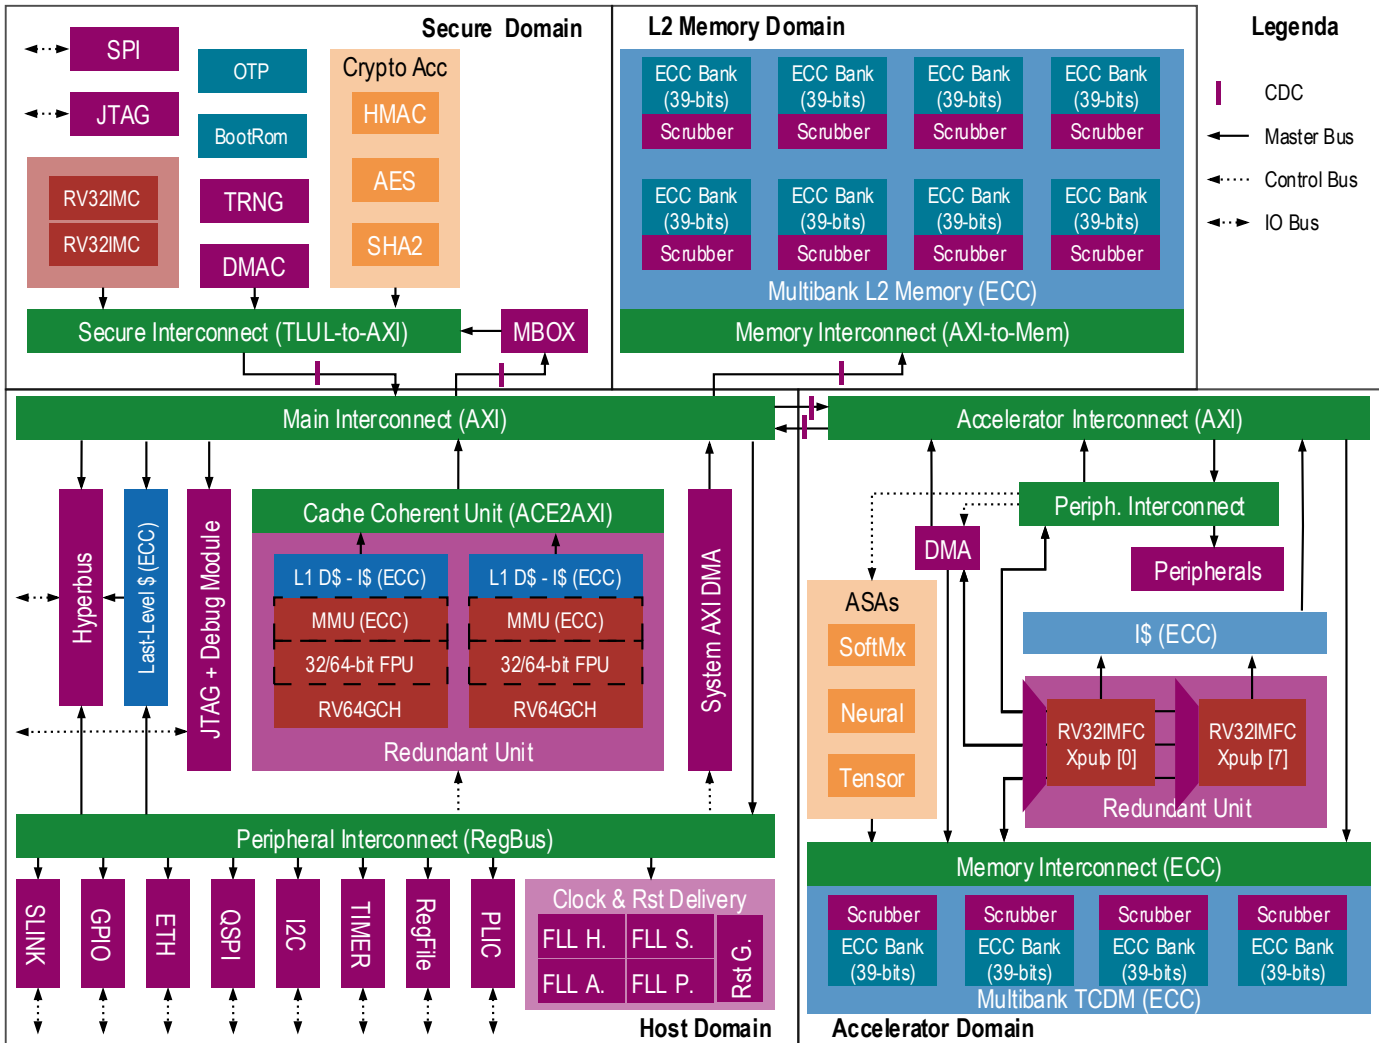
Efficient computing

- Support for various data types
- Vector units for different applications
- Heterogeneous computing, adding configurable accelerators

Scaling up compute to 100s and 1000s of cores

- Data transport solutions: NoC, working with sparse data, transforming data in transit


Astral: PULP's DSA for Space



- Fully open RV platform for AI in space
- Reliable → configurable modular redundancy + ECC and real-time fault recovery (< 100 clock cycles)
- Safe, Secure → secure domain
- Efficient AI → 14.5x and 9x acceleration on FP/INT AI workloads
- Flexible → Linux OS support with FPGA deployment flow
- Scalable → designed for modular scale-up



Astral: PULP's DSA for Space



Technology	Global Foundries 12 nm LP+
Area	5 mm ²
Power	380 mW (Host) 350 mW (Accelerator)
Frequency	780 MHz (Host) 620 MHz (Accelerator)
INT Performance	578 GOPS
INT Energy Efficiency	4.7 TOPS/W
FP Performance	58 GFLOPS
FP Energy Efficiency	557 GFLOPS/W
E2End Appl. Latency	17.5 ms (CloudViT) 1.5 ms (Reliable An. Det.)
Fault Recovery Latency	38.6 ns



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Chips-it
FONDAZIONE

ThalesAlenia
a Thales / Leonardo company
Space

ETH zürich

SpaceX Transporter-13 launch, March 15th 2025



Carried our first PULP chip
Trikarenos to space



Aboard the ALICE
experiment by ARIS





<http://pulp-platform.org>



@pulp_platform

Looking forward to more space missions



<https://open-source-chips.eu/>