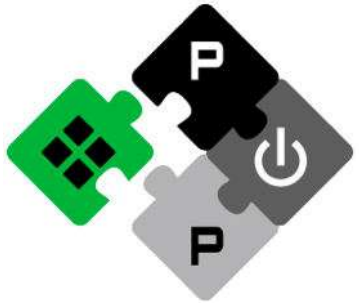

PULP PLATFORM

Open Source Hardware, the way it should be!

From Nano-Drones to Cars

A Precompetitive RISC-V Open Platform for next-generation Vehicles

Luca Benini <luca.Benini@unibo.it,lbenini@ethz.ch>

ETH zürich

http://pulp-platform.org

@pulp_platform

https://www.youtube.com/pulp_platform

Computing in Automotive

MEGATREND – E/E ARCHITECTURE TRANSFORMATION

Today, 1500 chips/car → by 2030, rise to 3000
(EU working document – Chips Act for Europe)

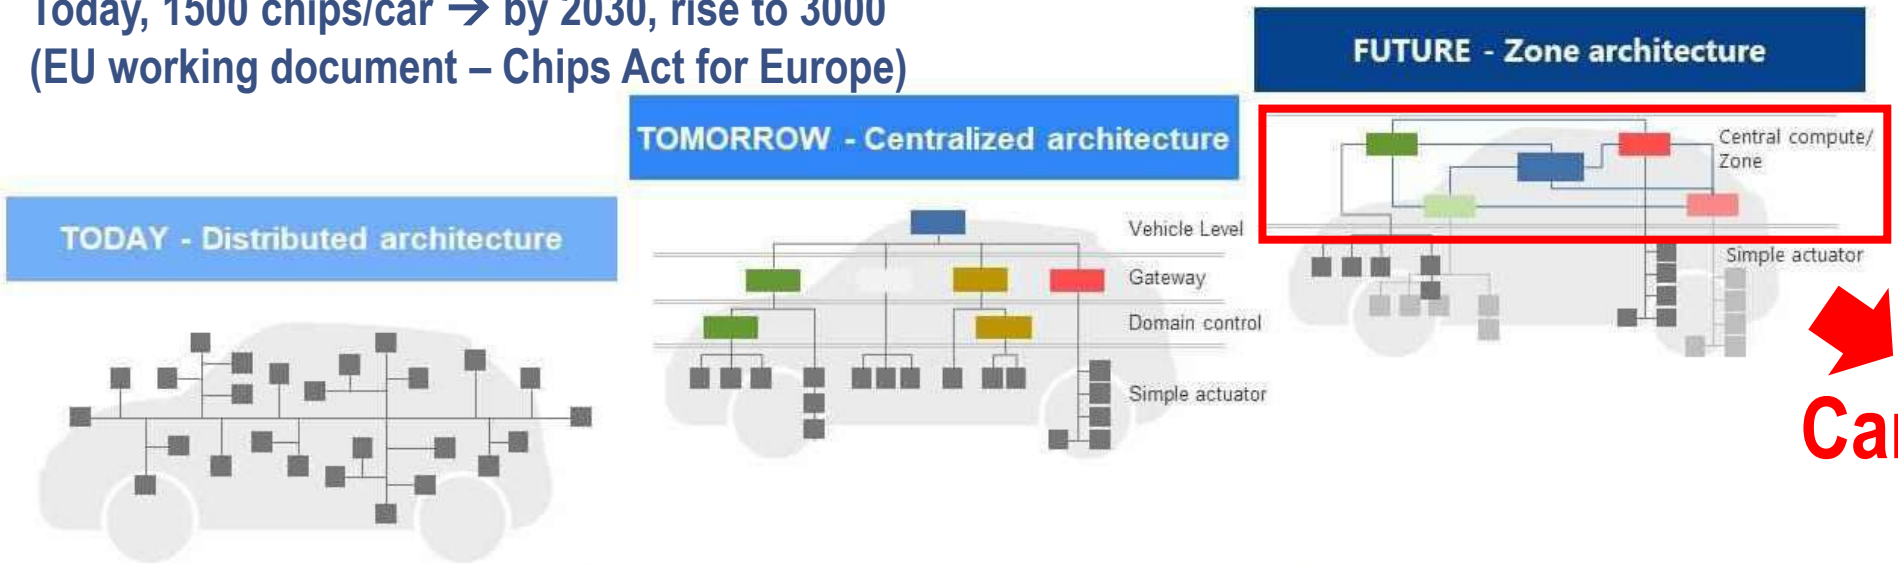

[renesas,bosh]

Cars need lots of computing

- Scalable & Easy to plug-in hardware
- Upgradable & Reusable software
- Safe, redundant and streamlined network

Computing → **SW** → **ISA** → **HW**

```
int a = 1;
int b = N;
do
  a   b;
  b - - 1;
} while (b != 0)
```

CPU

Open Source (Computing) Hardware

Hardware whose design is made publicly available so that anyone can study, modify, distribute, make, and sell the design or hardware based on that design

(source: Open Source Hardware (OSHW) Statement of Principles 1.0)

Very wide definition – includes PCBs and "makers" stuff

I will focus on Open Souce Computing Hardware (OSCHW)

https://www.linuxfoundation.org/research/world-of-open-source-europe-spotlight

OSCHW Needs an Open Source ISA

A modern, open, free ISA, extensible *by construction*
Endorsed by 3000+ Members & large SW ecosystem
An open ISA is a prerequisite!

Risc-V international moved to Zürich, CH
for international neutrality, 1/3 of members from EU

wide in RV32I, 64 in RV64I, and 128 in RV128I (x0=0). RV64I/128I add 10 instructions for the wider formats. The RVI base of <50 classic integer RISC instructions is required. Every 16-bit RVC instruction matches an existing 32-bit RVI instruction. See risc.org.

width matches the widest precision, and a floating-point control and status register fcsr. Each larger address adds some instructions: 4 for RVM, 11 for RVA, and 6 each for RVF/D/Q. Using regex notation, { } means set, so L{D|Q} is both LD and LQ. See risc.org. (8/21/15 revision)

Rich RISC-V Ecosystem Available Today

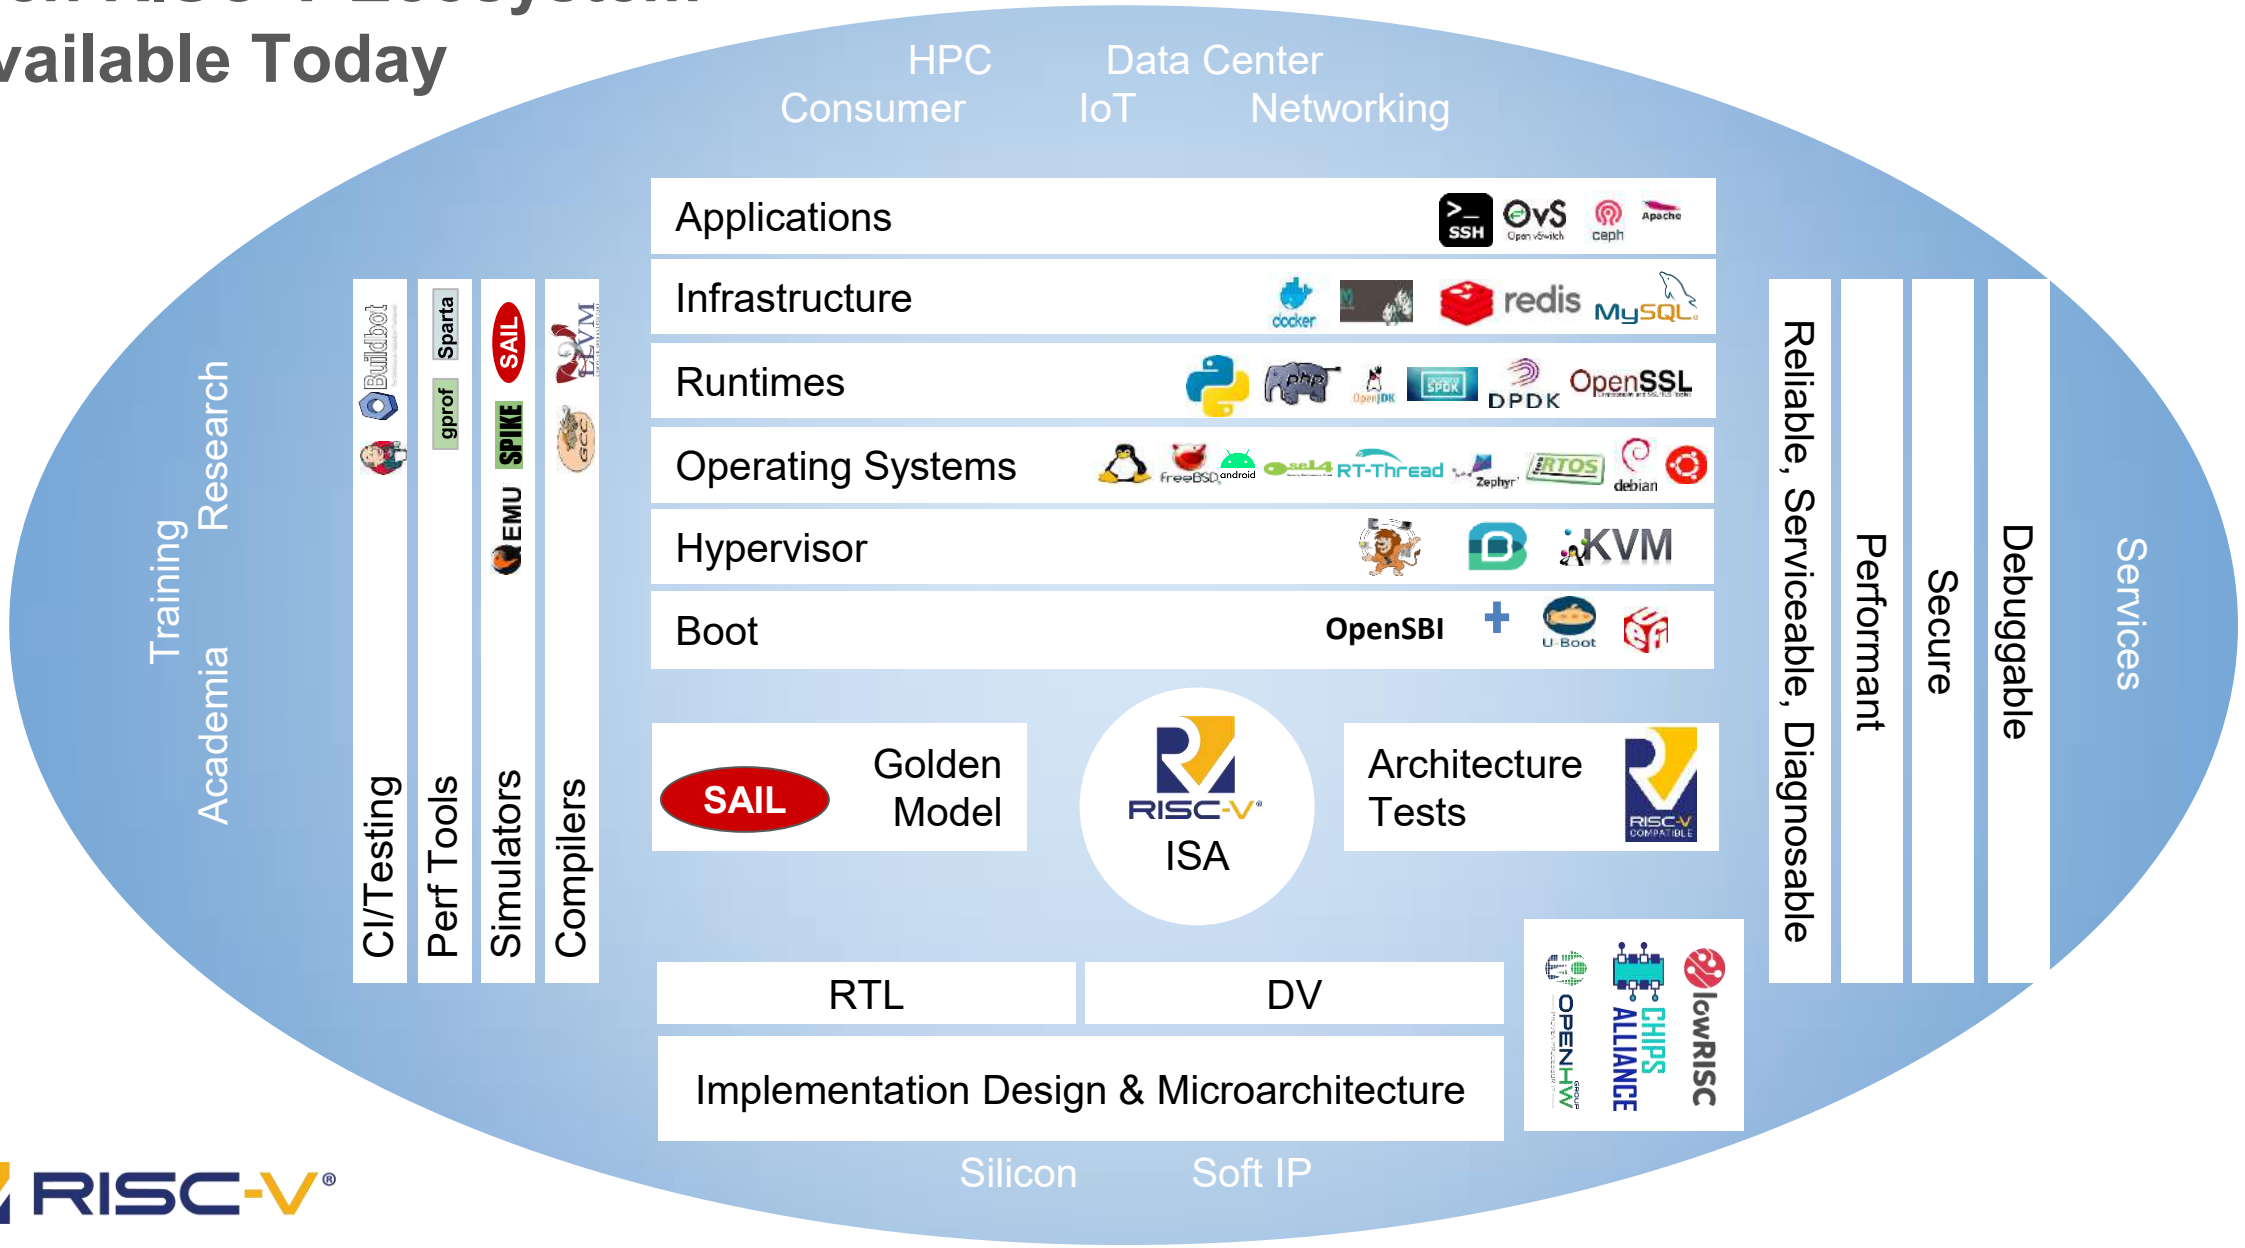

RISC-V CPU core market grows 114.9%
CPU cores by

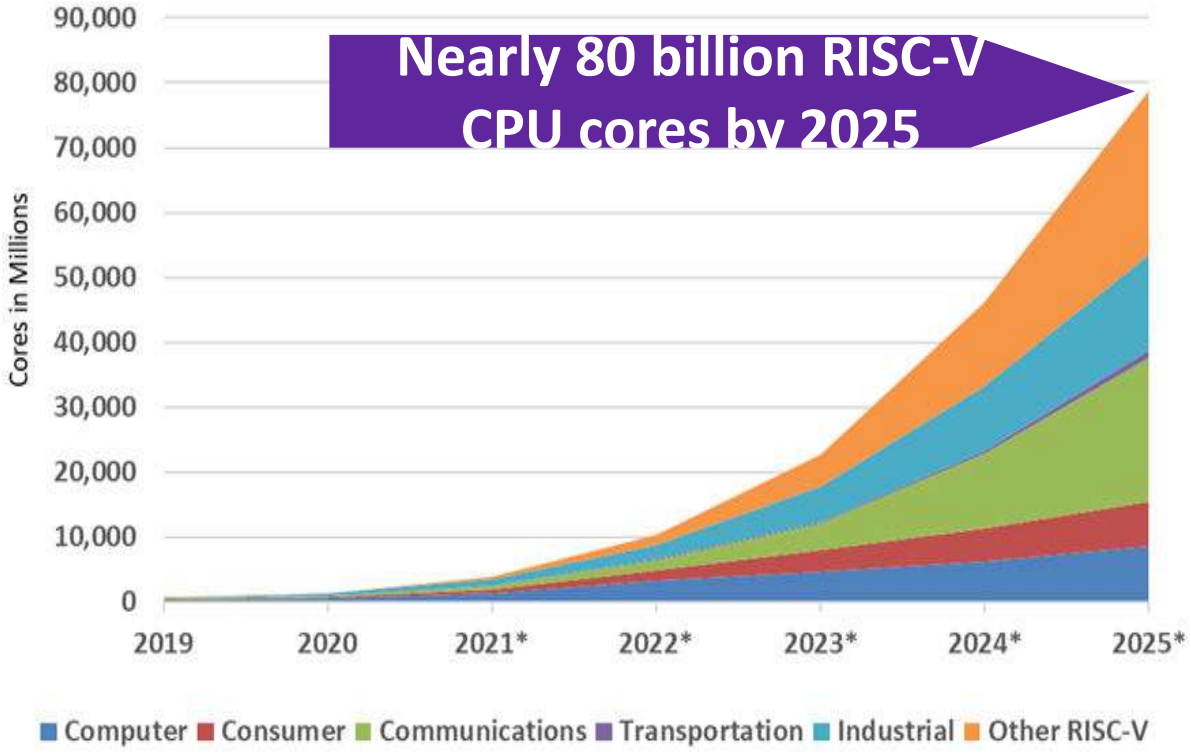

Source: Semico Research Corp, March 2021

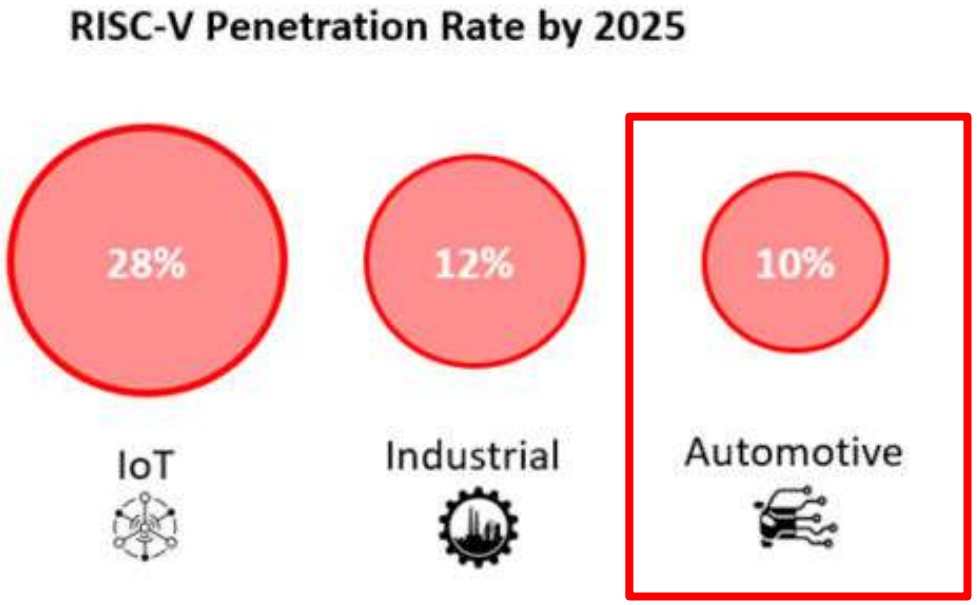

"The rise of RISC-V cannot be ignored... RISC-V will shake up the $8.6 Billion semiconductor IP market."

-- William Li, Counterpoint Research

Source: Counterpoint Research, September 2021

Automotive

MobileEye EyeQ Ultra vision advanced driver assist systems chips for 176 trillion ops per second with 12 RISC-V CPU cores.

Andes ISO 26262 Functional Safety ASIL D Dev Process Certification for RISC-V embedded safety with Andes processors. Used in Renesas R9A02G020 MCU ASSP for motor control systems

Renesas and SiFive partner on high-end automotive applications. SiFive licenses RISC-V core IP to Renesas (Automotive™ E6-A, X280-A, and S7-A for infotainment, cockpit, connectivity, ADAS)

Imagination Technologies GPU linked by a RISC-V core for ASIL-B level designs with ISO 26262 safety critical certification.

Codasip Security and safety embedded by design in processor design automation tools

What about OSCHW?

Open ISA < OSCHW… But what is OSCHW, then?

- The first stage of the silicon production pipeline → **HDL source code**
- Later stages contain closed IP of various actors → not open source by default

HDL

```
module toplevel(clock,reset);
  input clock;
  input reset;

  reg flop1;
  reg flop2;

  always @ (posedge reset or posedge clock)
    if (reset)
      begin
        flop1 <= 0;
        flop2 <= 1;
      end
    else
      begin
        flop1 <= flop2;
        flop2 <= flop1;
      end
endmodule
```

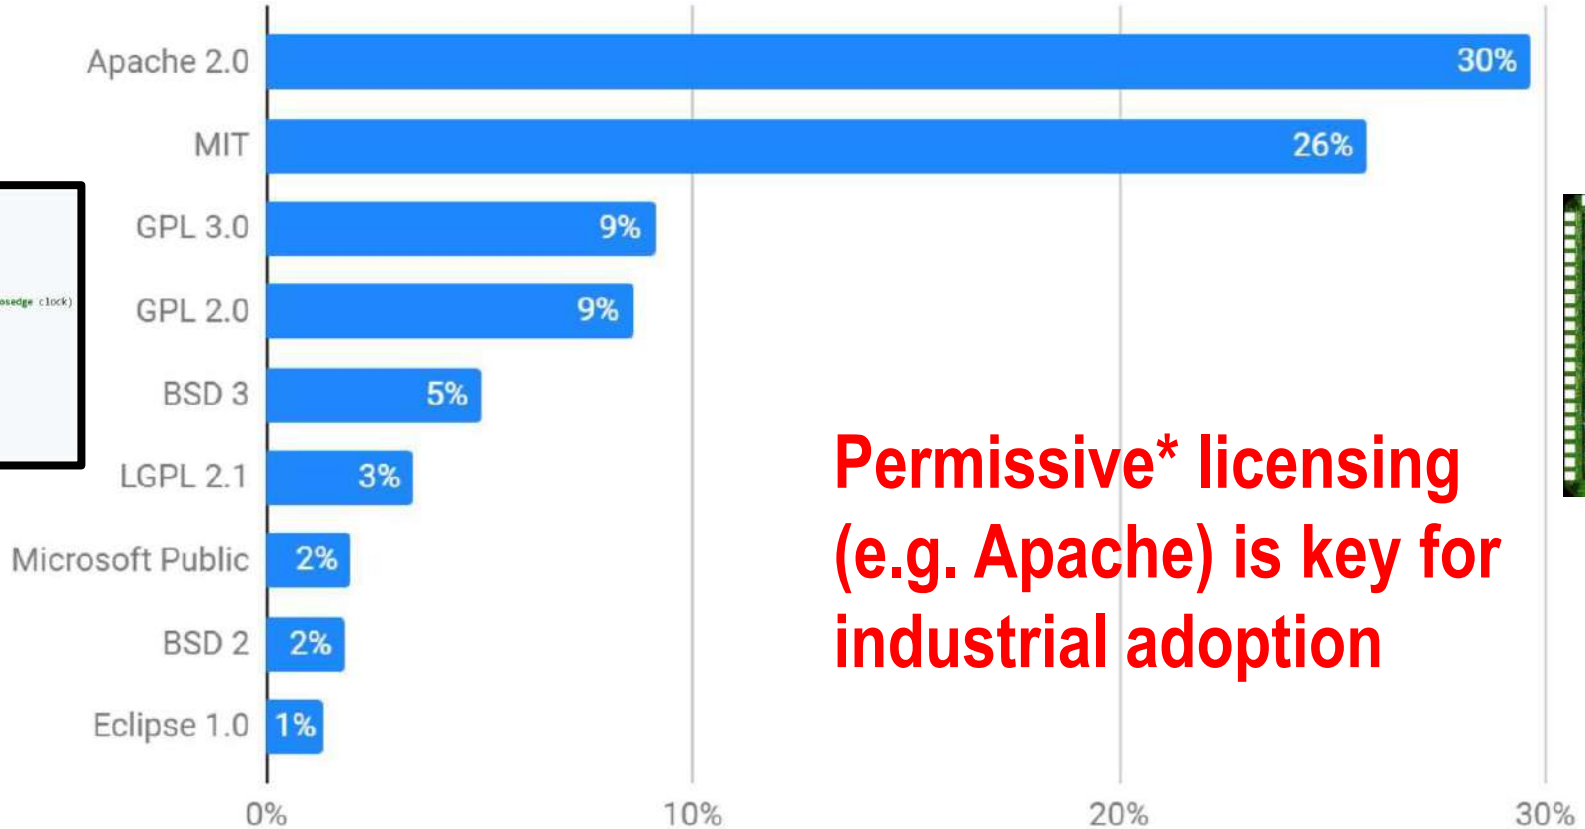

Chip

Permissive* licensing (e.g. Apache) is key for industrial adoption

*See: https://cern-ohl.web.cern.ch/ (CERN-OHL-S, -W, -P)

Cores, and more! Many IPs for a Platform

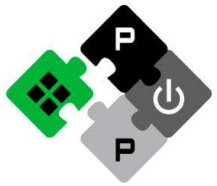

RISC-V Cores

RI5CY	Ibex	Snitch	Ariane + Ara
32b	32b	32b	64b

Peripherals

JTAG	SPI
UART	I2S
DMA	GPIO

Interconnect

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

Platforms https://github.com/pulp-platform/

Tens of active users, many use-cases
HW, SW specialization, verification, documentation, training
Cannot be sustained by one University, or two…

Single Core
- PULPino
- PULPissimo

Multi-core
- Open-PULP
- PULP-PM

Multi-cluster
- Hero
- Occamy

IOT → HPC

Accelerators

HWCE (convolution)	Neurostream (ML)	HWCrypt (crypto)	PULPO (1st ord. opt)

Industrial OSCHW: Feel the momentum!

Ibex RISC-V core, flash interface, communications ports, cryptography accelerators, and more.

35+ Contributors

1300+ Contributions

470 GitHub Issues

PULP Parallel Ultra Low Power → lowRISC

Zero-Riscy Ibex

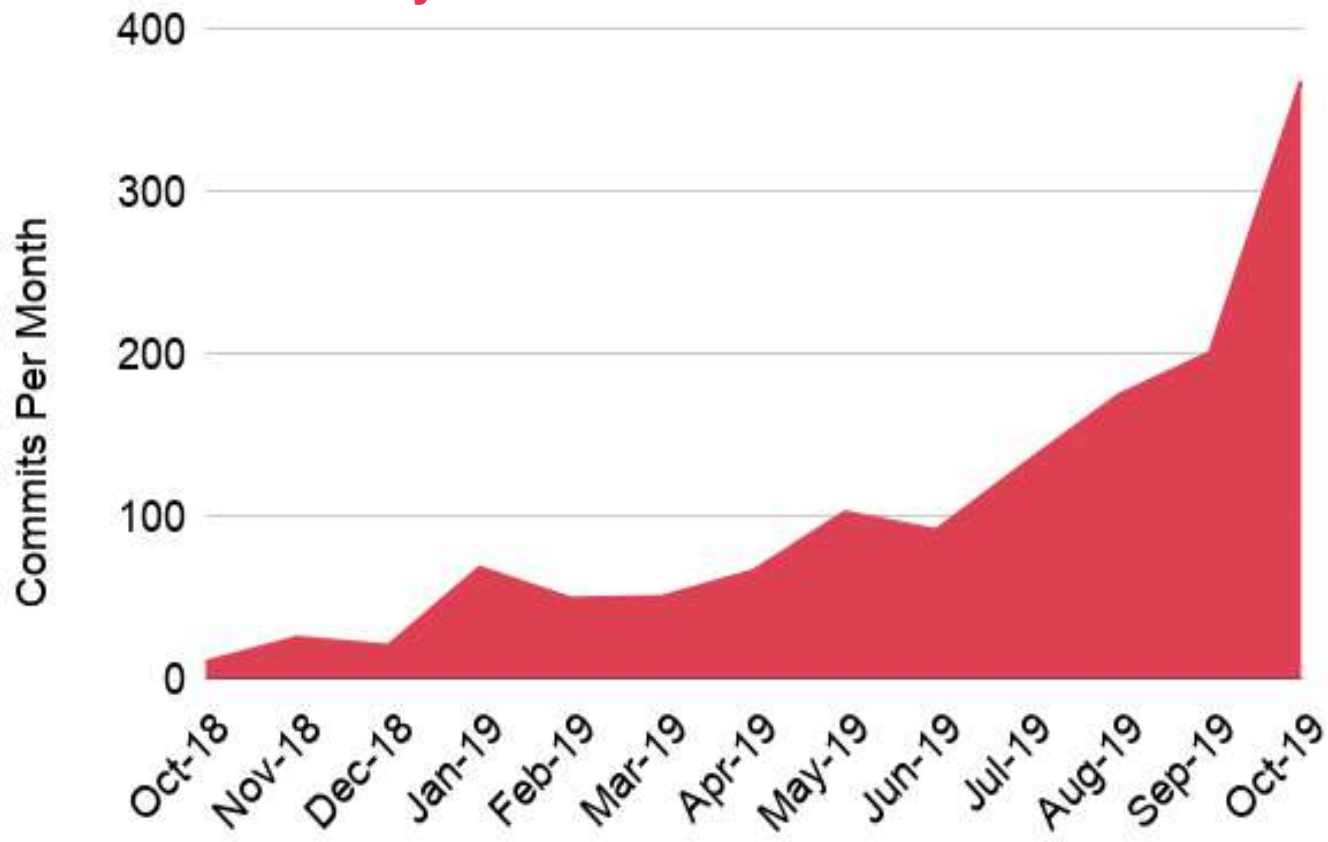

A vertical, application-focused open approach

- OpenTitan is the first OSCHW project building a transparent, high-quality reference design for silicon root of trust (RoT) chips.
- **Founding** partners

HWG: A Fast-Growing Precompetitive Industrial Ecosystem

OPENHW GROUP™ PROVEN PROCESSOR IP

Rick O'Connor (OpenHW CEO, former RISC-V foundation director)

- **OpenHW Group** is a not-for-profit, global organization (EU,NA,Asia) where HW and SW designers collaborate in the development of open-source cores, related IP, tools and SW such as the **Core-V** family
- **OpenHW Group** provides an infrastructure for hosting high quality open-source HW developments in line with industry best practices.

PULP
Parallel Ultra Low Power

RI5CY, ARIANE

CORE-V™

80+ members!

Strengthening the Ecosystem

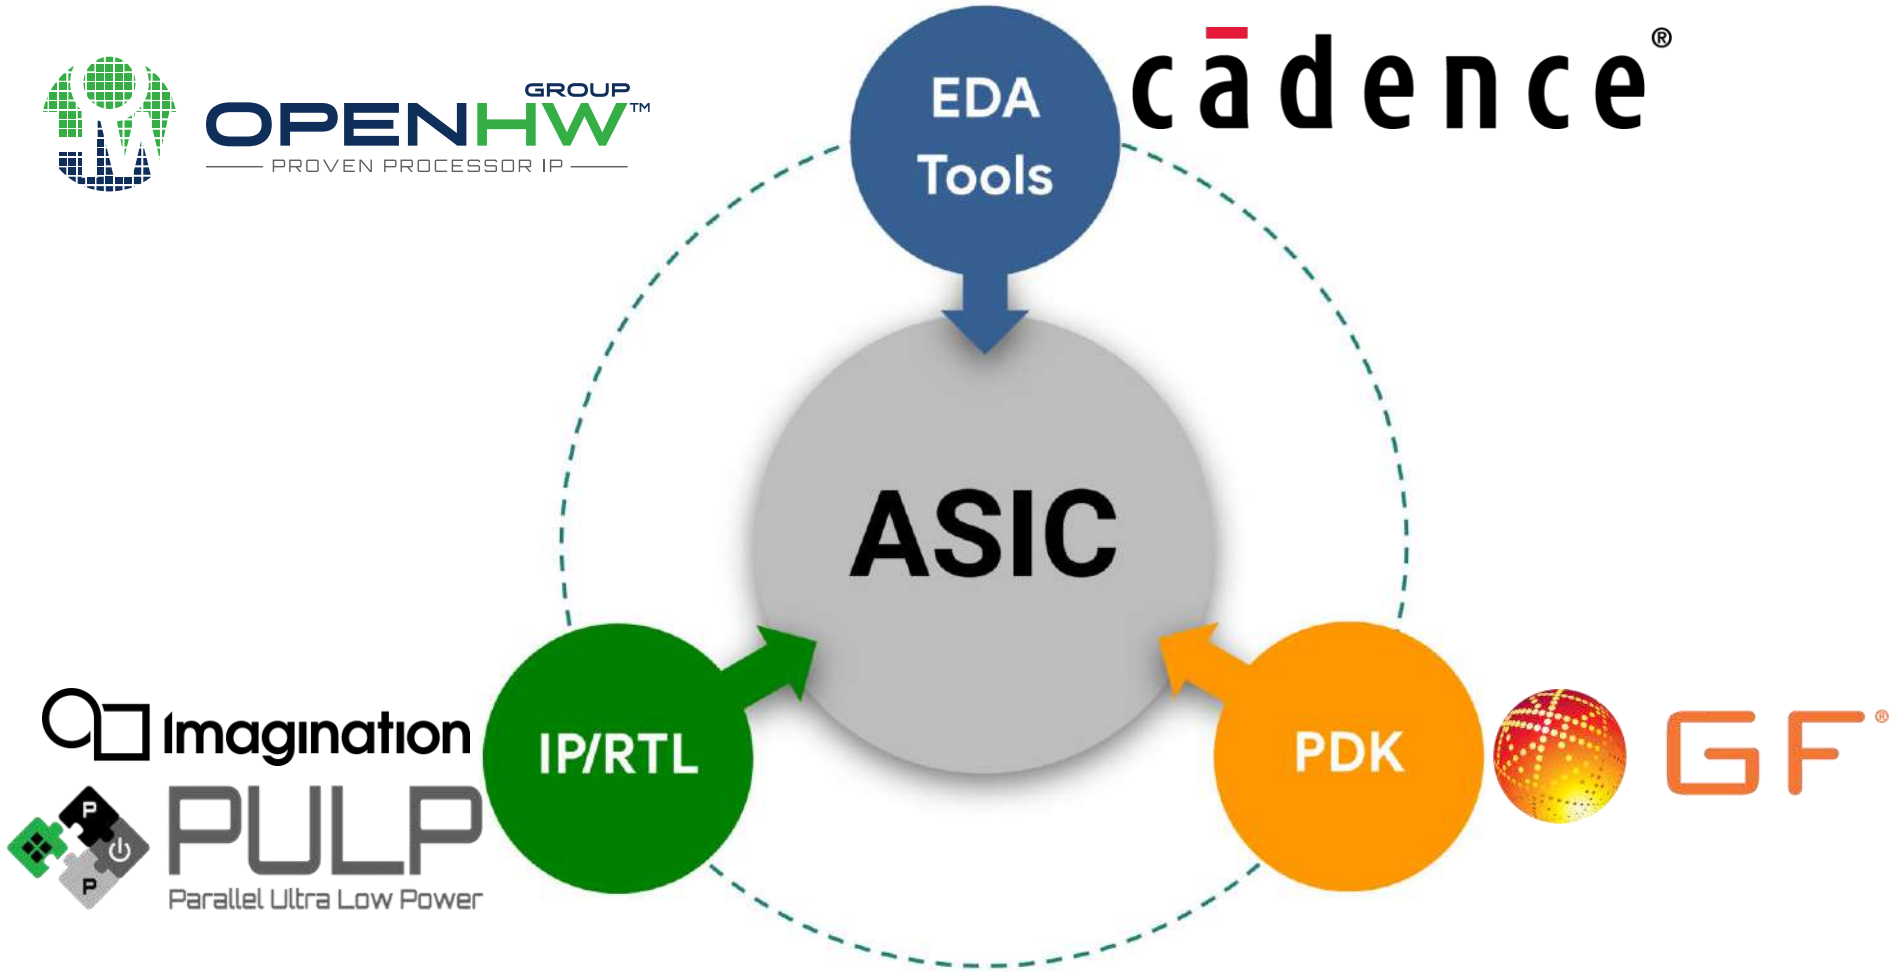

Start Small: Open Platform for Autonomous Nano-Drones

Advanced autonomous drone

[1] A. Bachrach, “Skydio autonomy engine: Enabling the next generation of autonomous flight,” IEEE Hot Chips 33 Symposium (HCS), 2021

https://www.skydio.com/skydio-2-plus

- 3D Mapping & Motion Planning
- Object recognition & Avoidance
- 0.06m2 & **800g of weight**
- Battery Capacity **5410mAh**

Nano-drone

https://www.bitcraze.io/products/crazyflie-2-1

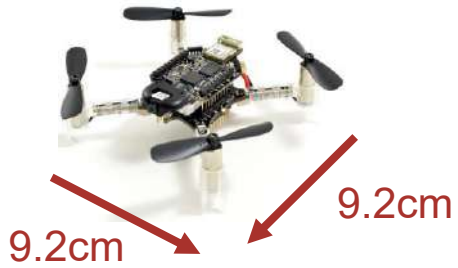

- Smaller form factor of 0.008m2
- Weight **27g (30X lighter)** KG
- Battery capacity **250mAh (20X smaller)**

Can we fit sufficient intelligence in a 30X smaller payload, 20X lower energy budget?

ETH zürich

Achieving True Autonomy on Nano-UAVs

Multiple, complex, heterogeneous tasks at high speed and robustness **fully on board**

Object detection

Obstacle avoidance & Navigation

Environment exploration

Multi-GOPS workload at extreme efficiency → P_{max} 100mW

Multiple Heterogeneous Accelerators

Brain-inspired: Multiple areas, different structure different function!

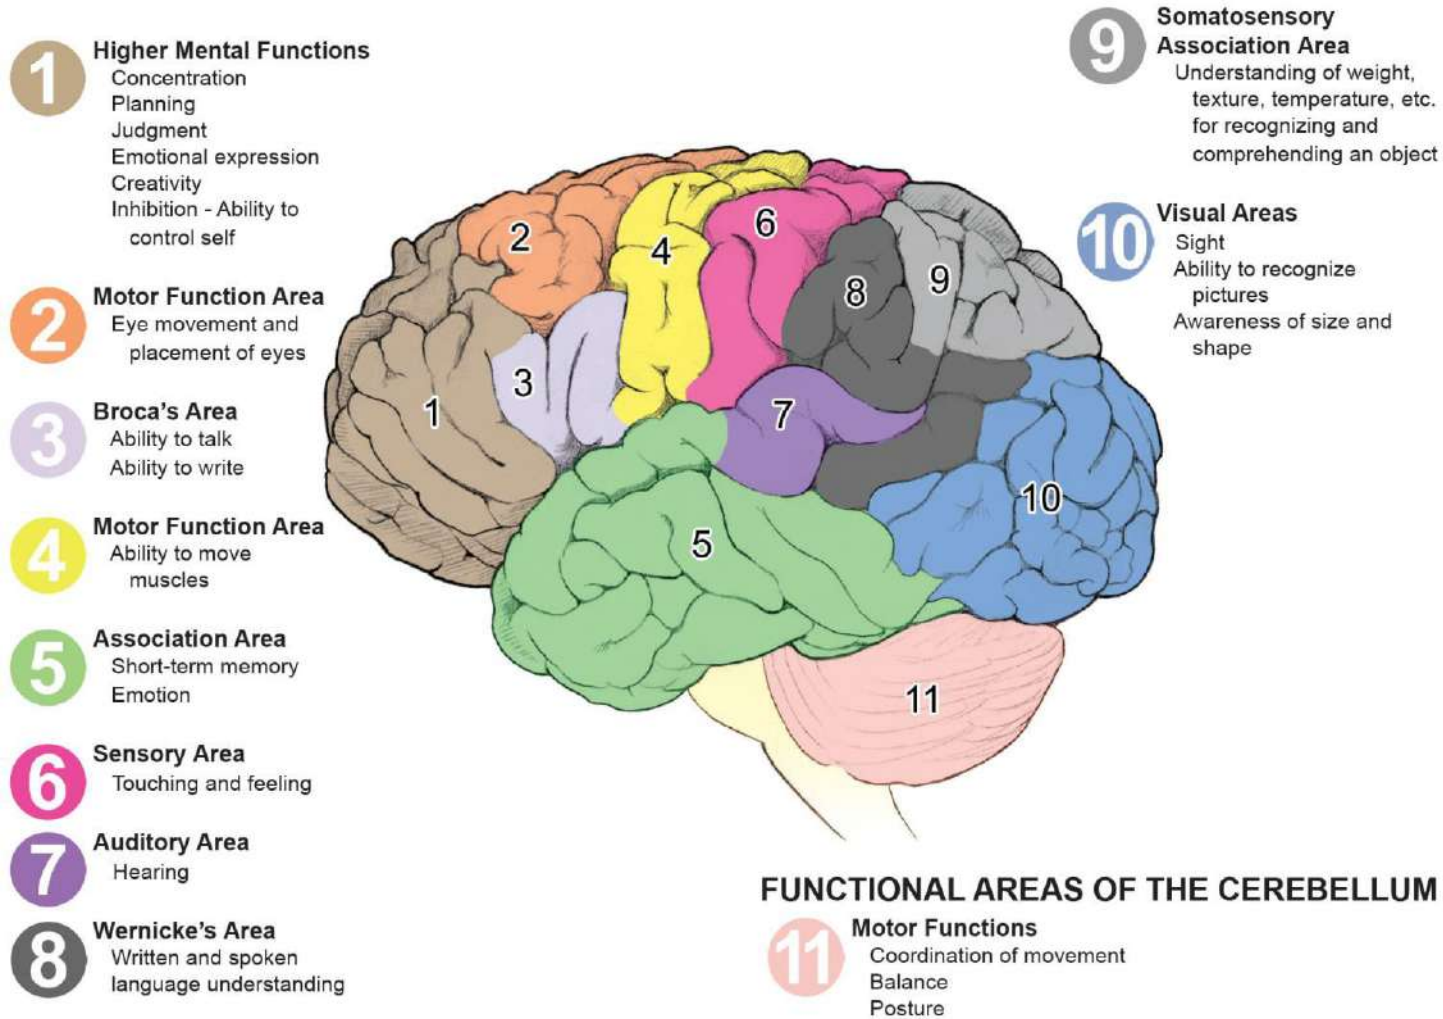

Multiple Heterogeneous Accelerators

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**
- PULPO – Floating point online optimizer (advanced control, state estimation…)

$$\mathbf{z} = \mathbf{A}\mathbf{x} + \mathbf{y} \quad (\mathcal{M})$$
$$\mathbf{z} = \mathbf{y} - \tau\mathbf{A}^H\mathbf{x} \quad (\mathcal{H})$$
$$\mathbf{z} = \mathrm{prox}(\rho\mathbf{x}) \quad (\mathcal{P})$$

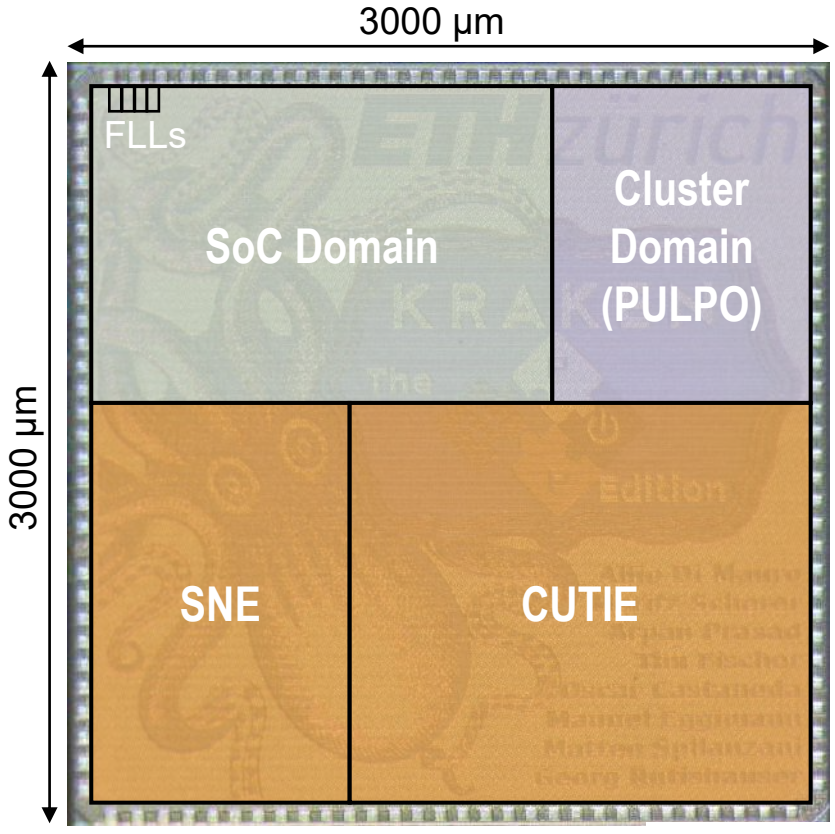

[Di Mauro HotChips22]

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**

CUTIE: Minimize Switching Activity & Data Movement

[Scherer et al. TCAD22]

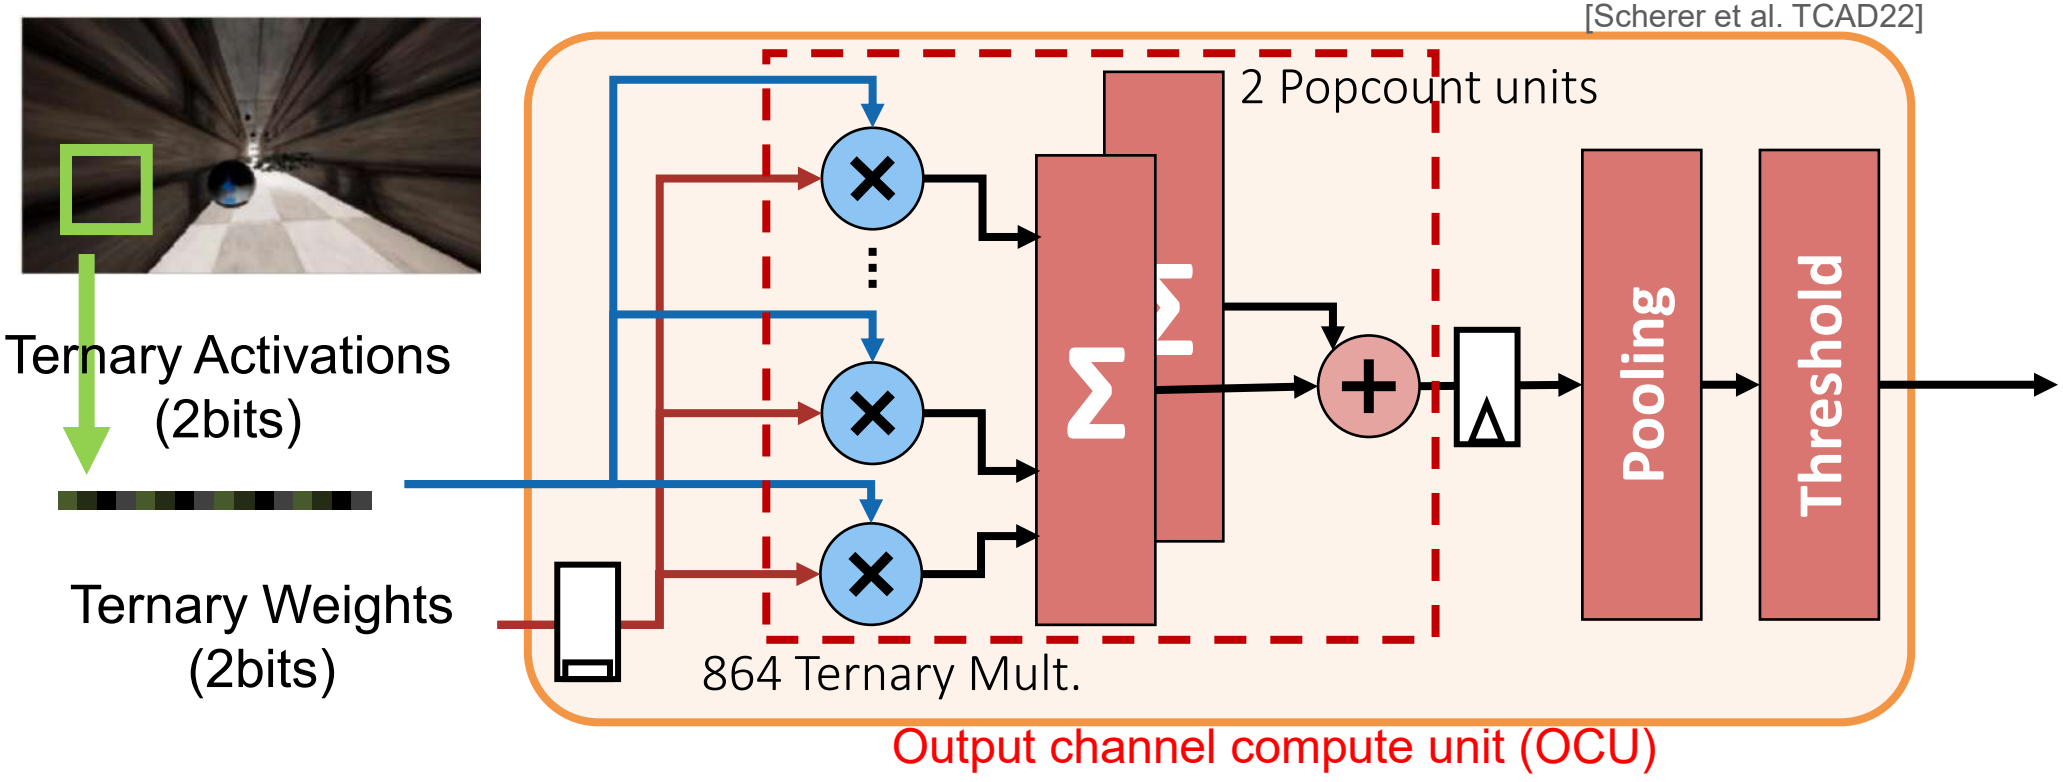

- KxK window on all input channels unrolled, cycle-by-cycle sliding
- Completely unrolled inner products one output activation per cycle!
- Zeros in weights and activations, spatial smoothness of activations reduce switching activity
- 96 OCUs, 96 Input channels, 3x3 kernels: 96 * 96 * 3 * 3 = 82'944 TMAC/cycle

Aggressive quantization and full specialization

Different Sensor Type, different Acceleration Engine

Event Sensors:
DVS
Ultra-low latency
Energy- proportional interface

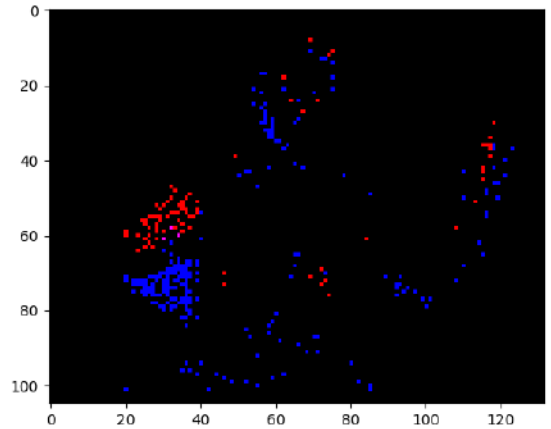

[Di Mauro et al. DATE22]

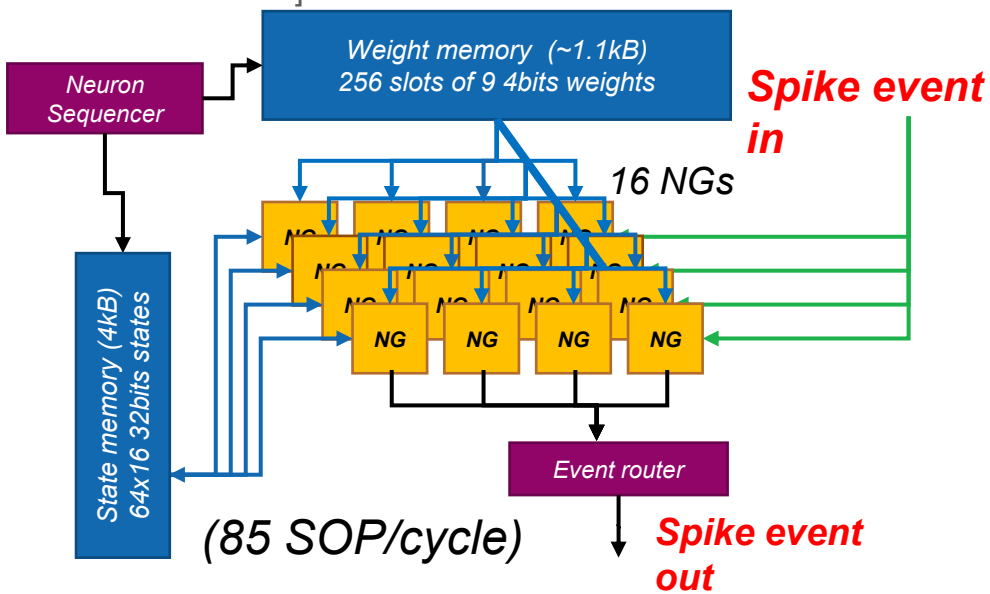

- **SNE Engine**: 16 Adaptive-LIF neuron data paths (NG). A NG executes one Synaptic Operation (SOP) per cycle
  - 1 SOP = 1 4b-ADD + 2 8b-MUL + 1 8b-ADD + 1 8b-CMP
- For fully connected layers one NG is time-shared for 64 virtual neurons
- Optimized buffering and neuron state update for 64x16 neurons in just 12 cycles for a 3x3 event receptive field
  - Equivalent number of 85 SOP/cycle per engine (682 SOP/cycle on 8 engines)

SNE works seamlessly with DVS (event-based) sensors

Not only Perception: SLAM, Planning

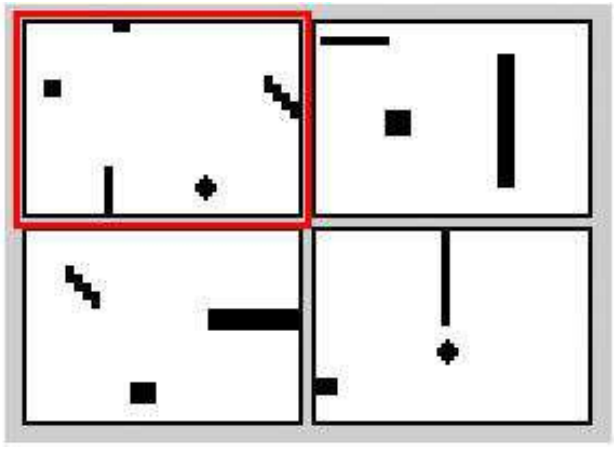

Particle filter-based

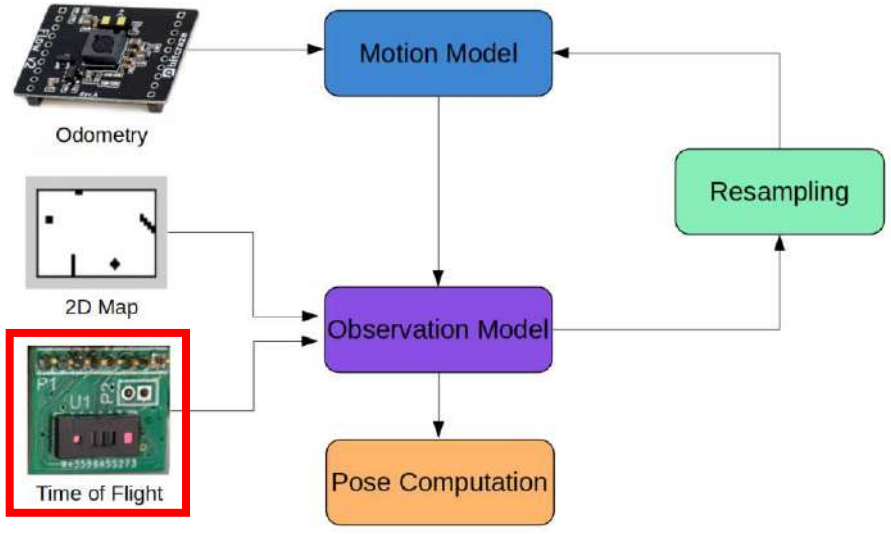

Convergence + Low ATE for Npart > 1024, 2ToF, FP16 acceptable

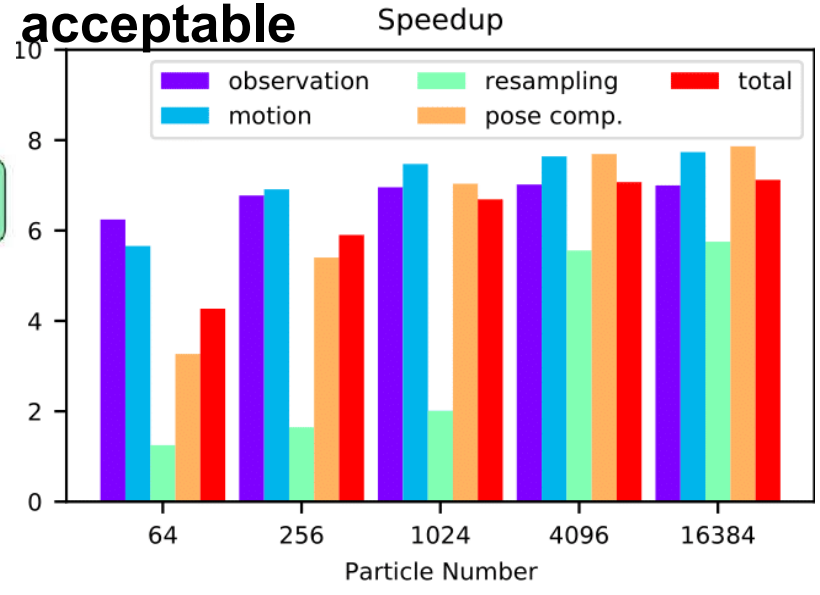

12MHz, 1Kpart. 13mW, 60msec
400MHz, 1Kpart 61mW, 1msec
400MHz 16Kpart 61mW, 30msec

ETH zürich

Advancing the SOA on all tasks

RISC-V Cluster

- Comparable 32bits-8bits SOA Energy efficiency to other PULPs [7]
- **The highest energy efficiency on sub-byte SIMD operations (4b-2b)**

SNE

- **1.7X** higher than SOA [5] energy/efficiency

CUTIE

- **2X** higher energy efficiency improvement over SOA [6]

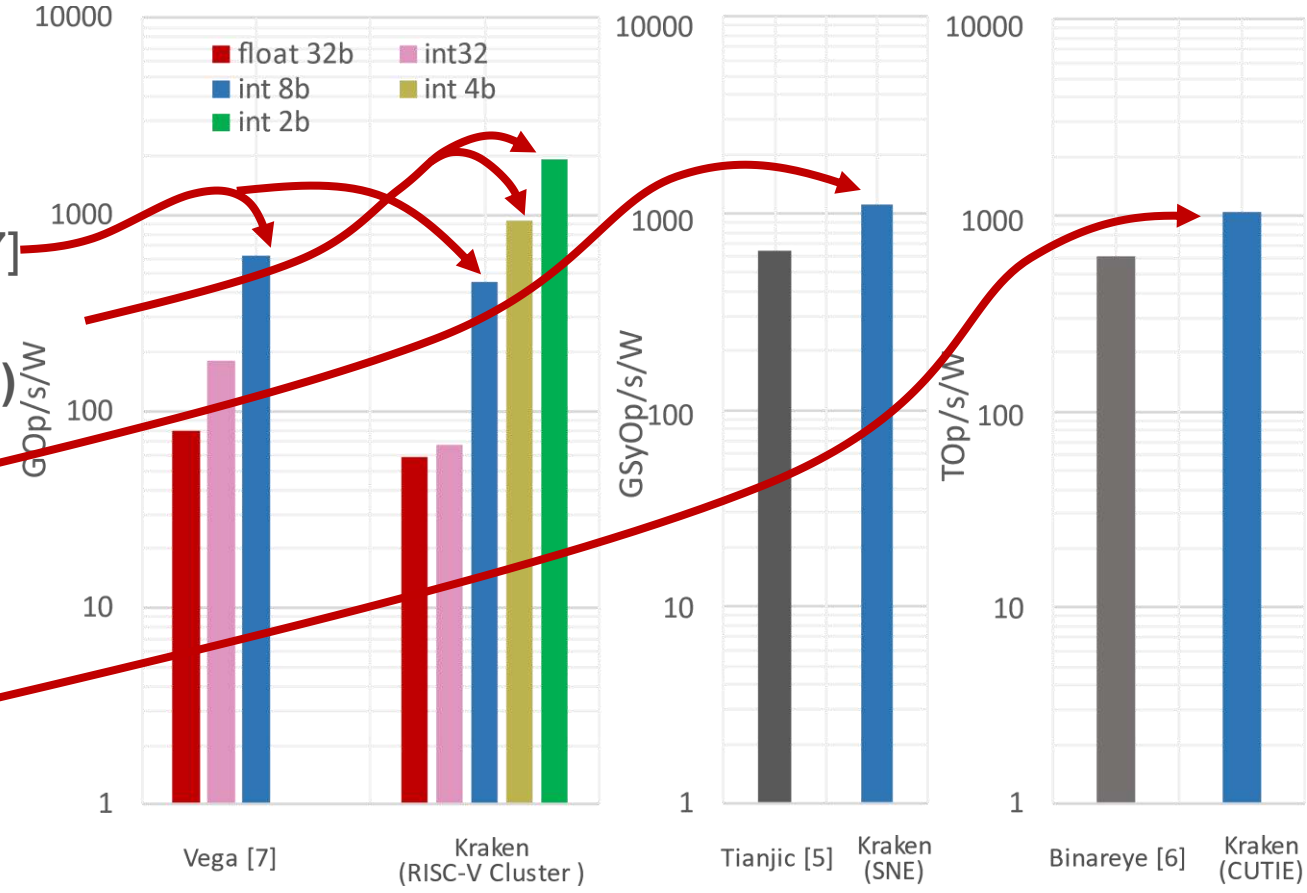

CUTIE, SNE can work concurrently for SNN + TNN "fused" inference (never done so far)

[5] L. Deng et al., "Tianjic," JSSC 2020
[6] B. Moons et al., "Binareye," CICC, 2018
[7] D. Rossi et al., "Vega," JSSC 2022.

Specialization in perspective

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

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

ISA-based 10-20x →**1-5pJ (8bit)** → **XPULPV2 &V3**

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

Highly specialized DP 10-20x →**1-5fJ (ternary)** → **SNE, CUTIE**

From Pre-competitive to commercial: PULP →GAP9

GAP9

- Cluster: 9 cores
- Fabric Ctrl: 1 core
- Hardware FPU
- L1: 128 KB
- L2: 1.5 MB
- Interface: CSI2

RISC-V®

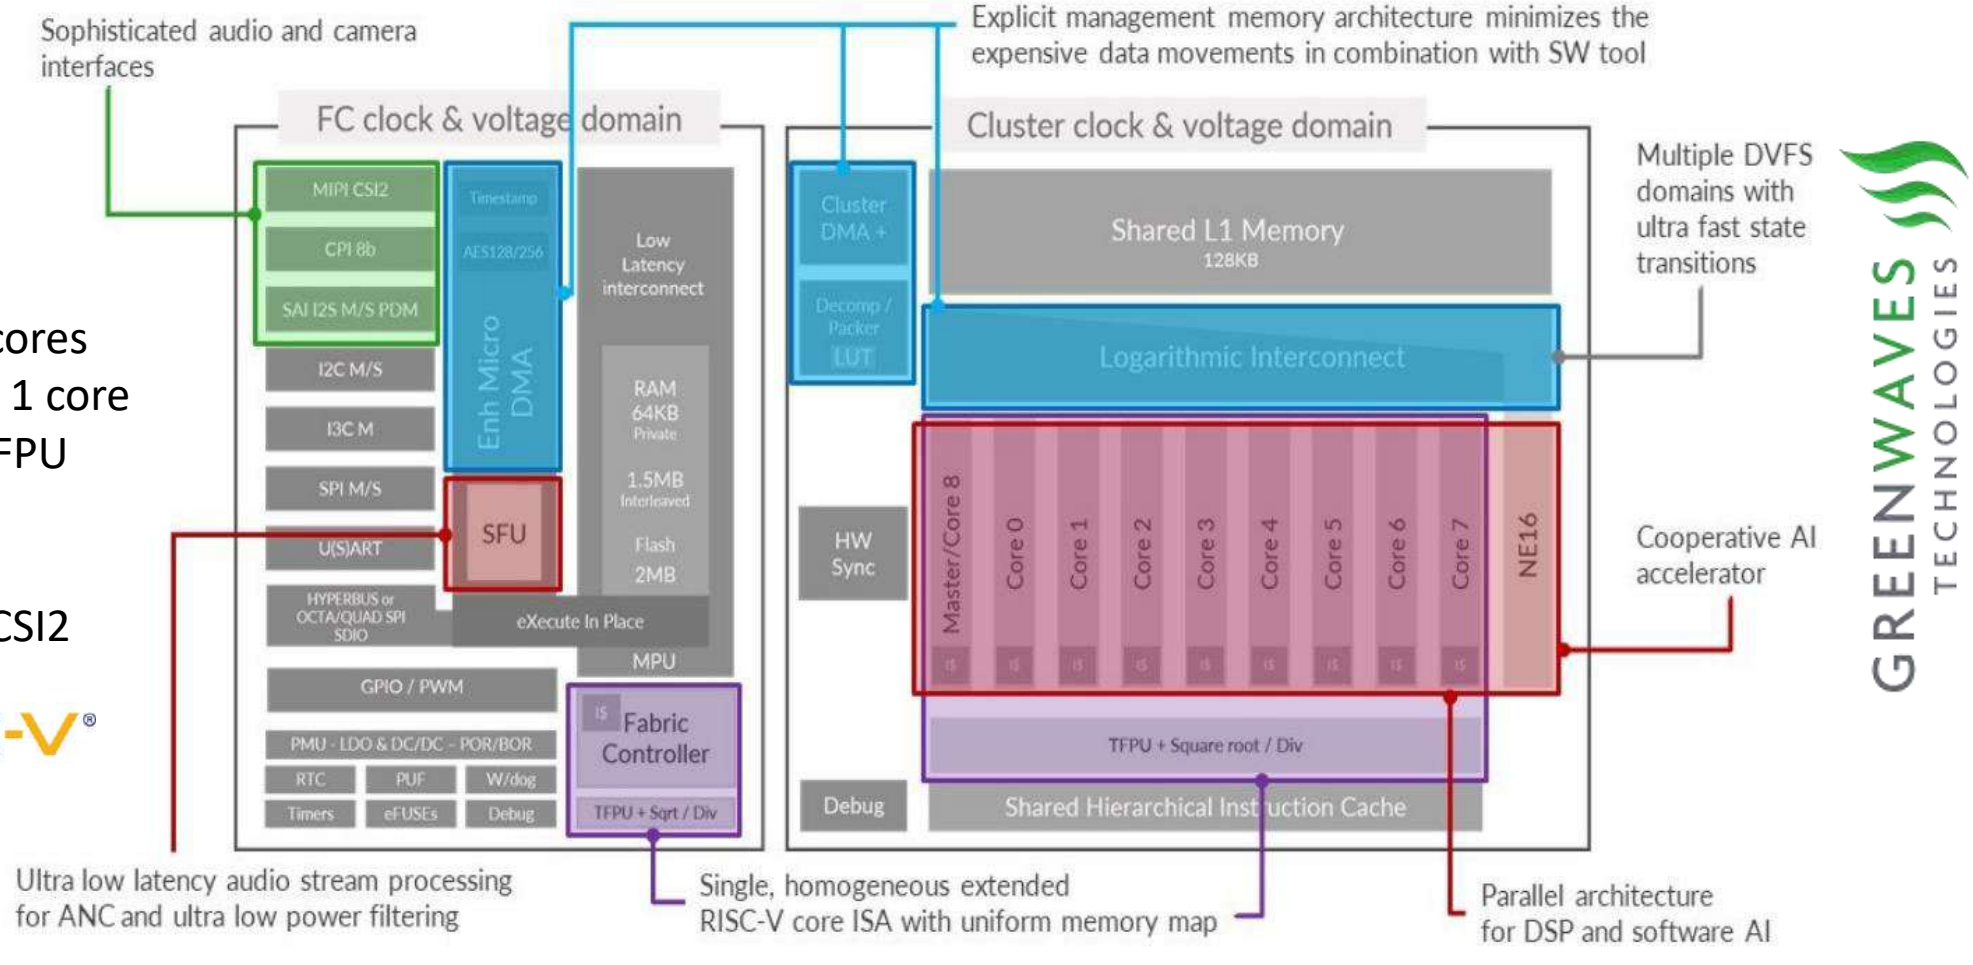

GREENWAVES TECHNOLOGIES

From Drones to Cars: Stepping up

- **Microcontroller class of devices**
  - Infineon AURIX Family MCUs
  - **Control tasks**, **low-power sensor acquisition & data processing**
    Features: lockstepped **32-b HP TriCore CPU** , HW I/O monitor, dedicated accelerators
- **Powerful real-time architectures**
  - ST Stellar G Series (based on ARM Cortex-R cores)
  - **Domain controllers and zone-oriented ECUs**
  - Features: HW-based virtualization, Multi-core **Cortex-R52** (+ NEON) cluster in split-lock, vast I/Os connectivity
- **Application class processors**
  - NXP i.MX 8 Family
  - **ADAS, Infotainment**
  - Features: Cortex-A53, **Cortex-A72,**HW Virtualization, **GPUs**

Carfield: Efficiency + Safety, Security, Predictability

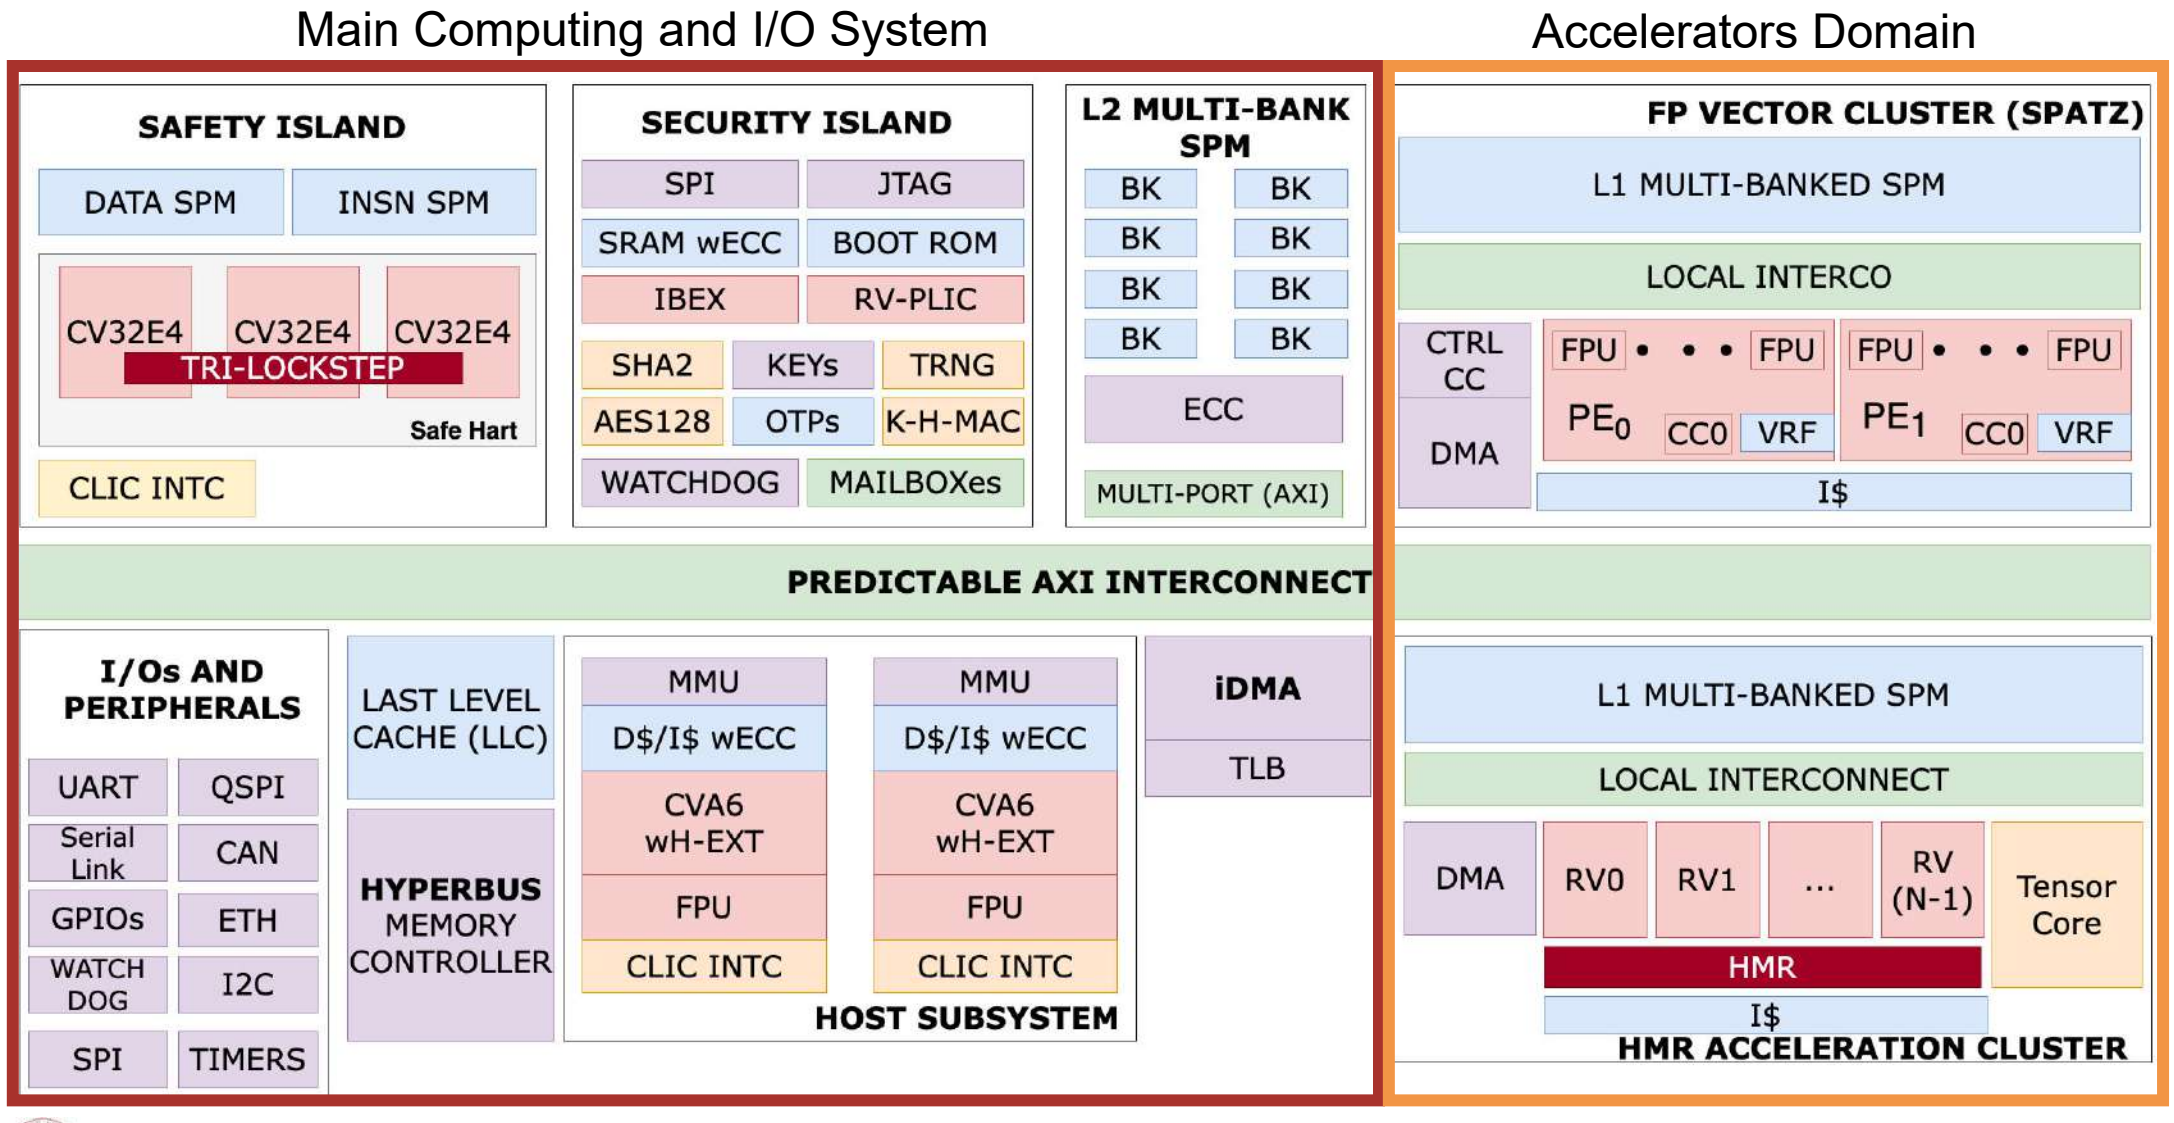

Precompetitive Partnership Buildup

ETH zürich

intel

BOSCH

life.augmented

ETH zürich

- Project Leaders
- Digital Systems Design, PULP, Open-Source, RISC-V
- Processors/Ips/Interconnects/Interrupts/HW Acceleration
- SW stack, compilers, runtime and optimized routines

- Real-Time (RT) Systems and On/Off-Chip RT Communication
- Safe/Secure Cyber-Physical Systems

- Virtualization-assisted systems, OS, Hypervisors, RISC-V
- Security of Cyber-Physical Systems

- Intel16 FinFet technology (for the first prototype)
- ASIC design support and packaging
- Supporters: STMicroelectronics, BOSCH

Moonshot: Toward Self-Driving Cars

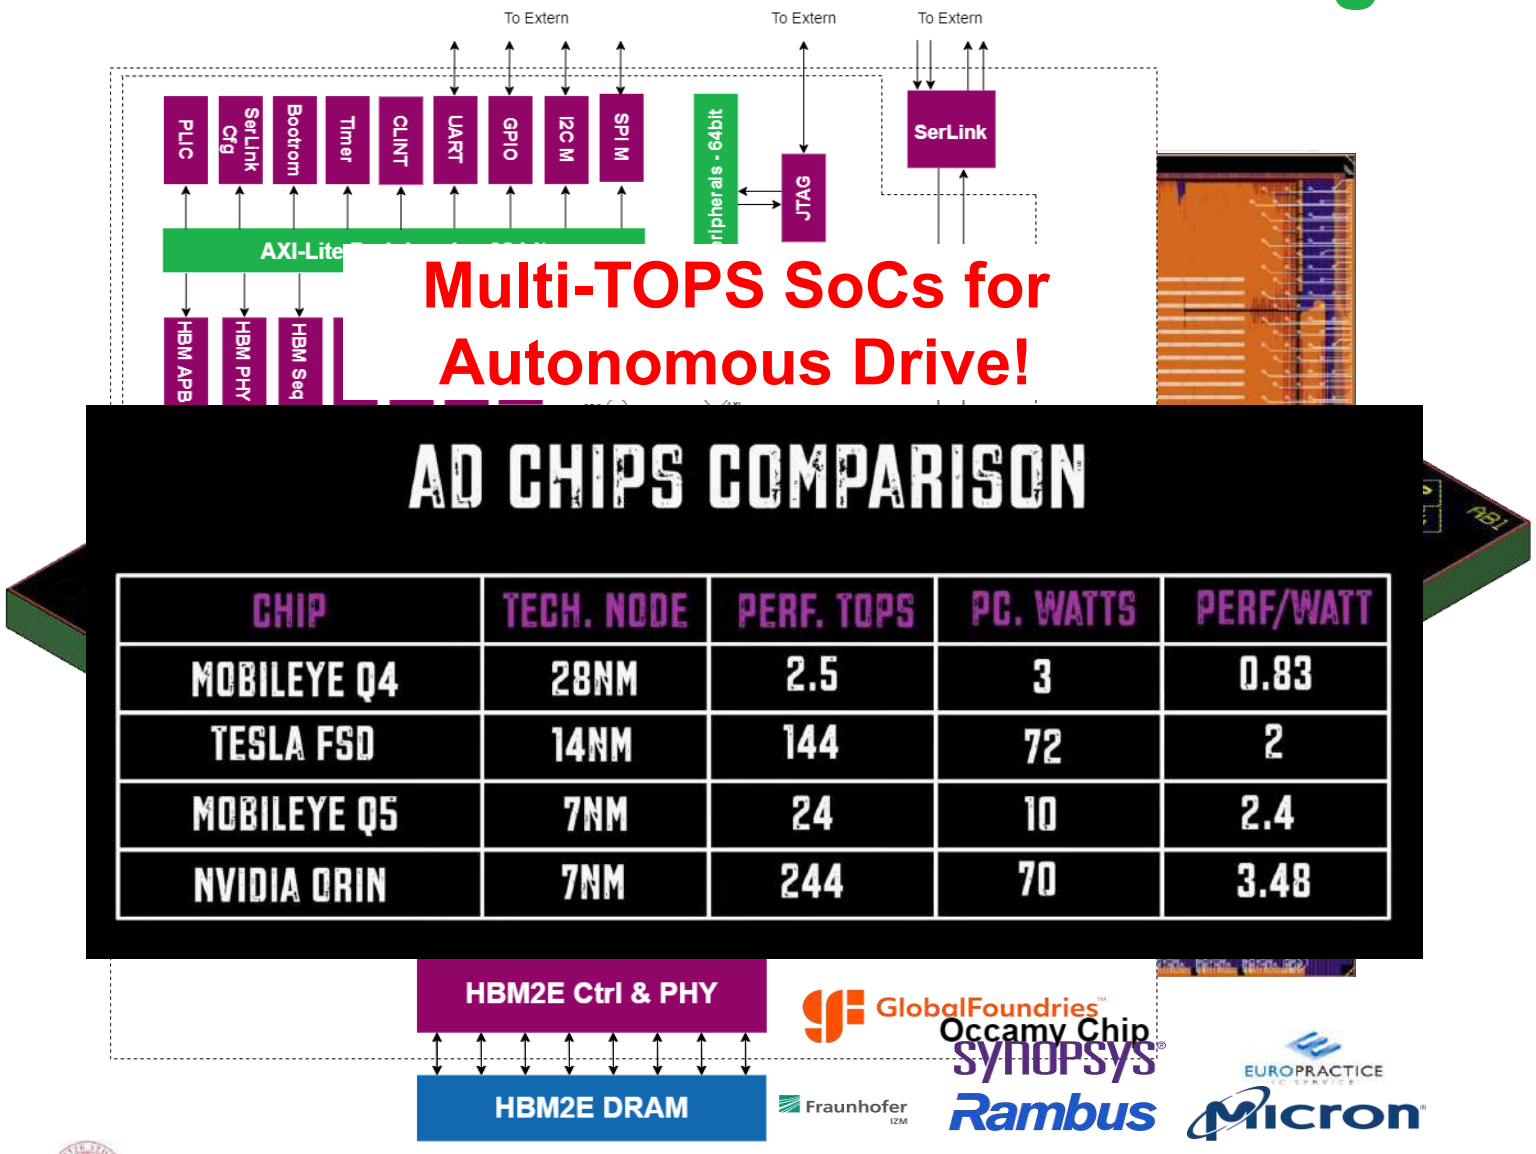

- GF12, target **1GHz** (typ)
- 2 AXI NoCs (multi-hierarchy)
  - 64-bit
  - 512-bit with “interleaved” mode
- Peripherals
- Linux-capable manager core CVA6
- 6 Quadrants: 216 cores/chiplet
  - 4 cluster / quadrant:
    - 8 compute +1 DMA core / cluster
    - 1 multi-format FPU / core (FP64,x2 32, x4 16/alt, x8 8/alt)
- 8-channel HBM2e (8GB) **512GB/s**
- D2D link (Wide, Narrow) **70+2GB/s**
- System-level DMA
- SPM (2MB wide, 512KB narrow)

Peak 384 GDPflop/s per chiplet

ETH zürich

Final Thoughts

- **Open Source Computing HW is Happening**
  - RISC-V Momentum is growing fast
  - Stronger Ecosystems emerge
  - RISC-V automotive products are in production (e.g. renesas)
- **Europe is NOT late**
  - Strong support from Public Authorities
  - Key EU OEMs (NXP, Infineon, STM) are active
  - Academic world is strongly committed
- **Nurture the industrial OSCHW ecosystem**
  - Automotive electronics is a key priority area
  - Avoid fragmentation
  - Create a stronger vertical ecosystem in the automotive value chain

https://digital-strategy.ec.europa.eu/en/library/recommendations-and-roadmap-european-sovereignty-open-source-hardware-software-and-risc-v

PULP

Parallel Ultra Low Power

Luca Benini, Alessandro Capotondi, Alessandro Ottaviano, Alessio Burrello, Alfio Di Mauro, Andrea Borghesi, Andrea Cossettini, Andreas Kurth, Angelo Garofalo, Antonio Pullini, Arpan Prasad, Bjoern Forsberg, Corrado Bonfanti, Cristian Cioflan, Daniele Palossi, Davide Rossi, Fabio Montagna, Florian Glaser, Florian Zaruba, Francesco Conti, Georg Rutishauser, Germain Haugou, Gianna Paulin, Giuseppe Tagliavini, Hanna Müller, Luca Bertaccini, Luca Valente, Manuel Eggimann, Manuele Rusci, Marco Guermandi, Matheus Cavalcante, Matteo Perotti, Matteo Spallanzani, Michael Rogenmoser, Moritz Scherer, Moritz Schneider, Nazareno Bruschi, Nils Wistoff, Pasquale Davide Schiavone, Paul Scheffler, Philipp Mayer, Robert Balas, Samuel Riedel, Segio Mazzola, Sergei Vostrikov, Simone Benatti, Stefan Mach, Thomas Benz, Thorir Ingolfsson, Tim Fischer, Victor Javier Kartsch Morinigo, Vlad Niculescu, Xiaying Wang, Yichao Zhang, Frank K. Gürkaynak,
all our past collaborators and many more that we forgot to mention

http://pulp-platform.org @pulp_platform