ETH zürich

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Leveraging the PULP Platform to build Reliable Multicore SoCs for Space Applications

Michael Rogenmoser michaero@iis.ee.ethz.ch
Yvan Tortorella yvan.tortorella@unibo.it
and the PULP team

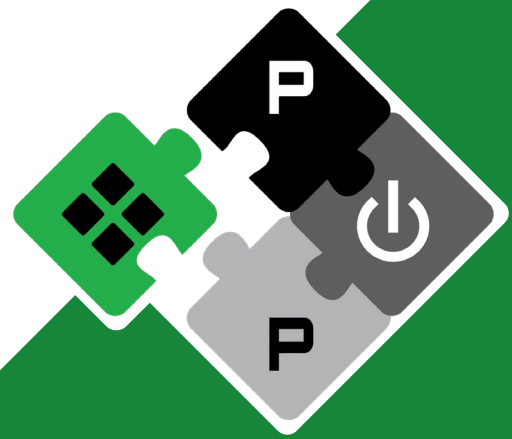

@pulp_platform

pulp-platform.org

youtube.com/pulp_platform

PULP Platform
Open Source Hardware, the way it should be!

Overview

1. The PULP Platform

2. Leveraging the PULP cluster for a Fault-Tolerant System

3. Using Heterogeneity for High-Performance Systems

Who are we and who is behind PULP?

Michael

Yvan

ETH zürich

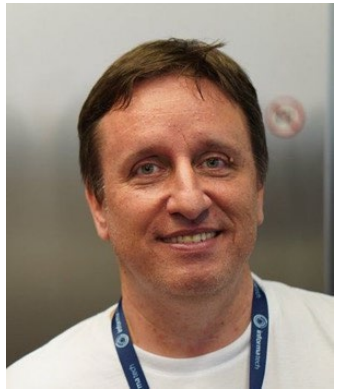

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Dr. Frank Gurkaynak Prof. Luca Benini Prof. Davide Rossi

In total about 90 people in Zurich and Bologna

https://pulp-platform.org/team.html

PULP: Energy-Efficient Computing from Core to Platform

- Core
  - Improving core efficiency with ISA and uAch extensions
- Cluster
  - Efficient shared-mem cluster
  - From a few to thousand processing element
- Full platform
  - Heterogeneity: host, processor, accelerator
  - One or multiple accelerators
  - From chips to chiplets (2D to 3D)

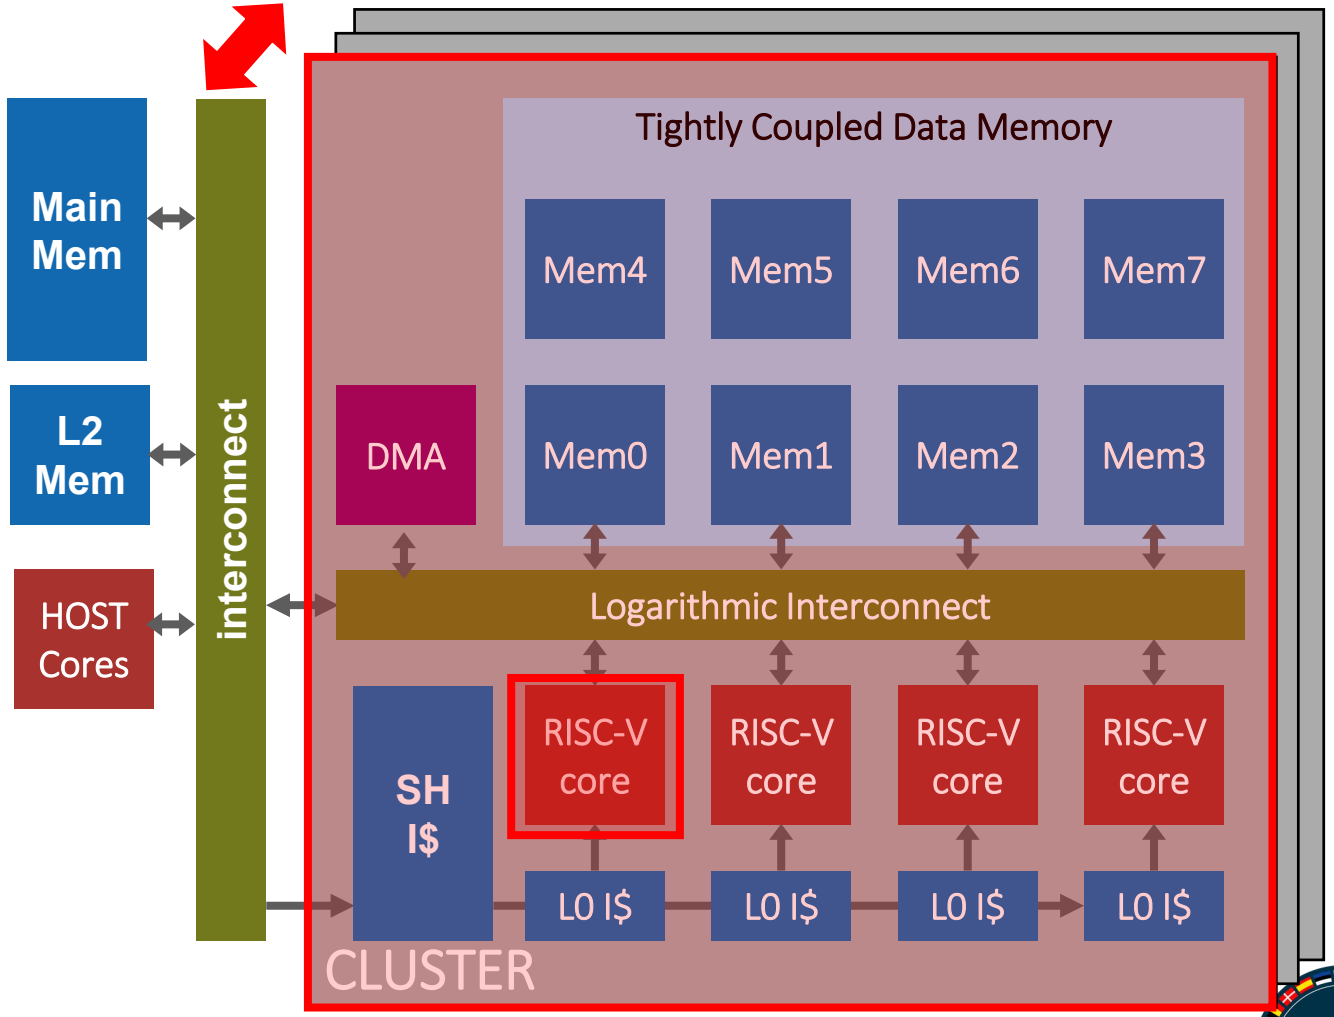

PULP uses a **permissive** open source license

- All our development is on GitHub
  - HDL source code, testbenches, software development kit, virtual platform

https://github.com/pulp-platform

- Allows anyone to use, change, and make products without restrictions.

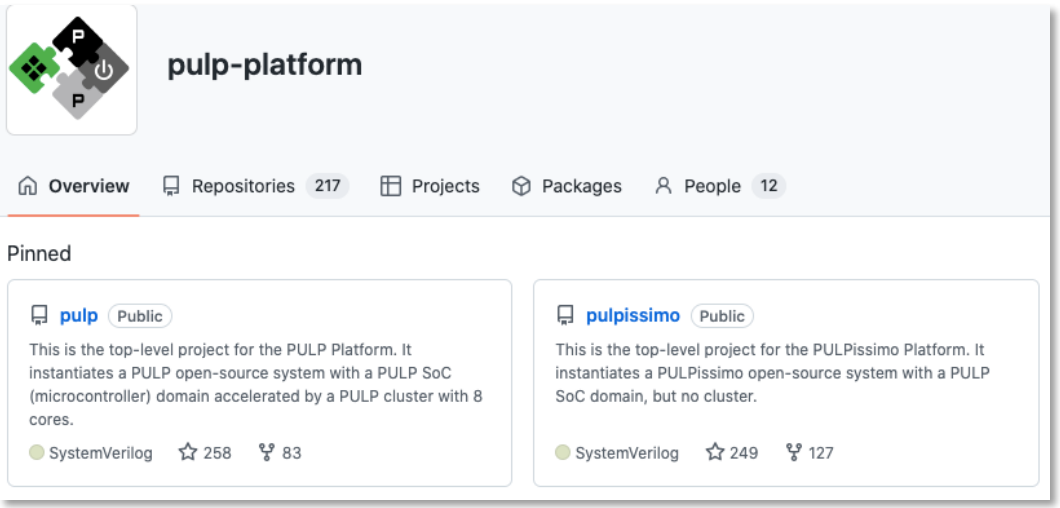

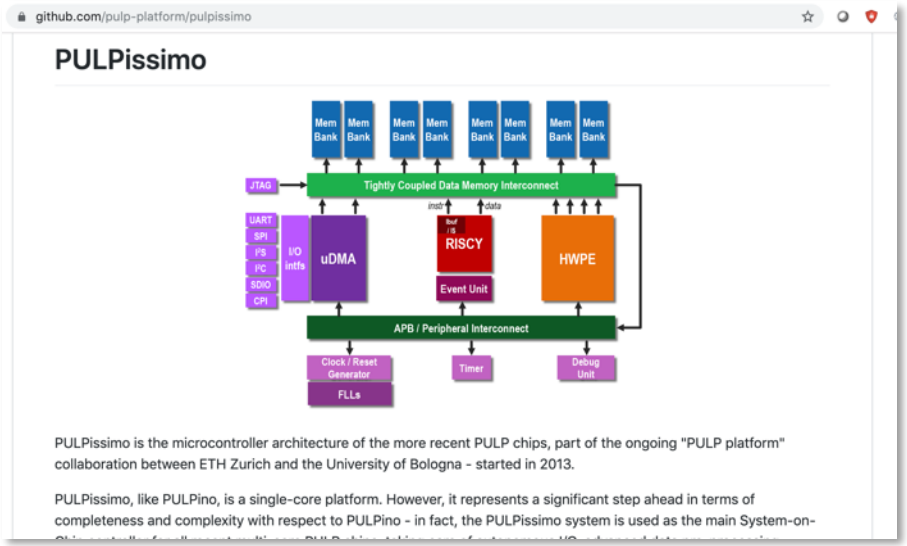

CV32E40P –aka RI5CY – RISC-V Core for edge AI

3-cycle ALU-OP, 4-cyle MEM-OP→IPC loss: LD-use, Branch

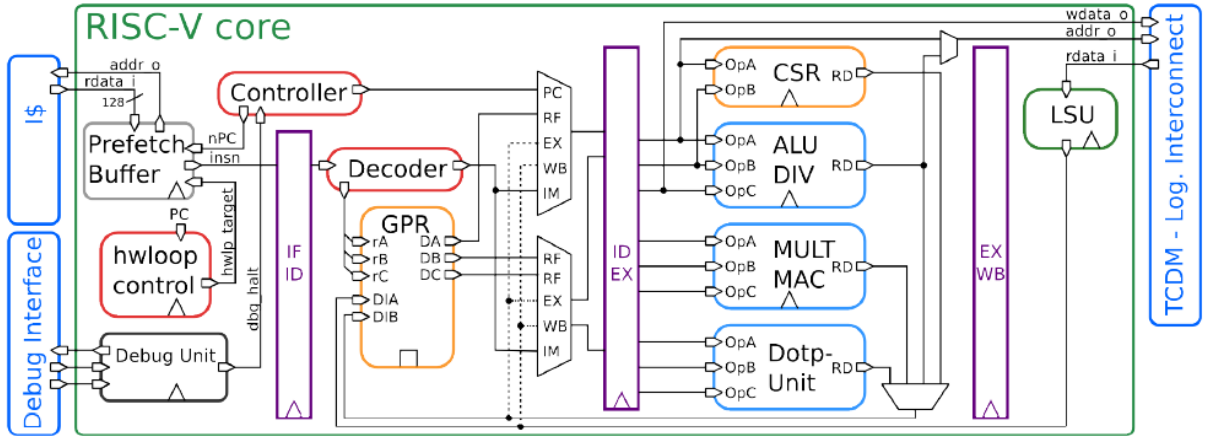

RISC-V

ISA is extensible *by construction* (great!)

```
for (i = 0; i < 100; i++)
    d[i] = a[i] + b[i];
```

Baseline 11 cyc./out

```
mv    x5, 0
mv    x4, 100
Lstart:
  lb     x2, 0(x10)
  lb     x3, 0(x11)
  addi   x10,x10, 1
  addi   x11,x11, 1
  add    x2, x3, x2
  sb     x2, 0(x12)
  addi   x4, x4, -1
  addi   x12,x12, 1
bne      x4, x5, Lstart
```

Auto-incr 8 cyc./out

```
mv    x5, 0
mv    x4, 100
Lstart:
  lb     x2, 0(x10!)
  lb     x3, 0(x11!)
  addi   x4, x4, -1
  add    x2, x3, x2
  sb     x2, 0(x12!)
bne      x4, x5, Lstart
```

HW Loop 5 cyc./out

```
lp.setupi 100, Lend
  lb     x2, 0(x10!)
  lb     x3, 0(x11!)
  add    x2, x3, x2
Lend:  sb x2, 0(x12!)
```

SIMD 1.25 cyc./out

```
lp.setupi 25, Lend
  lw     x2, 0(x10!)
  lw     x3, 0(x11!)
  pv.add.b x2, x3, x2
Lend: sw x2, 0(x12!)
```

Scaling performance: the case for parallelization

- As **VDD** decreases, **operating speed** decreases as well.
- However **efficiency** increases→ more work done per Joule
  - Until leakage effects start to dominate
- Put more units in parallel to get performance up and keep them busy with a parallel workload

[Rossi et al. IEEE Micro 2017]

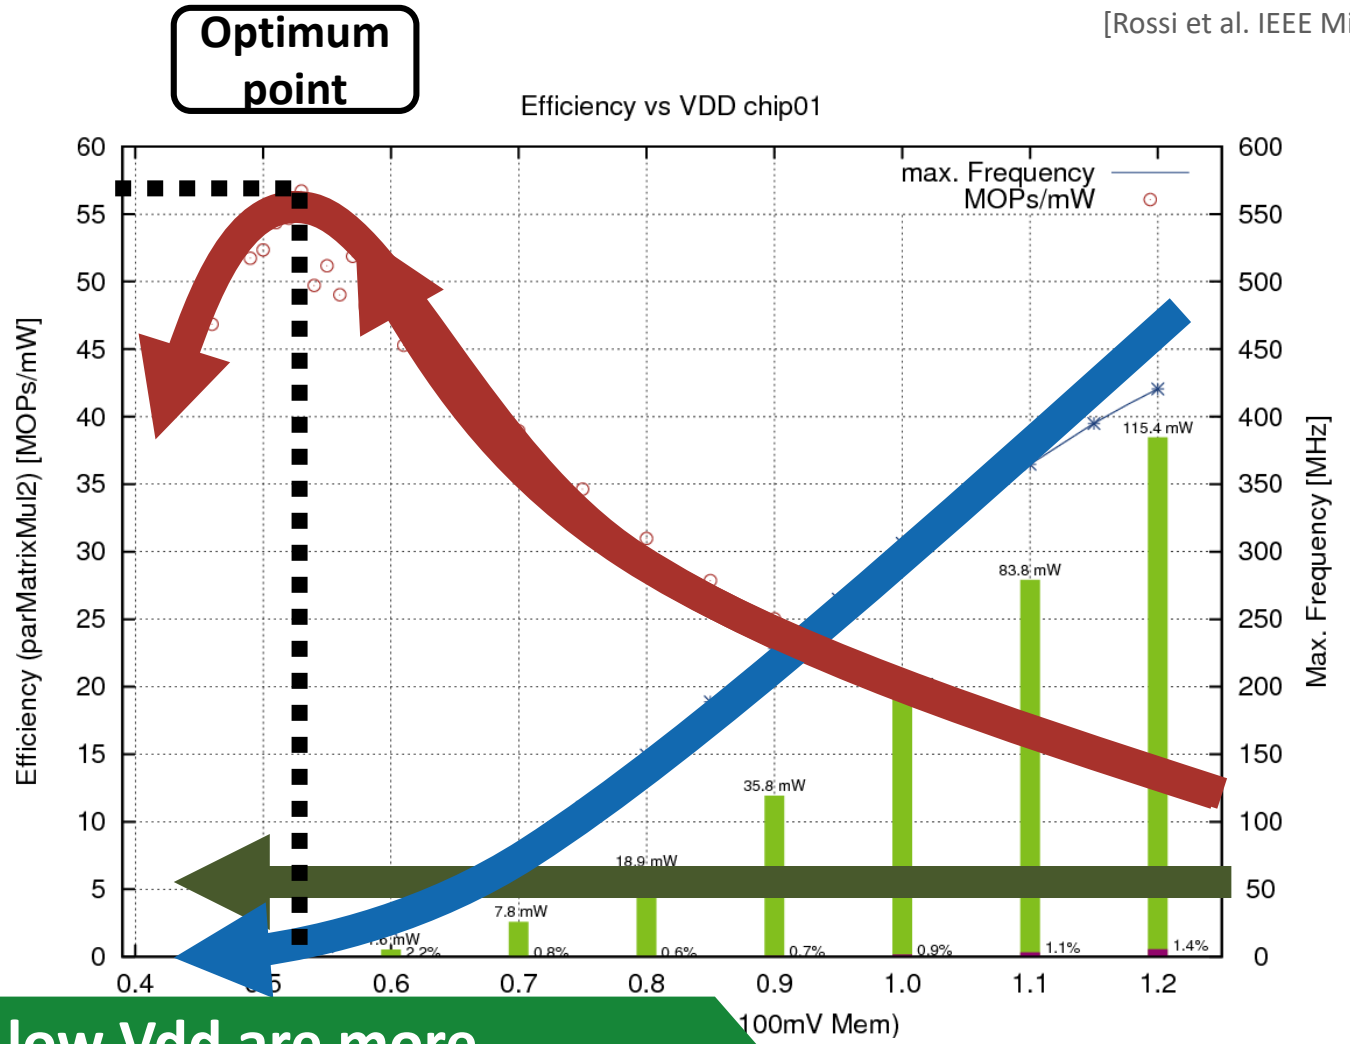

N cores running at moderate f, low Vdd are more energy efficient than a single core at N×f, high Vdd

Let's design a cluster with multiple (4-16) RISC-V cores

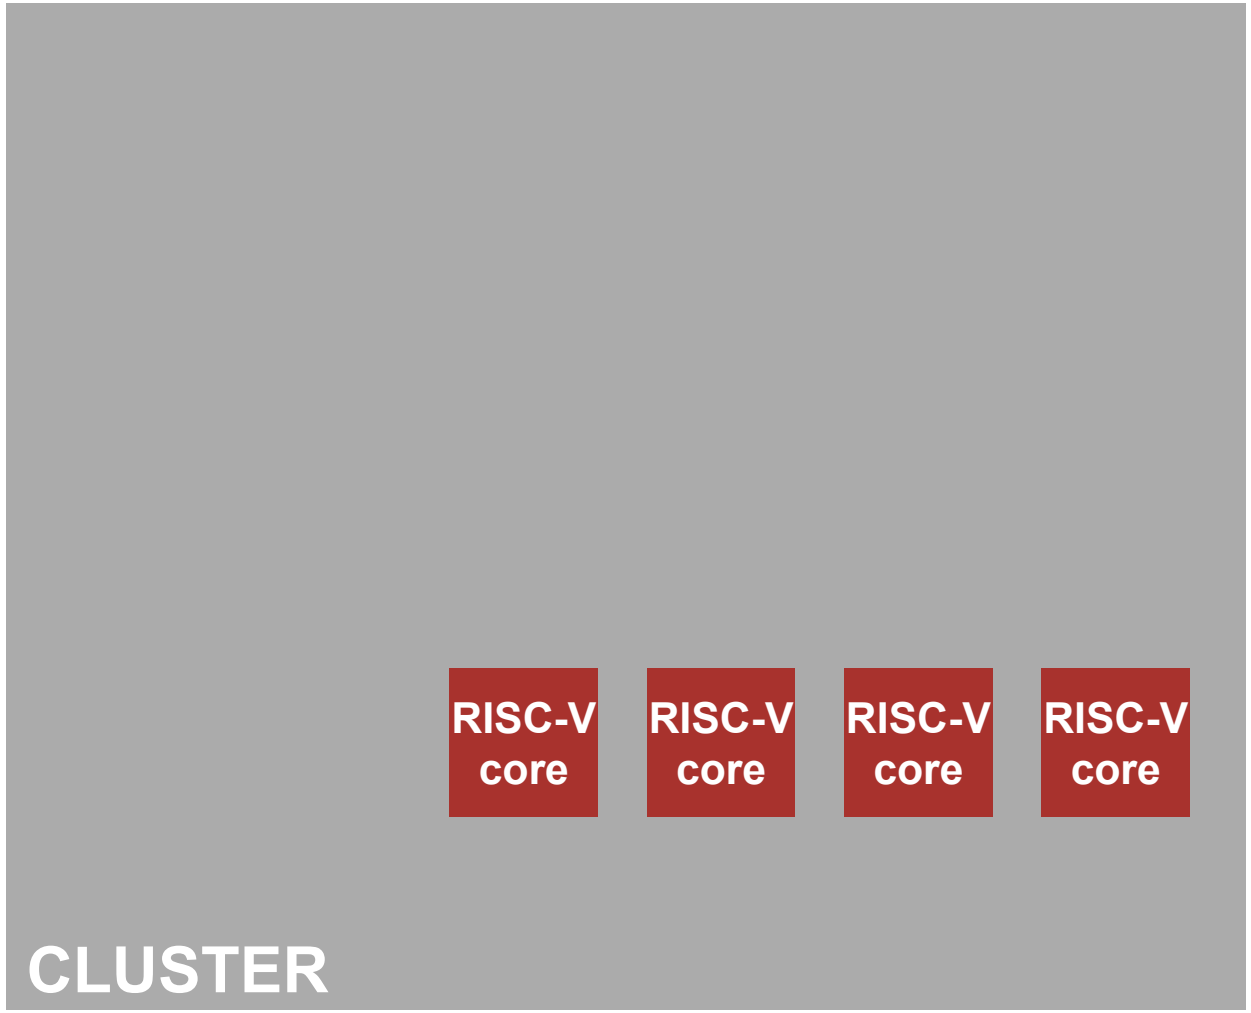

Low-Latency shared Tightly Coupled Data Memory (TCDM)

- Parallel memory access with low contention
  - Multi-banked, address-interleaved L1
- Fast interconnect with physical design awareness
  - Logarithmic depth of combinational switchboxes

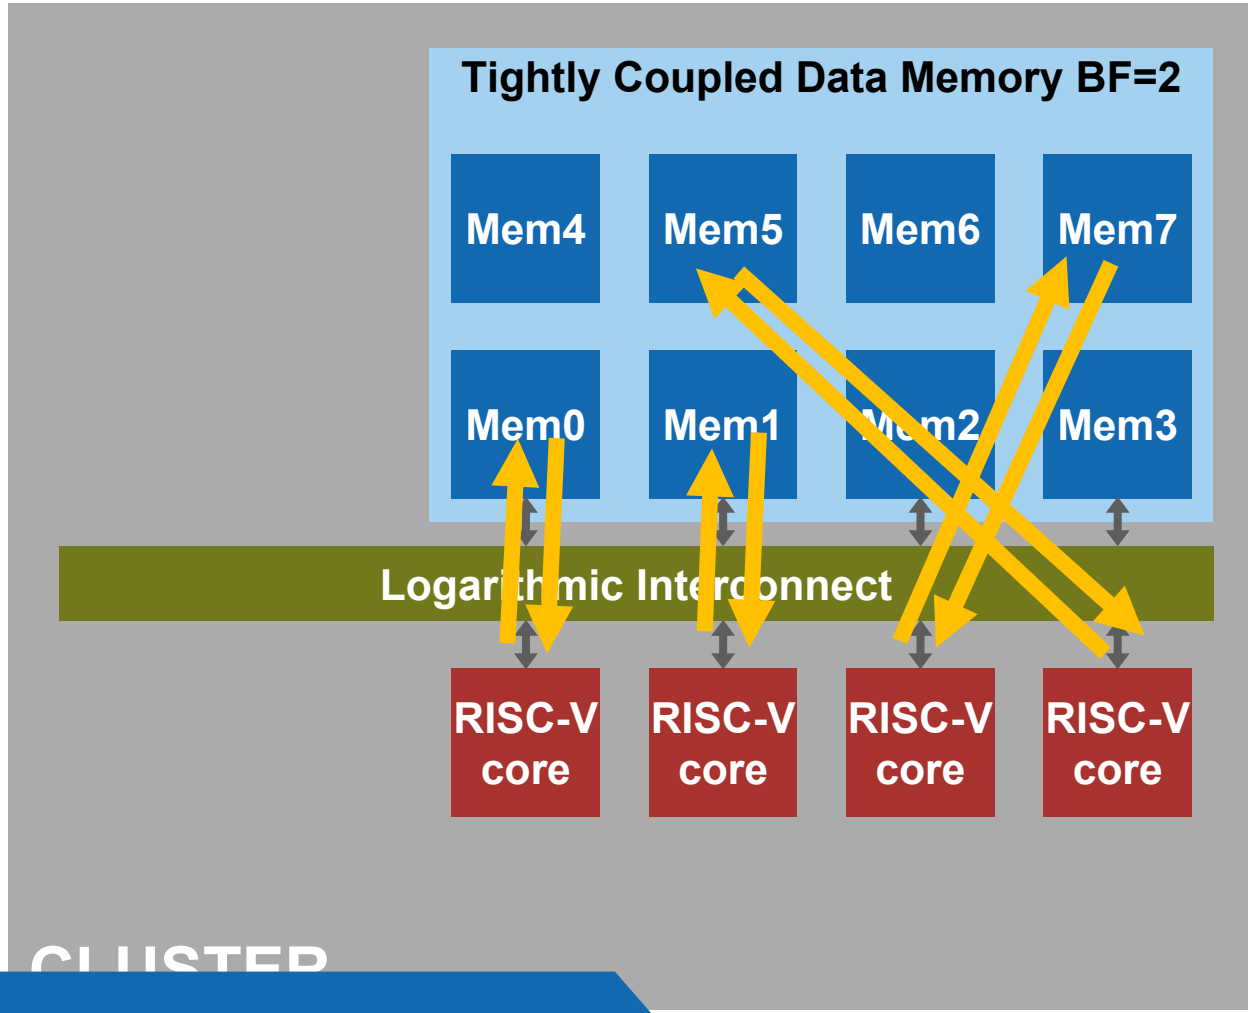

Trade-off between memory size and latency

Full Cluster

- Hierarchical Instruction Cache
  - Speeds up instruction fetching
- DMA
  - Facilitates memory transfers
- Event Unit
  - Facilitates core synchronization
- Ports to a host

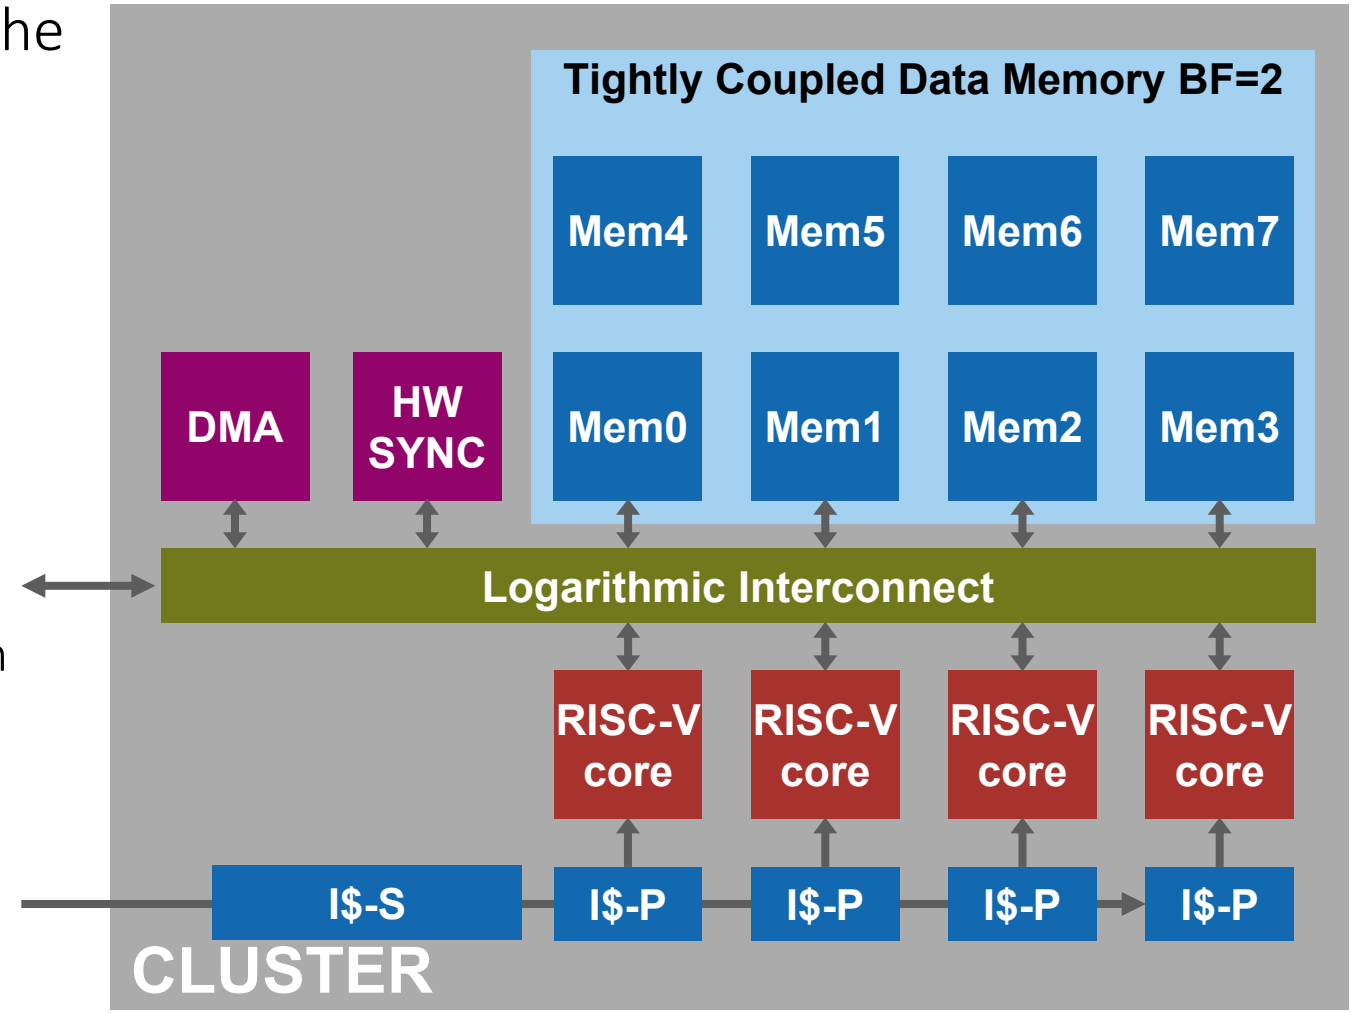

Combining ISA extension + Efficient parallel execution

- 8-bit convolution
  - Open source DNN library
- 10x through xPULP
  - Extensions bring real speedup
- Near-linear speedup
  - Scales well for regular workloads
- 75x overall gain
- 7-8 GMACs
  - 250MHz
  - 4 MAC/Cycle (8bit)

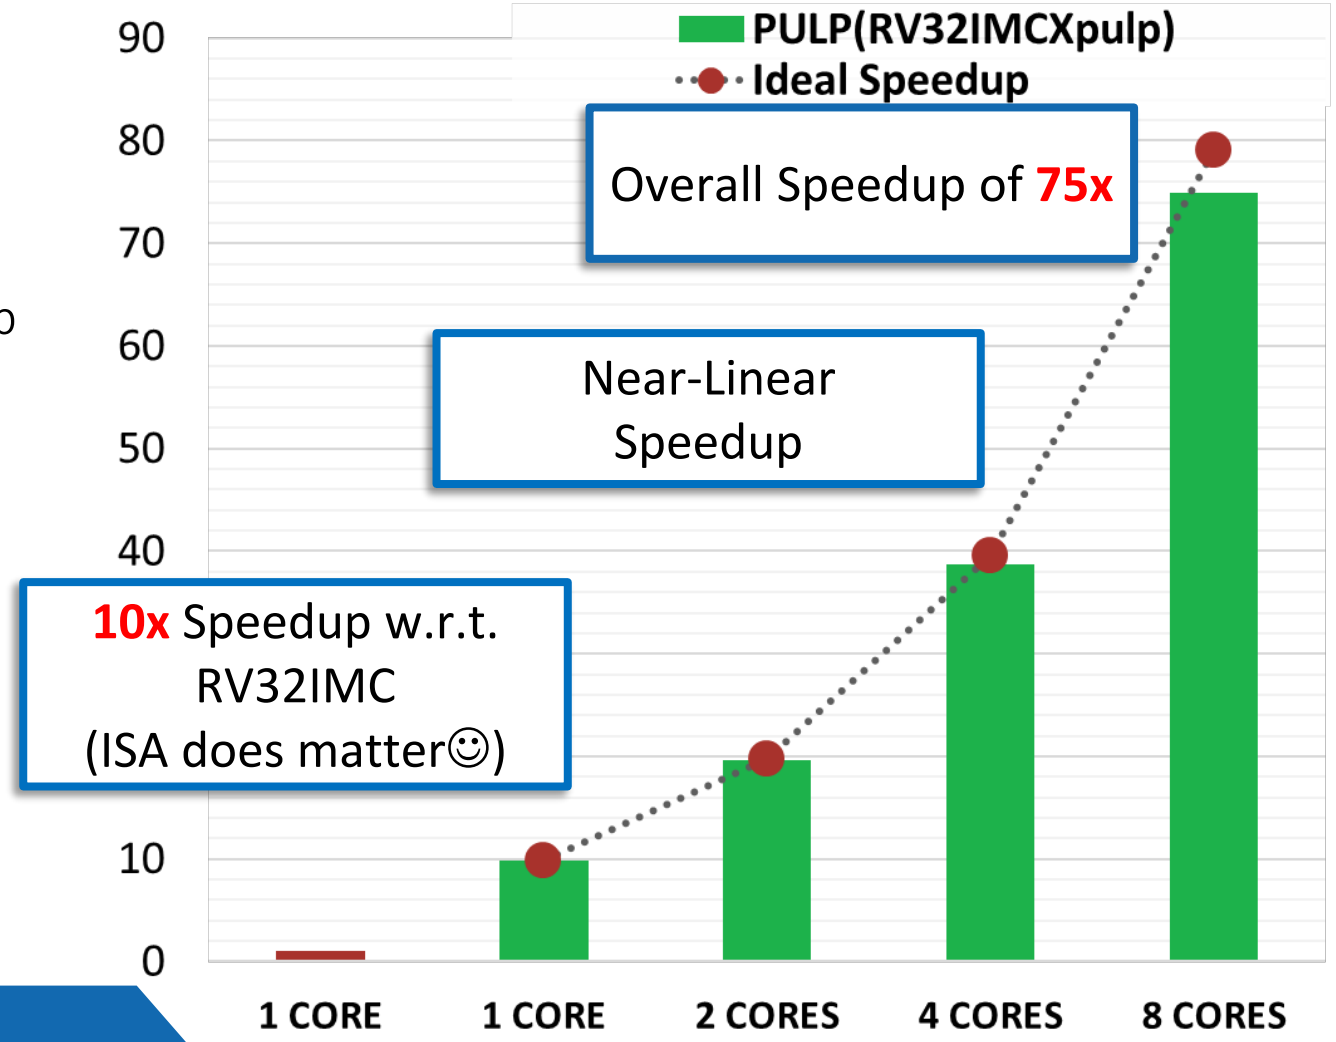

[Garofalo et al. Philos. Trans. R. Soc 20]

More GOPS for less power

Overview

1. The PULP Platform

2. Leveraging the PULP cluster for a Fault-Tolerant System

3. Using Heterogeneity for High-Performance Systems

Towards a space-proof PULP cluster

- Cores
- Memory
- Inteconnect, DMA, EU

Component-based
vs. Architectural approach

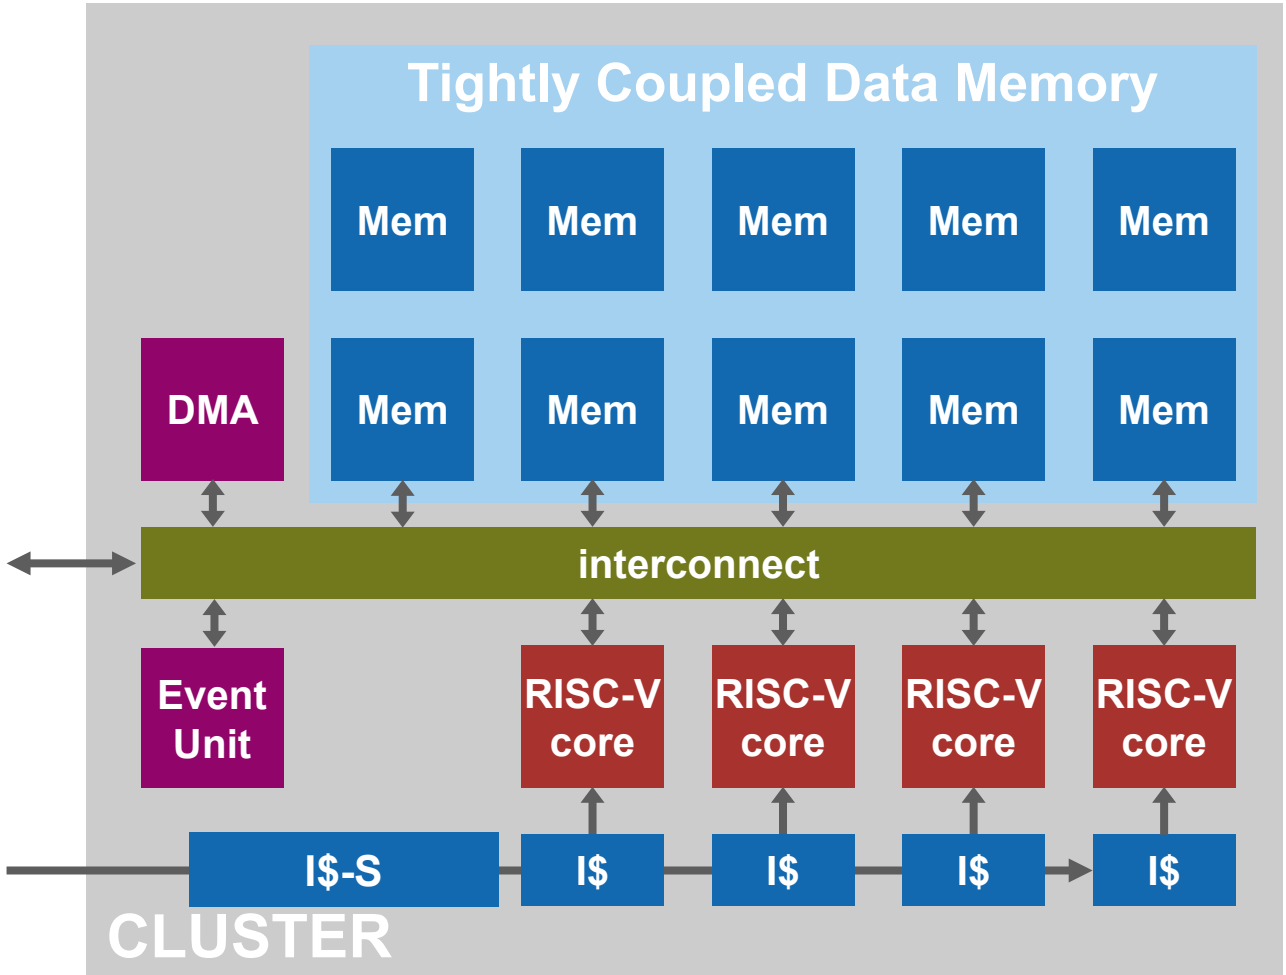

Leveraging Parallelism

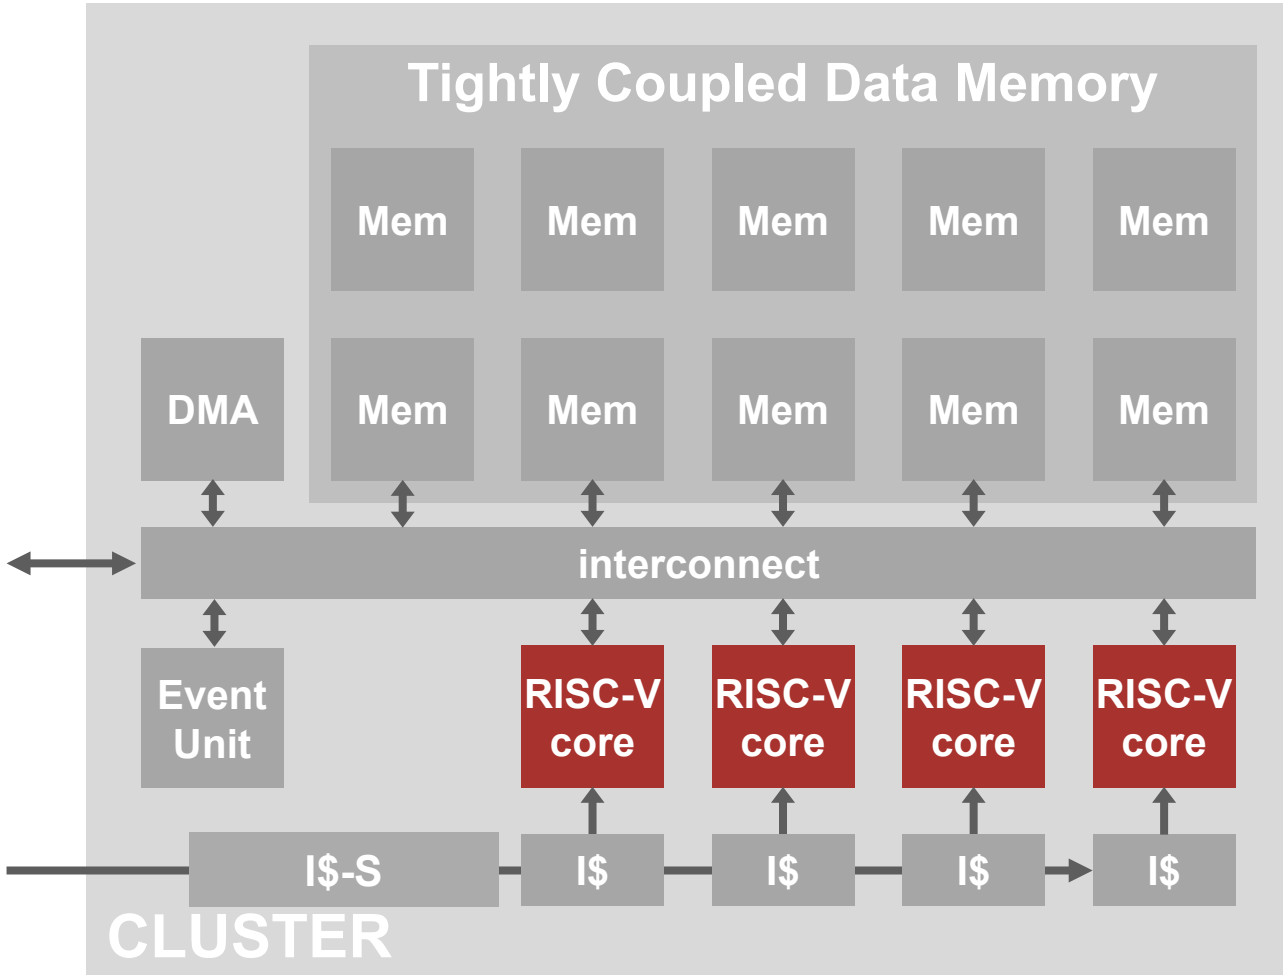

Protecting the Cores

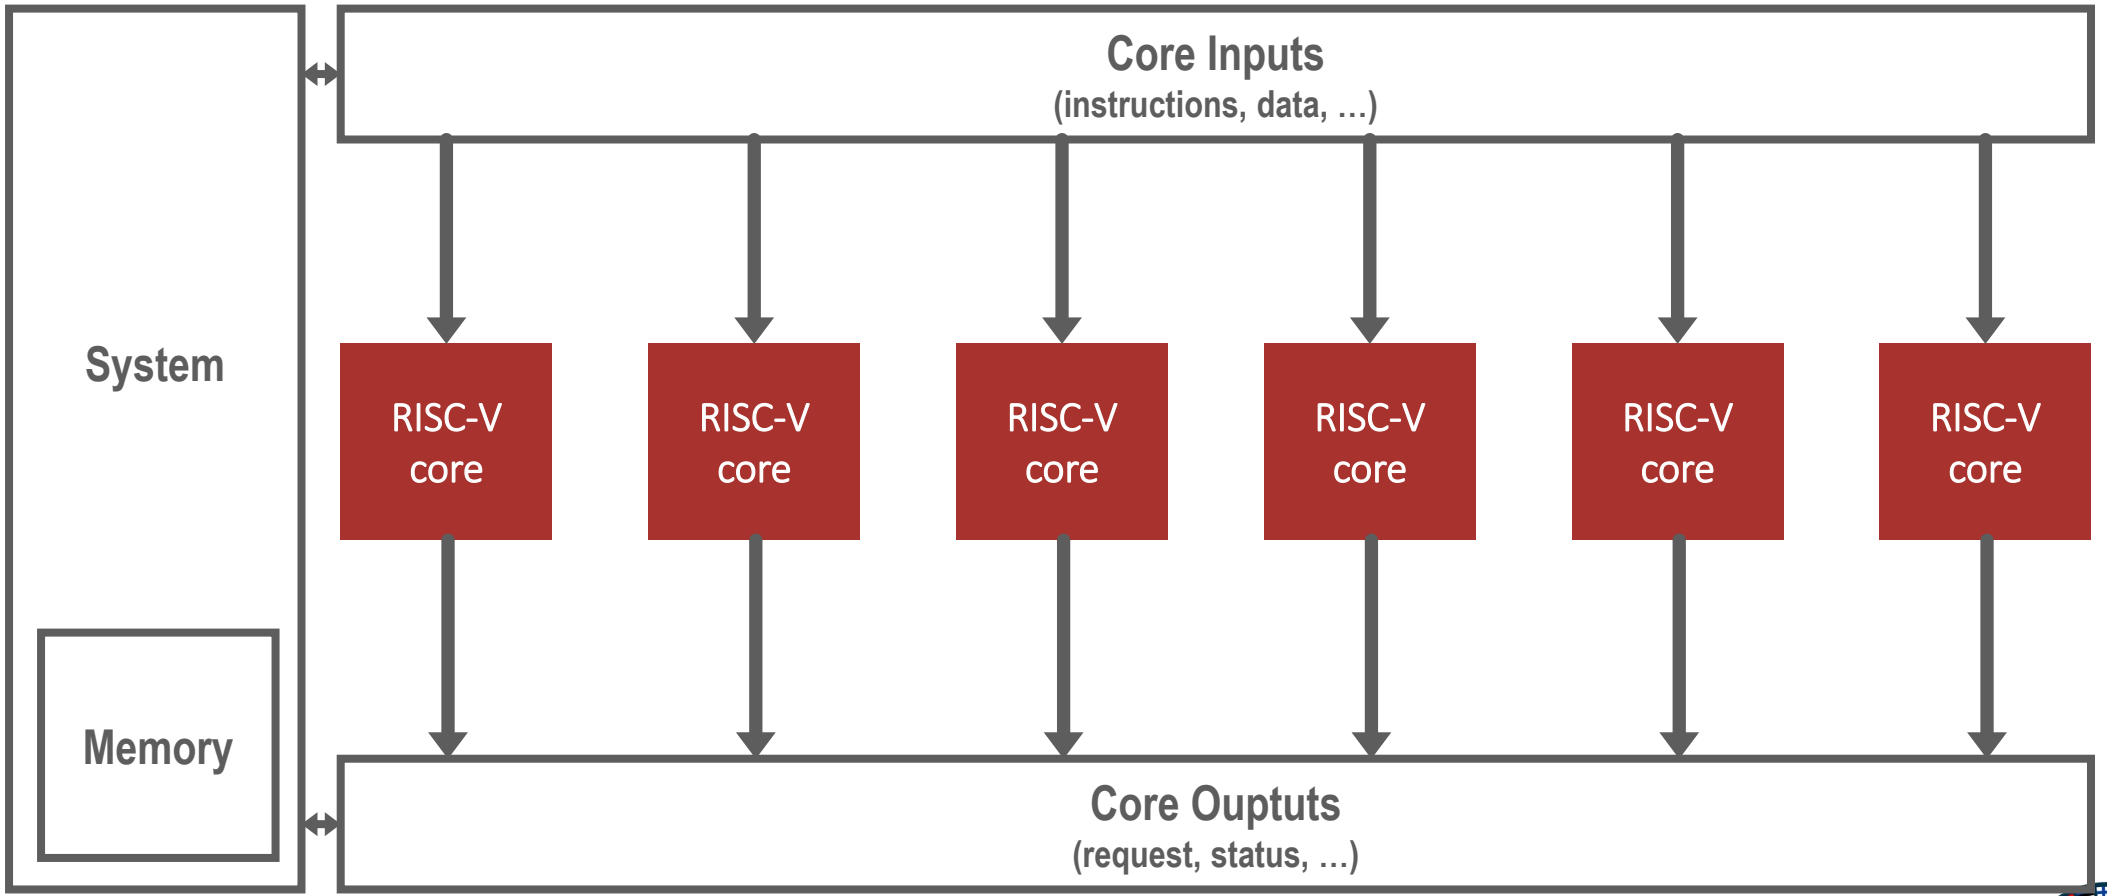

Protecting the Cores

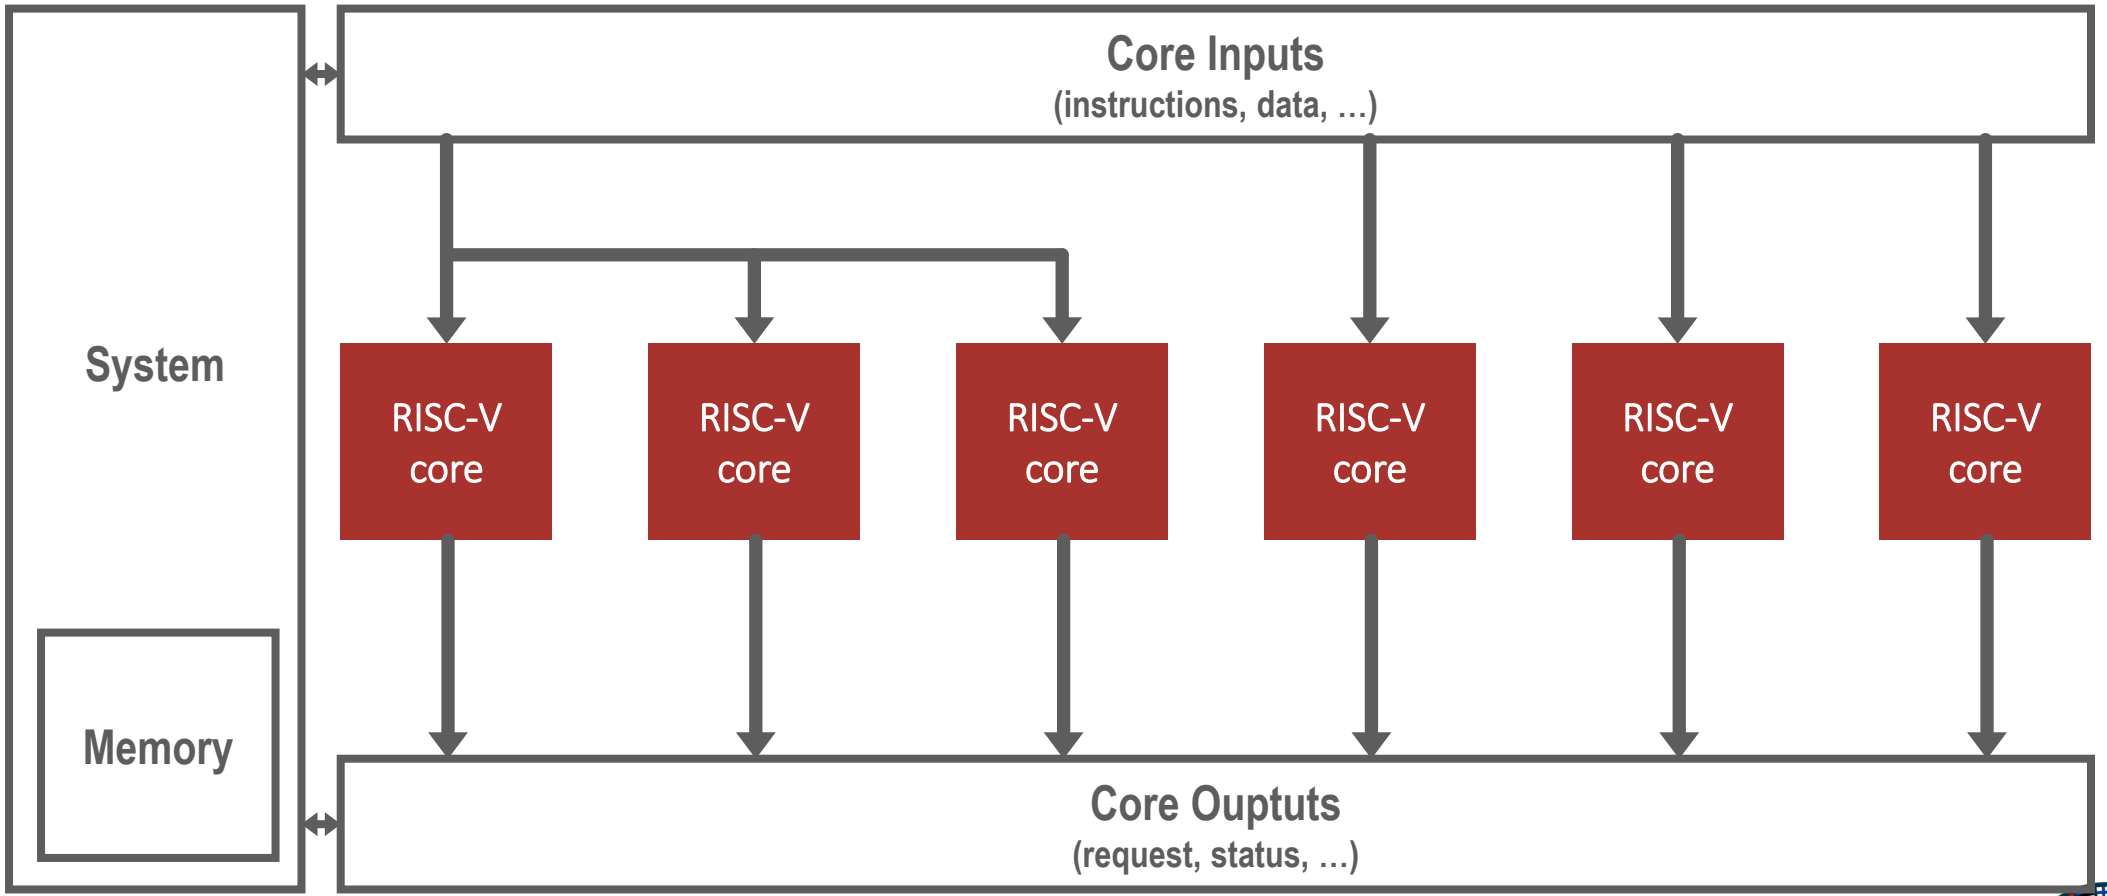

Protecting the Cores

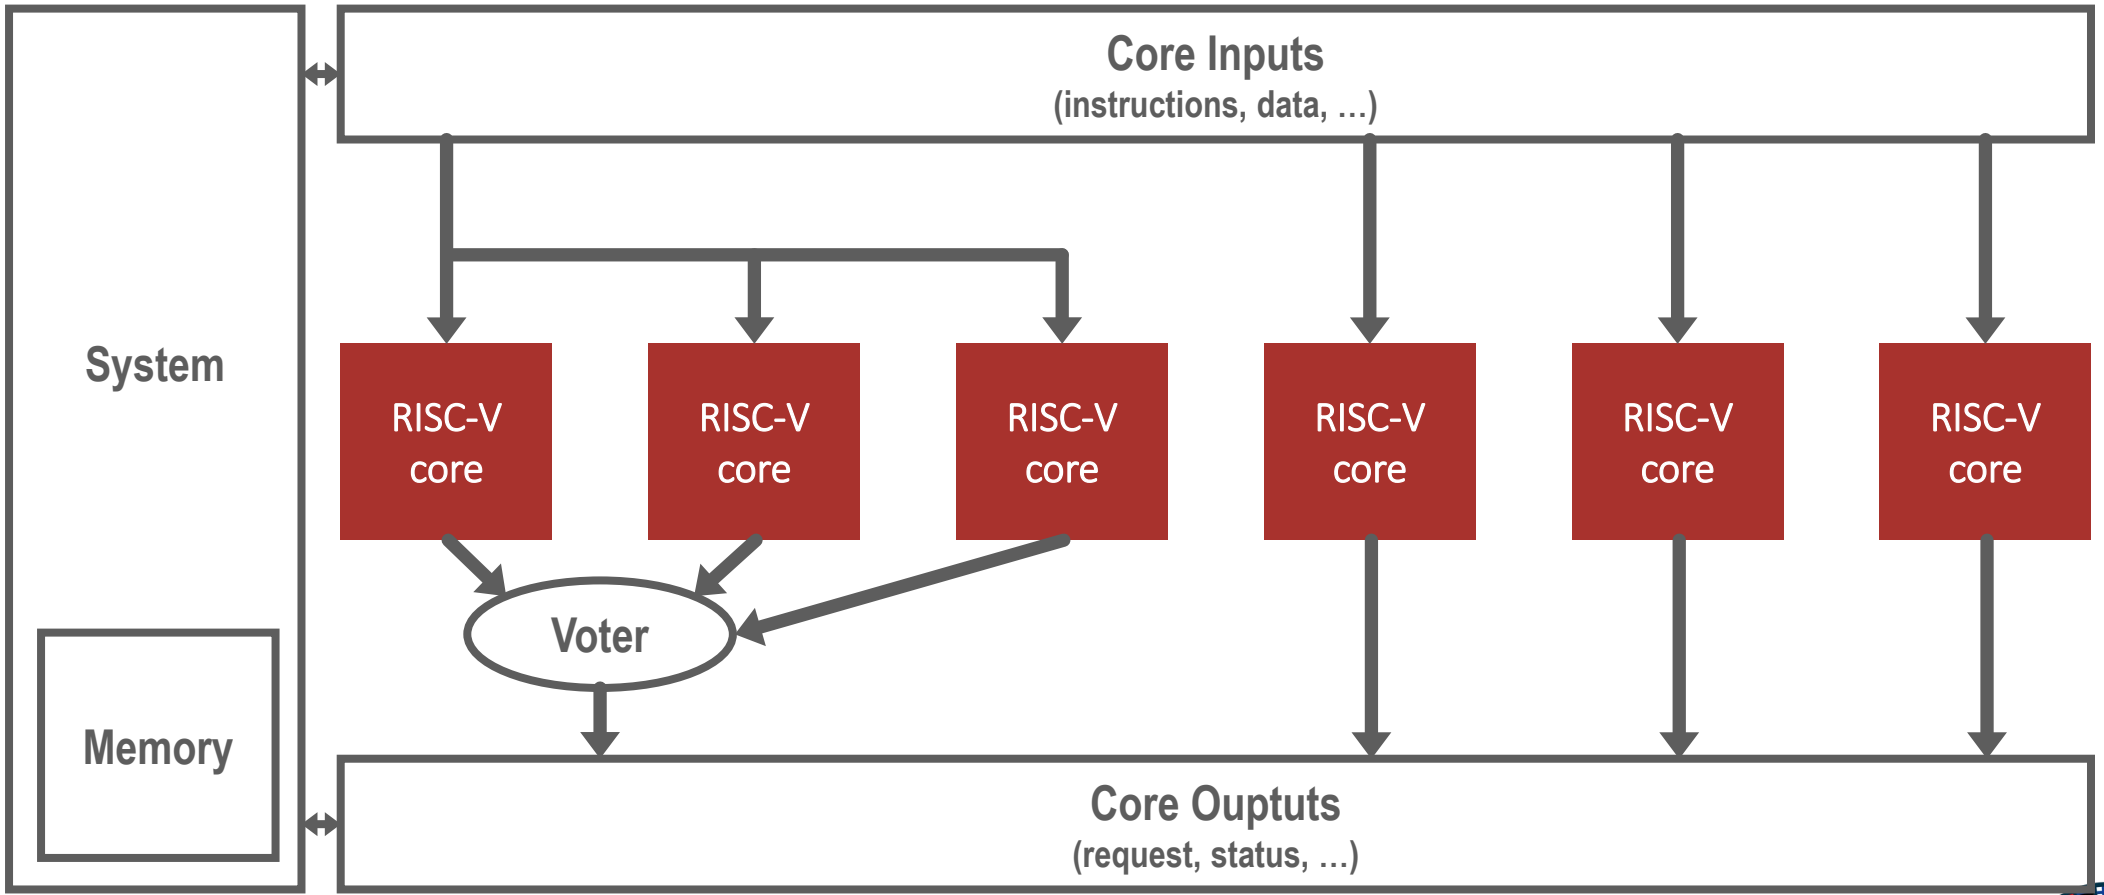

Triple-Core Lock Step

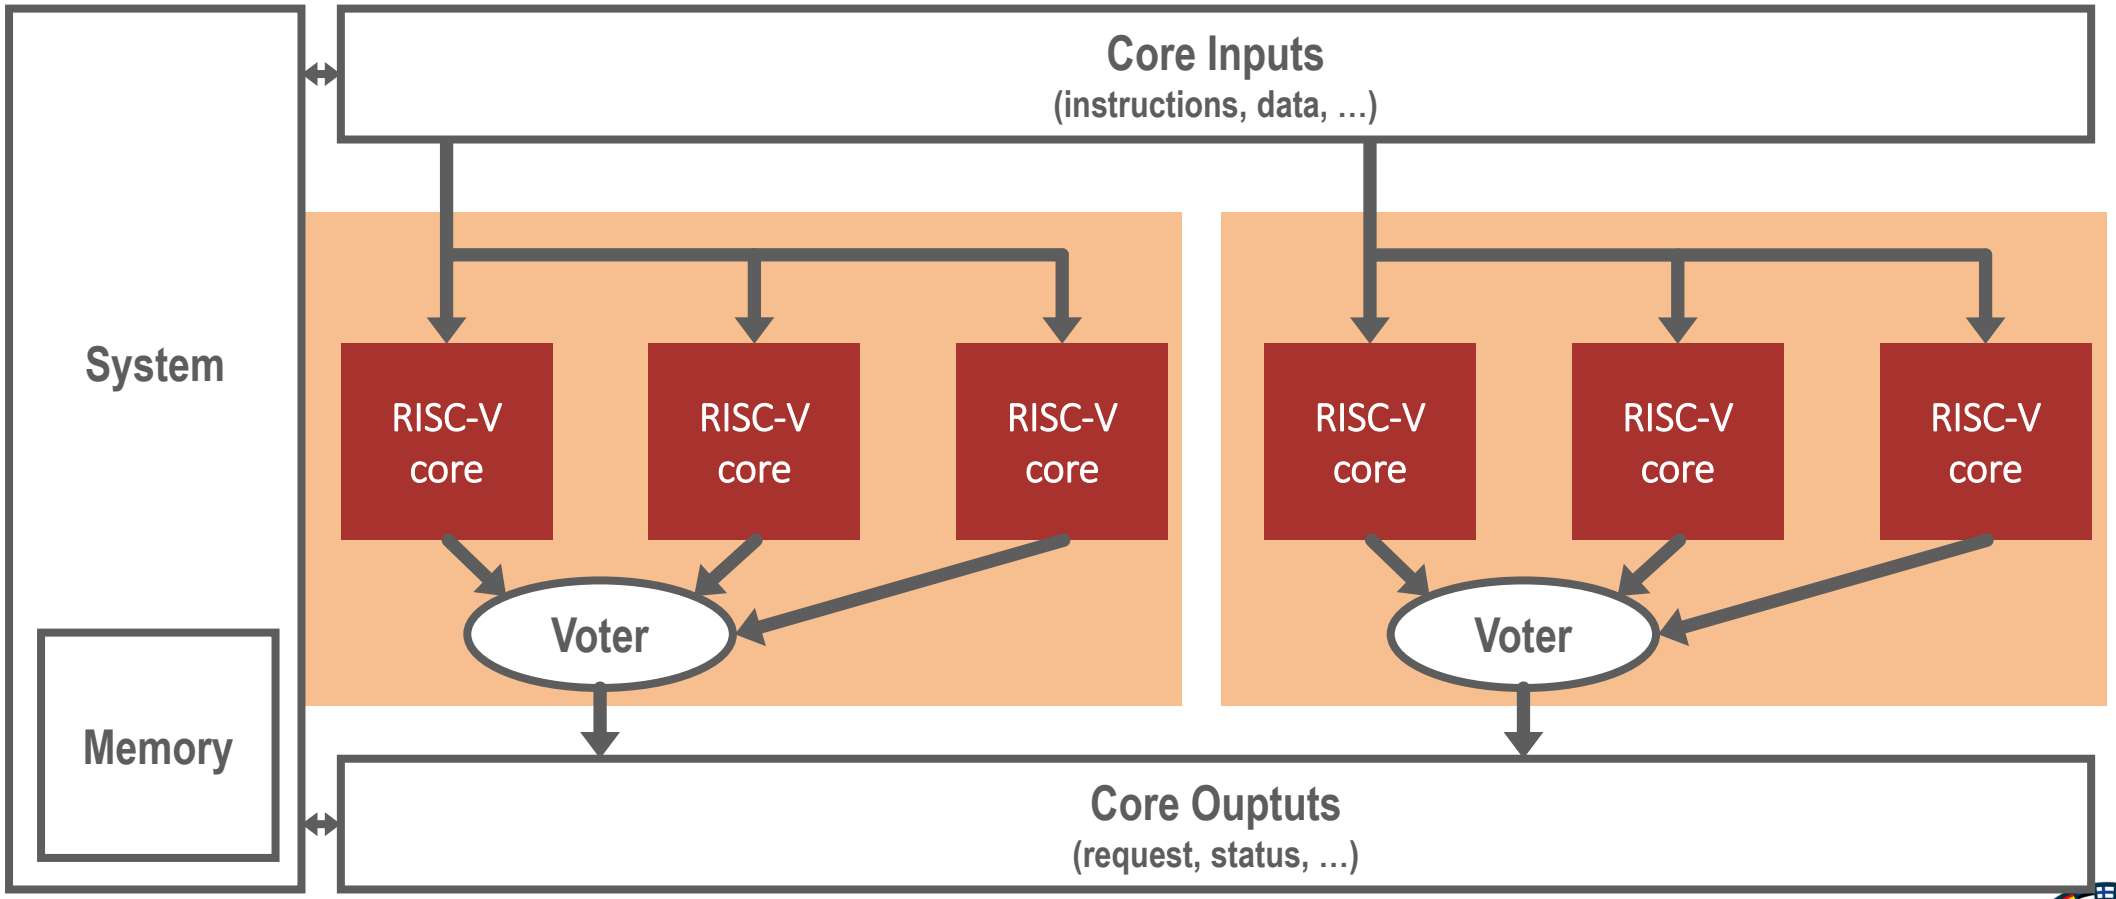

Redundancy Requirement

- Mission-Critical Application
  - Must be correct
- Data Processing Algorithm
  - Individual errors can be tolerated

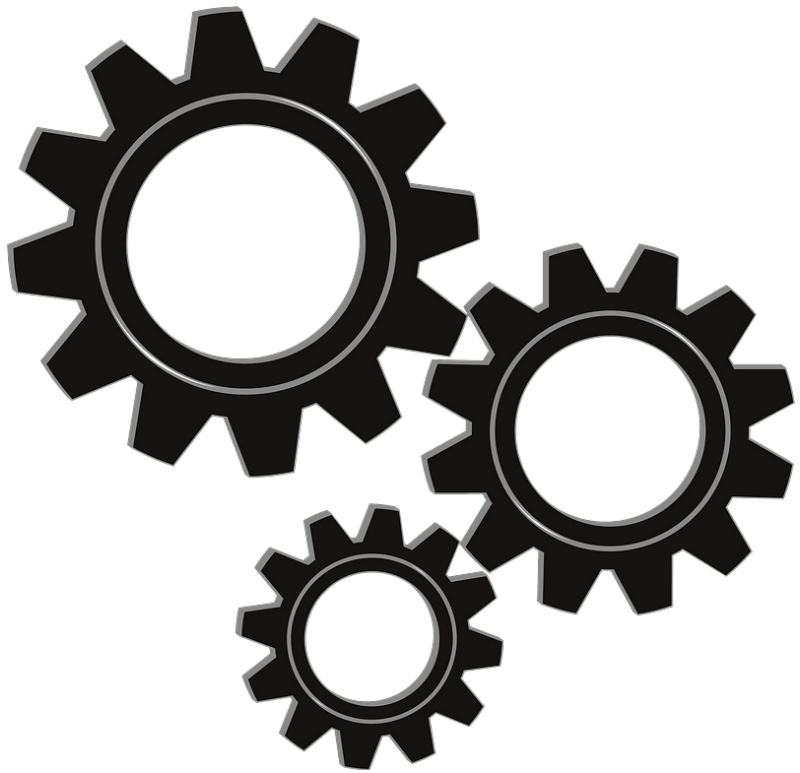

On-Demand Redundancy Grouping

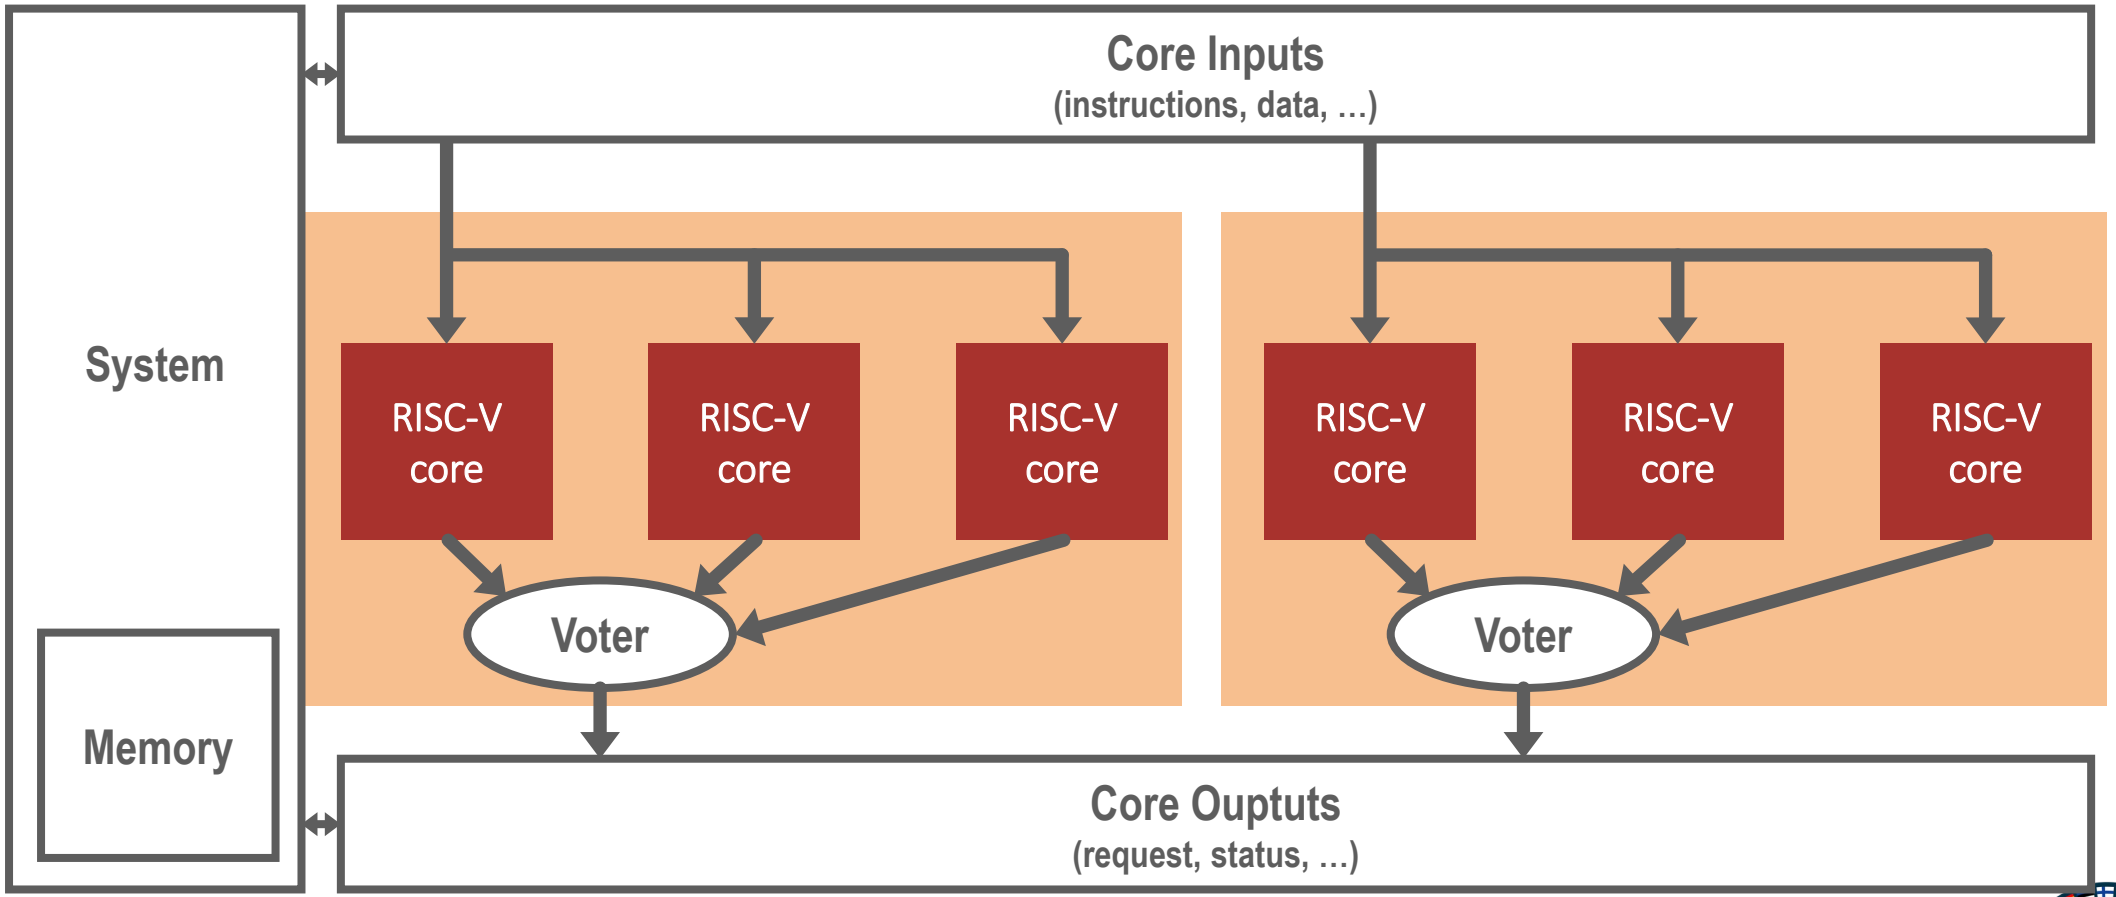

On-Demand Redundancy Grouping

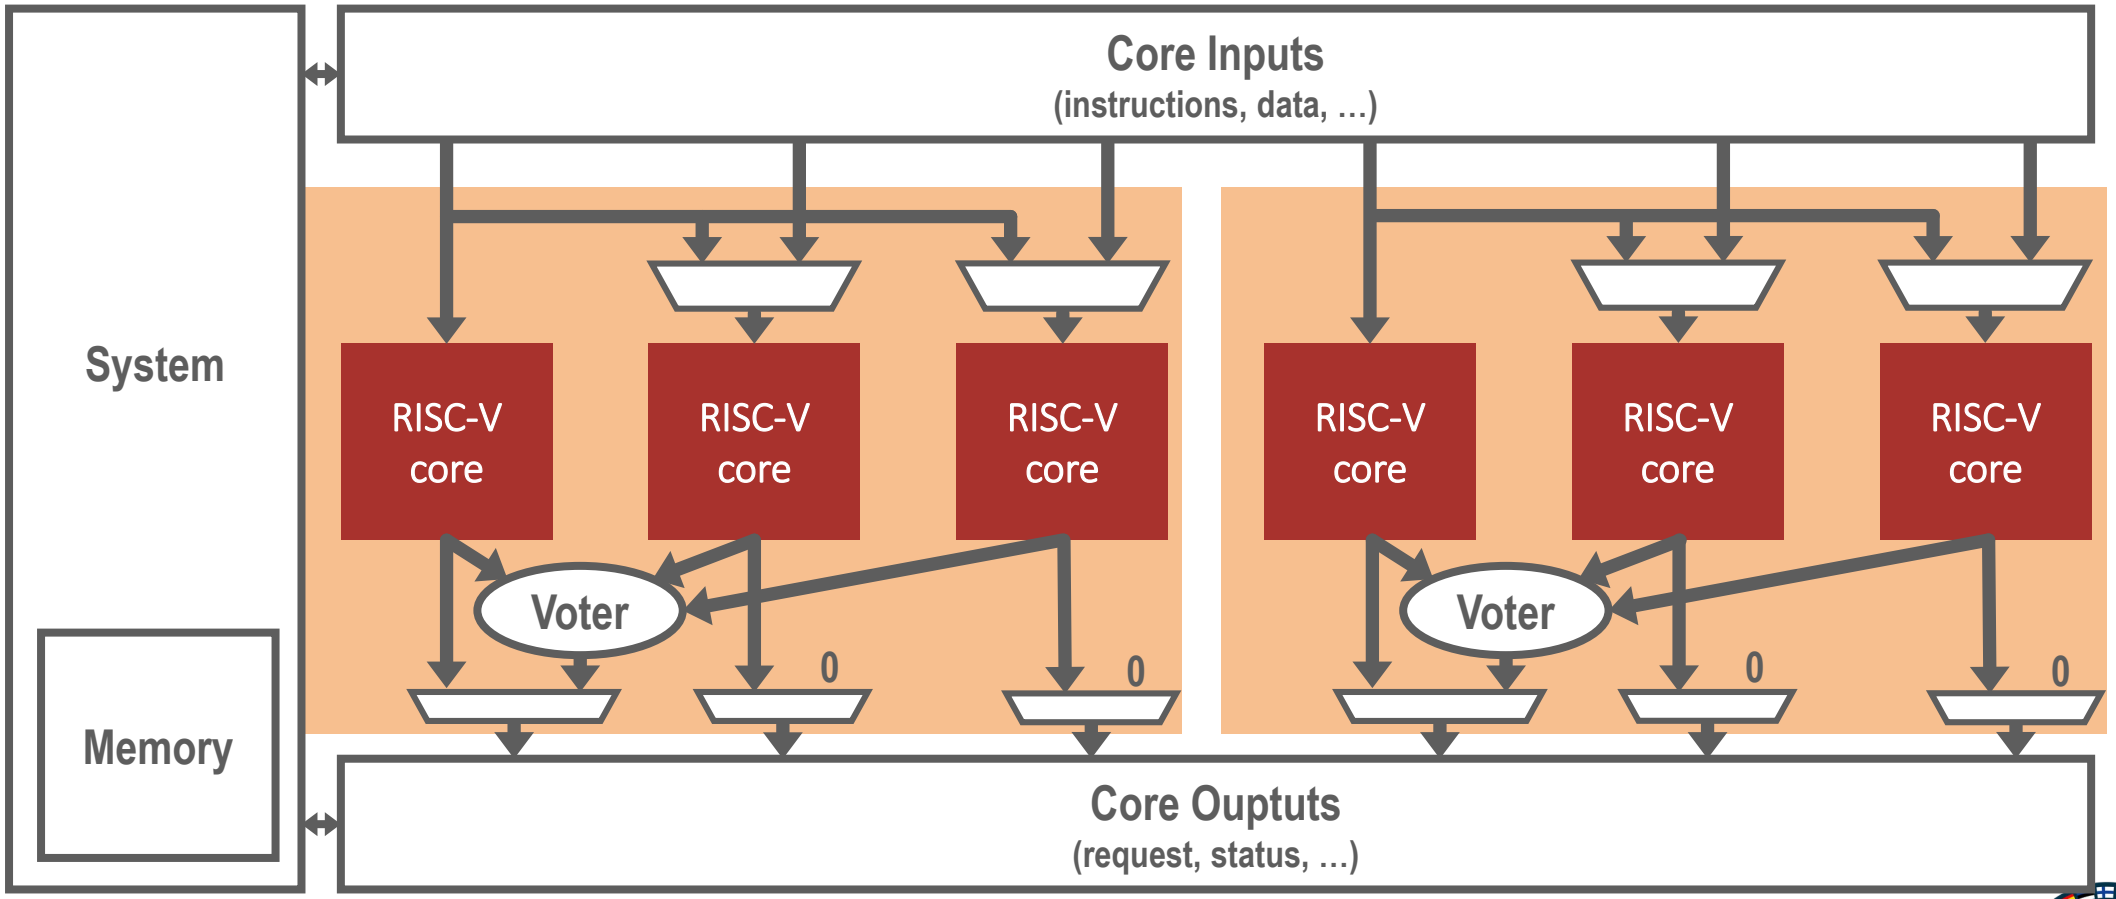

On-Demand Redundancy Grouping

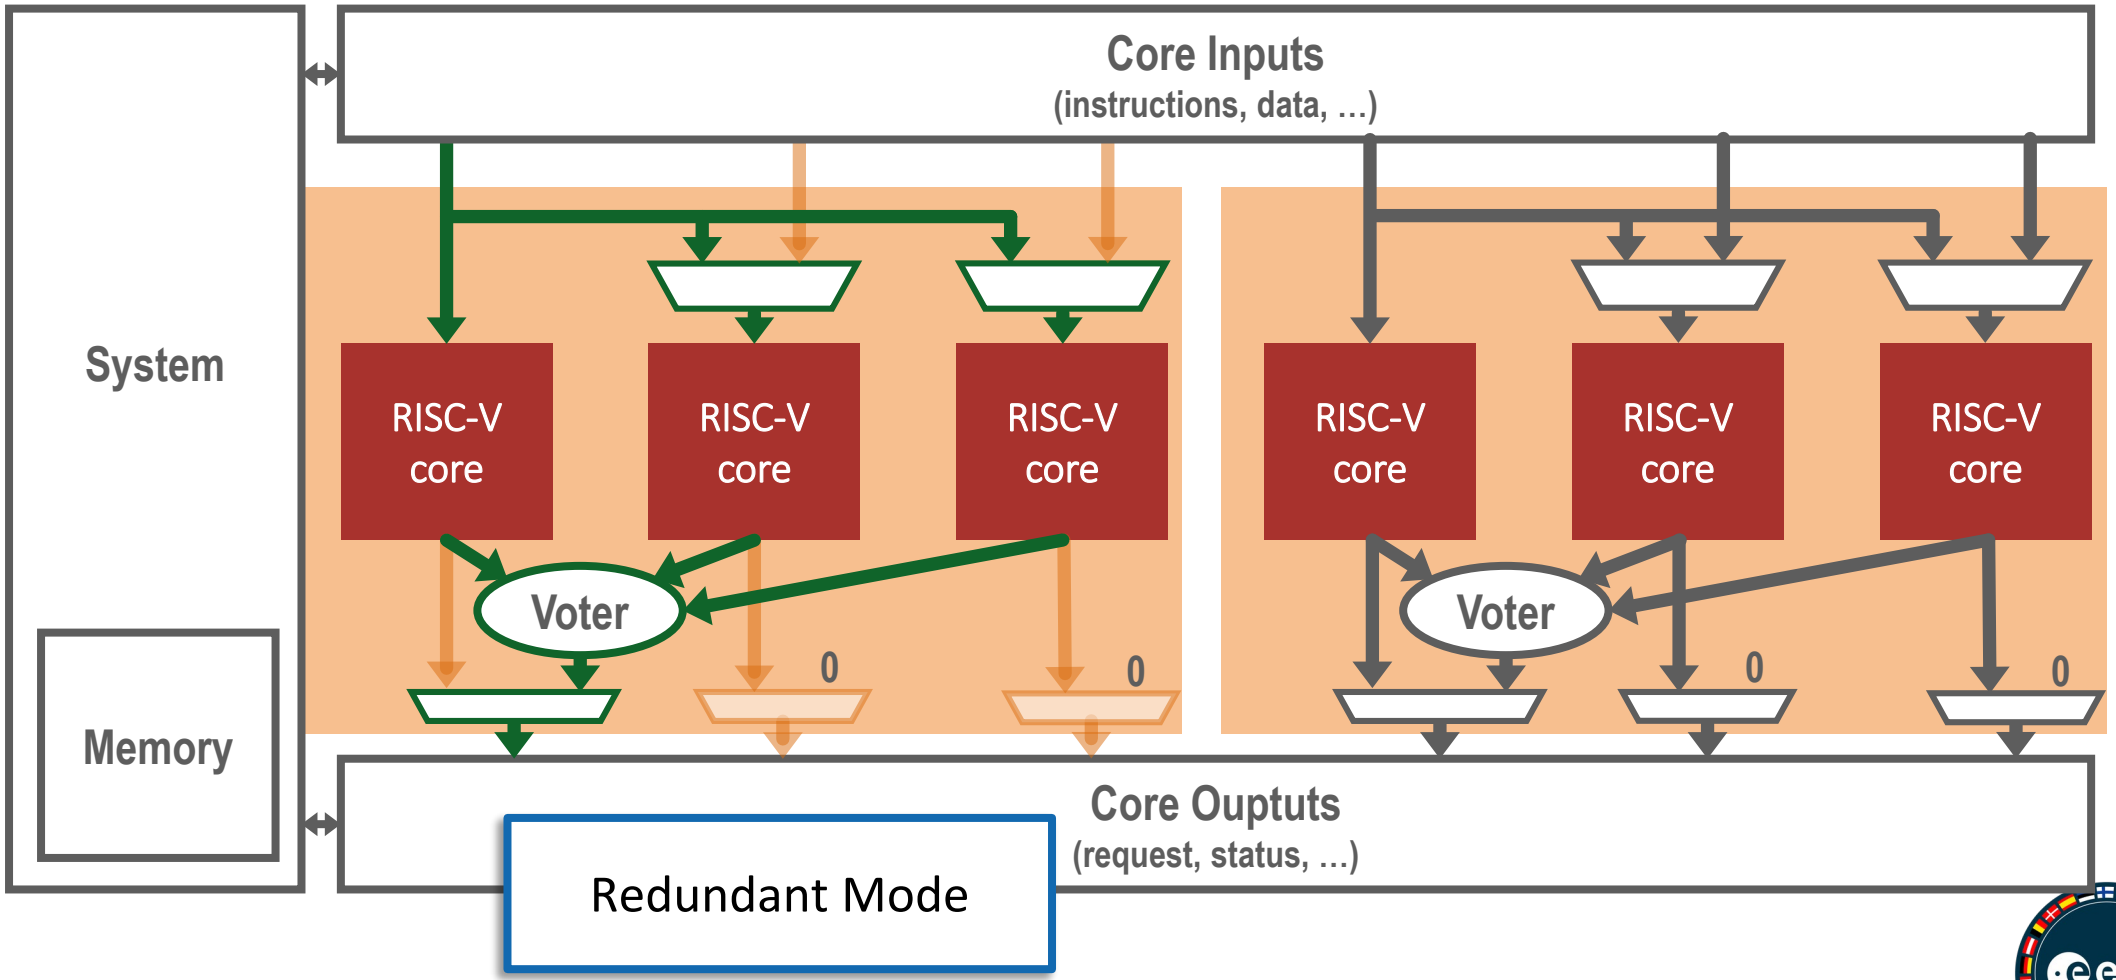

On-Demand Redundancy Grouping

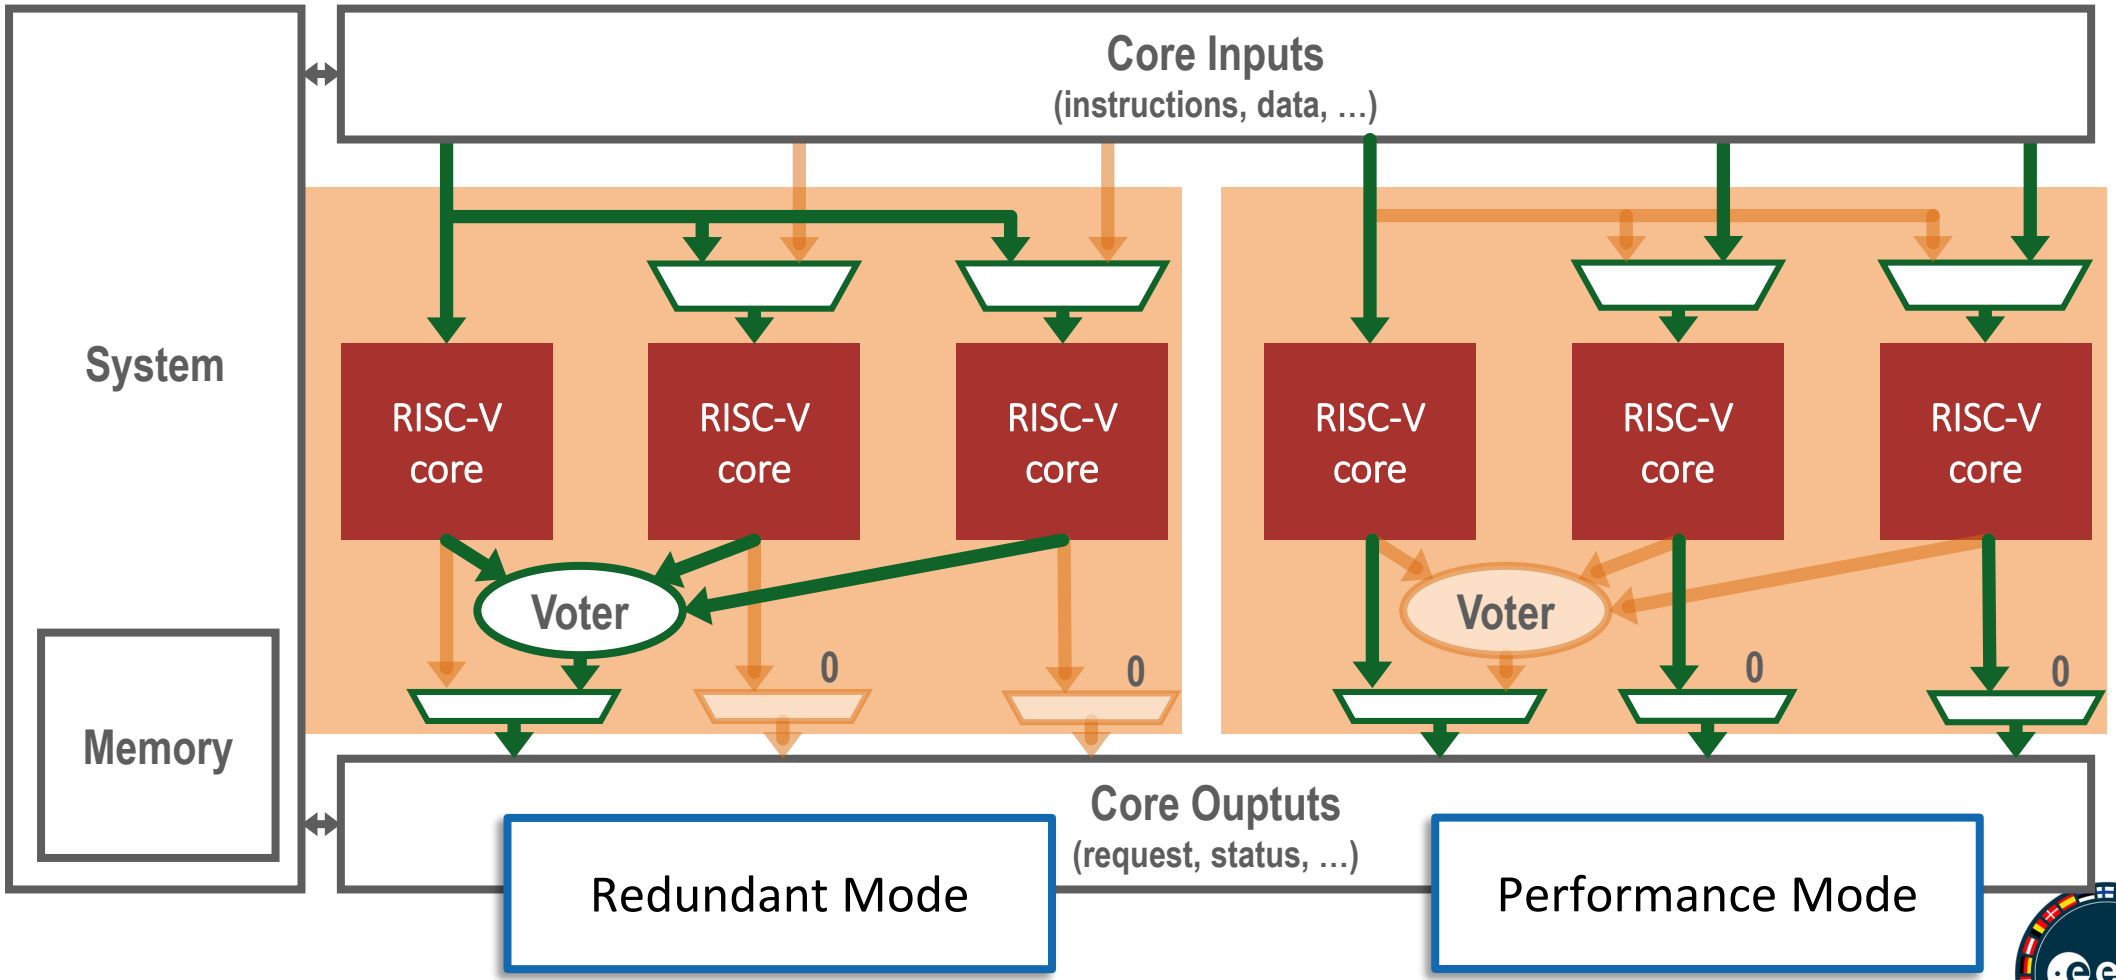

On-Demand Redundancy Grouping

- Up to 3x speedup for non-critical tasks
- <60’000 cycles to switch modes
- <1% area overhead for the cluster
  - ODRG for 3 cores requires ~11% area of a single core

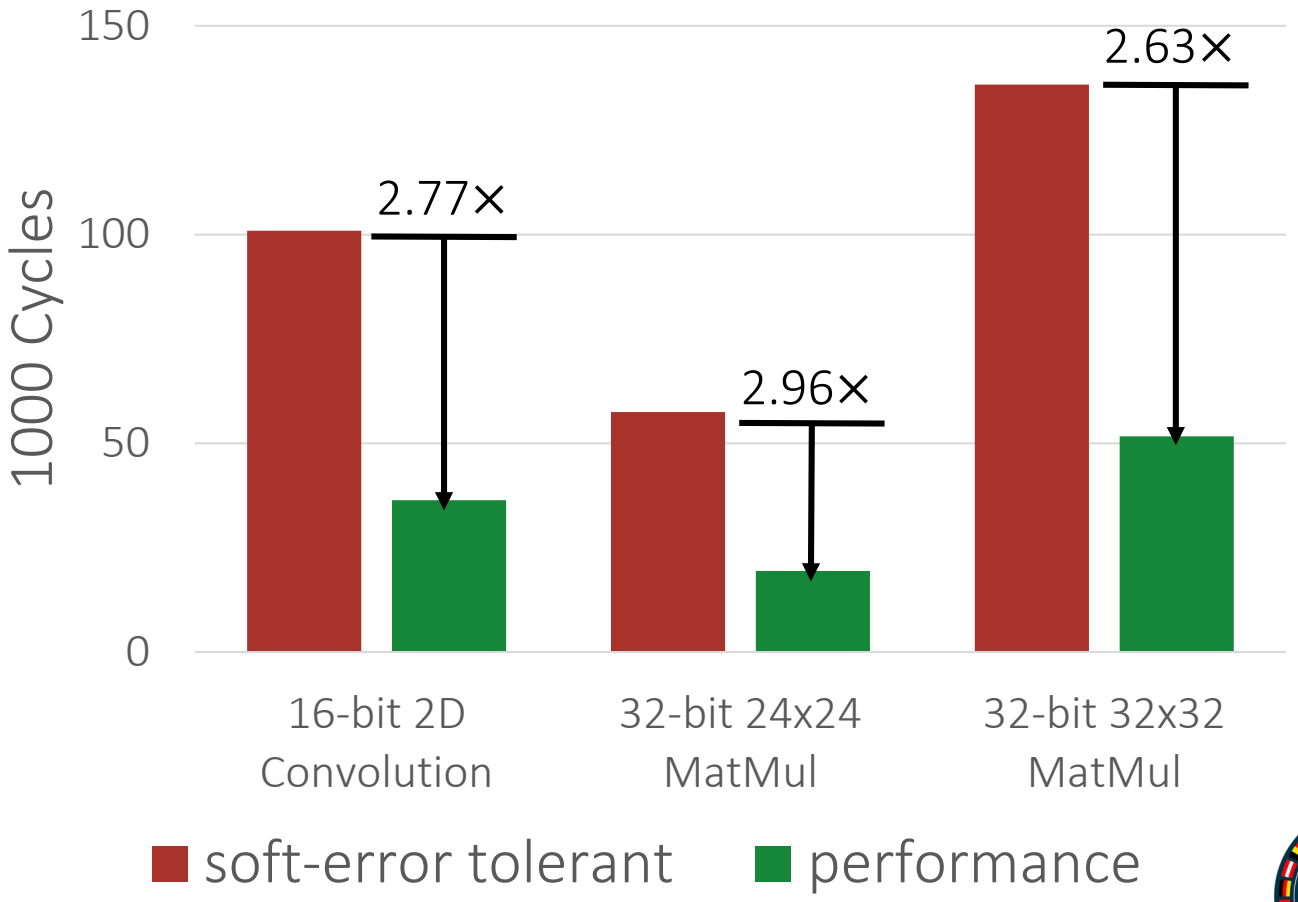

From Cores…

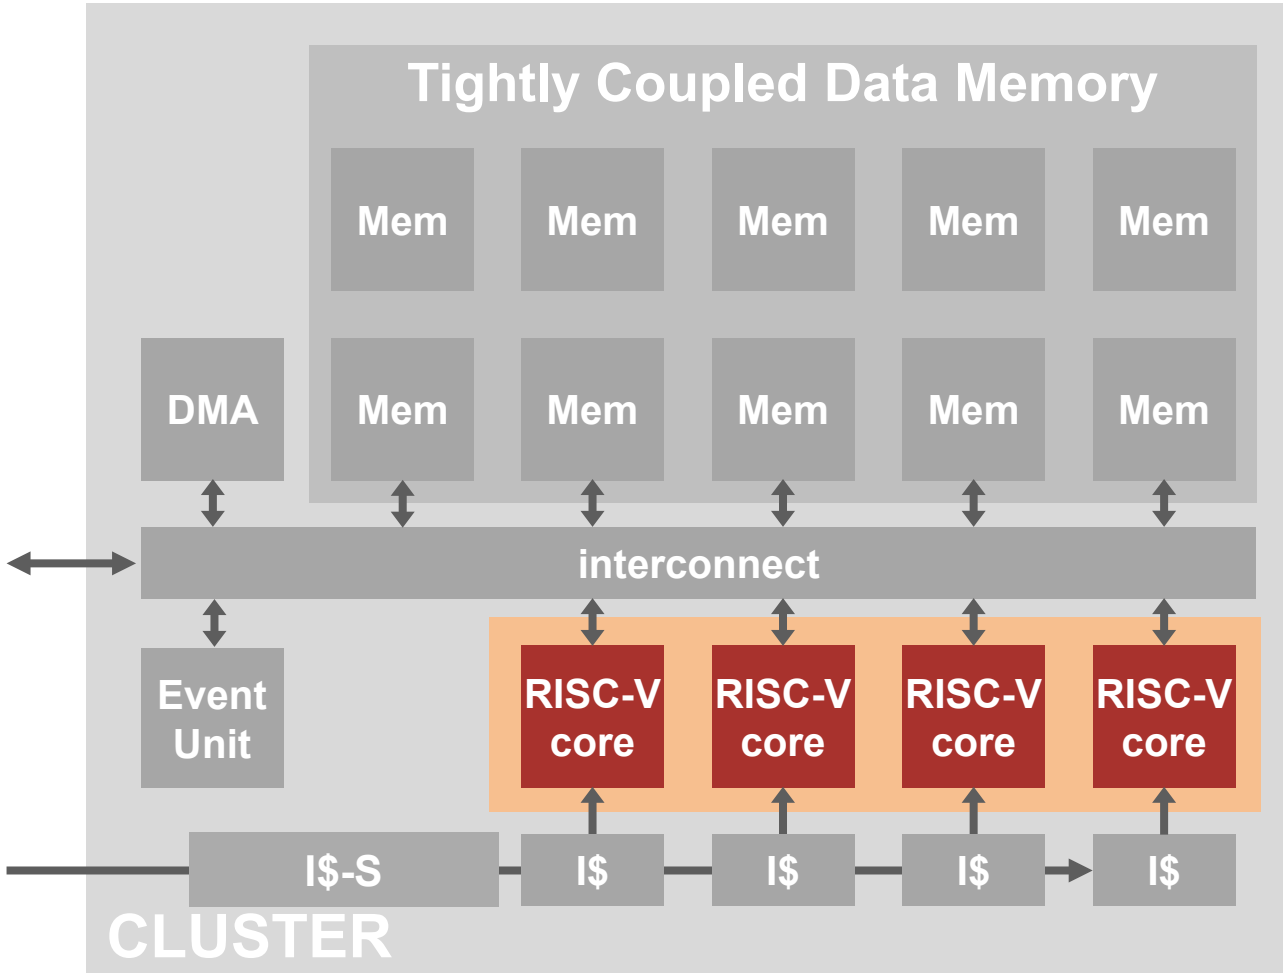
Tightly Coupled Data Memory
Mem
Mem
Mem
Mem
Mem
DMA
Mem
Mem
Mem
Mem
Mem
interconnect
Event Unit
RISC-V core
RISC-V core
RISC-V core
RISC-V core
I$-S
I$
I$
I$
I$
CLUSTER

To Memory

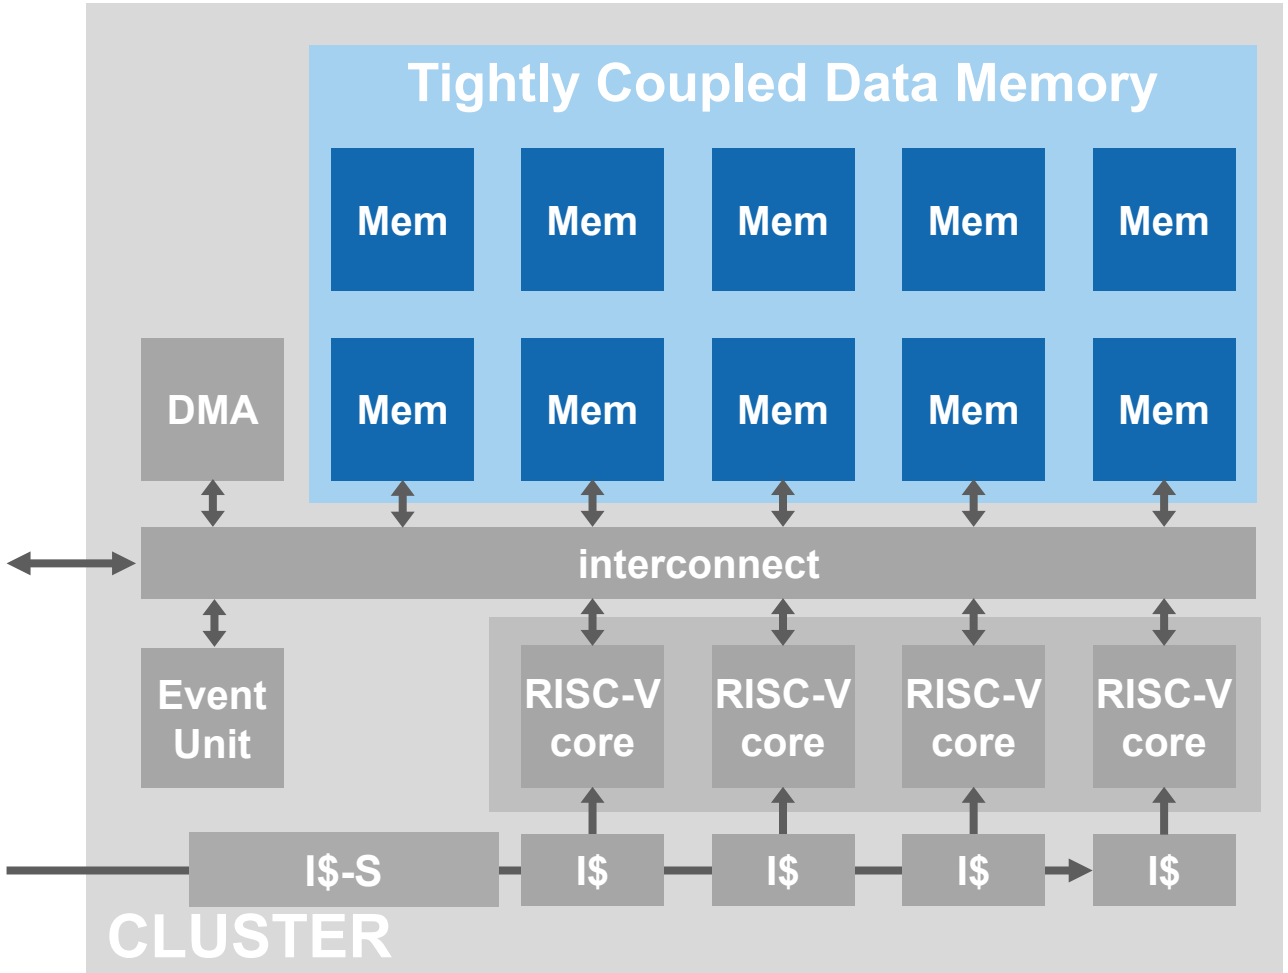
Tightly Coupled Data Memory
Mem
Mem
Mem
Mem
Mem
Mem
Mem
Mem
Mem
Mem
DMA
interconnect
Event Unit
RISC-V core
RISC-V core
RISC-V core
RISC-V core
I$-S
I$
I$
I$
I$
CLUSTER

ECC Memory

- 39-32 bit Hsiao Code
- Efficient byte-store
  - Coremark has <10% non-word stores
  - <1% overhead

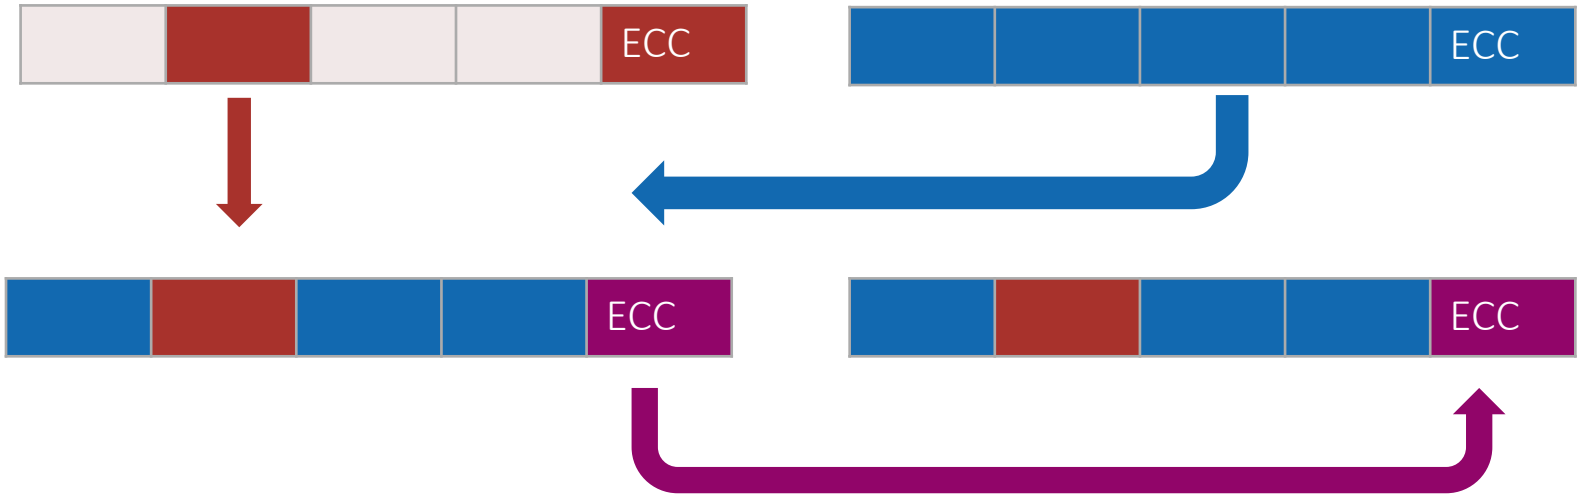

- Scrubber per memory bank

Trikarenos

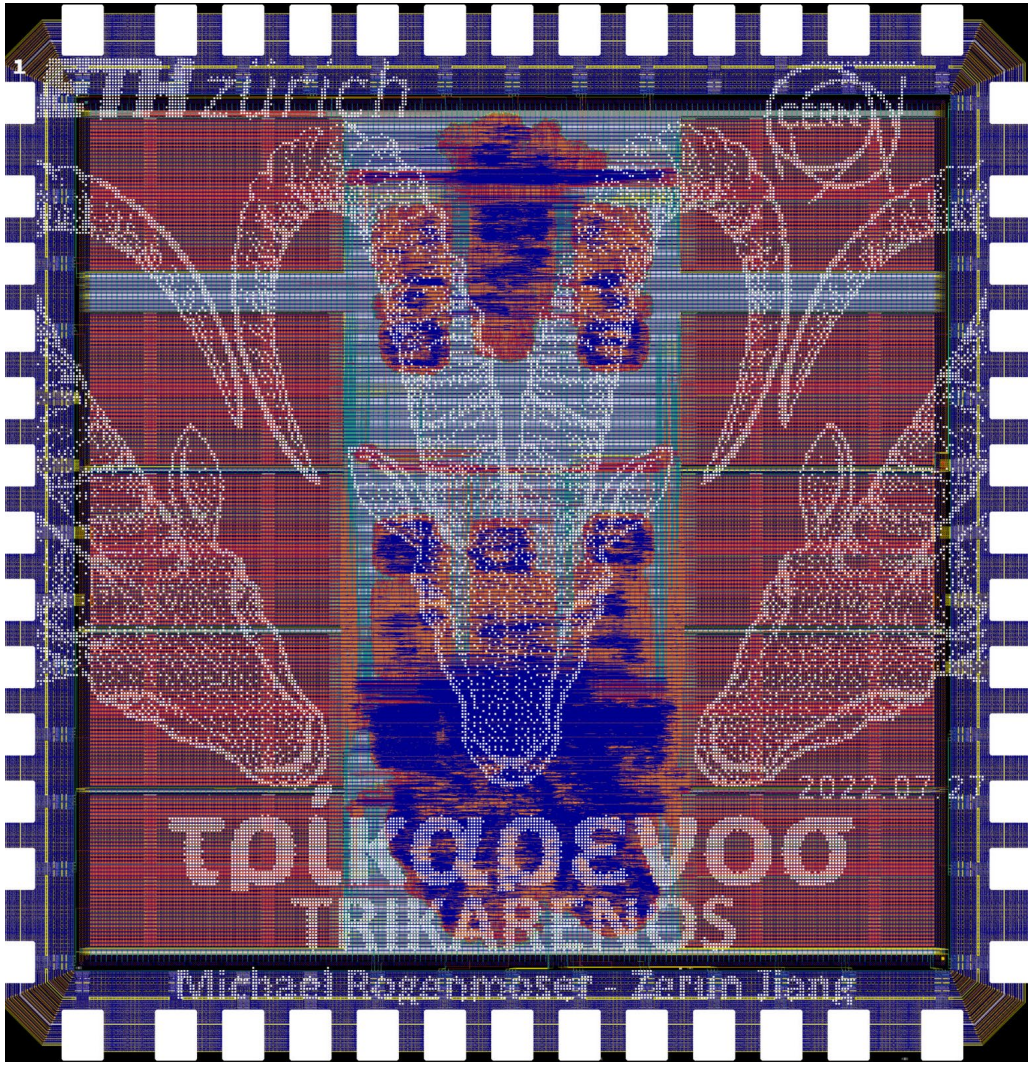

- Fault tolerance features implemented in an ASIC
- TSMC 28nm demonstrator
- Incorporates testing features

- Planned to fly on ARIS' CubeSat

aris
space to grow

Overview

1. The PULP Platform
2. Leveraging the PULP cluster for a Fault-Tolerant System
3. Using Heterogeneity for High-Performance Systems

Heterogeneous + Parallel… Why?

- Processors can do two kinds of useful work:

Decide (jump to different program part)

- Modulate flow of **instructions**
- Mostly sequential decisions:
  - Don't work too much
  - Be clever about the battles you pick (latency is king)
- **Lots of decisions**
  Little number crunching

Compute (plough through numbers)

- Modulate flow of **data**
- Embarassingly data parallel:
  - Don't think too much
  - Plough through the data (throughput is king)
- **Few decisions**
  Lots of number crunching

- Today's workloads are dominated by "Compute":
  - Tons of data, few (as fast as possible) decisions based on the computed values,
  - Data-Oblivious Algorithms (ML, or better DNNs are so!)
  - Large data footprint + sparsity

From Cluster to full Platform

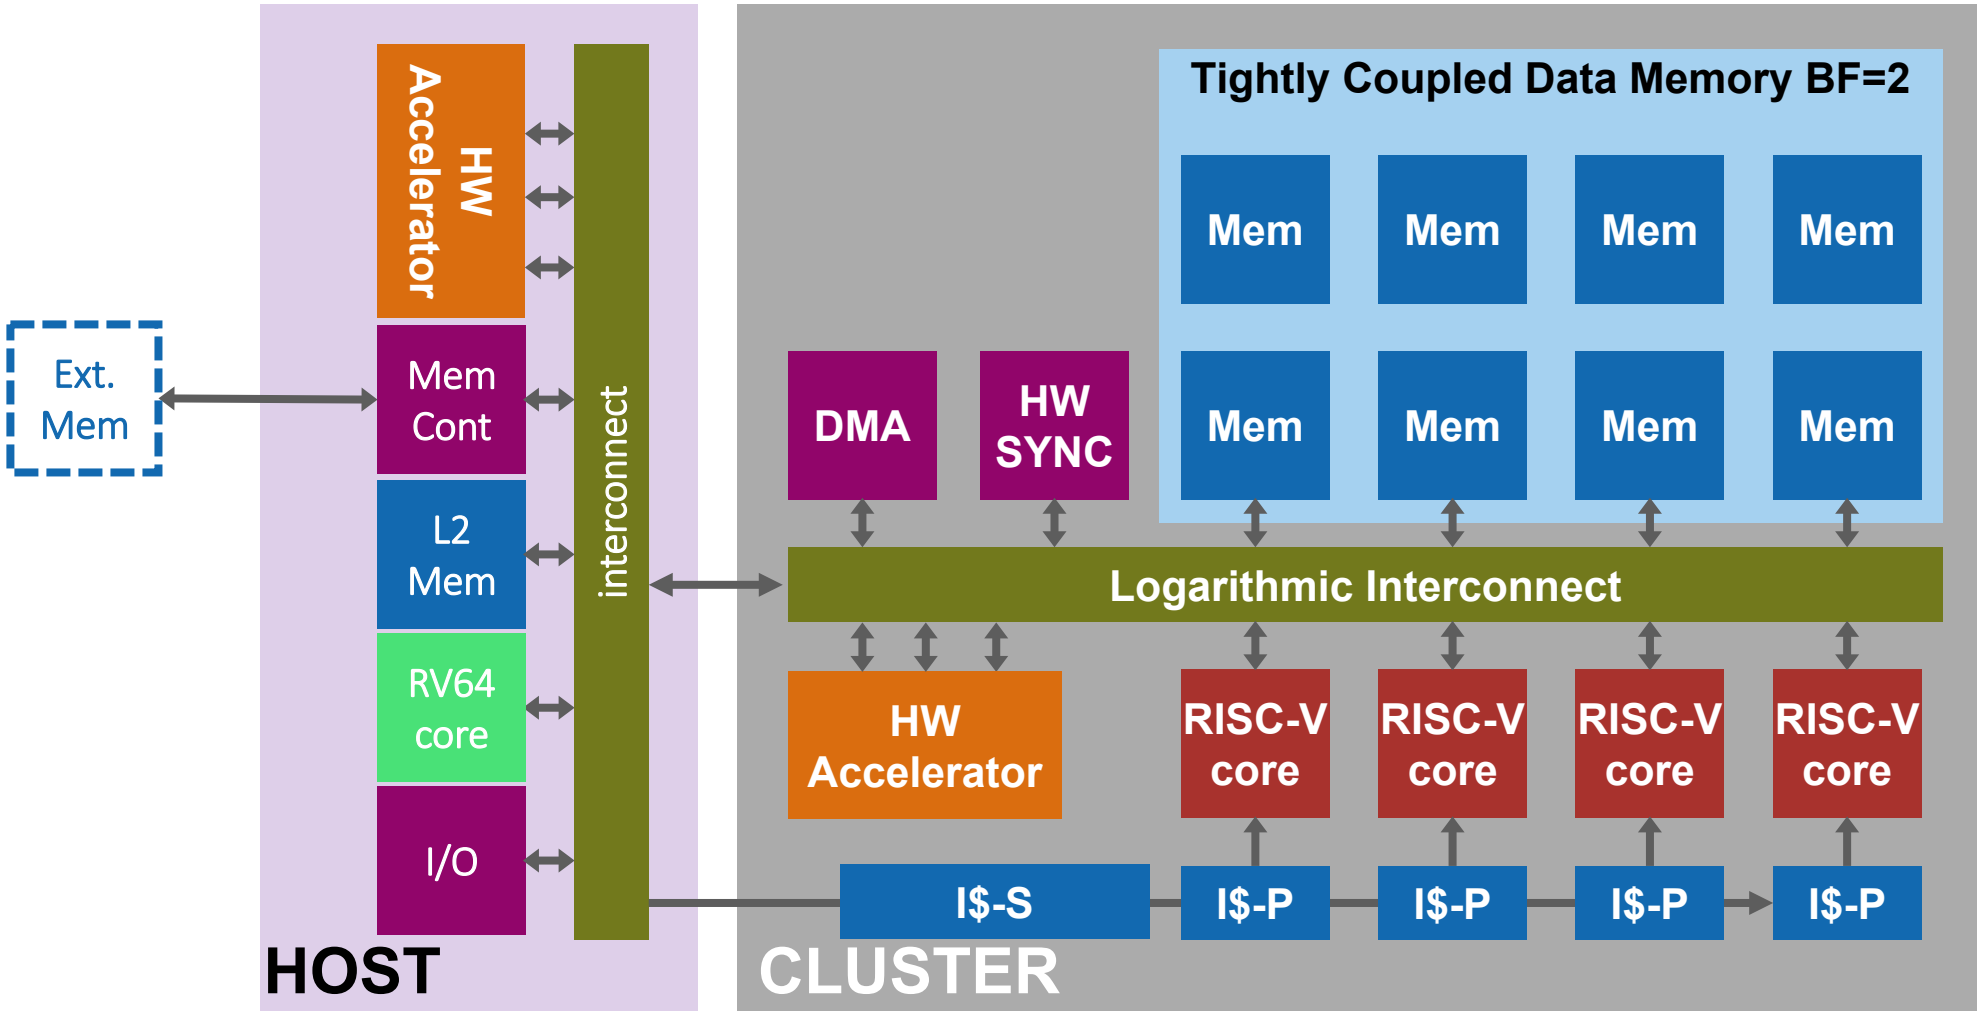

HW Accelerator Case Study: RedMulE

Reduced-Precision Matrix **Mul**tiplication **E**ngine

- Accelerates general matrix-matrix operations (**GEMM-Ops**) using mixed-precision floating point (**FP16/FP8**) arithmetic
- Enables **on-chip training** of **generalized** deep learning models
- Accelerates **graph** and **control-flow-based** algorithms (all-pairs shortest paths, minimum spanning tree, and others)

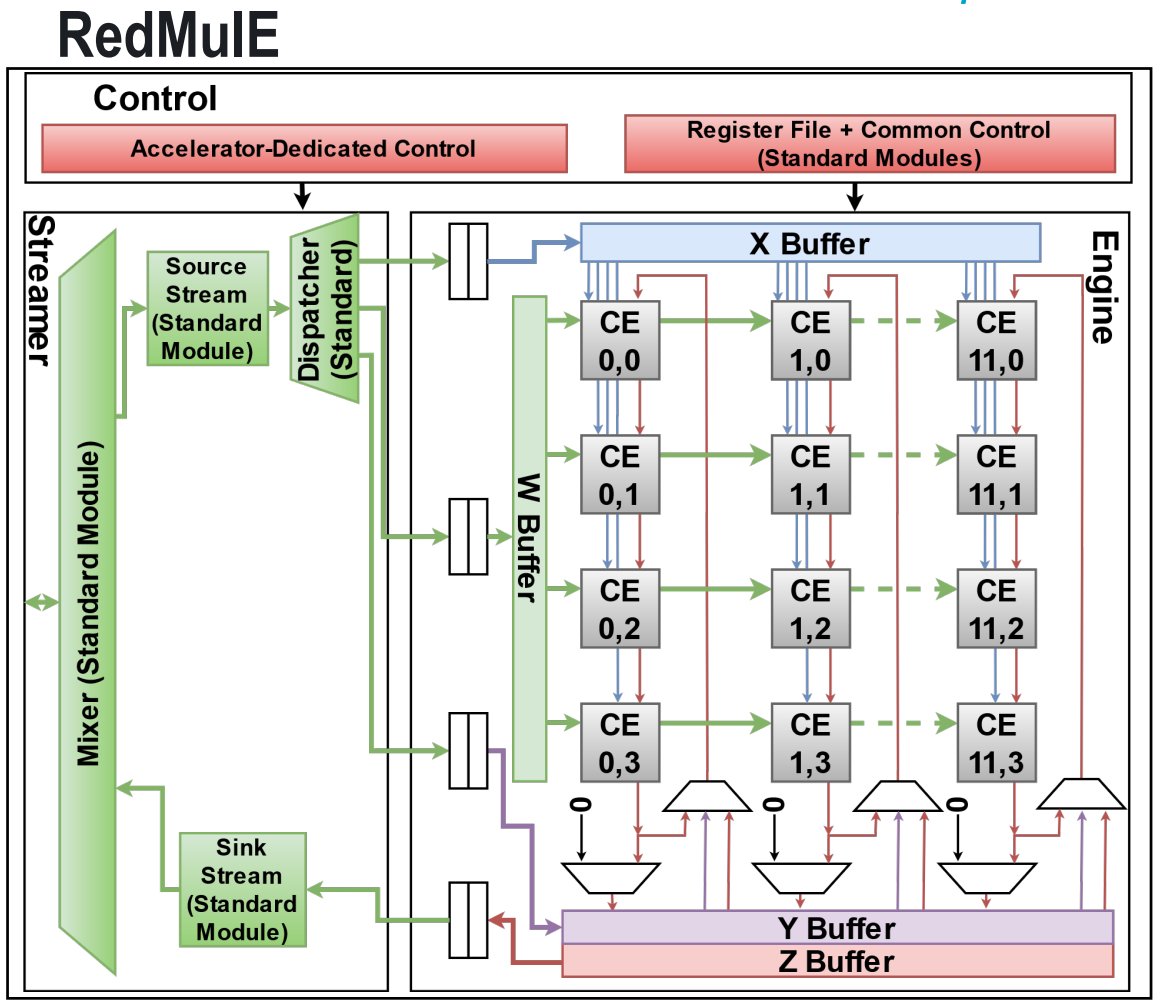

HW Accelerator Case Study: RedMulE

- Flexible template for hardware accelerators design
- Interface designed by composition of open-source IPs for **control** + **memory interfaces** through *HWPE-streams* → specialized streamers convert from streams to mem (TCDM)
- Memory access capabilities **equivalent or superior** to SW (misaligned access, delayed stores, high bandwidth) on part or all of the memory map

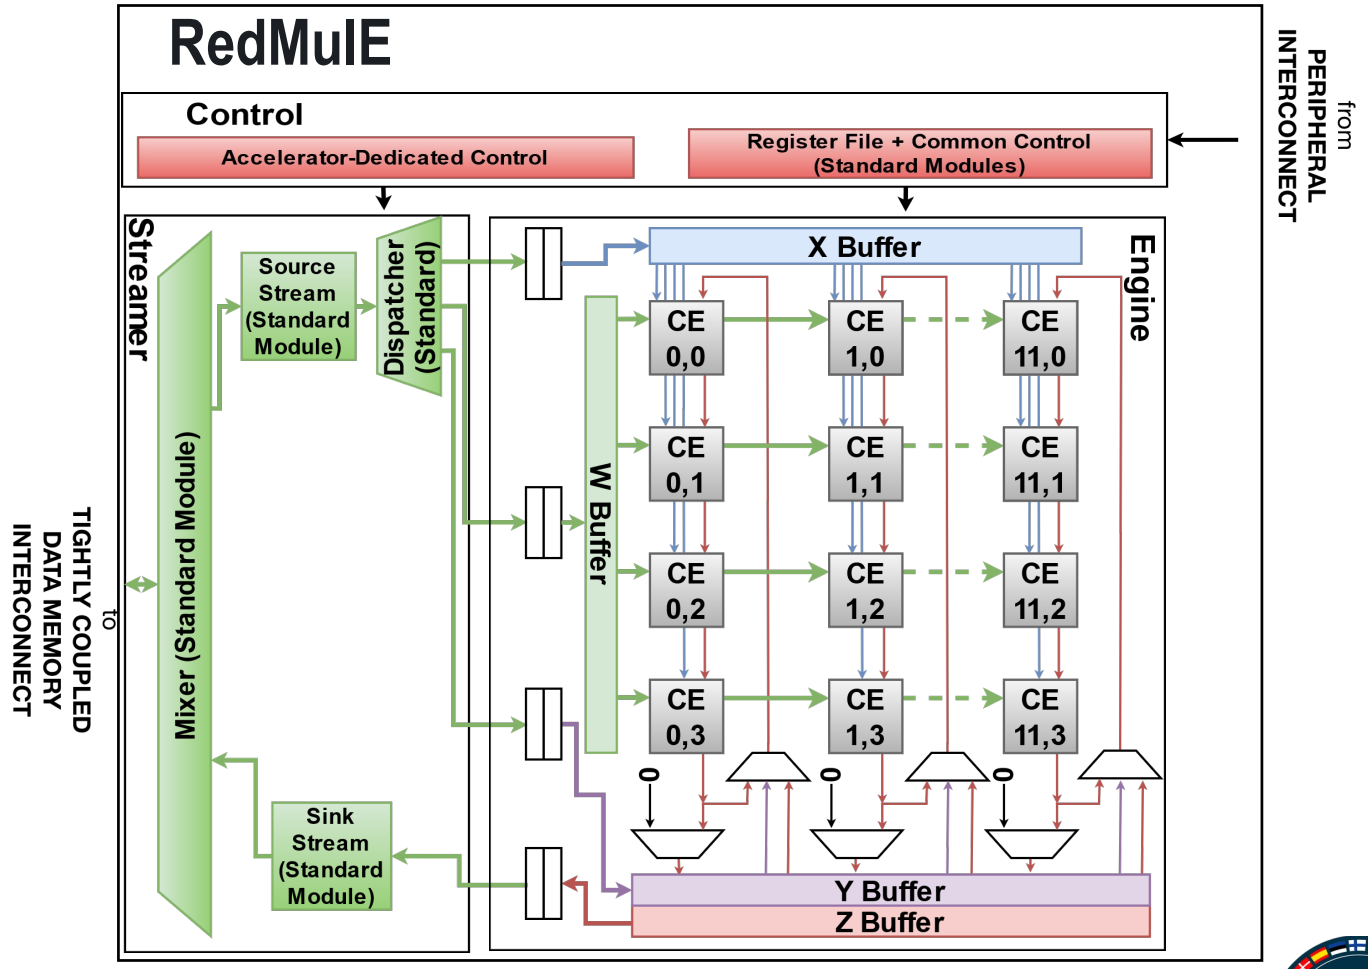

Heterogeneous SoC

- 64-bit **Linux-capable** RISC-V core
- Shared memory among the two existing domains
- Interaction with **external** environment (memory, IO)
- **Cross-domain** communication: **mailbox** + **TLB** for virtualization

Features

- **Ease of programmability** of Single-Board Computers **(Linux)**
- **Performance** and **efficiency** of parallel multi-core cluster

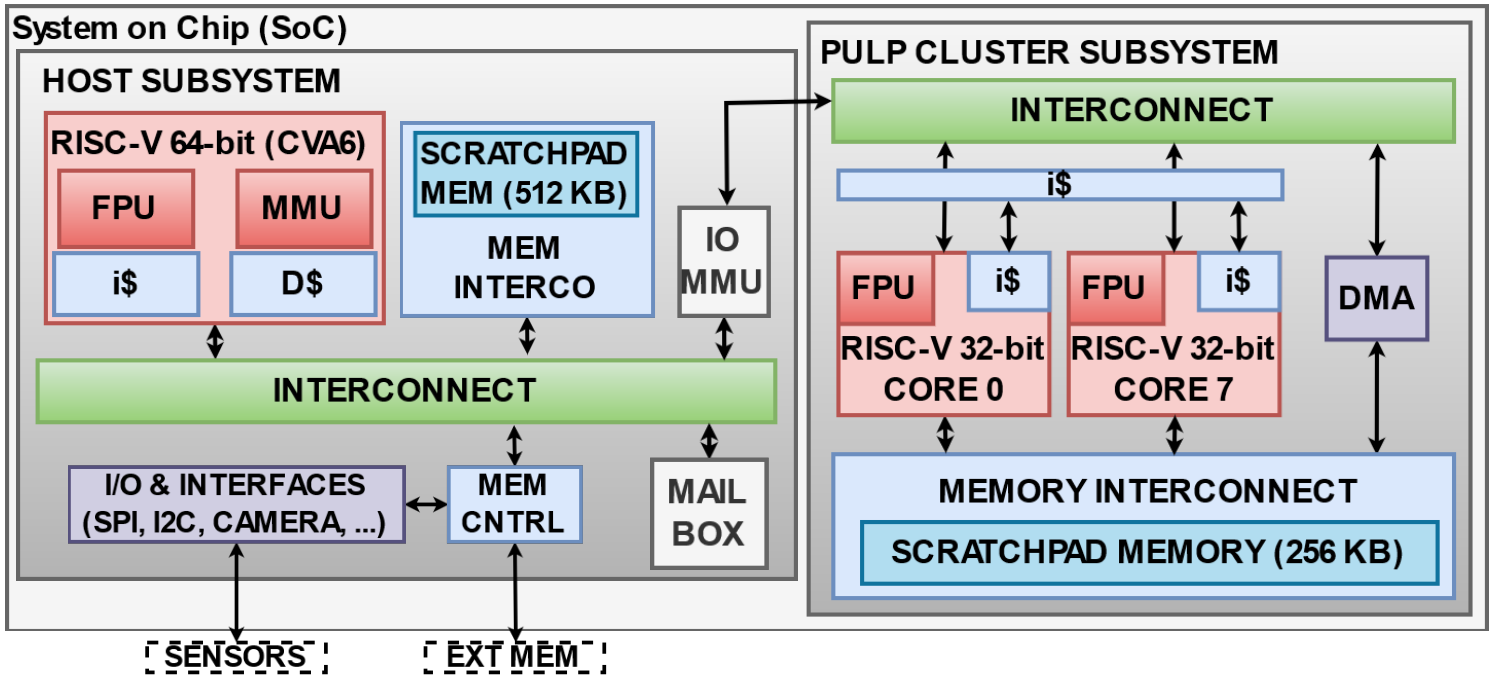

Heterogeneous SoC for Space Applications

OpenMP Offloading

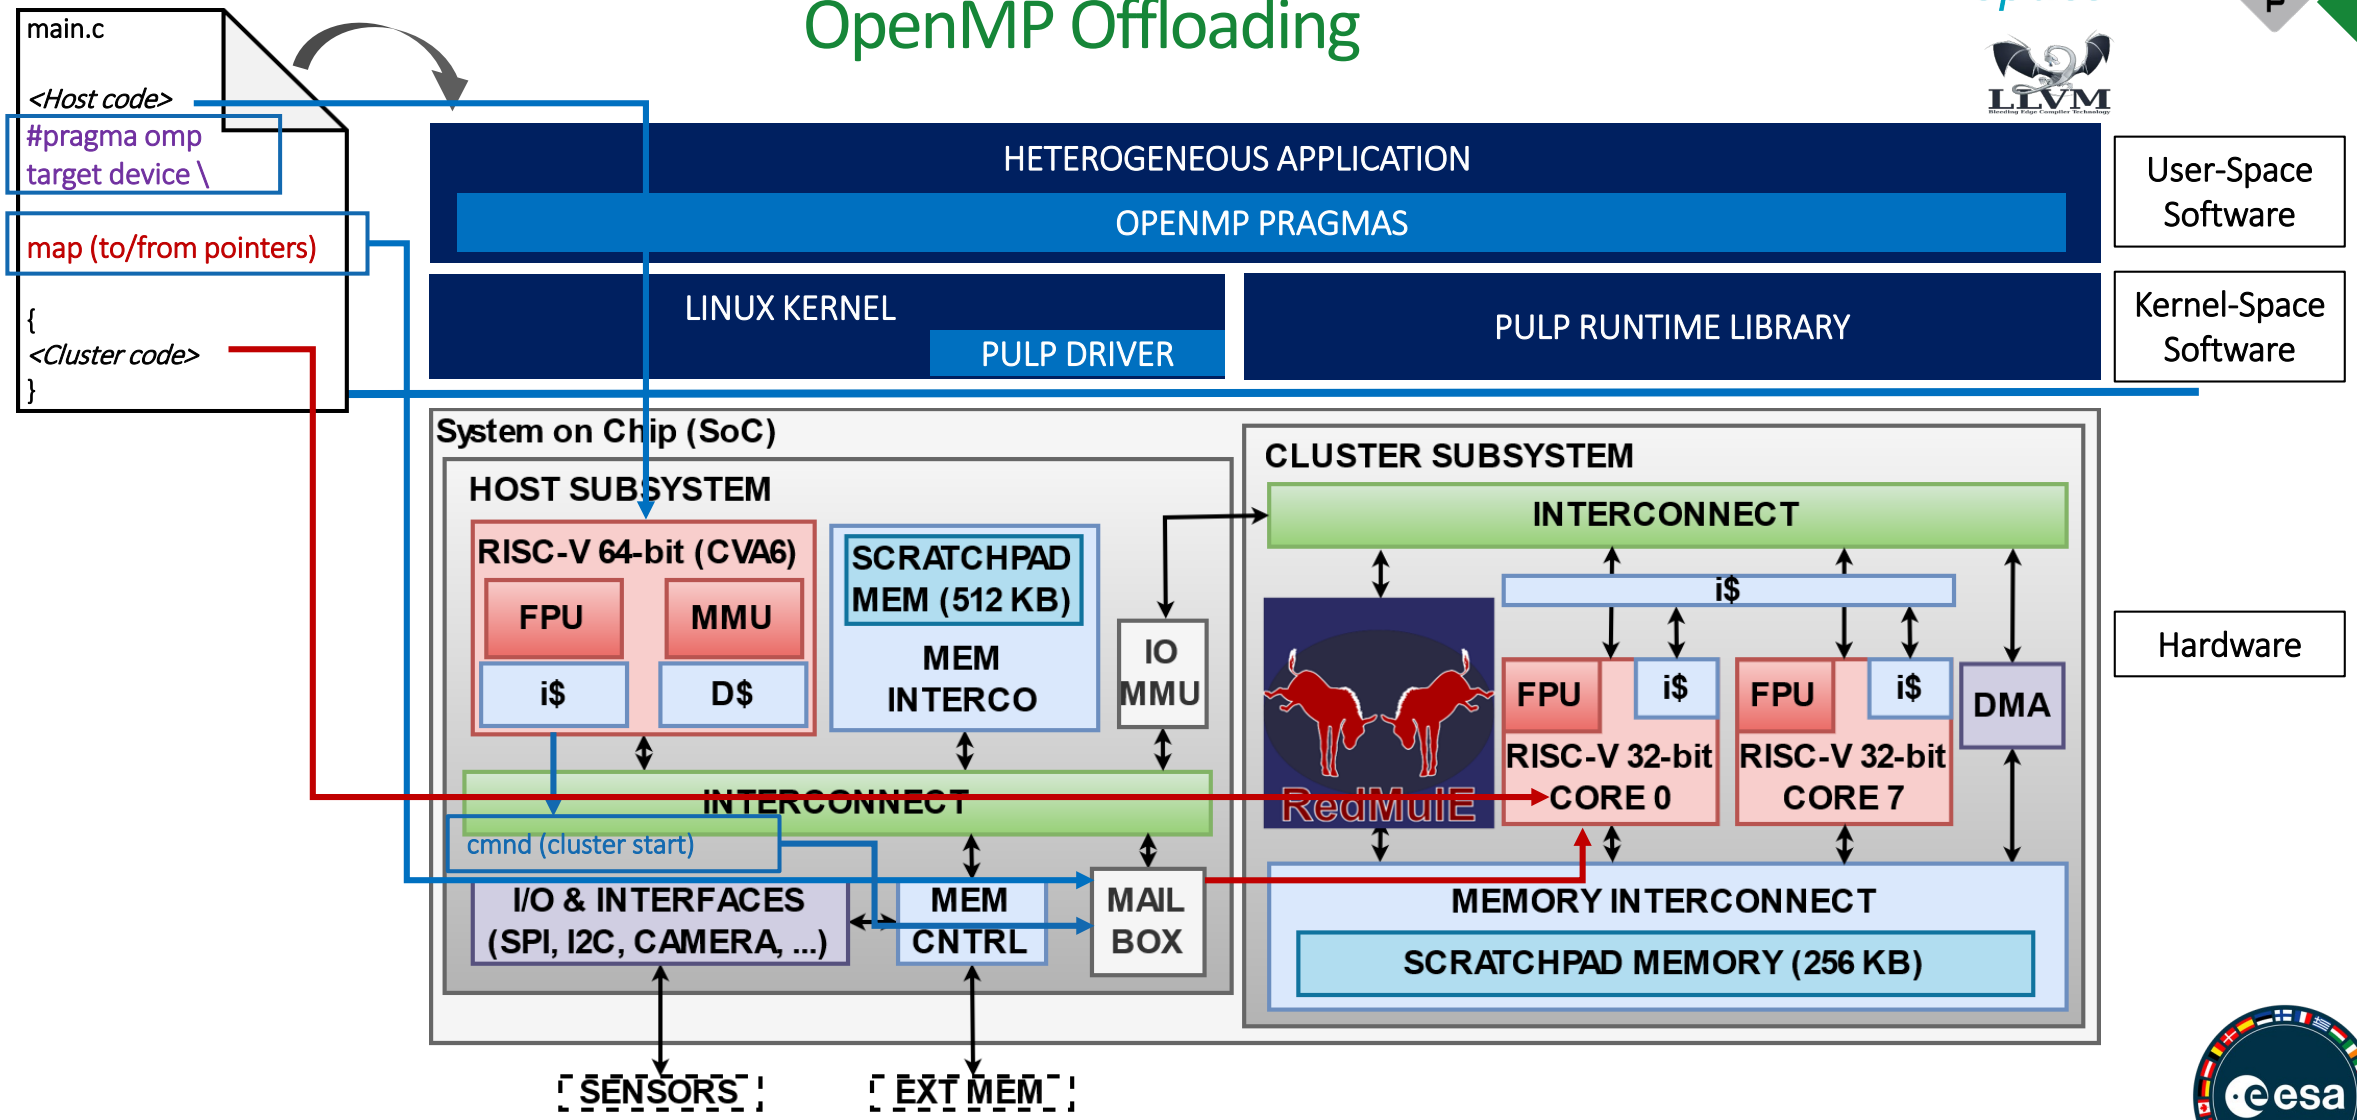

Heterogeneous SoC for Space Applications

Hybrid DSP Kernel-Set (Integer + FP)

Performance (@ $500\ MHz$):

- Integer: Cluster speedup up to **55x** with respect to Host
- FP: Cluster speedup up to **20x** with respect to Host

Energy Efficiency ($22\ nm, 0{,}8\ V$)

- Integer: up to $62\ GOPS/W$ (**12,7x** better than Host)
- FP: up to $15\ GFLOPS/W$ (**3,19x** better than Host)

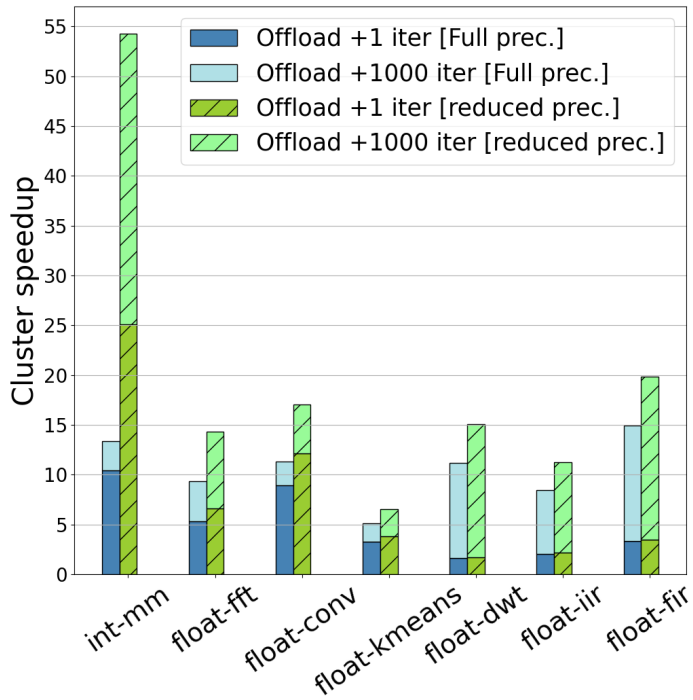

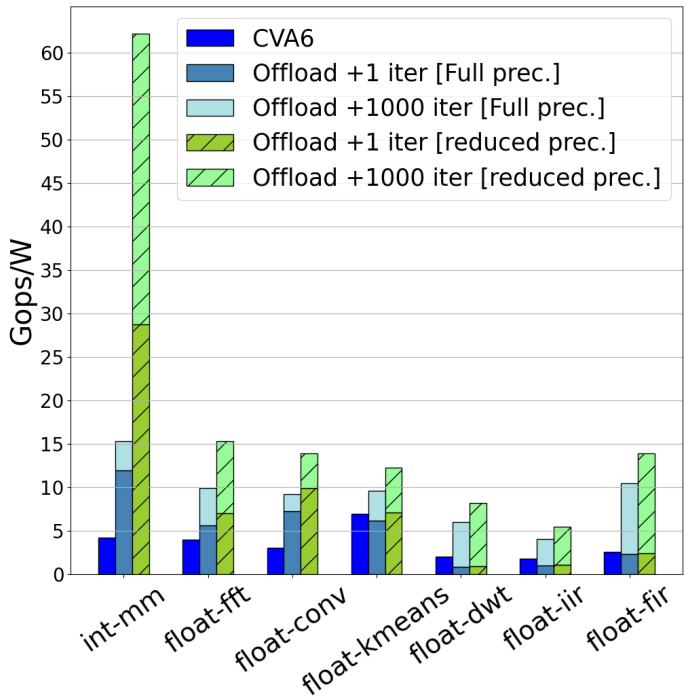

Future Developments

- Hybrid Modular Redundancy (**HMR**) in PULP clusters: **selectable** dual or triple core **lockstep**
- **Host** Core microarchitecture **protection** (redundancy, error correction)
- Reliable **on-chip interconnect**
- Reliable **peripherals communication** (interfaces redundancy)
- Reliability of **DMA** and **Hardware Accelerators**

Conclusion: What is PULP?

- Core
  - Improving core efficiency with ISA and uAch extensions
- Cluster
  - Efficient shared-mem cluster
  - From a few to thousand processing element
- Full platform
  - Heterogeneity: host, processor, accelerator
  - One or multiple accelerators
  - From chips to chiplets (2D to 3D)

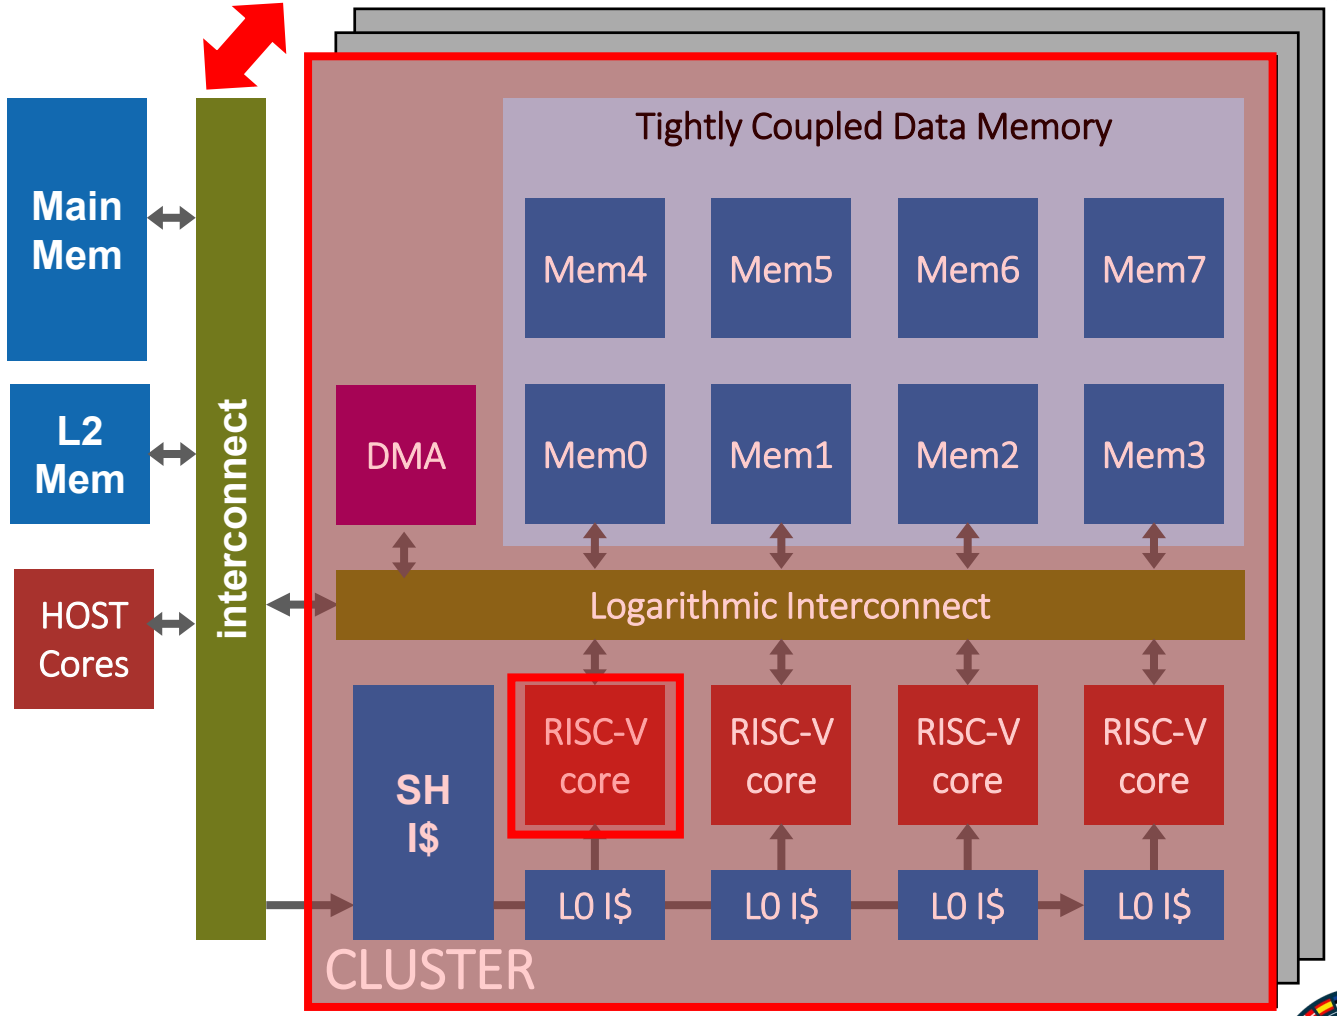

PULP

Parallel Ultra Low Power

Luca Benini, Ahmad Mirsalari, Alessandro Capotondi, Alessandro Nadalini, Alessandro Ottaviano, Alessio Burrello, Alfio Di Mauro, Andrea Borghesi, Andrea Cossettini, Angelo Garofalo, Arpan Prasad, Chi Zhang, Corrado Bonfanti, Cristian Cioflan, Cyril Koenig, Daniele Palossi, Davide Rossi, Fabio Montagna, Florian Glaser, Francesco Conti, Georg Rutishauser, Germain Haugou, Gianna Paulin, Giuseppe Tagliavini, Hanna Müller, Jannis Schoenleber, Lorenzo Lamberti, Luca Bertaccini, Luca Colagrande, Luca Valente, Maicol Caini, Manuel Eggimann, Manuele Rusci, Marco Bertuletti, Marco Guermandi, Matheus Cavalcante, Matteo Perotti, Mattia Sinigaglia, Michael Rogenmoser, Moritz Scherer, Moritz Schneider, Nazareno Bruschi, Nils Wistoff,, Paul Scheffler, Philipp Mayer, Robert Balas, Samuel Riedel, Segio Mazzola, Sergei Vostrikov, Simone Benatti, Thomas Benz, Thorir Ingolfsson, Tim Fischer, Victor Javier Kartsch Morinigo, Victor Jung, Viviane Potocnik, Vlad Niculescu, Xiaying Wang, Yichao Zhang, Yvan Tortorella, Frank K. Gürkaynak, all our past collaborators

and many more that we forgot to mention

http://pulp-platform.org

@pulp_platform

ETH zürich

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Additional Slides

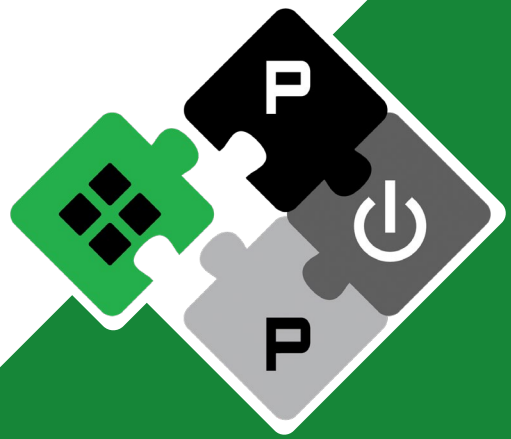

@pulp_platform

pulp-platform.org

youtube.com/pulp_platform

PULP Platform

Open Source Hardware, the way it should be!

Re-synchronization

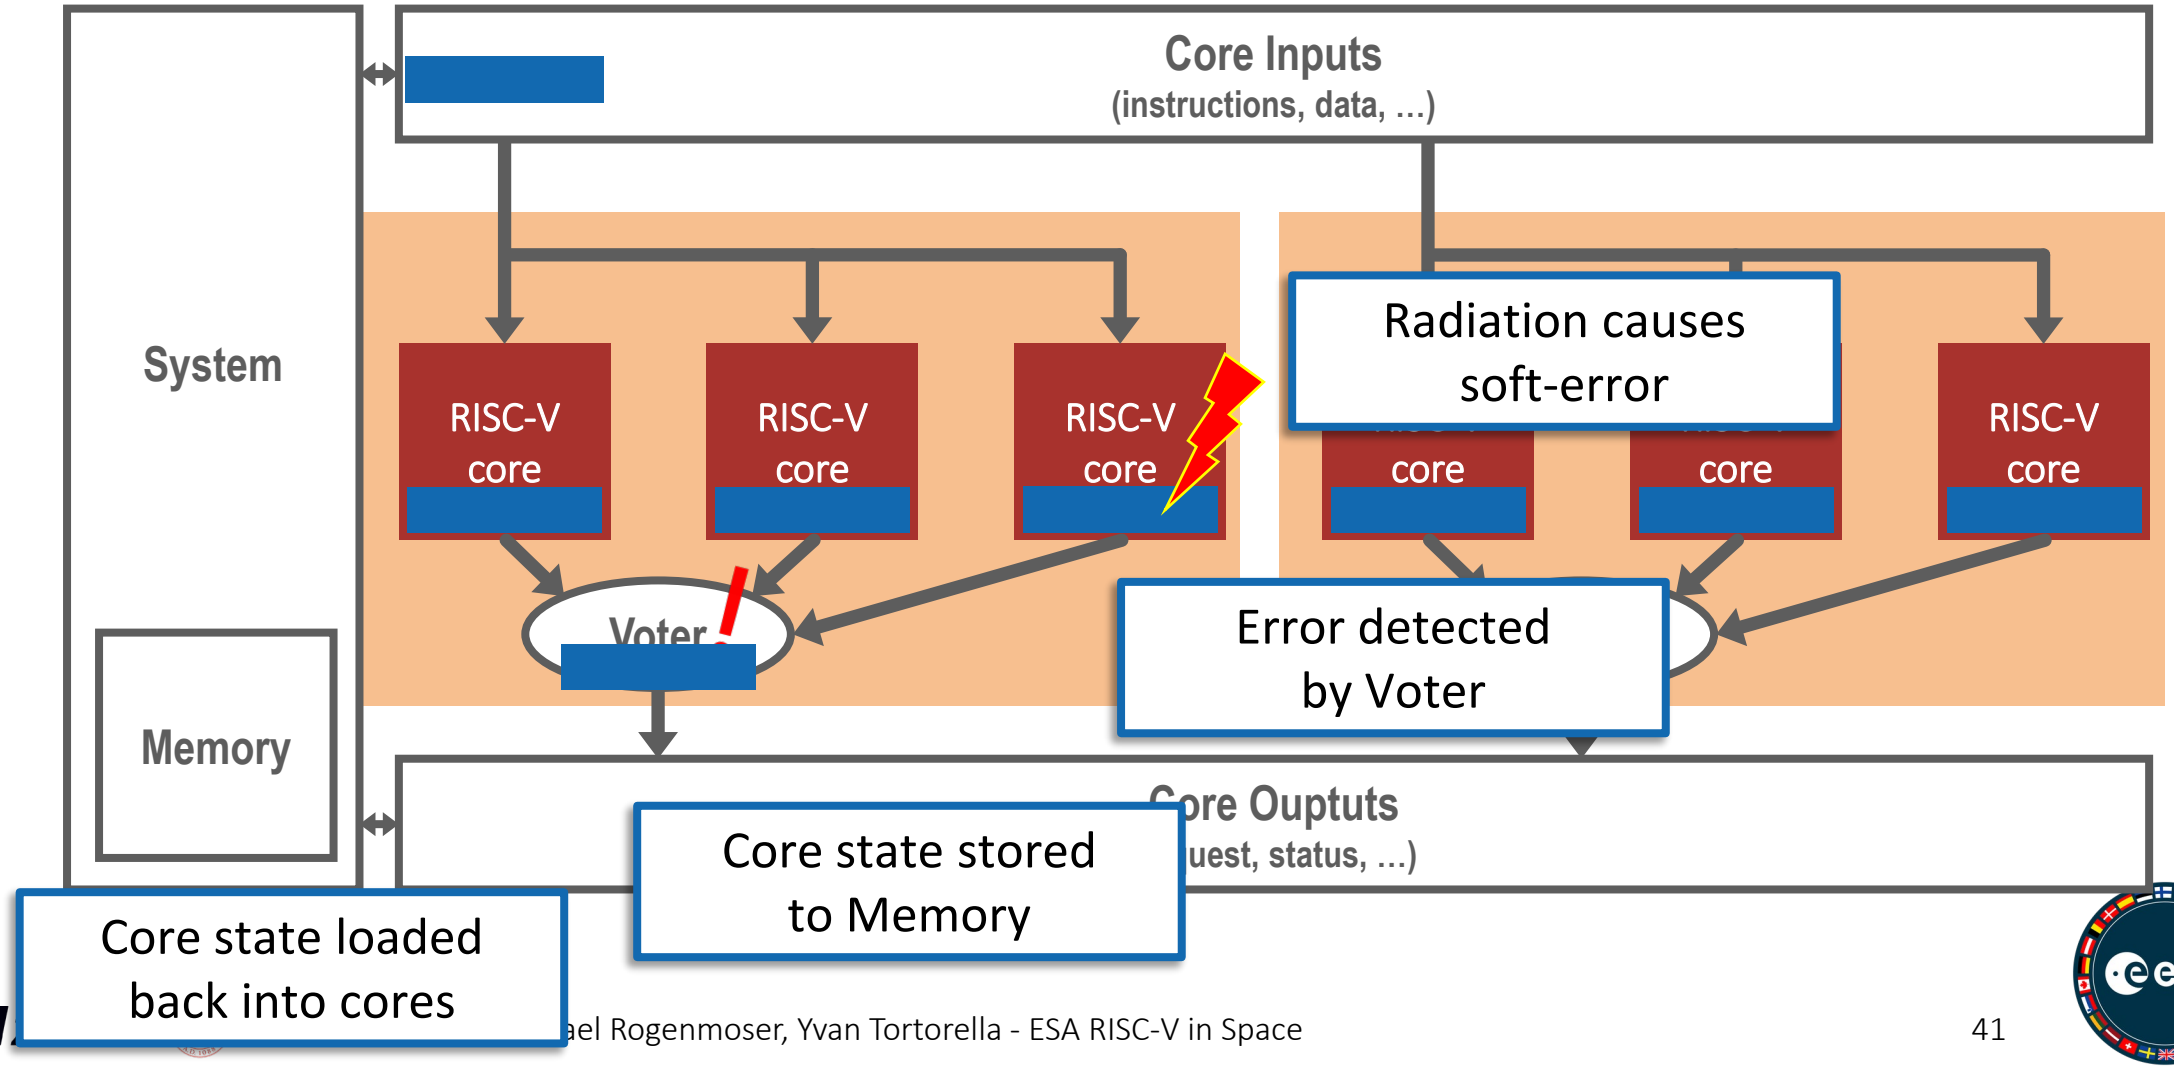

RISC-V cores developed by PULP team

32 bit			64 bit
Low Cost Core	**Core with DSP support**	**Core for Streaming**	**Linux capable Core**
▪ **IBEX** ▪ Zero-riscy ▪ Micro-riscy ▪ **2 options** ▪ RV32ICM ▪ RV32EC *ARM Cortex-M0+*	▪ **CV32E40P** ▪ RV32-ICMXF ▪ SIMD ▪ HW loops ▪ Bit manipulation *ARM Cortex-M4F* ▪ Fixed point	▪ **Snitch** ▪ RV32-IMAFD ▪ Small Int core ▪ Big 64bit FPU *novel* ▪ Extensions for	▪ **Ariane/ CVA6** ▪ RV64-ICMAFD *ARM Cortex-A55*

Maintained By LowRISC

Maintained By OpenHW

Brand New Core

Maintained by OpenHW

PULP is Cores+Interco+IO+Accelerators → Open Platform

RISC-V Cores

Ibex RV32-ICM/CE	CV32E40P RV32-ICMXF	Snitch RV32-IMAFD	CVA6 RV64-ICMAFD

Peripherals

JTAG	SPI
UART	I2S
DMA	GPIO

Interconnect

- Logarithmic interconnect
- APB – Peripheral Bus
- AXI4 – Interconnect

Platforms

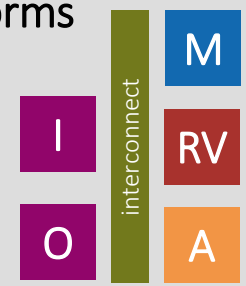

Single Core
- PULPissimo

Open-Source Hardware > Open-Source ISA:
- core implementations
- system-level integration
- peripherals, accelerators

Everything designed in System Verilog

- Mr. Wolf
- Vega

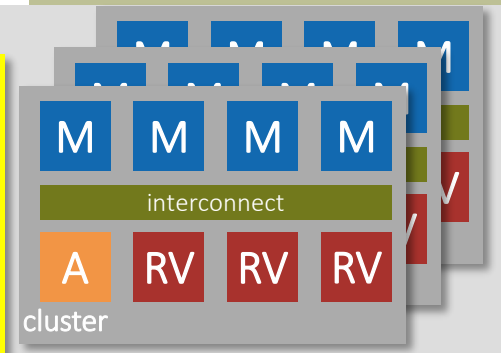

Multi-cluster
- HERO
- Open Piton

IOT → HPC

Hardware Accelerators

NE (DNN)	NTX (ML)	FFT	RedMulE (MatMul)

Heterogeneity of the PULP Echosystem

Heterogeneity in the PULP Echosystem can be exploited at **multiple levels**

- **Cluster** Level: parallel **general-purpose** RISC-V cores + **application-specific** hardware accelerators
- **Address Space** Level: **64-bit** controller core + **32-bit** general-purpose multi-core (or multi-cluster) accelerator

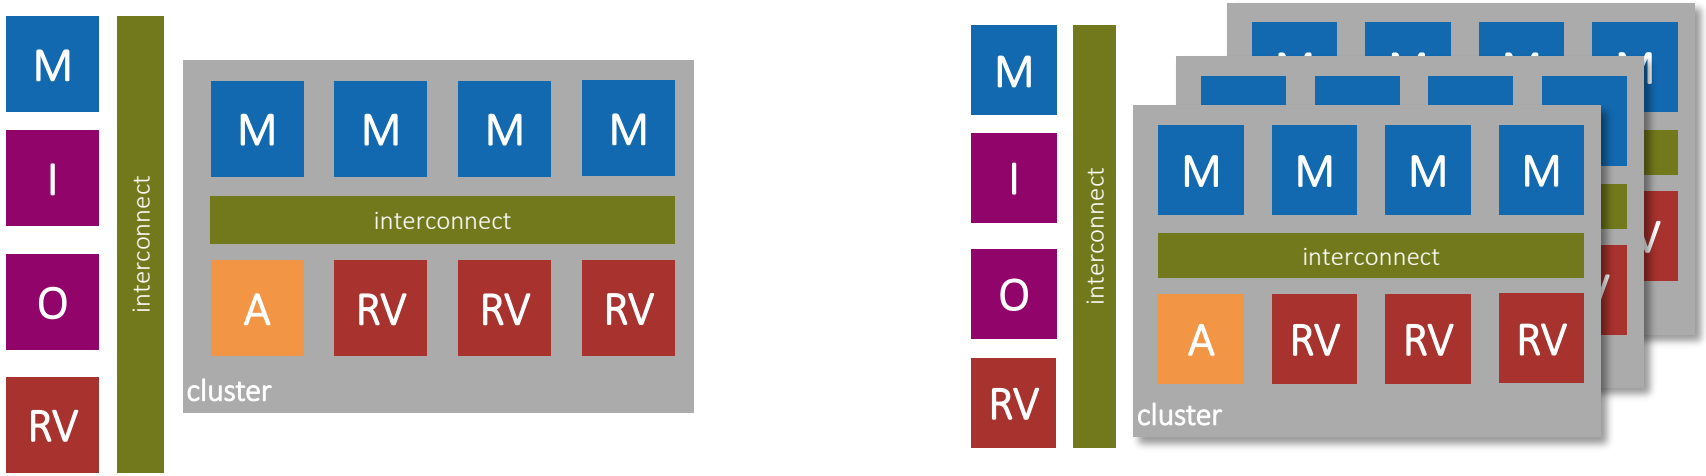

Timeline of Parallel Ultra Low Power (PULP) project

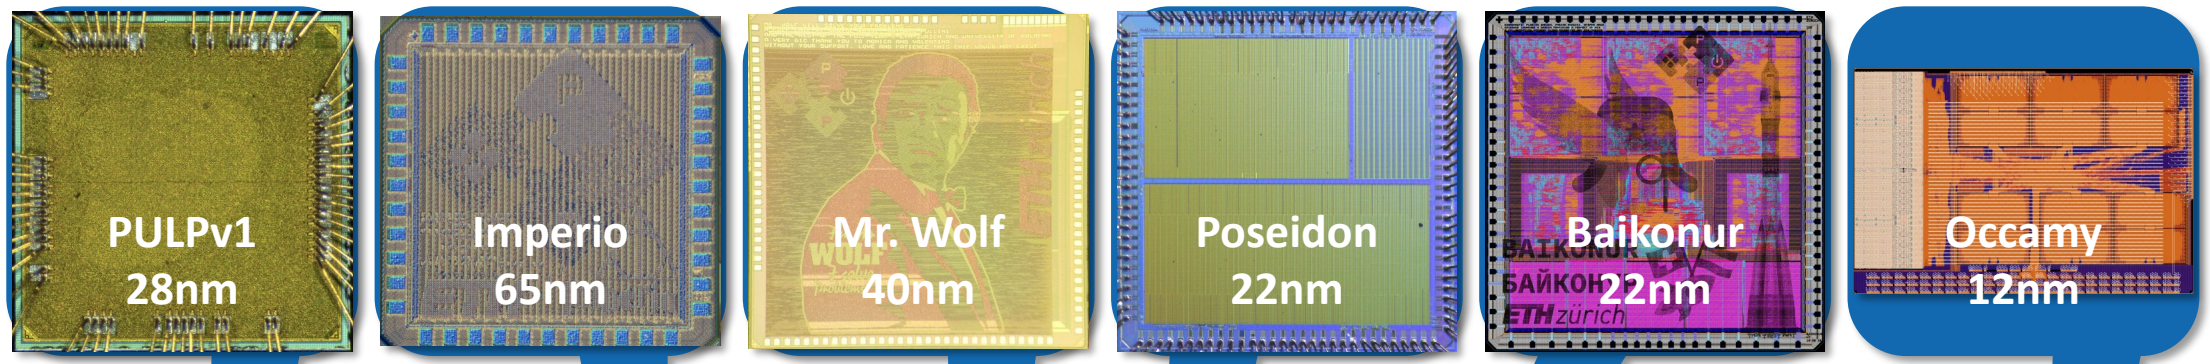

Manycore cluster

Heterogenous Acceleration

Application chips

64 bit cores

Vector Units

Snitch core

Scaling PULP for HPC

Start of PULP

Adopted RISC-V

First OSH Release

OpenHW Group

ZRH RISC-V Workshop

OSH within EC projects

European Processor Initiative

2013 2014 2015 2016 2017 2018 2019 2020 2021 2022 2023

Heterogeneous SoC for Space Applications

Bare-through-Linux Offloading

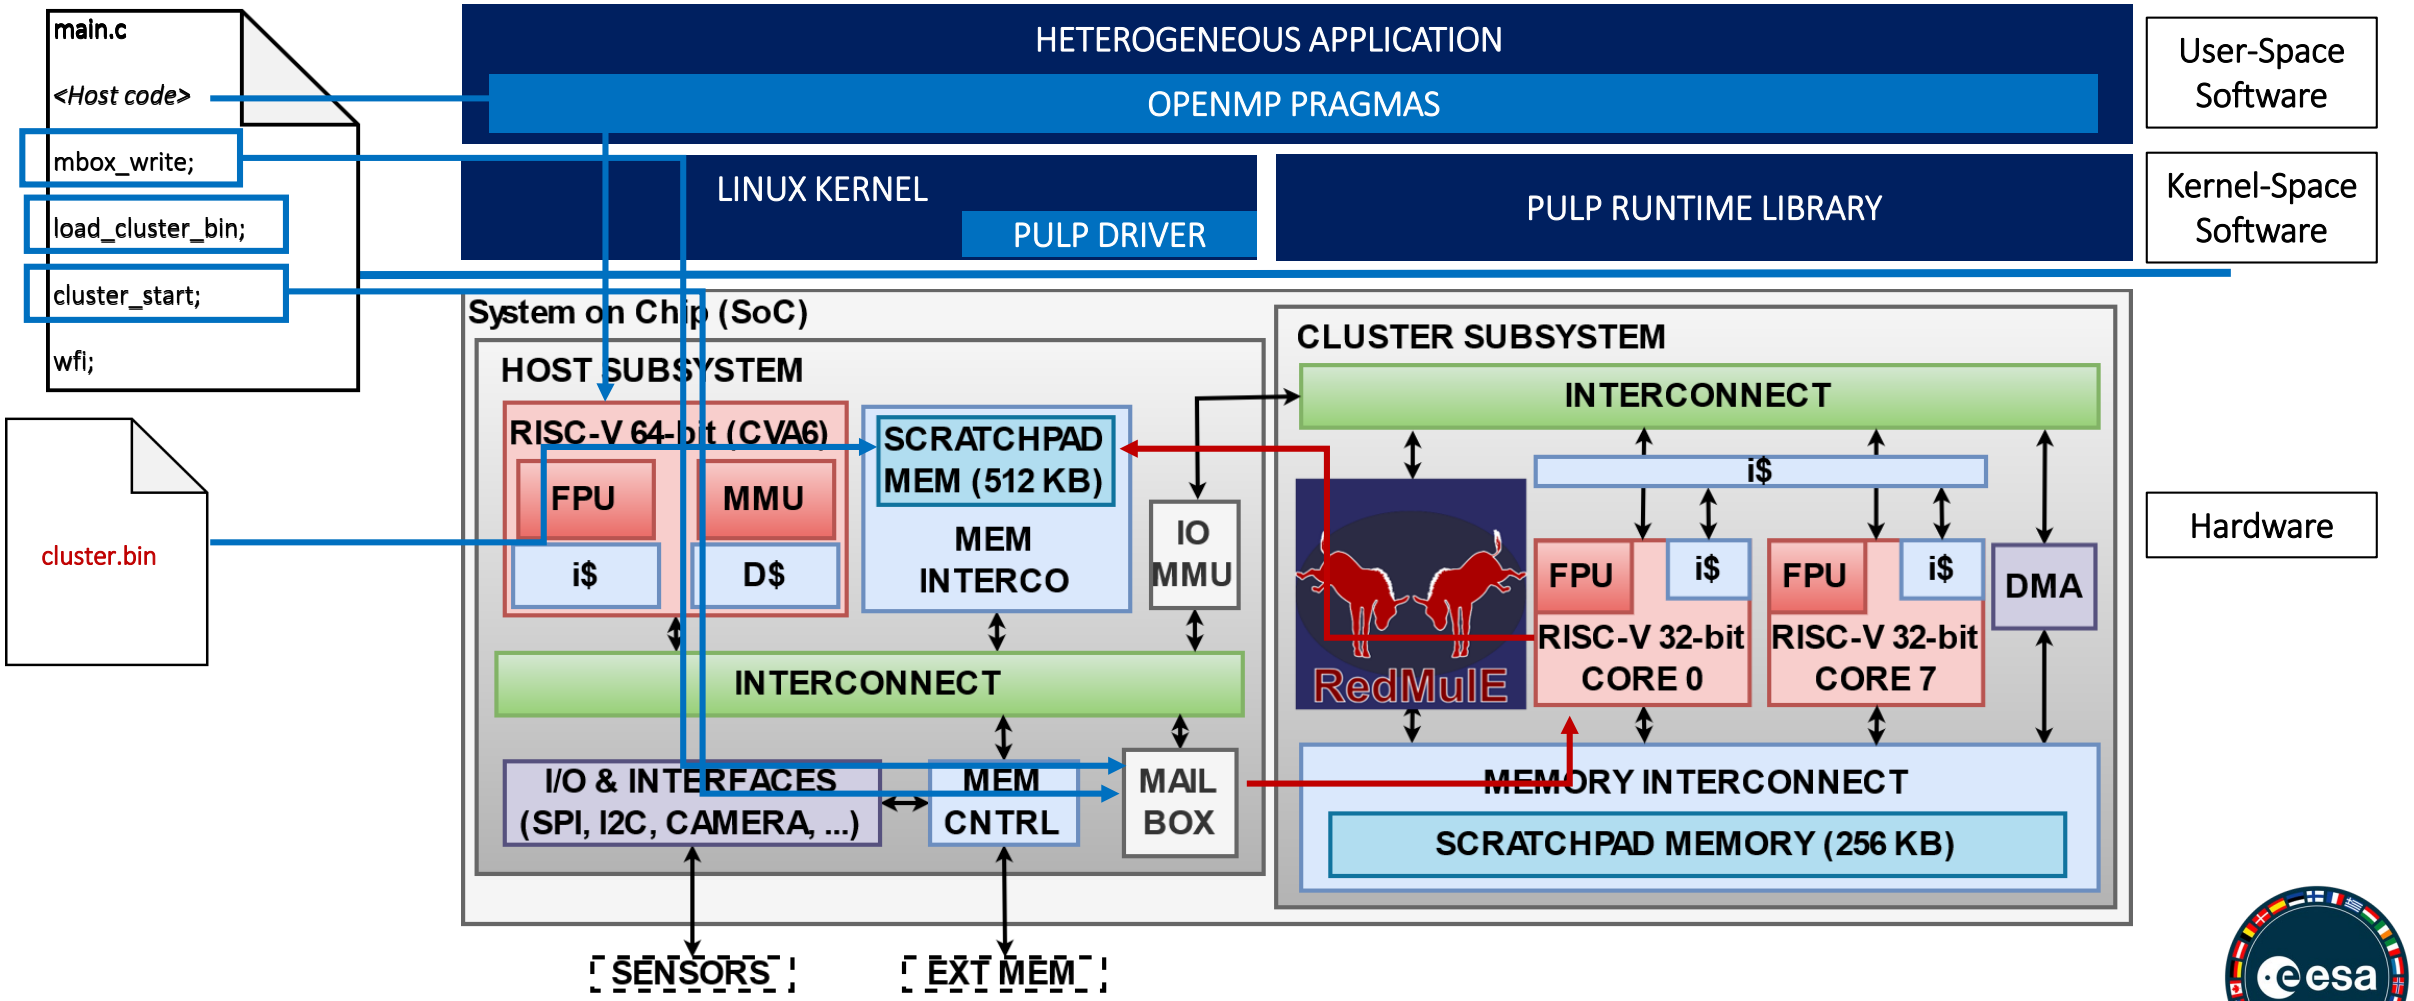

Open source collaboration scheme explained

Direct research collaborators on PULP

- Politecnico di Torino
- University of Cambridge
- USI Lugano
- TU Kaiserslautern
- University of Cagliari
- IBM Research Zurich
- EPF Lausanne
- CSEM Neuchatel
- Princeton University
- Technische Universität Graz
- CEA-Leti Grenoble
- Fraunhofer-Gesellschaft
- Sapienza Università di Roma

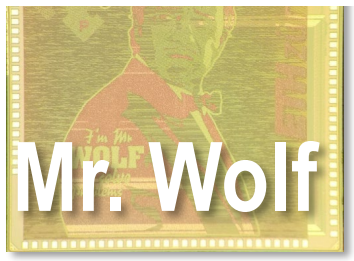

commits

commits

GitHub

Academic users we are aware of

- Università di Genova
- Politecnico di Milano
- Fondazione Bruno Kessler
- Lund University
- Stanford University
- UC Los Angeles
- UC San Diego
- Columbia University
- Universitat Bar-Ilan
- İstanbul Teknik Üniversitesi
- NCTU Hsinchu
- University of Zagreb, FER
- TUT Tampere
- RWTH Aachen
- IST University of Lisboa
- UFRN Rio Grande do Norte
- TU Darmstadt
- Universität Bremen
- Hongik University Seoul
- IIT Kharagpur
- LIRMM Montpelier
- University of Stuttgart
- University of Tübingen
- TU Münich
- FORTH Hellas
- Kyoto University
- Chalmers Göteborg
- FAU Erlangen-Nürnberg
- NTNU Trondheim

The open model led to successful industry collaborations

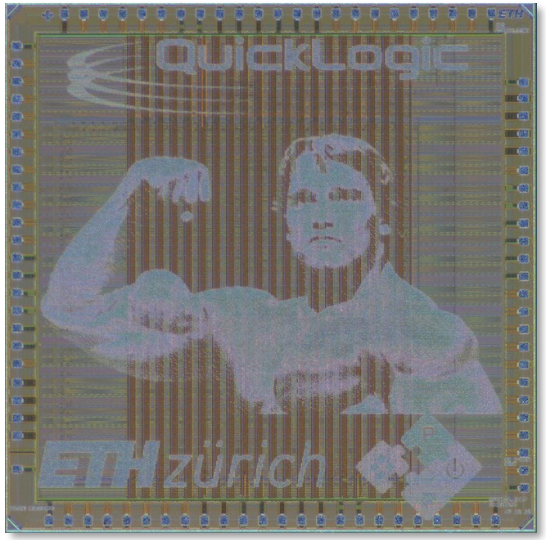

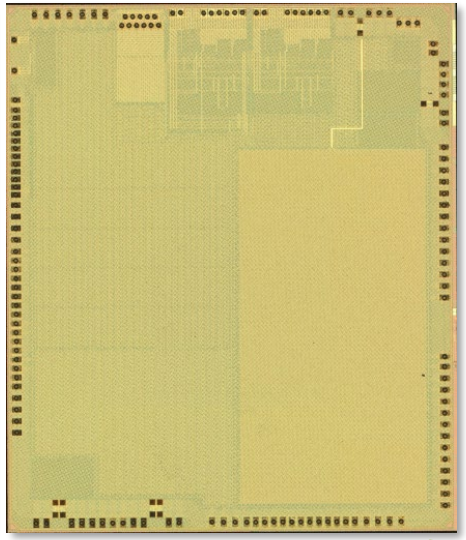

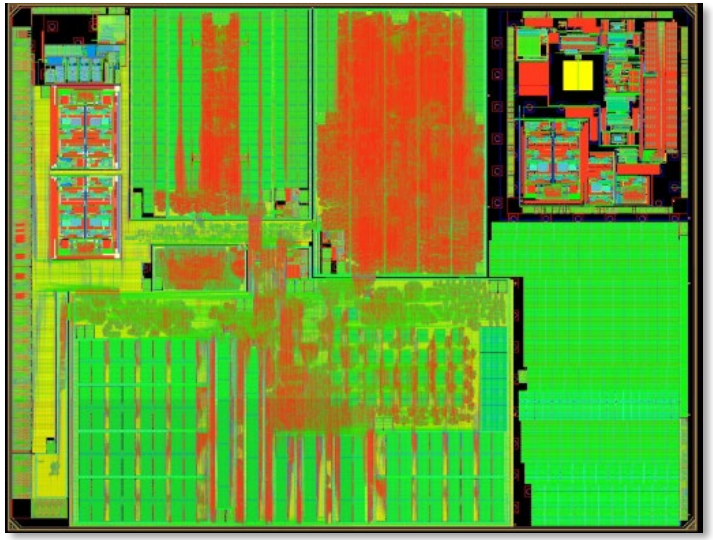

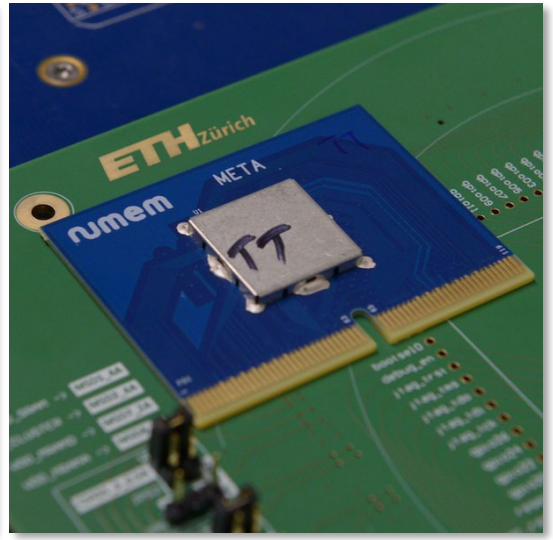

QuickLogic®

GREENWAVES TECHNOLOGIES

DOLPHIN INTEGRATION
The enabler of low-power Systems-on-Chip

Meta numem

Arnold (GF22)
eFPGA with RISC-V core

Vega (GF22)
IoT Processor with ML acceleration

Marsellus (GF22)
IoT Processor with low power modes and event based computing

Siracusa (TSMC16)
IoT Processor with NVM technology