

EPI TUTORIAL: FIRST STEPS TOWARDS A MADE-IN-EUROPE HIGH- PERFORMANCE MICROPROCESSOR

HIPEAC 2020

BOLOGNA, ITALY

22 JANUARY 2020



**European
Processor
Initiative**



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

AGENDA

SECTION	FROM	TO	TOPIC
Intro General Overview Accelerators Power PCIe	14:00		Introduction to tutorial Josip Knezovic, University of Zagreb, Croatia
	14:05		General EPI overview Denis Dutoit, CEA, France
	14:20		Common Platform & Rhea 1st Implementation Denis Dutoit, CEA, France
	14:50		Accelerators & Power aspects Mauro Olivieri, BSC, Spain Andrea Bartolini, University of Bologna, Italy
	15:15		EPI PCIe daughter card as software development vehicle Fabrizio Magugliani, E4, Italy
	15:30	16:00	BREAK
EPI and RISC-V	16:00	17:30	Bringing up EPI RISC-V Vector architecture Software Filippo Mantovani, BSC, Spain Roger Ferrer Ibanes, BSC, Spain

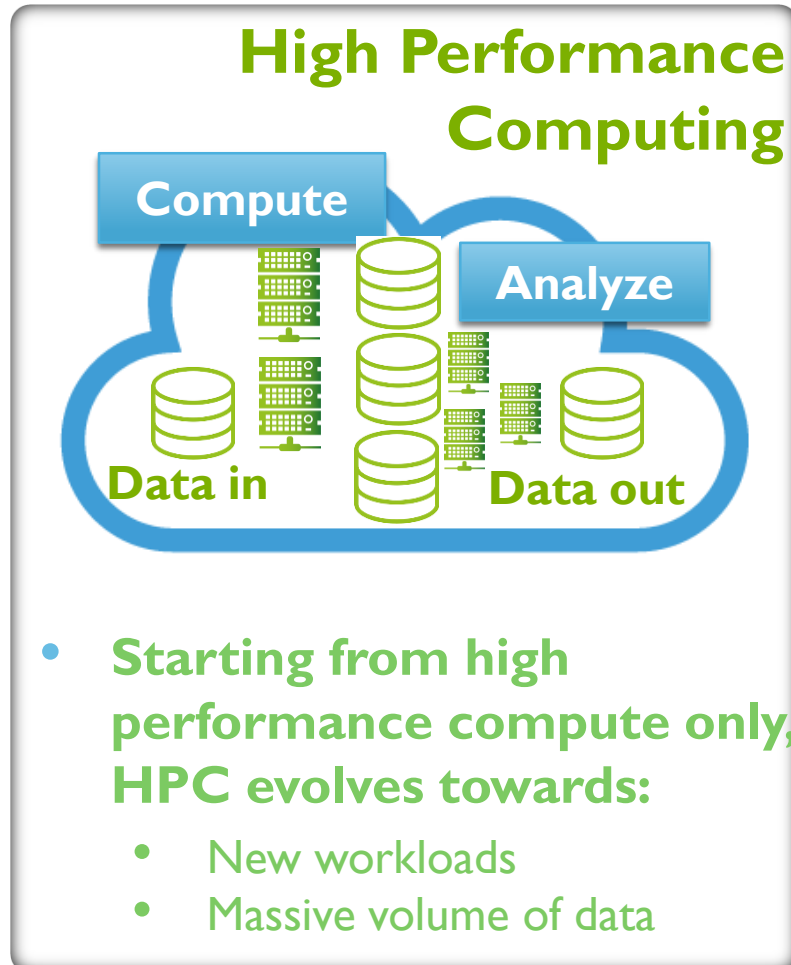


EPI OVERVIEW

DENIS DUTOIT (CEA)

HIGH PERFORMANCE COMPUTING EVOLUTION

Source CEA



New drivers	Requirements	Solutions
New workloads	<ul style="list-style-type: none"> - More computing performance (Ops per second), also for simple operations (FP16, FP8, INT...). - Energy efficiency (Ops per Watt). 	Heterogeneity: Generic processing + accelerators Low power design
Massive volume of data	<ul style="list-style-type: none"> - Increased Bytes per Flops. - High bandwidth/low latency access to more and more data. 	New type of memories and associated integration

- ➔ 10x energy efficiency improvement every 4 years
- ➔ 4x data amount every 4 years

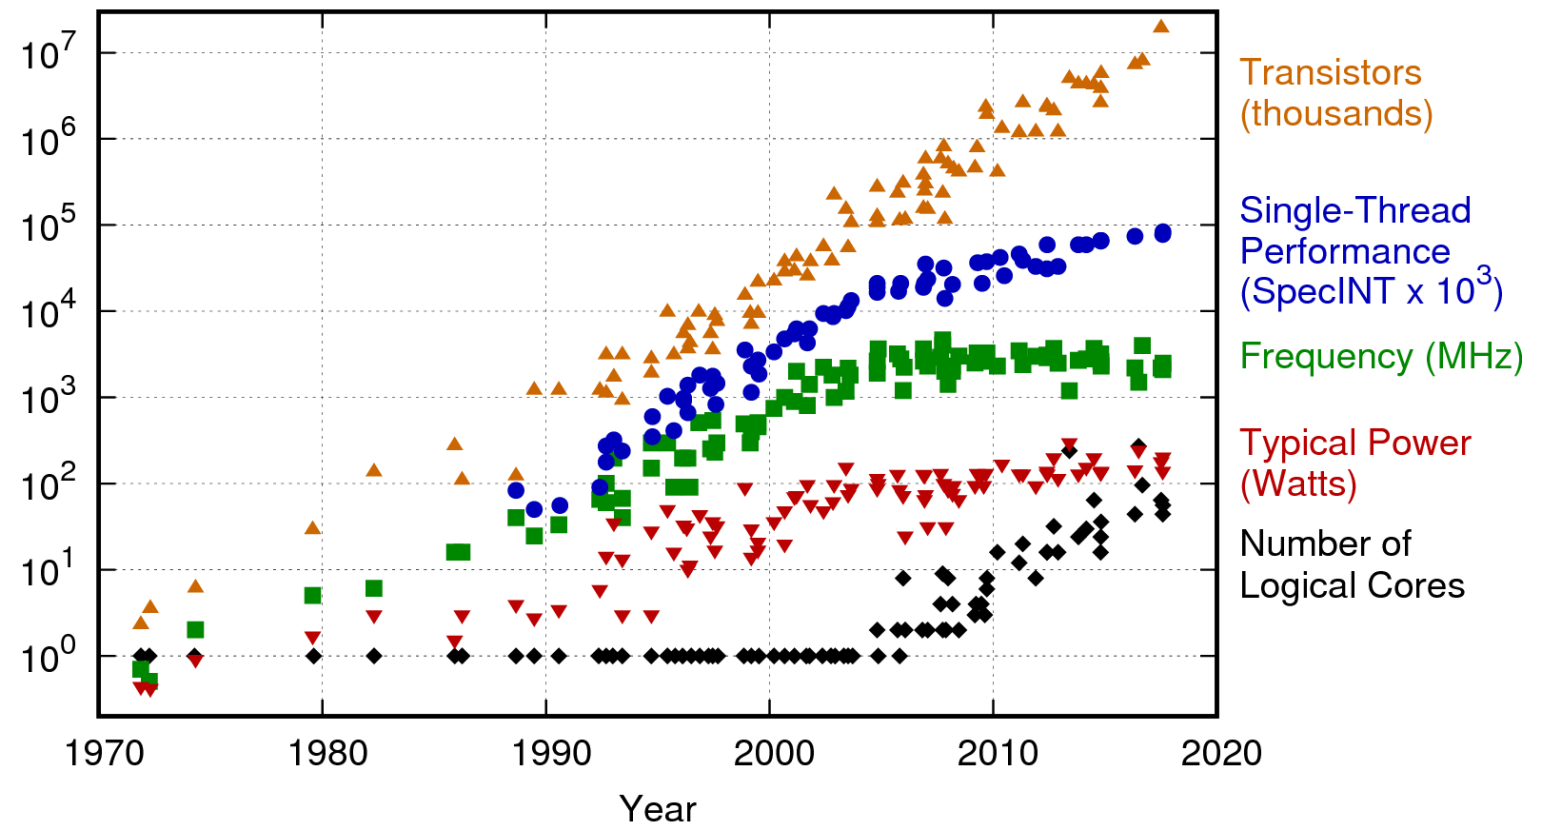


TERA1000 - CEA

TECHNOLOGY SCALING TRENDS

- Processor performance is fueled by semiconductor technology scaling.
- But silicon technology reached, reaches and will reach successive limits or walls: frequency, power, transistor density.
- Architecture evolution is the answer for recovering from technology walls.
- But it causes complex system design (HW and SW)
- And design cost is increasing

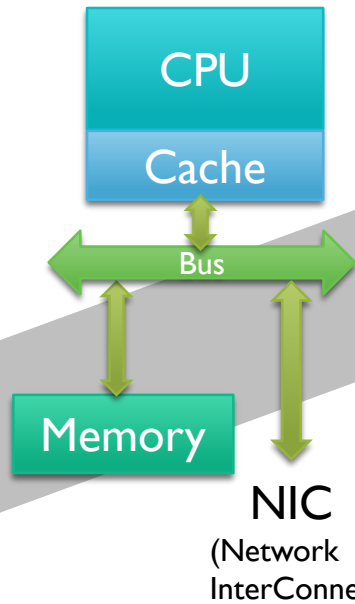
42 Years of Microprocessor Trend Data



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

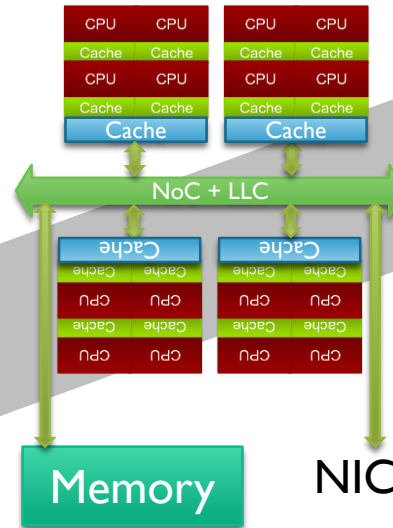
COMPUTE NODE ARCHITECTURE EVOLUTION

Performance = ~frequency



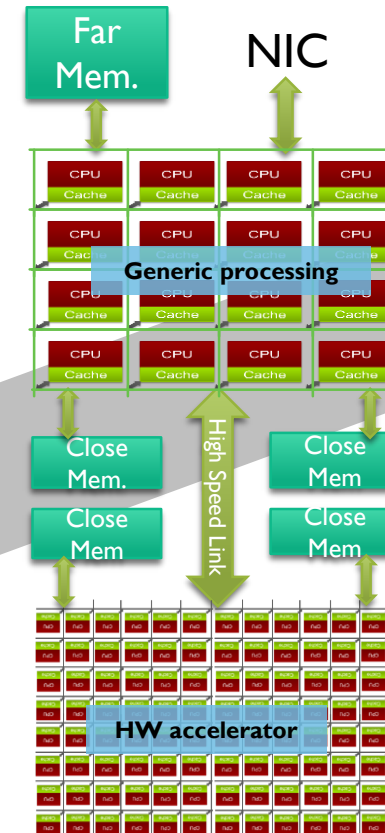
2005
Frequency wall

Performance = ~nb cores



2015
Power wall

Performance = ~architecture



2025
Moore's law slow-down
Cost wall

X86 cores, RISC cores, Co-pro extension, Accelerator, GPU, FPGA, Real Time processing, Homogeneous, Heterogeneous, Data centric...

- Today: processor architecture choice is driven by energy efficiency

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 

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

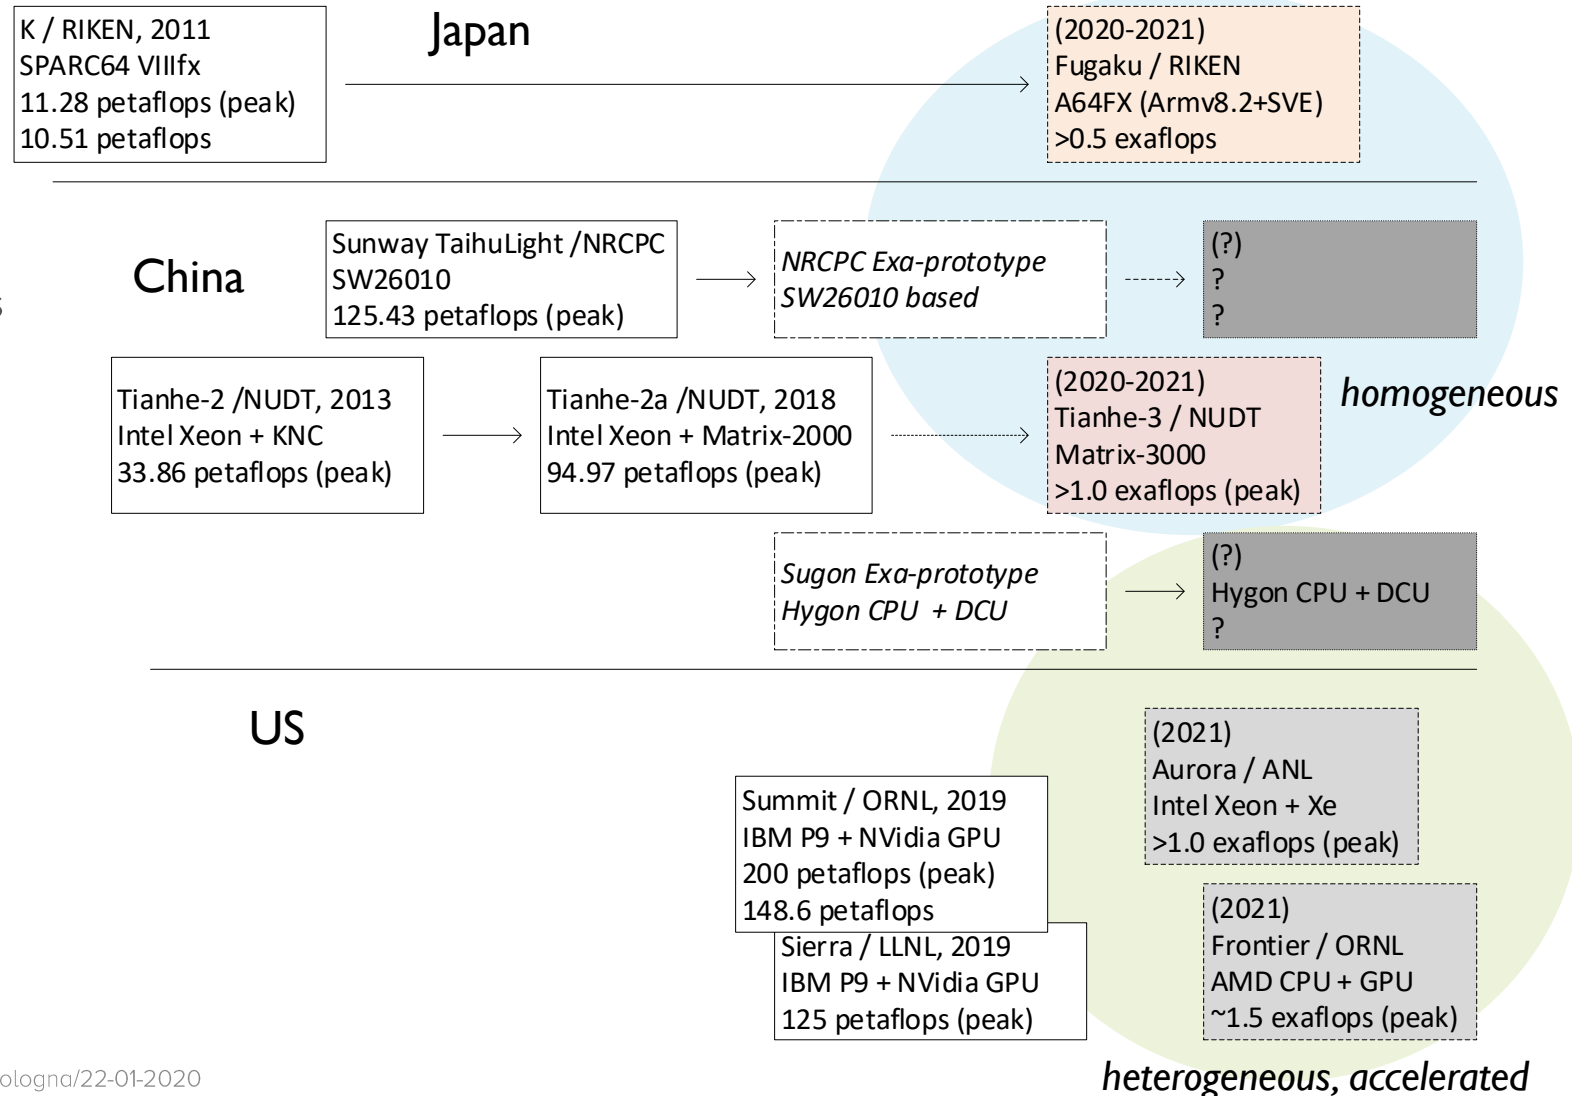
■ Processor design & technology:

Chip	Design	Manuf.
IBM POWER9		
NVIDIA Volta GV100		
Sunway SW26010		

How to bring back
Europe into processor
race ?

RACE TO EXASCALE

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



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

**SGAI
(EPI Phase I)**

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

Coordinator: ATOS

27 partners, 10 countries

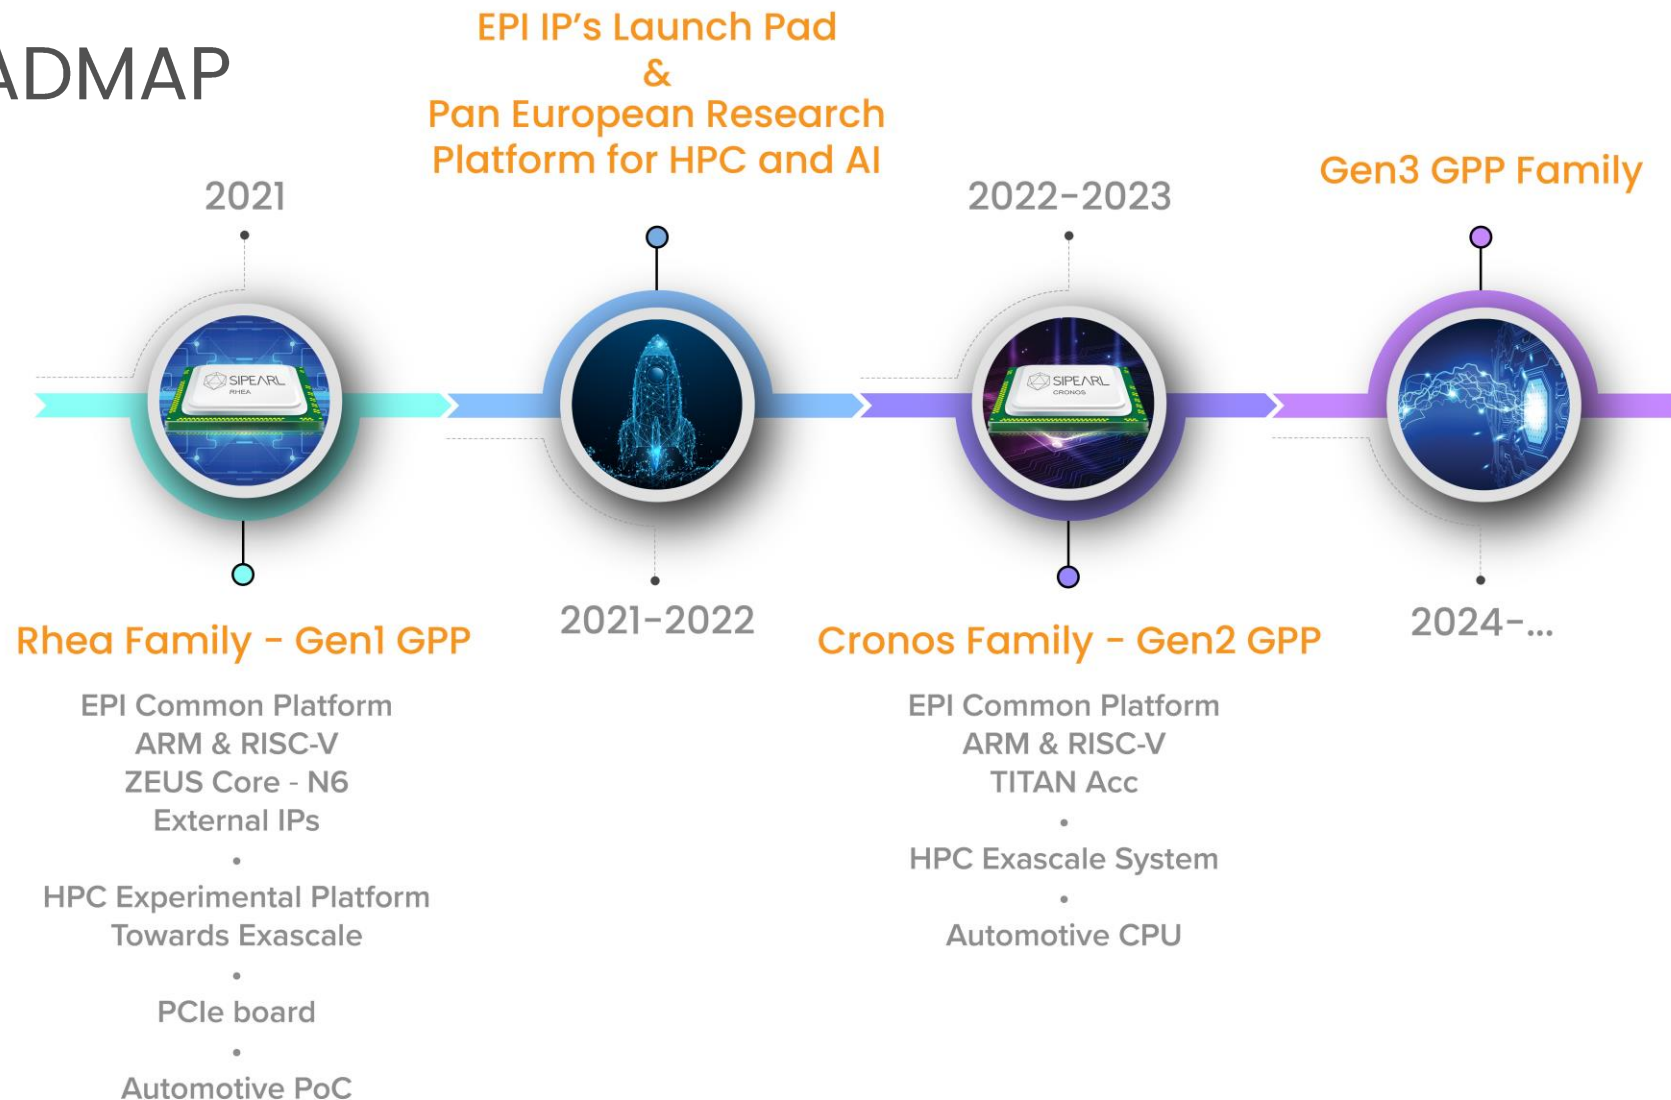
Budget: 80 M€



EPI TECHNOLOGIES

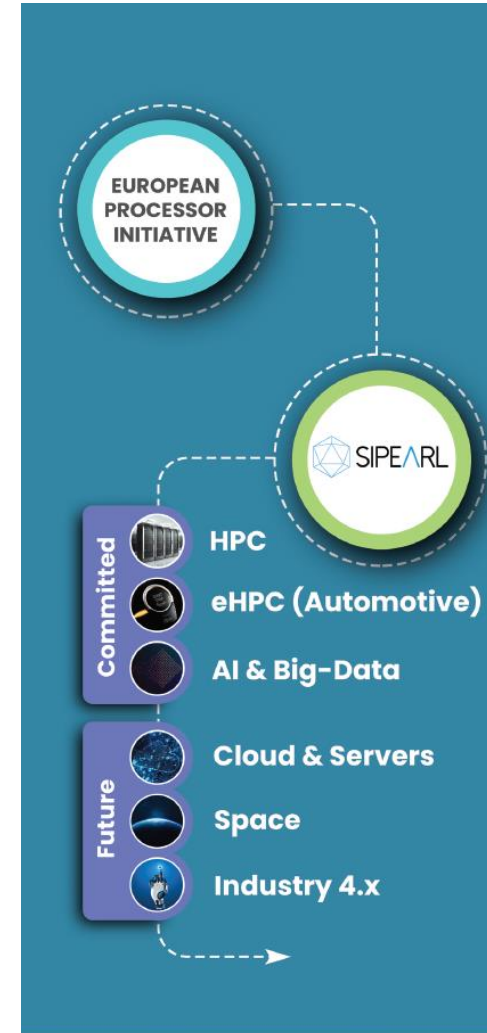
- Energy Efficiency
 - Adopt Arm general-purpose CPU core with SVE / vector acceleration in the first EPI chip
 - Develop power management solutions for the EPI chip
 - Develop acceleration technologies based on RISC-V for better DP GFLOPS/Watt performance
 - Inclusion of MPPA for real-time application 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





EPI COMMON PLATFORM & PROCESSOR

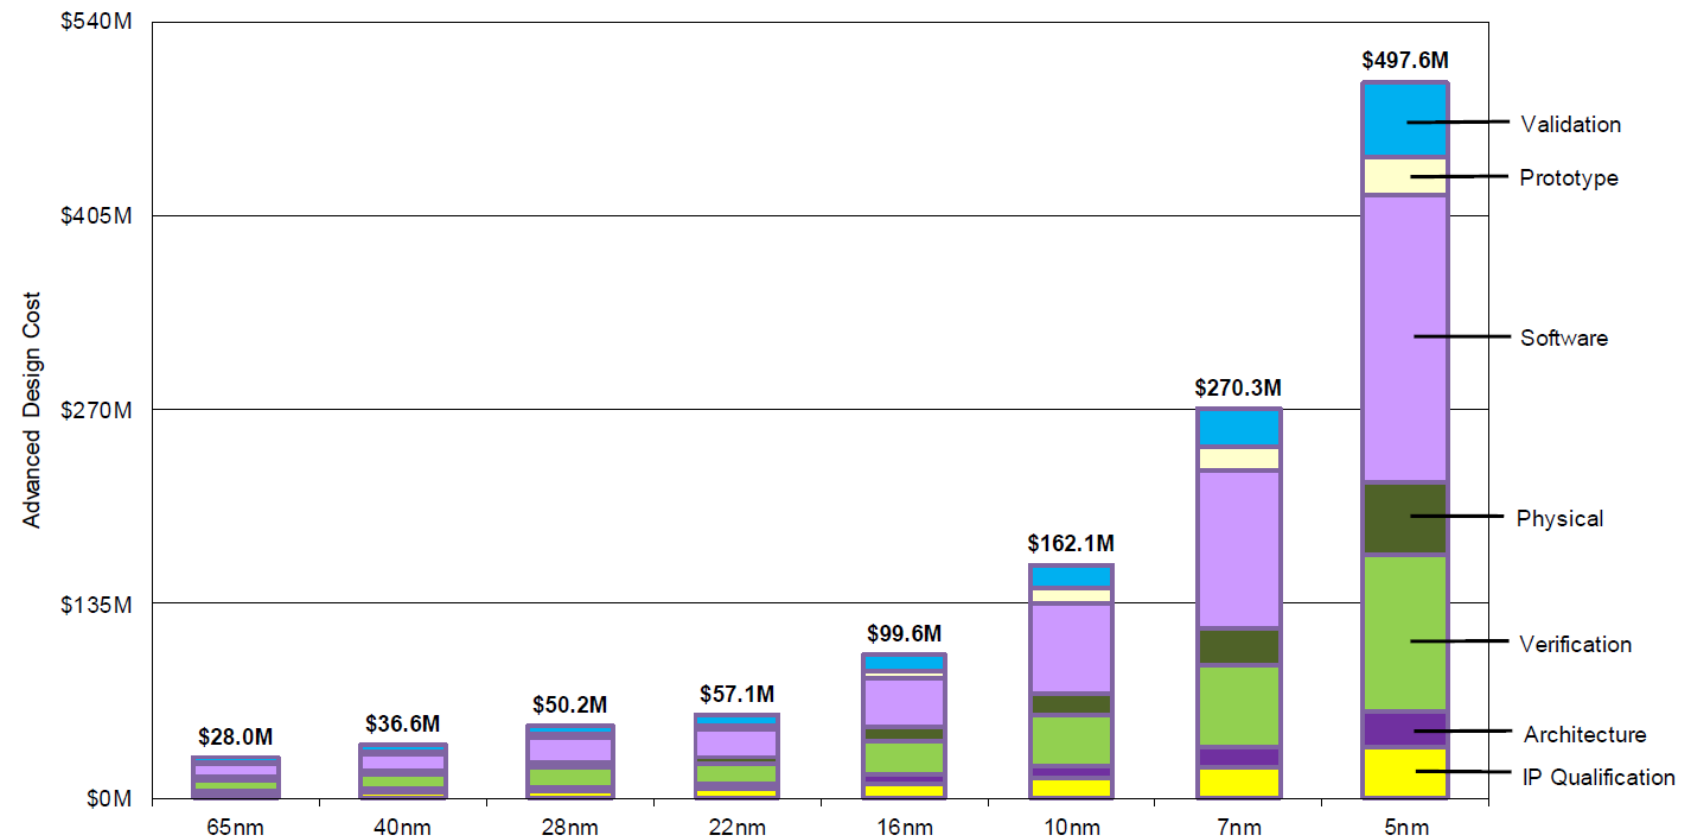
DENIS DUTOIT (CEA)

WHY DO WE NEED A COMMON PLATFORM ?

- Specialization is key for HPC to reach exascale.
- Specialization requires several advanced chip designs: generic processor, accelerator...
- Cost of advanced designs reaches \$500M.
- Very high value from defining a **design methodology** that gives reductions in design costs and eases incremental roadmap implementation.

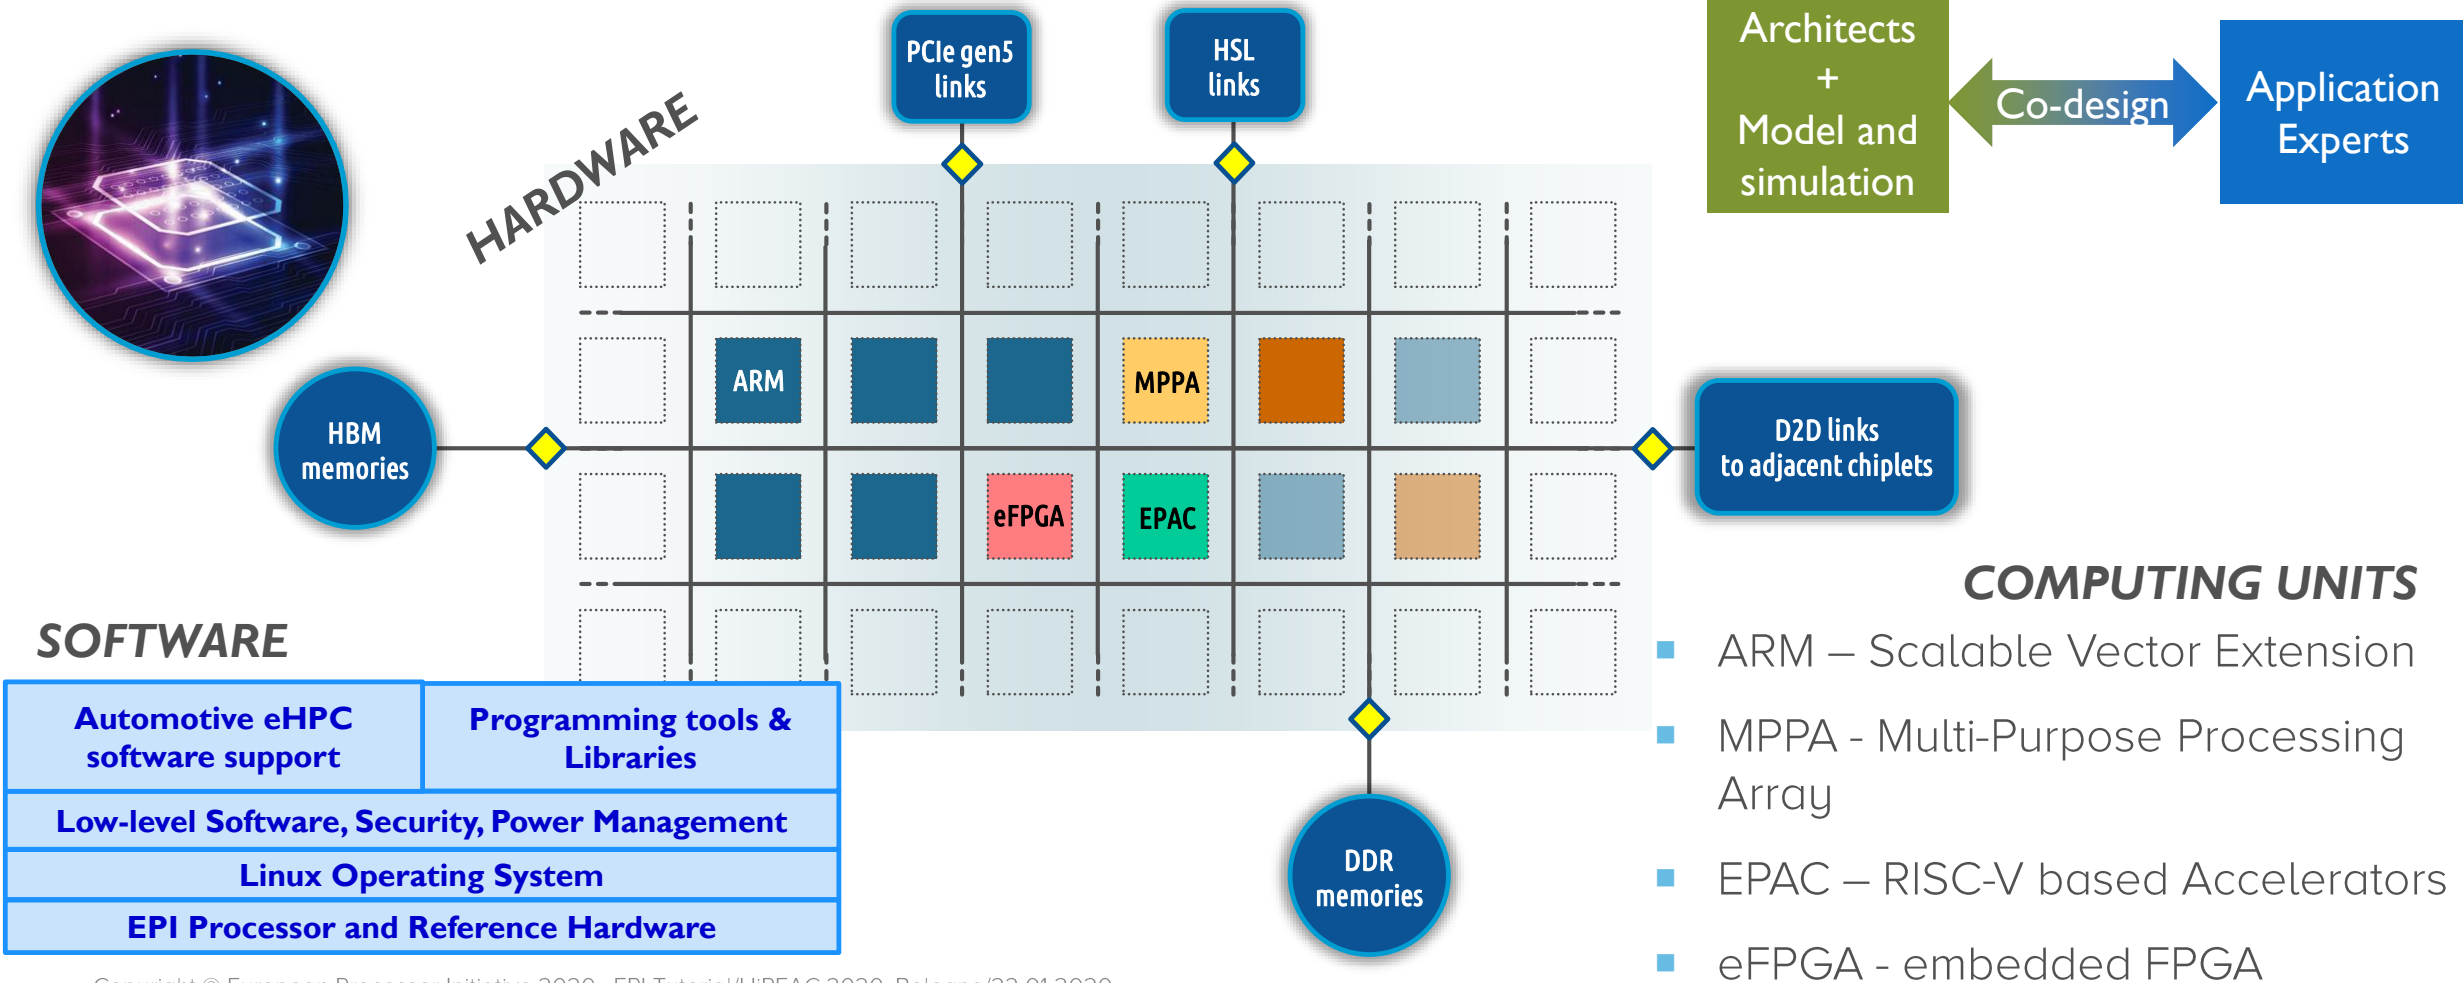
➤ EPI Common Platform

Cost of Advanced Designs



Source IBS, July 2019

COMMON PLATFORM TO HARMONIZE THE HETEROGENEOUS COMPUTING ENVIRONMENT



HETEROGENEOUS COMPUTING UNITS

*Binary compatible scalable performance implementations
and leveraging state-of-art designs*

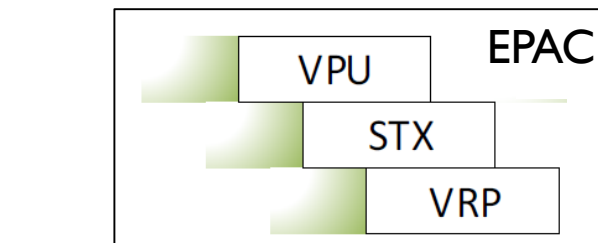
General Purpose



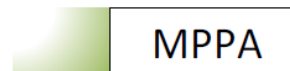
arm

Specialization

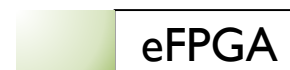
Common platform to enable application specializations



 **RISC-V®**



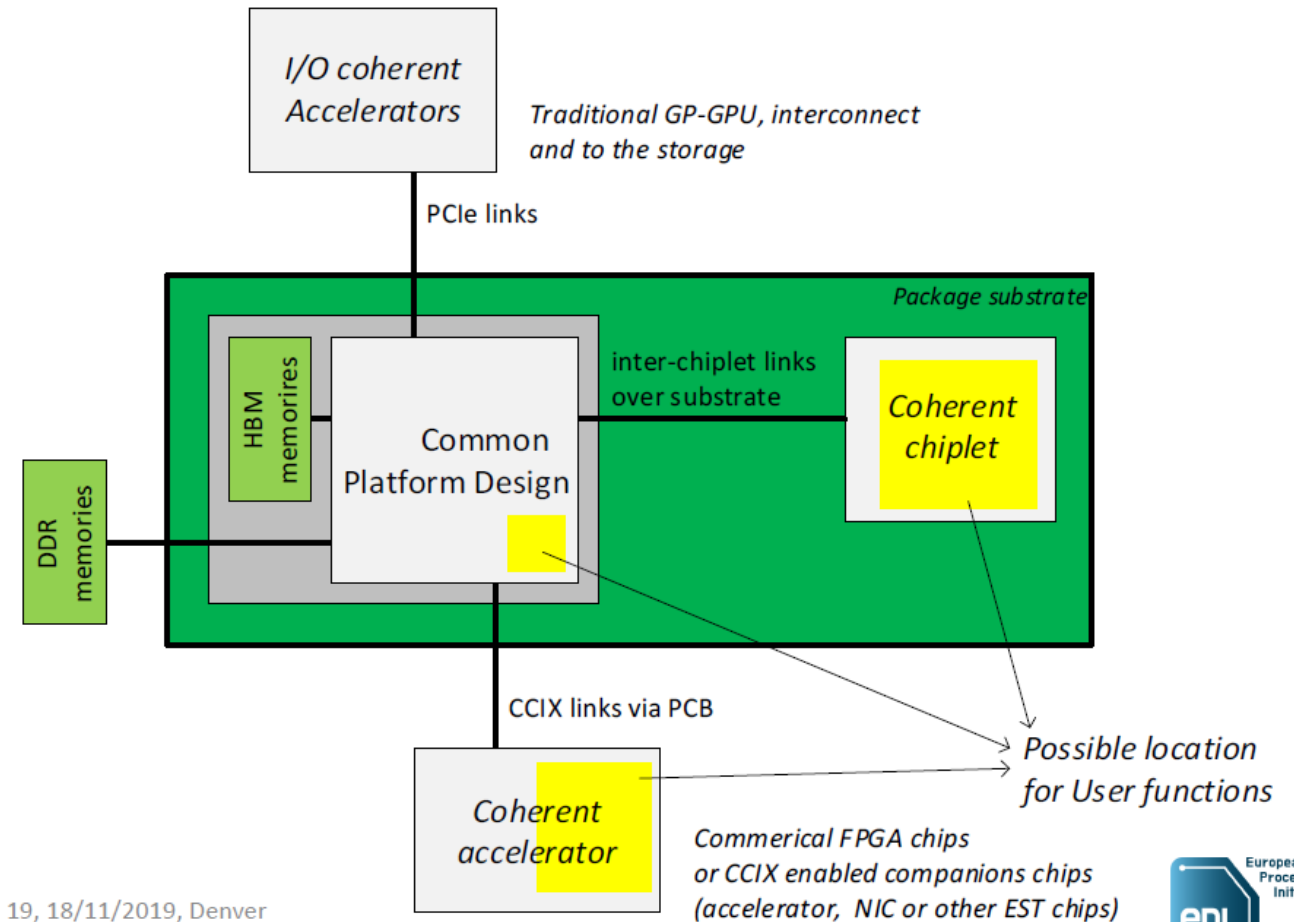
 **KALRAY**




menta

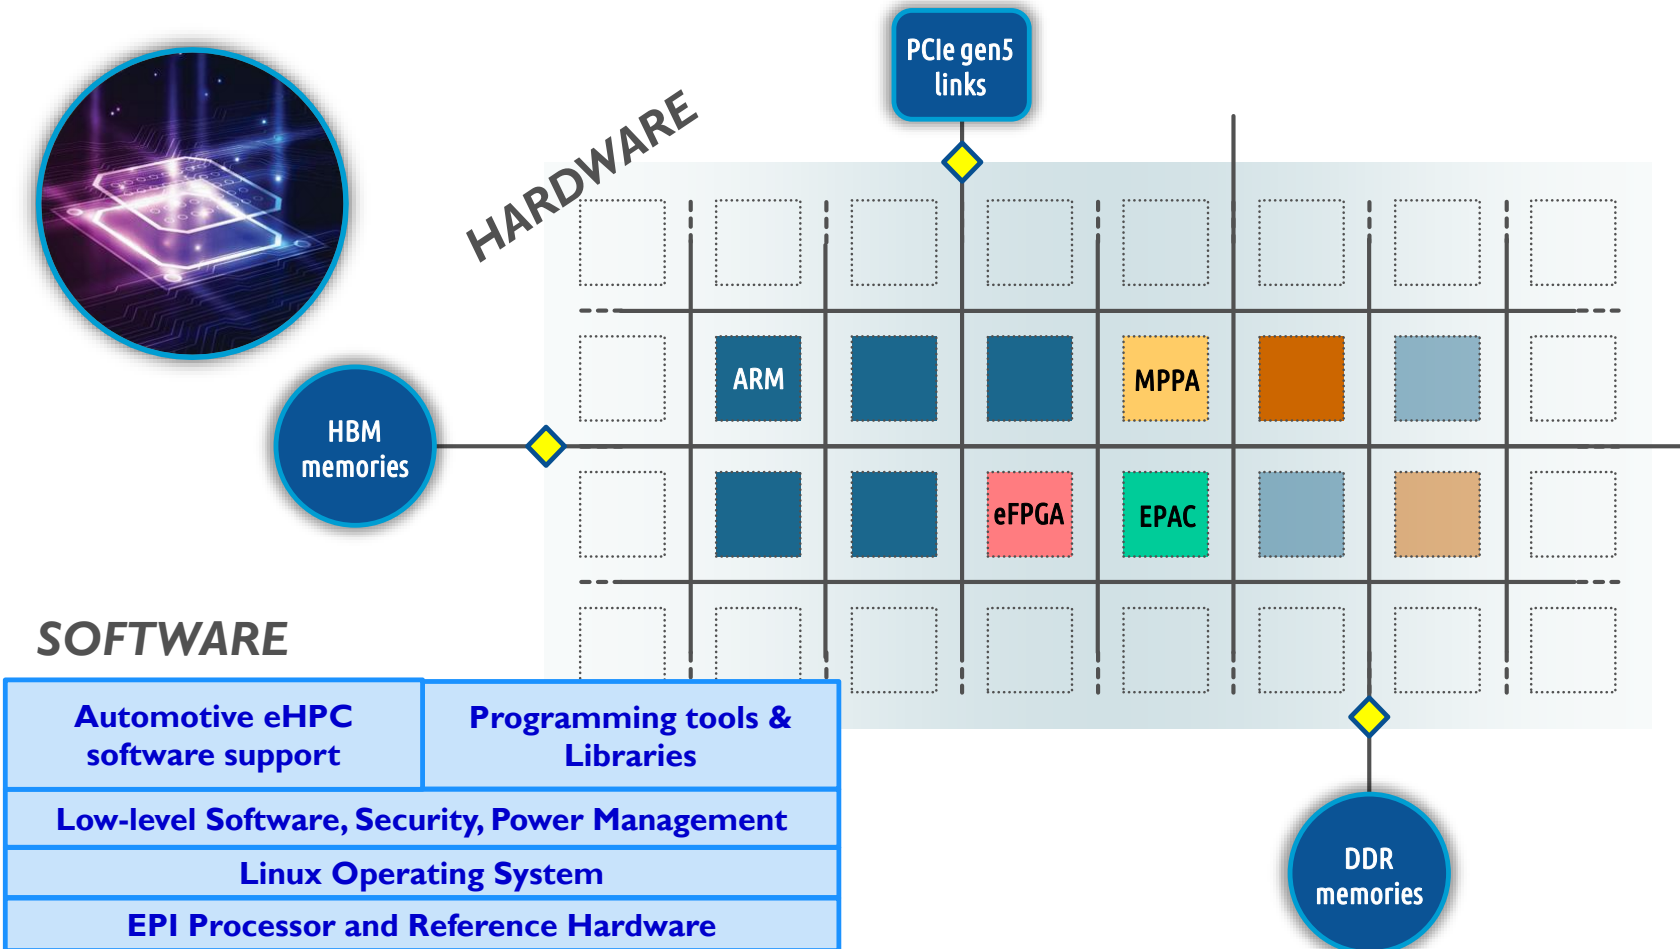
HETEROGENEOUS INTEGRATION

- Integrating customized functions at different level



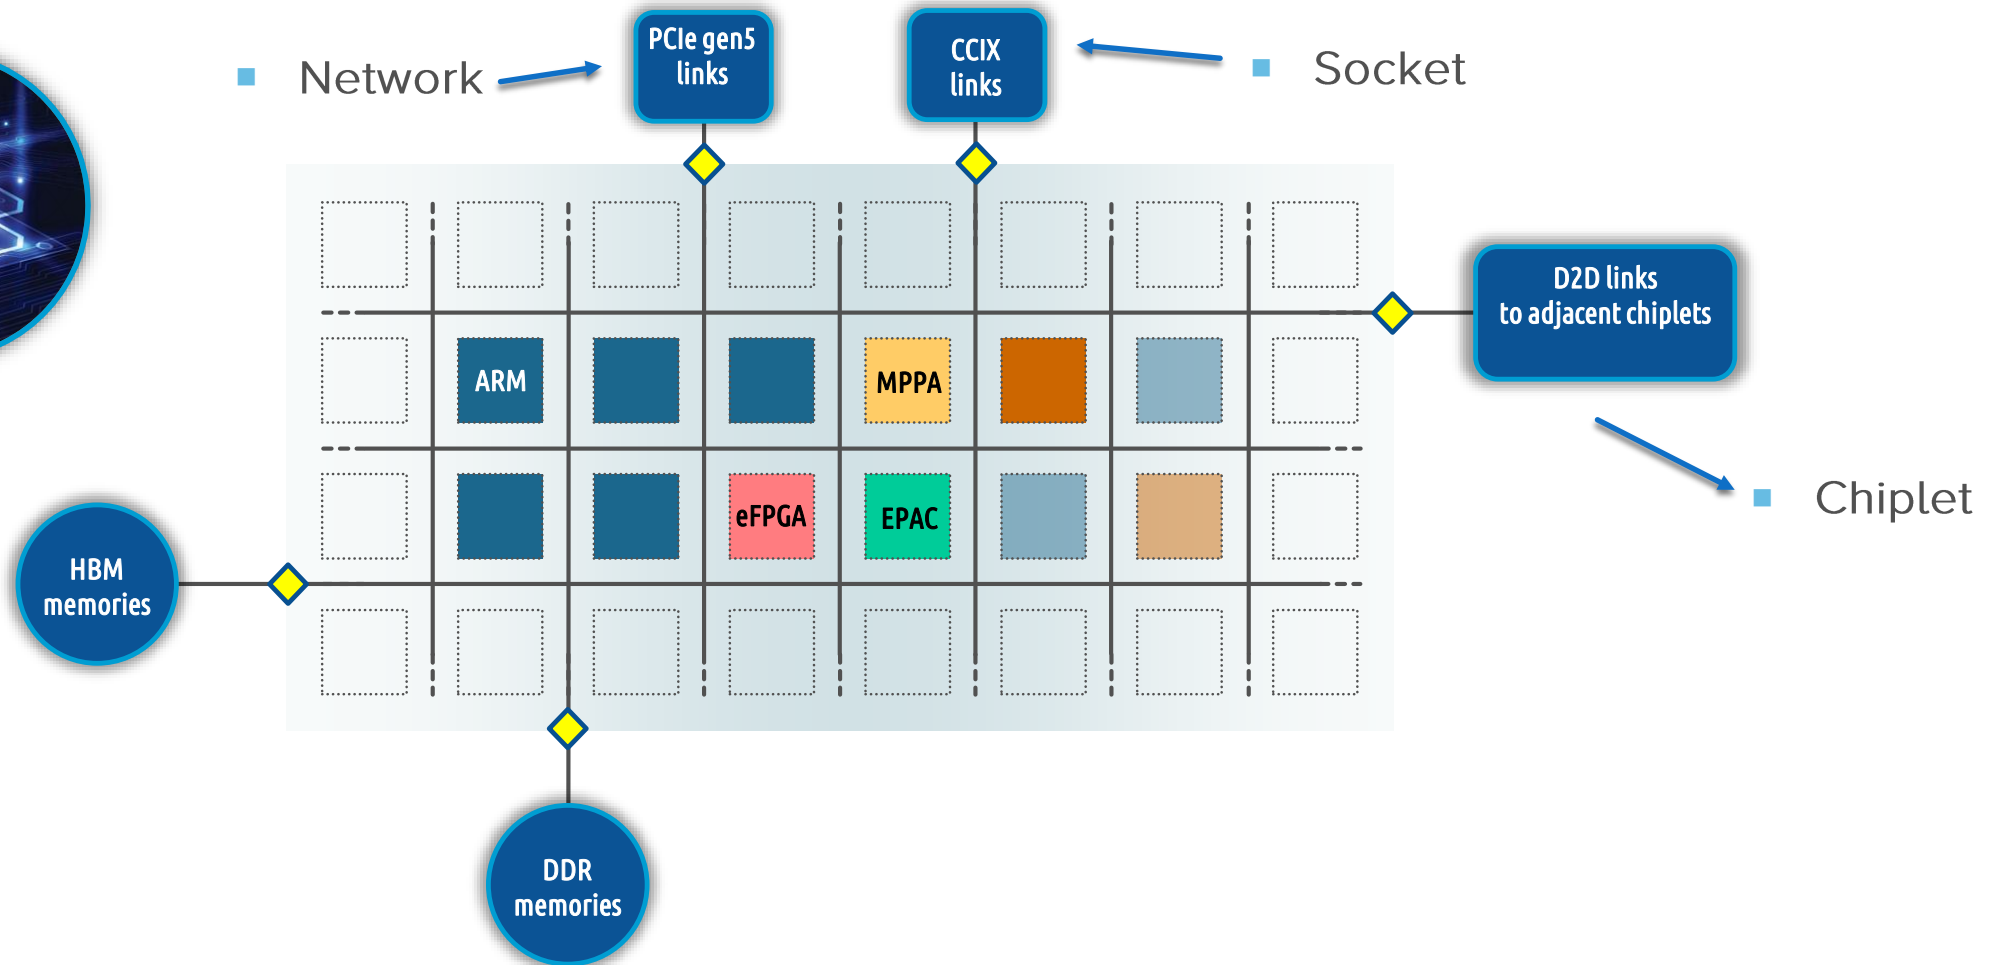
19, 18/11/2019, Denver

ON-CHIP HETEROGENEOUS INTEGRATION

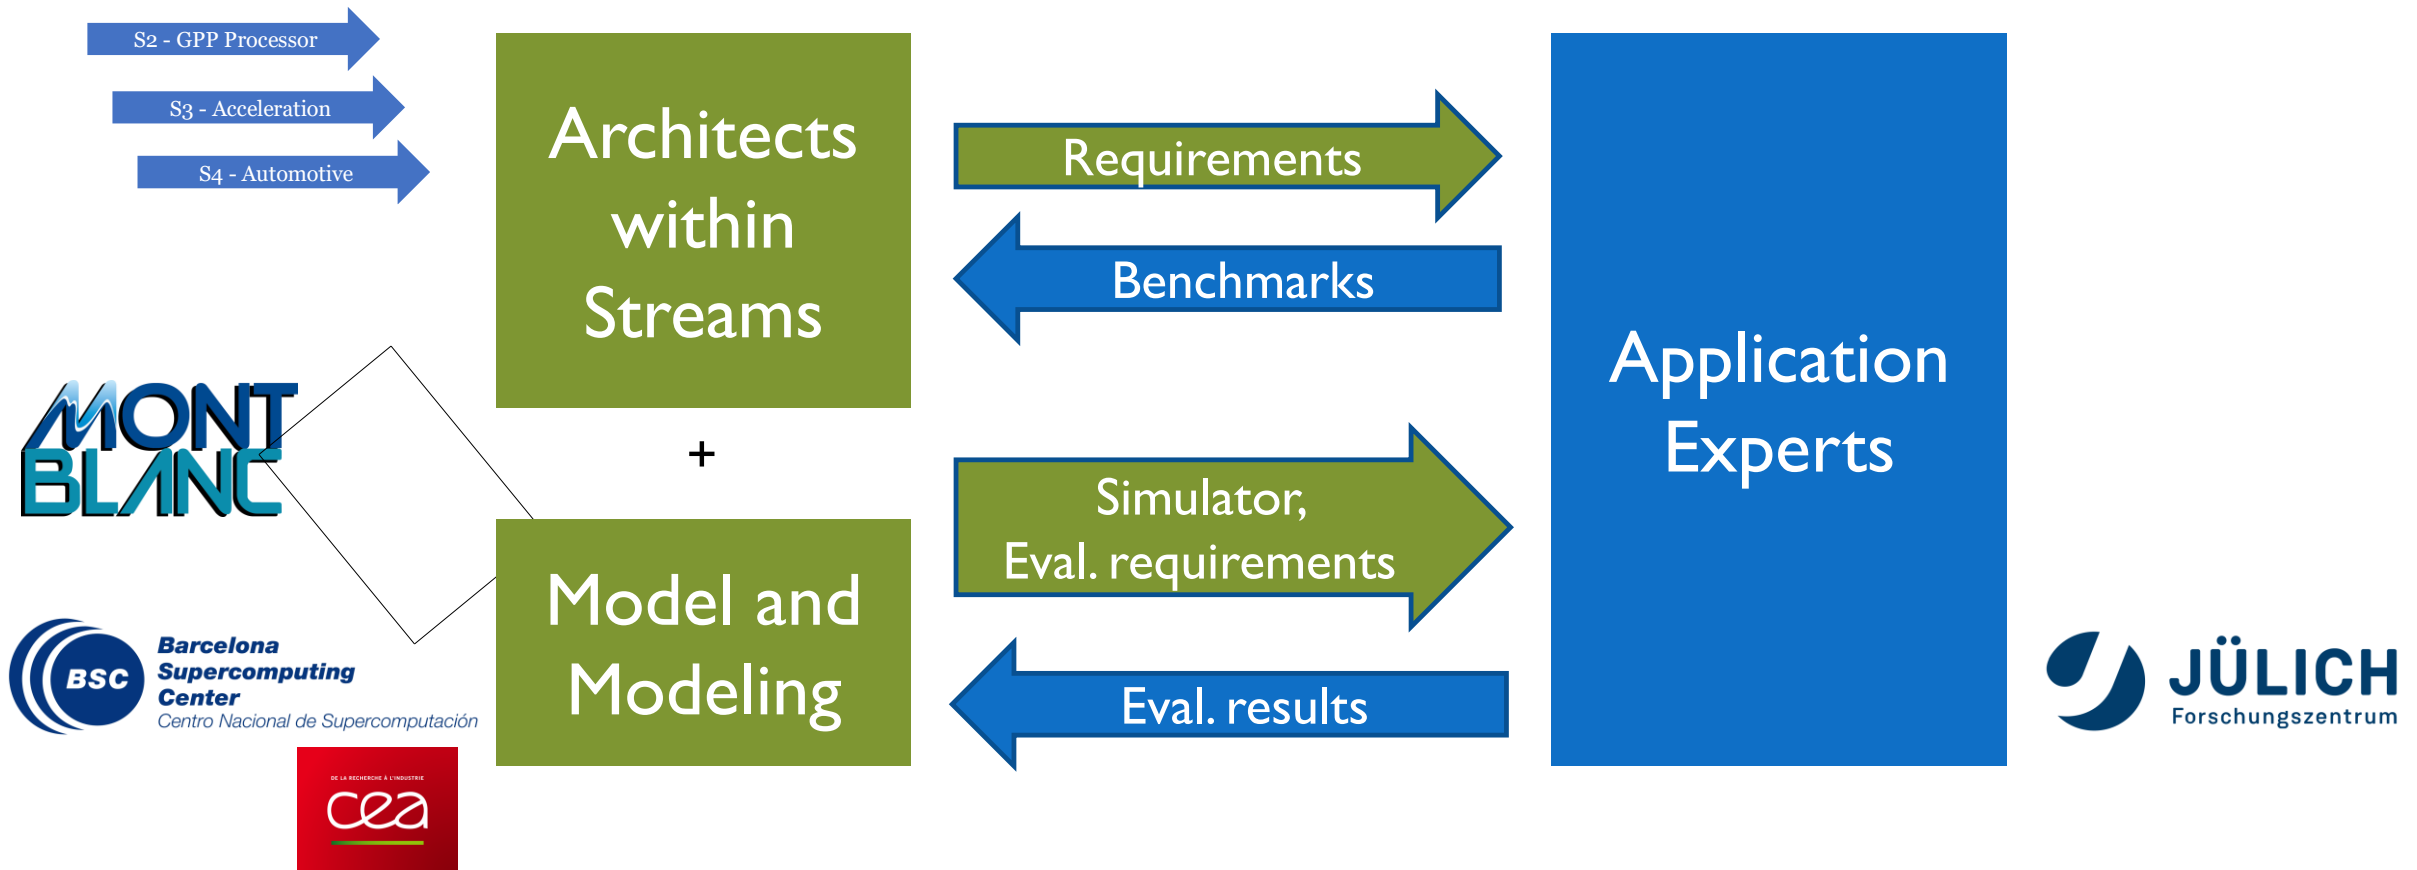


- 2D-mesh Network-on-Chip (NoC) to connect computing units: Arm, EPAC, MPPA, eFPGA.
- Common software environment between heterogeneous computing tiles to harmonize their integration with the external environment such as memories (DDR, HBM) and loosely coupled accelerators (through PCIe).

OFF-CHIP HETEROGENEOUS INTEGRATION

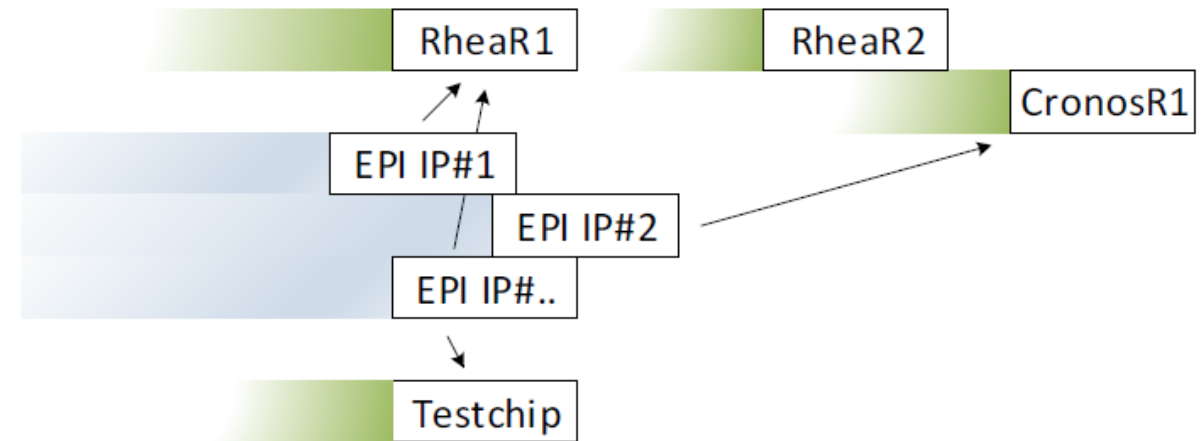


EPI CO-DESIGN



RHEA PROCESSOR

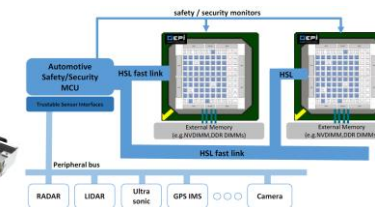
- Rhea is the first EPI General Purpose Processor
- Rhea targets HPC application
- Rhea is the first « instantiation » of EPI Common Platform
- Rhea design is led by SiPEARL (the EPI fabless company) and joint-developed by EPI partners.
- Rhea chip will be integrated into test platforms in order to validate the hardware units, develop the software, and run applications.



HPC blade



Automotive PoC



PCIe daughter card

SEMICONDUCTOR TECHNOLOGY NODE

Market Demand by Node

28/22-nm: 25% of Presales

MARKETS

- Digital Home
 - Set Top Box
 - Wireless Combo
 - Personal Assistants
- Internet of Things
 - Wireless MCU
 - Industrial
 - Health
 - Surveillance
 - Smart Home

16/14-nm: 30% of Presales

MARKETS

- Mobile
 - Mainstream mobile
 - Augmented Reality
 - Drones
- Data Center
 - Storage
- Automotive
 - V2X & Infotainment
- Digital Home
 - Smart TV/Home
 - Gateways
 - Surveillance
 - STB

7/5-nm: 20% of Presales

MARKETS

- Data Center
 - Servers
 - Cryptocurrency
 - SAN/NAS (Enterprise)
 - AI Accelerators (Training/Inference)
- Mobile
 - Mobile AP
 - Mobile BB
 - Network Routers
- Automotive
 - High-end ADAS

RHEA DESIGN

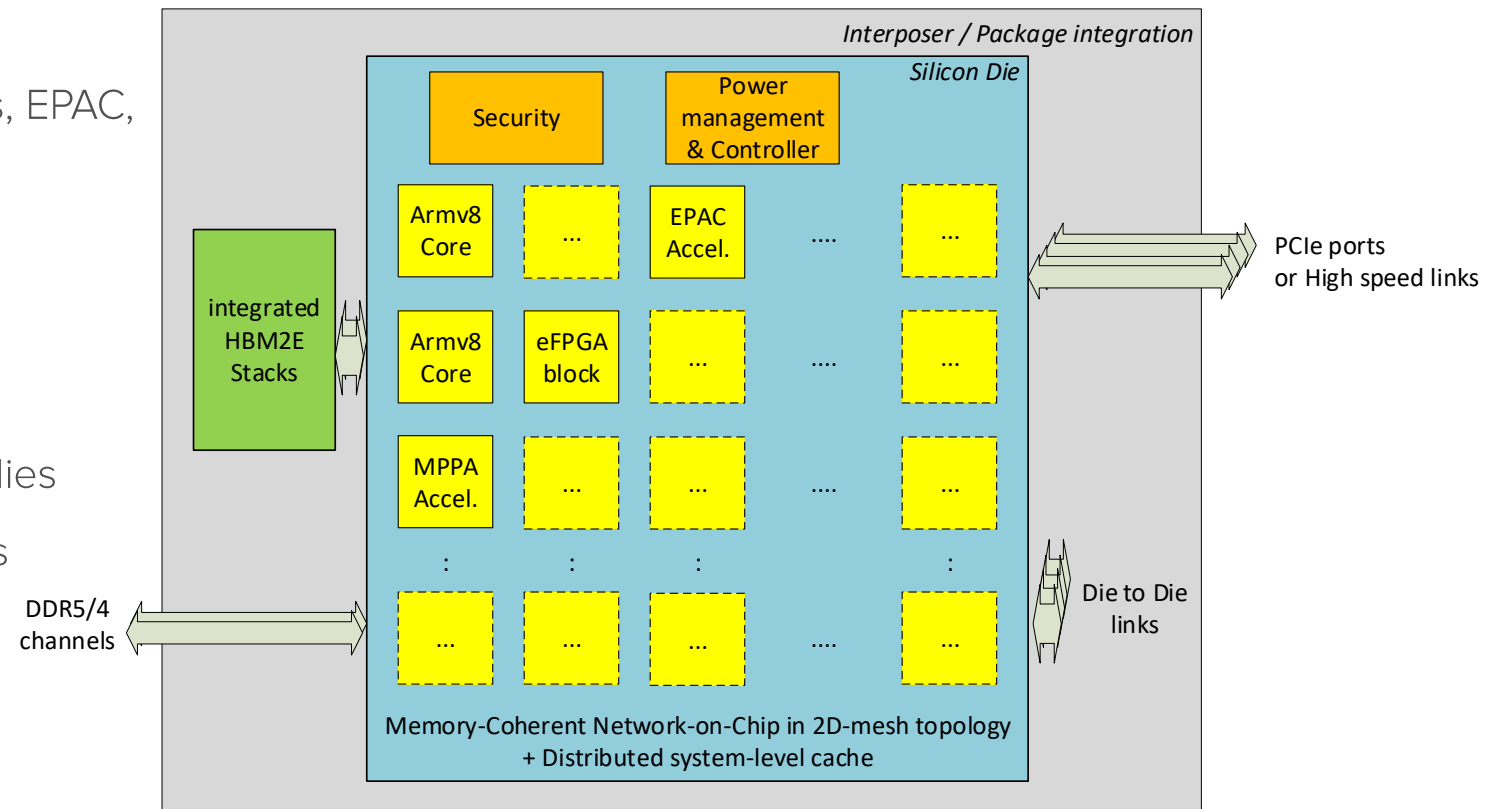
- Generic processing multi-core backbone:
 - Multi-core Arm Zeus processor with SVE engines for pre-ExaScale level generic processing.
 - Coherent NoC with distributed system level cache to keep the data local.
- Prototypes of High energy-efficient accelerator tiles:
 - RISC-V based acceleration (EPAC) 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

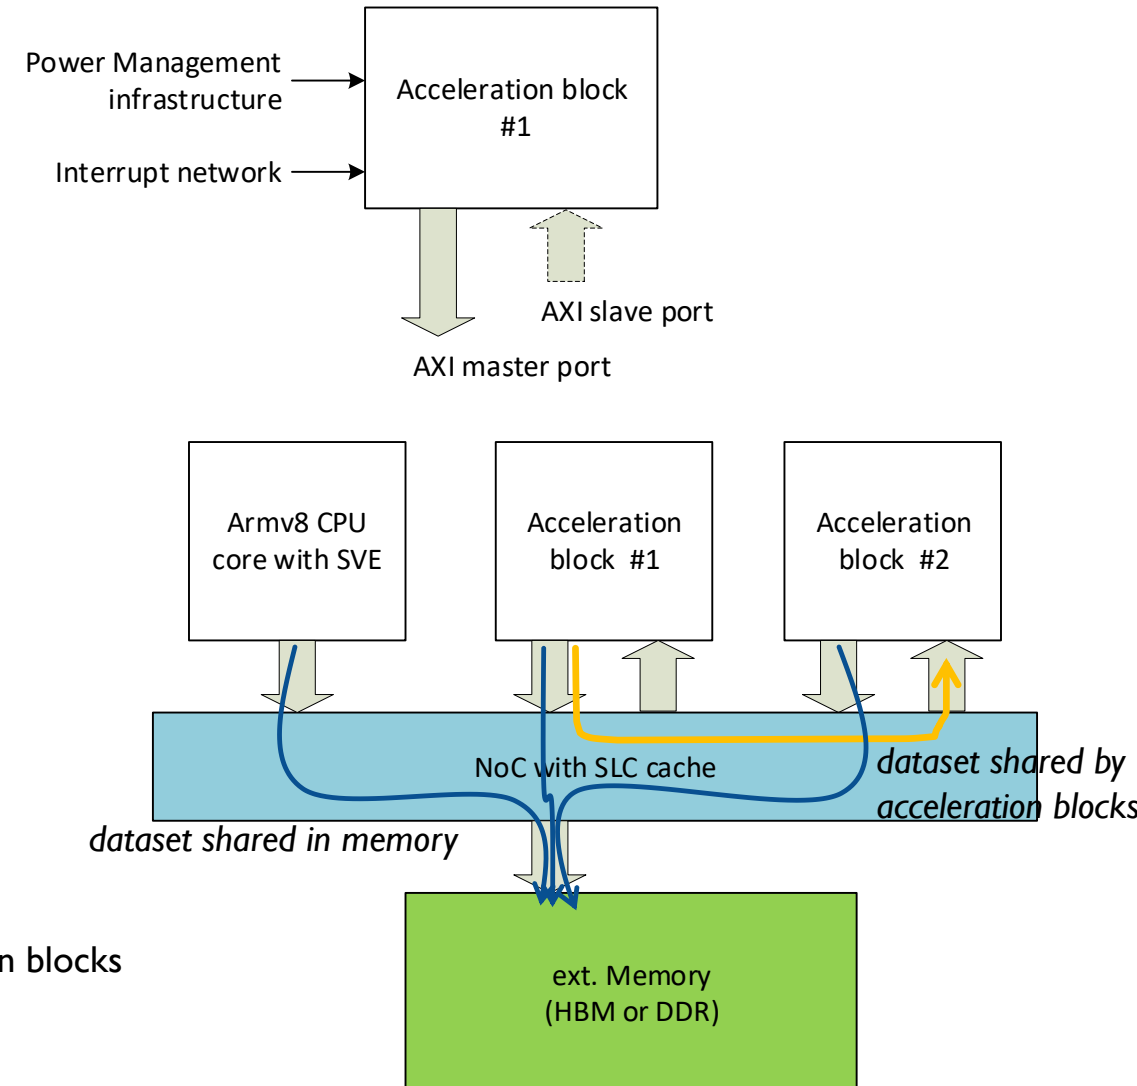
NoC: network on chip

HSL: High speed links (with memory coherent support)



INTEGRATION OF ACCELERATORS

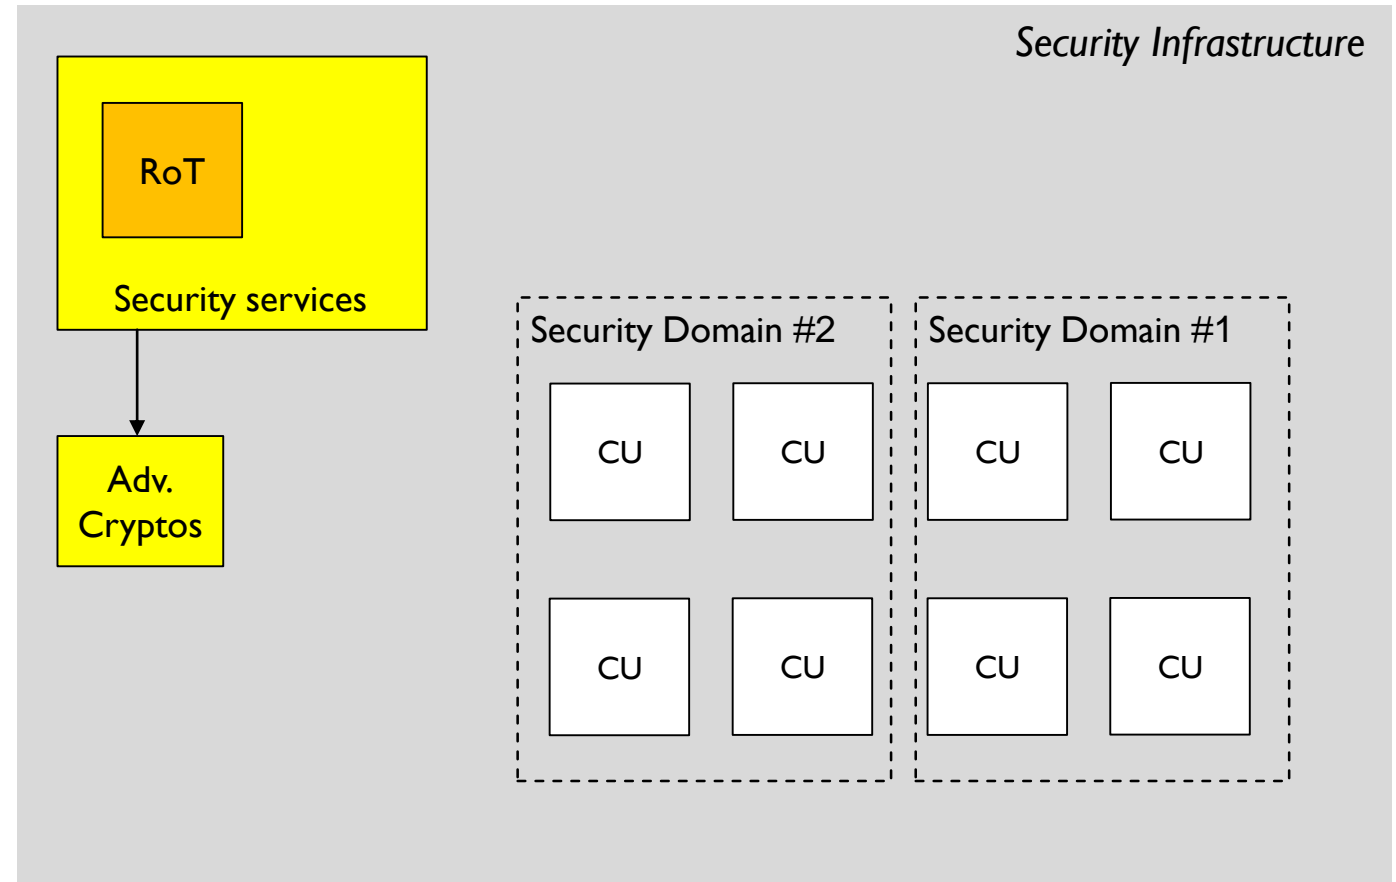
- Interfaces to connect acceleration functions to the NoC
 - Data access and sharing through AXI ports
 - Receiving interrupt
 - Power management
- Enable memory-centric computations
 - Same copy of dataset is shared by multiple CUs
 - In the ext. memory (DDR or HBM) cached by SLC cache
 - In the local scratch memories near or local the acceleration blocks
 - System MMU to provide same virtual memory view



CU: Computing Unit; either Arm v8 core with SVE or the EPAC/MPPA acceleration blocks
SLC: System Level Cache; a last-level cache before ext. memories

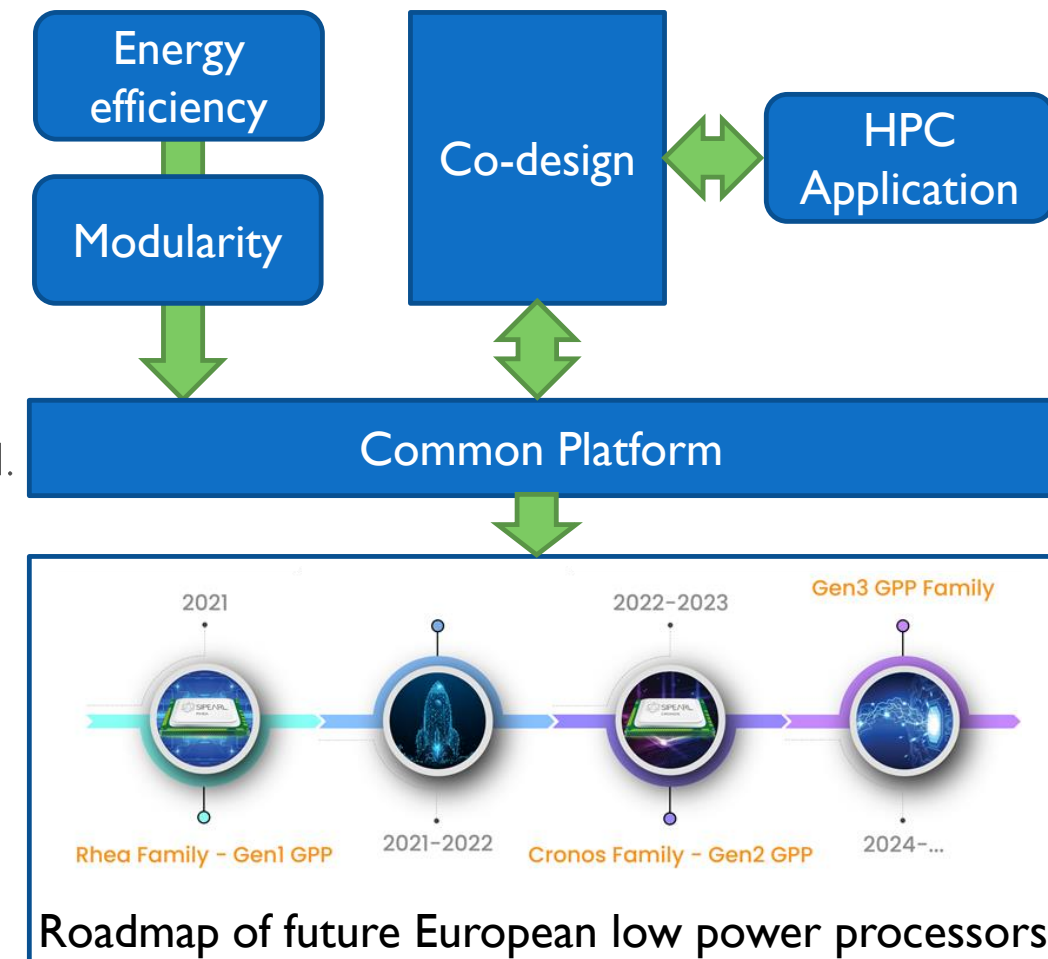
SECURITY IN EPI PROCESSOR

- Security aspects taken into account since the start of architecture definition
- Root of trust
- Secure Boot
- Security domains
- Security services isolated
- Advanced cryptographic functions
- Various monitors for fault-injection, physical intrusion and other conditions



EPI TODAY AND RHEA DESIGN STATUS

- The project finished its first year.
- Updated roadmap.
- Presence at key events.
- Introduction of the EPI Common Platform.
- Co-design process defined.
- The first-generation processor family, named Rhea, specified.
- Silicon process and advanced package integration selected: 6nm and 2.5D.
- Individual IP (EPAC, MPPA, eFPGA, power management, ...) design on track: some first RTL revision delivered.
- Top-level integration will start soon.





THANK YOU