ETH zürich ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA

Open-source SoC design using PULP

Austrochip 2025, Linz

Frank K. Gürkaynak kgf@iis.ee.ethz.ch

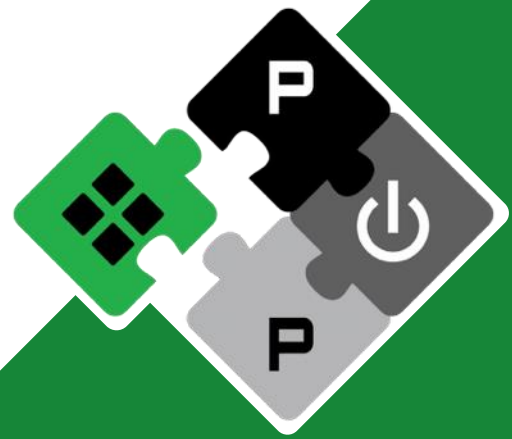

@pulp_platform

pulp-platform.org

youtube.com/pulp_platform

PULP Platform

Open Source Hardware, the way it should be!

Team of 100 people in ETH Zürich – University of Bologna

- **Research on open-source energy-efficient computing**

In 12 years PULP team has designed more than 60 chips

RISC-V and open-source hardware have been instrumental in our success

Including complex chiplet based designs like this one

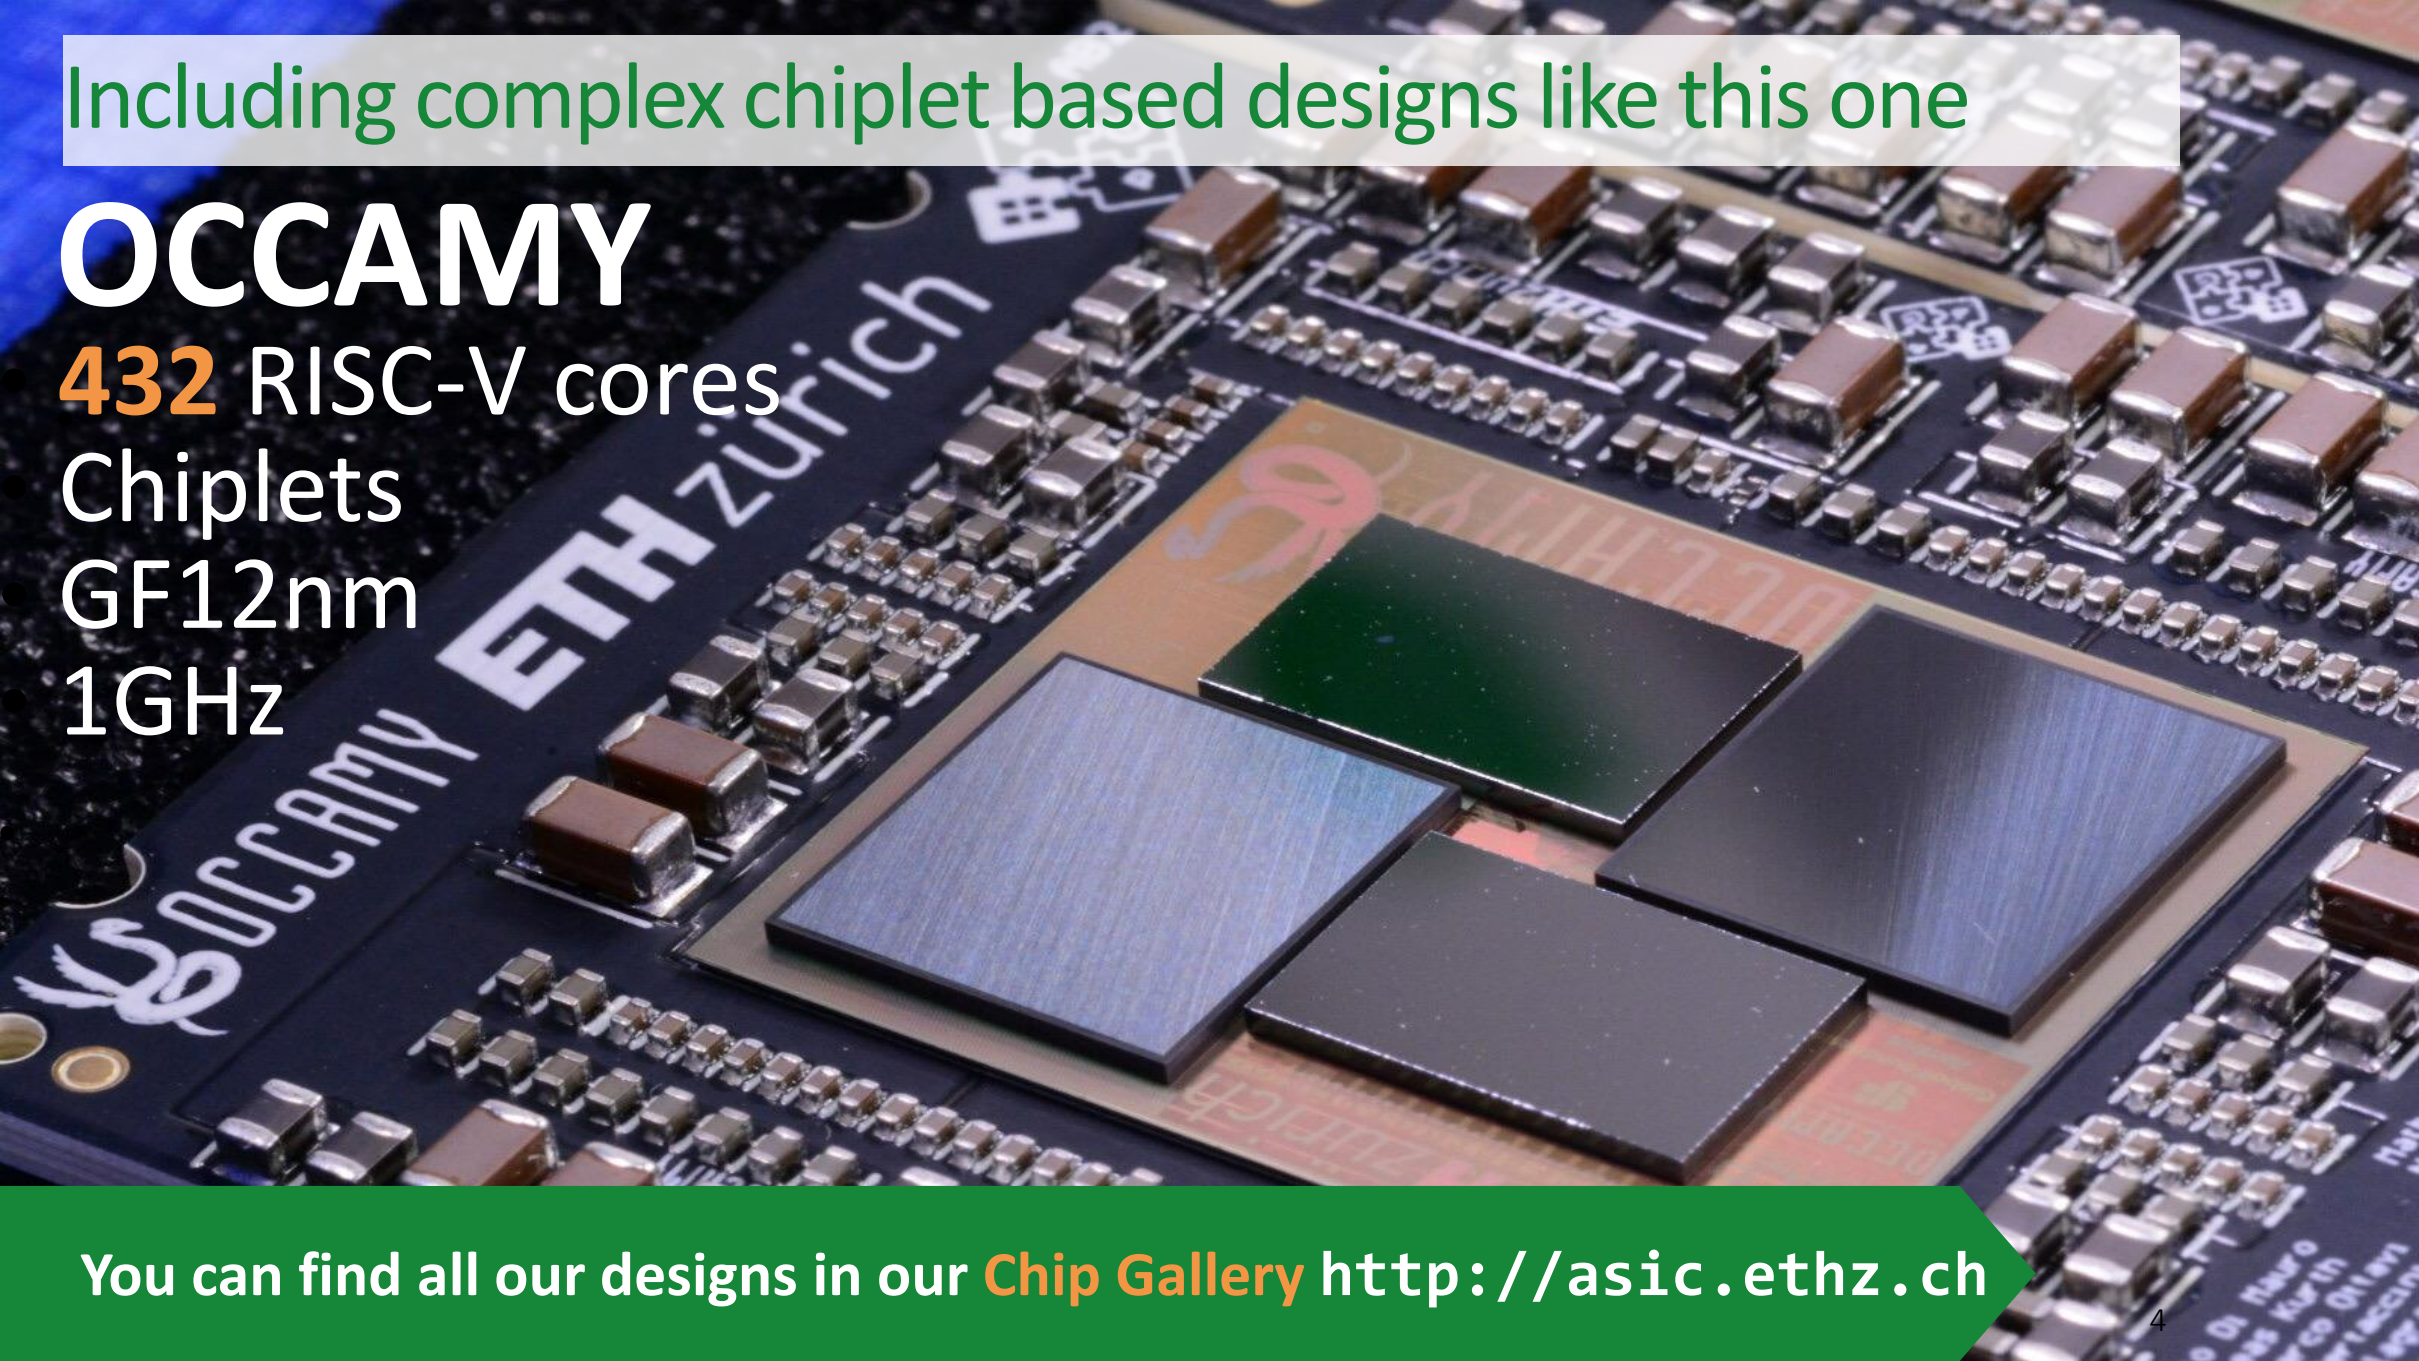

OCCAMY

- **432** RISC-V cores
- Chiplets
- GF12nm
- 1GHz

You can find all our designs in our Chip Gallery `http://asic.ethz.ch`

What I want to share today

Why open source?

- Managing complexity, Collaboration, Exploitation, Education

A sample of our projects

- Croc, Cheshire, Kraken, Occamy, Picobello

What is next for open-source hardware

- Challenges and opportunities

Why open source?

- Managing complexity
- Collaboration
- Exploitation
- Education

In the last 20 years IC Design has changed a lot

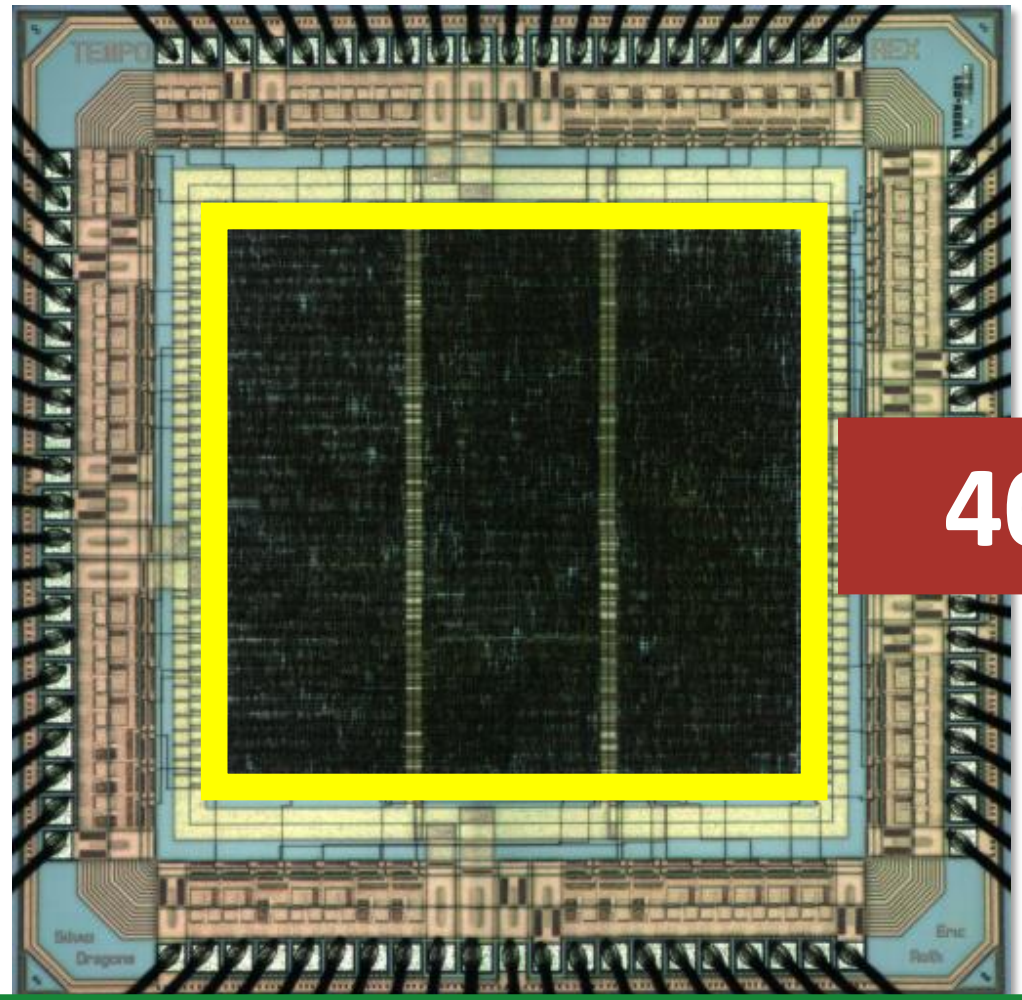

4000 x

80 MGE

What used to be a complete chip is now a small part of a SoC !

There is so much that makes up a modern SoC

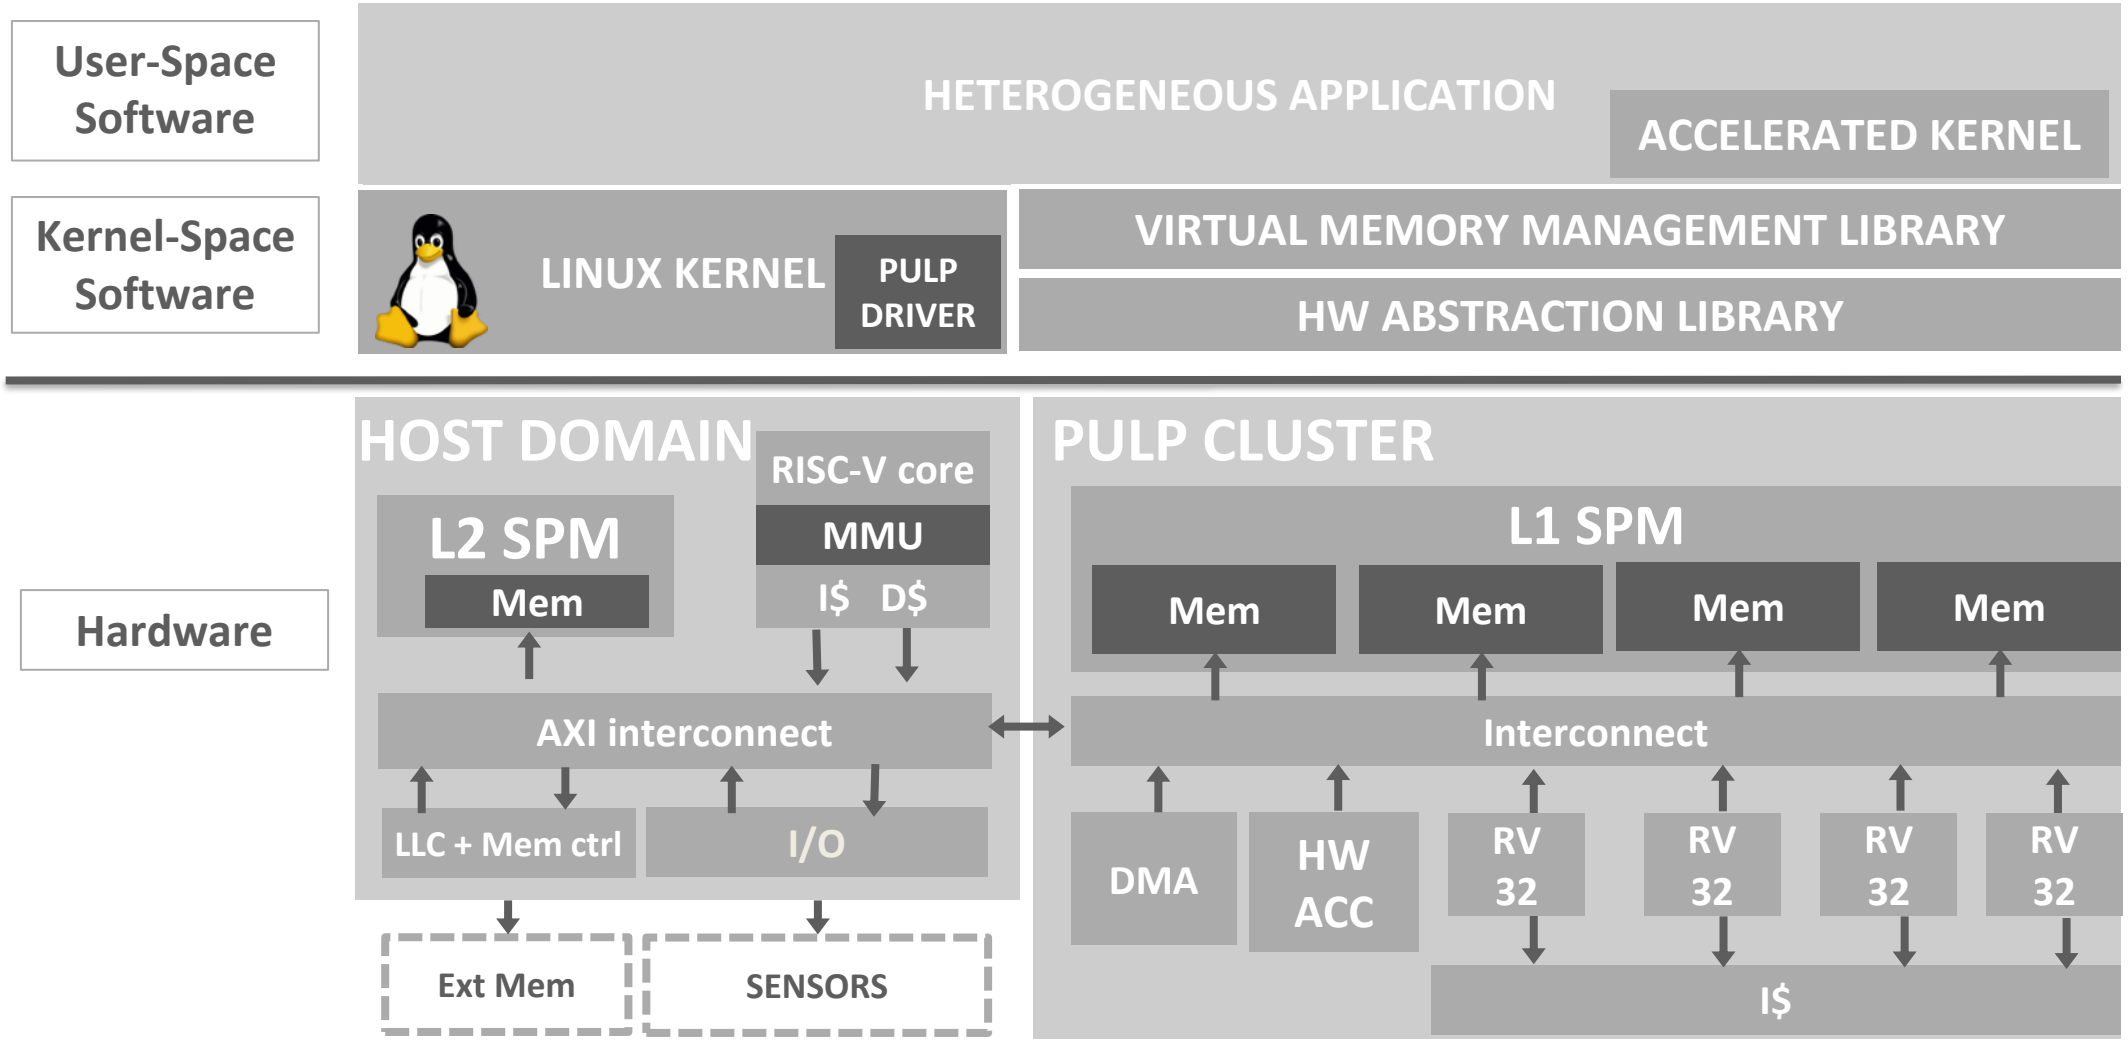

In a typical design, innovation is only in a limited scope

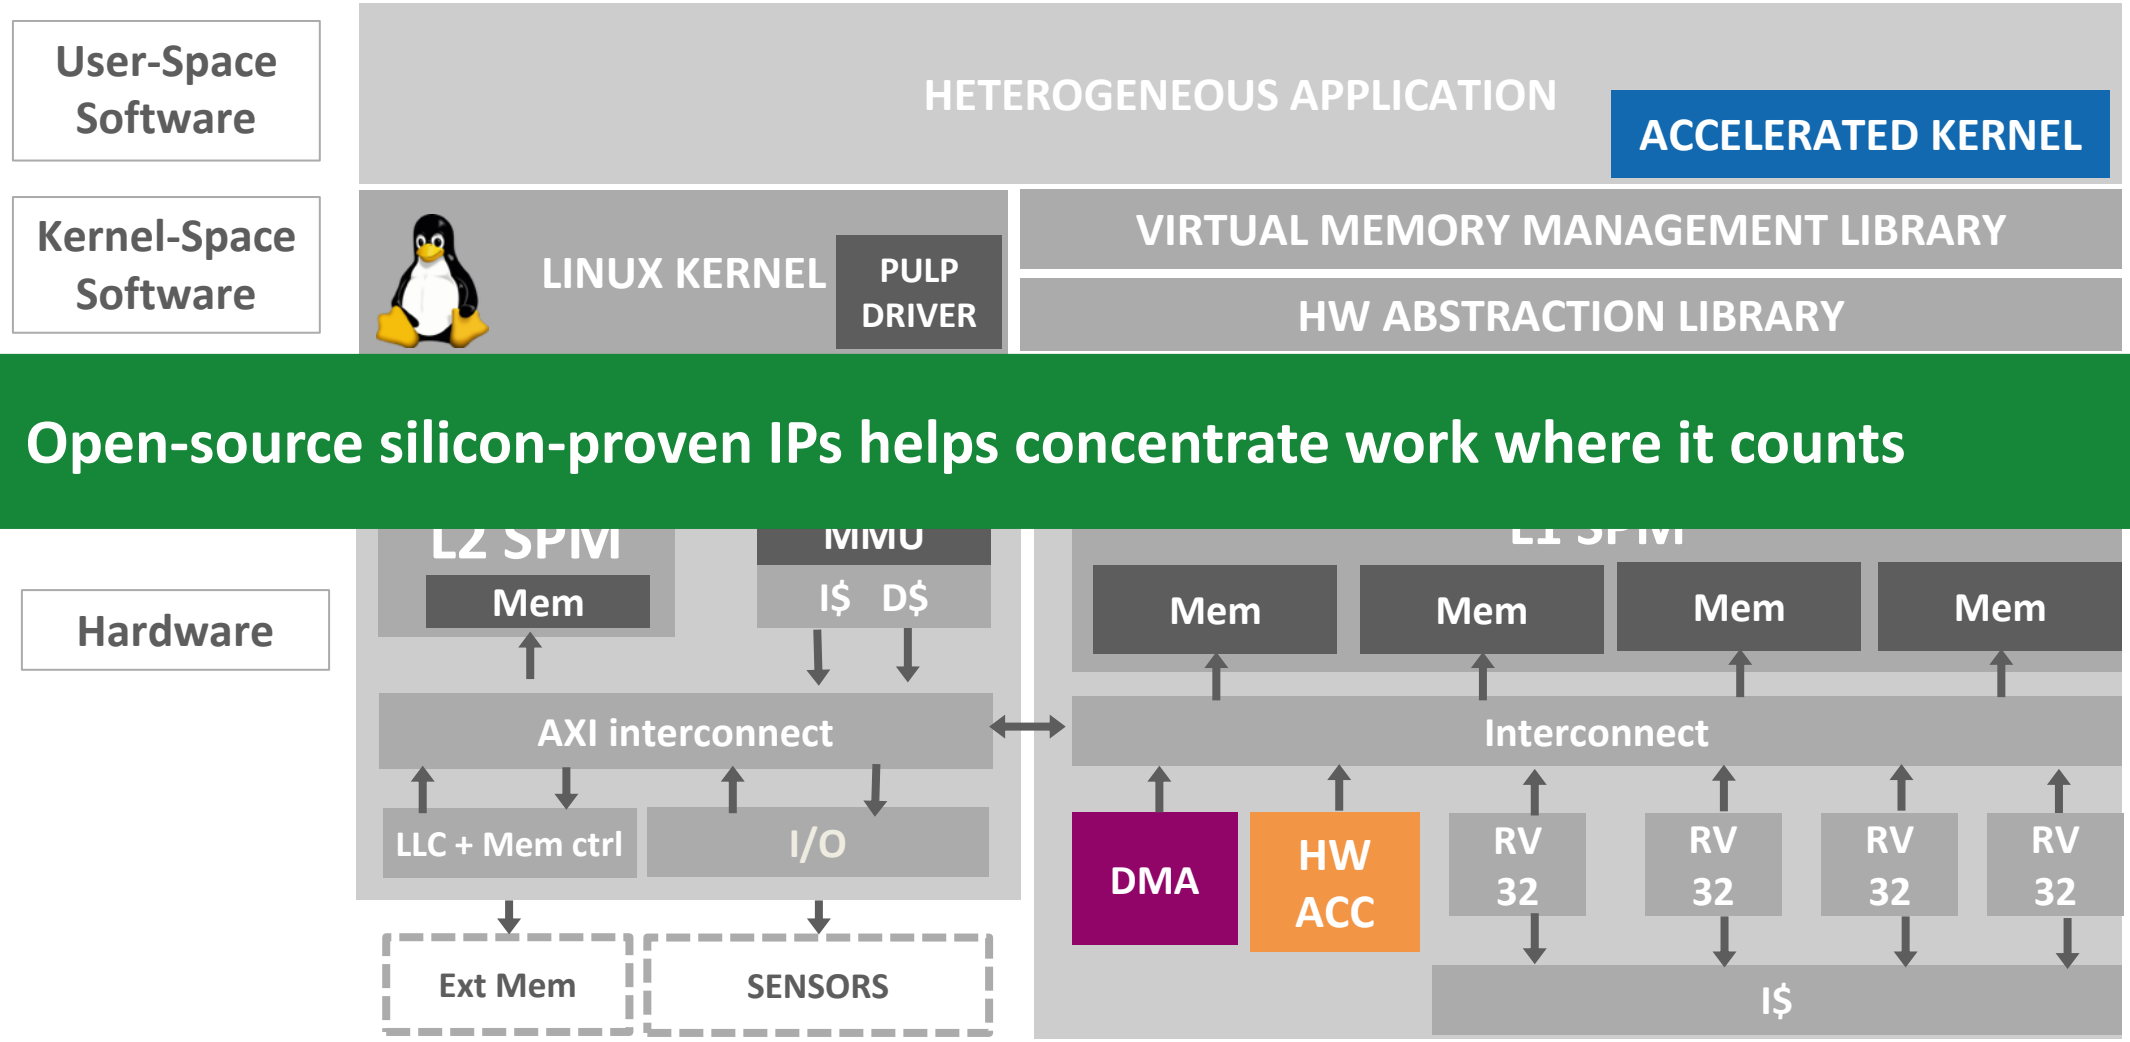

All of our designs are open-source hardware

- **All our development is on GitHub using a permissive license**
  - HDL source code, testbenches, software development kit, virtual platform

https://github.com/pulp-platform

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

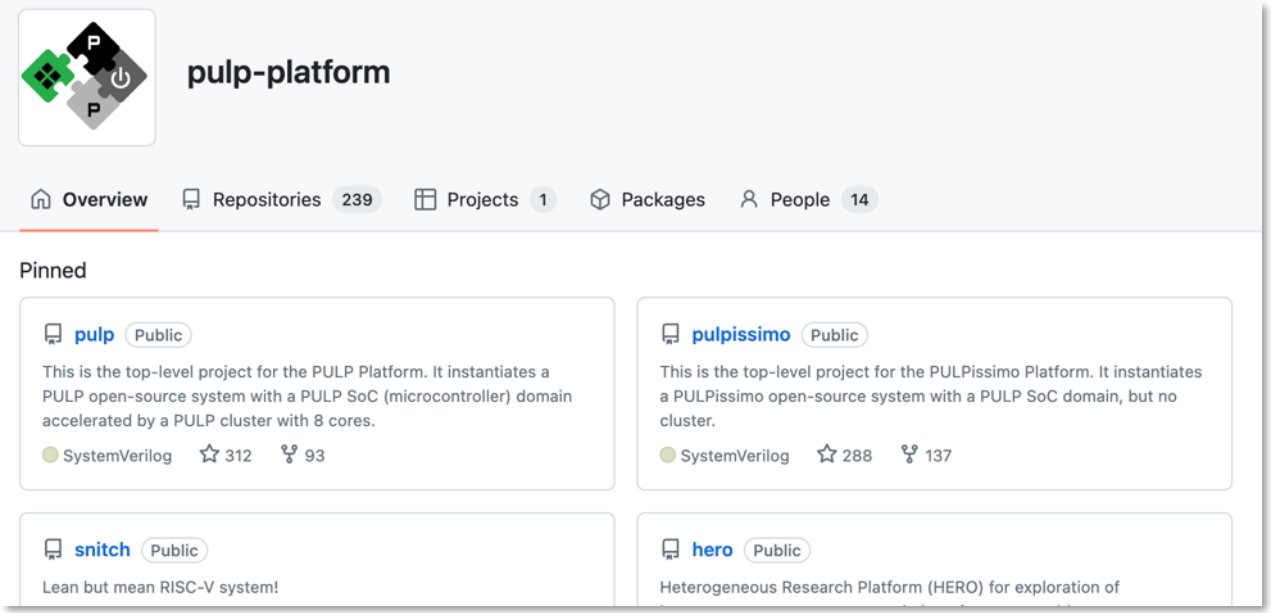

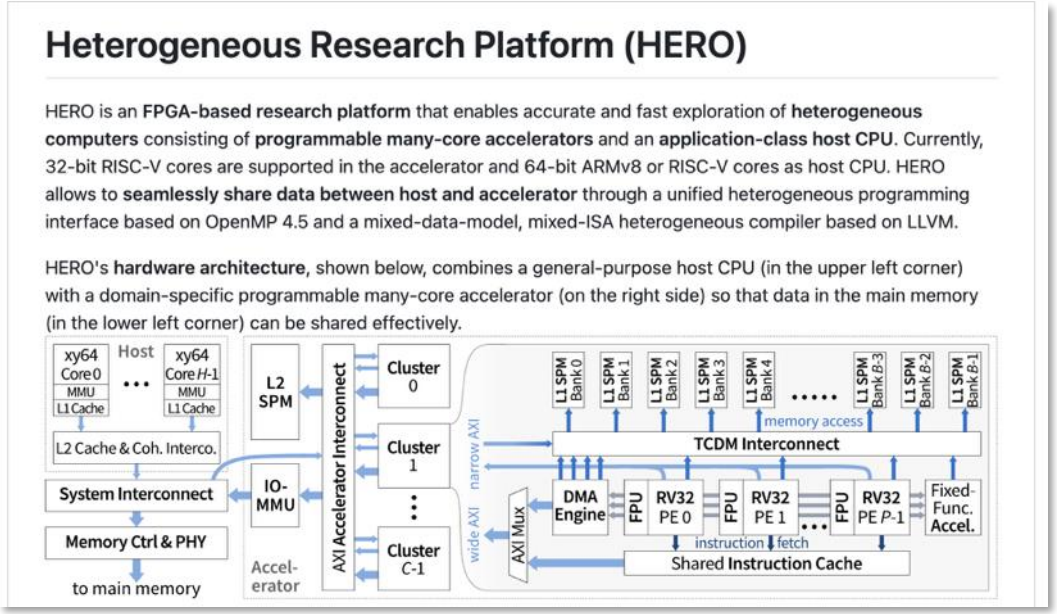

What PULP provides is a box of building blocks

RISC-V Cores and Vector Units

RI5CY *CV32E*	Zero R *Ibex*	Snitch	Spatz	Ariane *CVA6*	ARA
RV32	RV32	RV32	RVV	RV64	RVV

Peripherals

JTAG, SPI, UART, I2S, DMA, GPIO

Interconnects

LIC, HCI, APB, FlooNoC, AXI4

Platforms

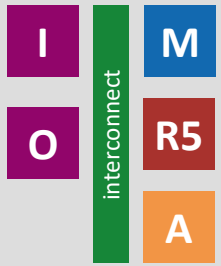

Single core
- Croc, PULPissimo
- Cheshire

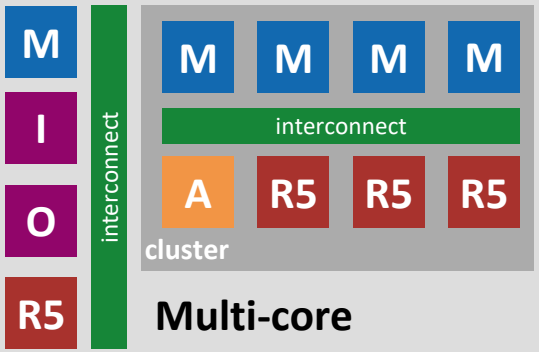

Multi-core
- OpenPULP
- ControlPULP

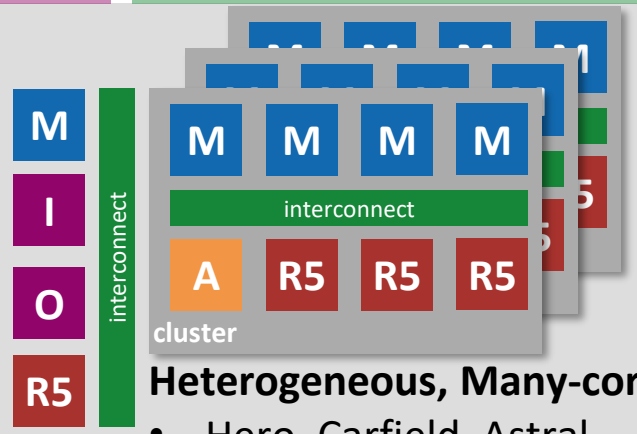

Heterogeneous, Many-core
- Hero, Carfield, Astral
- Occamy, Mempool

IOT → **HPC**

Accelerators and ISA extensions

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

How does PULP collaborate with 3rd parties?

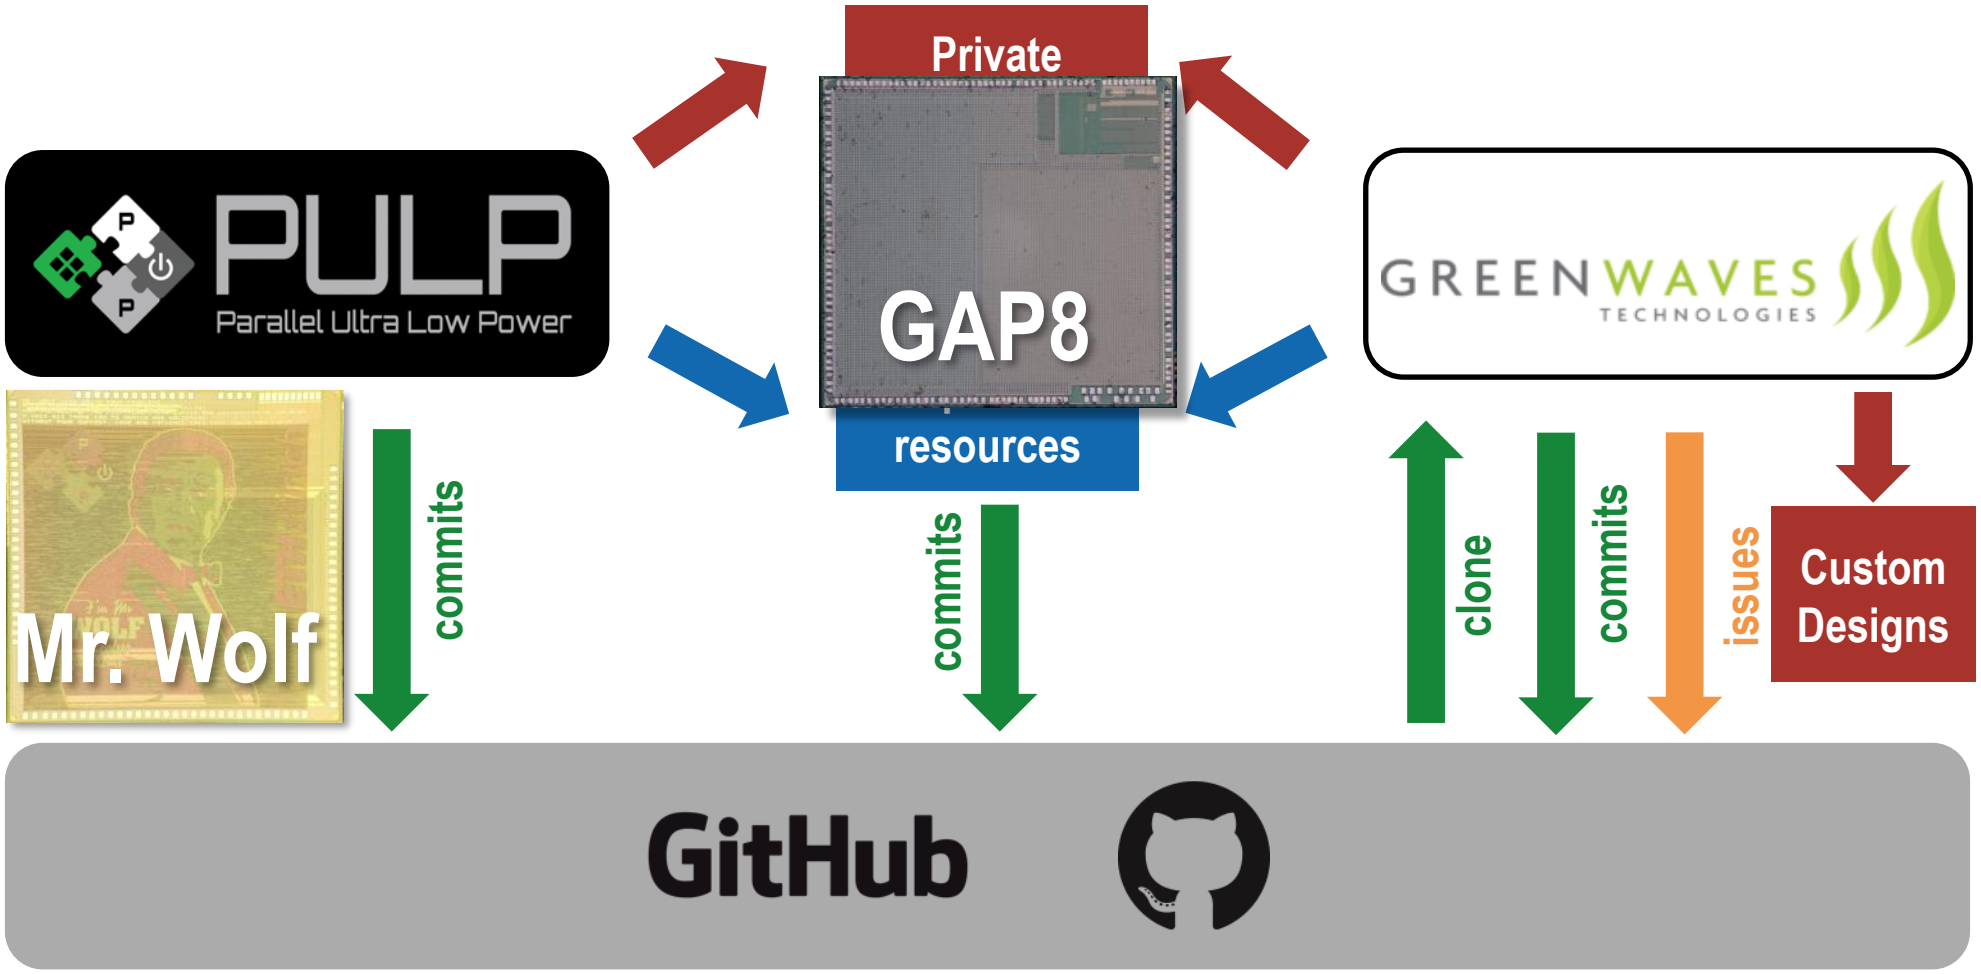

Diverse set of open source based industry collaborations

GF22 (2018)

Arnold

eFPGA coupled with a RISC-V microcontroller.

In one year from agreement to actual tapeout

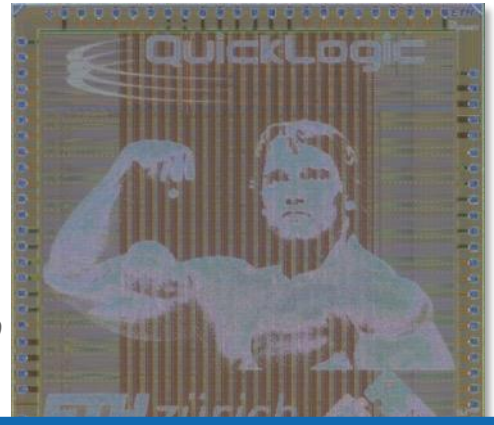

GF22 (2022)

Marsellus

Heterogeneous IoT processor With Aggressive voltage scaling

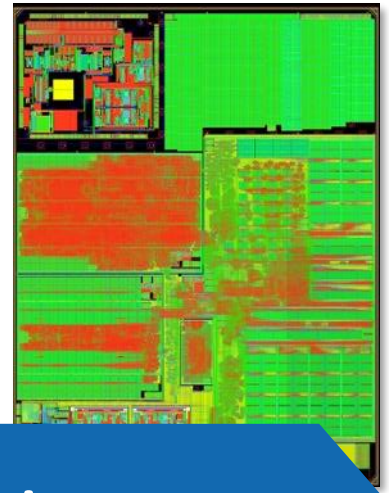

Permissive open-source licensing key to our industrial relationships

Siracusa

SoC for Extended Reality visual processing

Meta

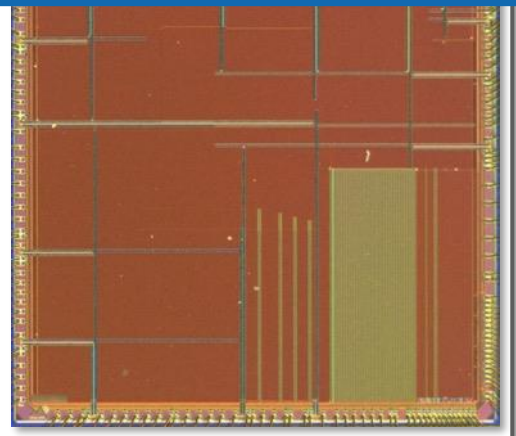

Carfield

Open-Research platform for safety, resilient and time-predictable systems

intel

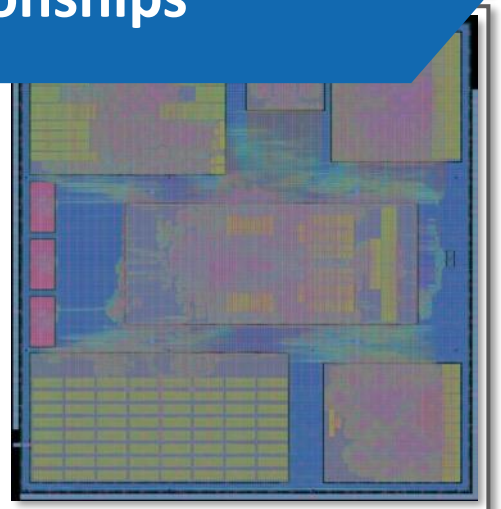

And many continue to use our work for their research

RISC-V week Barcelona 2018

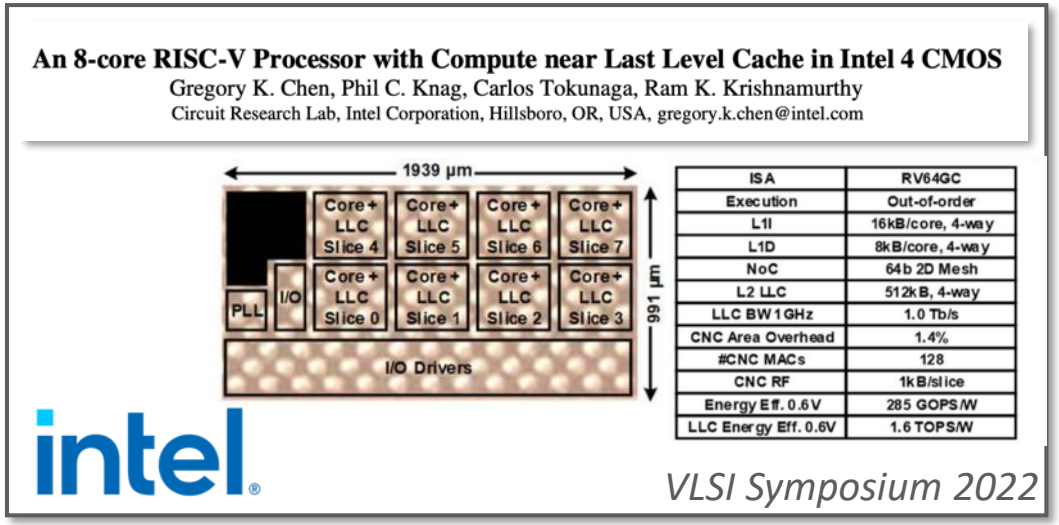

VLSI Symposium 2022

ISSCC Keynote 2020 – Nature 2020

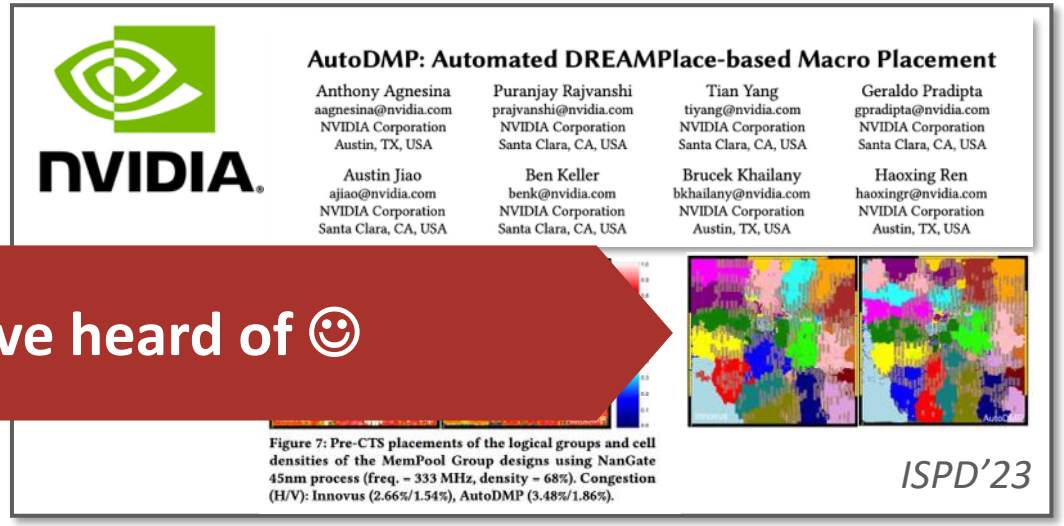

ISPD'23

Some smaller companies you might have heard of ☺

We rely more and more on open-source HW for teaching

- **Open source RTL is standard for all student projects**
  - Our research projects have been based on open source RTL anyway
  - **Legal aspects**: In CH we need the *permission* of students to make their code openly available
- **Open source EDA and PDK allows exercises to be transferred anywhere**
  - We can share our exercises with others
  - No need to rely on costly infrastructure to provide students a place to run exercises

Thanks to the great IIC-OSIC tools from JKU

and the open-source PDK from IHP

At ETHZ, IC Design teaching now uses open source HW

- **In Spring 2025, our IC Design course switched to (mostly) open source**
  - Using IHP 130, Yosys and OpenROAD
    - Parts for backannotated simulation, test pattern generation, DRC/LVS, still use proprietary tools
    - Will be gradually replaced by open tools

https://vlsi.ethz.ch

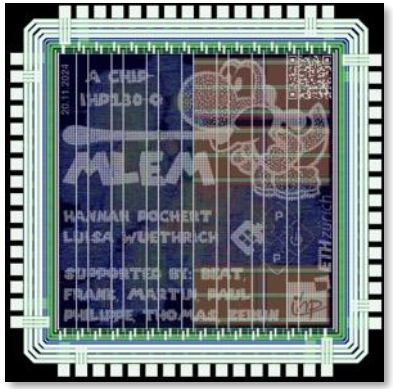

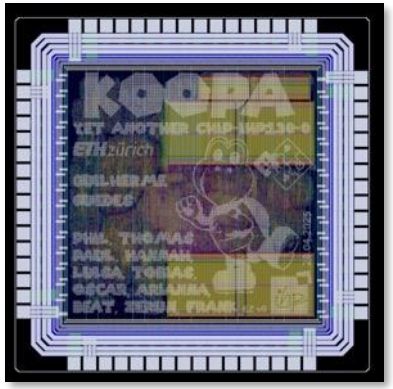

- **Project based grading**
  - Students (in groups of two) will have to modify the Croc reference design
  - Best five designs will be taped-out
- **72 students enrolled**
  - Projects finished in summer
  - Taped-out 5 designs in IHP130

And the students delivered: 33 valid designs, 5 taped-out

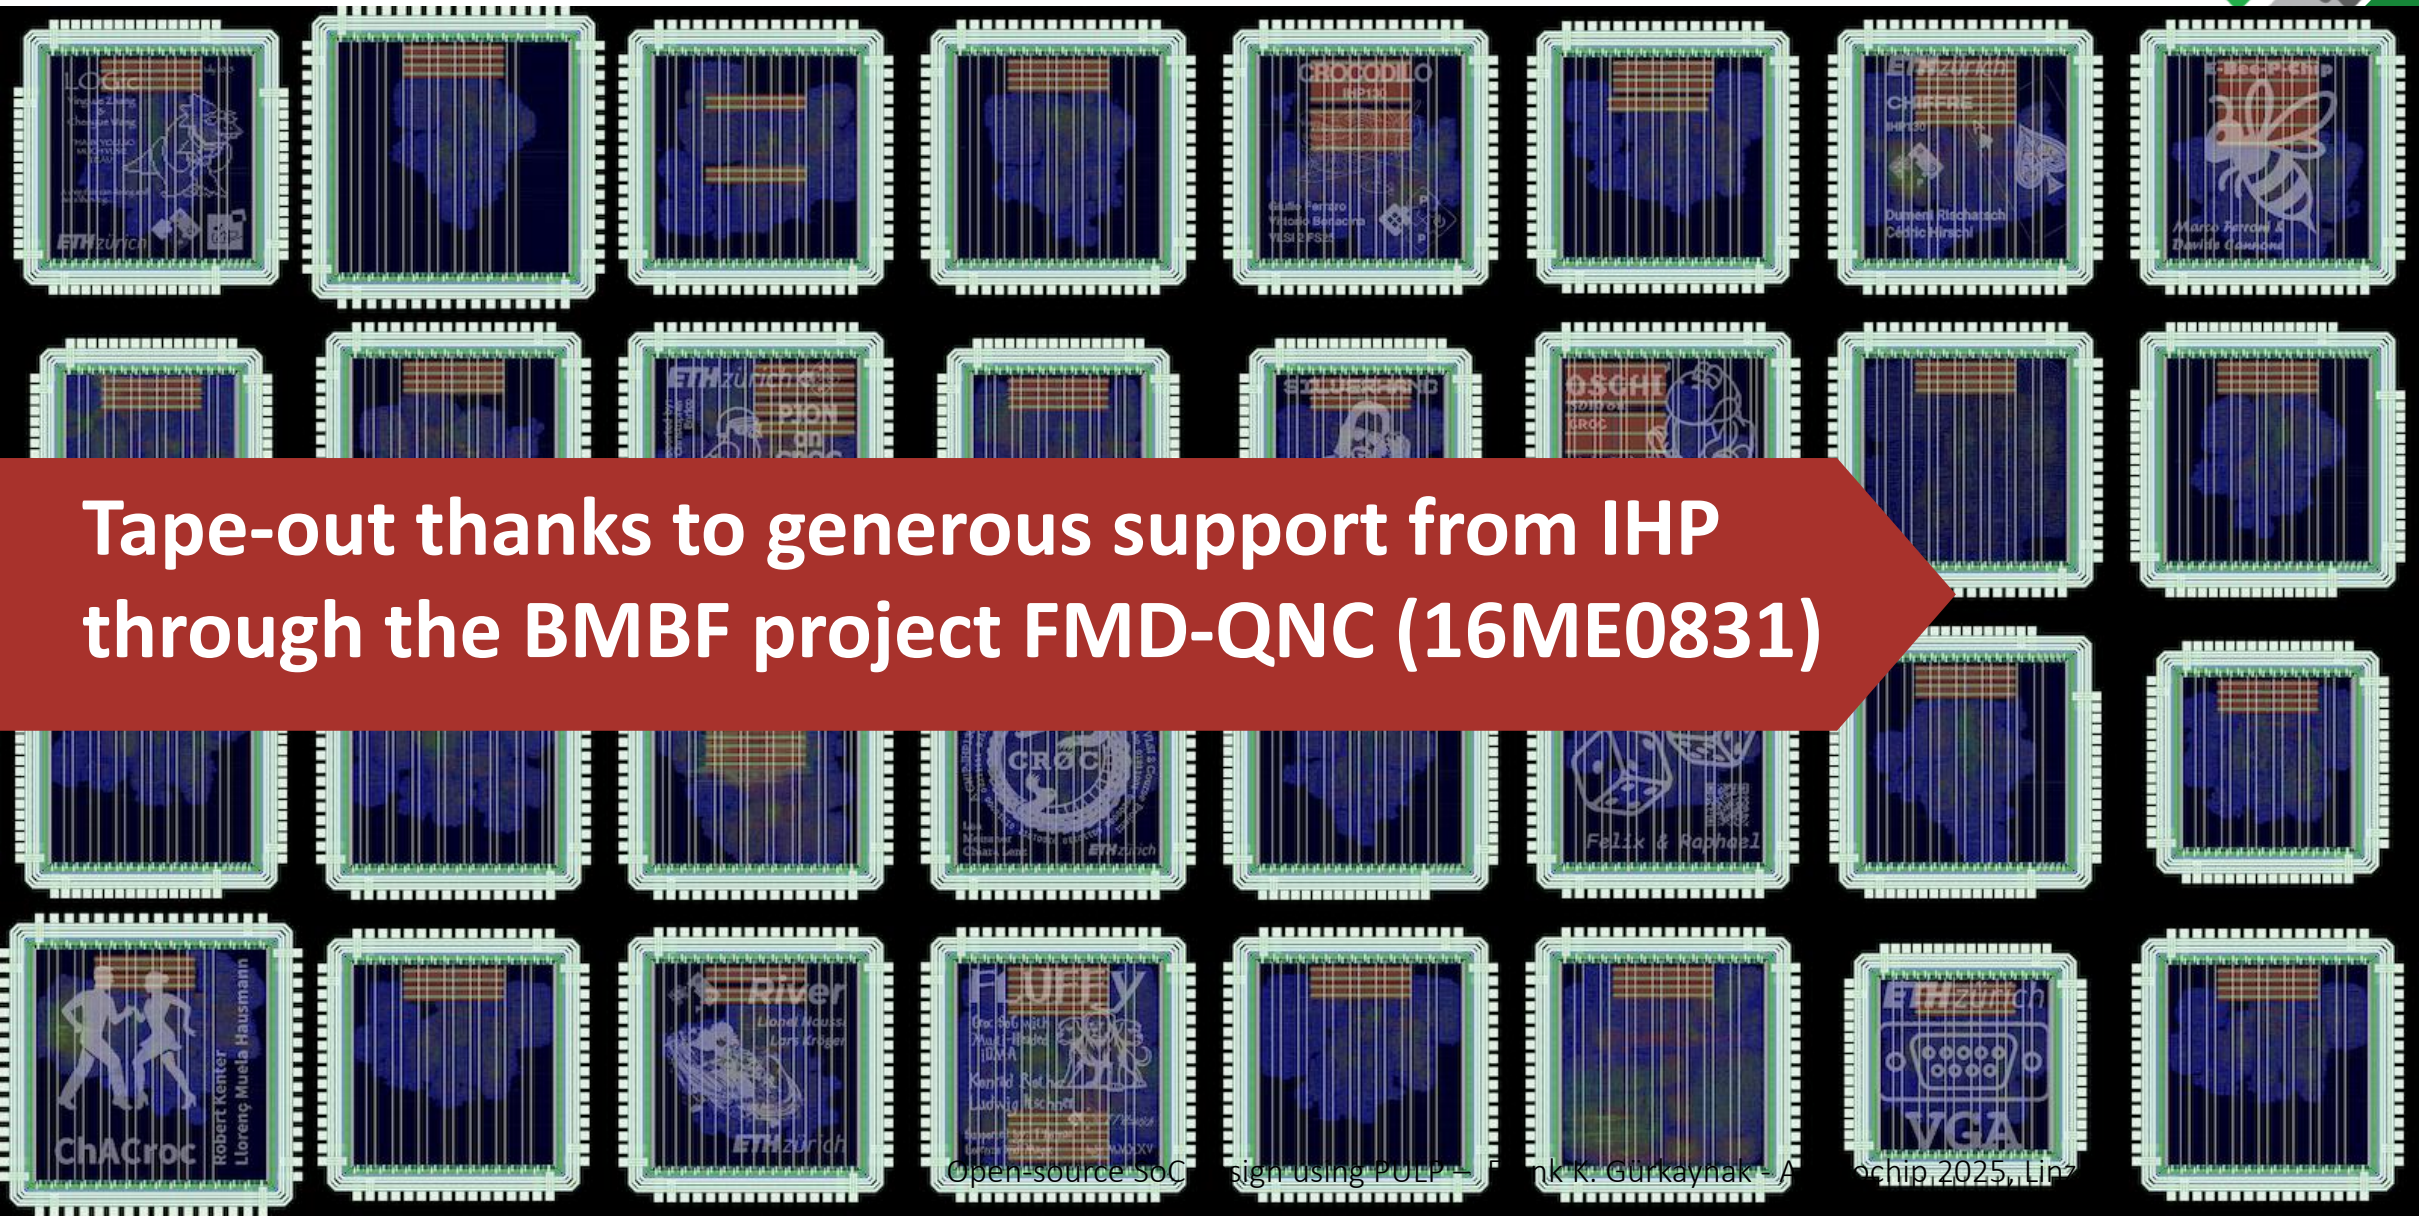

Tape-out thanks to generous support from IHP through the BMBF project FMD-QNC (16ME0831)

A sample of our projects

- Croc
- Cheshire
- Kraken
- Occamy
- Picobello

Croc our end-to-end open SoC for teaching and training

- **Scalable ULP design**
- **32-bit RISC-V Core**
  - Complete SoC
  - Simple “Raspberry Pi”
- **Rich Peripherals**
- **Ready for Acceleration**
  - Digital-only interface
- **Silicon-proven**
  - Tapeouts with open & commercial EDA

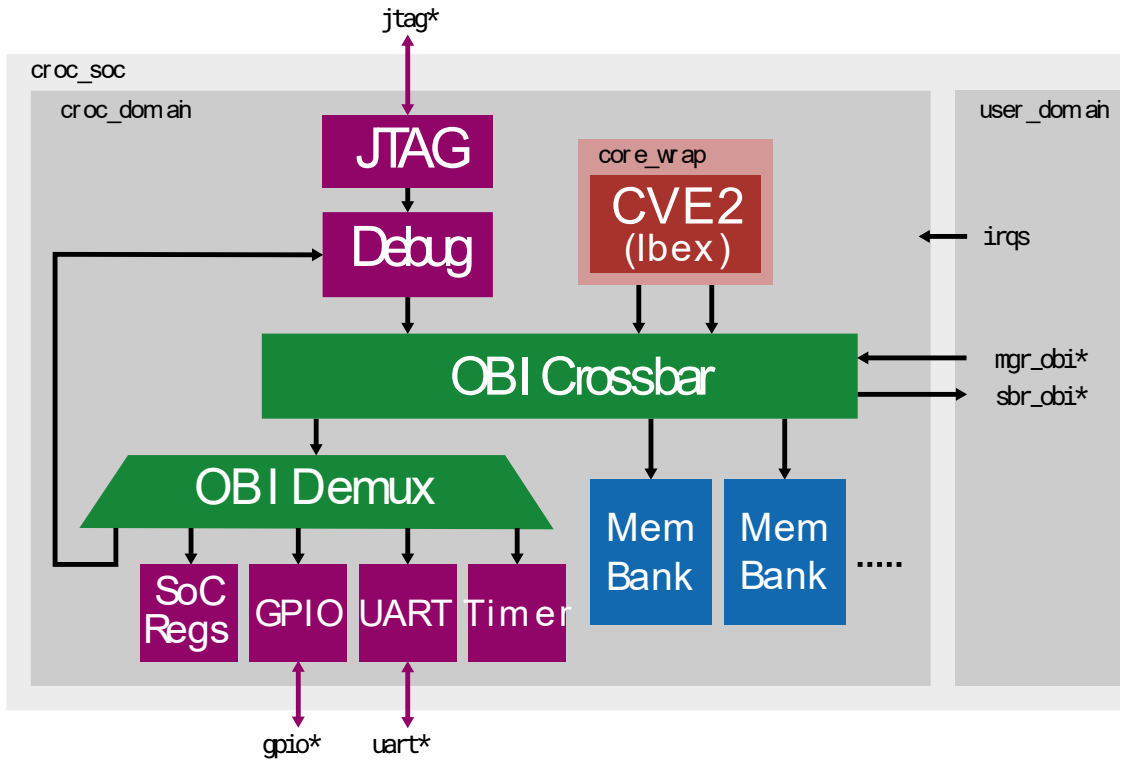

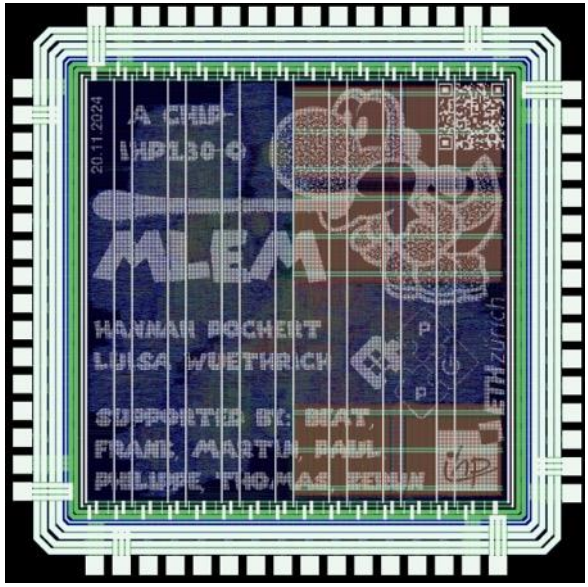

github.com/pulp-platform/croc

CVA6 (aka Ariane): The standard 64bit core

- Open-source 64-bit application-class RISC-V processor
- Boots Linux
- Developed by PULP team at ETH Zürich (as “Ariane”)
- Now owned and maintained by OpenHW Group
- Widely used in academia and industry

RTL Simulation

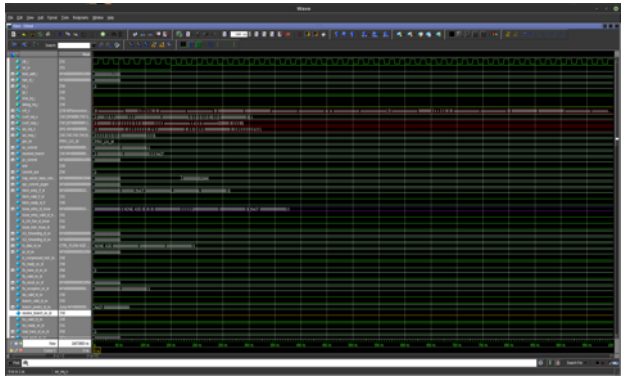

FPGA Emulation

ASIC

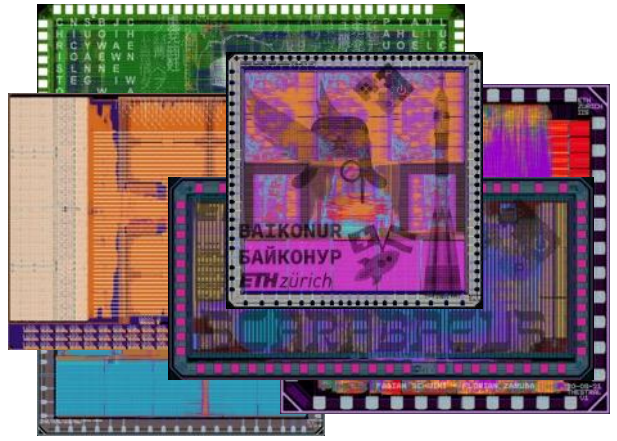

https://github.com/openhwgroup/cva6

CVA6: Many optional features added over time

- XLEN 32/64 bits (RV32/RV64)
- WB/WT/HPDcache L1 data cache
- Embedded/application class (MMU, U/S mode)
- Vector extension (V)
- Hypervisor extension (H)
- Code-size extensions (Zcb, Zcmp)
- Bit-manip extensions (Zba, Zbb, Zbc, Zbkb Zbkx, Zbs)
- Scalar crypto extensions (Zknd, Zkne, Zknh)
- ECC support
- Fast interrupts (CLIC)
- ACE support
- …

Cheshire SoC: A SoC around CVA6

- **Permissive licensed RTL from various groups**
  - PULP Platform (Cheshire, AXI, Hyperbus...)
  - OpenTitan (SPI, I2C)
  - OpenHW Group (CVA6 core)
  - SpinalHDL (USB)
- **Architecture**
  - **RV64GC**-compliant CVA6 core
  - **Cheshire SoC** provides vital peripherals
  - 4-way 16KiB L1-D and L1-I cache
  - 4-way 64KiB LLC / Scratchpad memory
  - **Hyperbus DRAM controller** achieving transfer speeds up to 124MB/s

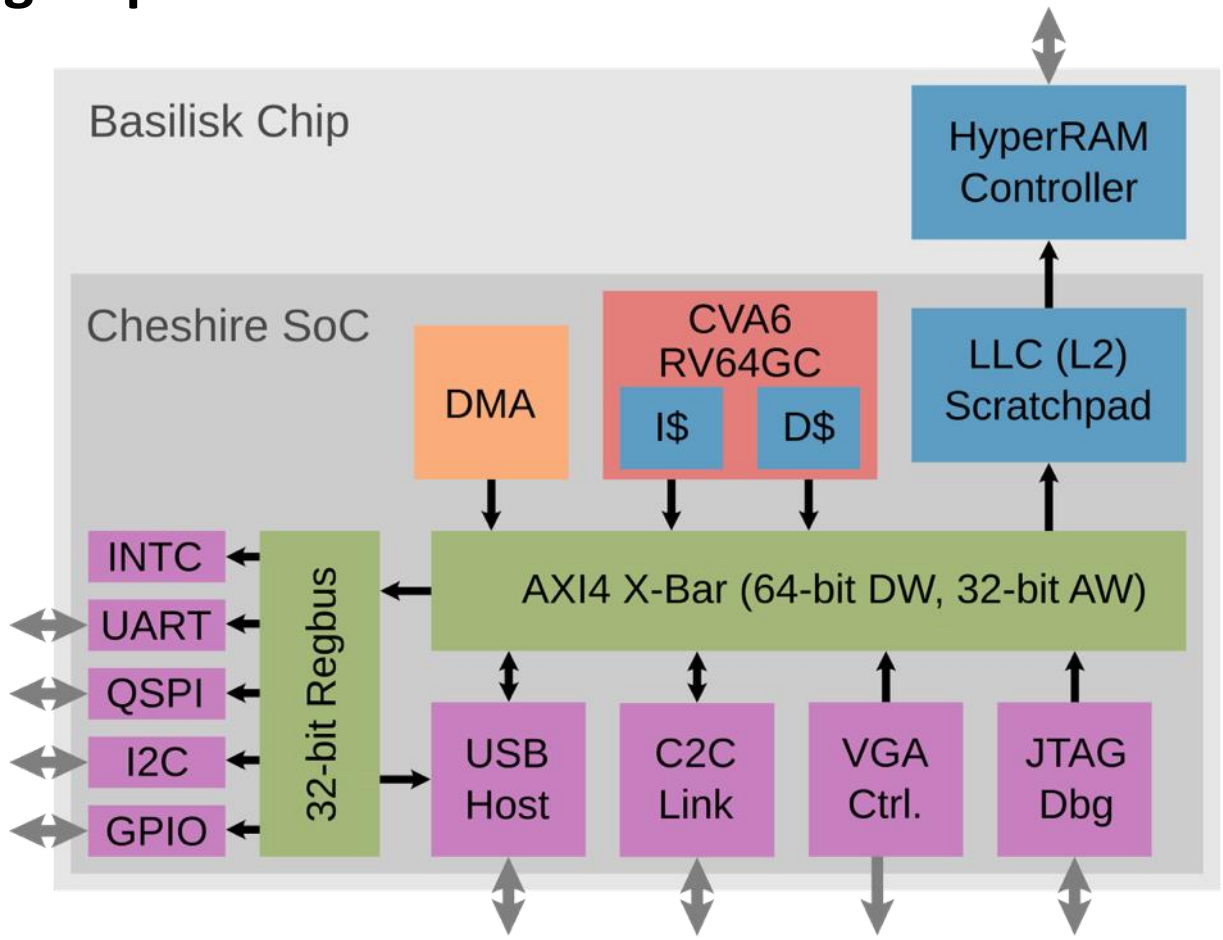

github.com/pulp-platform/cheshire

Meet **Basilisk**: Open **RTL**, Open **EDA**, Open **PDK**

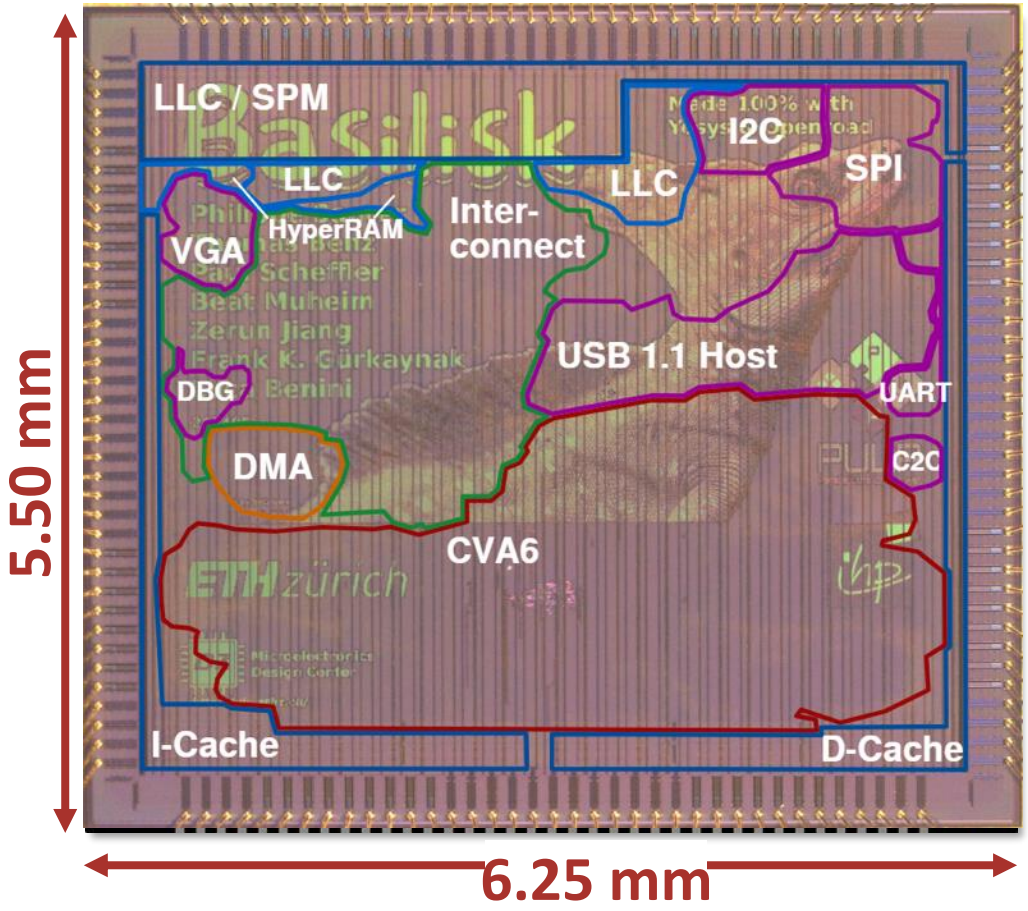

- **Designed in IHP 130nm OpenPDK**
  - **34mm²** (6.25mm x 5.50mm)
  - **~5× larger** than previous end-to-end OS designs
  - 2.7 MGE total, 1.14MGE logic
  - 24 SRAM macros (114 KiB)
  - **62MHz** at nominal voltage (1.2V)
- **RV64GC design runs Linux**
- **Active collaboration with**

github.com/pulp-platform/cheshire-ihp130-o

Basilisk is a complete Linux-capable SoC

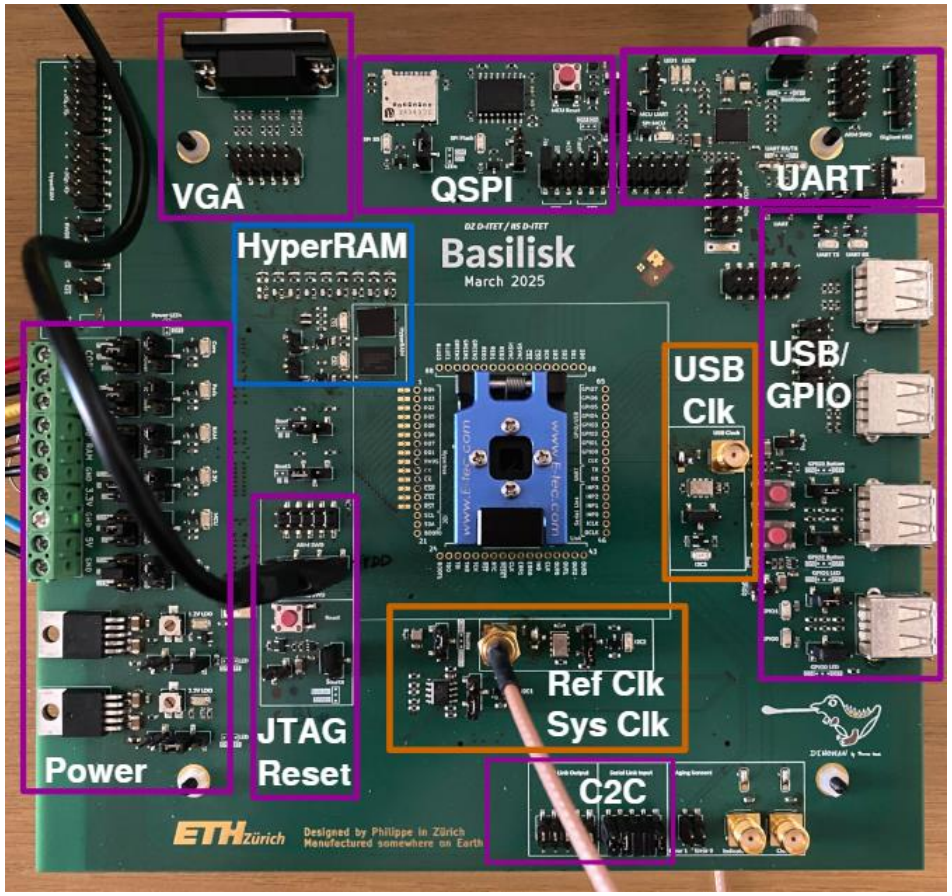

- **64-bit RISC-V core**
- **Rich peripherals:**
  - HyperRAM controller **@154MB/s**
  - C2C AXI-Link **@77MB/s**
- **Automatic boot via scratchpad**

Watch the video

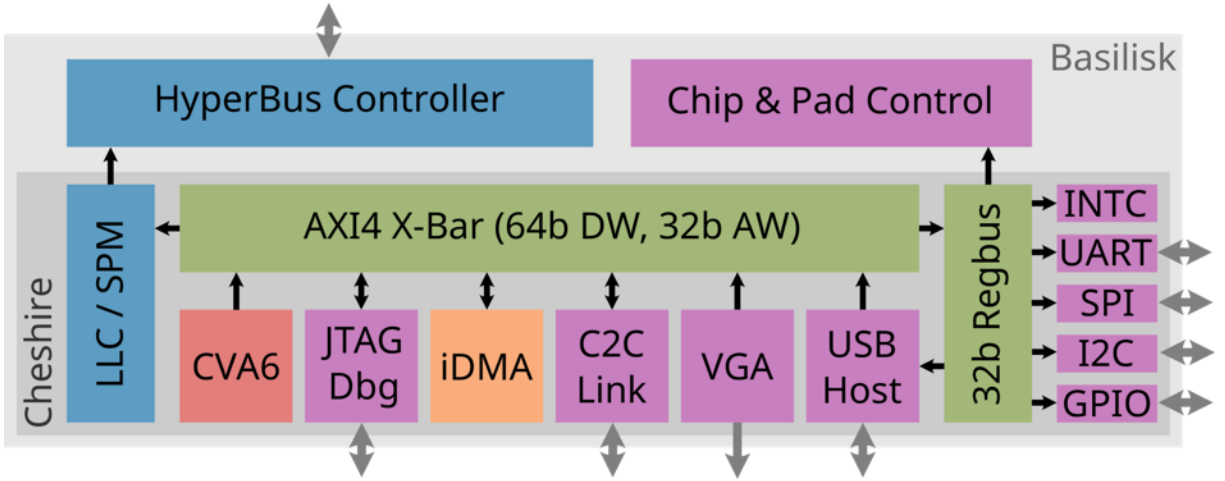

arxiv.org/pdf/2505.10060

Our research focus: cluster-based many-core accelerators

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

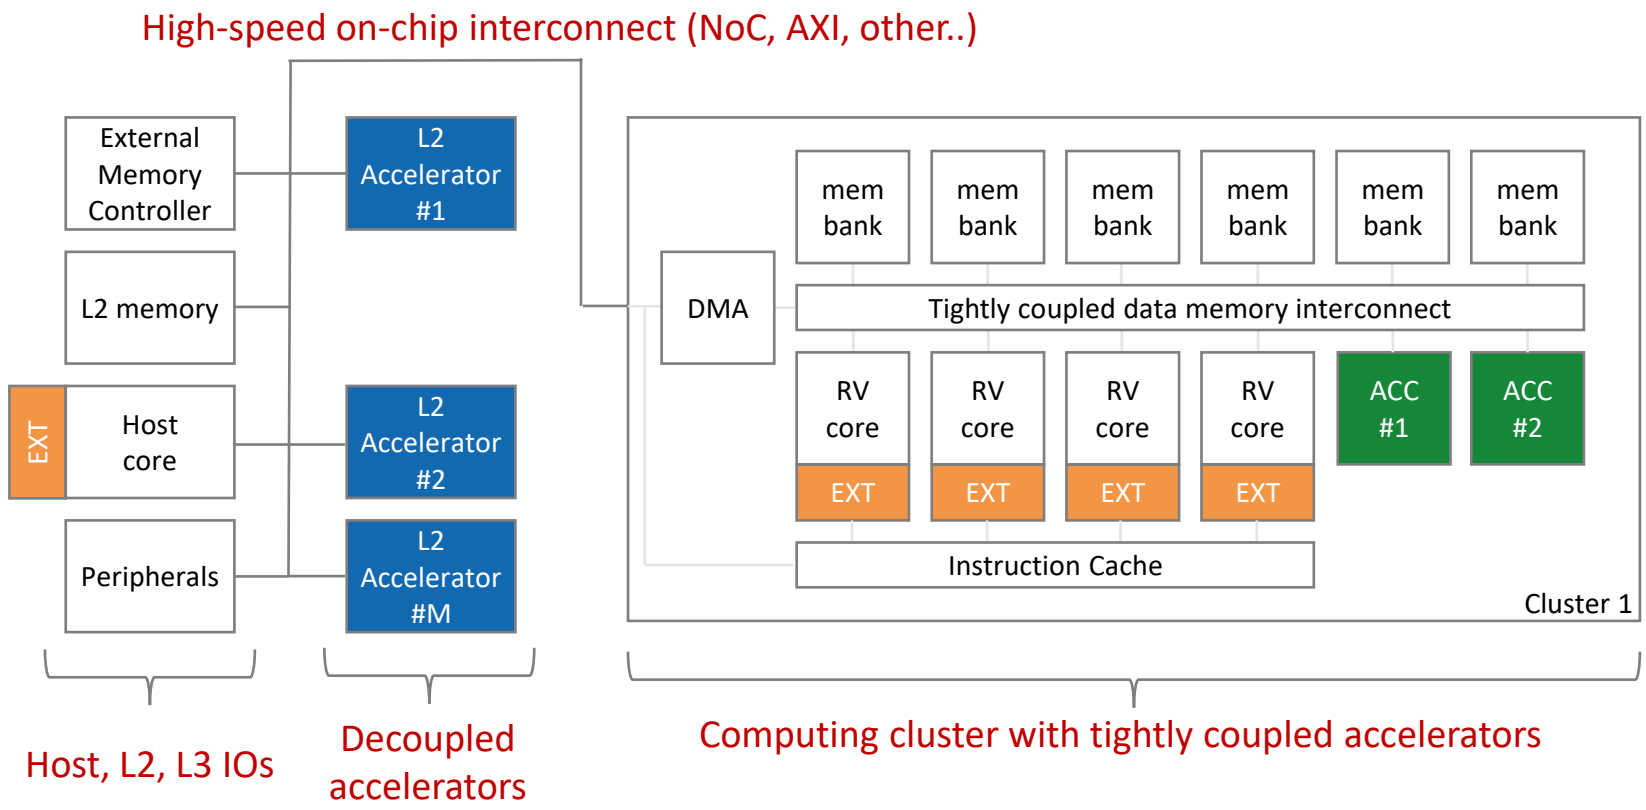

RISC-V is a key enabler → max agility, enabling SW build-up, without vendor lock-in

2013: Perf. + Efficiency + Flexibility ← **Parallelism**

- **As VDD decreases, operating speed decreases as well.**
- **However efficiency increases→ more work done per Joule**
  - Until leakage effects start to dominate
- **More units in parallel**
  - Get performance up (if you can keep them busy)
  - Energy efficiency stays high!

[Rossi et al. IEEE Micro 2017]

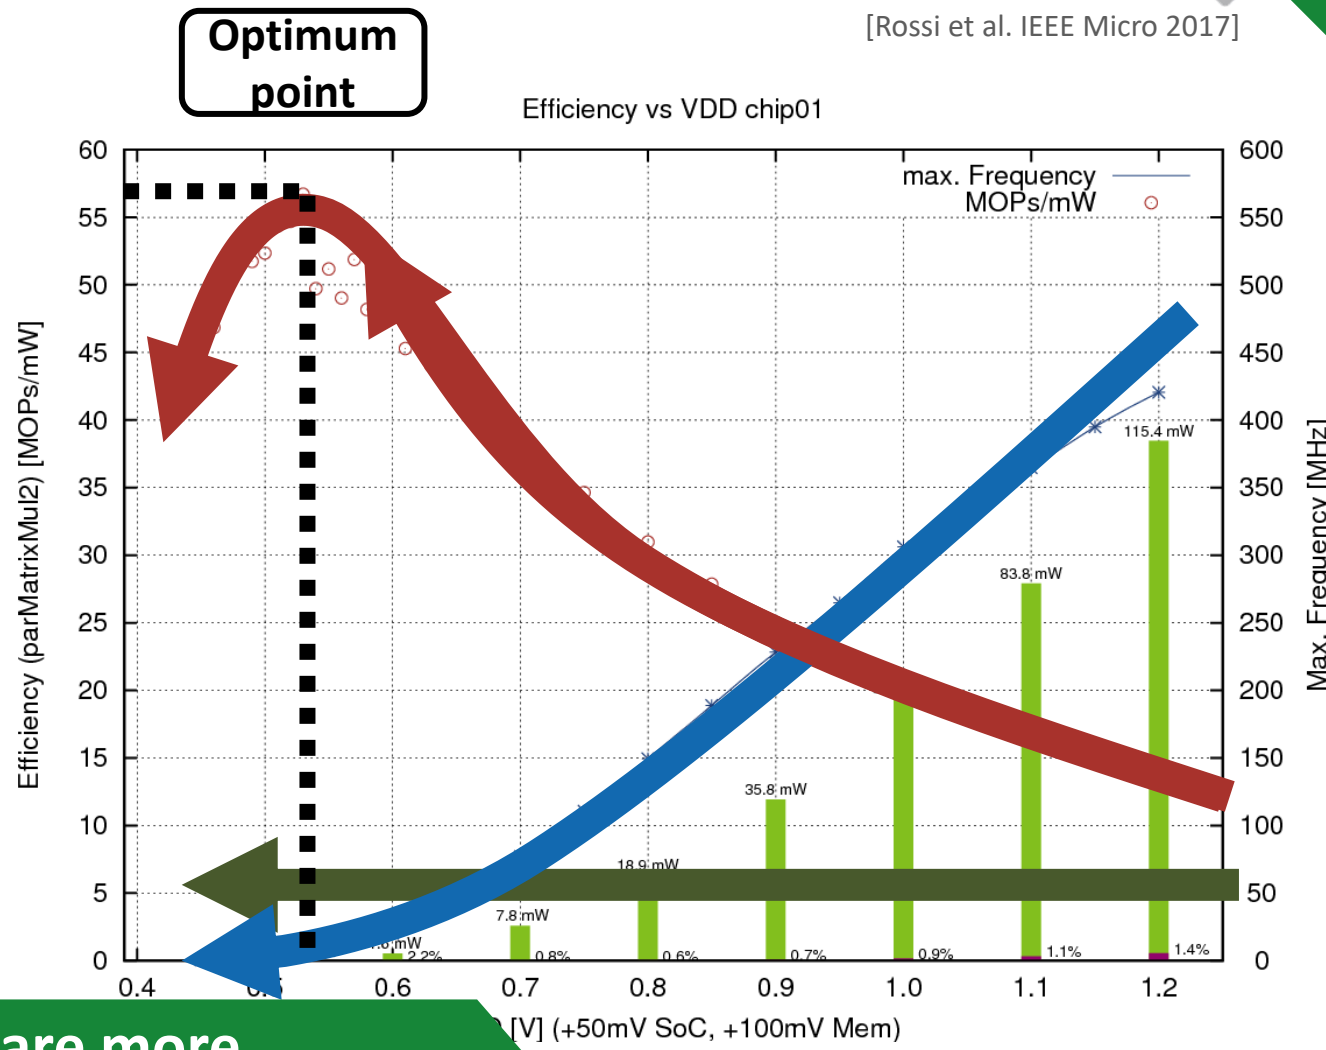

N cores running at moderate f, low Vdd are more energy efficient than a single core at N×f, high Vdd

PULP Paradigm: Multiple Cores (2-16)

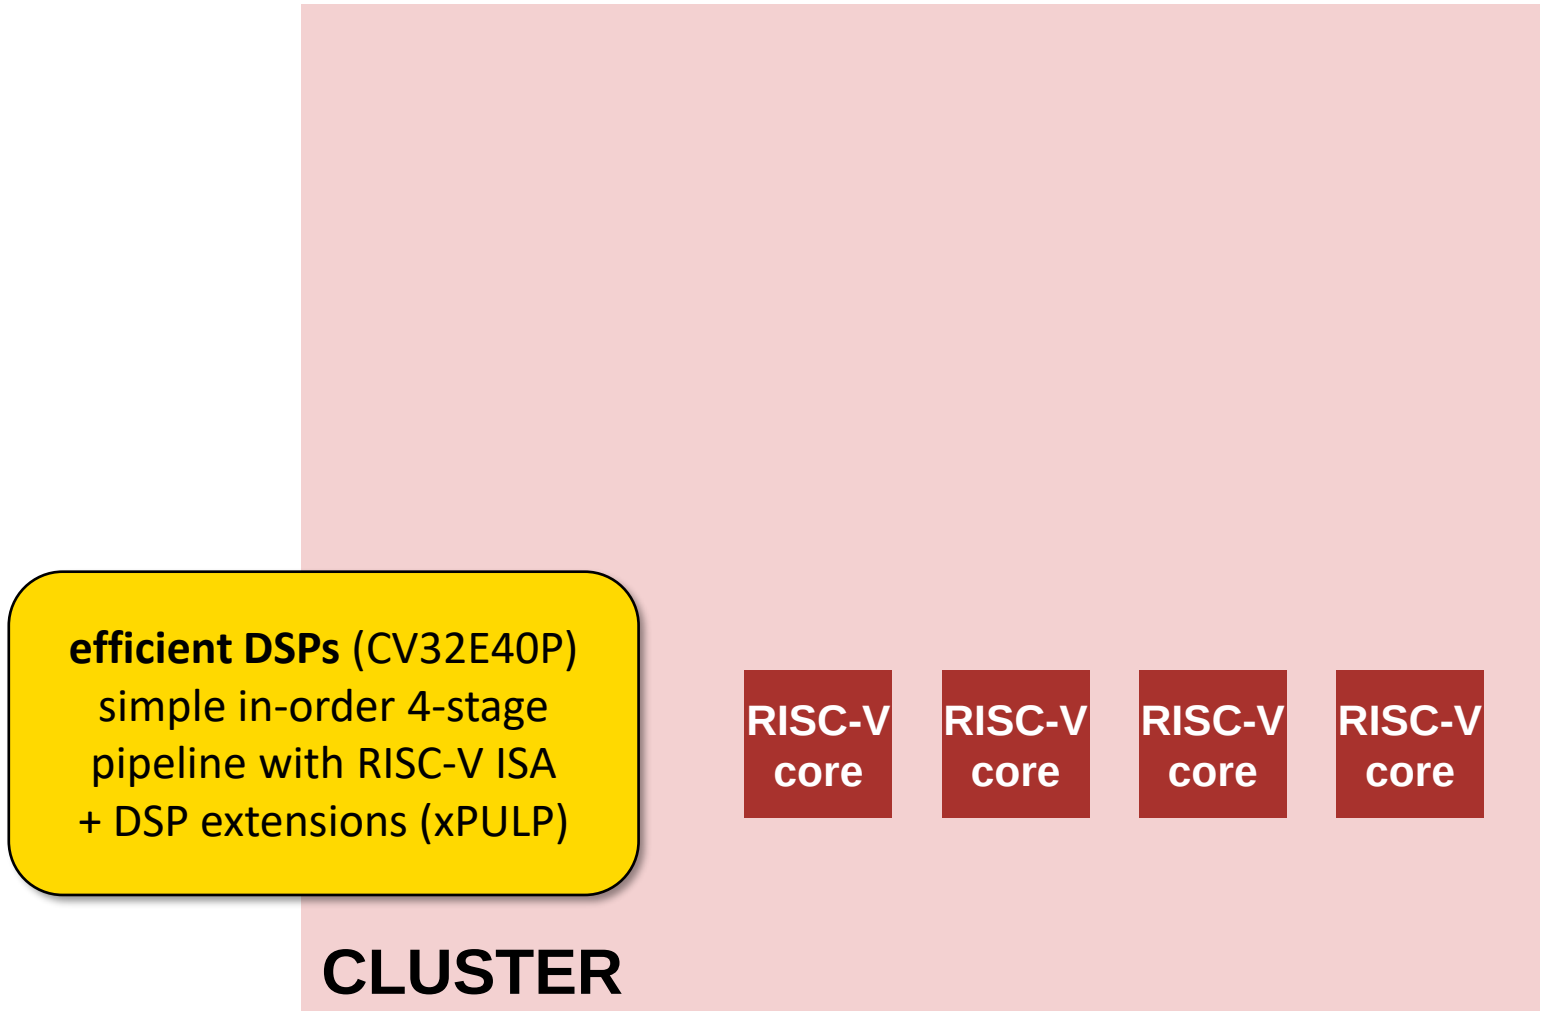

PULP Paradigm: Low-Latency Shared TCDM

bank interleaving to maximize available bandwidth in typical parallel computing scenarios

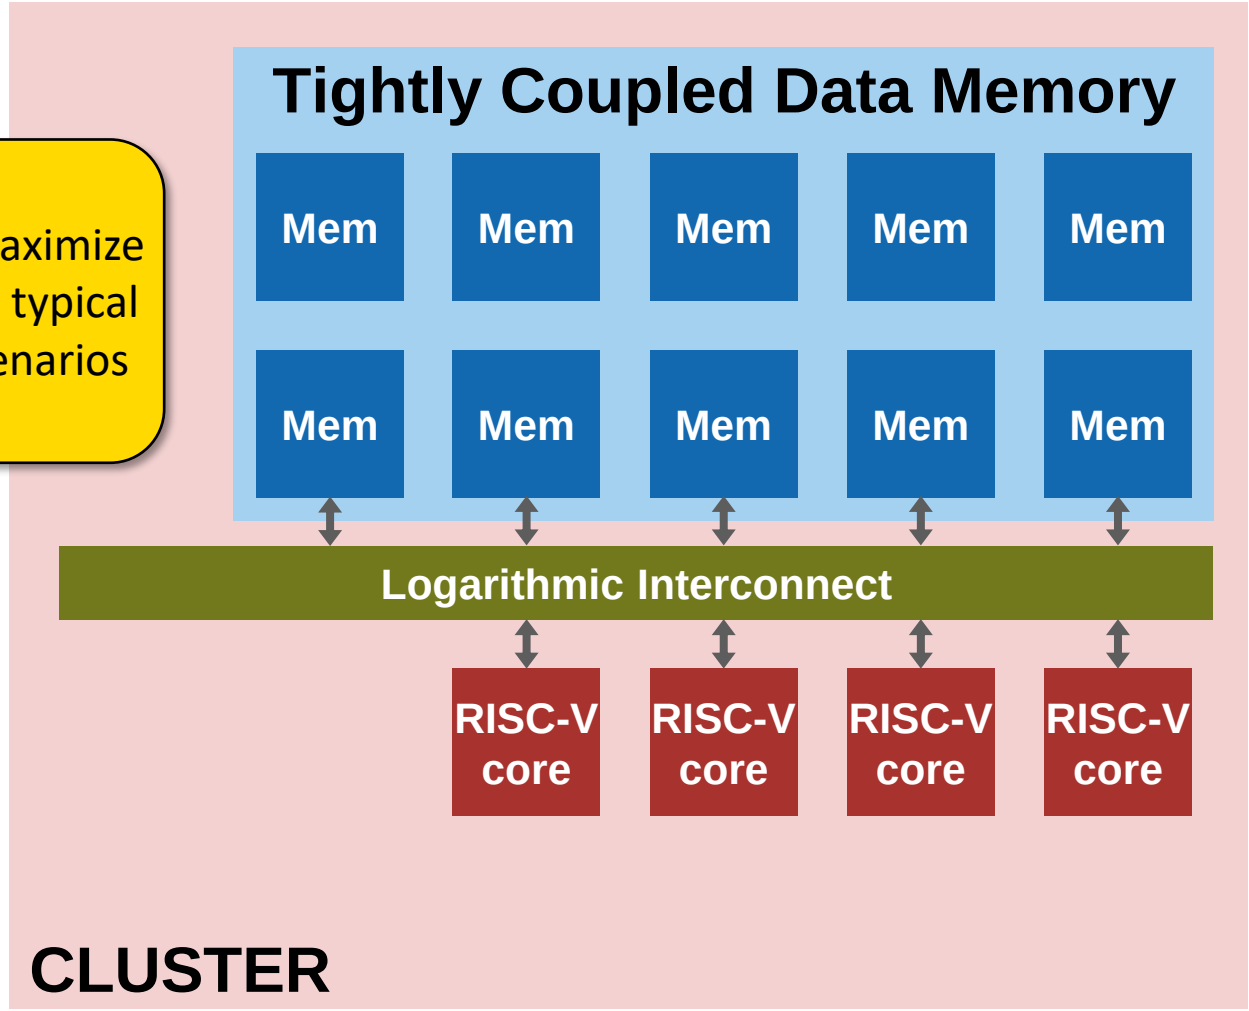

PULP Paradigm: DMA and I$ to talk with ext. memory

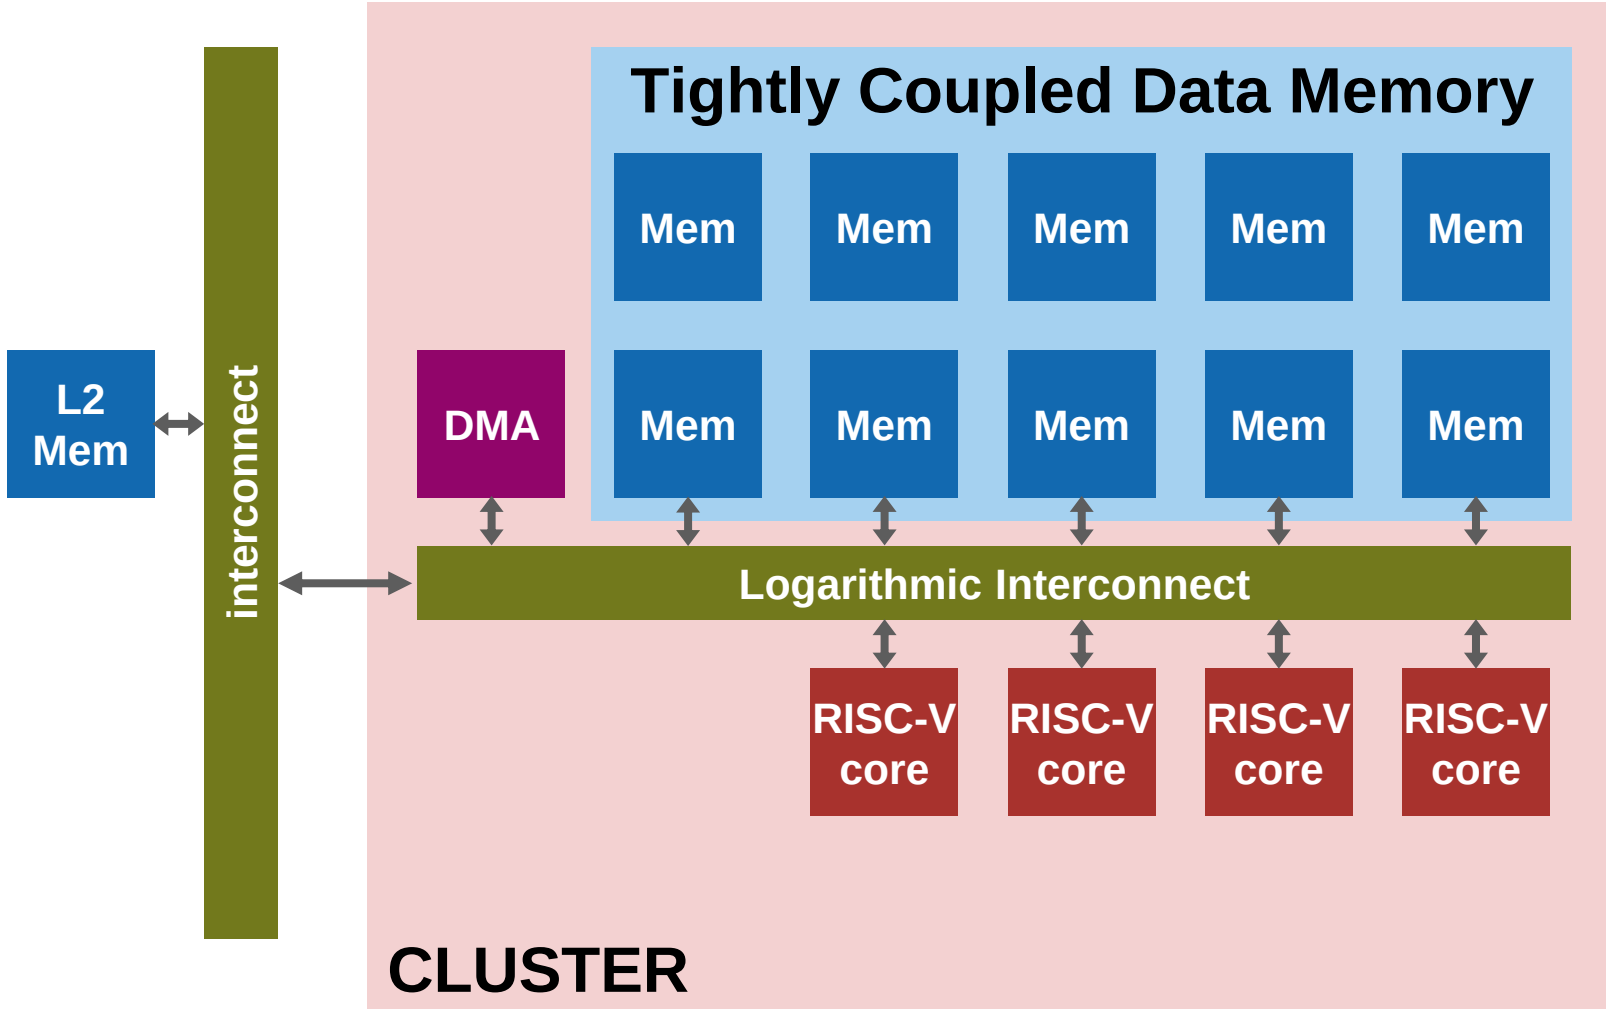

PULP Paradigm: The cluster accelerates a host system

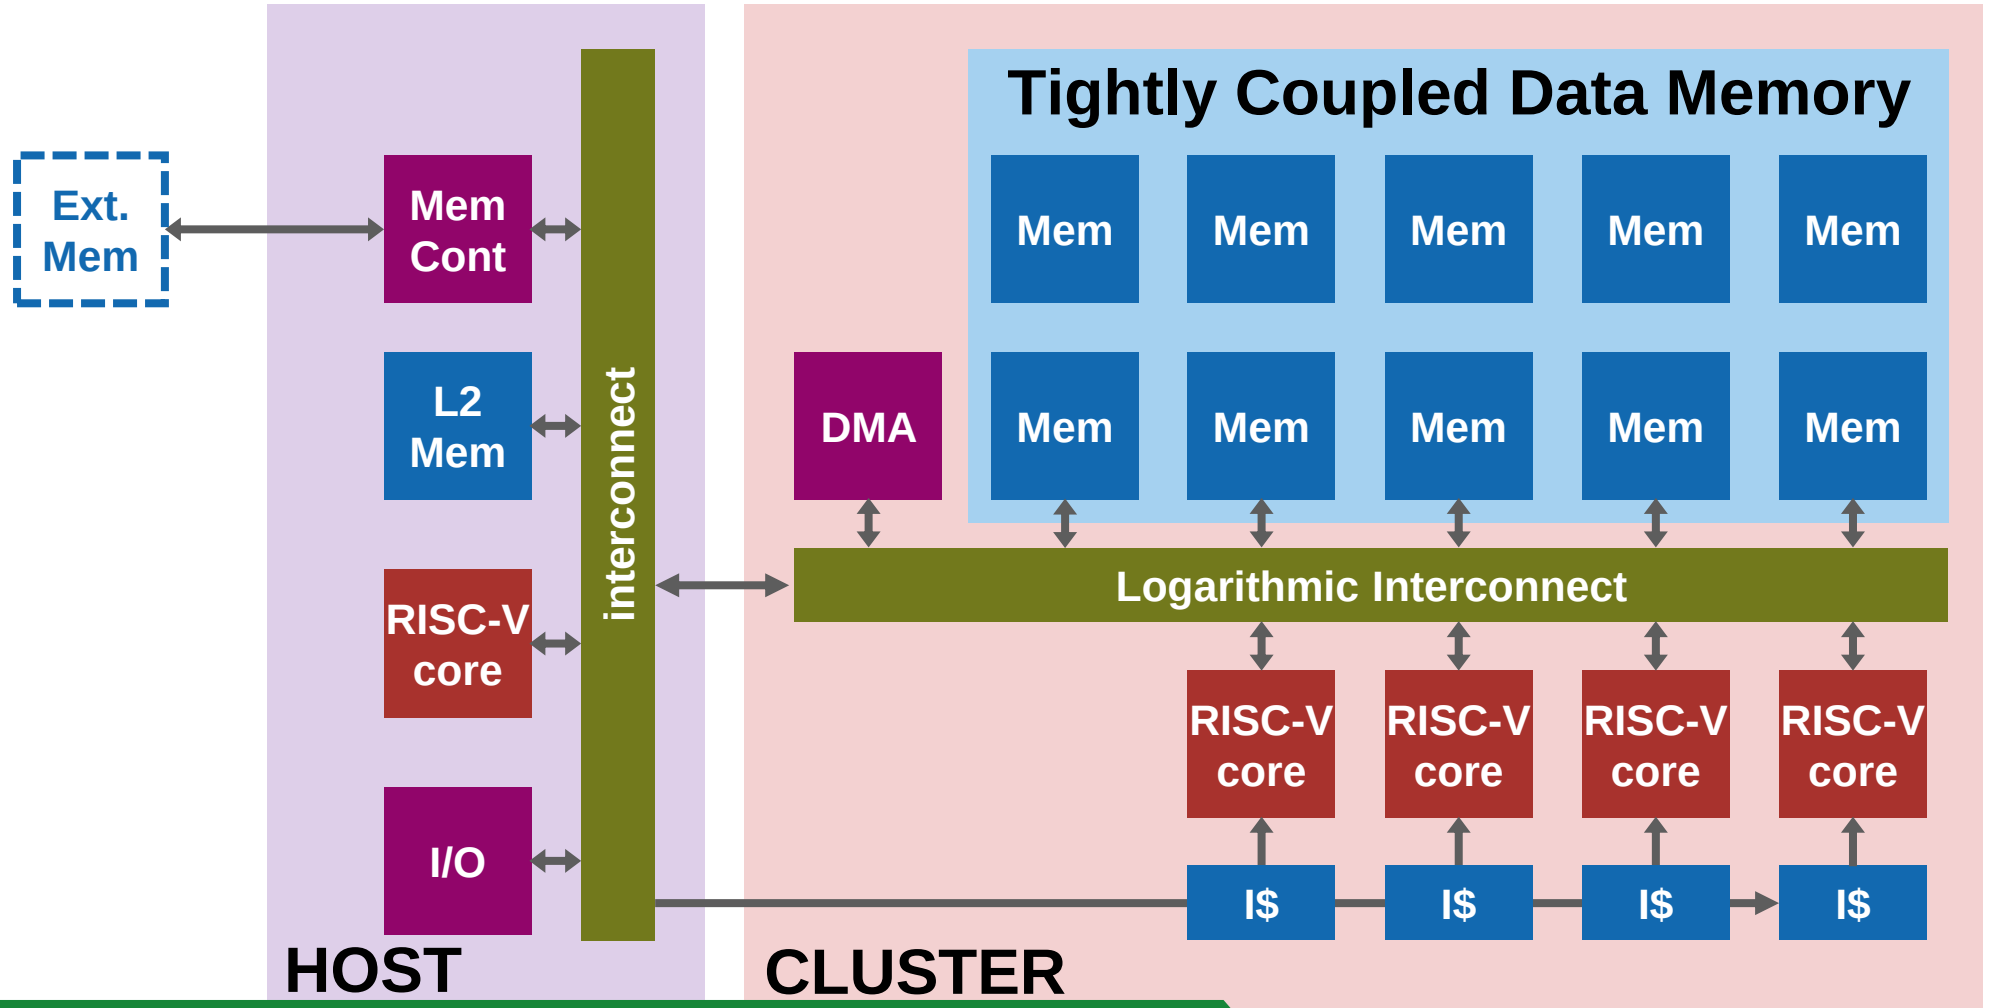

github.com/pulp-platform/pulp

PULP is a template for heterogeneous parallel SoCs

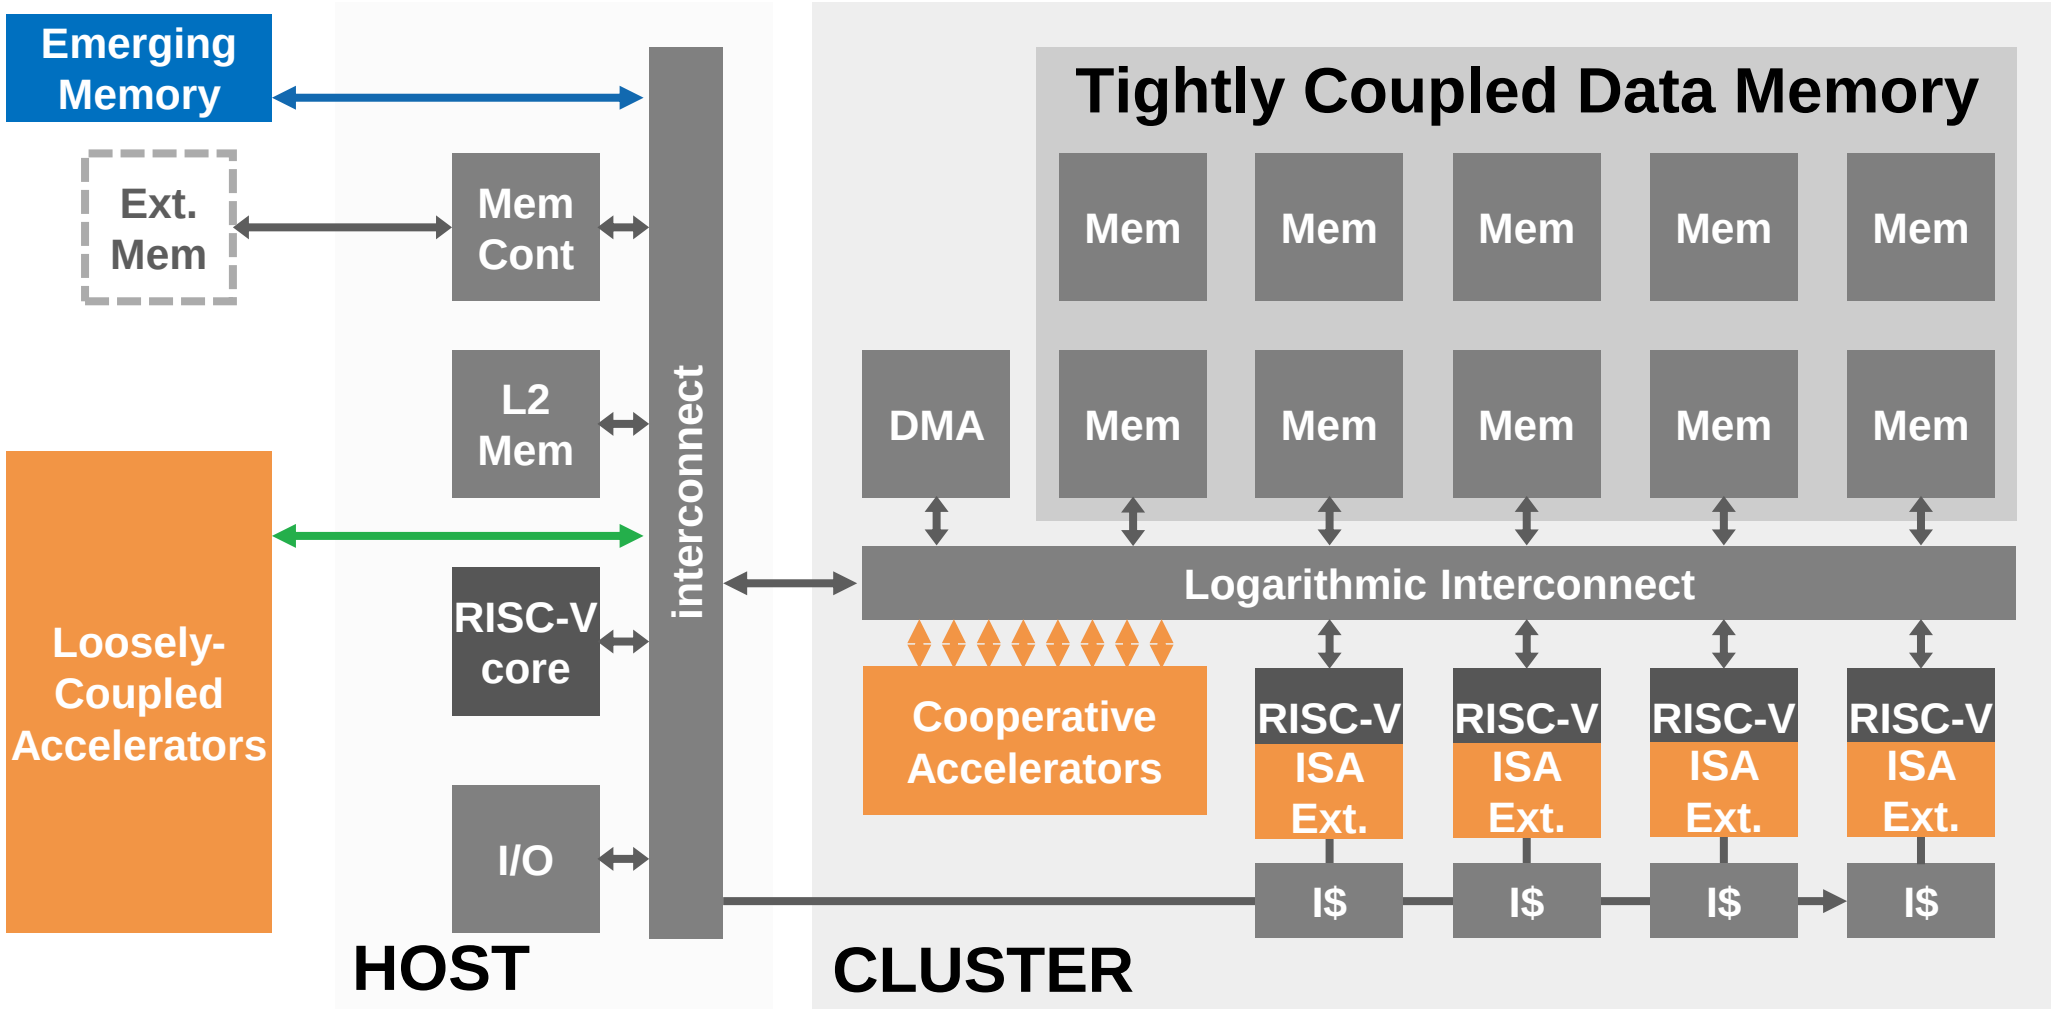

Kraken: Multiple Heterogeneous Accelerators

The ***Kraken***: an "Extreme Edge" Brain

- **RISC-V Cluster**
  (8 Cores + 1)
- **CUTIE**
  Dense ternary neural network accelerator
- **SNE**
  Energy-proportional spiking neural network accelerator

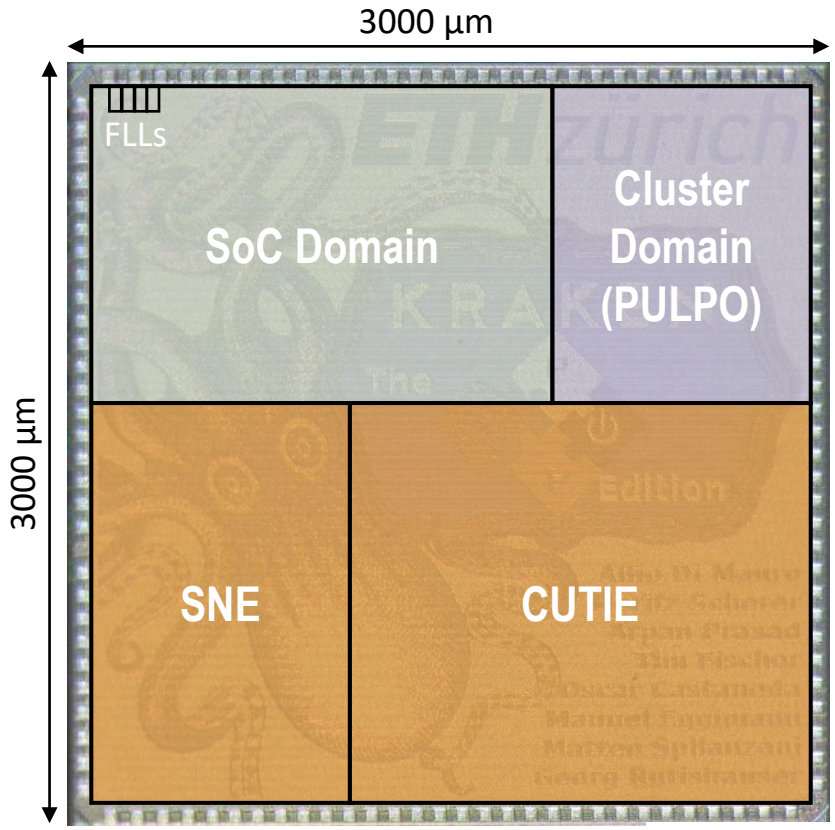

[Di Mauro HotChips22]

Technology	22 nm FDSOI
Chip Area	9 mm^2
SRAM SoC	1 MB
SRAM Cluster	128 KB
VDD range	0.55 V - **0.8 V**
Cluster Freq	**~370MHz**
SNE Freq	**~250MHz**
CUTIE Freq	**~140MHz**

Specialization in perspective

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

Energy-Efficient RV 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**

Fully Open-Source Deployment Flow!

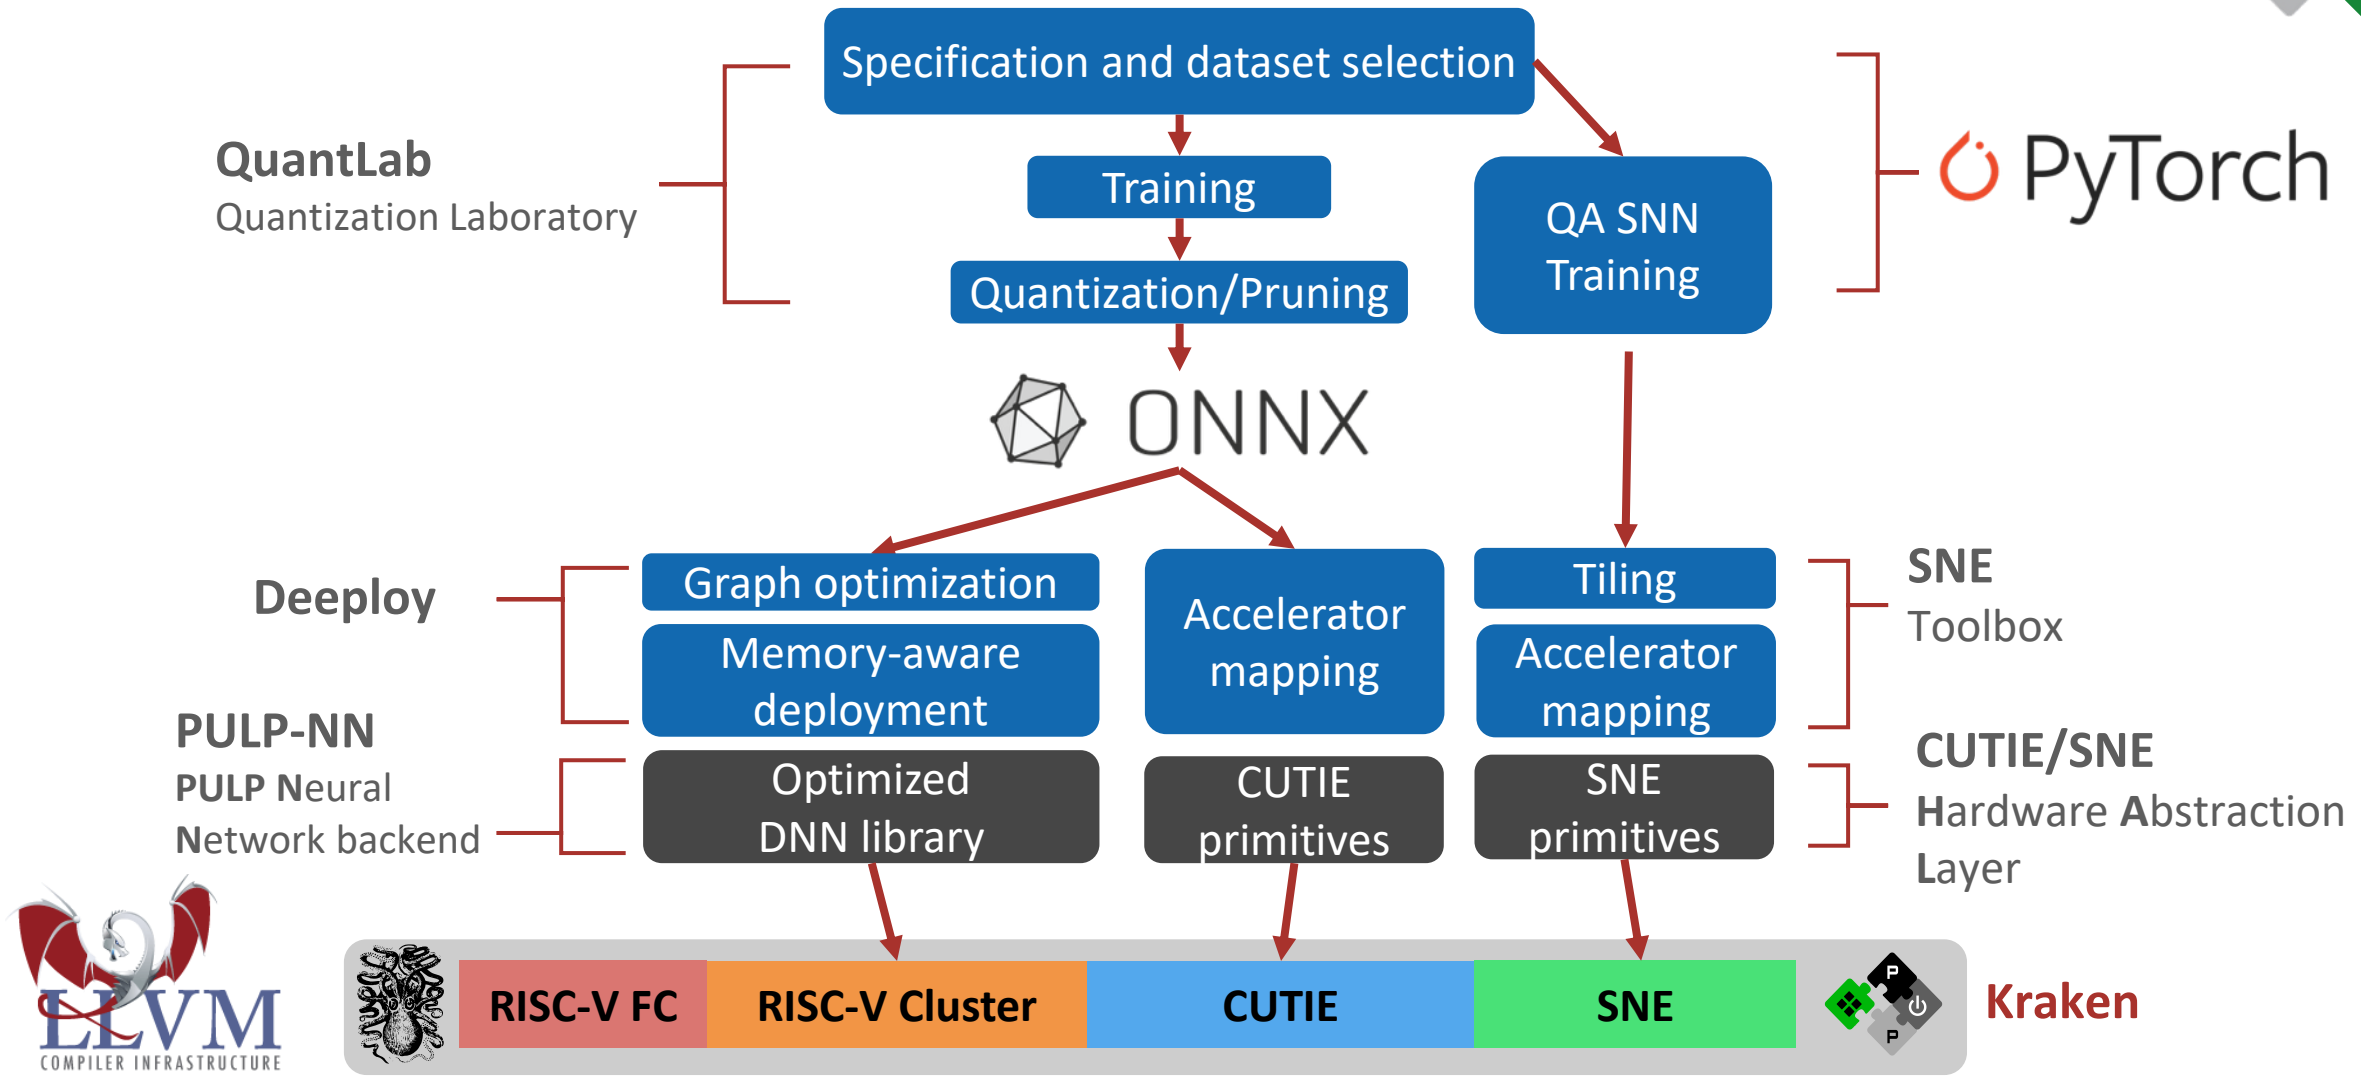

Where do we go from here: *Scale-Up* vs. *Scale-Out*

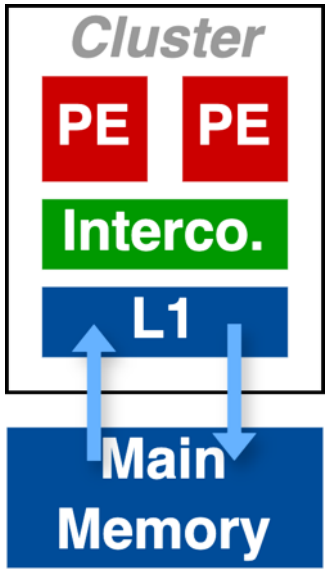

Scale-up

Increase the number of PEs in a cluster

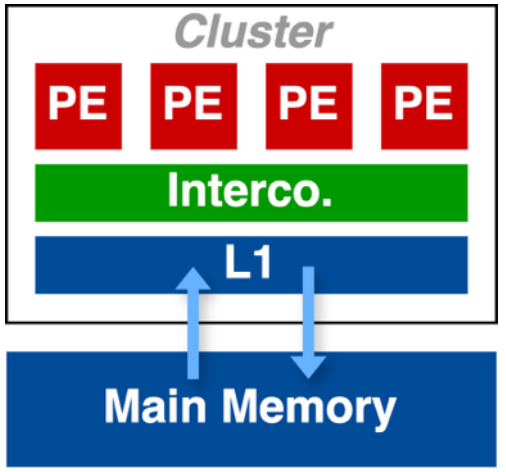

Scale-out

Increase the number of parallel Clusters

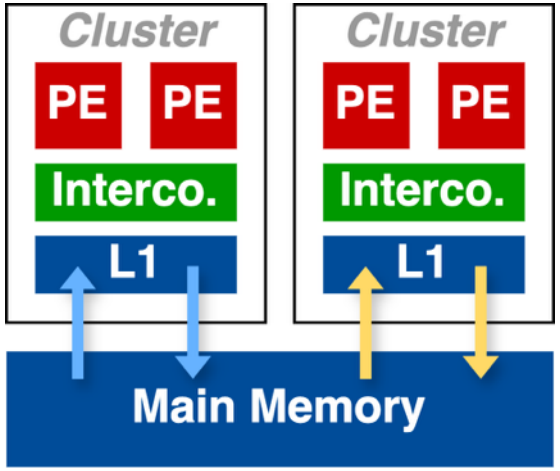

Cluster: A group of **PEs** tightly coupled with a shared **L1 memory** through low-latency **interconnect**

Achieving Scale through Hierarchical Design

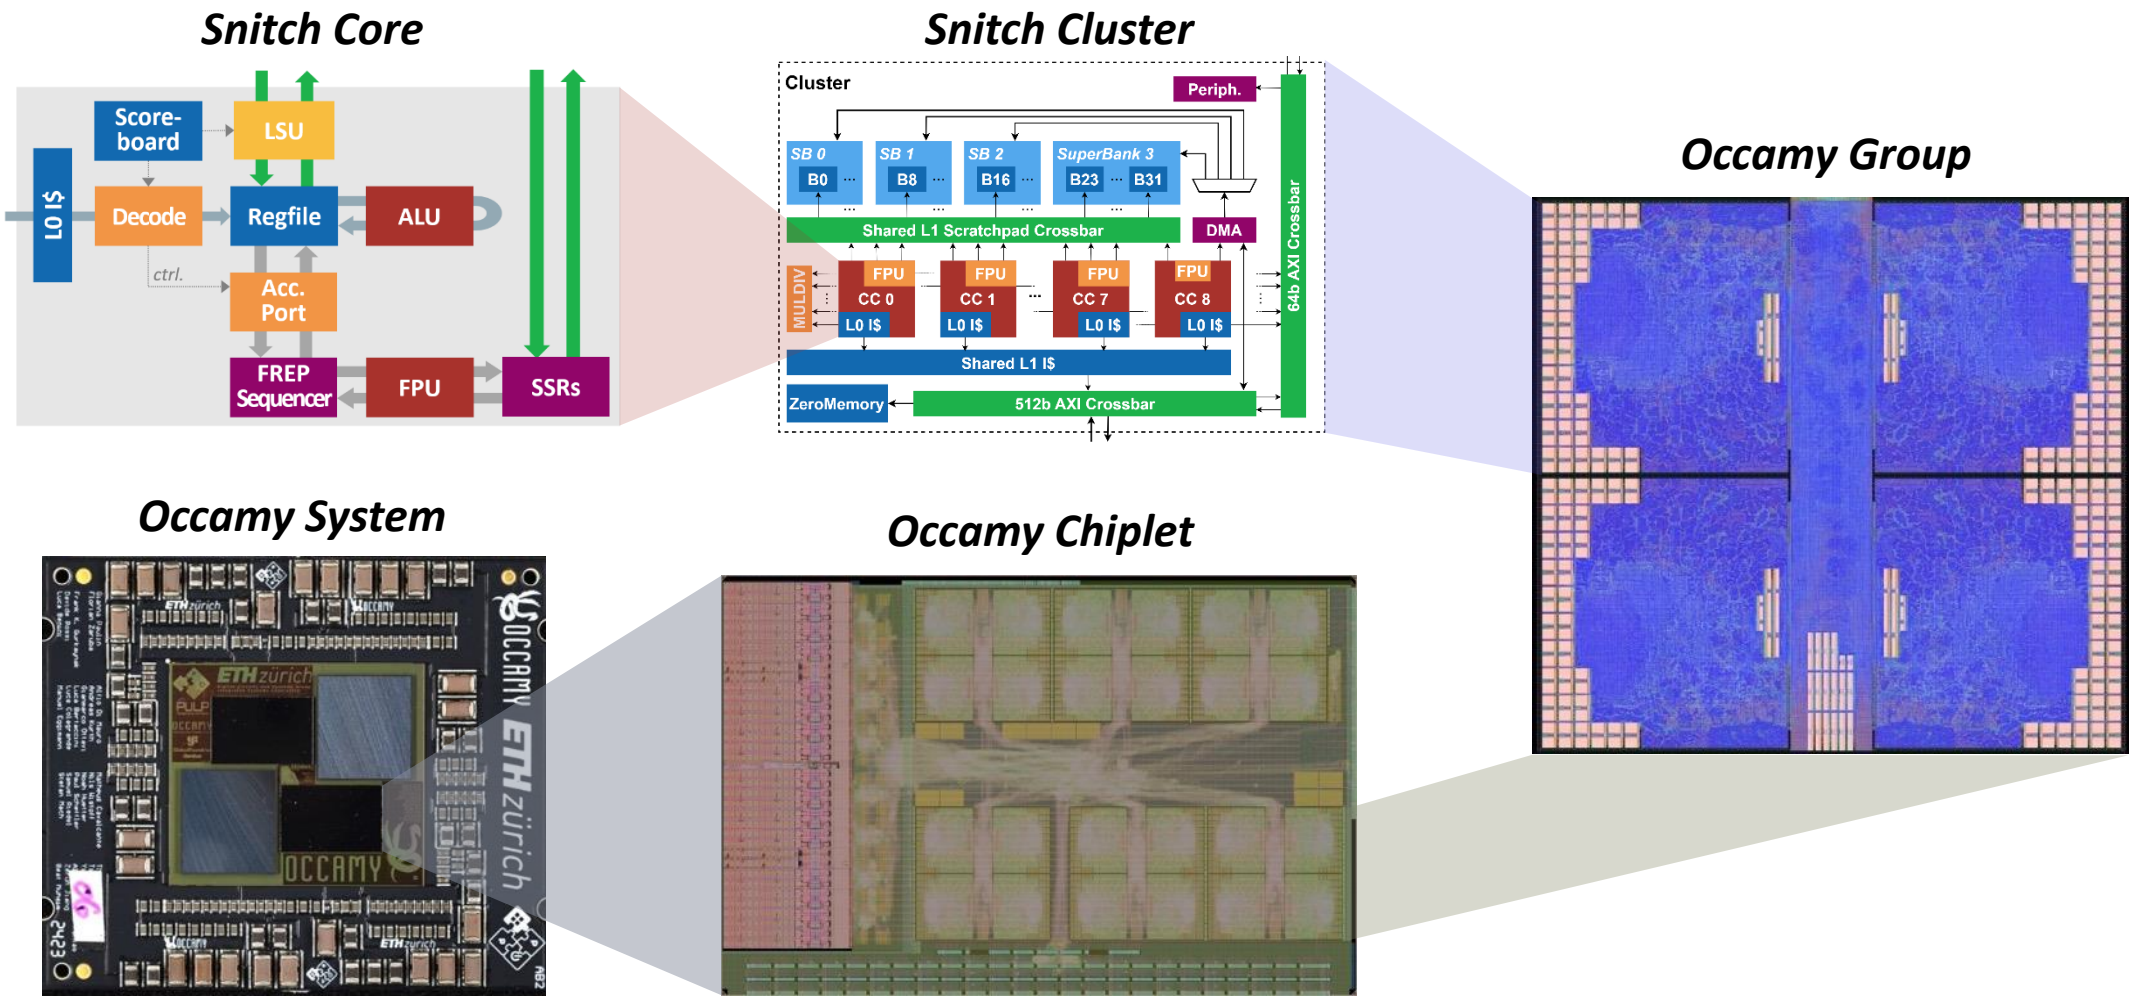

Occamy Chiplet: Six Groups with HBM and D2D Link

- **6 fully connected groups**
  - 24 clusters, 216 cores total
  - 512b data, 64b for messages interconnect
- **Autonomous 64b host domain**
  - CVA6 RV64GC Linux-capable core
  - Rich peripherals (SPI, I2C, UART…)
- **16 GiB, 410 GB/s HBM2E**
  - Optional page-level interleaving
- **12.8 GiB/s die-to-die link**
  - Fully digital and fault-tolerant

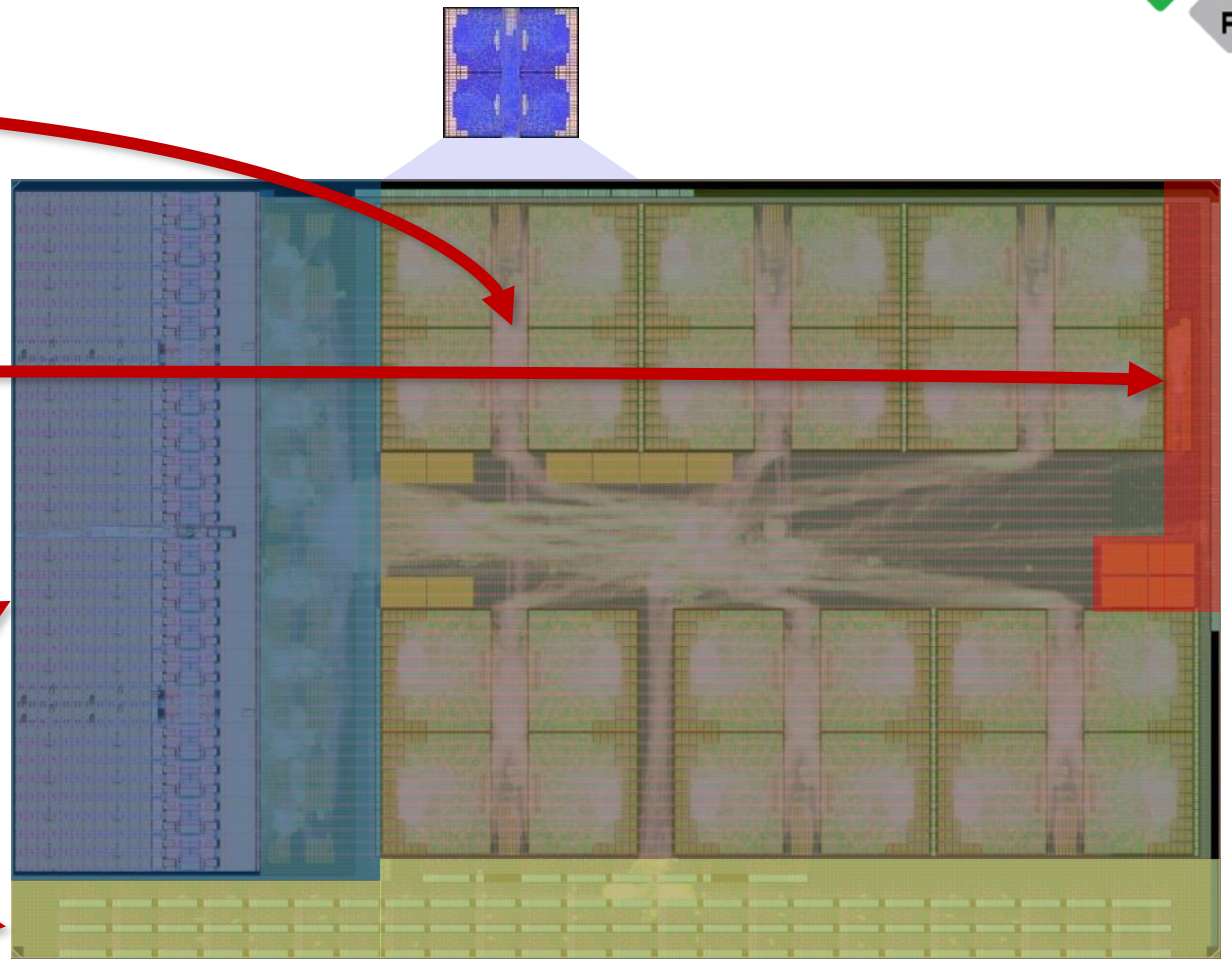

github.com/pulp-platform/snitch_cluster

Addressing interconnect scalability

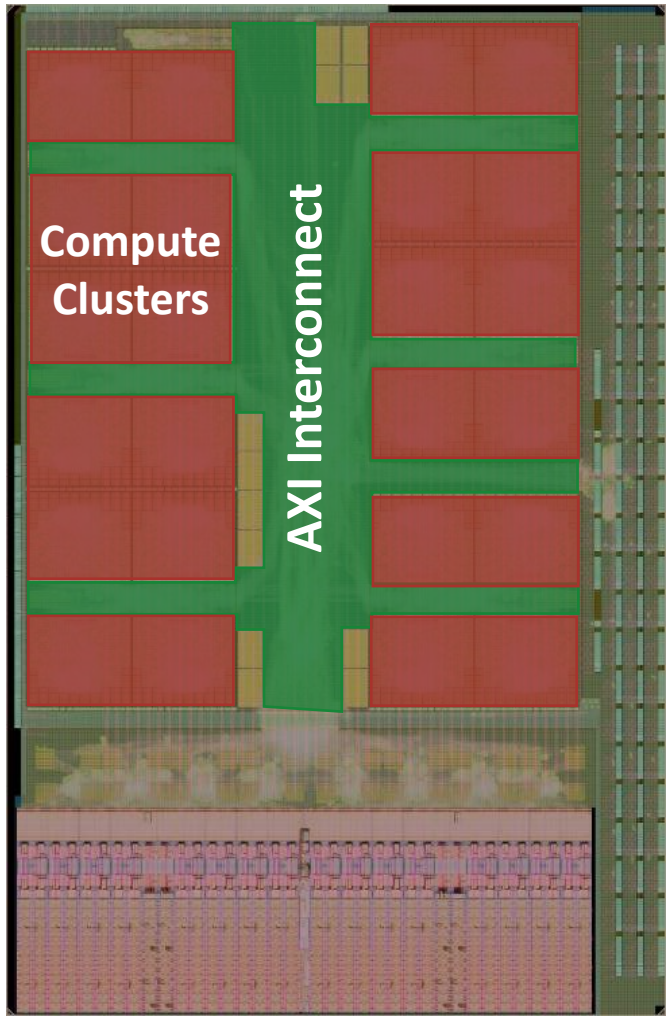

- **AXI interconnect was very challenging for PD**
  - AXI has severe scalability issues
  - Top-level Xbar had to be split up
  - Still, interconnect takes up almost **40%***
- **Working on NoC solution, *FlooNoC***
  - Fully AXI4 compatible
  - Solves AXI4 **scalability issues**
  - Designed with awareness of physical design
  - **Wide** & **physical** channels

`github.com/pulp-platform/FlooNoC`

**HBM & C2C excluded*

Picobello Next generation many-core architecture in 7nm

- **64bit host**
  - Linux capable CVA6
- **16 clusters**
  - 9x Snitch (RV32) cores with Sparse SSRs
  - MatrixMul
- **8 MB L2 memory**
- **FlooNoc interconnect**
- **Taped out in Aug 2025**

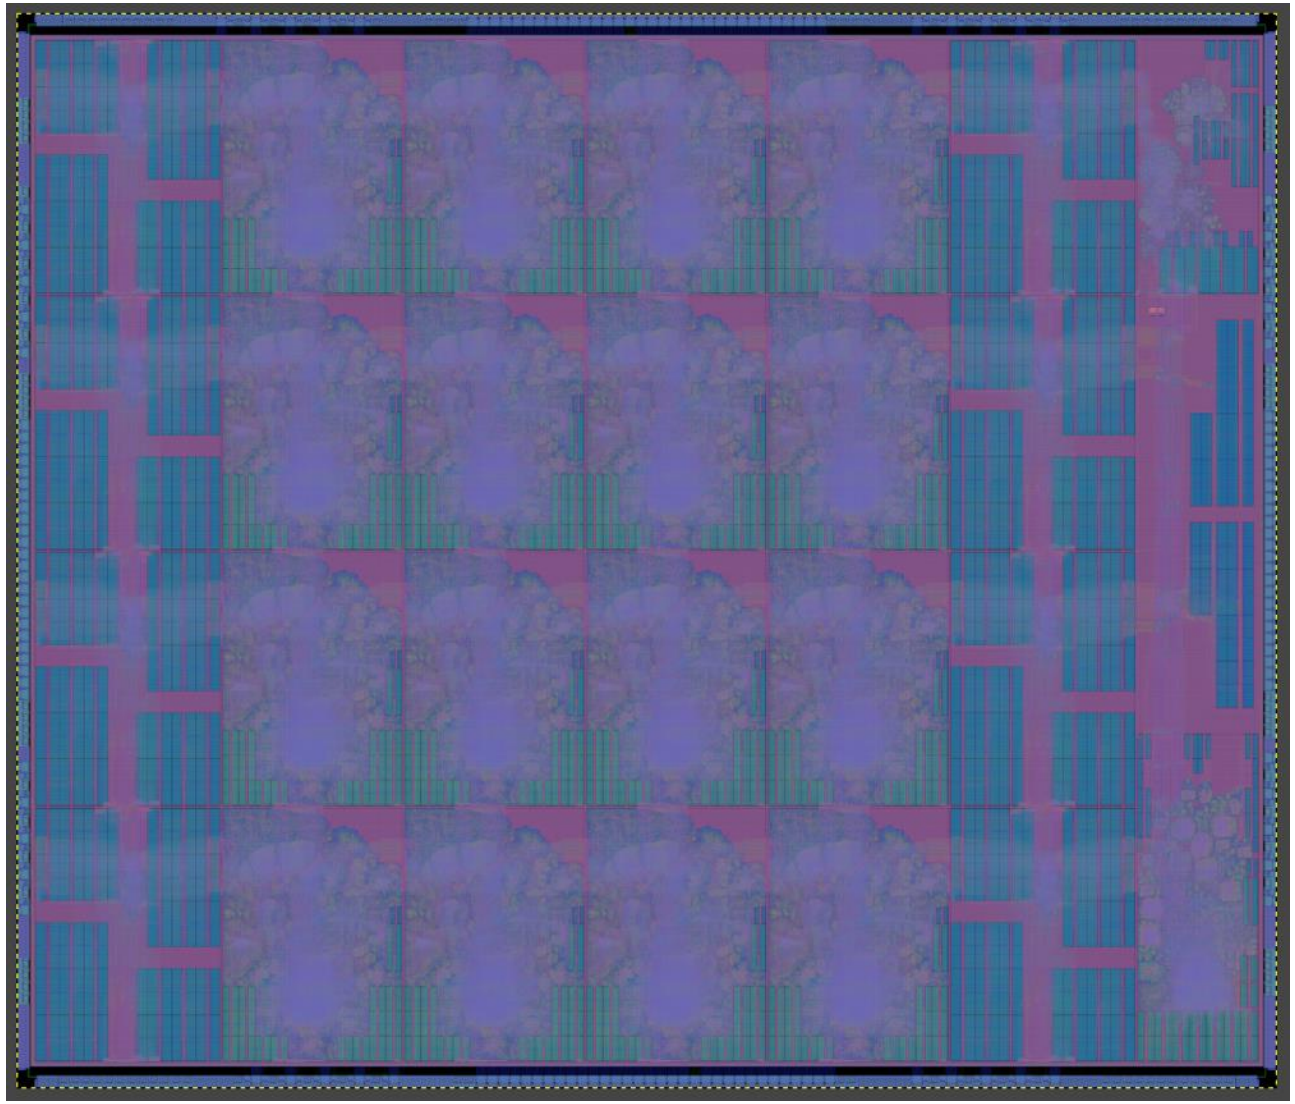

THE EUPILOT

ETH zürich ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA

Austrochip Workshop on Microelectronics

What is next for open-source hardware

- Challenges
- Opportunities

End-to-end Open-Source IC Design is possible today!

Design: from PULP

`github.com/pulp-platform`

Tools: from Johannes Kepler University (JKU)

Reliable VM with large collection of open-source tools

`github.com/iic-jku/IIC-OSIC-TOOLS`

Manufacturing: IHP130nm

`github.com/IHP-GmbH/IHP-Open-PDK`

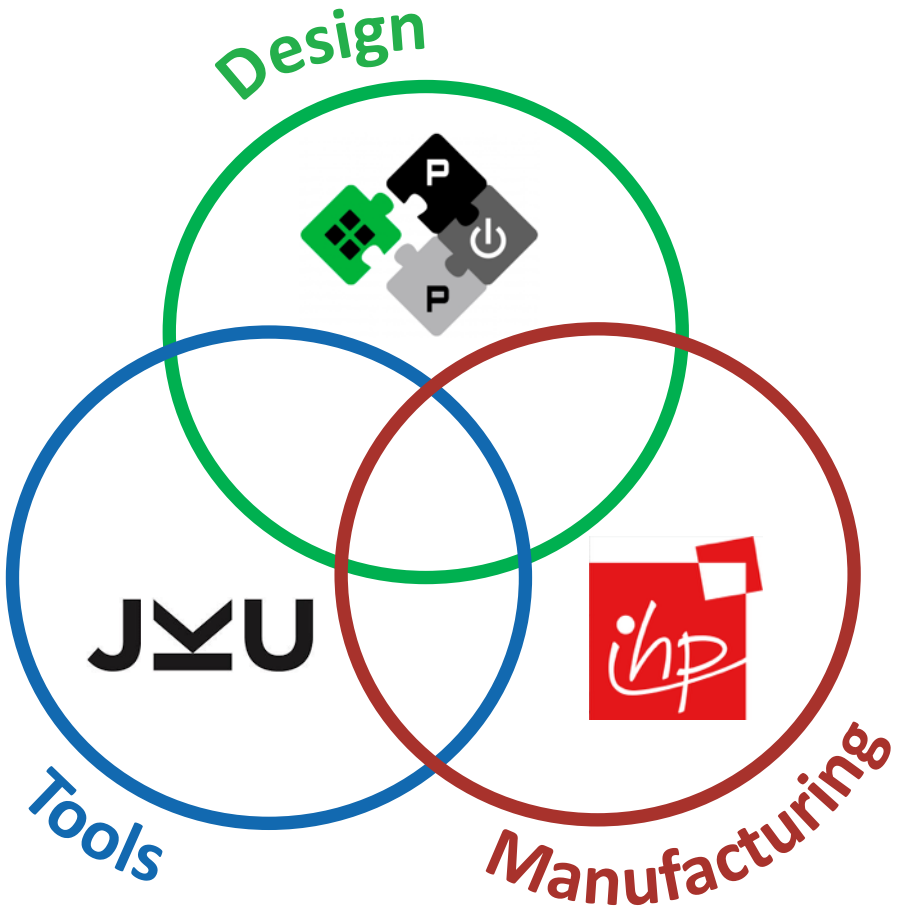

End-to-end Open-Source allows sharing of design data

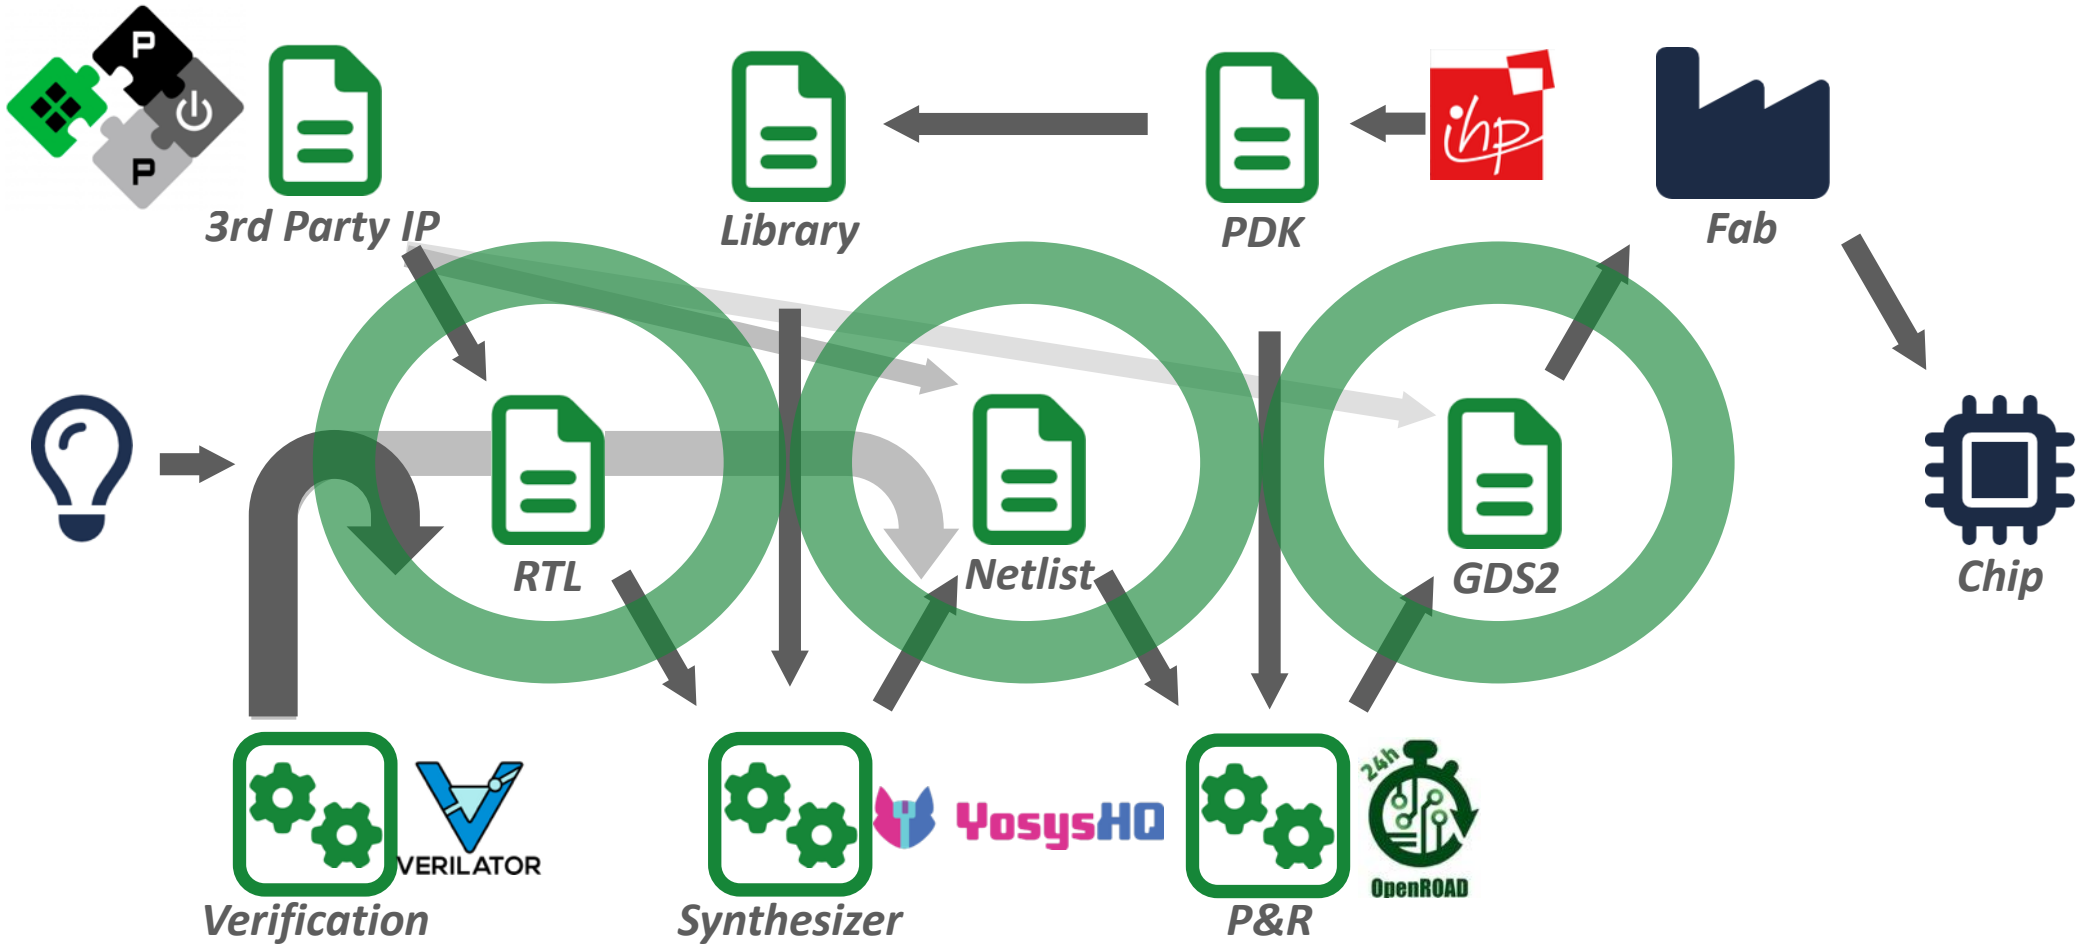

Closing the PPA gap to commercial EDA

6

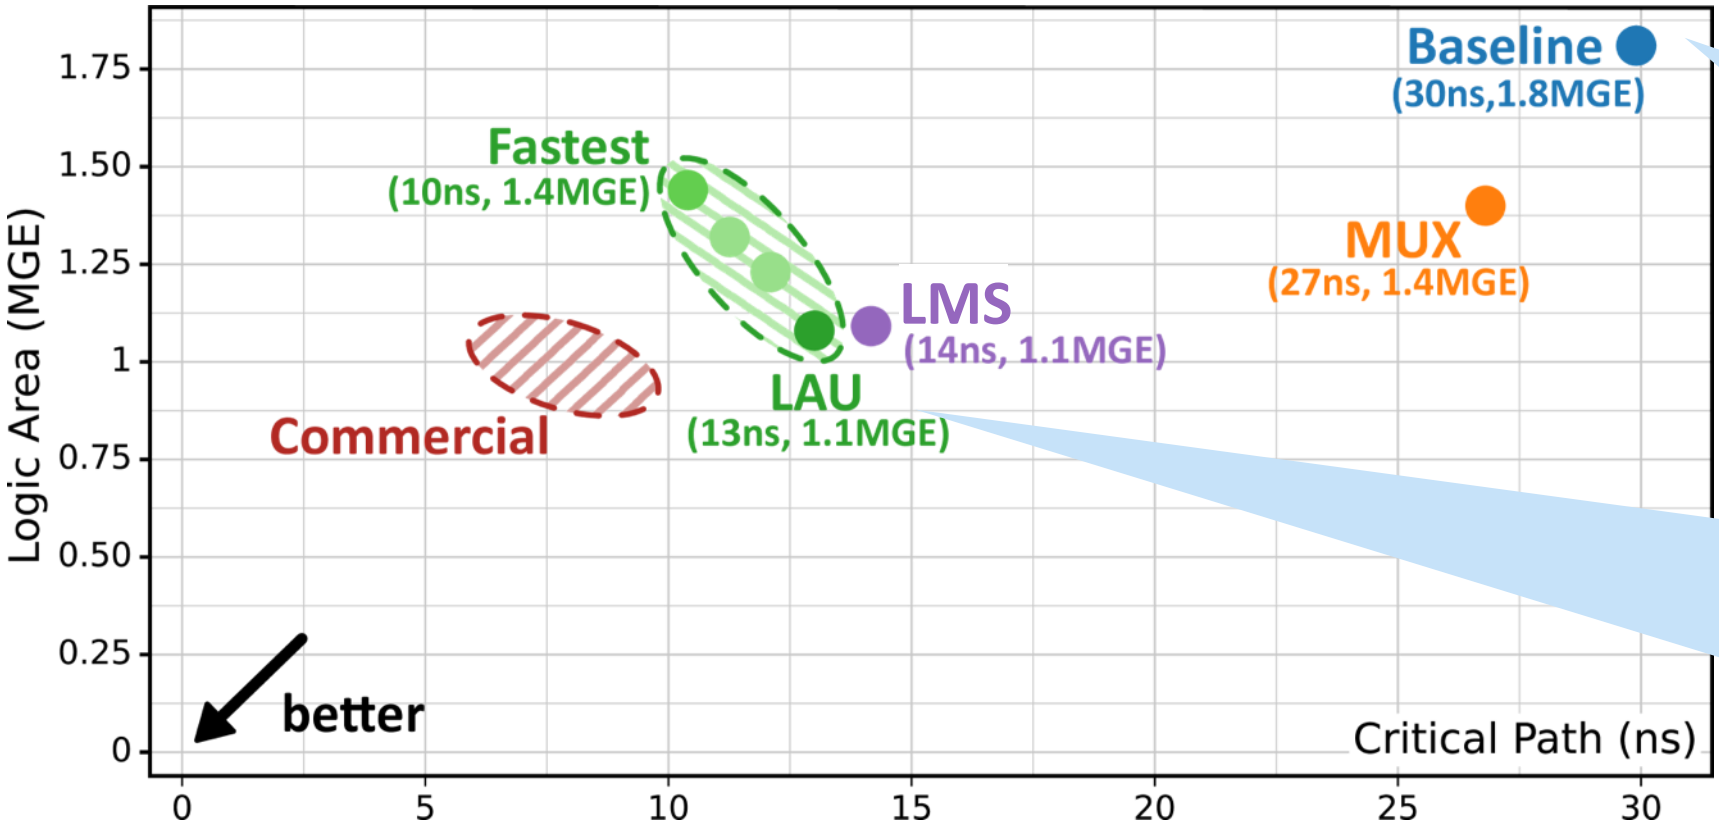

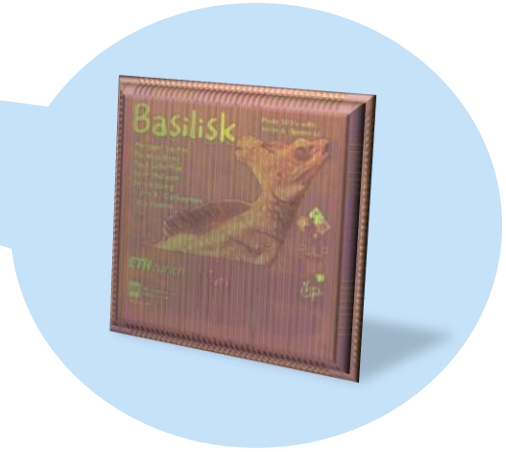

Yosys-slang full Sysverilog Frontend: @ <6sec runtime (from minutes)

Yosys synthesis: 1.1 MGE (1.6×) @ 77 MHz (2.3×), 2.5× less runtime, 2.9× less RAM

OpenROAD P&R: tuning -12% die area, +10% core utilization

Yes but why? Specialization + EDA multiplicative effect!

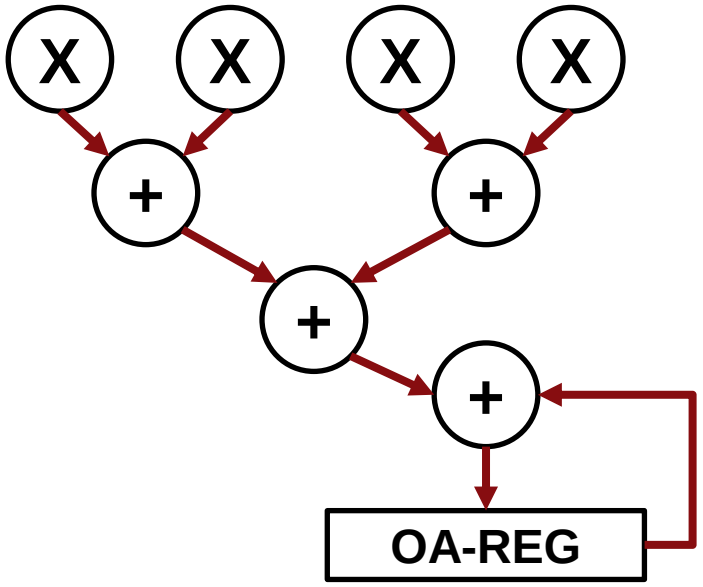

Inner Product

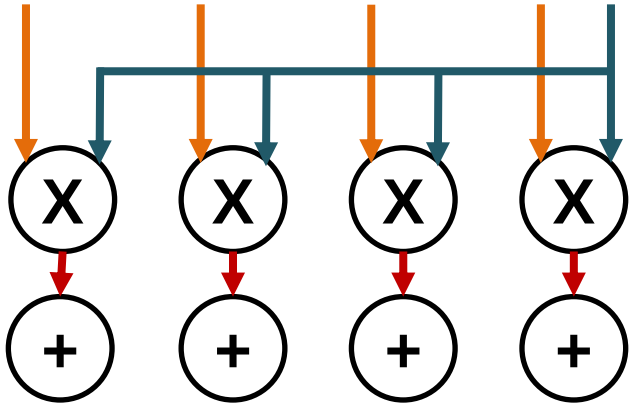

Outer Product

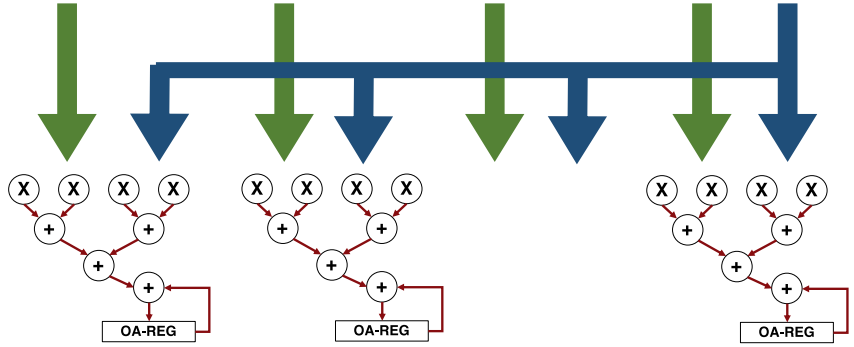

Mixed

Precision tuning – OP/Mem tuning - deep arithmetic optimization – operand network tuning...

Co-Specialize SW, HW, EDA & Technology is the frontier

Open EDAs Heterogeneous Chips in Advanced CMOS?

Extreme Performance + Energy Efficiency is required!

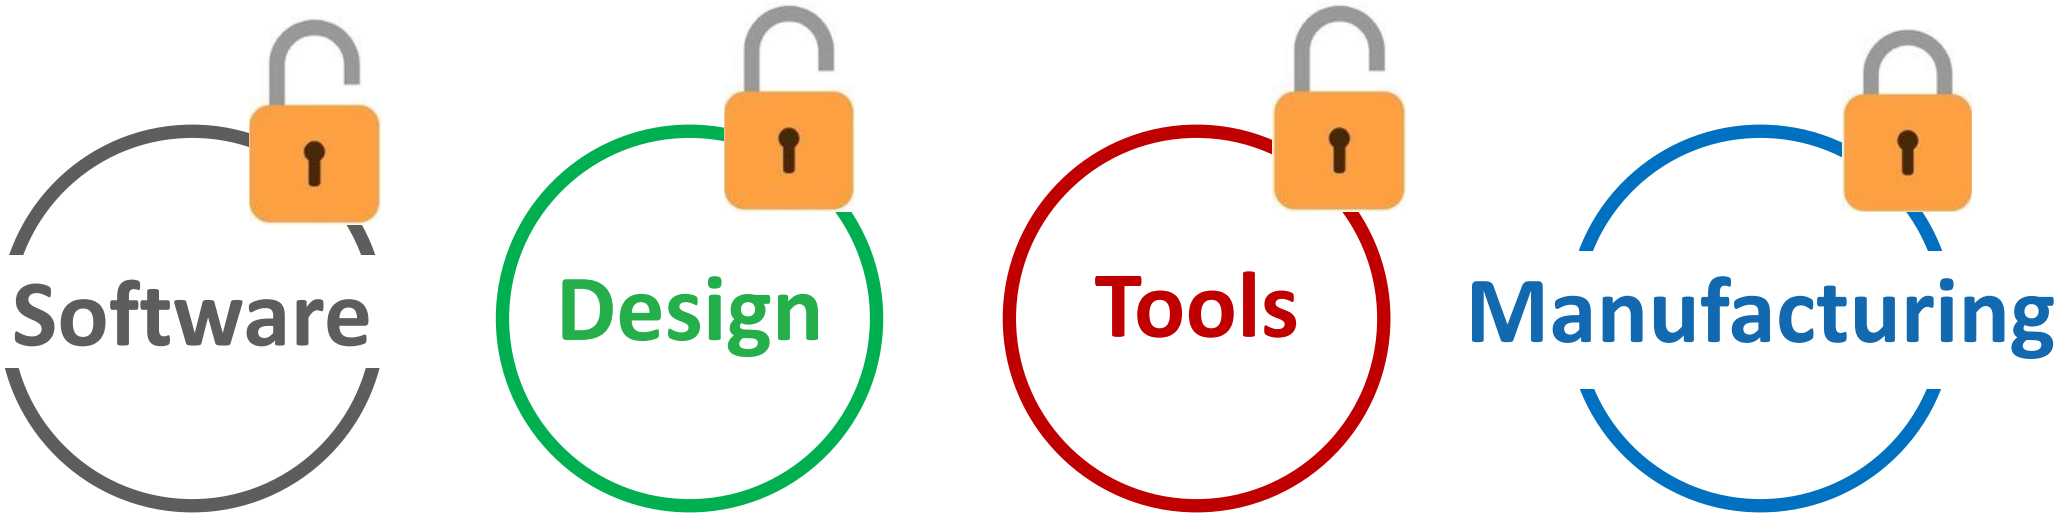

HW-SW Co-Design

HW-SW & Tools Co-Optimization

Get the best from advanced nodes!

On the Horizon: **SeyrITA** – GF22 with Open-Source Tools

- **RISC-V Linux host platform**
- **Transformer accelerator**
  - Targeting BERT, mobileBERT, DEiT-T
  - Leveraging microscaling quantization
  - MXINT and MXFP32 formats
- **10x larger!**
  - 20-40MGE (SeyrITA) vs 2-3MGE (Basilisk)
- **500MHz target frequency**
- **1-2TFLOP/s**

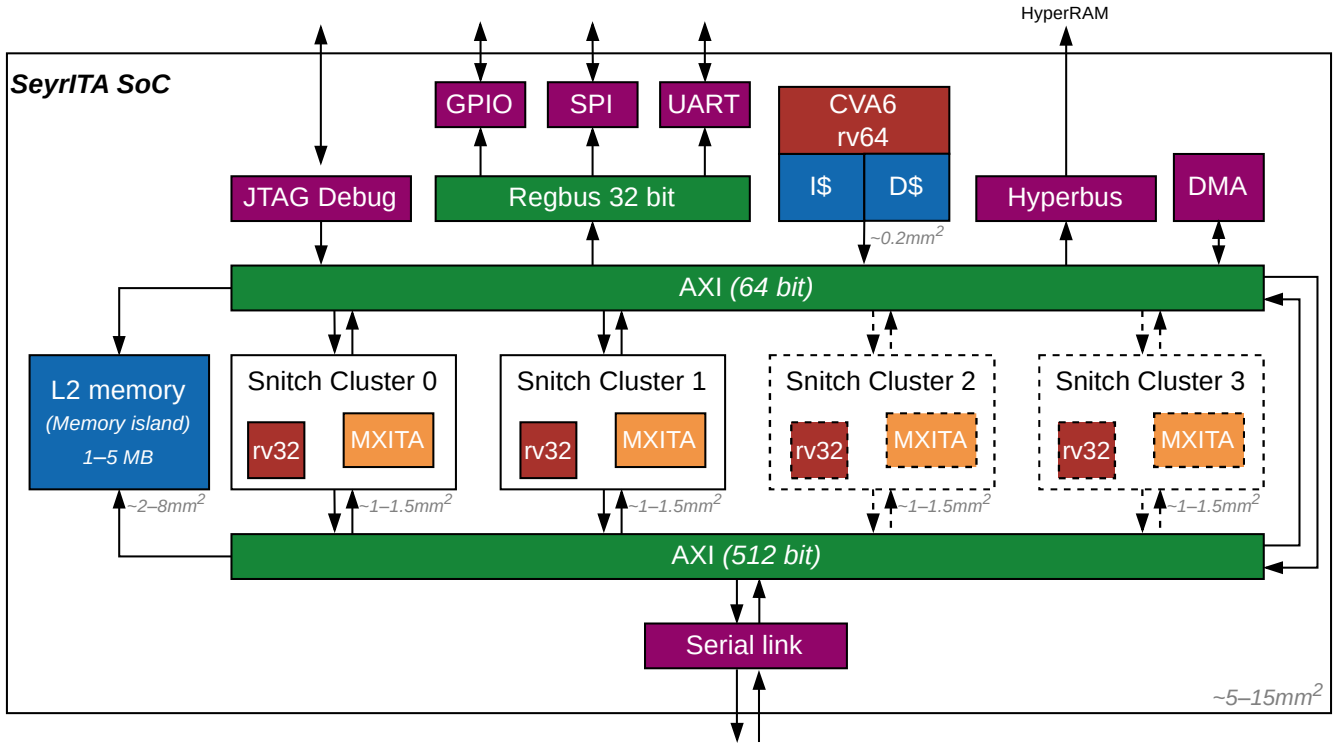

On the Horizon: SeyrITA – Top Level

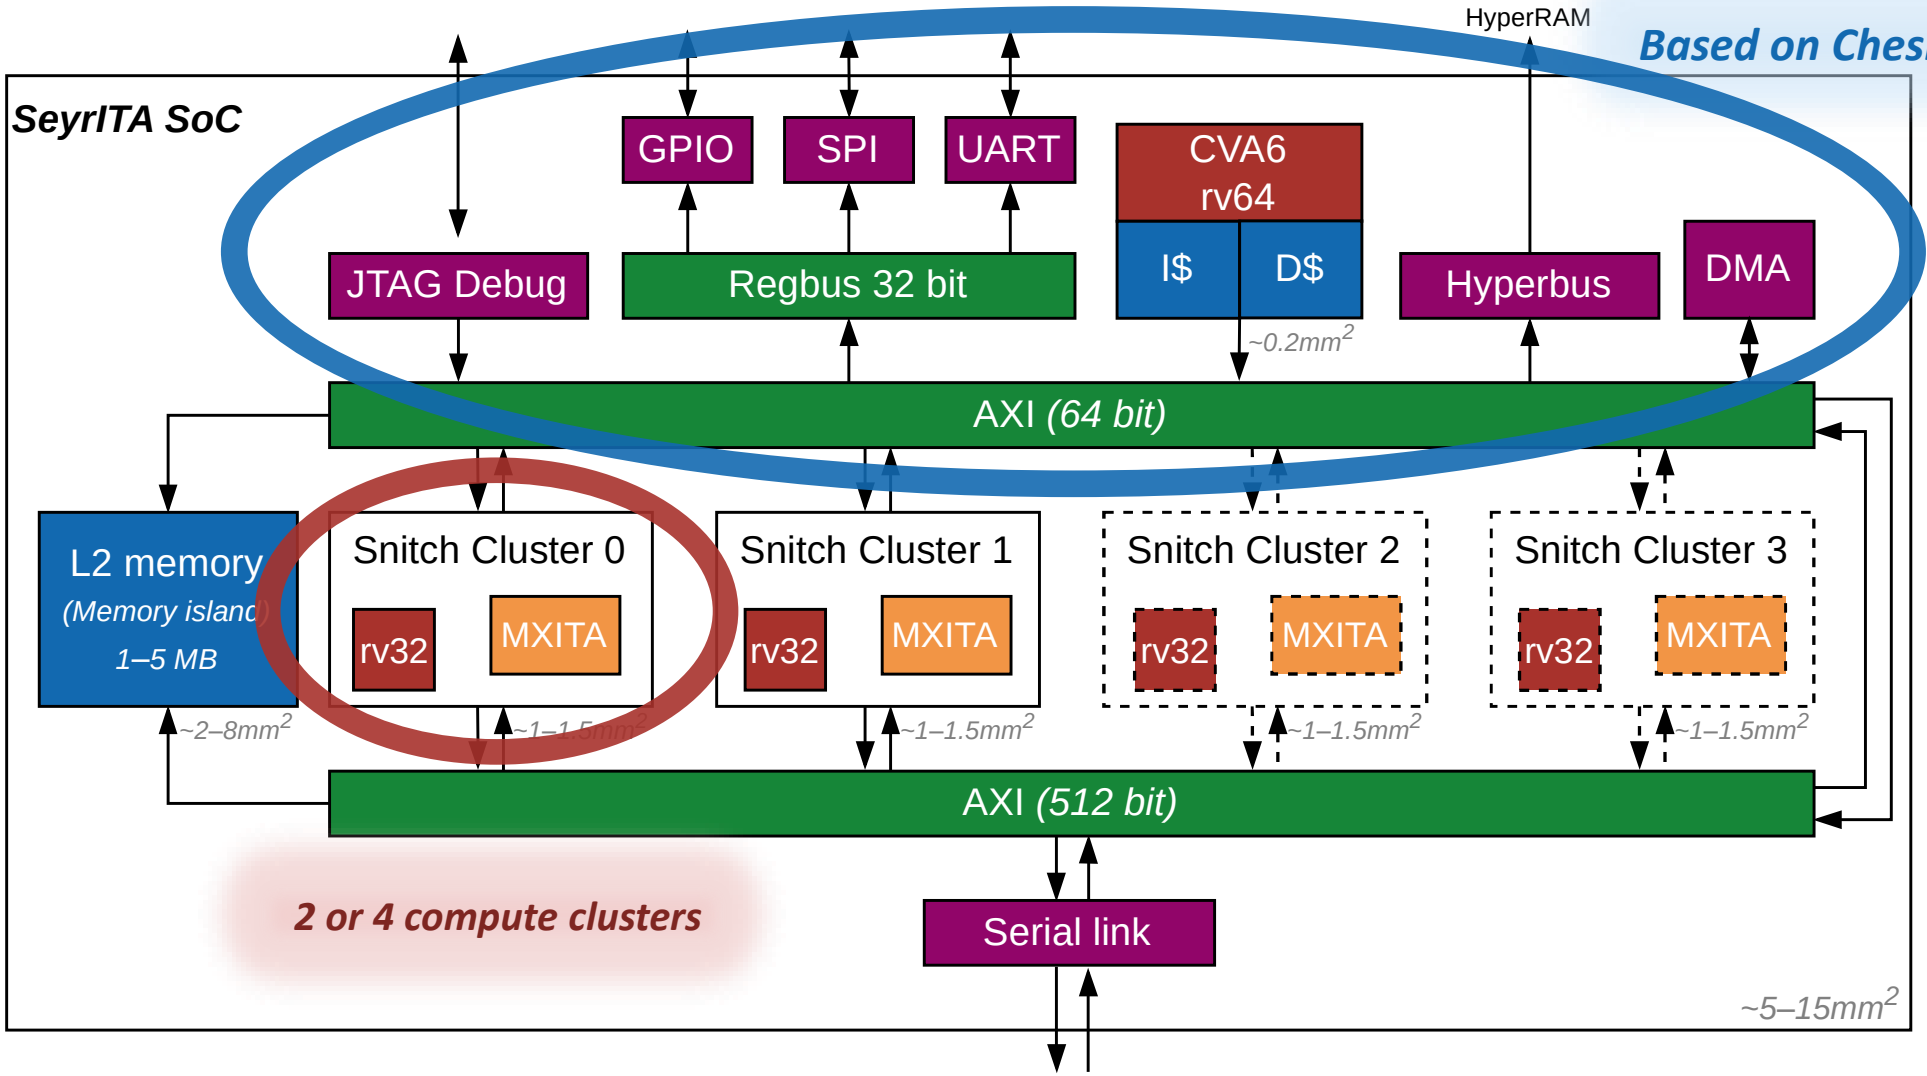

On the Horizon: SeyrITA – Working on the Tapeout

- **Demonstrate** a large **22nm tapeout** with open-source tools
- **Improve tools** and close the **performance** gap
- Identify and **implement missing features** along the way
- **Active Collaboration with**

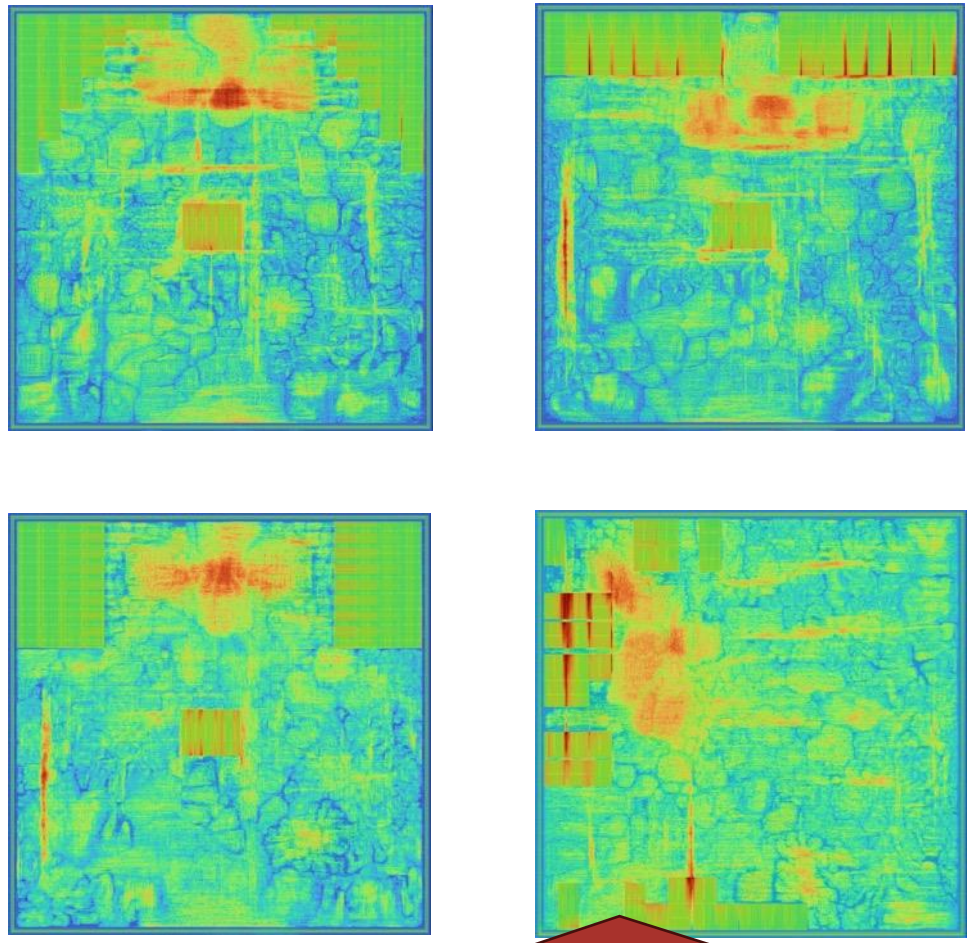

Snitch cluster floorplan exploration

Open PDKs are key to success of open source HW

- **We need more open PDKs**
  - Something in the 65nm – 28nm range would be a **game changer**
  - Drafted an open letter to raise awareness with 300+ signatures

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

- **One possible avenue open Proxy PDKs for a selected technology**
  - A PDK that is *not the same* as the original PDK, developed independently and openly.
  - But designs made using this proxy PDK could be *translated*' to be manufactured using the original PDK through a service center.
  - **Important:** We need to ensure that the technology provider is not *'against*' this idea
    - better yet is supportive
  - This will be tricky, but we think it is doable.

Next steps: **ODE4EC** is an opportunity

- **ODE4EC proposal**
  - Three sub projects: Digital, Analog, Productivity
  - HORIZON-JU-CHIPS-2025-IA-EDA-two-stage proposal,
    - PO stage passed,
    - FPP submitted 19th of September !
- **Austria has major parts in the proposal**
  - JKU, Yosys are vital parts of the open source HW developments
  - We need national co-funding in Austria to make all parts of the proposal supported
- **If accepted funding will start April/May 2026**
  - IHP, JKU, ETHZ will have major parts
  - And work together with other partners

Finally, let's not forget the Austrian colleagues of PULP

From Florian Zaruba's talk on 5th RISC-V Symposium

Philipp

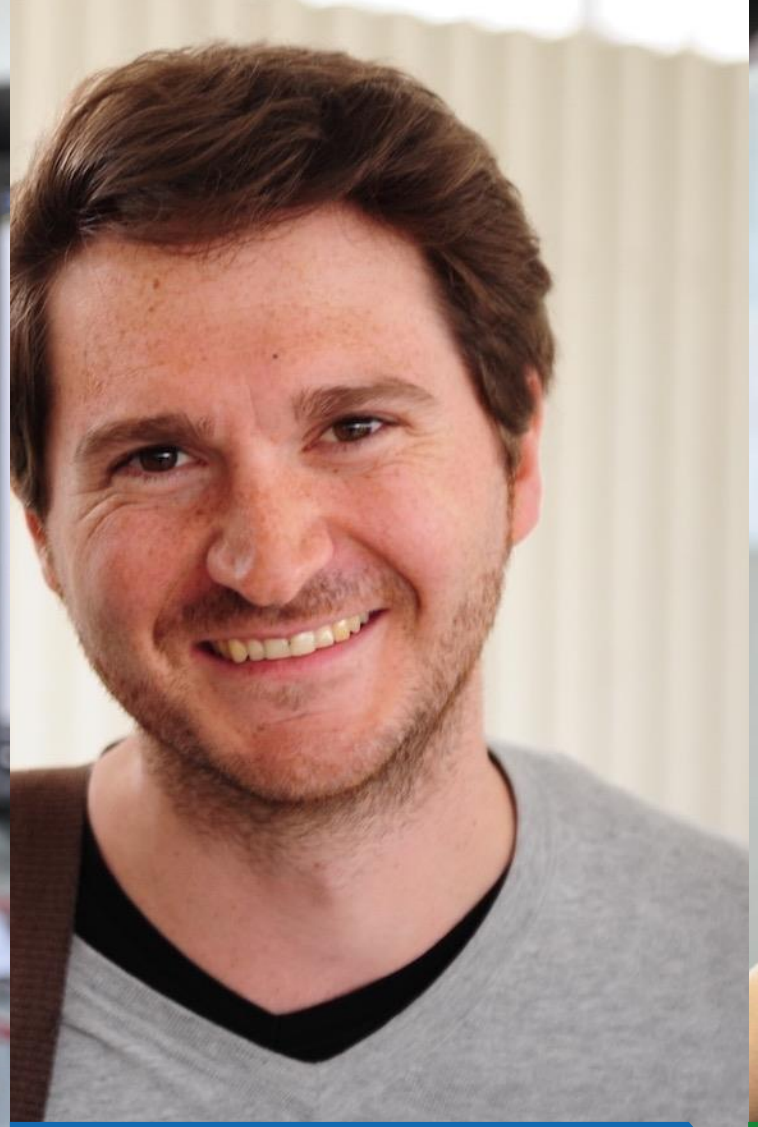

Michael

Florian

Our WWW page contains a wide collection of talks/papers

Most of our talks: https://pulp-platform.org/conferences.html

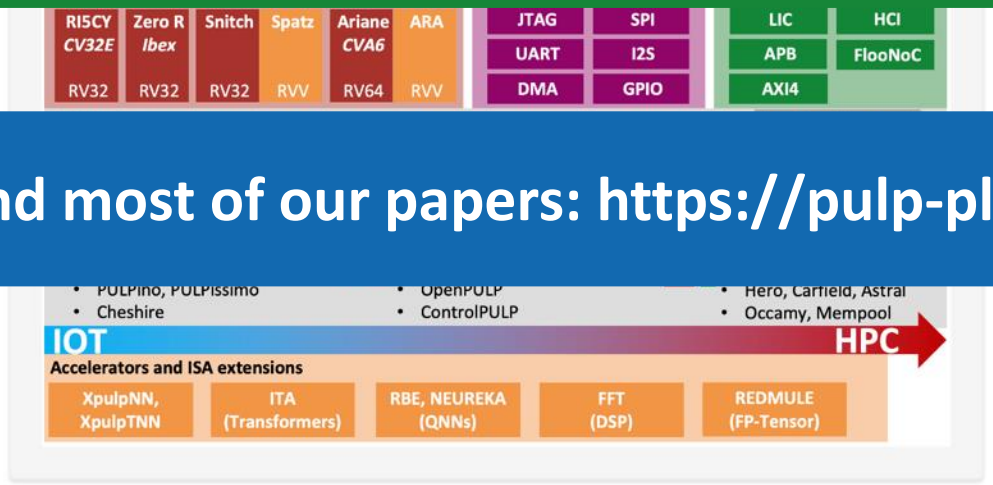

And most of our papers: https://pulp-platform.org/publications.html

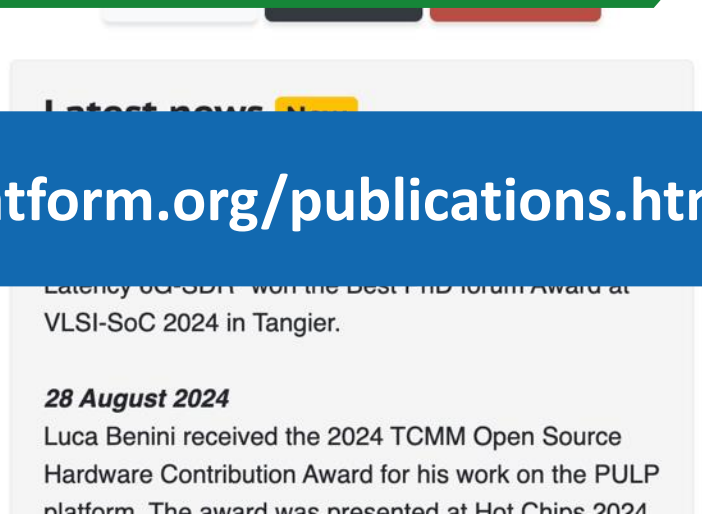

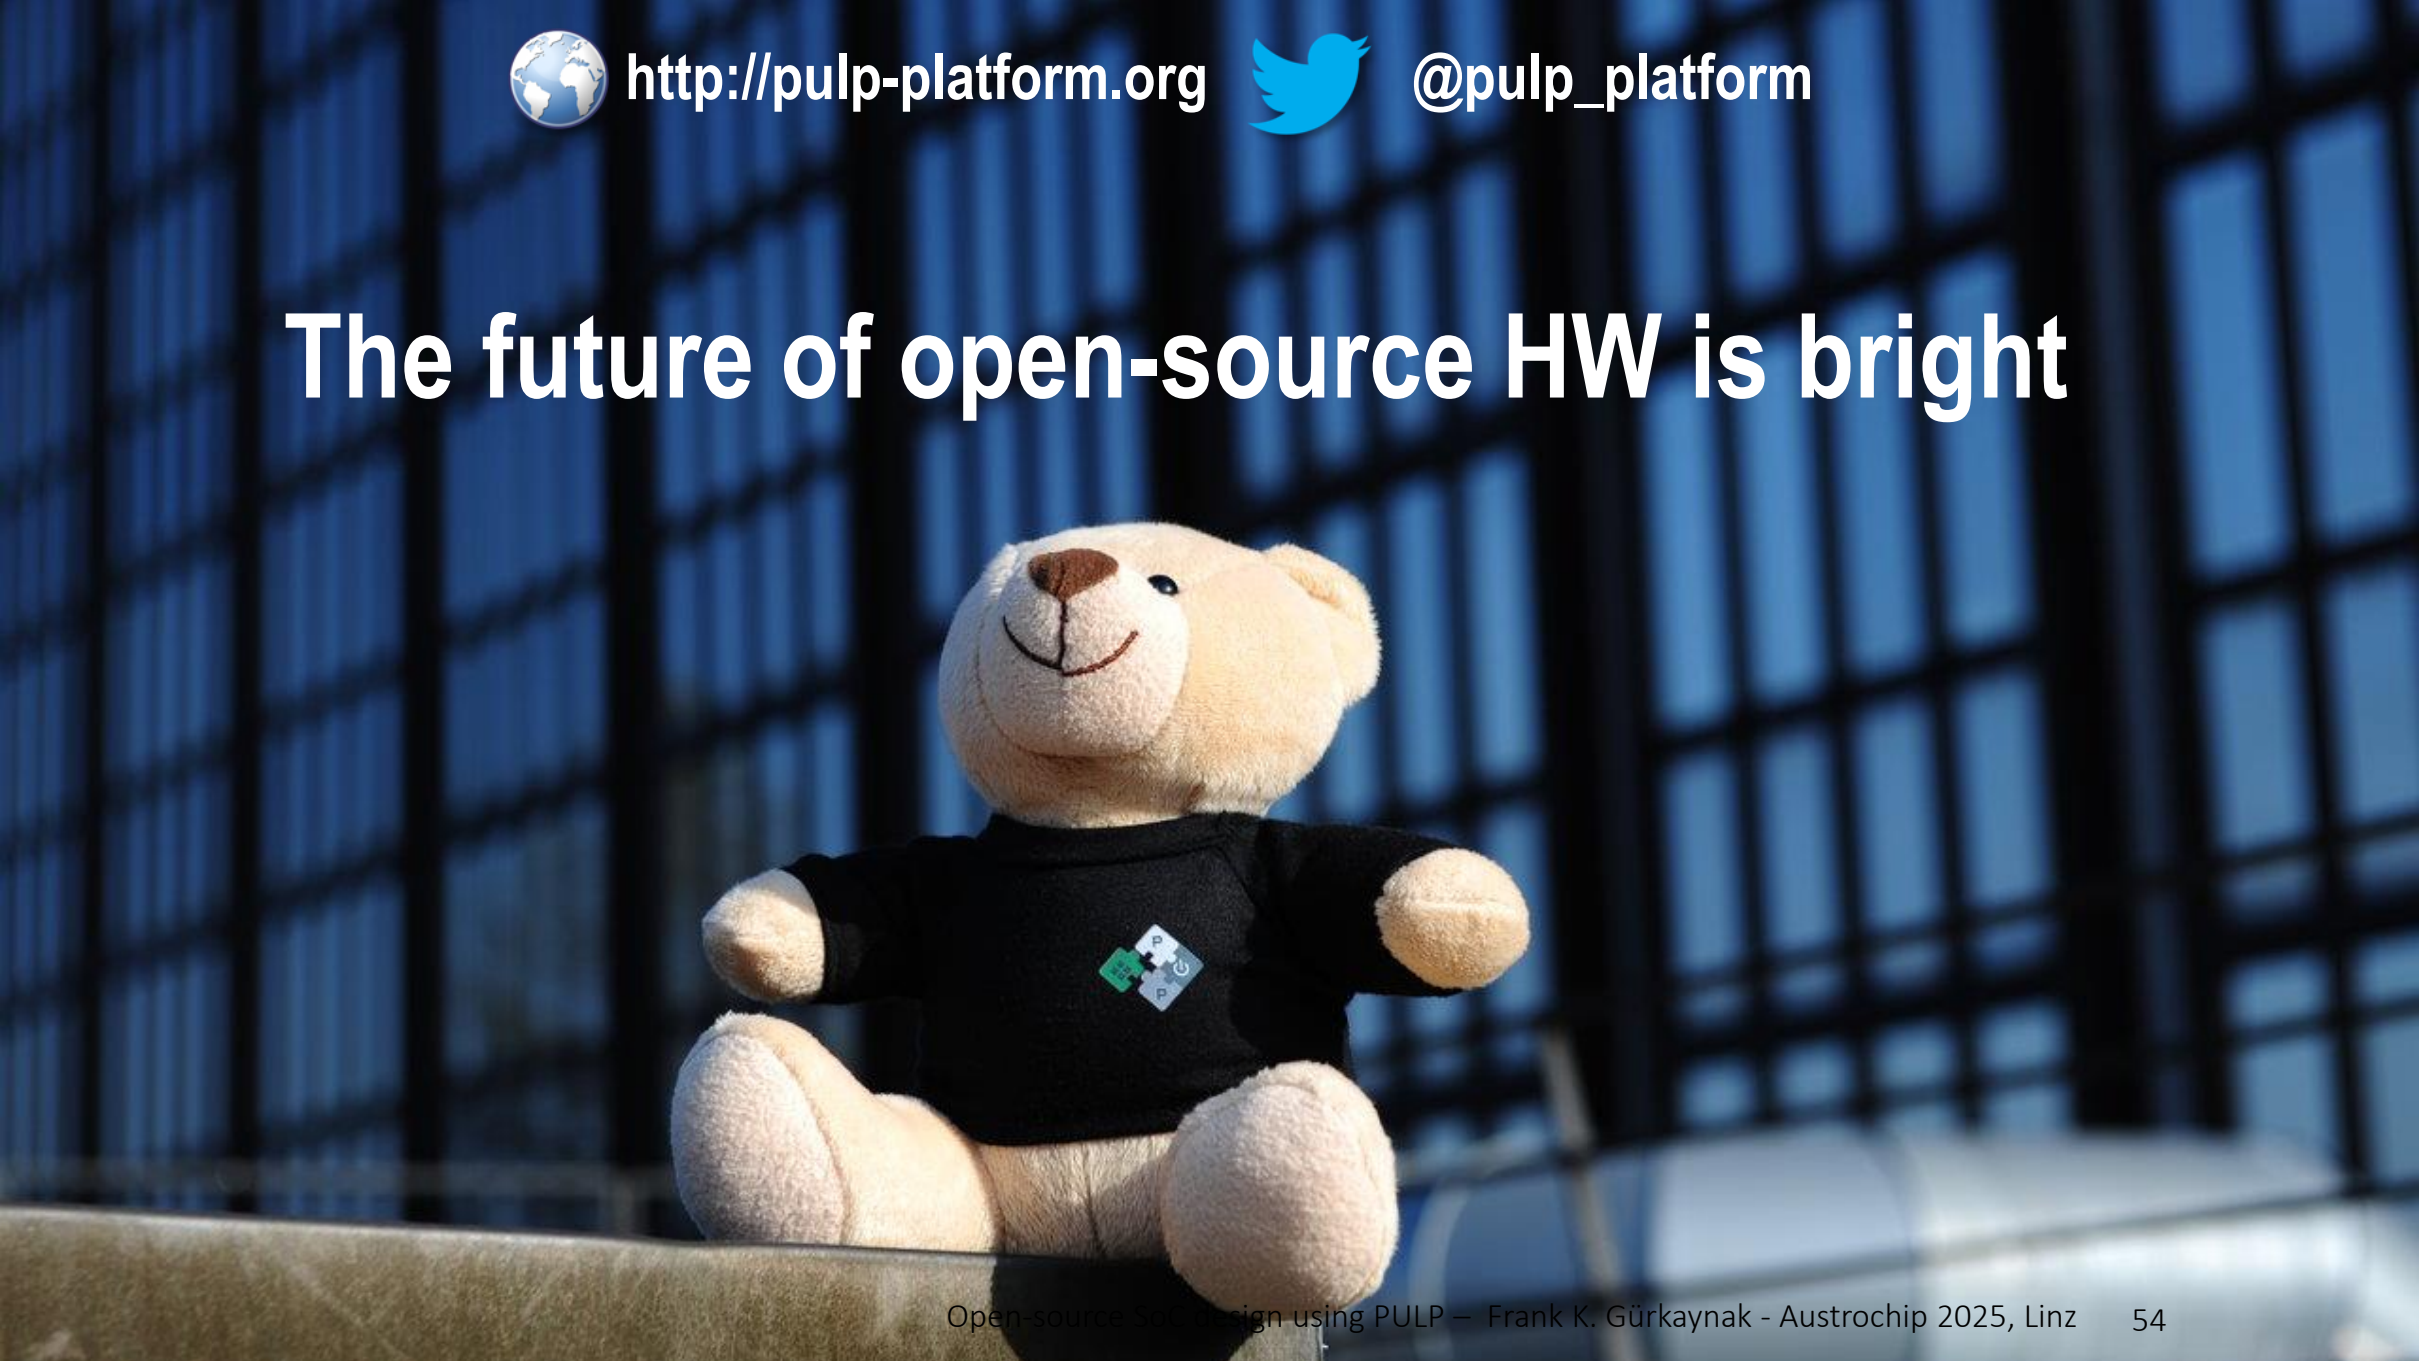
http://pulp-platform.org
@pulp_platform
The future of open-source HW is bright