


Carfield: The Open-Research Platform for Safety, Resilient and Predictable Systems

Integrated Systems Laboratory (ETH Zürich)

University of Bologna

Angelo Garofalo agarofalo@iis.ee.ethz.ch
& the Carfield team



PULP Platform

Open Source Hardware, the way it should be!

[@pulp_platform](https://twitter.com/pulp_platform) 


pulp-platform.org 

youtube.com/pulp_platform 


Why An Automotive-Oriented PULP System?




TODAY Distributed Architecture



TOMORROW Domain Architecture



FUTURE Zone Architecture



- ECUs are moving towards domain, zone architectures (more than simple MCUs)
- Systems require higher **real-time** processing **performance**, enhanced **safety & security** features
→ Interesting challenges to be addressed (...and not only limited to Automotive world 😊)



Automotive Trends: Platforms on the Market



- Microcontroller class of devices
 - Infineon AURIX Family MCUs, TC32x (Low-Power) – TC39x (ADAS)
 - **Control tasks, motor control, sensor acquisition and data processing** (including Radar processing)
 - Features: 1-to-6 **32-b HP TriCore CPU** , lockstepped cores, 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**, Cortex-M4, HW Virtualization, **GPUs**



Goals of The Project

- Develop a pre-competitive Automotive SoC
 - Scalable and configurable architectural template
 - SW stack to address requirements of RISC-V based automotive applications
- Collaborative research roadmap for automotive-driven computing architectures
 - Functional safety
 - Hardware/Software Acceleration for critical kernels
 - Real-Time and Predictable Systems
 - Fast-Interrupts and Virtualizations
- Open-Source
 - Academic: • Reduce “getting up to speed” overhead for partners; • Enables fair and well controlled benchmarking
 - Industry: • Reduces the NRE, • faster innovation path, • helps exchange of information across NDA walls, • great for Marketing & Training



We Started Joining Forces



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ETH zürich



Universidade do Minho

intel®

- Project's 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



Architectural Pillars of The Project



- We start from our PULP architectural ball-park
- Host architecture
 - 64-bit CVA6 processors → non-critical applications (GPOs, e.g. Linux)
 - 32-bit CV32E lock-stepped cores → safety-critical applications (RTOSs)
 - Pluggable Hardware Root-of-Trust (RoT) based on OpenTitan (lowRISC open-source proj)
 - Wide set of peripherals: SPI, I2C, I2S, ETHERNET (RMII), CAN, HyperRAM..
- Pluggable Heterogeneous Accelerators
 - Heterogeneous parallel programmable clusters [Spatz, Snitch/PULP clusters, MemPool..]
 - ISA extensions (FP, INT, SIMD, Vector) and custom data-paths (TensorCores, DNN,..)
- Flexible Programming model
 - Linux + OpenMP for non-critical code
 - RTOS (FreeRTOS, AUTOSAR, Erika) + bare-metal task offloading for safety-critical code

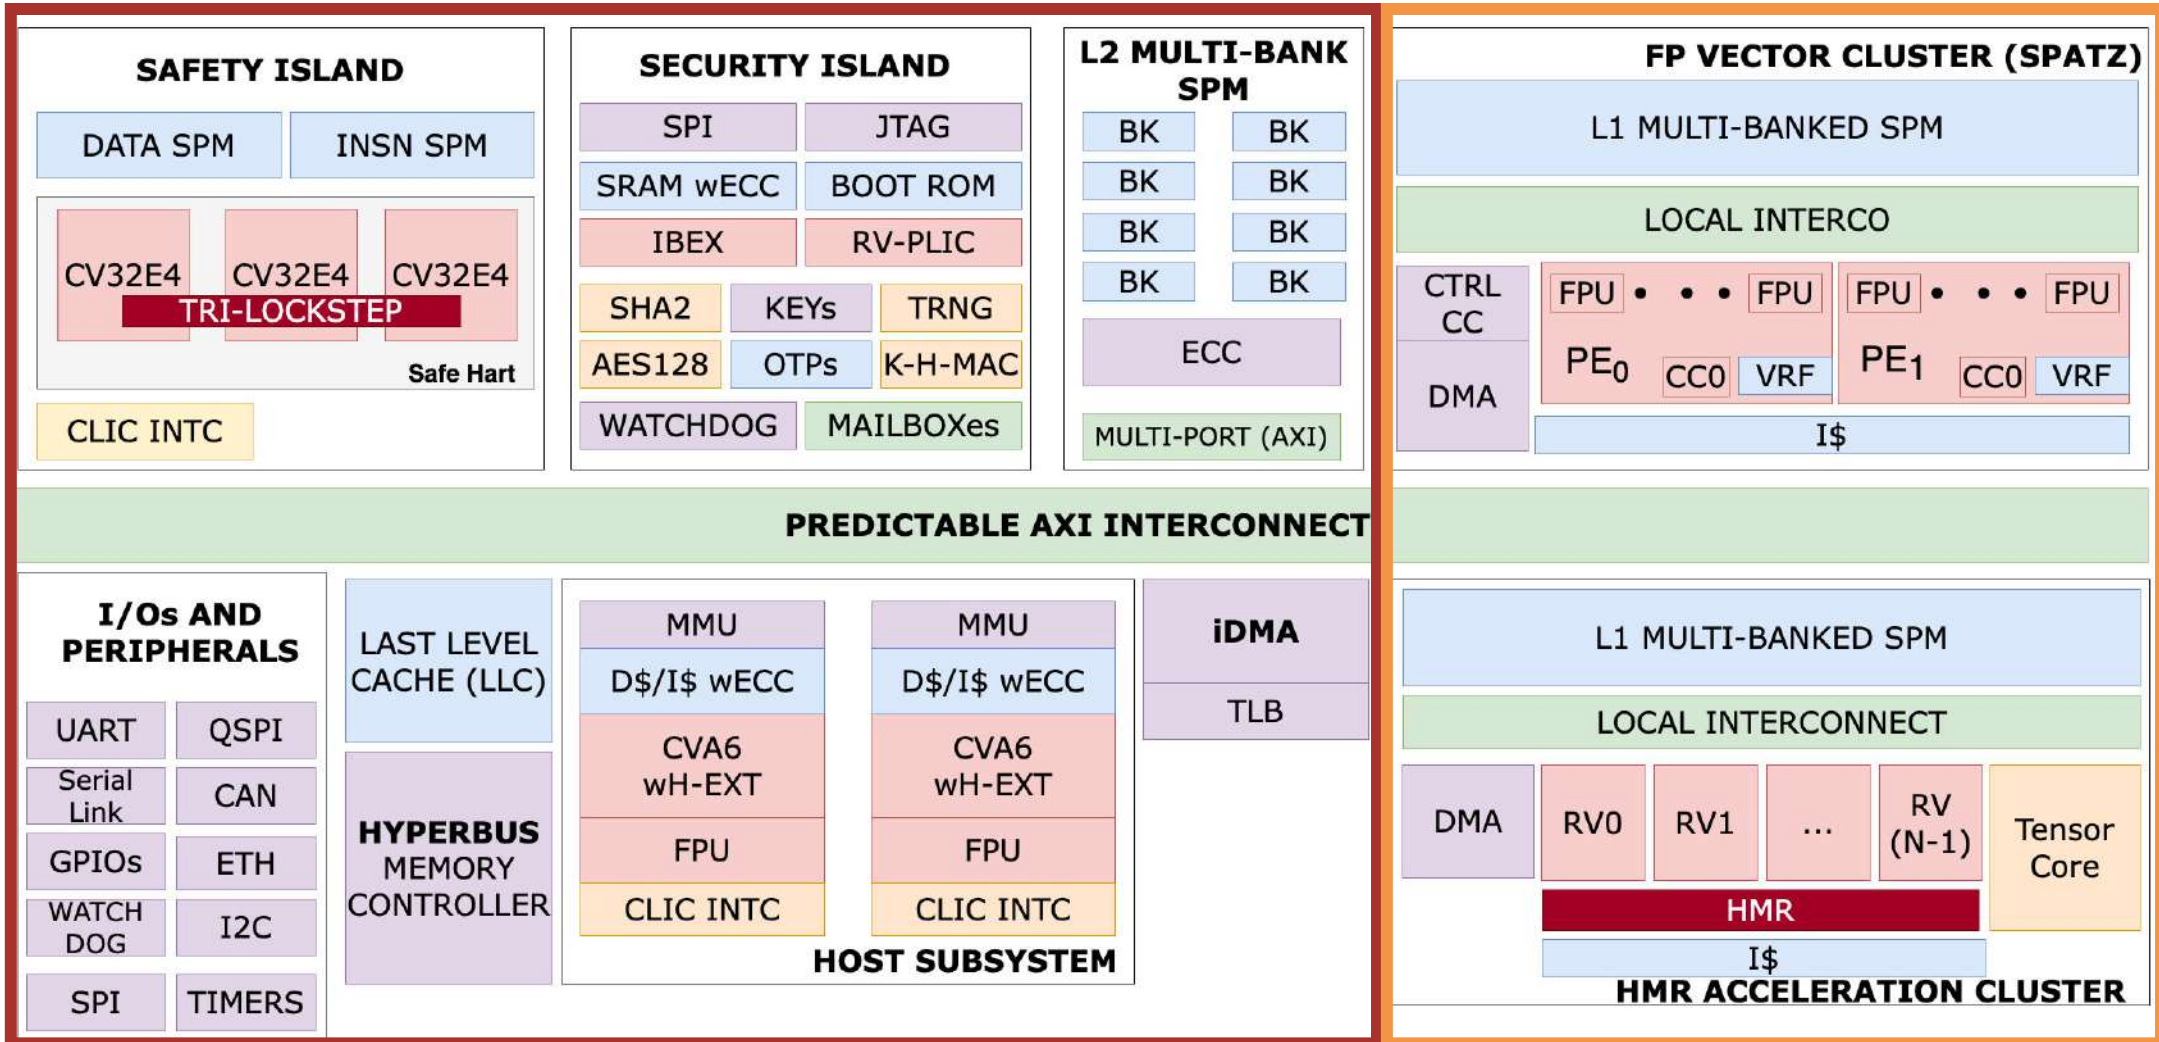


Carfield: Architectural Template Based on Fully Open IPs

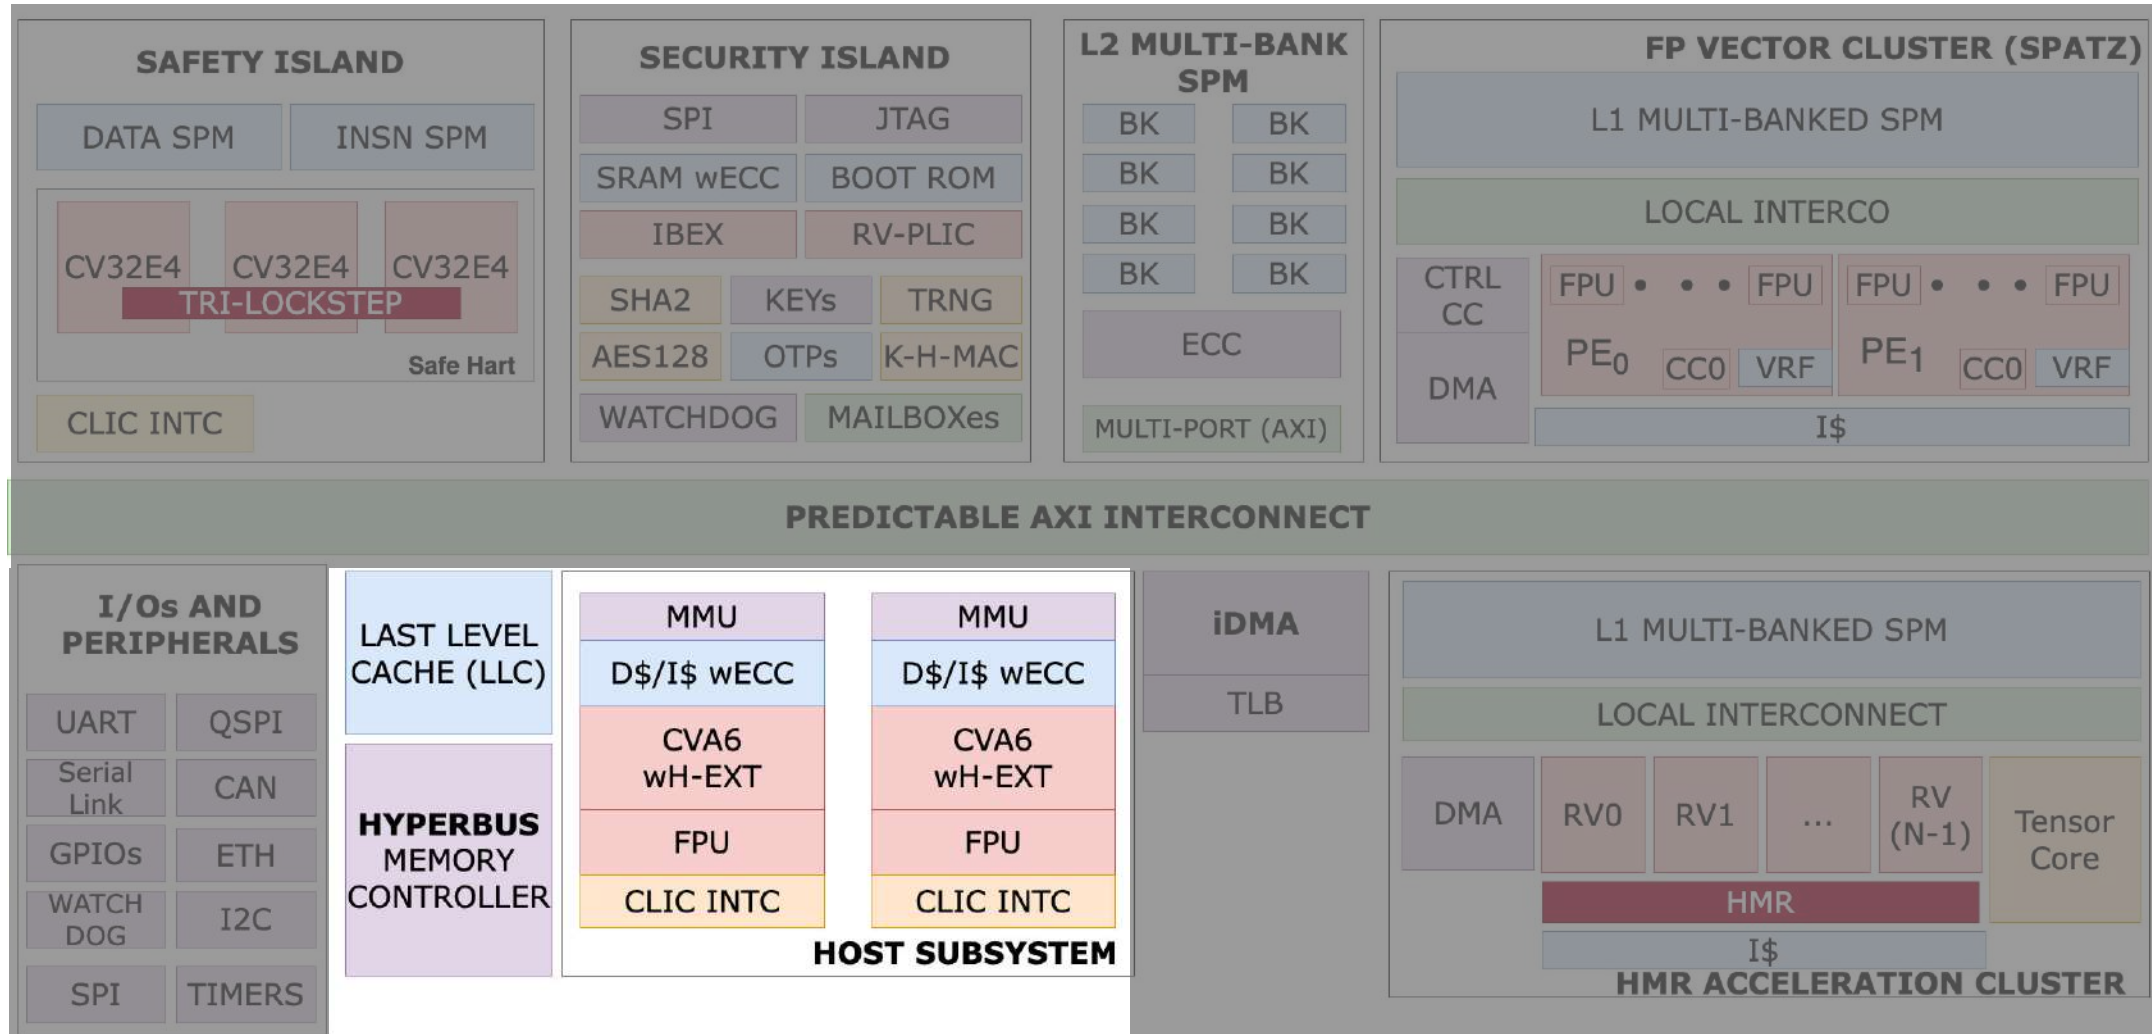


Main Computing and I/O System

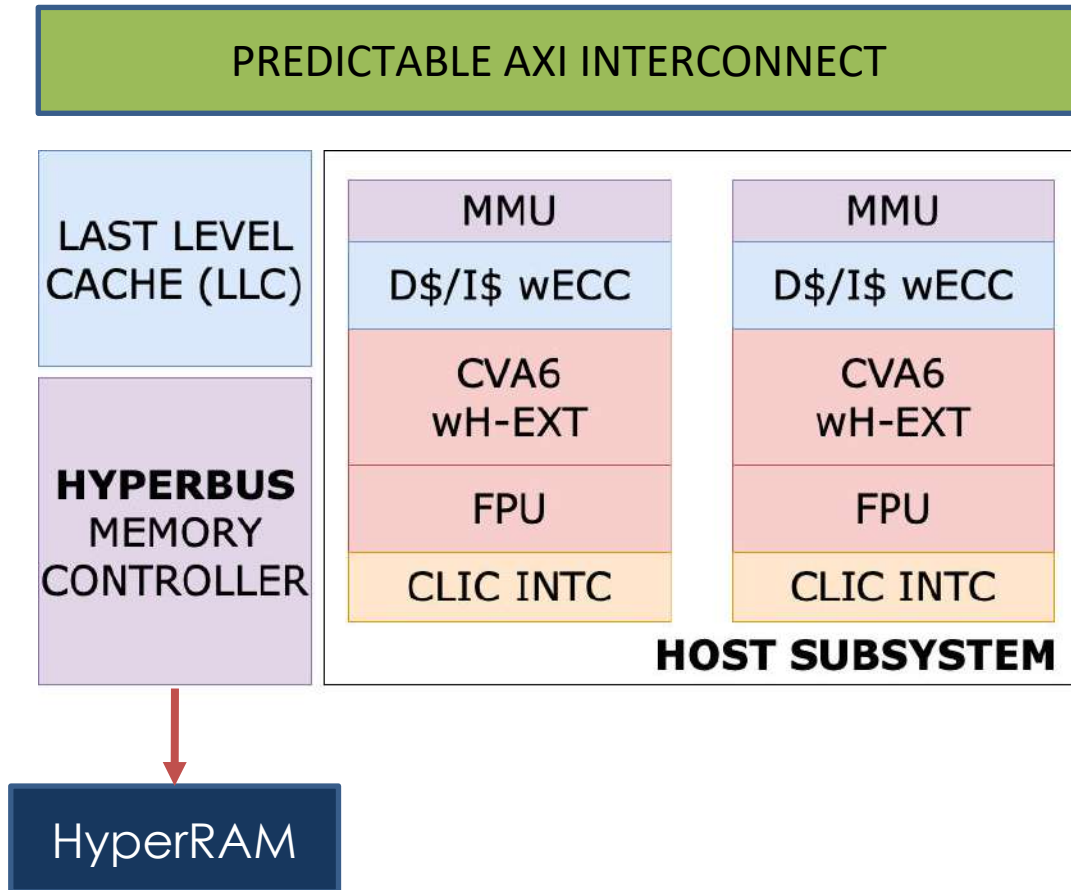
Accelerators Domain



Host-Domain for Non-Critical Linux-Based Applications



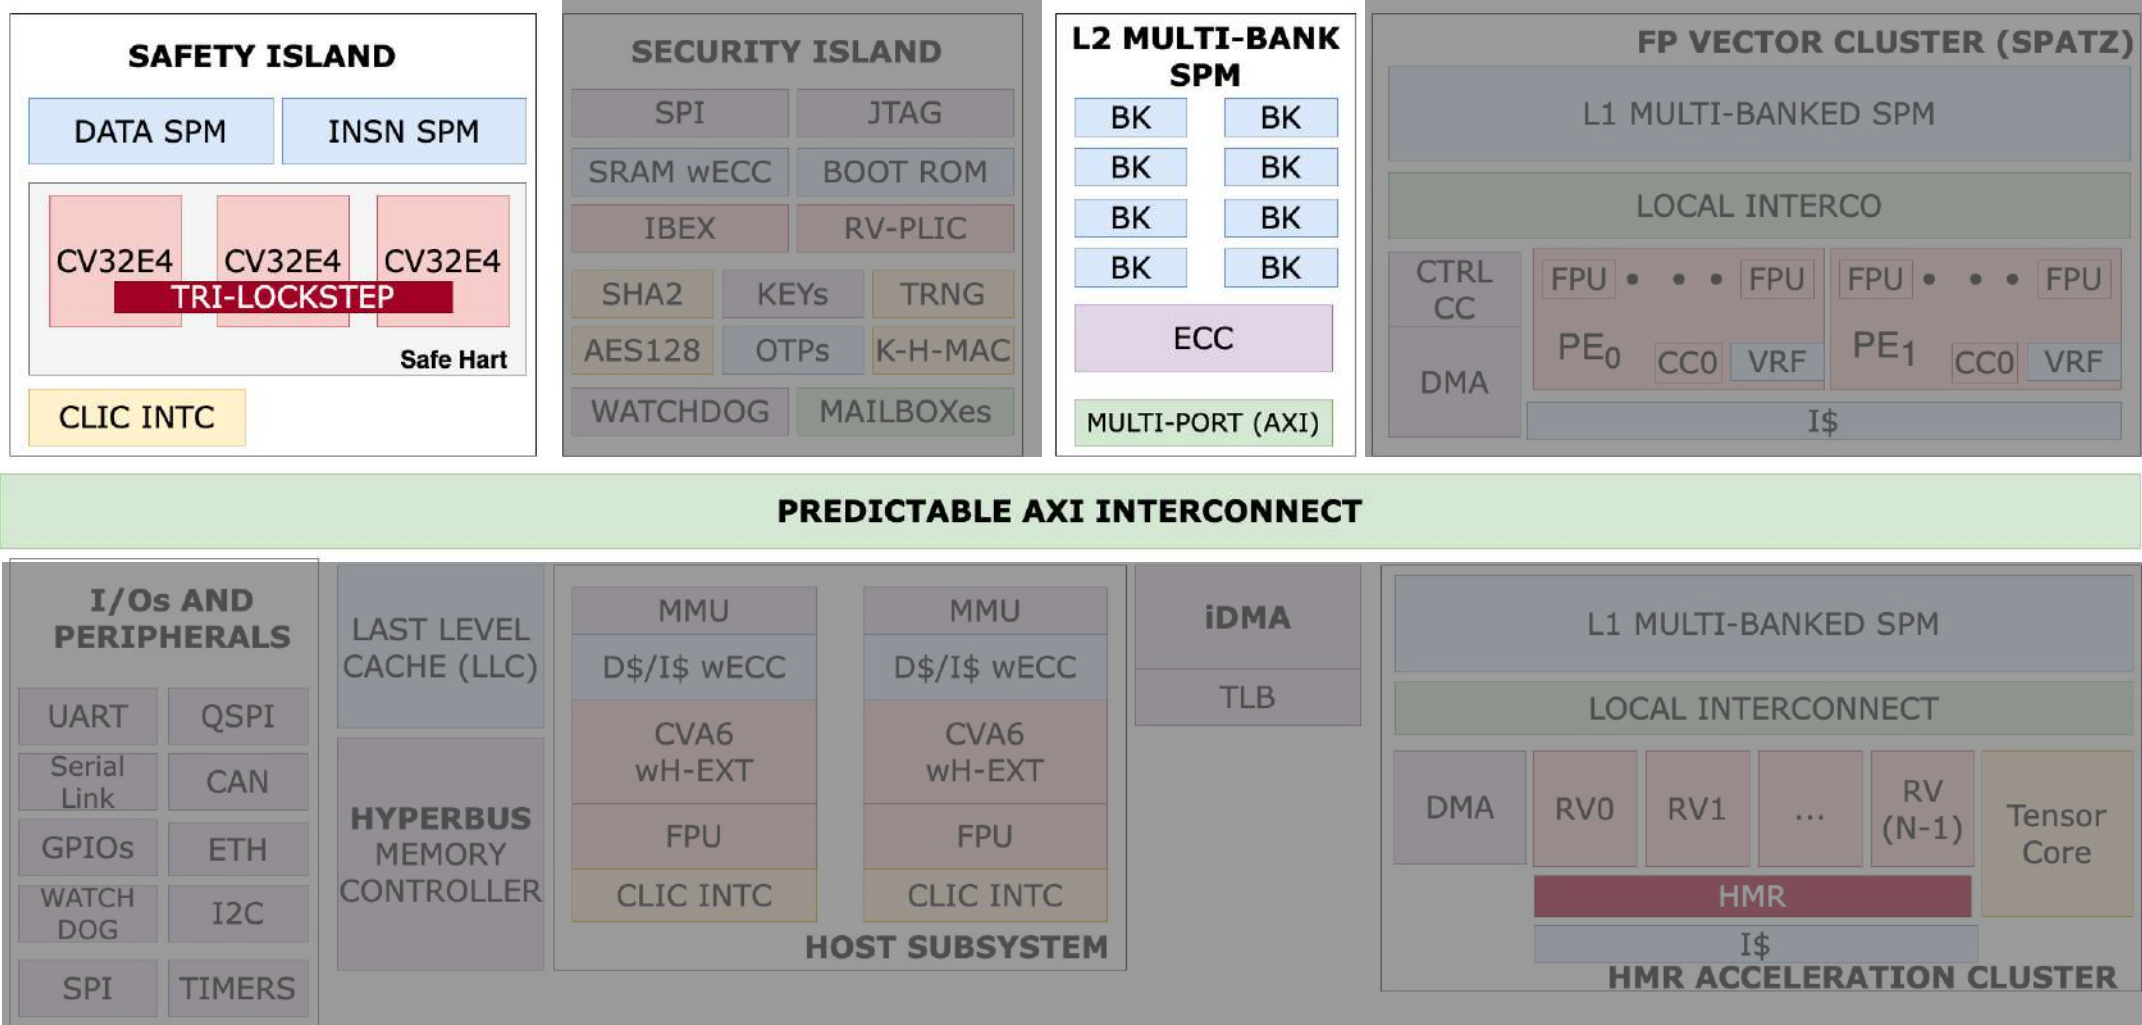
Host-Domain for Non-Critical Linux-Based Applications



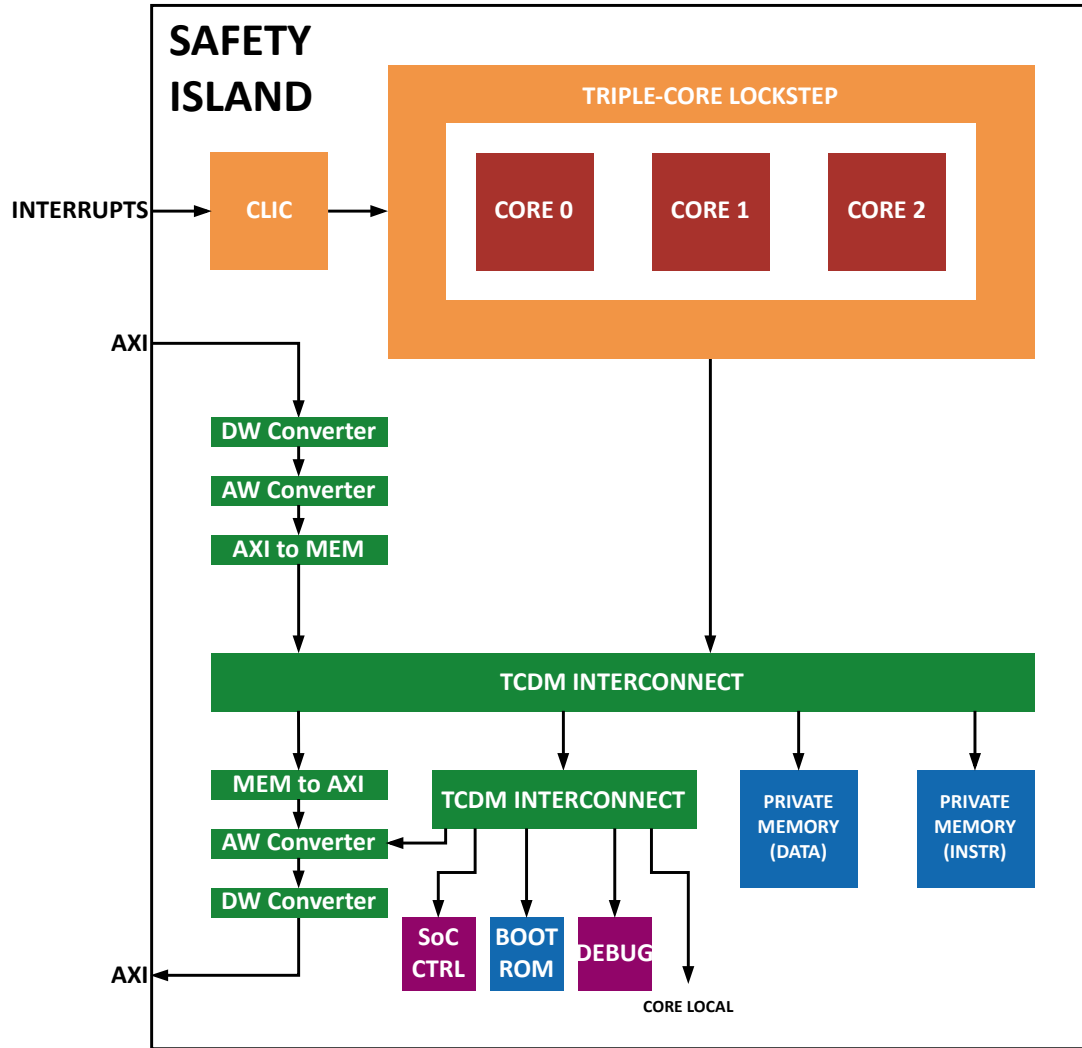
- Coherent Cluster based on a Custom Interconnect targeting 2 to 4 CVA6 processors
- CVA6: Application mid-end processor
 - Linux-capable processor
 - Supports 48-bit virtual memory MMU (Sv48)
 - M, HS (Hypervisor-extended Supervisor) and U privilege modes;
 - Tightly integrated D\$ and I\$;
 - (22 FDX) Freq.: 1+ GHz; Area: 1.5 MGE
- Shared Last Level [Data] Cache (LLC)
- HyperBUS Controller for off-chip HyperRAM access



How Do We Handle Safety-Critical and Real-Time Tasks?



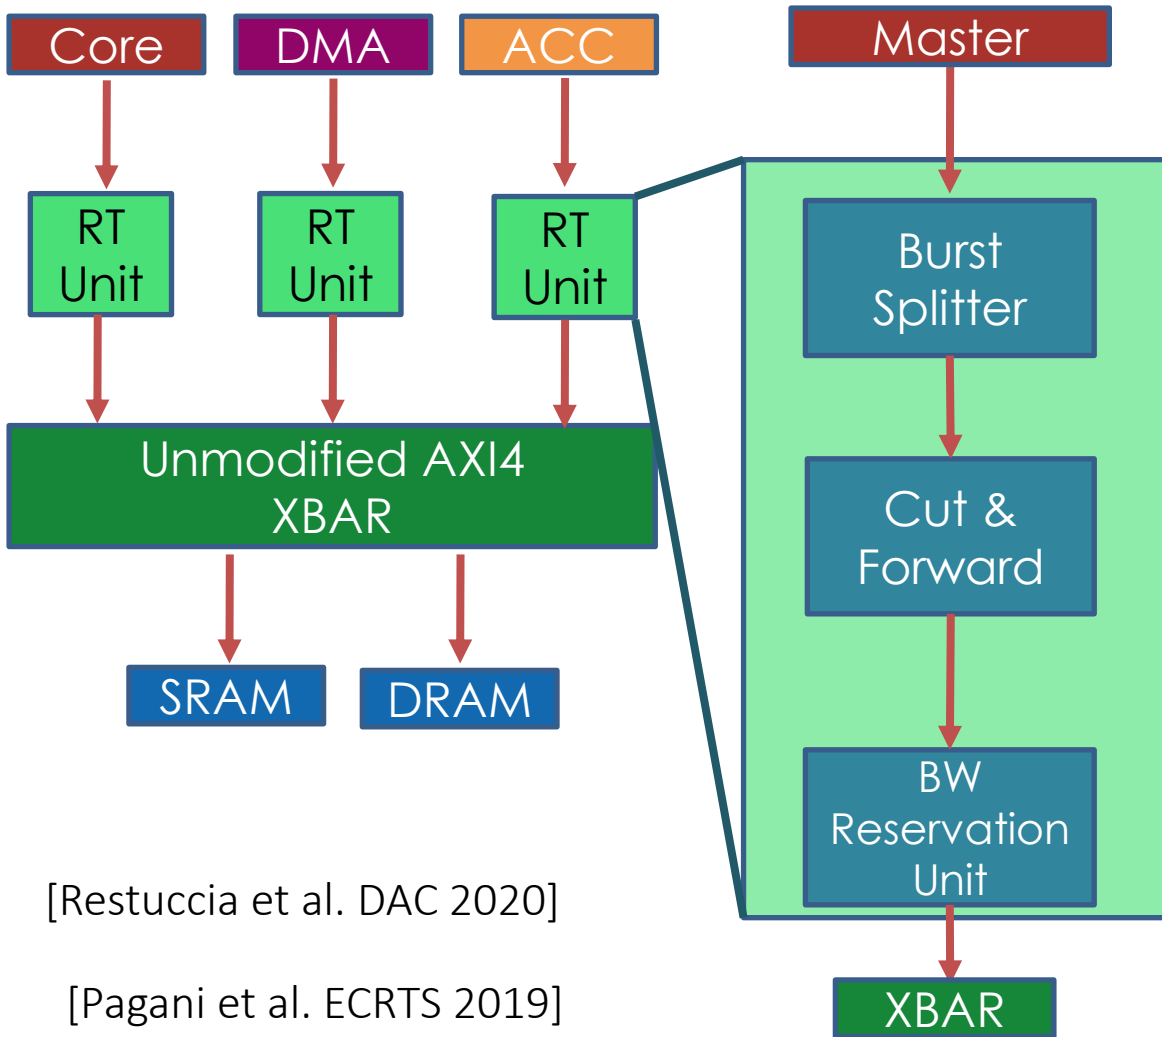
The Safety Island



- Safety-critical applications running on top of a RTOS
- Three CV32E40 cores in lockstep
- ECC protected scratchpad memories for instructions and data
- CLIC controller for fast and configurable interrupt handling
- AXI-4 port to communicate with system



Predictable On-Chip Communication (AXI RT)



[Restuccia et al. DAC 2020]

[Pagani et al. ECRTS 2019]


- Minimal Intrusive Solution
 - No huge buffering, limited additional logic
- **AXI Burst Splitter**
 - Equalizes length of transactions to avoid unfair BW distribution
- **AXI Cut & Forward**
 - Configurable chunking unit to avoid long transactions influencing access time to the XBAR
- **AXI Bandwidth Reservation Unit**
 - Predictably enforces a given max nr of transactions per time period (to each master)
 - Per-address-range credit-based mechanism
 - Periodically refreshed (or by user)
- Solution verified in systematic worst-case Real Time Analysis




Contention-Free Shared L2 Scratchpad Memory



1. Two Address Mapping Modes

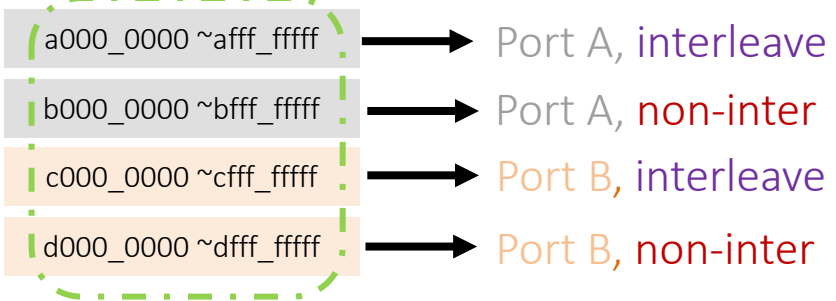


Non-interleaved



Interleaved

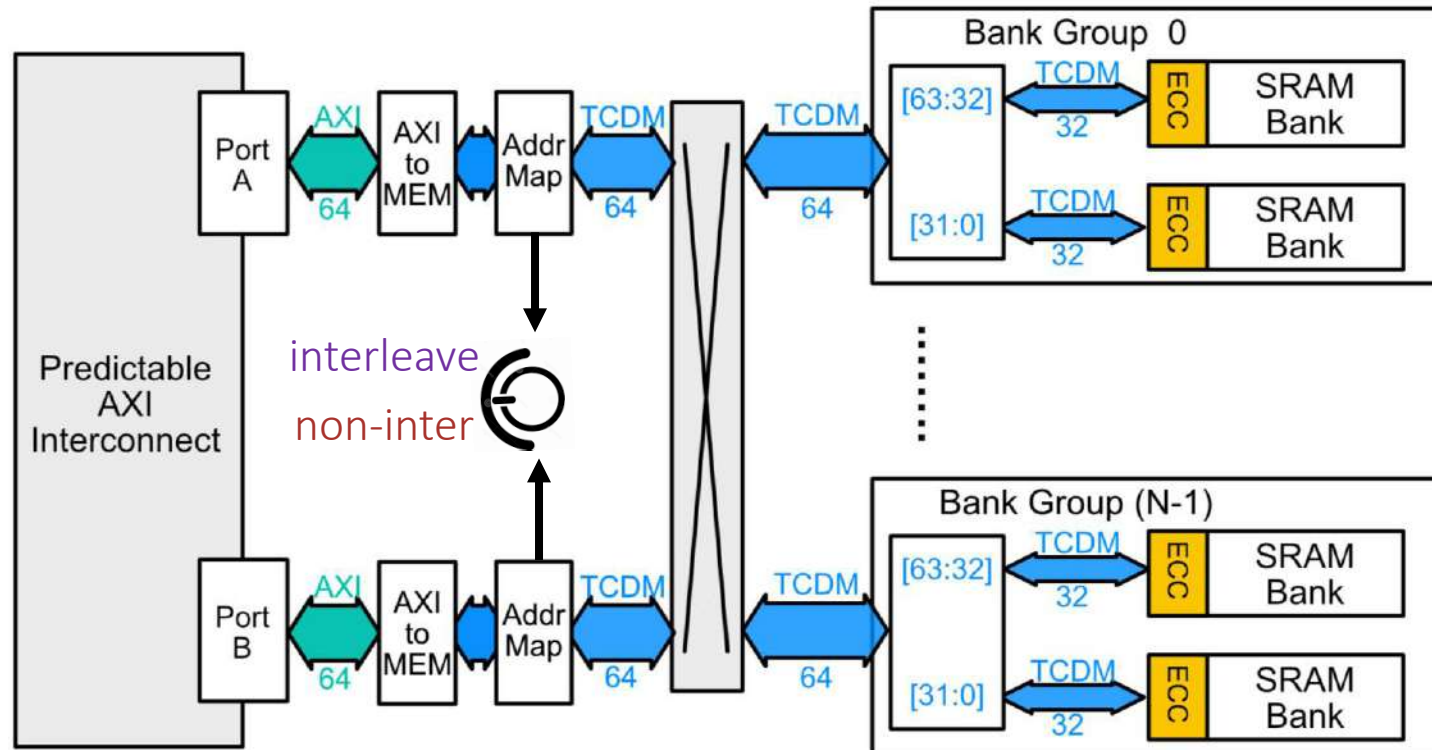
2. Dynamic Address Mapping by Address spaces, eg:



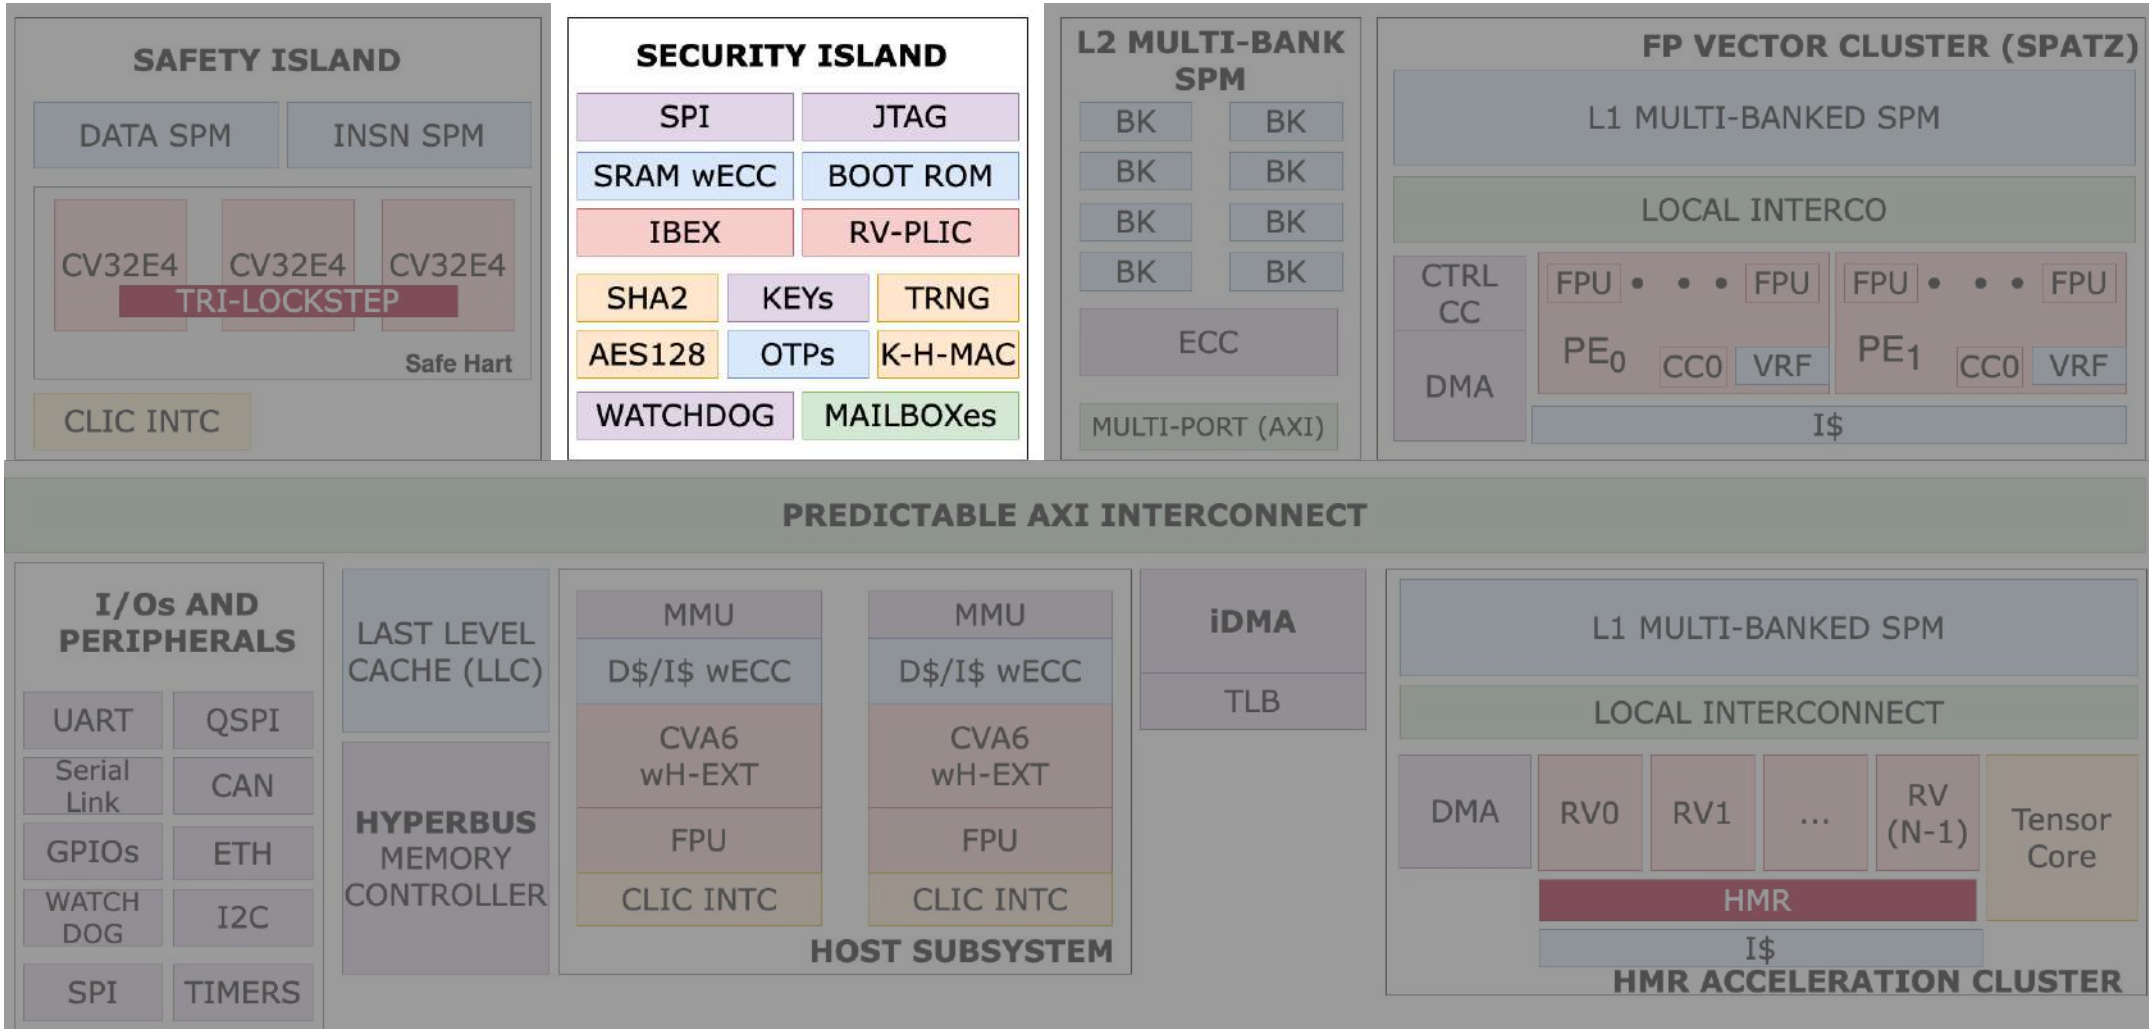
Point to the same L2 Mem space

3. L2 Mem diagram

- Currently we set 2 bank groups, namely 4 banks
- Software-side determine which port and which mode to use
 - By using different address space

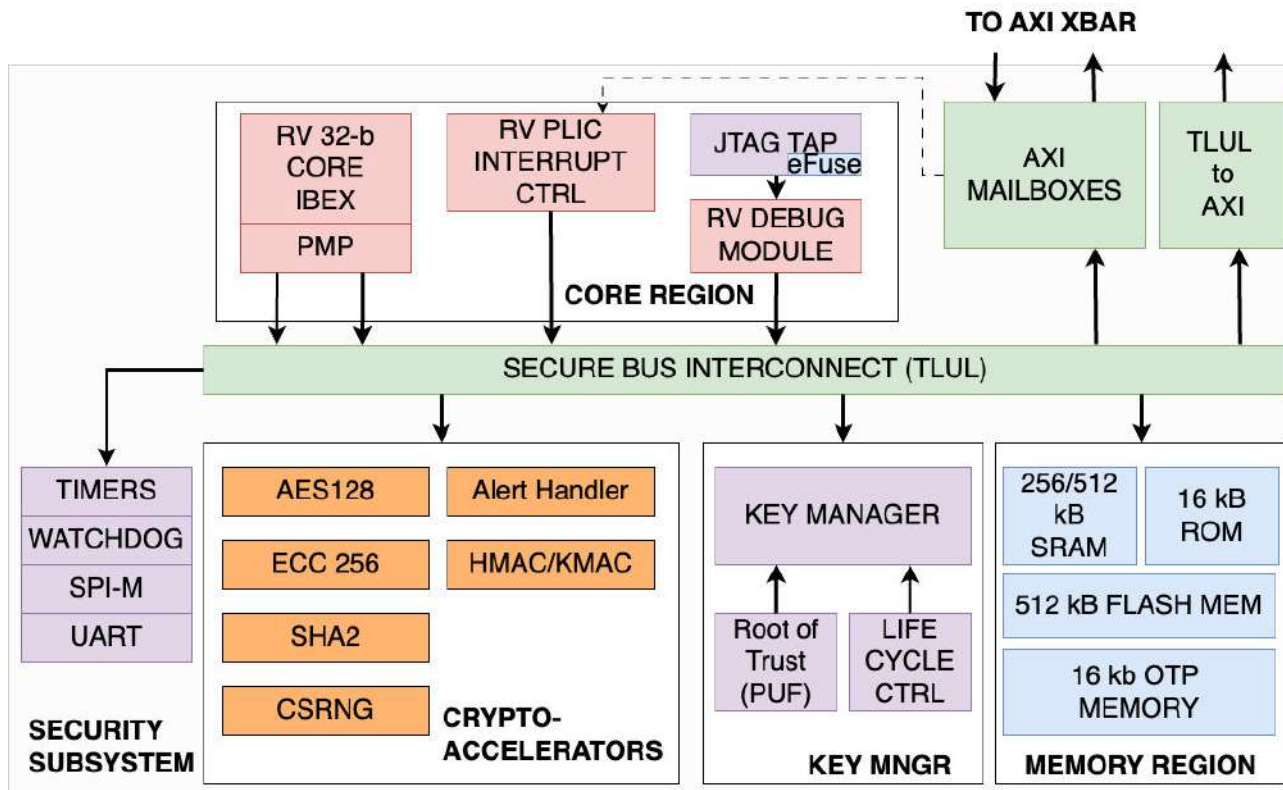


What About Security and Data Encryption/Decryption?



The Security Island

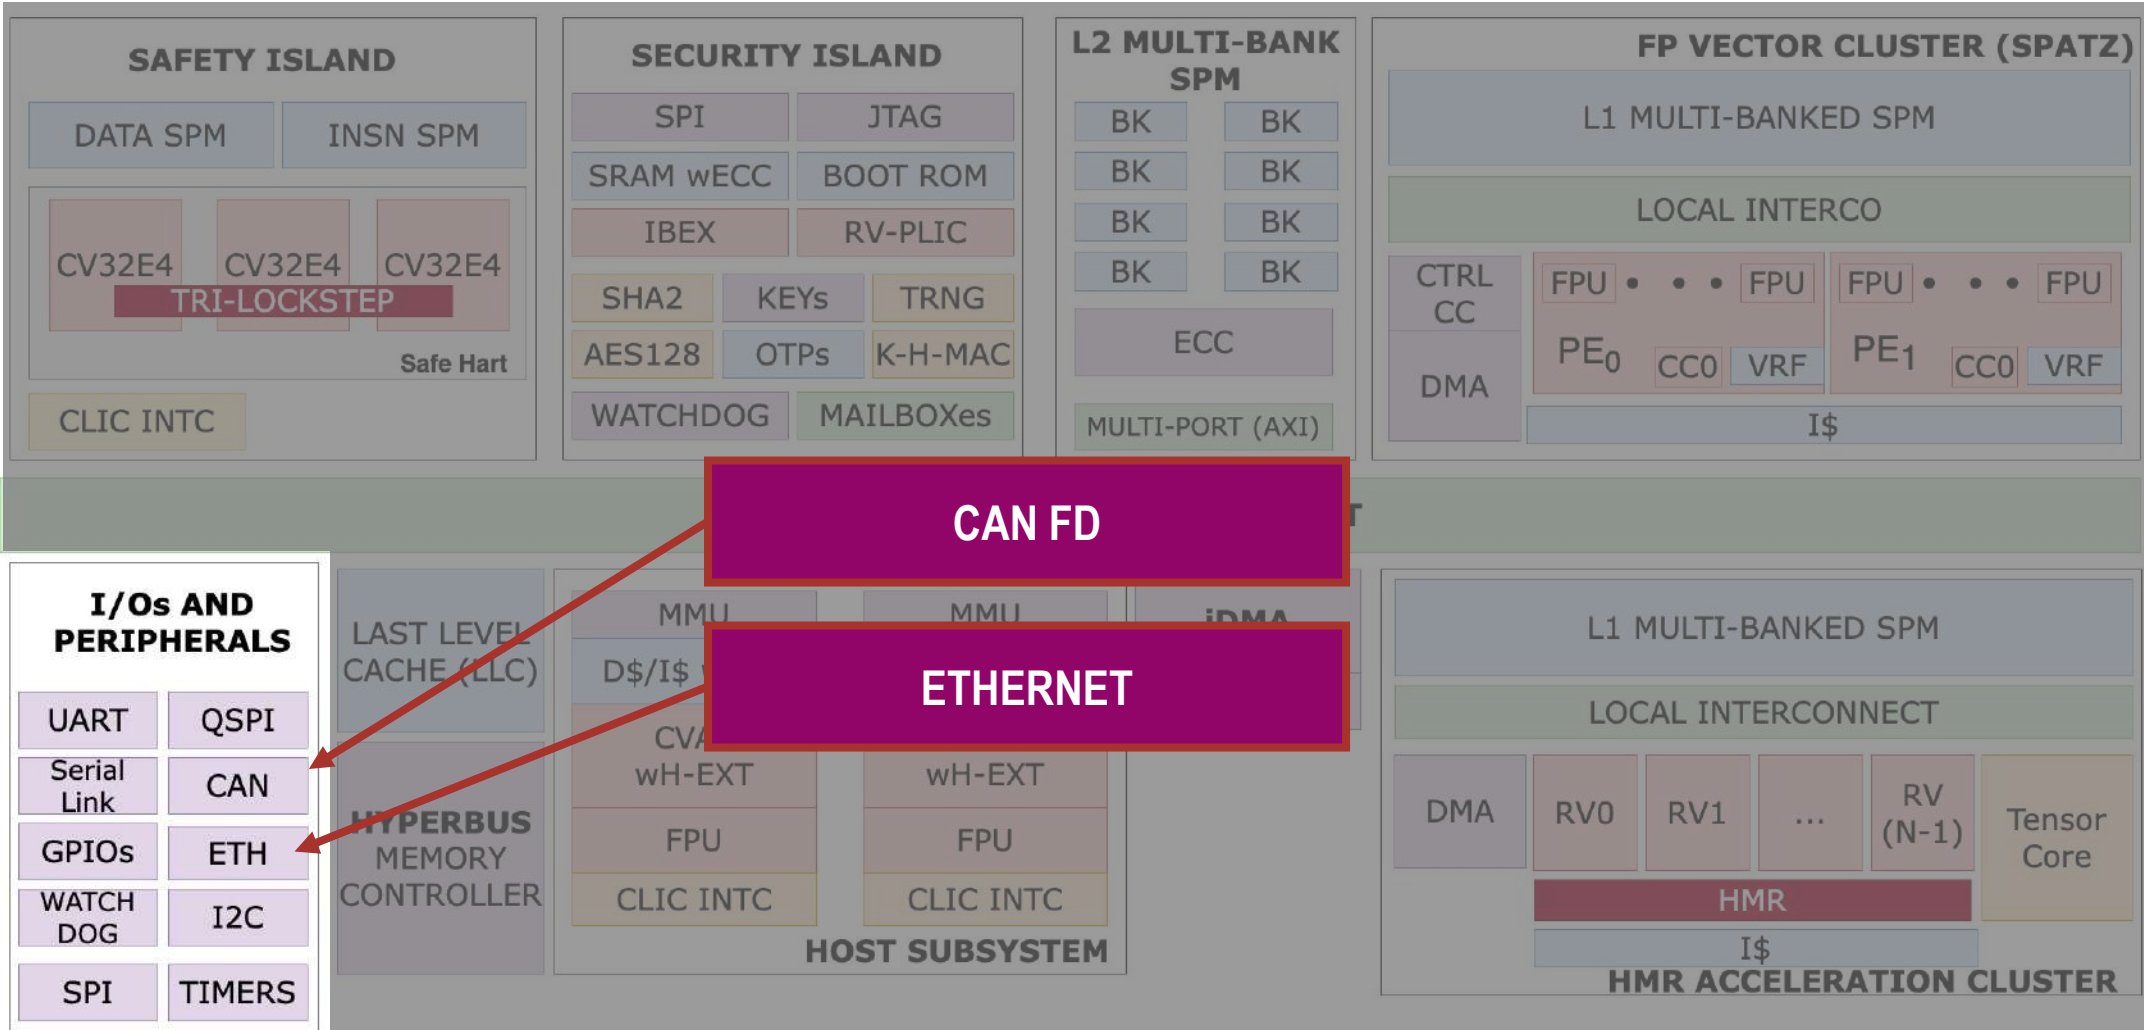
Derived from the OpenTitan project by lowRISC



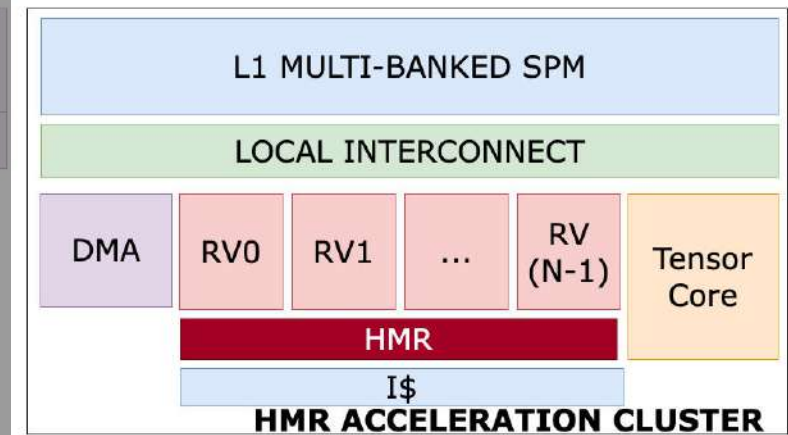
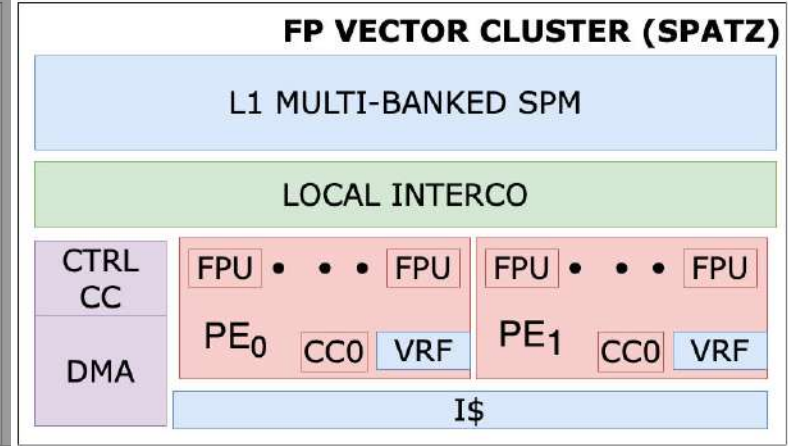
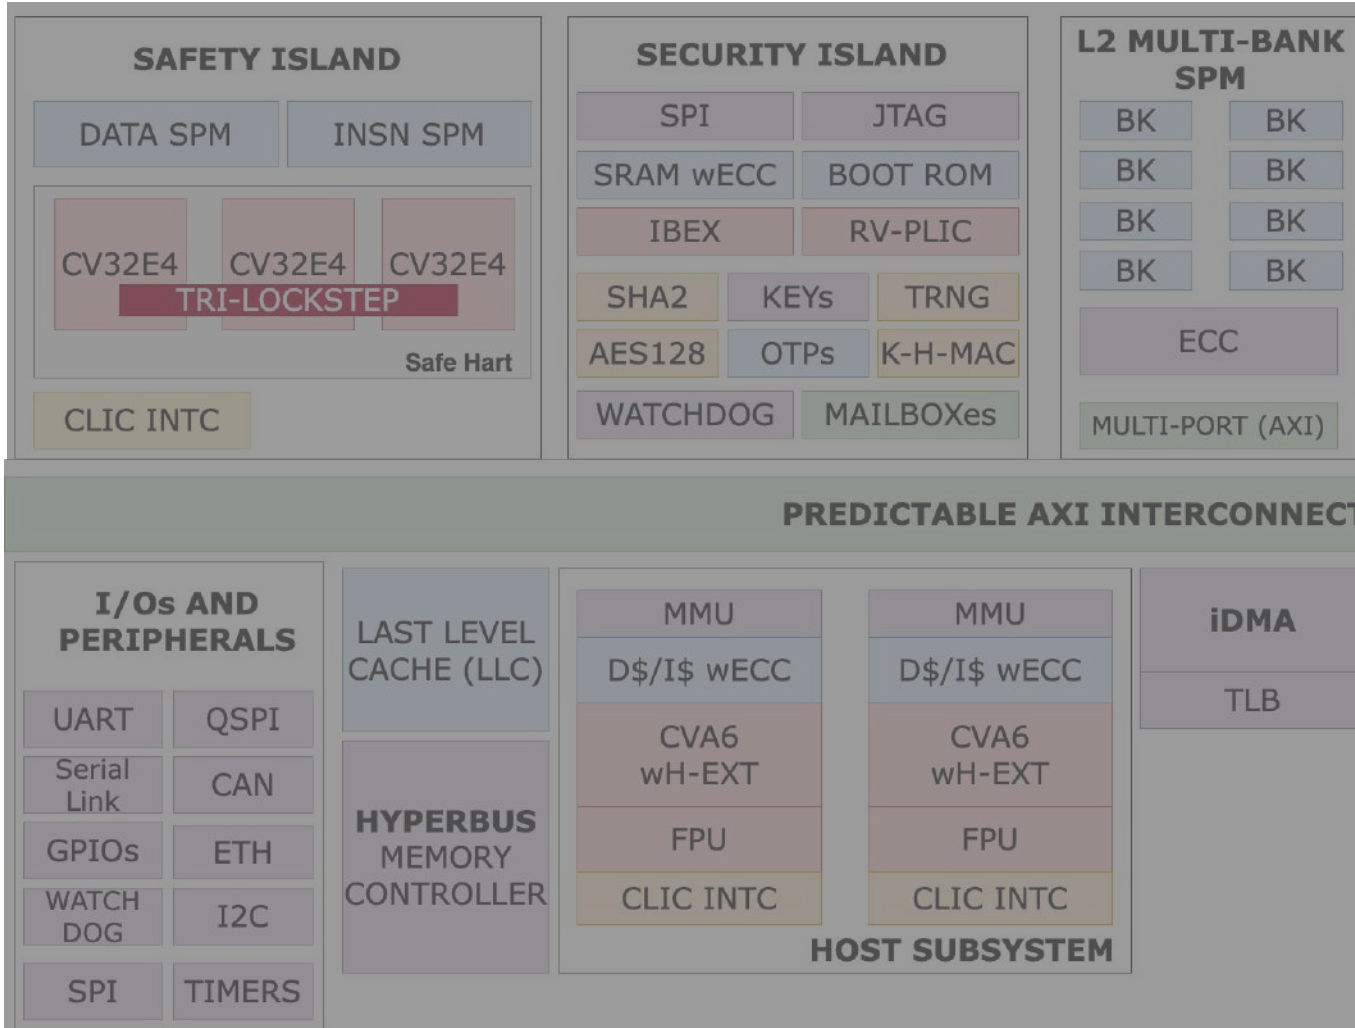
- Root of Trust
 - Stores Cryptographic Secrets
- Early Secure Boot Stages
 - Verify Cryptographic Signatures and Measurements before unlocking next boot stages
- Cryptographic Services
 - Available to the whole System through Mailboxes
 - Set of Crypto-Accelerators
 - SHA2, AES128, etc.



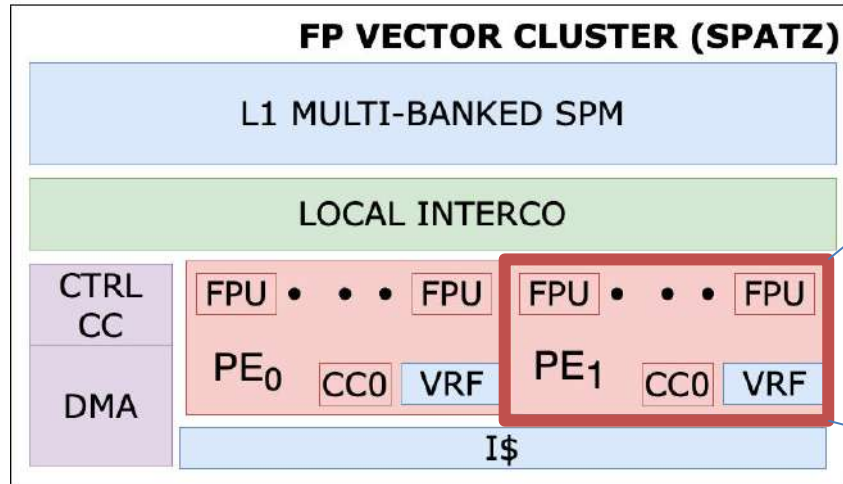
The I/O Communication



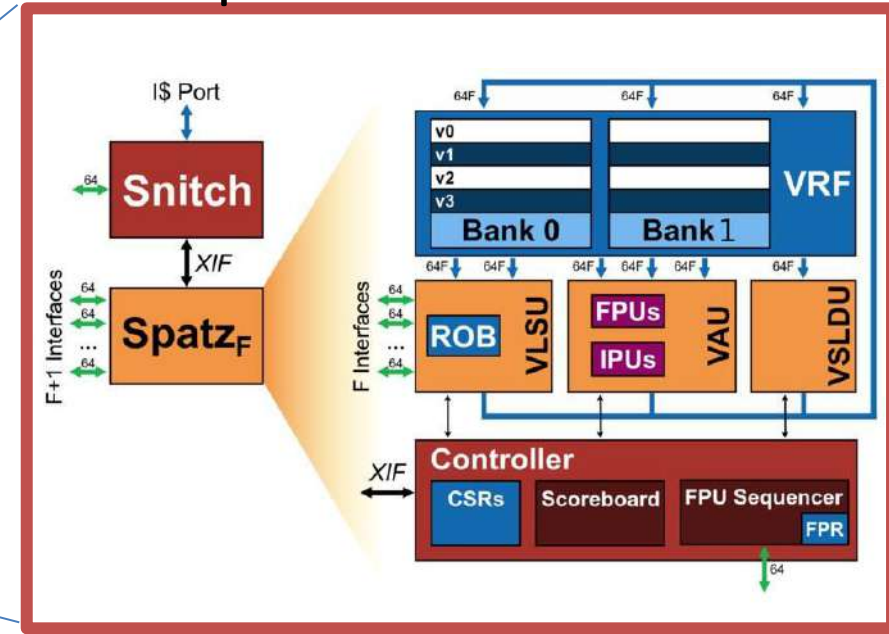
The Acceleration Clusters



The Spatz Cluster for FP Vector Workloads



Spatz Vector Processor



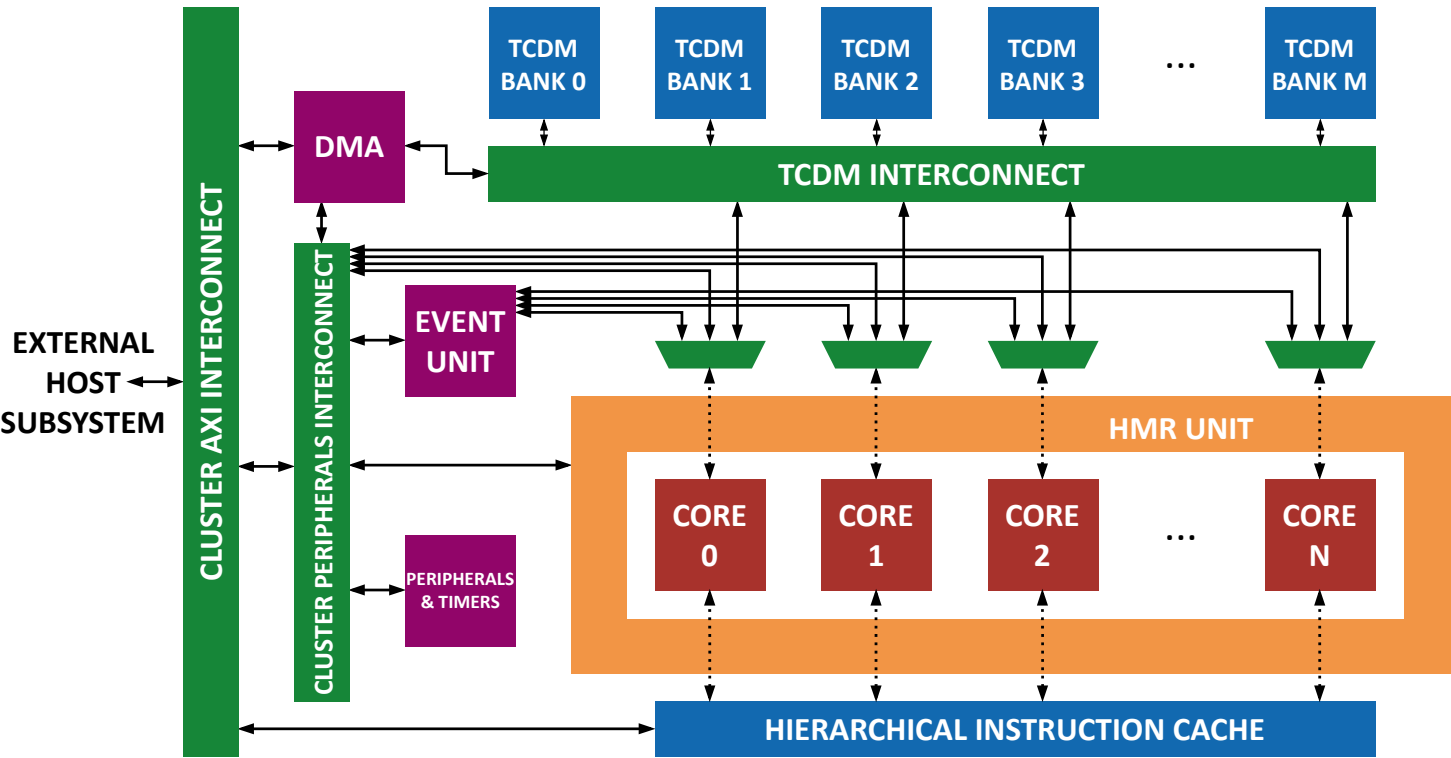
Spatz Configuration

	Spatz ₂	Spatz ₄
#MACUs	2	4
Vector length [bit]	256	512
VRF size [KiB]	1	2
Peak performance [OP/cycle]	4	8
Memory bandwidth [B/cycle]	8	16

- Compact 64-bit vector processing unit based on RVV Zve64d
- Multi-precision FPU support
 - FP64, FP32, FP16, FP8, SDOTP operations supported
- Physically-driven implementation: small footprint, high operating frequency, high scalability



The HMR Cluster for DNN-Oriented INT/FP Workloads



- 12x 32-bit RISC-V cores with support for DSP/QNN ISA Extensions
- Single-Cycle Multi-Banked Tightly-Coupled Data Memory
- Hardware Synchronizer
- DMA Controller for Explicit Memory Management
- L1-coupled TensorCore (RedMule)
- Runtime-configurable Dual/Triple core redundancy mode + hw/sw-based recovery




Carfield Goes to Intel FinFet Technology



March
2023

Q1
2024

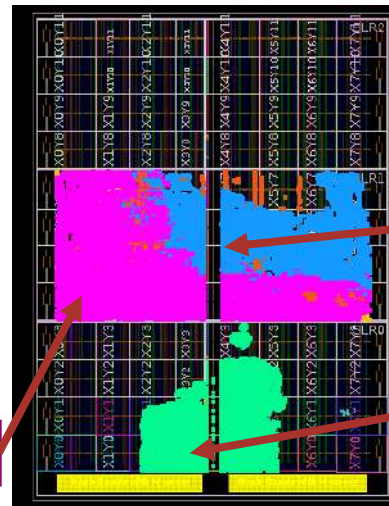


Spatz Cluster Macro

**NOW
Platform in
Place**

First test-chip out

- Intel16 FinFet technology
- BGA Flip-chip packaging



CHESHIRE

Safety Island


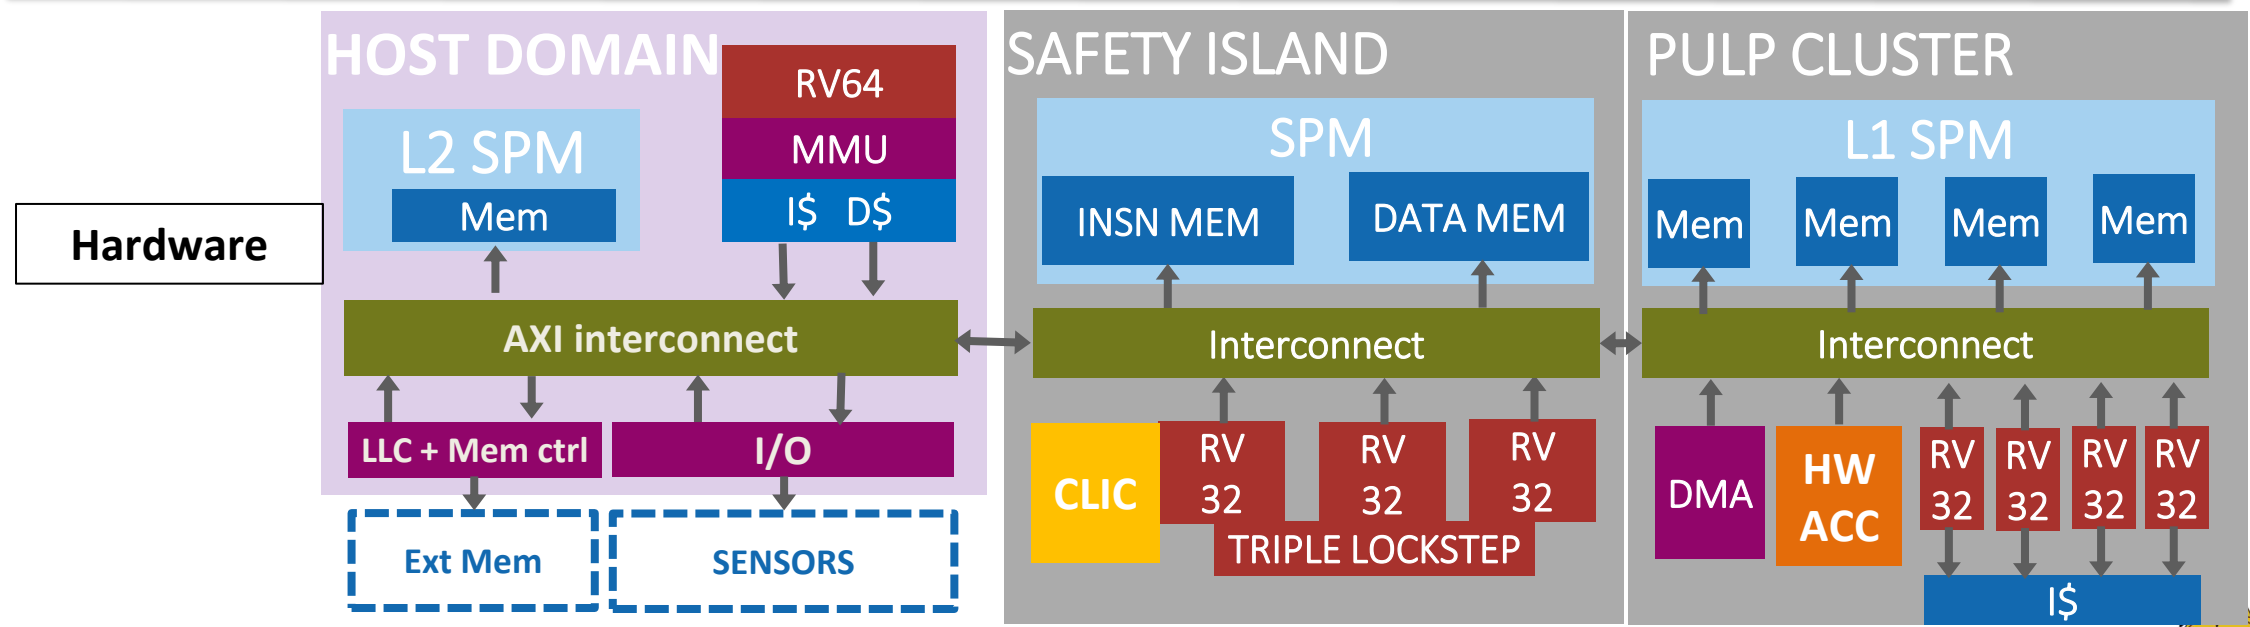
DRAM ctrl

FPGA Emulation available

**Tape-Out Date:
11 Nov 2023**



Software Stack



The (Near) Future Research Roadmap



- SW Stack Development
- Design Optimizations
- Concurrent OS Support for RTOS and GPOS
- Virtualization Assisted Processors
- Real-Time analysis of I/O communication
- Area-Optimized Safety solutions for RISC-V processors
- Reconfigurable Architectures for Image Processing
- Radiation tests



Conclusion

- Carfield: Open-Source Research Platform for Safety, Predictable and Secure Systems
 - Hardware Architecture based on fully open-source Ics
 - Complete Software stack (open-source)
- Collaborative Research opportunities among universities and industry
- First prototype soon to silicon-prove initial architecture and get feedback for next generation platform



github.com/pulp-platform/carfield



Thank you!

Q&A

References

- [Restuccia et al. DAC 2020]: Restuccia, F., Biondi, A., Marinoni, M., Cicero, G., & Buttazzo, G. (2020, July). Axi hyperconnect: A predictable, hypervisor-level interconnect for hardware accelerators in fpga soc. In *2020 57th ACM/IEEE Design Automation Conference (DAC)* (pp. 1-6). IEEE.
- [Pagani et al. ECRTS 2019] Pagani, M., Rossi, E., Biondi, A., Marinoni, M., Lipari, G., & Buttazzo, G. (2019, July). A bandwidth reservation mechanism for AXI-based hardware accelerators on FPGAs. In *31st Euromicro Conference on Real-Time Systems (ECRTS 2019)*.
- [Modica et al., IEEE ICIT 2018] Modica, P., Biondi, A., Buttazzo, G., & Patel, A. (2018, February). Supporting temporal and spatial isolation in a hypervisor for arm multicore platforms. In *2018 IEEE International Conference on Industrial Technology (ICIT)* (pp. 1651-1657). IEEE.
- [Sá et al., IEEE TC 2021] Sá, B., Martins, J., & Pinto, S. (2021). A first look at RISC-V virtualization from an embedded systems perspective. *IEEE Transactions on Computers*, 71(9), 2177-2190.

