FRAMEWORK PARTNERSHIP AGREEMENT IN EUROPEAN LOW-POWER MICROPROCESSOR TECHNOLOGIES

THIS PROJECT HAS RECEIVED FUNDING FROM THE EUROPEAN UNION'S HORIZON 2020 RESEARCH AND INNOVATION PROGRAMME UNDER GRANT AGREEMENT NO 826647

54TH EDITION OF THE TOP500 LIST (NOVEMBER 2019)

- Top#1 performance today:
  - 0.2 10^{18} Flop/s Peak
  - It is 1/5 of Exascale level of performance
- Users:

  #1-#2: US 

  #3: China 
- Processor design & technology:

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	DOE/SC/Oak Ridge National Laboratory United States	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband IBM	2,414,592	148,600.0	200,794.9	10,096
2	DOE/NNSA/LLNL United States	**Sierra** - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband IBM / NVIDIA / Mellanox	1,572,480	94,640.0	125,712.0	7,438
3	National Supercomputing Center in Wuxi China	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway NRCPC	10,649,600	93,014.6	125,435.9	15,371

Chip	Design	Manuf.
IBM POWER9		
NVIDIA Volta GV100		
Sunway SW26010		

RACE TO EXASCALE

- CPU architecture choice:
  - Japan approach: Arm/SVE (homogeneous)
  - China approach: Custom many-cores (homogeneous)
  - US approach: x86 + GPU (heterogeneous)

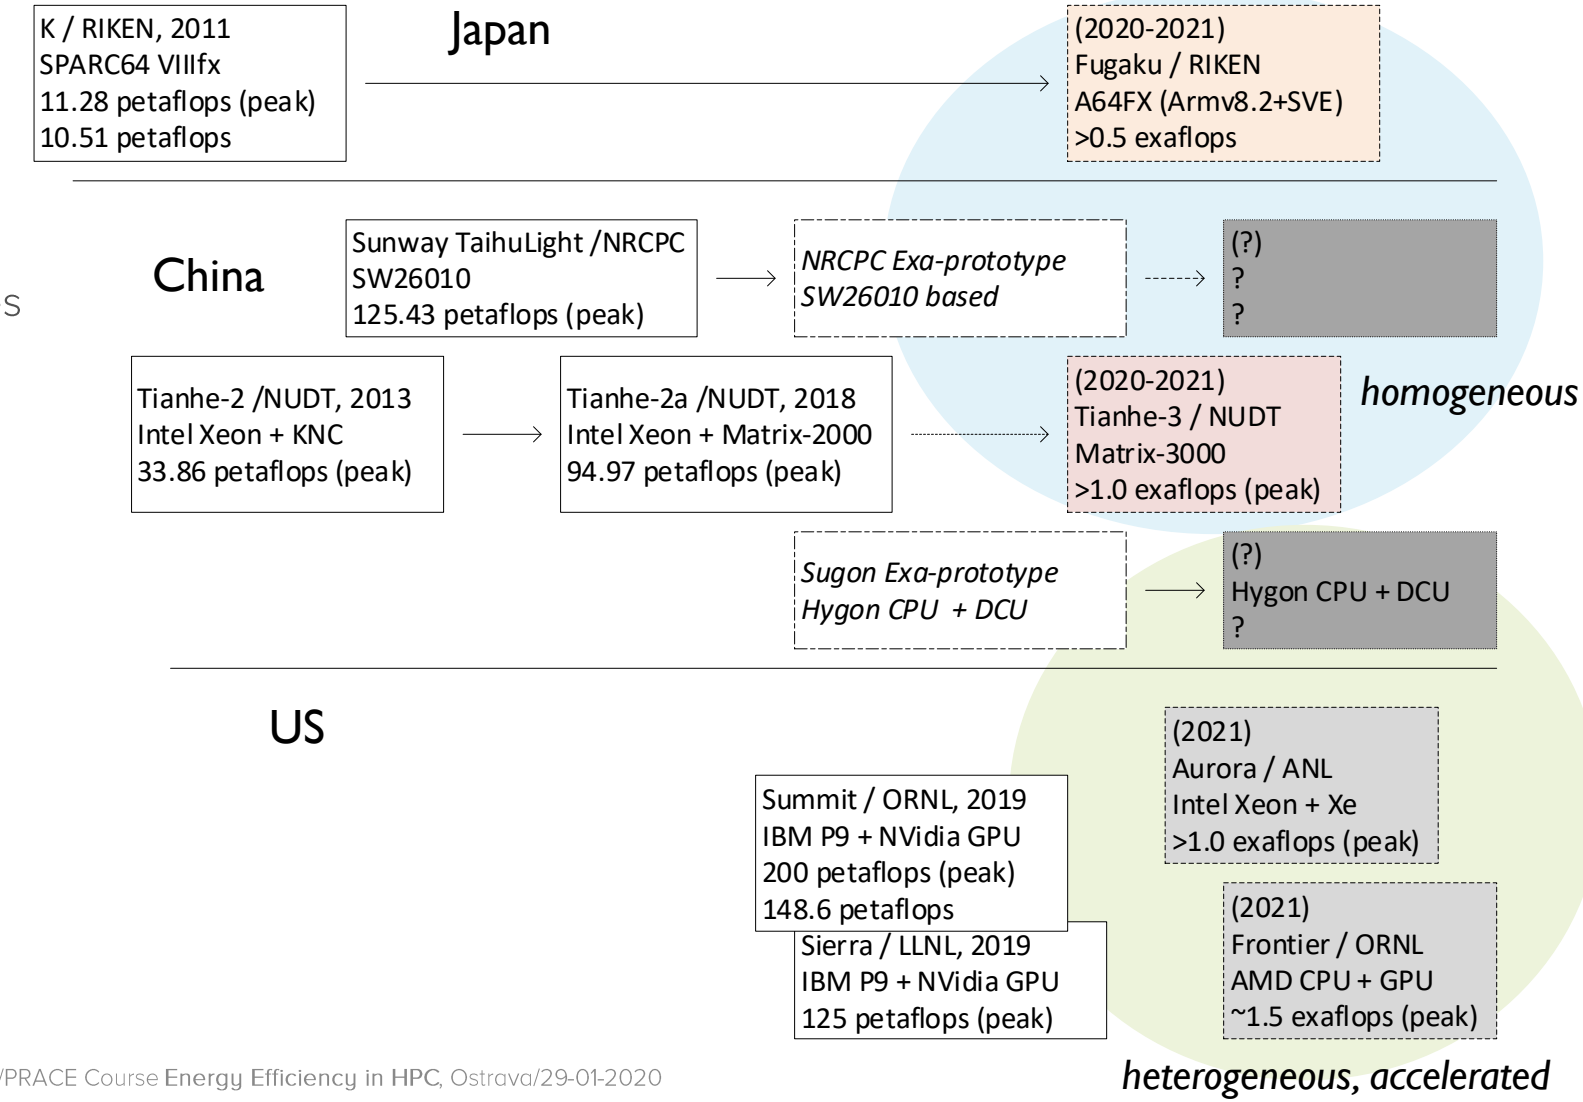

WHY EUROPE NEEDS ITS OWN PROCESSORS

- Processors now control almost every aspect of our lives
- **Security** (back doors etc.)
- Possible **future restrictions on exports to EU** due to increasing protectionism
- **A competitive EU supply chain** for HPC technologies will create jobs and growth in Europe
- Sovereignty (data, economical, embargo)

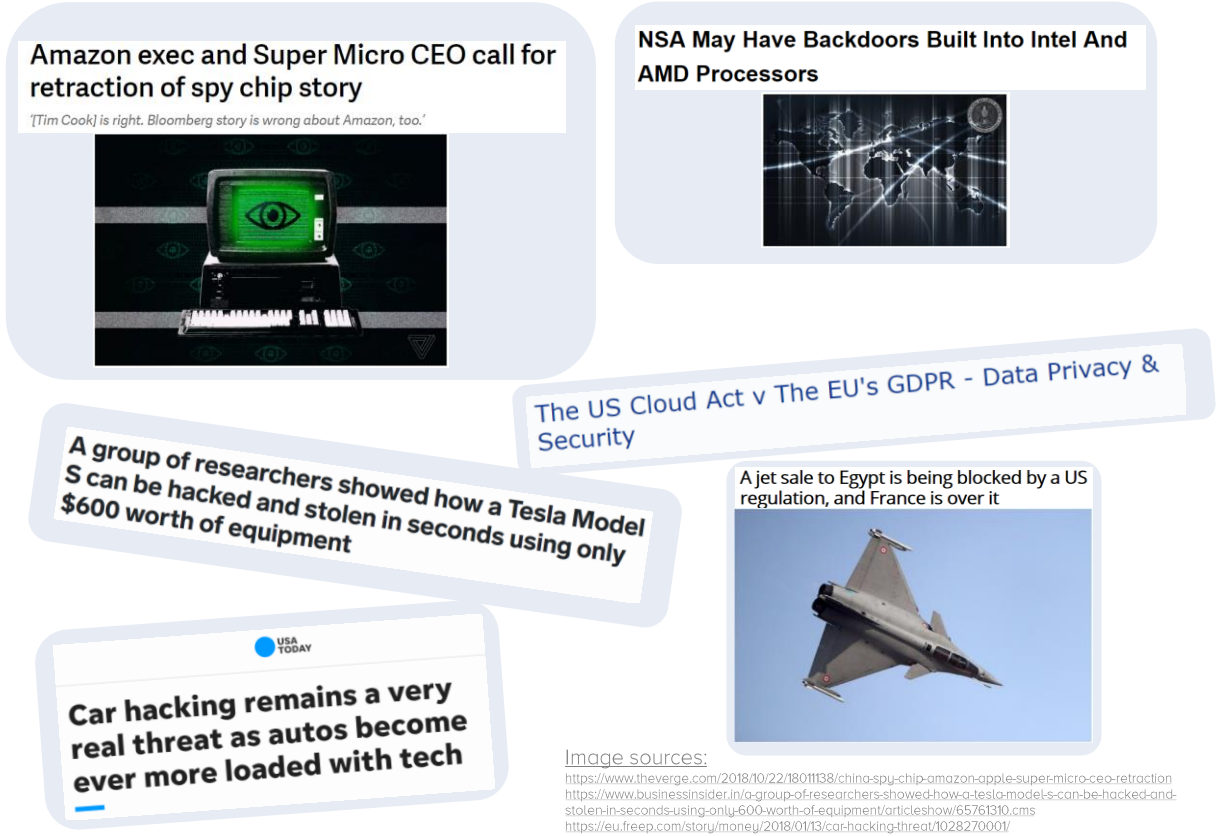

Image sources:

https://www.theverge.com/2018/10/22/18011138/china-spy-chip-amazon-apple-super-micro-ceo-retraction
https://www.businessinsider.in/a-group-of-researchers-showed-how-a-tesla-model-s-can-be-hacked-and-stolen-in-seconds-using-only-600-worth-of-equipment/articleshow/65761310.cms
https://eu.freep.com/story/money/2018/01/13/car-hacking-threat/1028270001/
https://www.eteknix.com/nsa-may-backdoors-built-intel-amd-processors/
https://www.pearse-trust.ie/blog/the-us-cloud-act-v-the-eus-gdpr-data-privacy-security
https://www.defensenews.com/global/europe/2018/08/01/a-jet-sale-to-egypt-is-being-blocked-by-a-us-regulation-and-france-is-over-it/

EUROPE'S AMBITION: EUROHPC

- Developing a new European supercomputing ecosystem: HPC systems, network, software, applications, access through the cloud
- Making HPC resources available to public and private users, including SMEs.
- Stimulating a technology supply industry

EUROPEAN PROCESSOR INITIATIVE

Design a roadmap of future European low power processors

Common platform	CoDesign, Platform for hardware and software, Power management, Modeling and Simulation
General purpose processor	High Performance General Purpose Processor for HPC
Accelerator	High-performance RISC-V based accelerator
Automotive	Computing platform for autonomous cars

SGA1 (EPI Phase 1)

H2020
36 months (Dec. 2018 → Nov. 2021)

Coordinator: ATOS

27 partners, 10 countries

Budget: 80 M€

EPI TECHNOLOGIES

- Energy Efficiency
  - Adopt Arm general-purpose CPU core with SVE / vector acceleration in the first EPI chip
  - Develop power management solutions for the EPI chip
  - Develop acceleration technologies based on RISC-V for better DP GFLOPS/Watt performance
  - Inclusion of MPPA for real-time application acceleration
  - Inclusion of eFPGA for reconfigurable logic
- Modularity
  - Supply sufficient Memory Bandwidth (Byte/FLOP)
  - Focus on programming models to include accelerations.
  - Develop a Common Platform to enable EPI accelerations and that eases incremental roadmap implementation

EPI ROADMAP

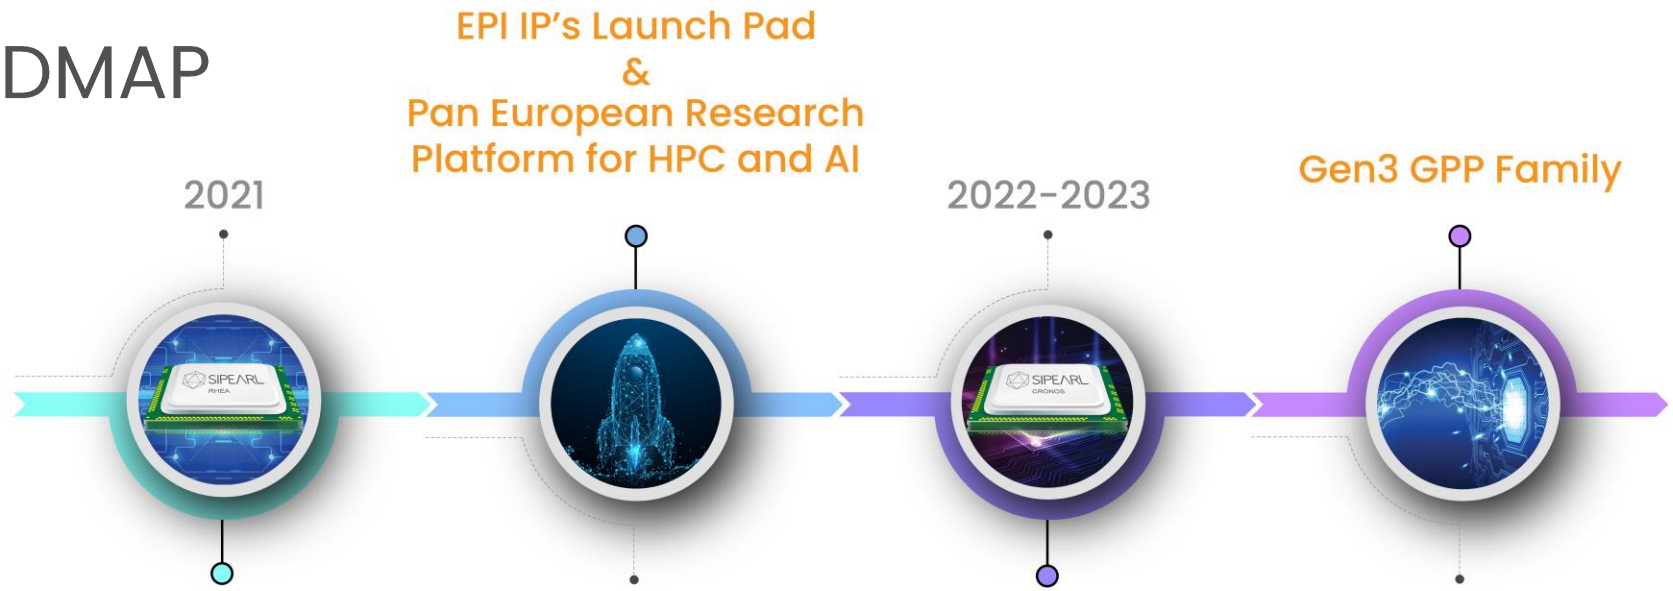

EPI FABLESS COMPANY

- EPI's Fabless company
  - licence of IPs from the partners
  - develop own IPs around it
  - licence the missing components from the market

  - generate revenue from both the HPC, IA, server and eHPC markets
  - integrate, market, support & sales the chip
  - work on the next generations

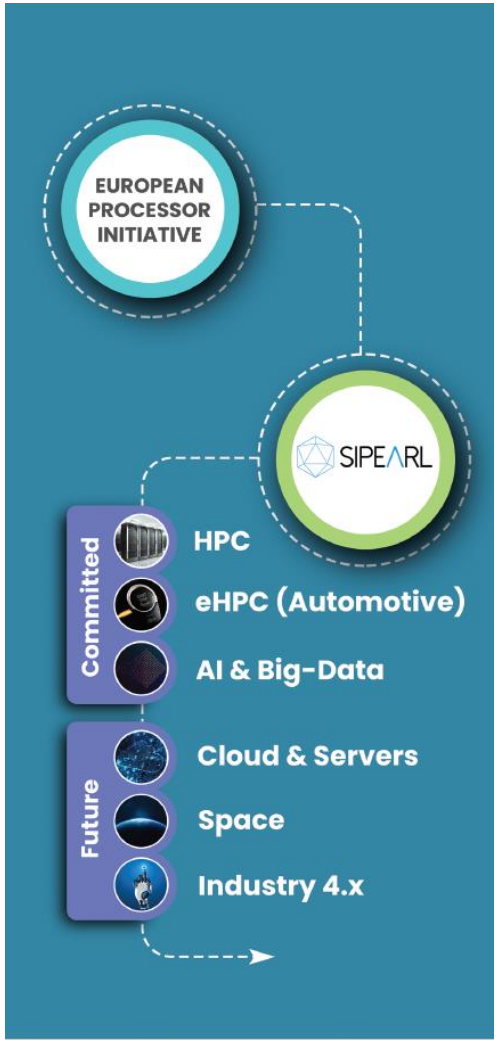

CONCLUSION

- HPC is crucial to resolve societal challenges and preserve European competitiveness
- Europe is going in the right direction with EuroHPC. This must be sustained in the long-term
- The chip design effort must continue for the EU's security and competitiveness, and should create a processor ecosystem covering IoT, servers, cloud, autonomous connected vehicles and HPC

www.european-processor-initiative.eu

@EuProcessor

European Processor Initiative

European Processor Initiative

GPP AND COMMON ARCHITECTURE

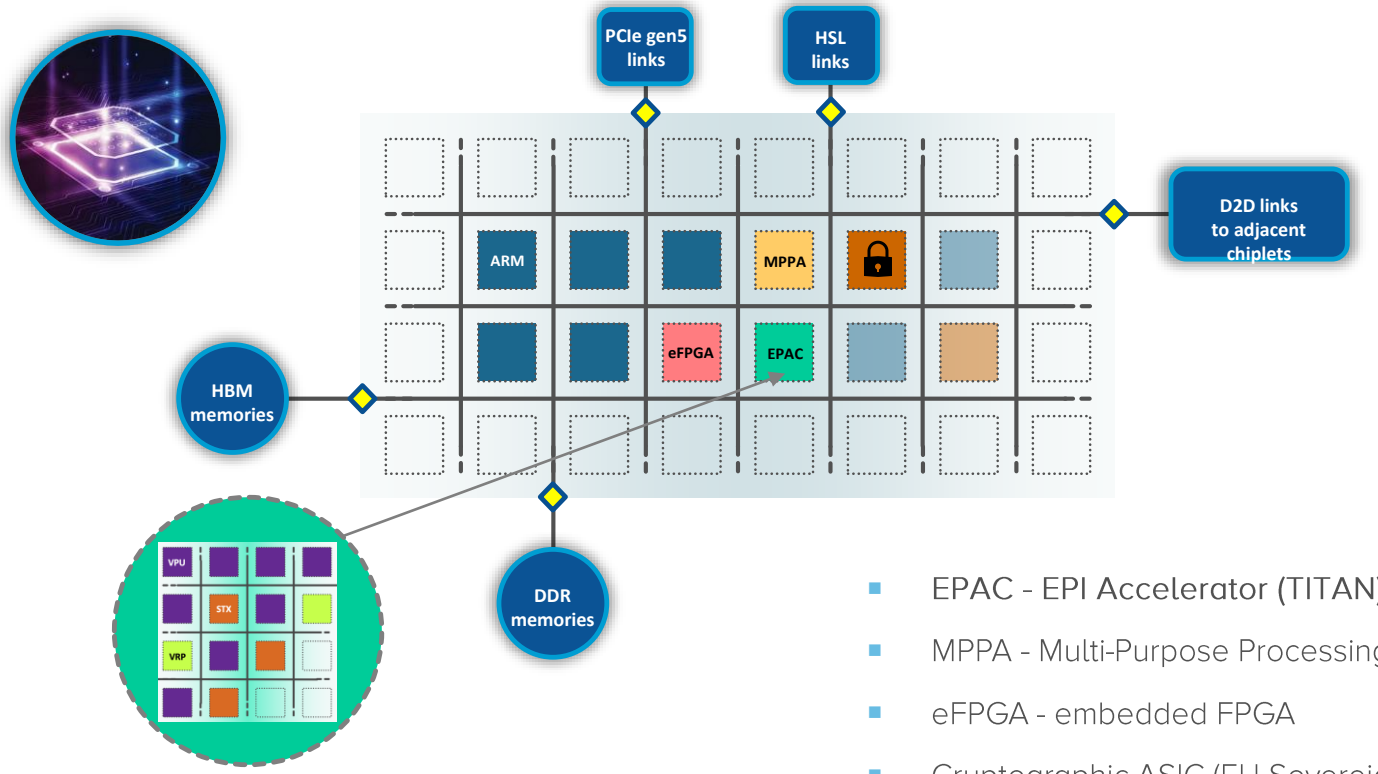

- EPAC - EPI Accelerator (TITAN)
- MPPA - Multi-Purpose Processing Array
- eFPGA - embedded FPGA
- Cryptographic ASIC (EU Sovereignty)

EPAC – RISC-V ACCELERATOR

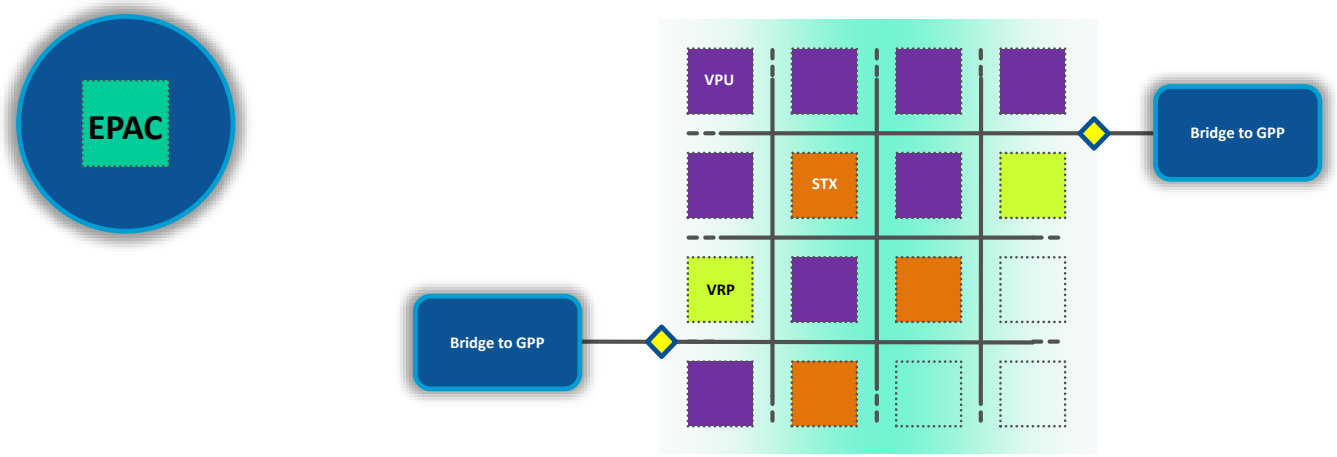

- EPAC – TITAN = EPI Accelerator
- VPU – Vector Processing Unit (plan of record)
- STX – Stencil/Tensor accelerator (PoR)

RHEA PROCESSOR

- Rhea is the first EPI General Purpose Processor
- Rhea targets HPC application
- Rhea is the first « instanciation » of EPI Common Platform
- Rhea design is led by SiPEARL (the EPI fabless company) and joint-developed by EPI partners.

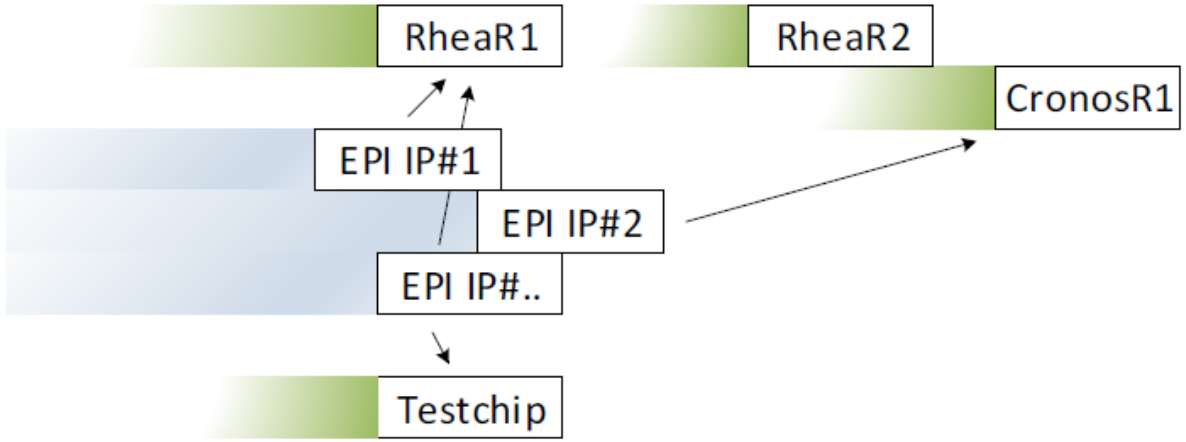

- Rhea chip will be integrated into test platforms in order to validate the hardware units, develop the software, and run applications.

Automotive PoC

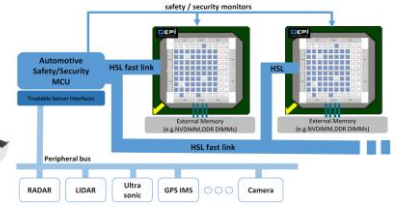

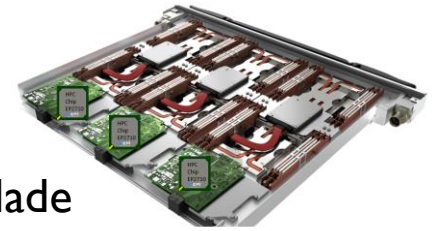

HPC blade

PCIe daughter card

RHEA DESIGN

- Generic processing multi-core backbone:
  - Multi-core Arm Zeus processor with SVE engines for pre-ExaScale level generic processing.
  - Coherent NoC with distributed system level cache to keep the data local.
- Prototypes of High energy-efficient accelerator tiles:
  - RISC-V based acceleration (EPAC) for better GFLOPS/Watt performance.
  - Multi-Purpose Processing Array (MPPA-Kalray) for real-time application acceleration.
  - eFPGA (Menta) reconfigurable logic for flexibility.
  - Accelerators work in I/O coherent mode and share the same memory view as the multi-core backbone.
- HBM2E, DDR5 memory support.
- PCIe gen5 support for loosely coupled accelerators.
- High speed links for SMP extension and tightly coupled accelerators.
- Power Management infrastructure with low voltage support for energy efficiency
- Security infrastructure.
- Peripherals to connect an automotive MCU for PoC purpose.
- First Rhea chip will be fabricated in 6nm technology aiming at the highest processing capabilities and energy efficiency

RHEA ARCHITECTURE

- Memory-coherent NoC connects
  - Array of computing units (CU): Arm cores, EPAC, MPPA, eFPGA
  - Memory and I/O controllers
  - Bridge to links
- High speed links
  - Die-2-Die links to connect on-package dies
  - HSL links to connect on-board packages
- Top level infrastructures
  - Power management & controller
  - Security

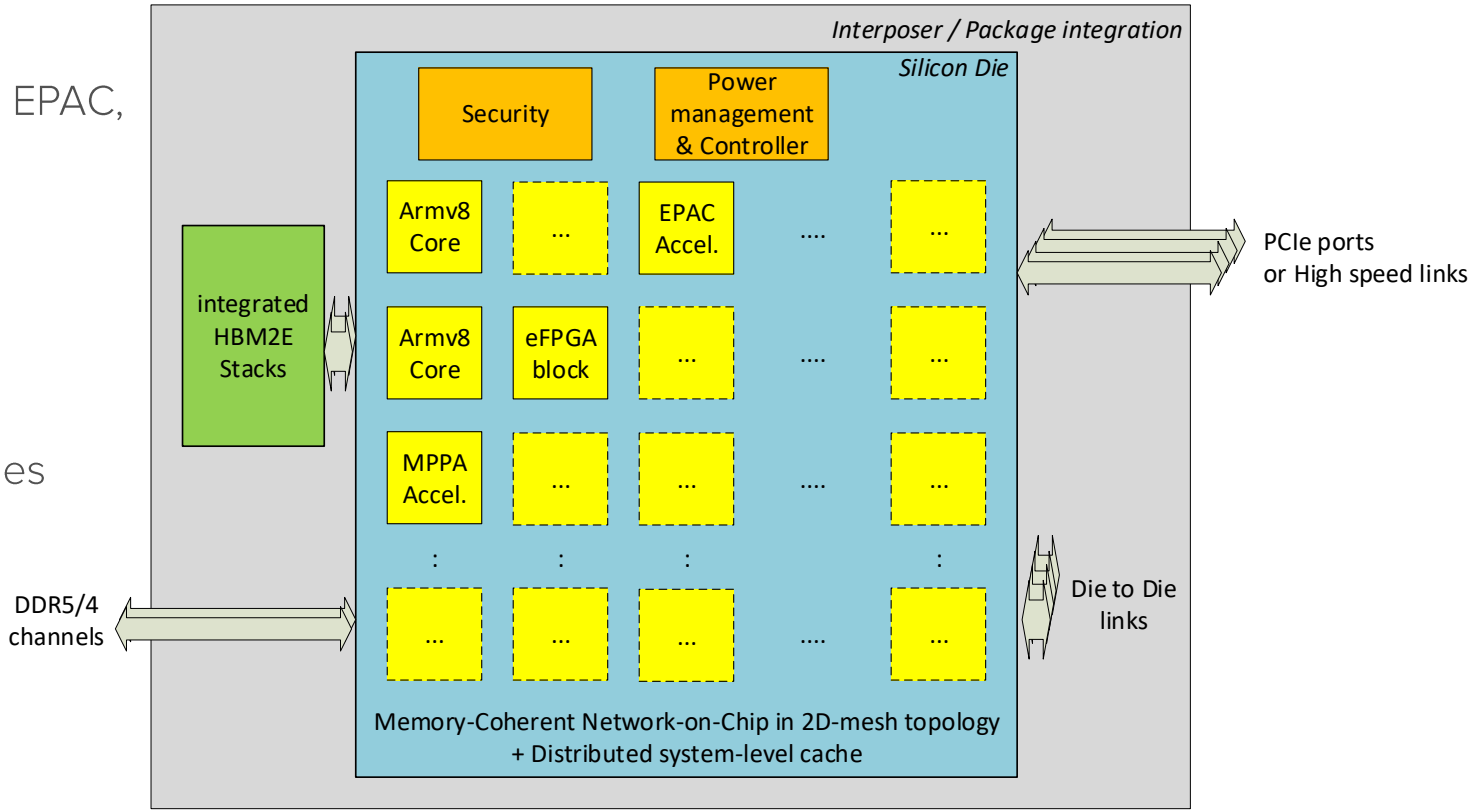

NoC: network on chip
HSL: High speed links (with memory coherent support)

POWER ASPECTS

ANDREA BARTOLINI (UNIBO) – POWER MANAGEMENT LEADER

WP3 – POWER MANAGEMENT & CONTROLLER

- Out-of-band – zero overhead telemetry
- Node Pcap – Max perf @ Pnode<Pmax
- RAS – error and conditions reporting

System Management / RM

- cpufreq/ cpuidle
- Based on O.S. metrics
- Slow & often unused

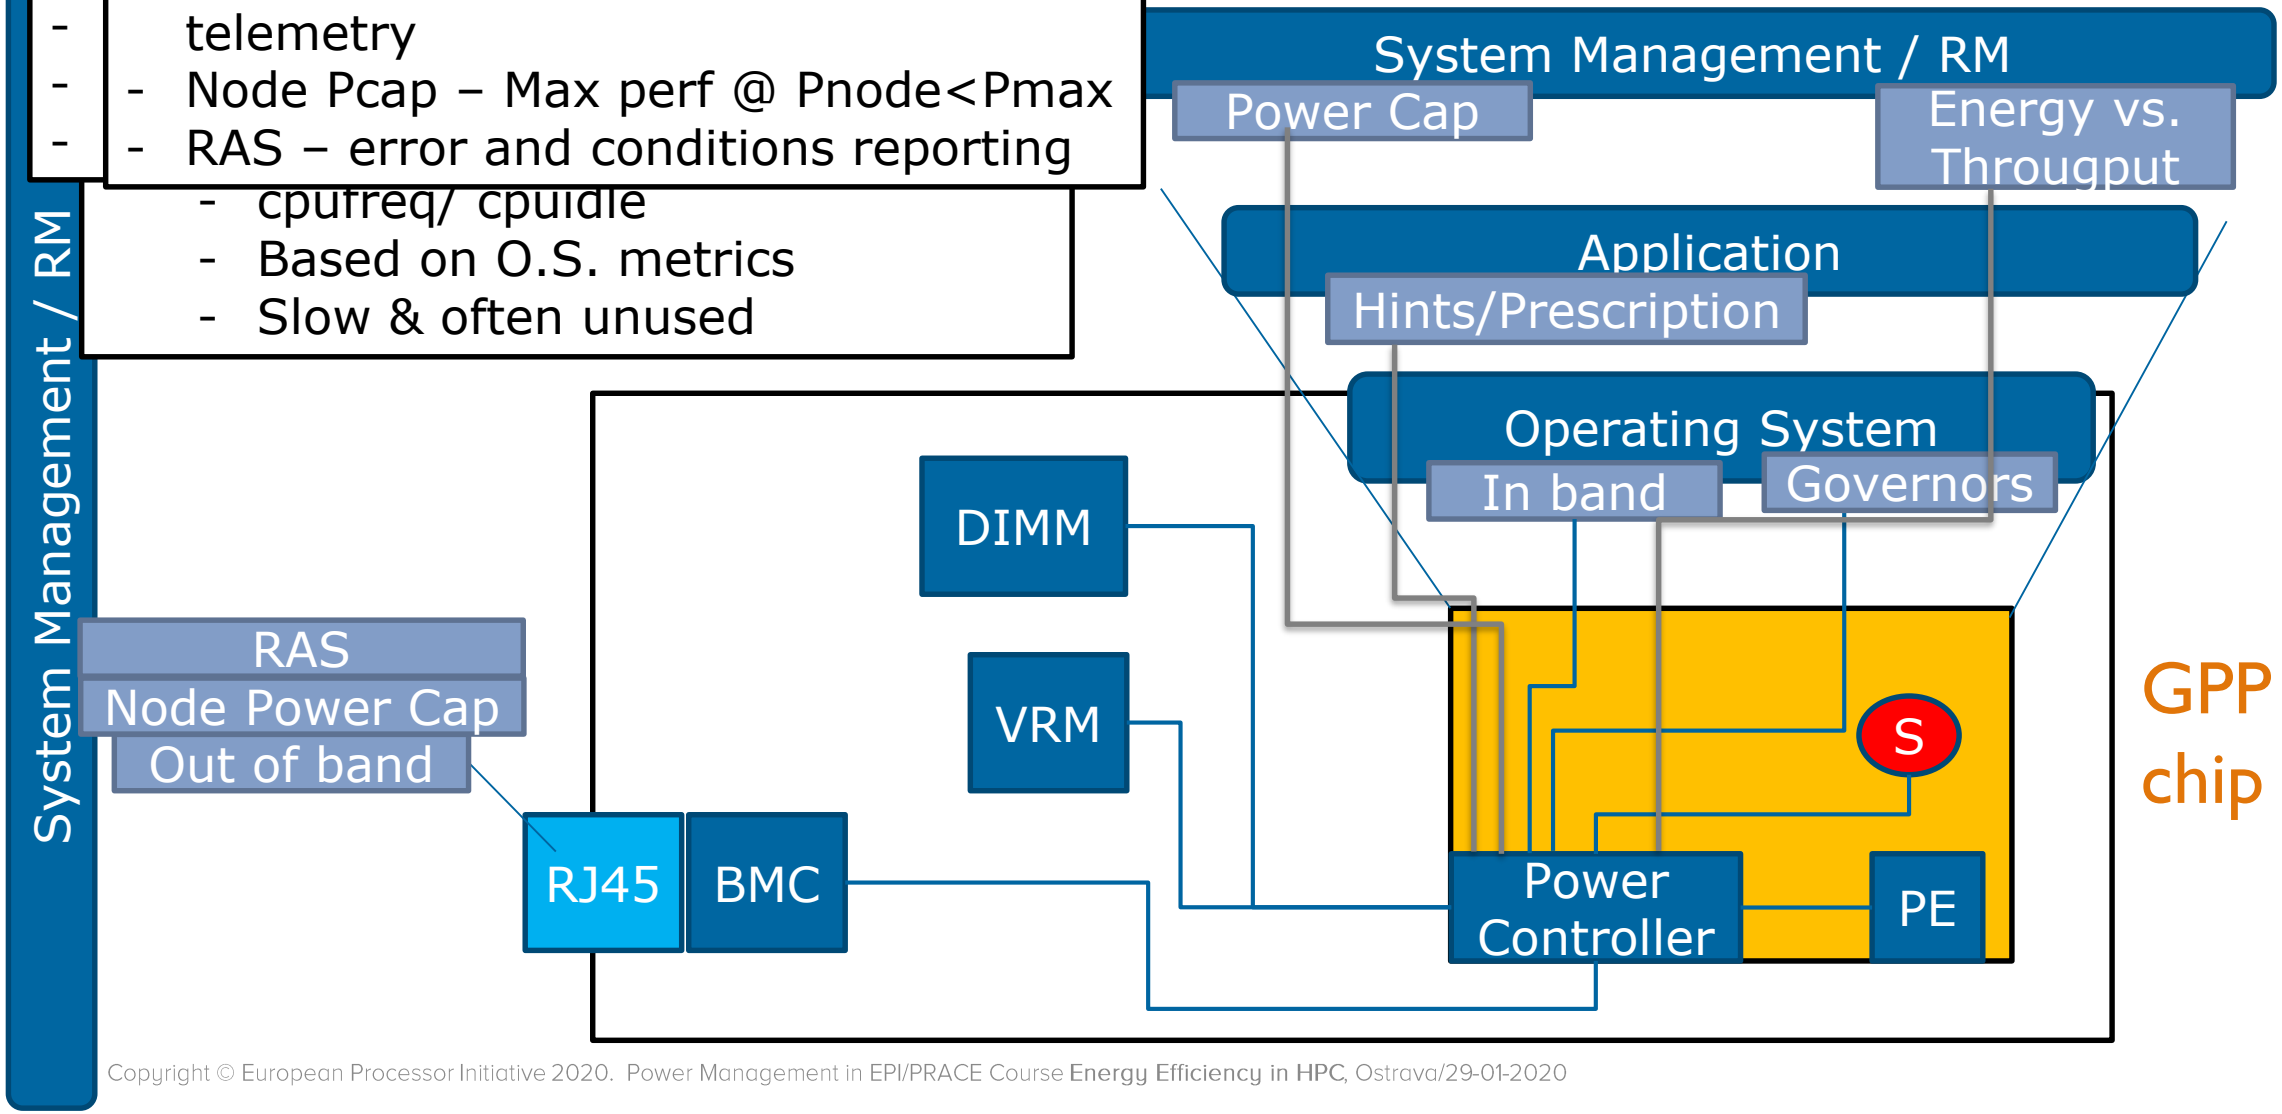

POWER MANAGEMENT SOA & REQUIREMENTS

	Intel	IBM	ARM	AMD	Cray	Fujitsu
Monitor (Domain,Granularity)	S, M, A, T **1ms**	N, S, M, A, T, U **500us**,10ms aggregation 16ms for T & U, 100ms aggregation	S, M, T 1-10KHz with SCP	N, S, M, A, T 1 sec (C), 1ms (G)	N, S, M, A, N OOB (100ms)	N, S, C, M **1ms** (N), ~ns - model based (C)
Control (Domain,Granularity)	S, M RAPL **1ms** (in-band), DVFS **500us**	N, S, M, A 10-100ms	S, M 1-10KHz (100ms to 1s)	N, S, M, A ~secs	N, S, M, A DVFS, RAPL, min-max range, 10-30s at job launch	S, C, M, **DVFS**, Decode Width, HBM2 B/W
Interfaces, Tools, etc	RAPL MSRS, msr-safe, libmsr, PAPI, likwid	**OpenBMC**, amester, **Memory Map**	ACPI, SCP (sys ctrl proc), IPA (intelligent allocator), PAPI	Likwid, PAPI, **Memory Map**	CapMC, PAPI, Cray **BMC interfaces**	Power API, PAPI

Source PowerStack19

Socket (S), Core (C), Memory (M), Accelerator (G), Node (N), Utilization (U), Temperature (T)

EPI power management design targets:

- Support for fine grain power monitoring, and control
- An higher performance power controller capable of supporting advanced power control algorithms.

GENERAL ARCHITECTURE

- Top level infrastructures
  - Power management & controller
  - Dedicated power management and control network
  - Security
- EPI Power Management Subsystem
- RISC-V ISA, Derived from the PULP platform
- Parallel processor w. DSP extensions
- Open-Source Design

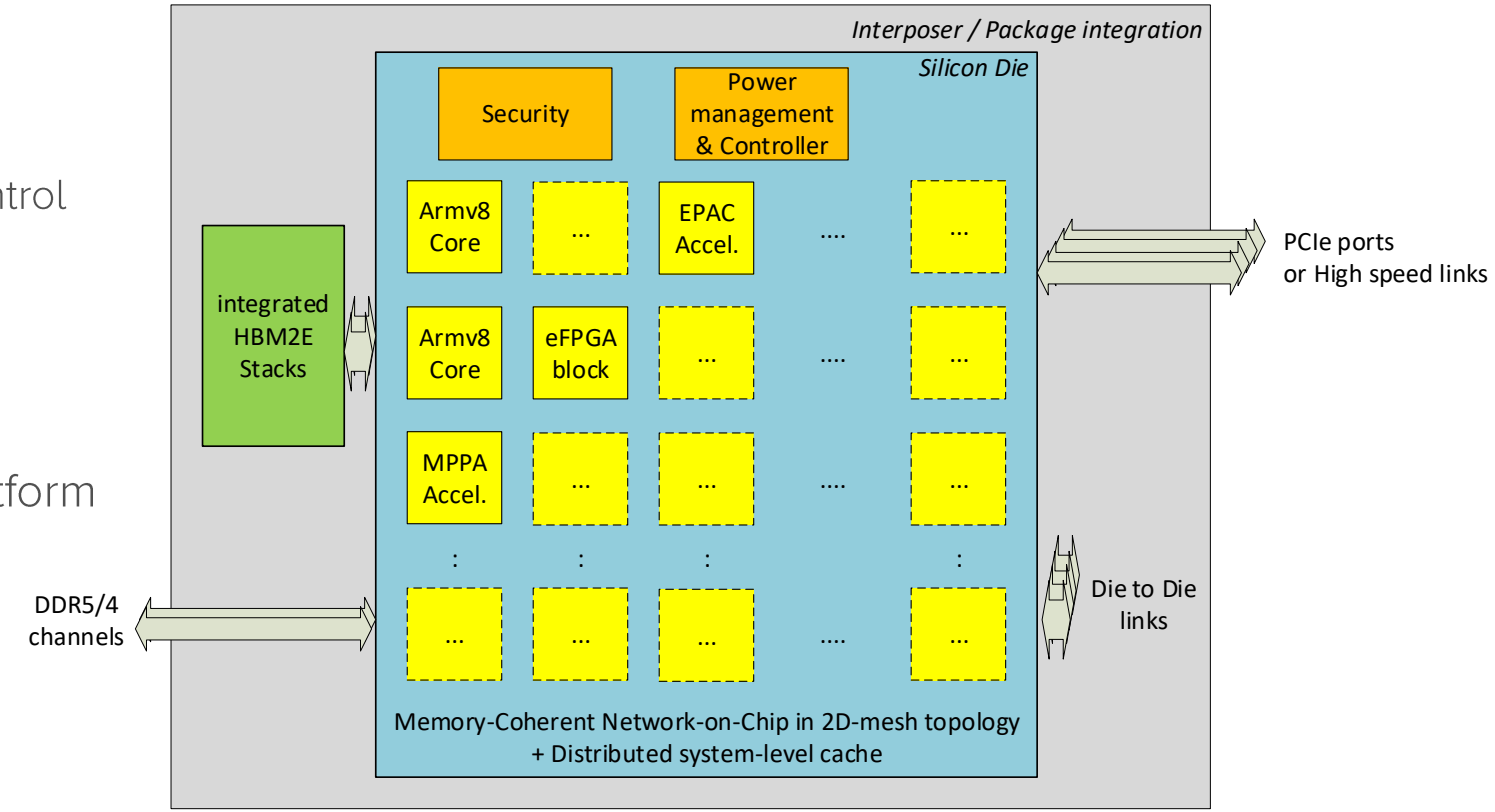

PULP OPEN SOURCE HARDWARE AND SOFTWARE STACK – PULP

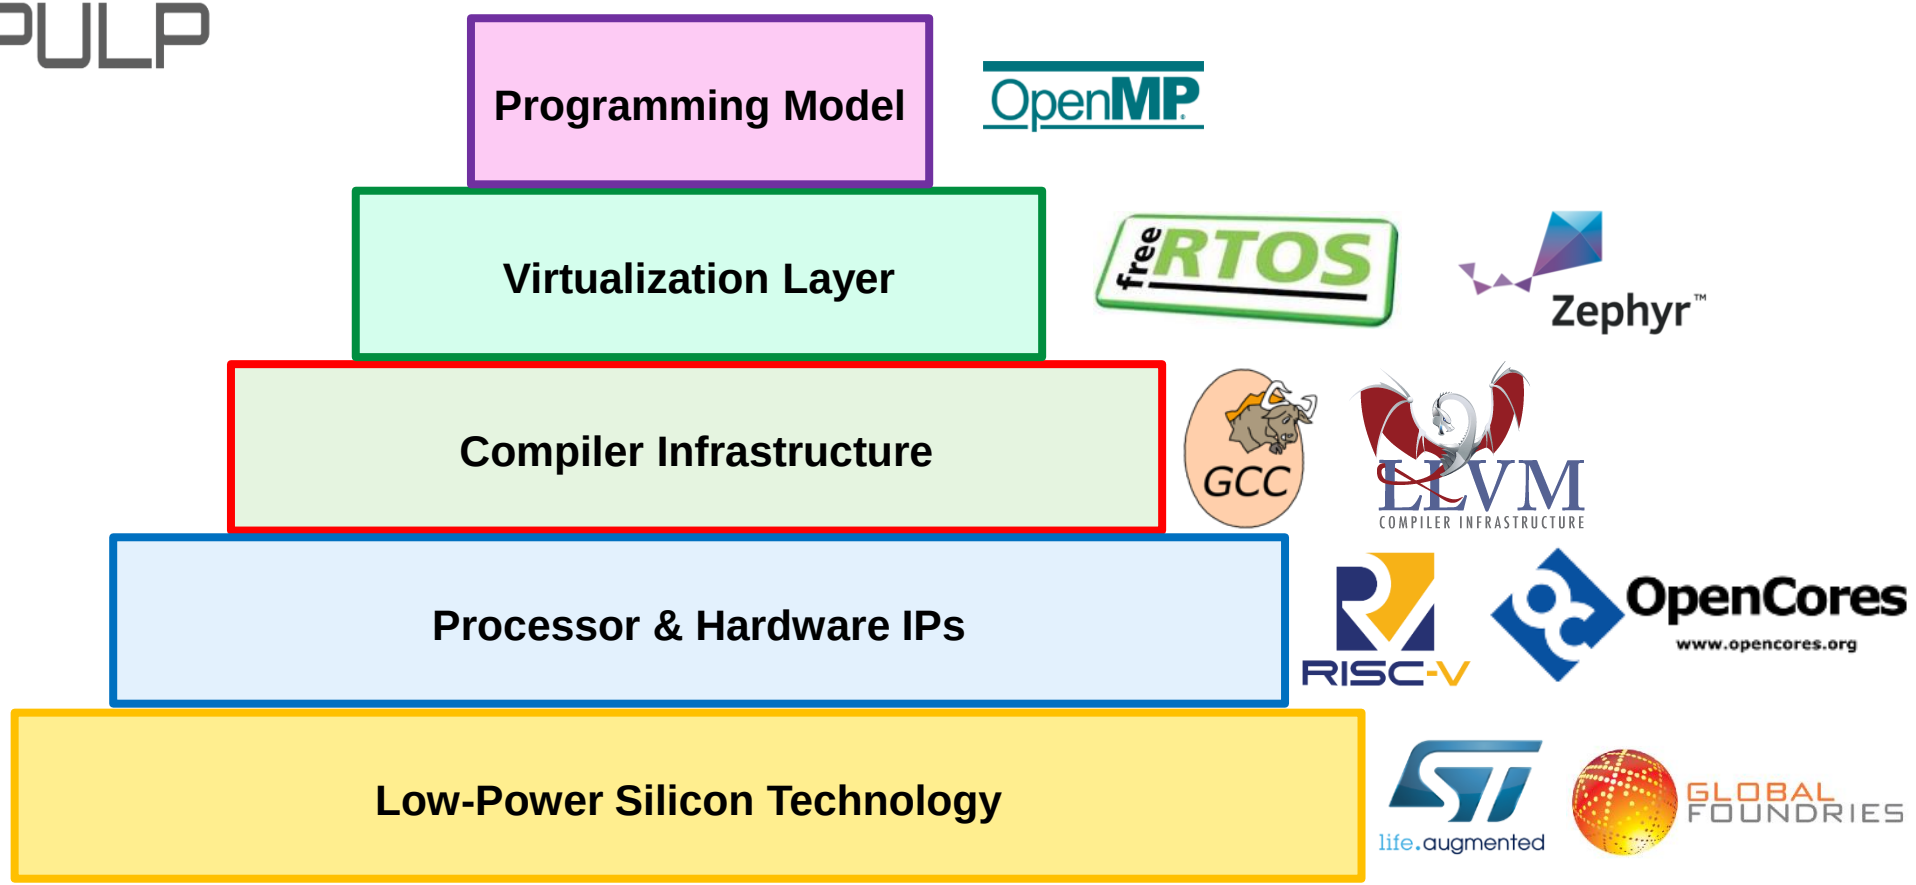

https://github.com/pulp-platform

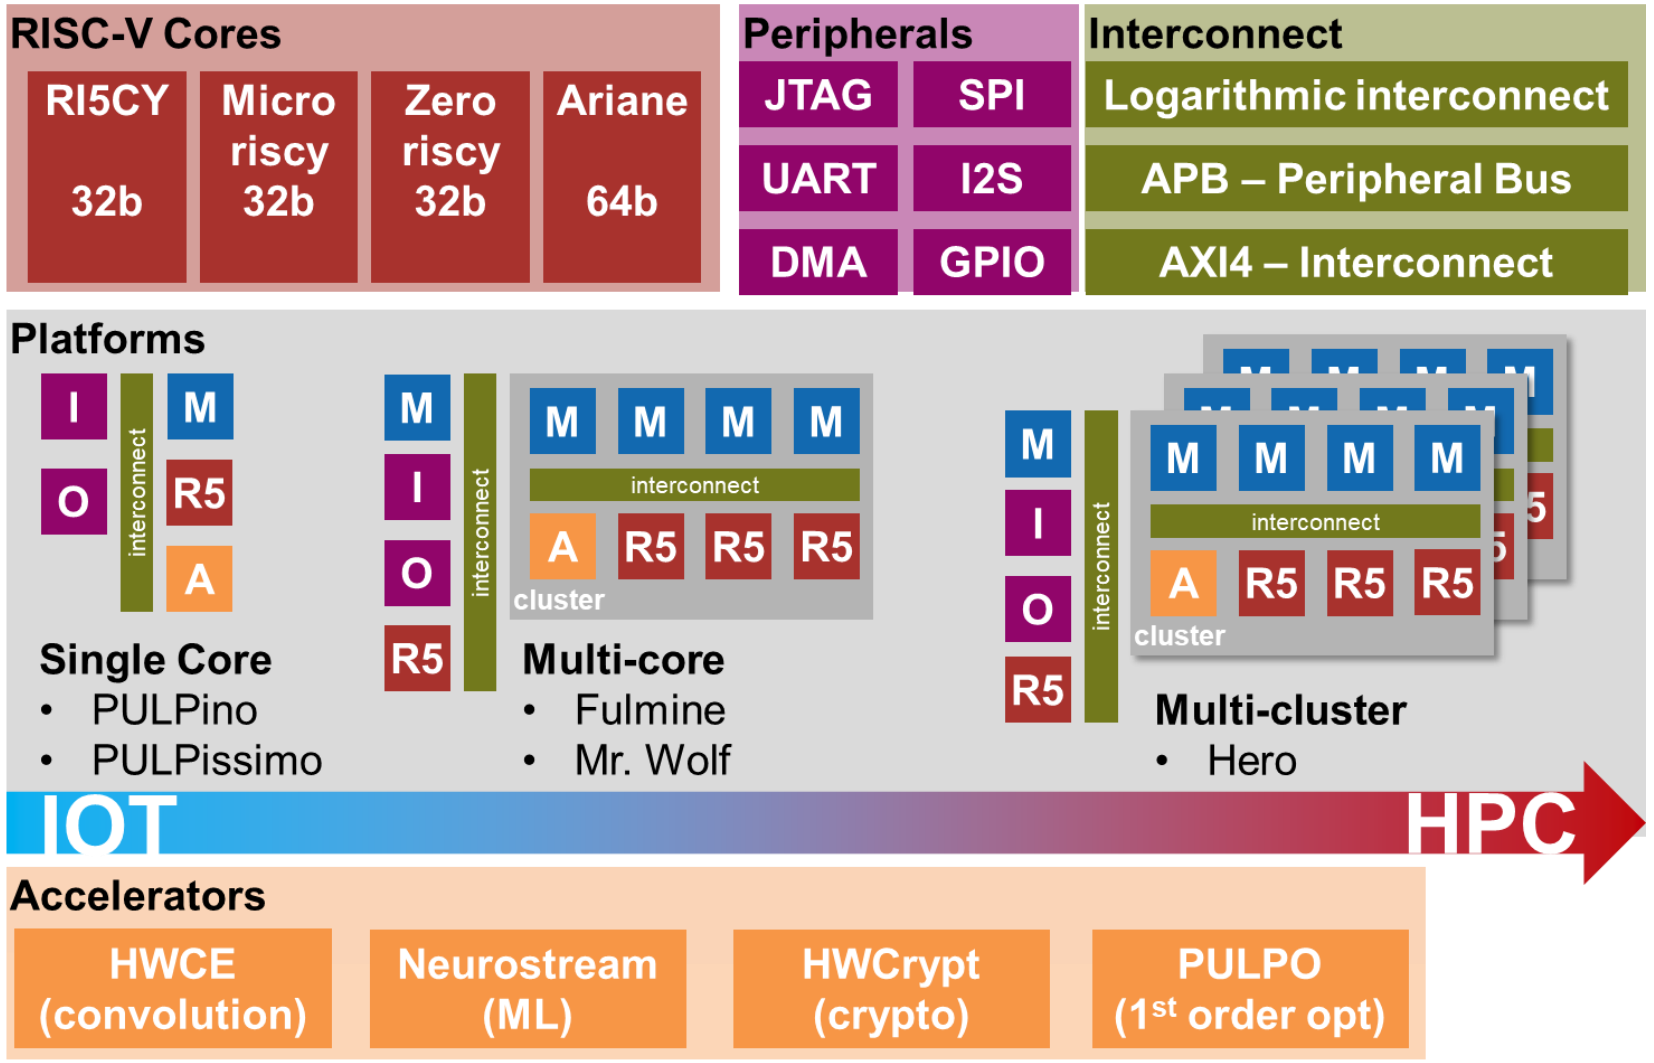

THE POWER CONTROLLER FIRMWARE

T control

Watchdog reset
Write power controller settings
Write to internal memory telemetry data
Read PVT sensors
Read workload from O.S.
Read **target P/C state settings, power budget**
Read Pending BMC requests
Compute controller settings

PM task

- Read voltage regulator, power, status (VR)
- Power model update

BMC

- Read pending command queue
- Decode Command/data
- Perform action:
  - Change **target P/C state, power budget**
    - Set pending BMC
  - Ask telemetry data

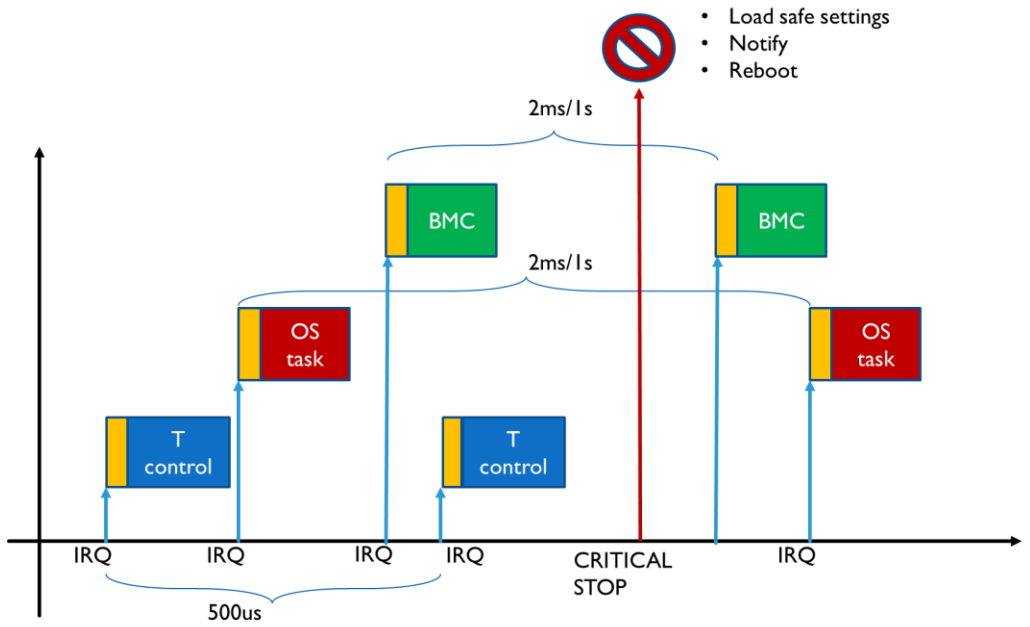

[ICECS19] A. Bartolini et al. A PULP-based Parallel Power Controller for Future Exascale Systems

CONCLUSIONS

- Power management is a key aspect of HPC processors
- Implemented by mean of a embeddded computing subsystem with extensions for interfacing with the power management IPs.
- EPI will leverage a best-in class power management subsystem based on parallel architecture with DSP extensions.

