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

CO-LOCATED WITH THE ACM 2019 SUMMER SCHOOL ON HPC ARCHITECTURES FOR AI AND DEDICATED APPLICATIONS

UNIVERSITAT POLITÈCNICA DE CATALUNYA, BARCELONA, SPAIN

17 JULY 2019

European Processor Initiative

Association for Computing Machinery

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

Start	Finish	Topic	Presenter
14:00	14:25	**HPC processor landscape**	Andrea Bartolini, UNIBO
		Main challenges for HPC processor	
		Architecture evolution towards heterogeneity	
		Semiconductor technology overview	
		European landscape and introduction to EuroHPC	
14:25		**European Processor Initiative (EPI)**	
14:25	14:35	Overview	Andrea Bartolini, UNIBO
14:35	14:50	Processor and general architecture	Andrea Bartolini, UNIBO
14:50	15:00	Co-design process and modeling	Andrea Bartolini, UNIBO
15:00	15:45	Accelerator	Mauro Olivieri, BSC
15:45	16:30	Software	Jesus Labarta, BSC
16:30	17:00	Automotive	Jaume Abella, BSC Francisco Cazorla, BSC
	17:00	End	

HPC PROCESSOR LANDSCAPE

ANDREA BARTOLINI (SLIDES PREPARED BY DENIS DUTOIT)

HPC PROCESSOR LANDSCAPE

Start	Finish	Topic	Presenter
14:00	14:25	**HPC processor landscape**	Andrea Bartolini, UNIBO
		Main challenges for HPC processor	Slides Prepared (Denis Dutoit)
		Architecture evolution towards heterogeneity	
		Semiconductor technology overview	
		European landscape and introduction to EuroHPC	

HIGH PERFORMANCE COMPUTING EVOLUTION

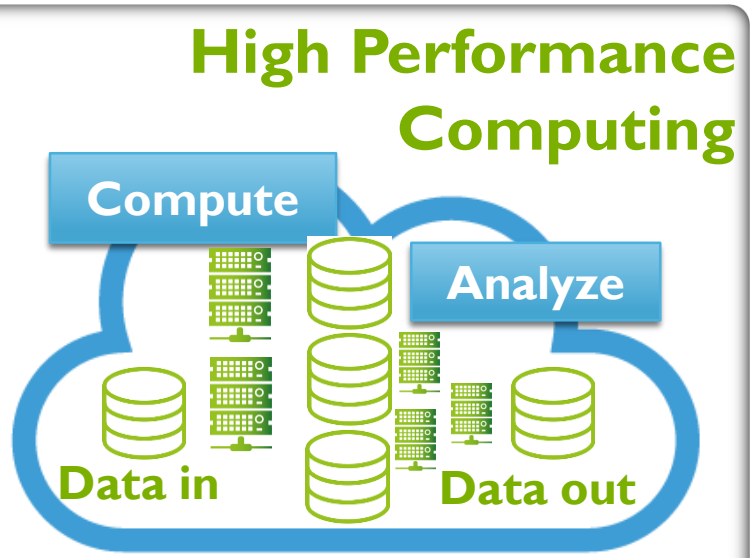

- **Starting from high performance compute only, HPC evolves towards:**
  - New workloads
  - Massive volume of data

New drivers	Requirements	Solutions
New workloads	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	Increased Bytes per Flops. High bandwidth/low latency access to all data.	***High Bandwidth Memories and 2.5D integration***

TERA1000 - CEA

CHALLENGES FOR ADVANCED COMPUTING

10x energy efficiency improvement every 4 years

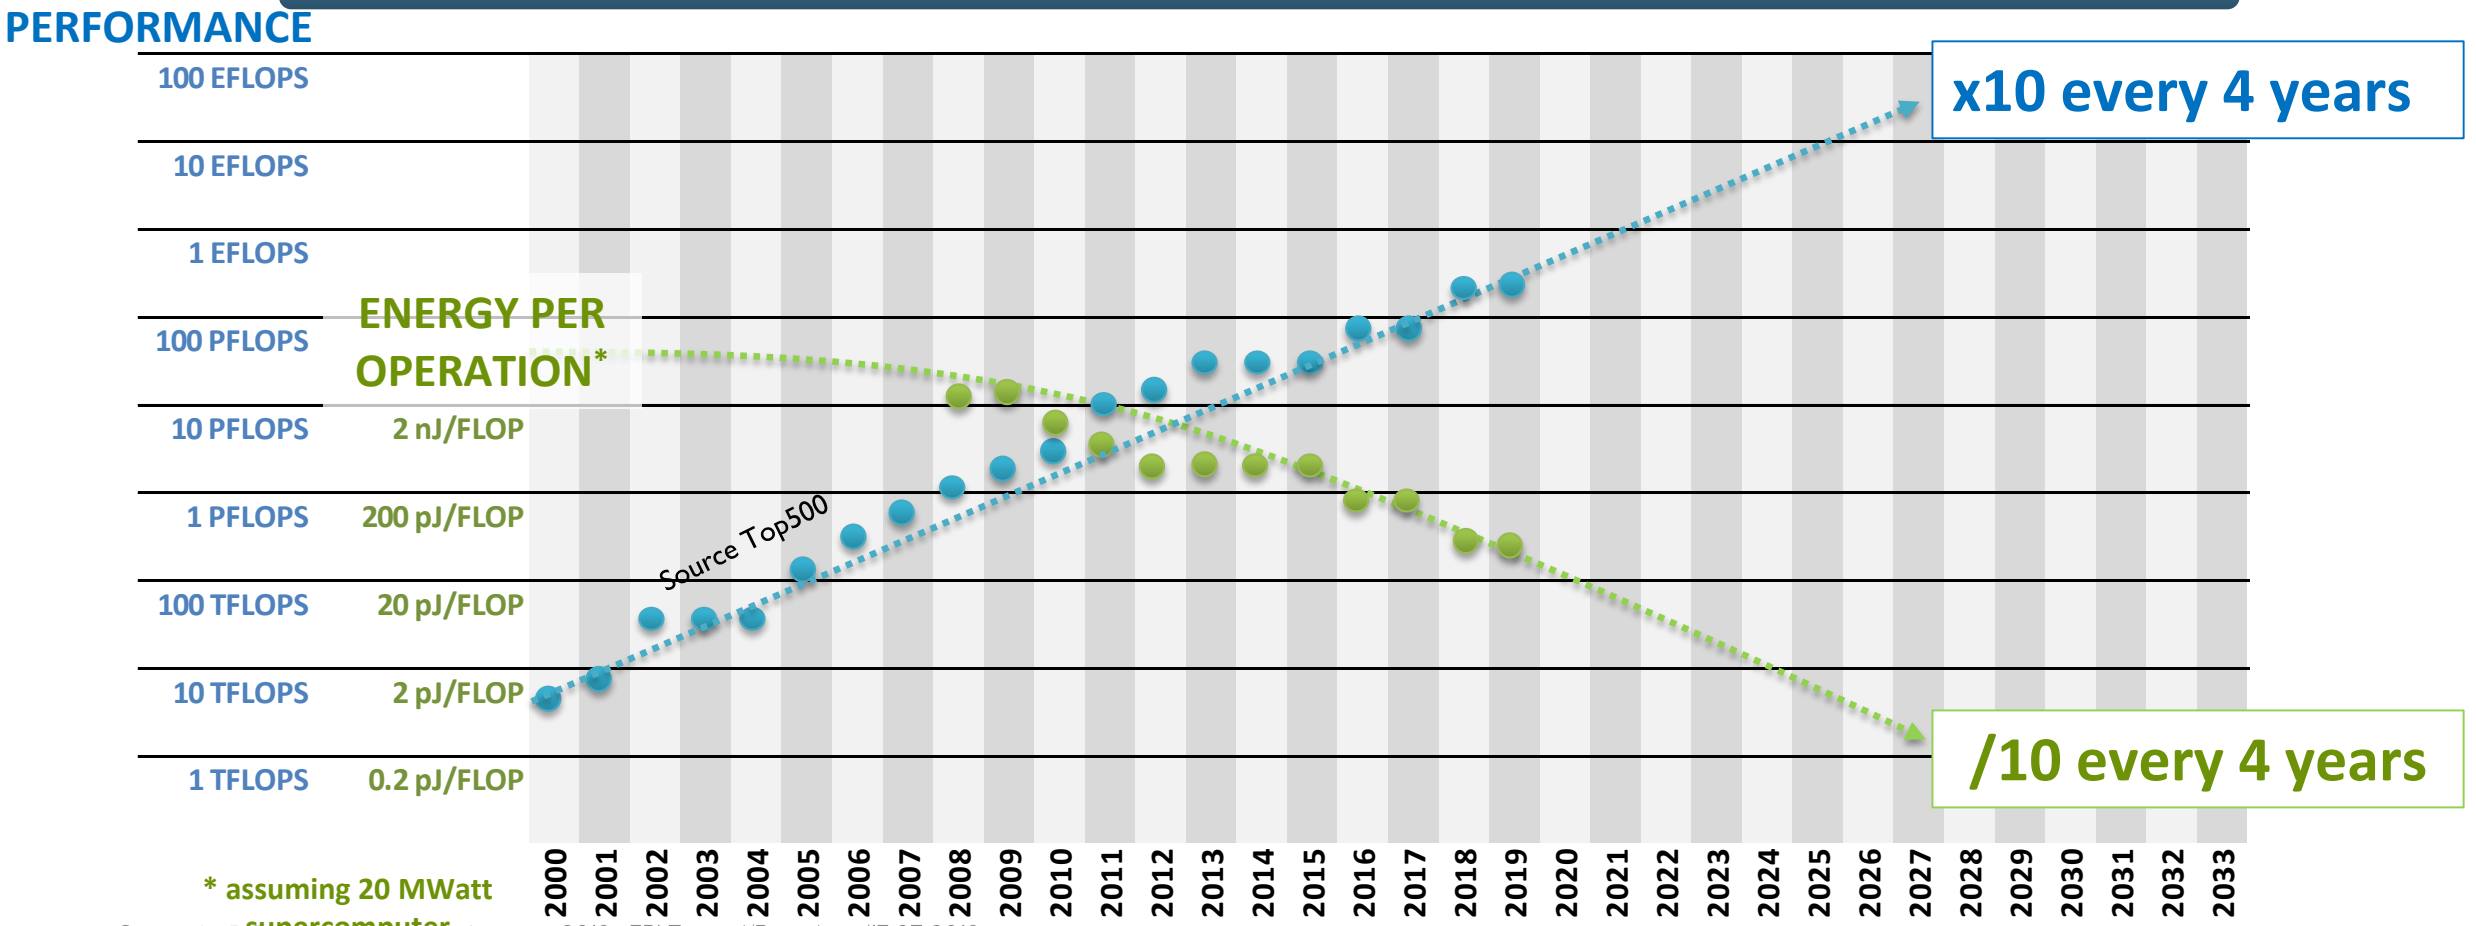

* assuming 20 MWatt supercomputer

TECHNOLOGY SCALING TRENDS

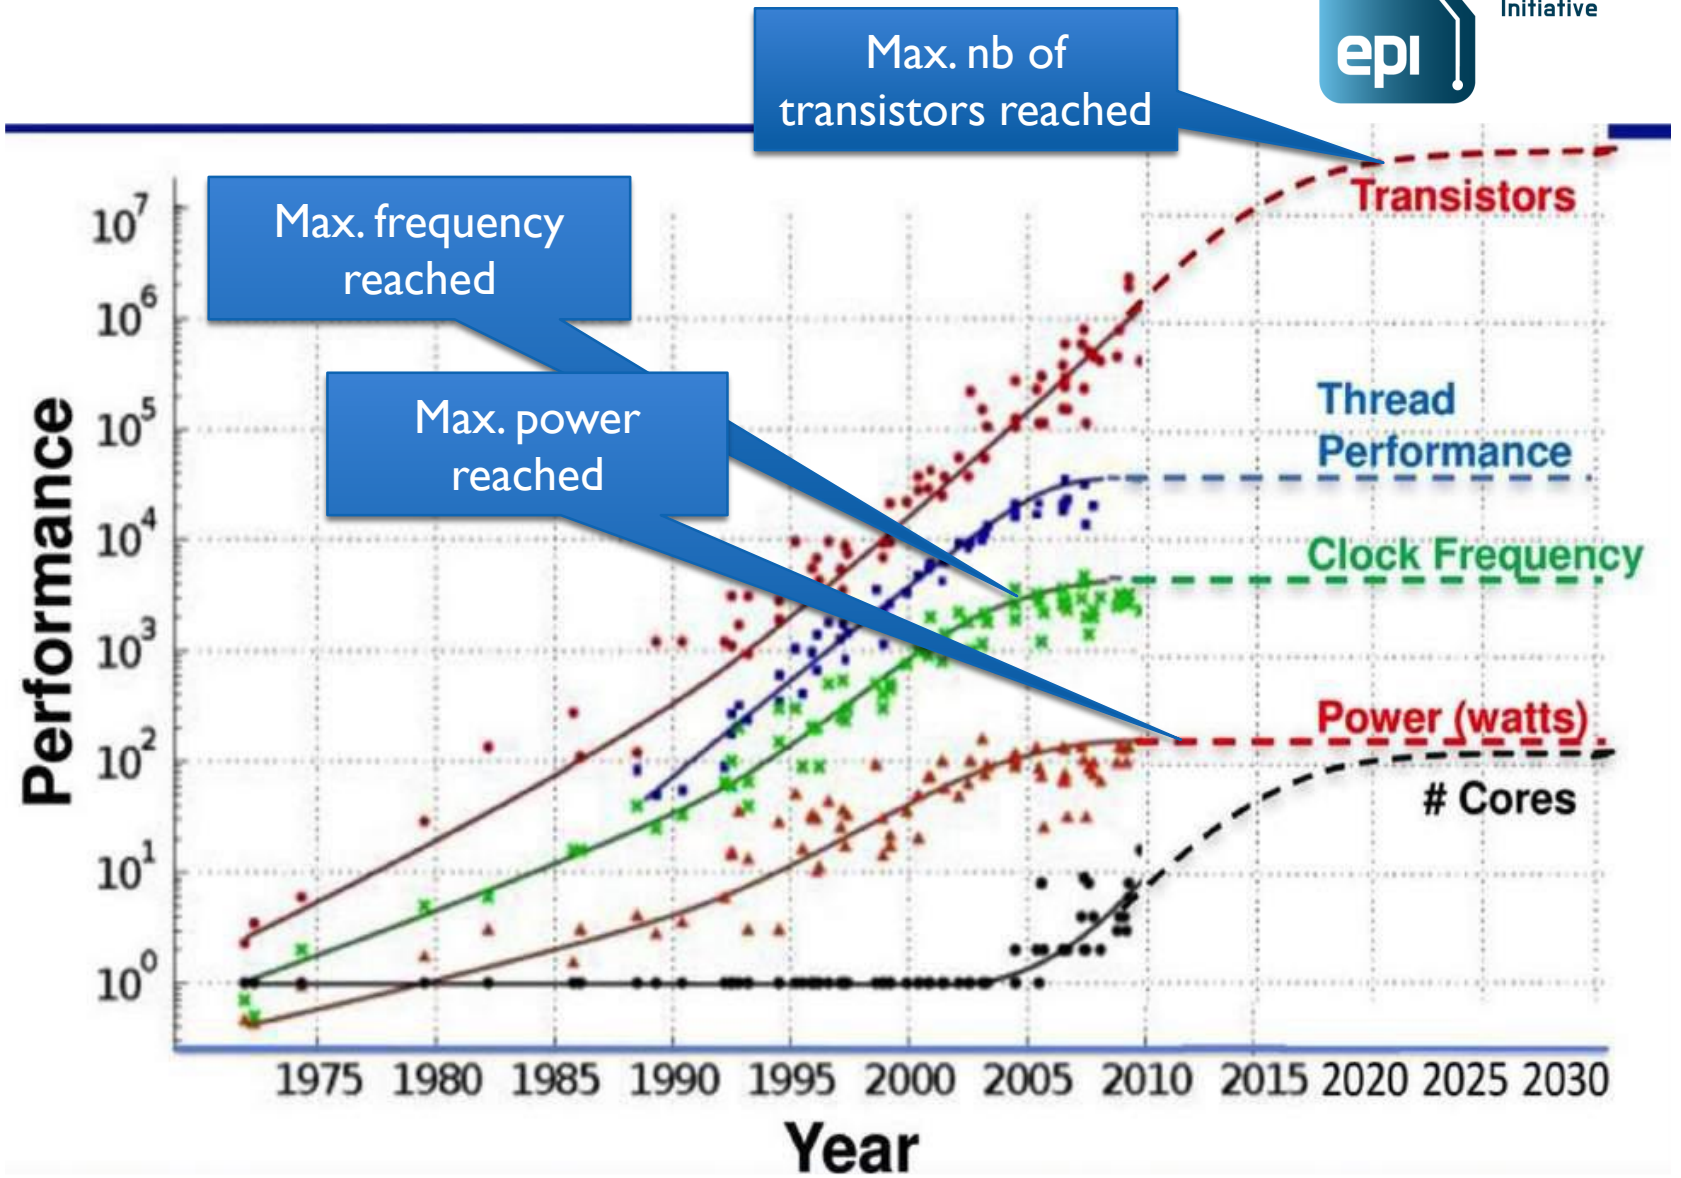

Source from Kunle Olukotun, Lance Hammond, Herb Sutter, Burton Smith, Chris Batten, and Krste Asanovic¸

HAPPY SCALING

Transistor nb ↗ ↗ ↗
Frequency ↗
Power density →

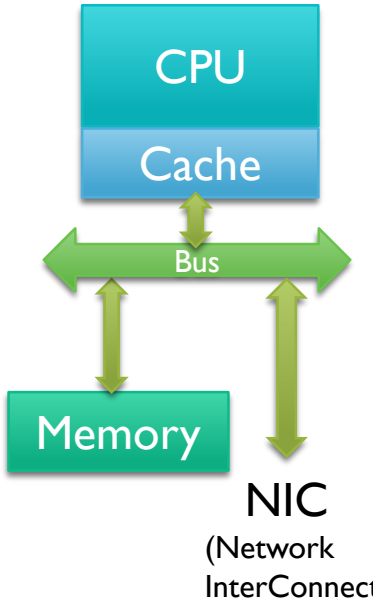

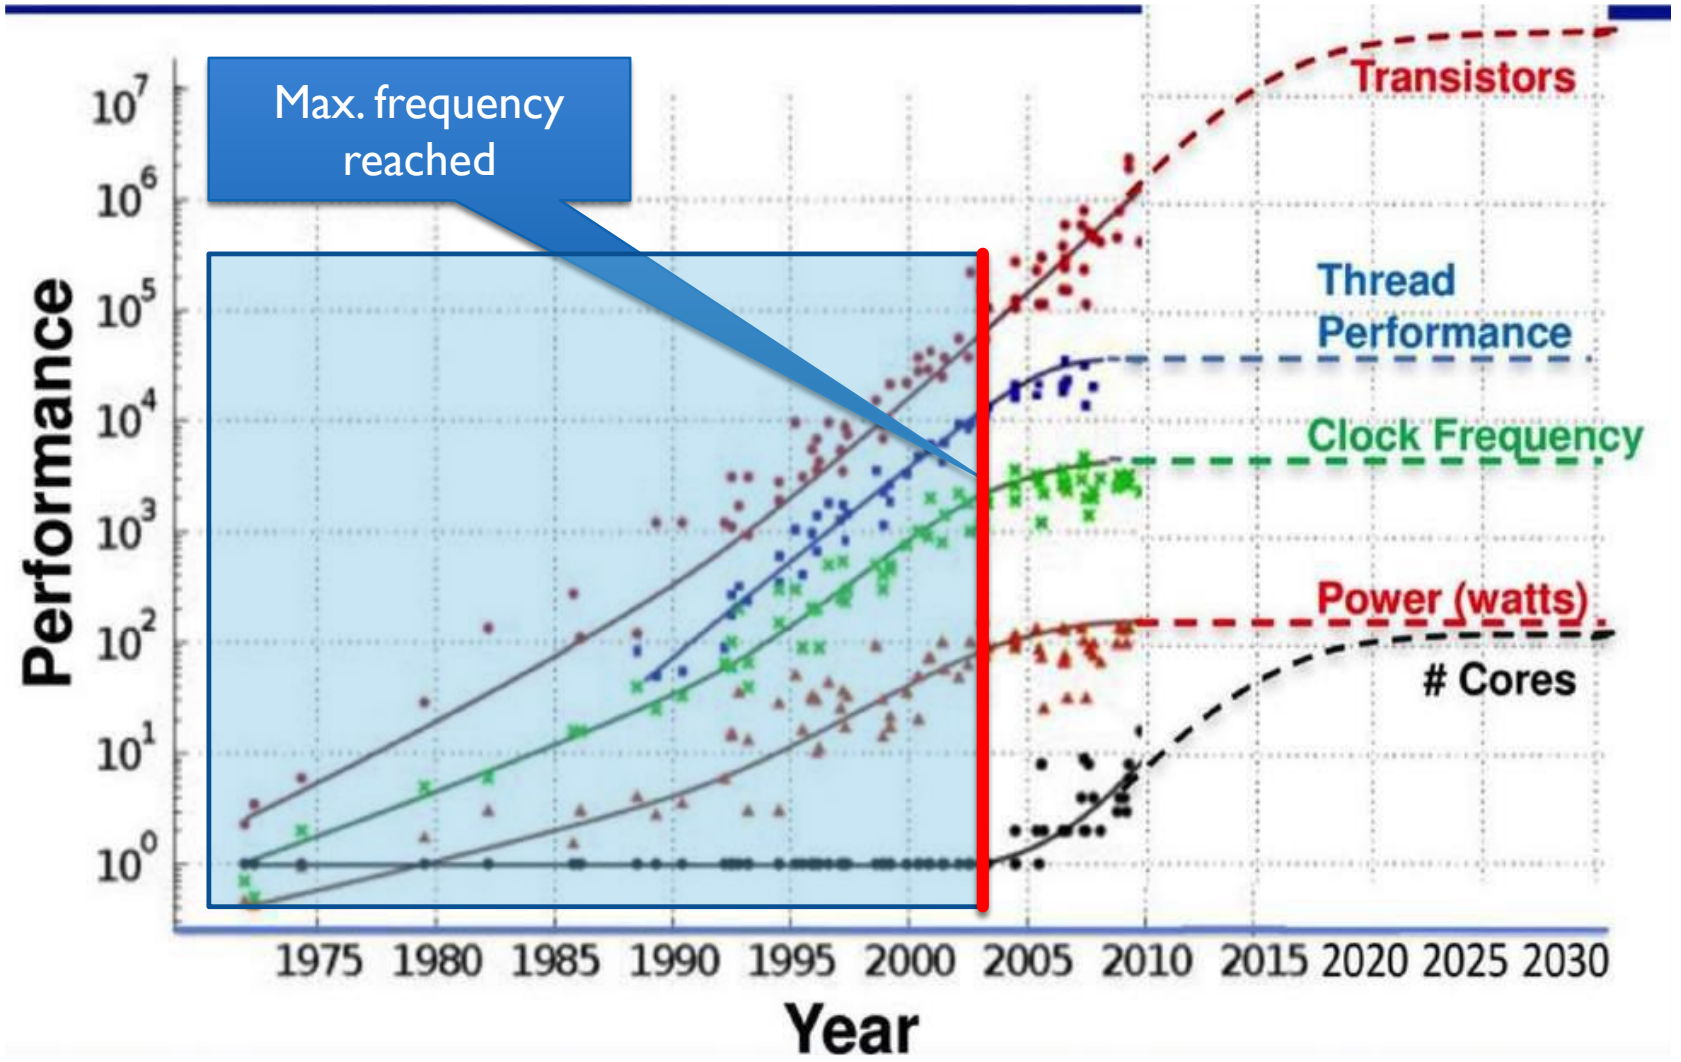

Source from Kunle Olukotun, Lance Hammond, Herb Sutter, Burton Smith, Chris Batten, and Krste Asanovic¸

MANY-CORES

Transistor nb ↗ ↗
Frequency →
Power density →

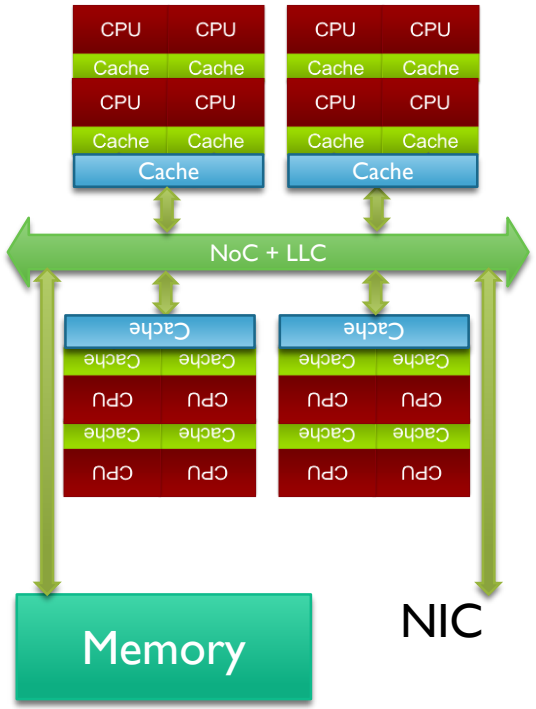

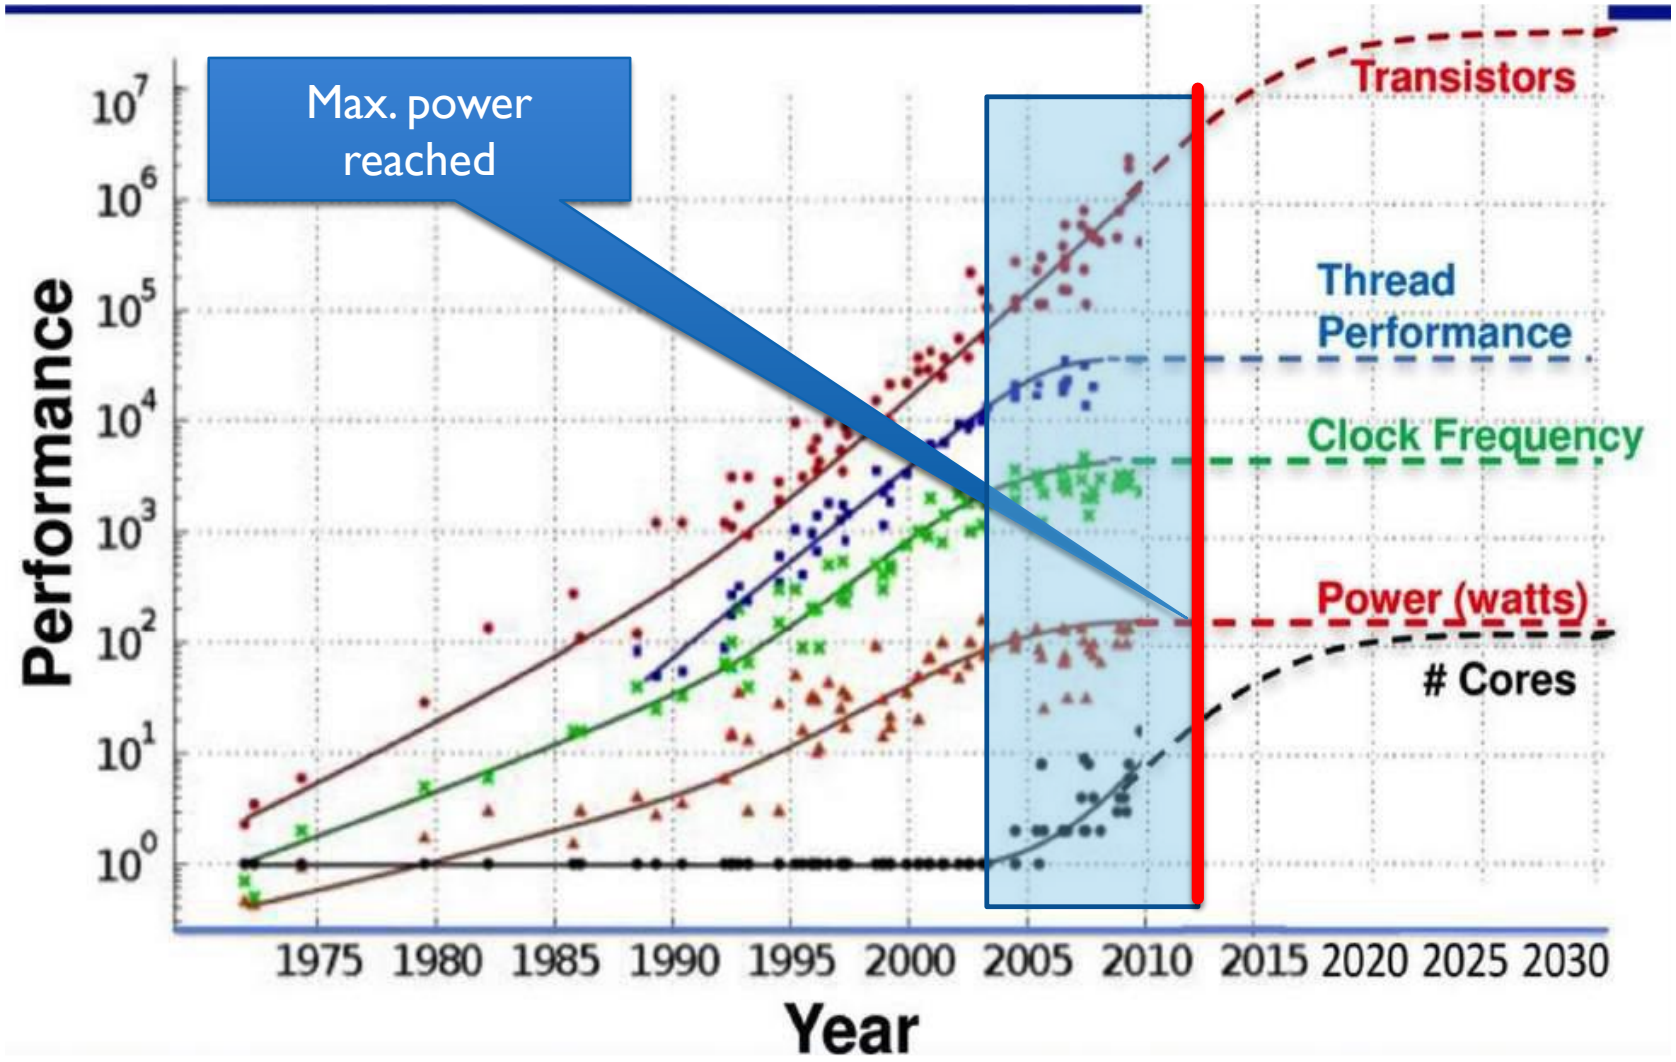

Source from Kunle Olukotun, Lance Hammond, Herb Sutter, Burton Smith, Chris Batten, and Krste Asanovic¸

HETEROGENEOUS ARCHITECTURES

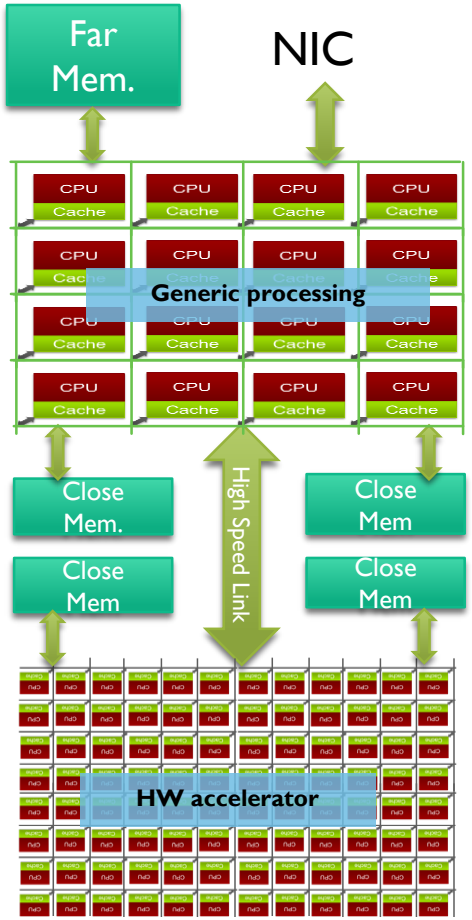

Transistor nb ↗
Frequency →
Total power →

Today and next generation

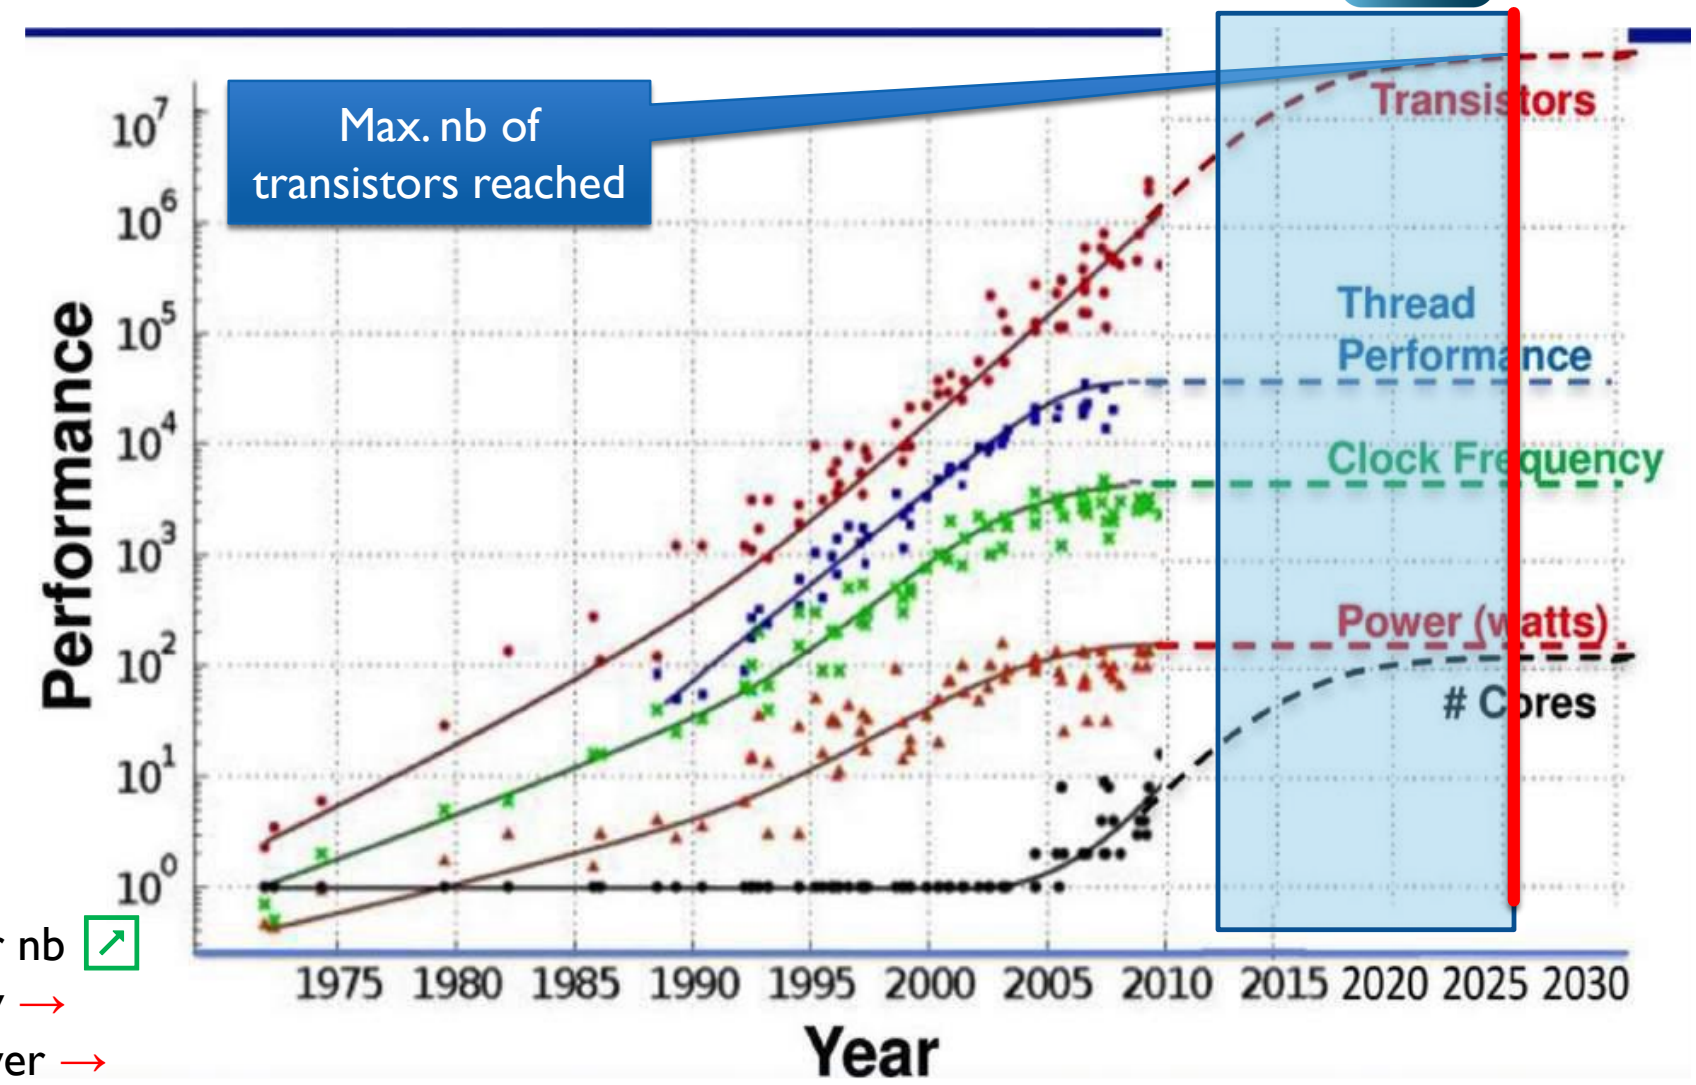

Source from Kunle Olukotun, Lance Hammond, Herb Sutter, Burton Smith, Chris Batten, and Krste Asanović

RACE TO EXASCALE

- CPU architecture choice
  - x86 + accelerator (heterogeneous)
  - Arm/SVE (homogeneous)
  - Others

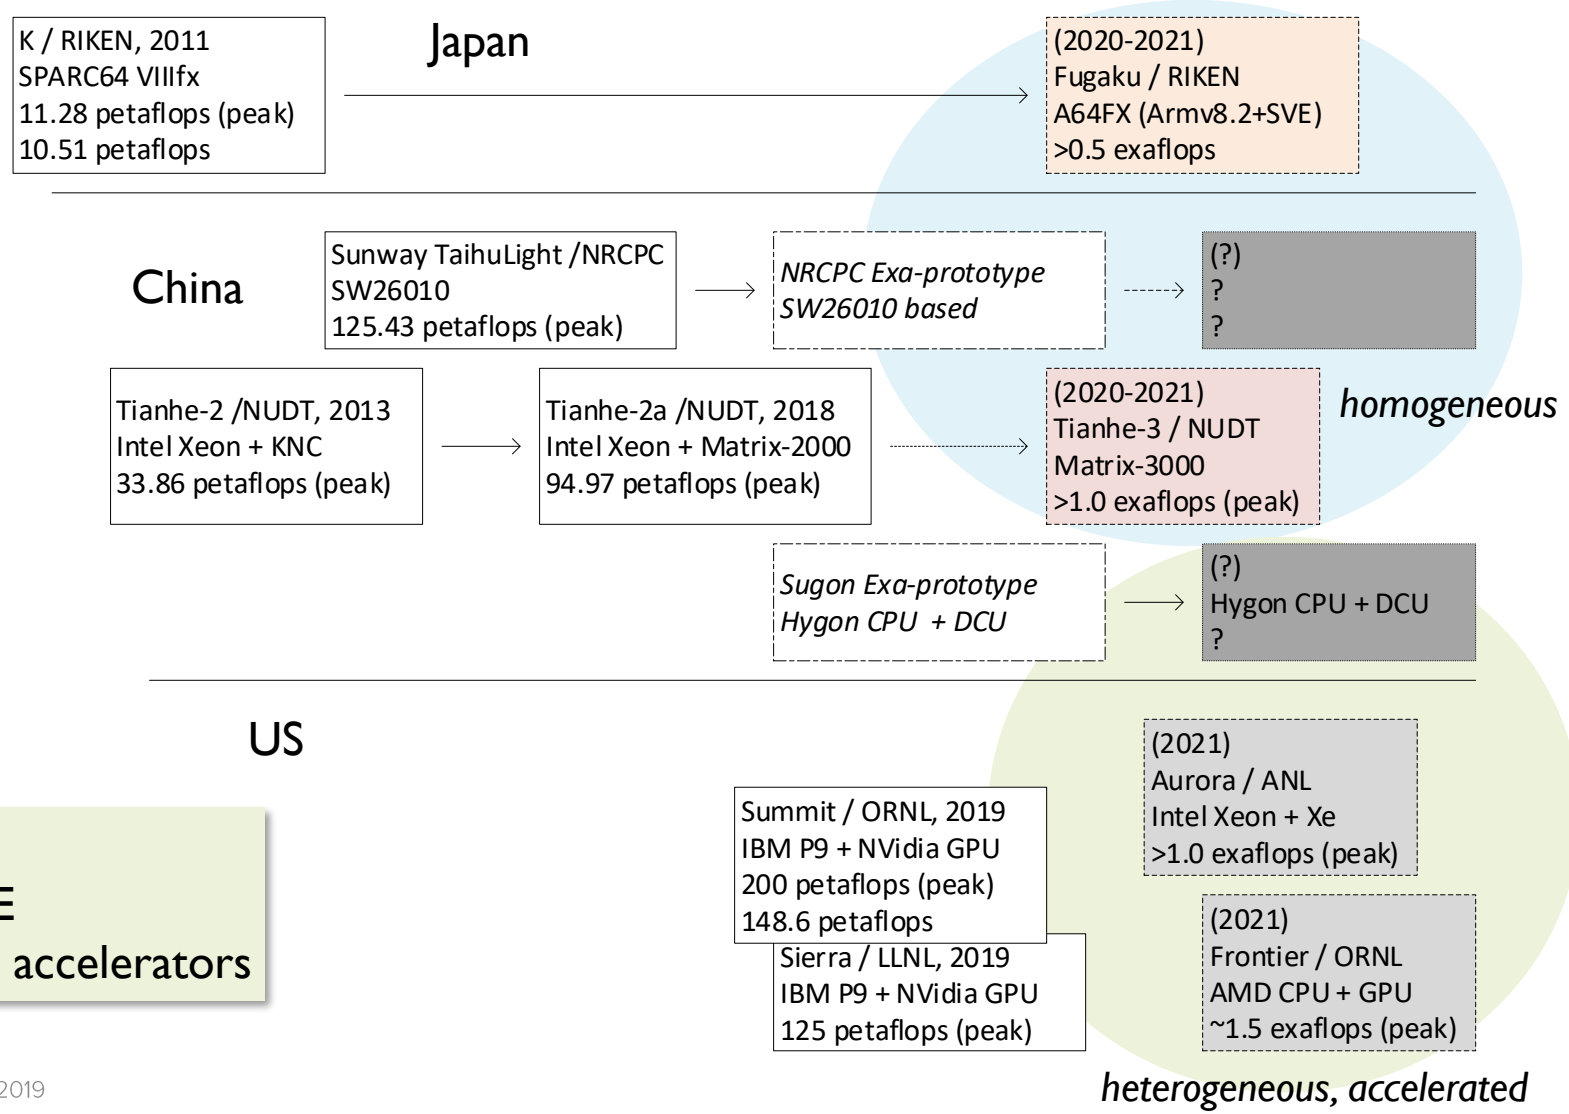

EPI takes 2-step approach
step#1 : homogeneous with Arm core+SVE
step#2 : heterogeneous with additional EPI accelerators

53RD EDITION OF THE TOP500 LIST (JUNE, 2019)

Heterogeneous integration ?

	Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
Yes	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
Yes	2	DOE/NNSA/LLNL United States	**Sierra** - IBM Power System S922LC, 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
No	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
Yes	4	National Super Computer Center in Guangzhou China	**Tianhe-2A** - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000 NUDT	4,981,760	61,444.5	100,678.7	18,482
Yes	5	Texas Advanced Computing Center/Univ. of Texas United States	**Frontera** - Dell C6420, Xeon Platinum 8280 28C 2.7GHz, Mellanox InfiniBand HDR Dell EMC	448,448	23,516.4	38,745.9	

TOP500 #1 & #2: NVIDIA TESLA V100 GPU + IBM POWER9 CPU

Source IBM

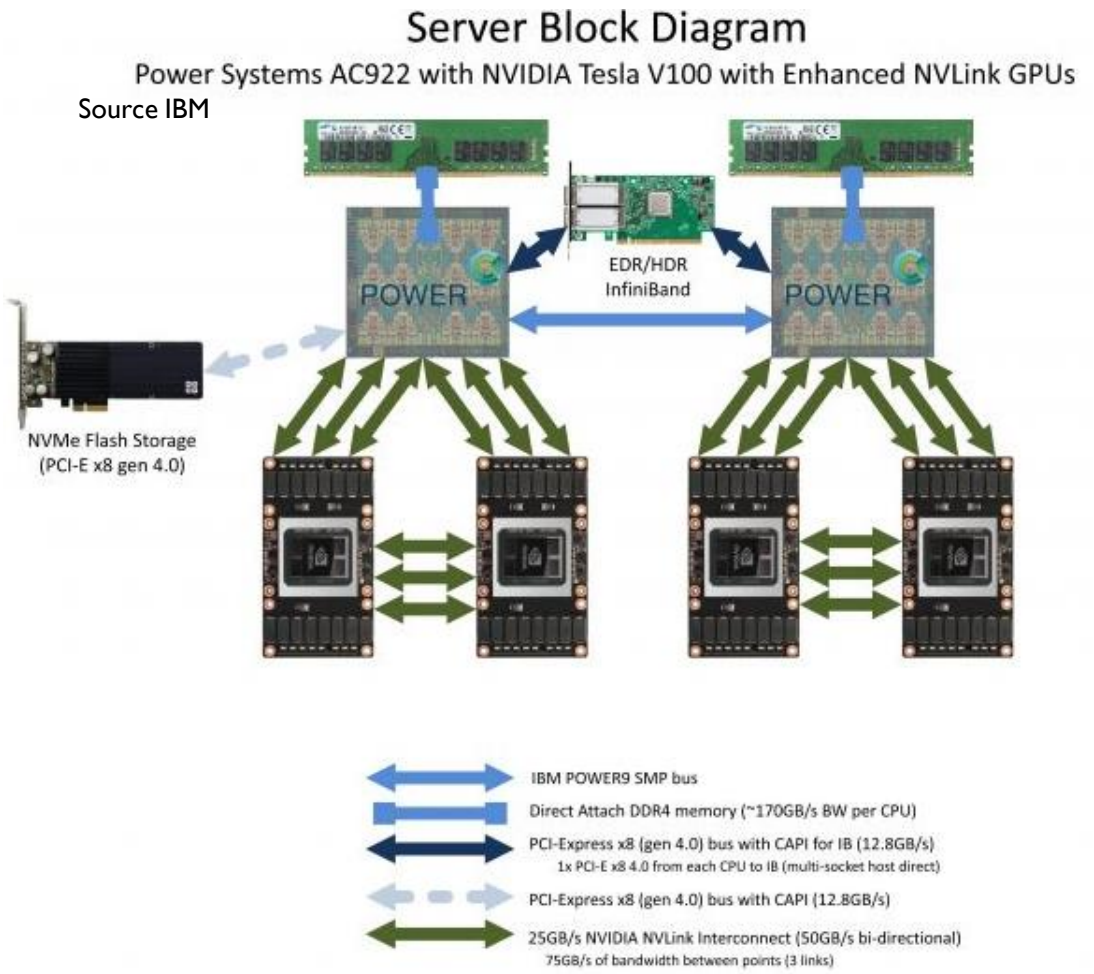

NVIDIA TESLA V100 SPECIFICATIONS

	Tesla V100 for NVLink
PERFORMANCE with NVIDIA GPU Boost™	DOUBLE-PRECISION 7.8 teraFLOPS SINGLE-PRECISION 15.7 teraFLOPS DEEP LEARNING 125 teraFLOPS
INTERCONNECT BANDWIDTH Bi-Directional	NVLINK 300 GB/s
MEMORY CoWoS Stacked HBM2	CAPACITY 32/16 GB HBM2 BANDWIDTH 900 GB/s

Source NVIDIA

AMD'S EPYC AND RADEON TO POWER EXASCALE SUPERCOMPUTER

https://www.amd.com/es/products/frontier

- Performance target: 1.5 exaflops; 40 MW; 37 GFLOPS/W
- Compute node: x1 CPU + x4 GPU + coherent fabric

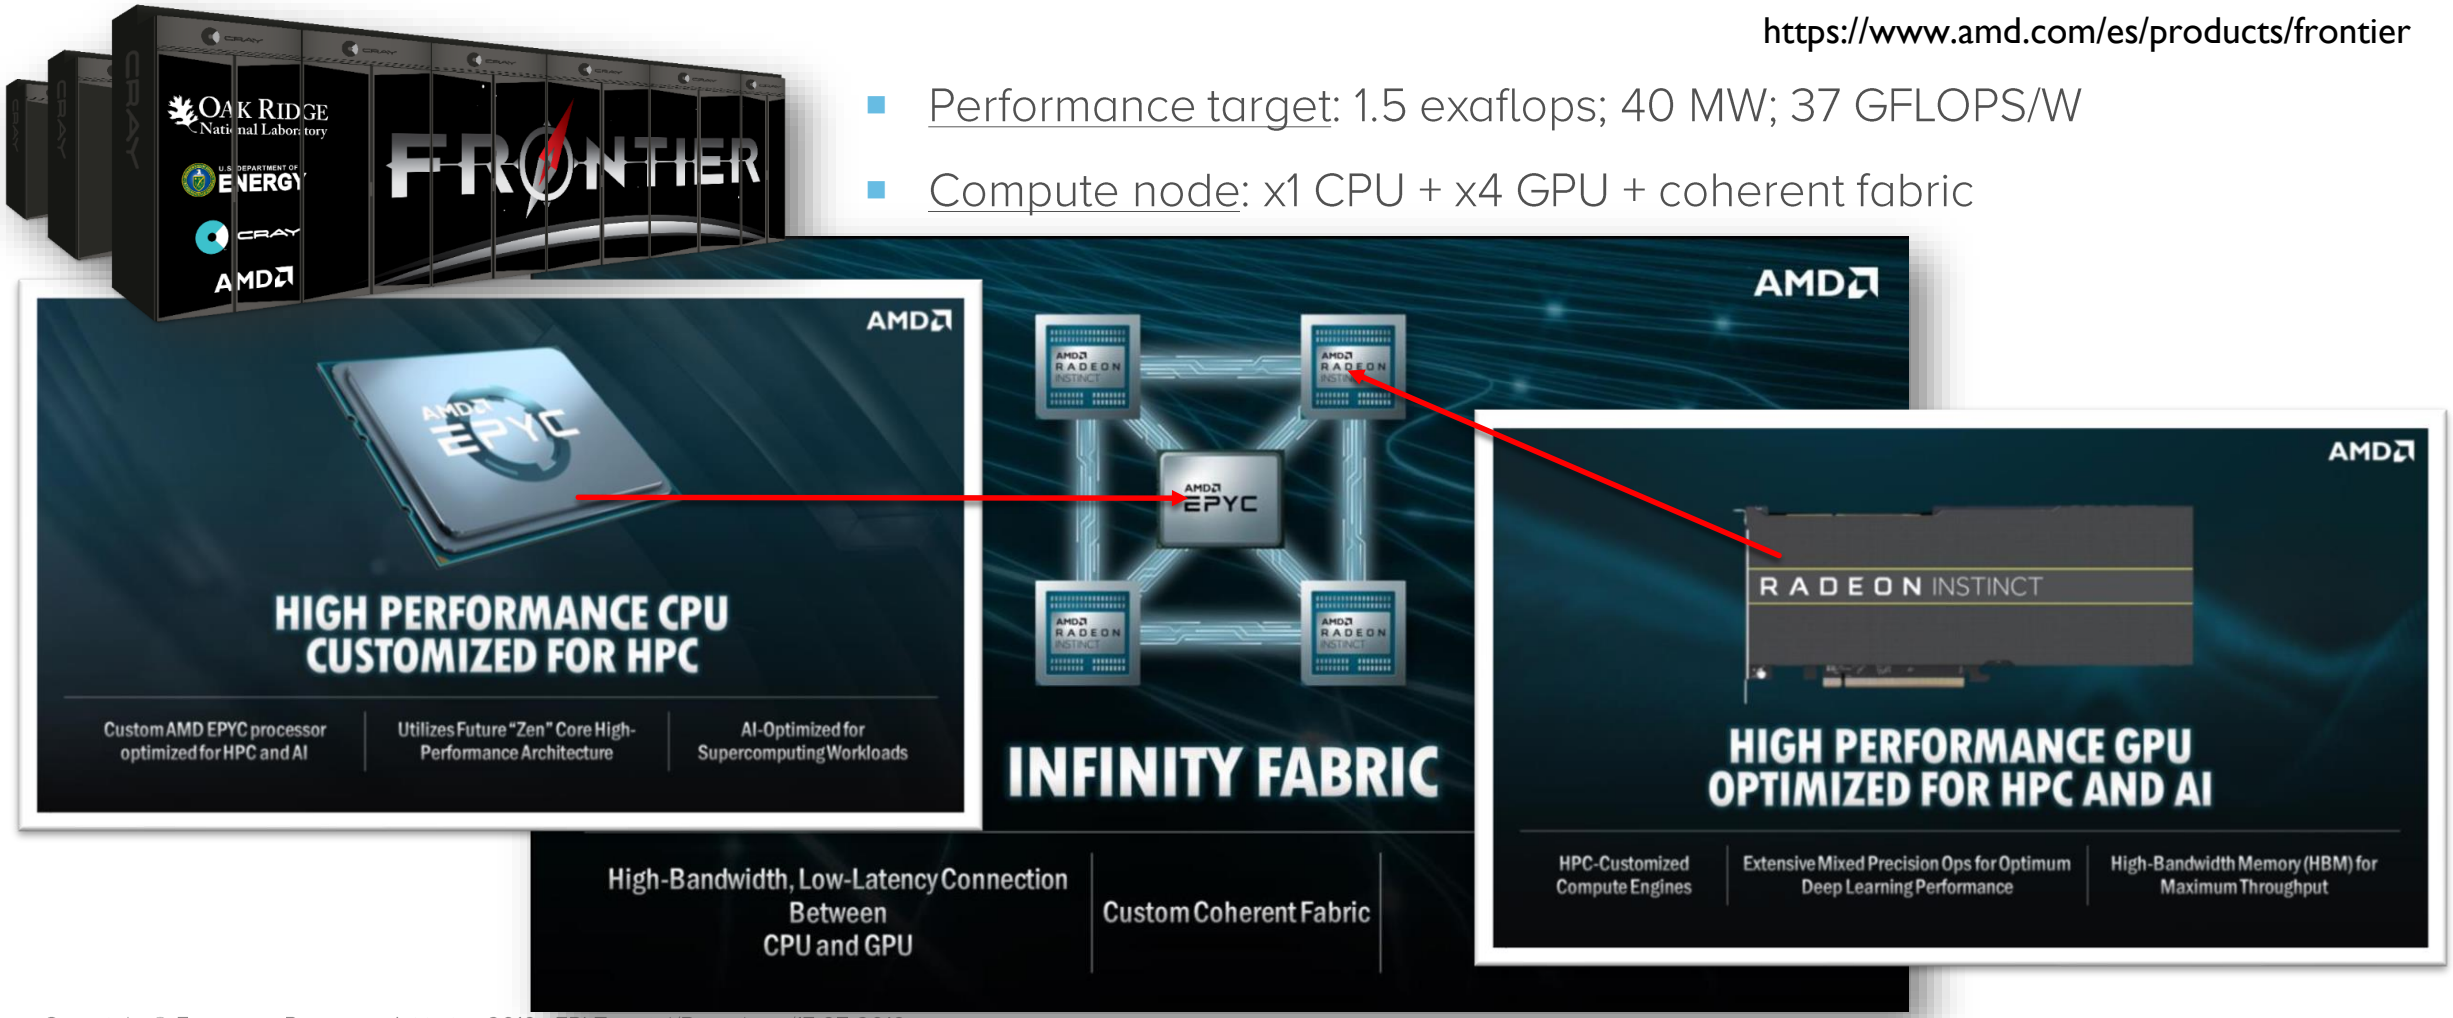

GENERIC PROCESSING WITH SCALABLE VECTOR EXTENSION: FUJITSU & ARM

- Generic processing is going towards ultra-high memory bandwidth

Source Fujitsu HotChips 2018

A64FX Chip Overview

FUJITSU

Architecture Features

- Armv8.2-A (AArch64 only)
- SVE 512-bit wide SIMD
- 48 computing cores + 4 assistant cores*

*All the cores are identical

- HBM2 32GiB
- Tofu 6D Mesh/Torus
  28Gbps x 2 lanes x 10 ports
- PCIe Gen3 16 lanes

7nm FinFET

- 8,786M transistors
- 594 package signal pins

Peak Performance (Efficiency)

- >2.7TFLOPS (>90%@DGEMM)
- Memory B/W 1024GB/s (>80%@Stream Triad)

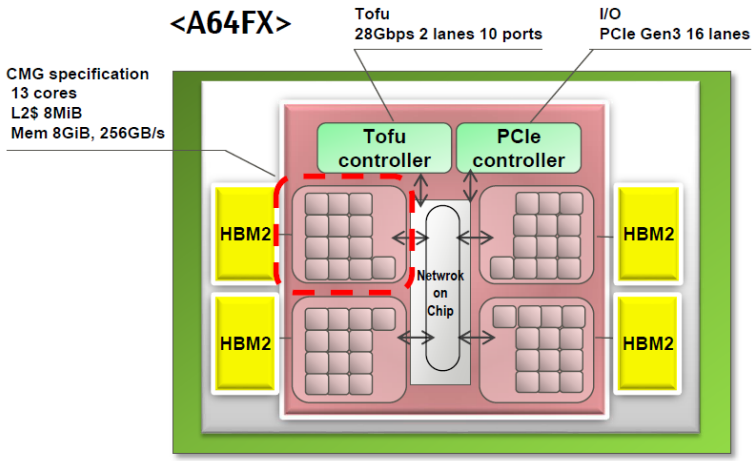

	A64FX (Post-K)	SPARC64 XIfx (PRIMEHPC FX100)
ISA (Base)	Armv8.2-A	SPARC-V9
ISA (Extension)	SVE	HPC-ACE2
Process Node	7nm	20nm
Peak Performance	>2.7TFLOPS	1.1TFLOPS
SIMD	512-bit	256-bit
# of Cores	48+4	32+2
Memory	HBM2	HMC
Memory Peak B/W	1024GB/s	240GB/s x2 (in/out)

6

SEMICONDUCTOR TECHNOLOGY EVOLUTION

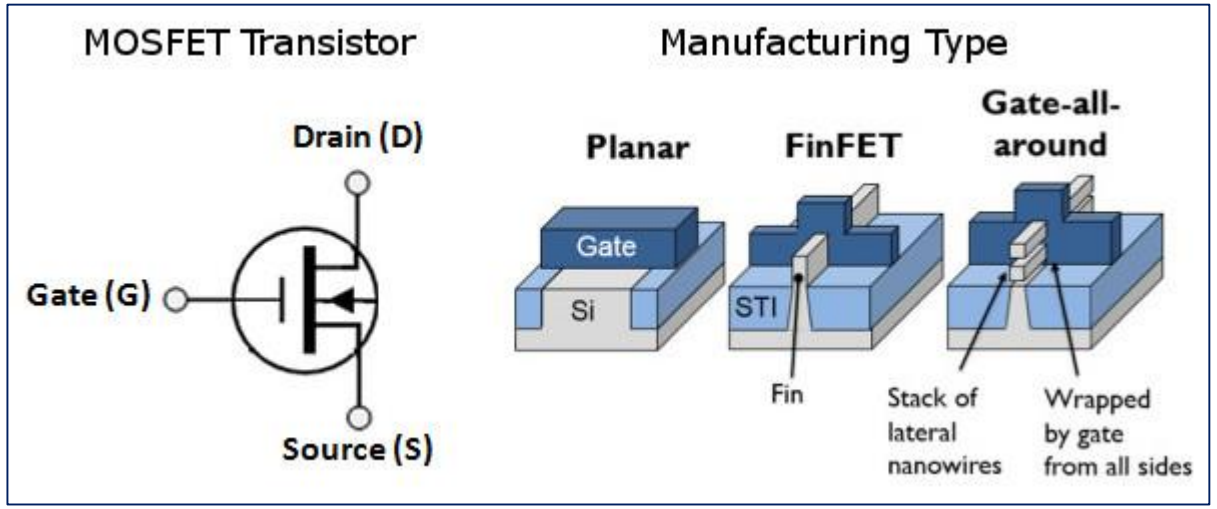

Source ScienceDirect

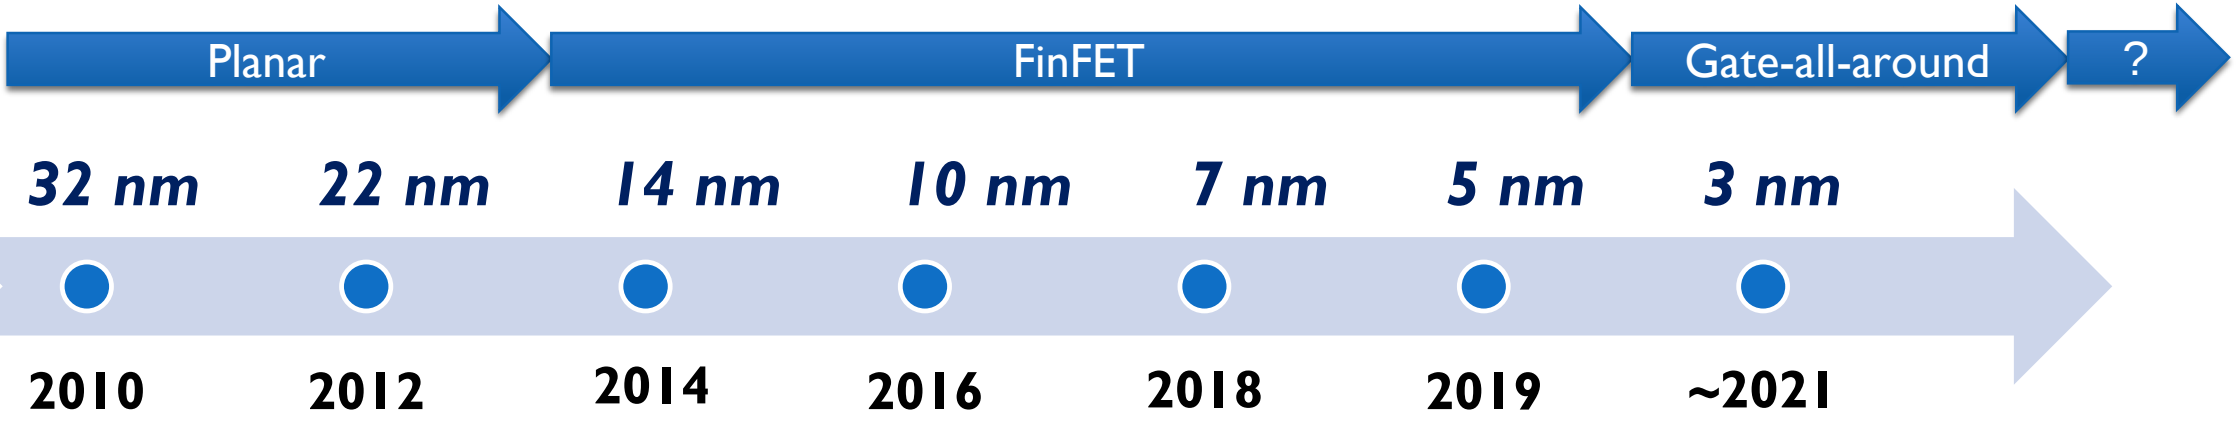

Source Wikipedia, WikiChip

SEMICONDUCTOR MANUFACTURING PROCESSES AND FOUNDRIES

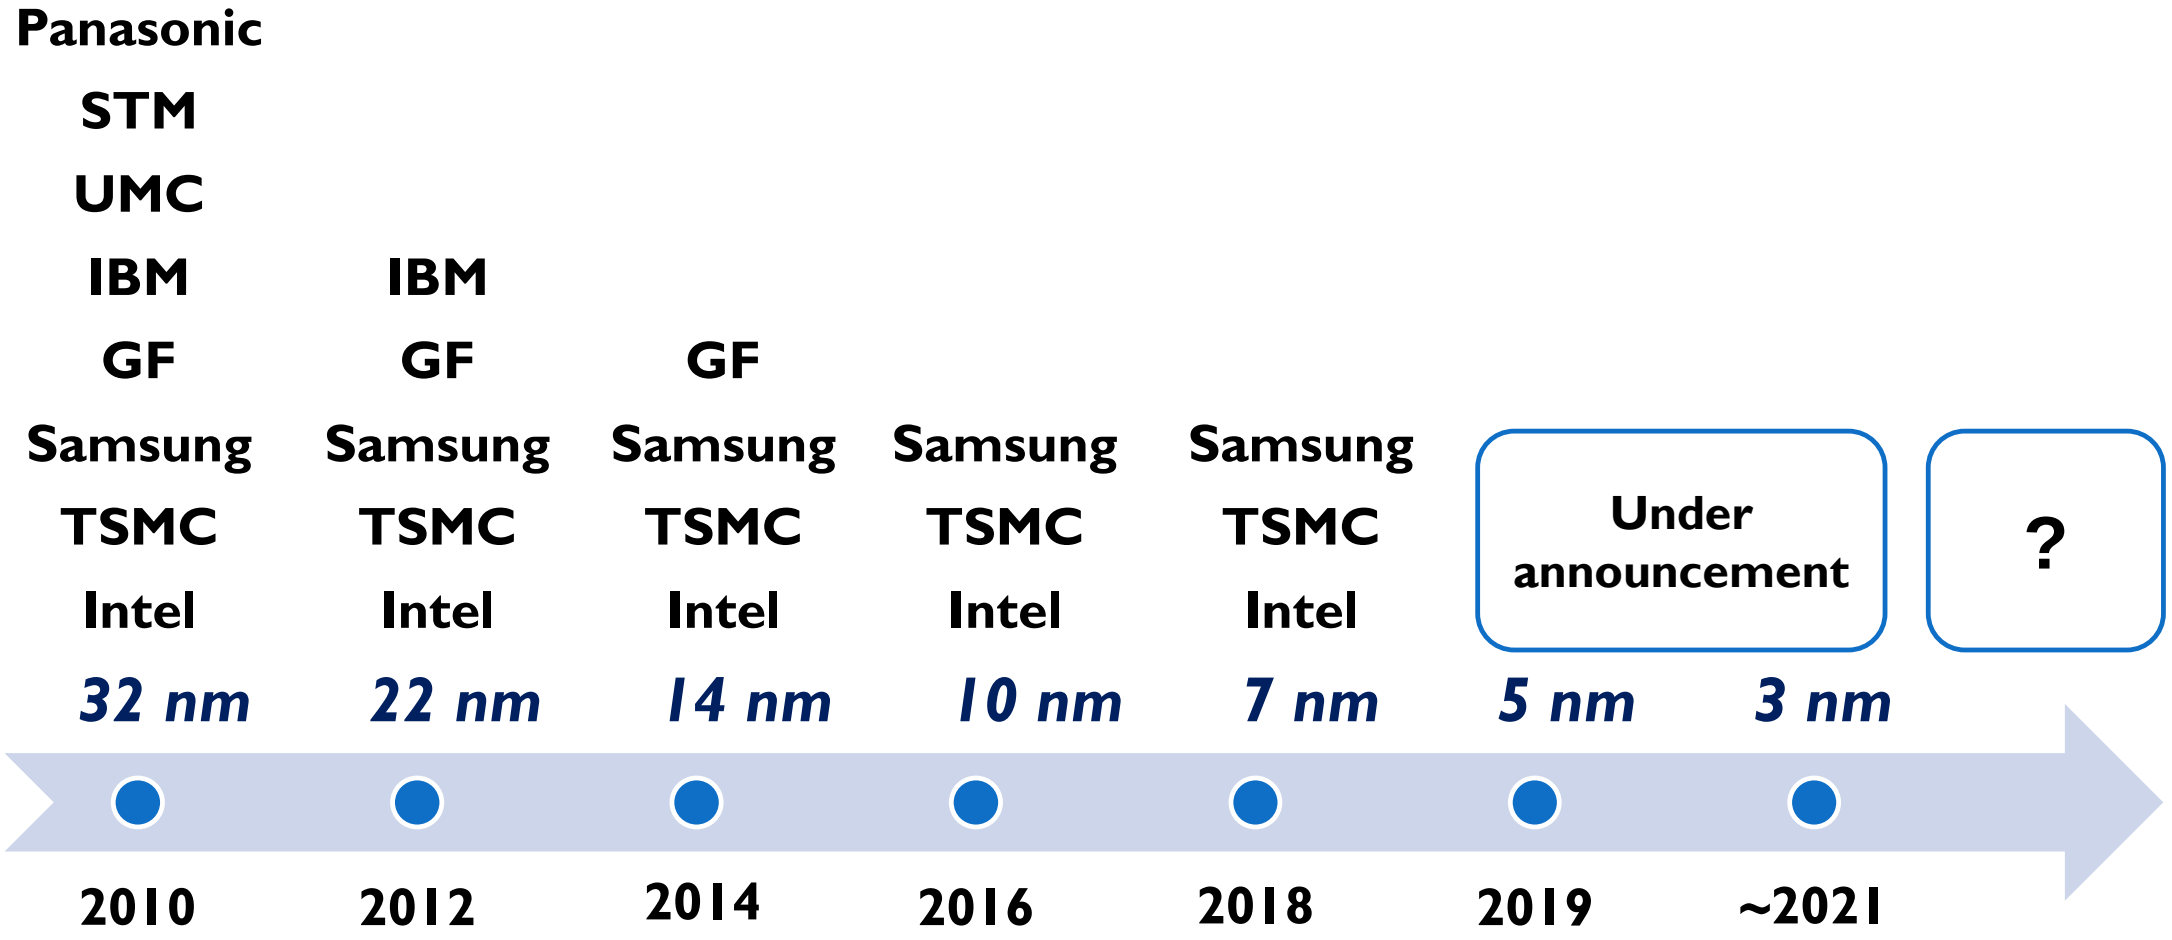

Source Wikipedia, WikiChip

FROM ADVANCED PACKAGING TECHNOLOGIES …. … TO CHIPLETS

Advanced Integration

SiP

Multi-Chip-Module

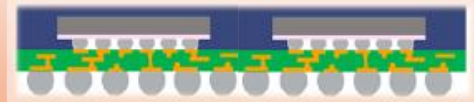

Interconnect density: 100µm x 100µm

3D

Die stacking

Interconnect density: 10µm x 10µm

2.5D

Interposer based

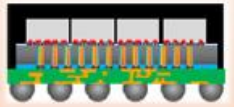

Interconnect density: 10µm x 10µm

CHIPLET partitioning

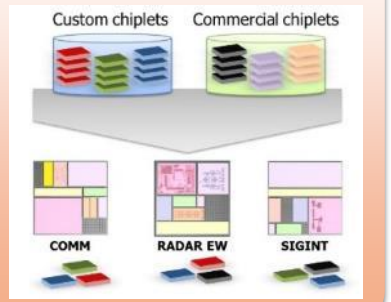

Source: DARPA

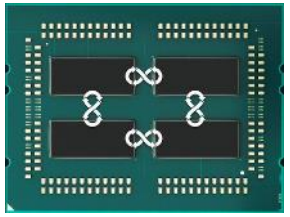

Source: AMD EPYC 7260, 4-chiplet chip

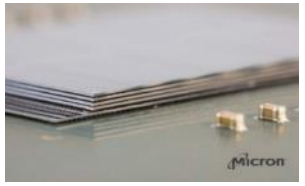

Source: Micron High-Bandwidth-Memory

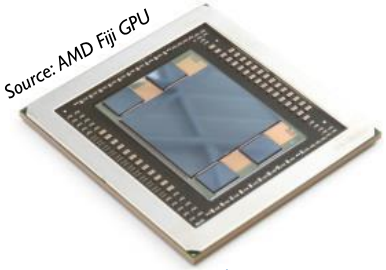

Source: AMD Fiji GPU

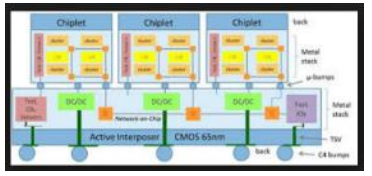

Source: LETI

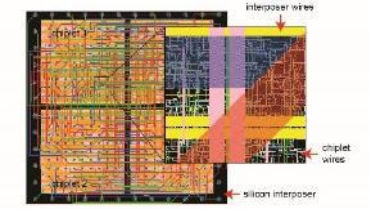

Source: GeorgiaTech

System-in-Package

3D Integrated-Circuit (3D IC)

MULTI-CHIP-MODULE: INTEGRATION WITH CHIPLETS

- AMD “Zen” architecture integrates upto 4 chiplets on a substrate; for scalable solution and more than reticle size silicon area in a chip.

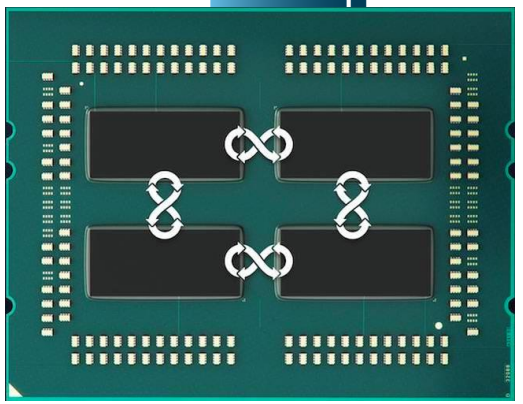

(2017.10) EPYC 7260, 4-chiplet chip

- AMD “Zen 2” architecture integrates upto 9 chiplets on a substrate

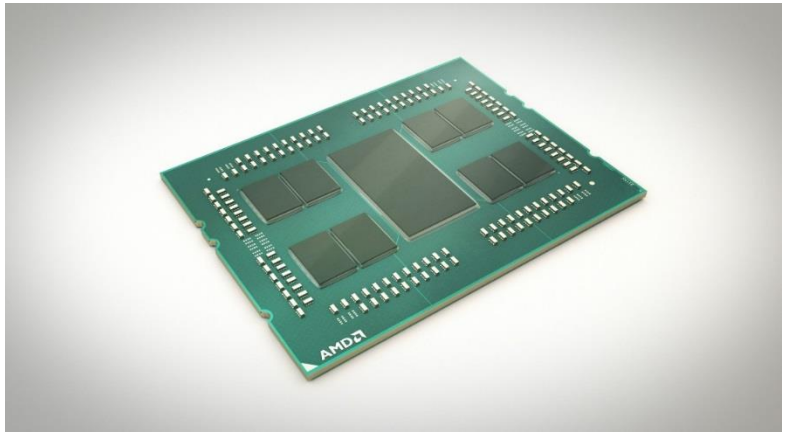

(2018.11) EPYC “Rome”, 9-chiplet chip

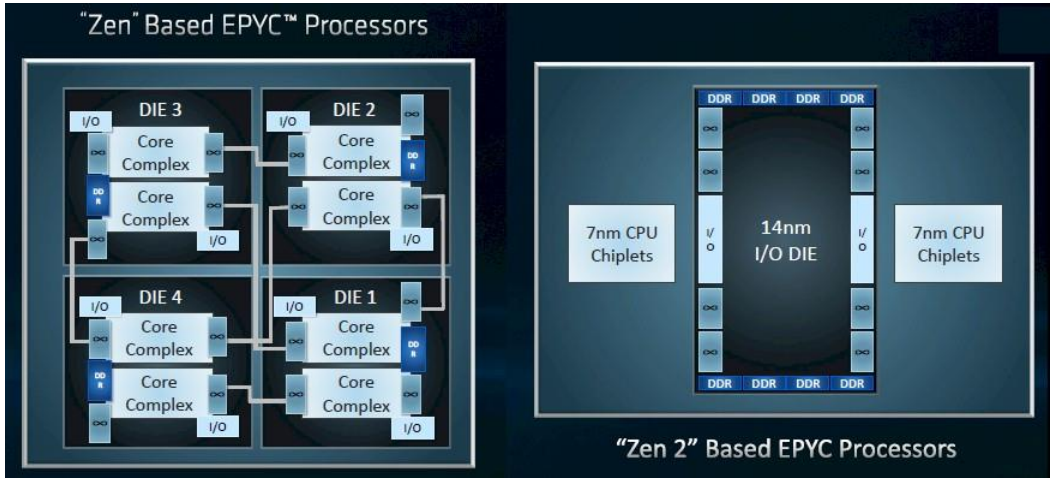

(2018.11) AMD Zen2 architecture;

2.5D INTERPOSER: HBM INTEGRATION FOR MEMORY BANDWIDTH

- NEC Aurora SX-10+
  - First product with 6x HBM2:
  - 1.2 TB/s total memory bandwidth
  - 2.45 TFLOPS
  - ~0.5 Byte/Flops

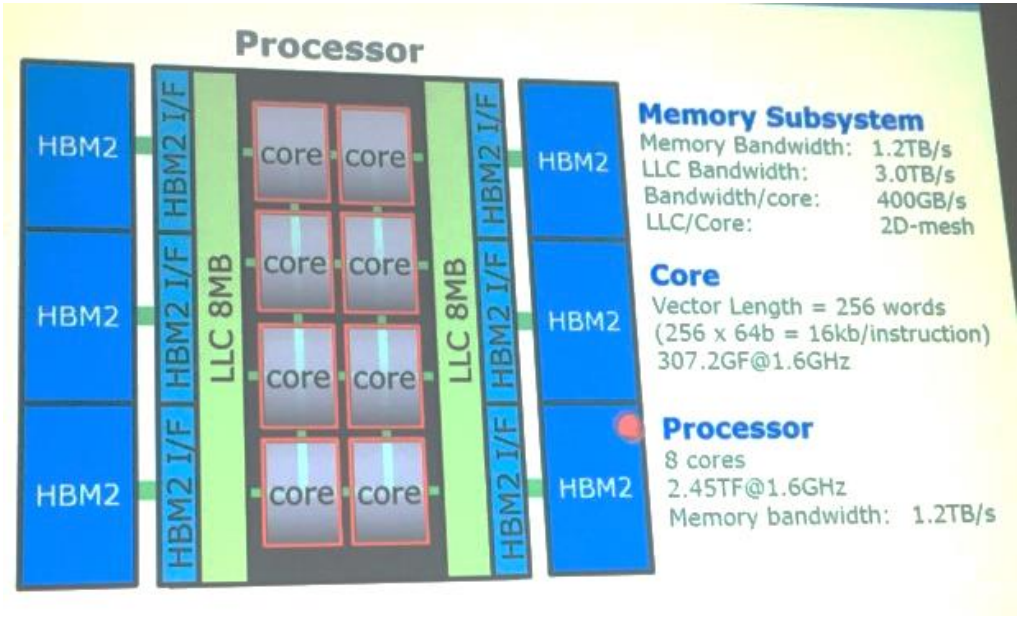

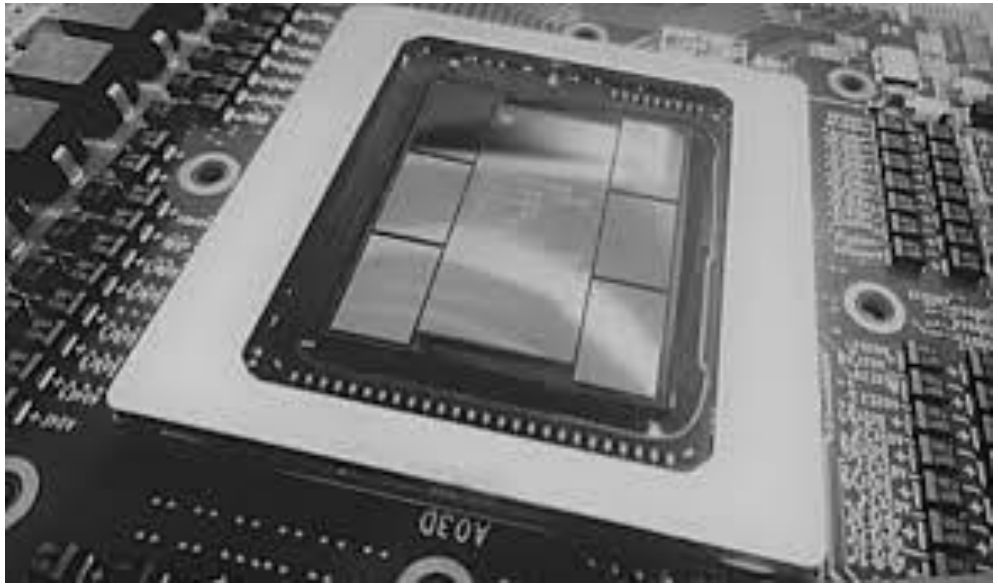

53RD EDITION OF THE TOP500 LIST (JUNE 2019)

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

  #1-#2: US

  #3-#4: China
- Processor design & technology:

Chip	Design	Manuf.
IBM POWER9		
NVIDIA Volta GV100		
Sunway SW26010		
Intel Xeon E5		

Rank	Site	System
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	DOE/NNSA/LLNL United States	**Sierra** - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband IBM / NVIDIA / Mellanox
3	National Supercomputing Center in Wuxi China	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway NRCPC
4	National Super Computer Center in Guangzhou China	**Tianhe-2A** - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000 NUDT

53RD EDITION OF THE TOP500 LIST (JUNE 2019)

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

  #1-#2: US

  #3-#4: China
- Processor design & technology:

Chip	Design	Manuf.
IBM POWER9		
NVIDIA Volta GV100		
Sunway SW26010		
Intel Xeon E5		

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)

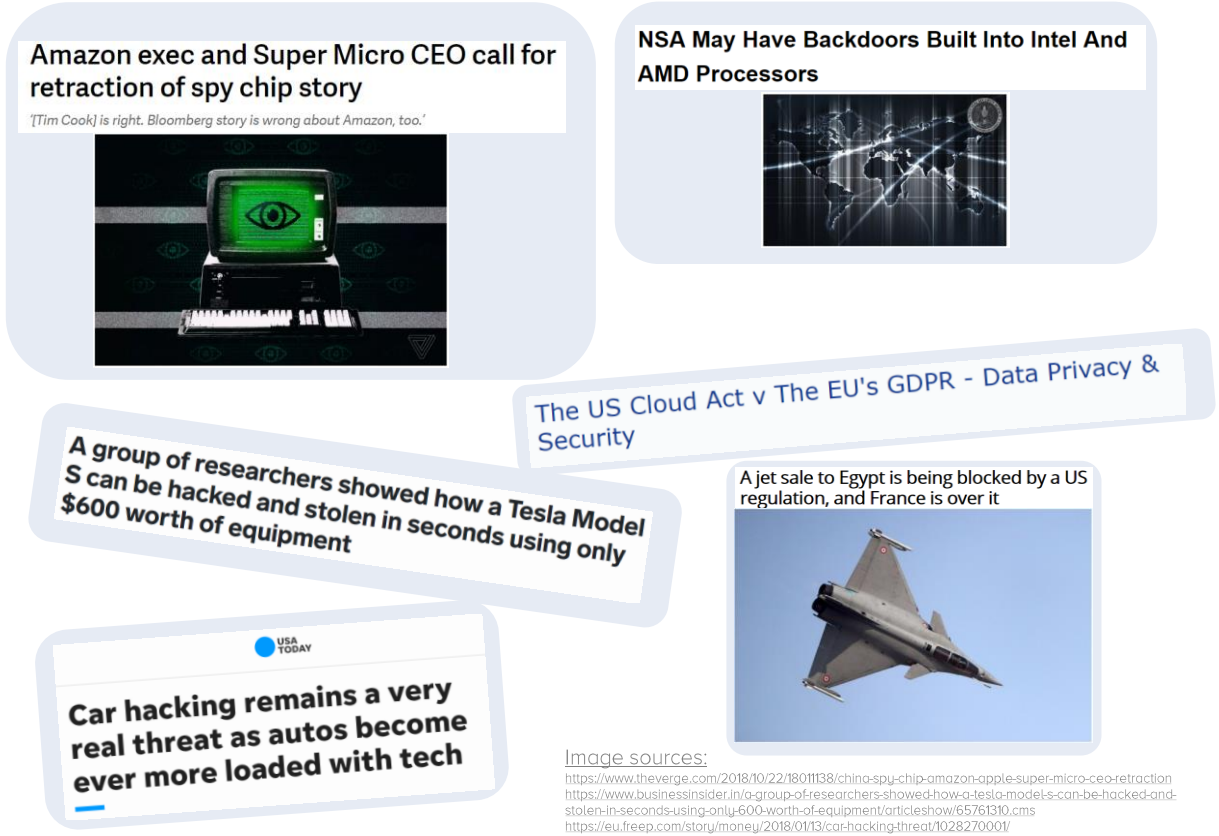

Image sources:

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/

HOW EUROHPC WILL HELP TO MAKE US STRONGER

- 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 (EPI)

ANDREA BARTOLINI (SLIDES PREPARED BY DENIS DUTOIT)

EUROPEAN PROCESSOR INITIATIVE (EPI)

Start	Finish	Topic	Presenter
14:25		**European Processor Initiative (EPI)**	
14:25	14:35	Overview	Andrea Bartolini, UNIBO
14:35	14:50	Processor and general architecture	Andrea Bartolini, UNIBO
14:50	15:00	Co-design process and modeling	Andrea Bartolini, UNIBO
15:00	15:45	Accelerator	Mauro Olivieri, BSC
15:45	16:30	Software	Jesus Labarta, BSC
16:30	17:00	Automotive	Jaume Abella, BSC; Francisco Cazorla, BSC
	17:00	End	

OVERVIEW

ANDREA BARTOLINI (SLIDES PREPARED BY DENIS DUTOIT)

EUROPEAN PROCESSOR INITIATIVE

- High Performance General Purpose Processor for HPC
- High-performance RISC-V based accelerator
- Computing platform for autonomous cars
- Will also target the AI, Big Data and other markets in order to be economically sustainable

EUROPEAN PROCESSOR INITIATIVE

European Processor Initiative

www.european-processor-initiative.eu

PROJECT PILLARS

- Common platform and global architecture stream
- HPC general purpose processor stream
- Accelerator stream
- Automotive platform stream

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 826647

EPI OBJECTIVES

- Architect of the common platform to accommodate the developed technologies
  - CoDesign Methodology, Platform for hardware and software, Power management, Modeling and Simulation
- Build a GPP processor chip ready for PreExascale level machines (RheaR1)
- Develop Accelerator technologies for HPC workload (EPAC)
- Implementation of a Real-time acceleration PoC based on the first EPI GPP Processor (MPPA)
- Interfacing with the Automotive MCU
- Development efficient power conversion technologies
- Software activities based on the platform built
- PoC systems (test-chip; ref. board, HPC blades, PCIe card and automotive PoC)
- Related research around the EPI project scopes

EPI KPIS

- Energy Efficiency

  * Pre-ExaScale level with general-purpose CPU core in the first EPI GPP chip
  * Develop acceleration technologies for better DP GFLOPS/Watt performance
  * Inclusion of MPPA for real-time application acceleration
  * Develop a Common Platform to enable EPI accelerations

- Easy to use

  * Adopt Arm general-purpose CPU core with SVE / vector acceleration in the first EPI chip
  * Supply sufficient Memory Bandwidth (Byte/FLOP) to support the GPP application
  * in SGA1, focus on programming models to include accelerations.

EPI STREAMS

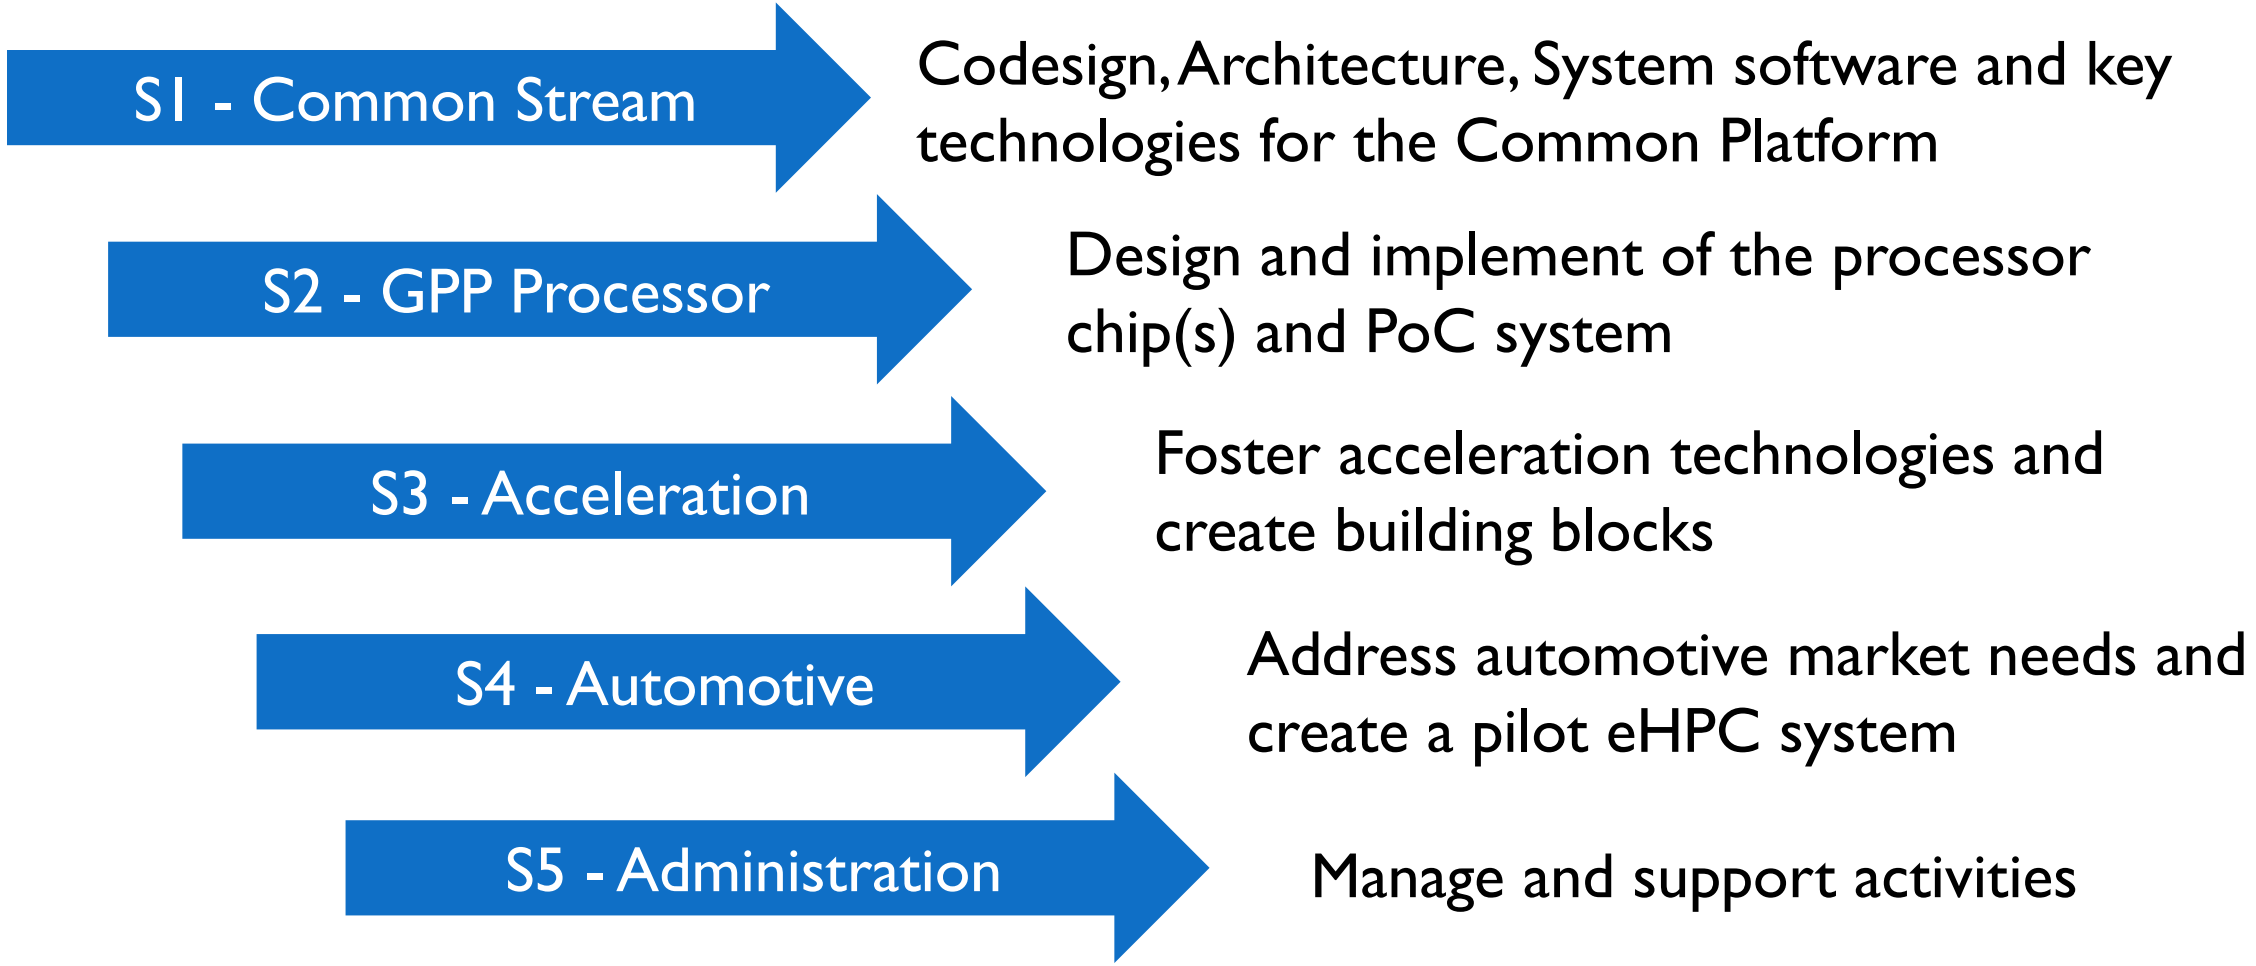

GPP AND COMMON ARCHITECTURE

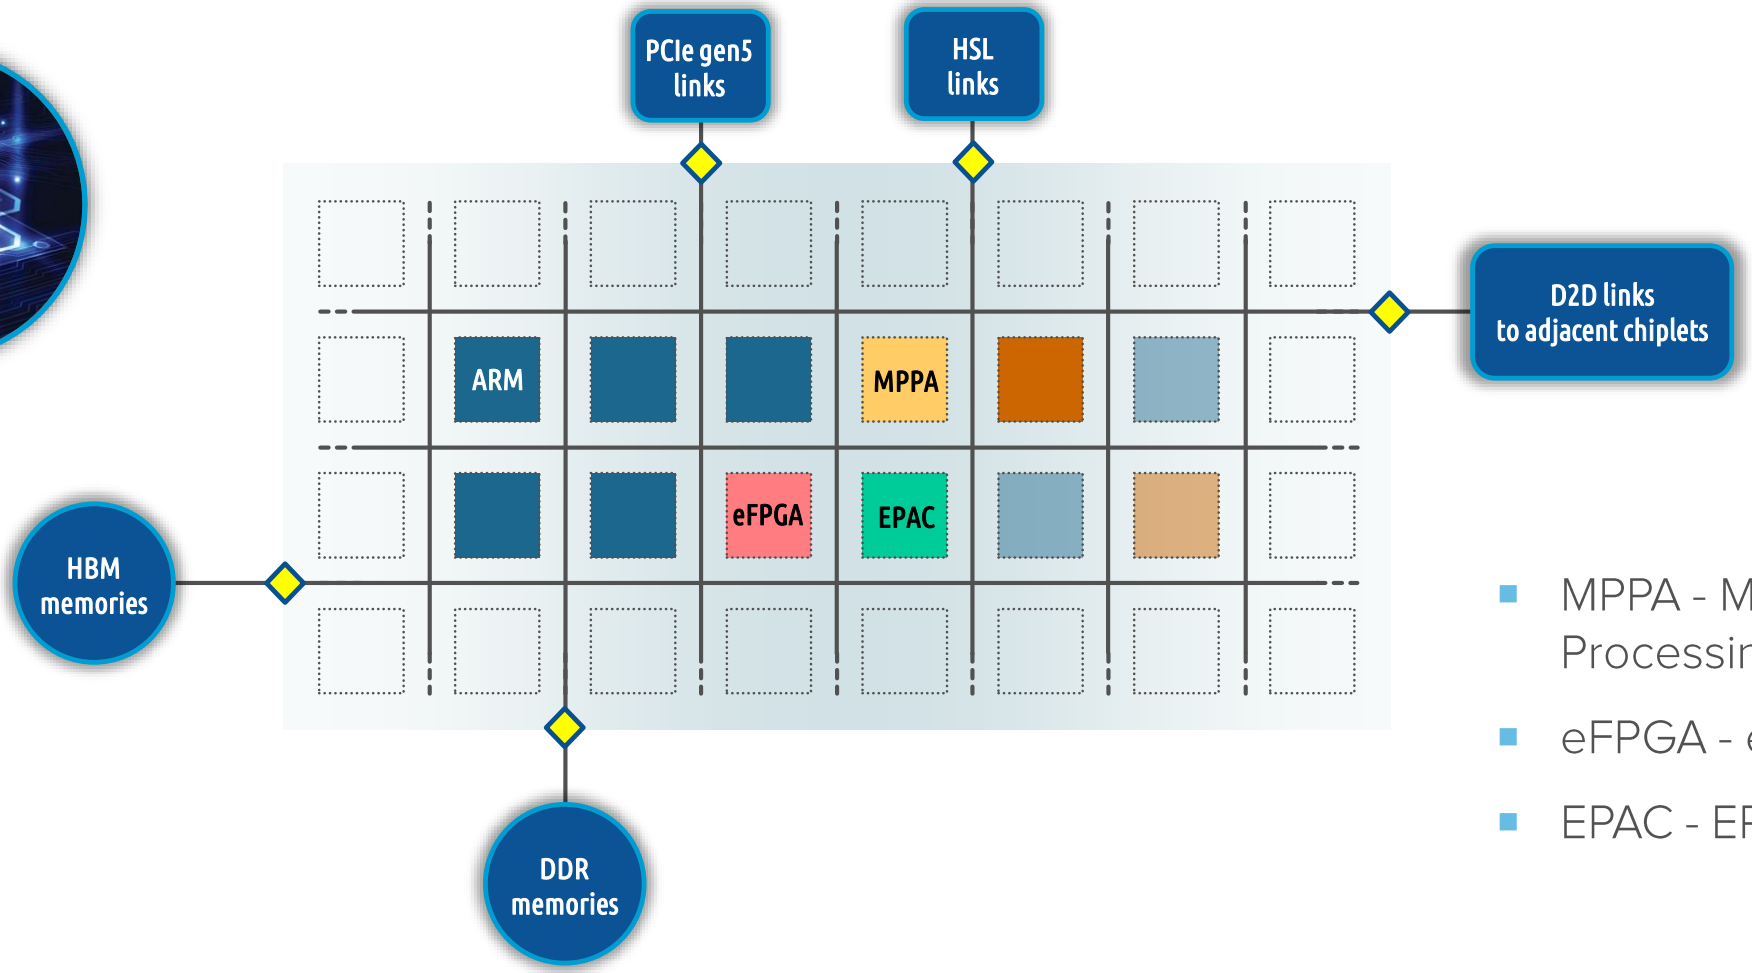

- MPPA - Multi-Purpose Processing Array
- eFPGA - embedded FPGA
- EPAC - EPI Accelerator

EPAC – RISC-V ACCELERATOR

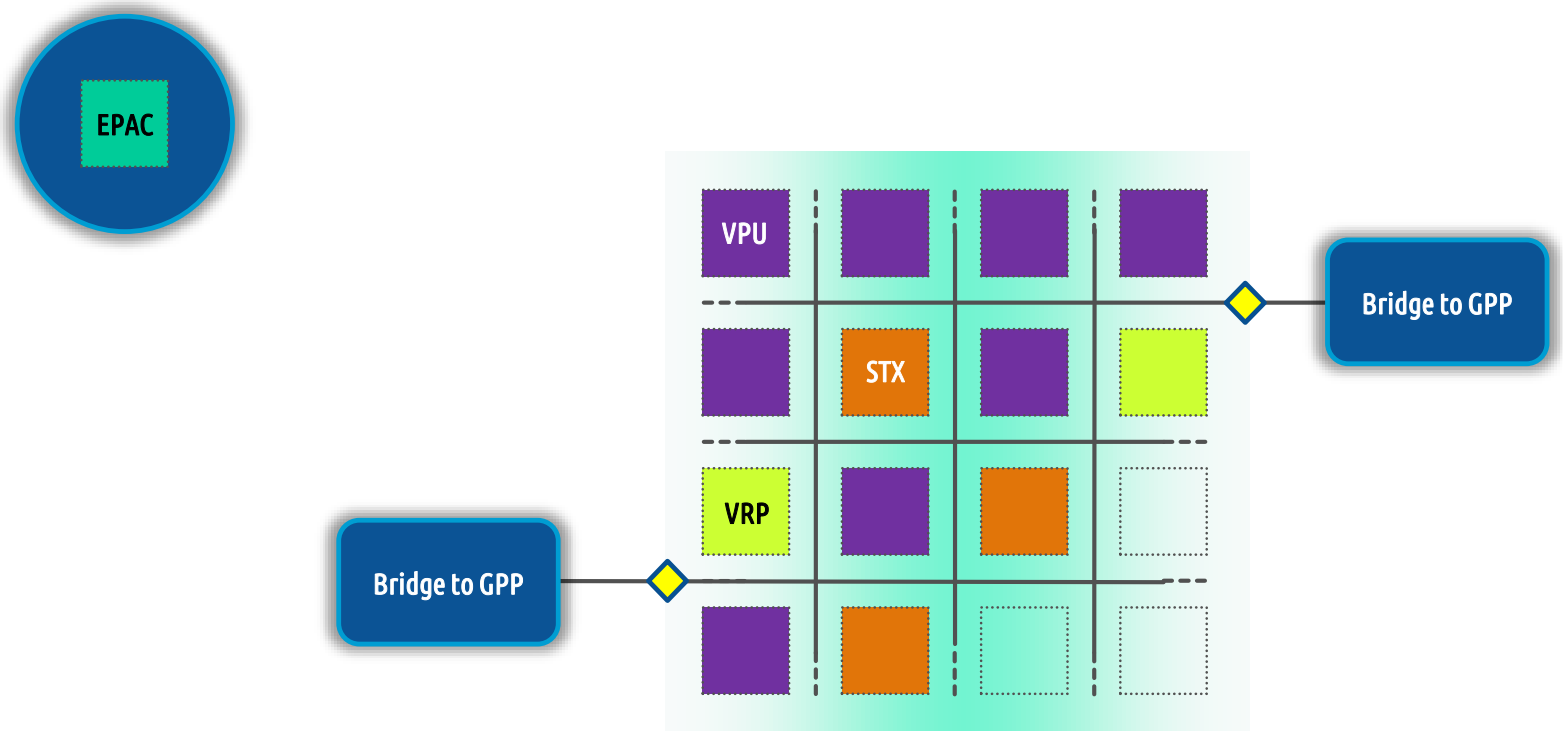

- EPAC - EPI Accelerator
- VPU – Vector Processing Unit
- STX – Stencil/Tensor accelerator
- VRP - VaRiable Precision co-processor

EPI AUTOMOTIVE

- Autonomous driving systems
- Connected mobility
- EPI: A powerful data fusion platform – the automotive embedded HPC platform
- EPI heterogeneous multicore architecture can provide enough performance and low power consumption in parallel

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

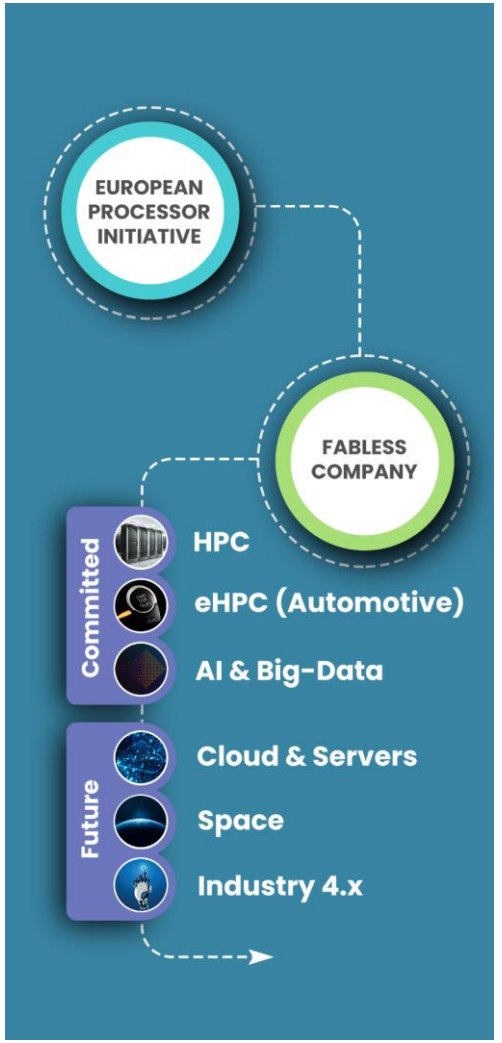

SCALABILITY ALLOWS WIDE MARKET POTENTIAL COVERAGE

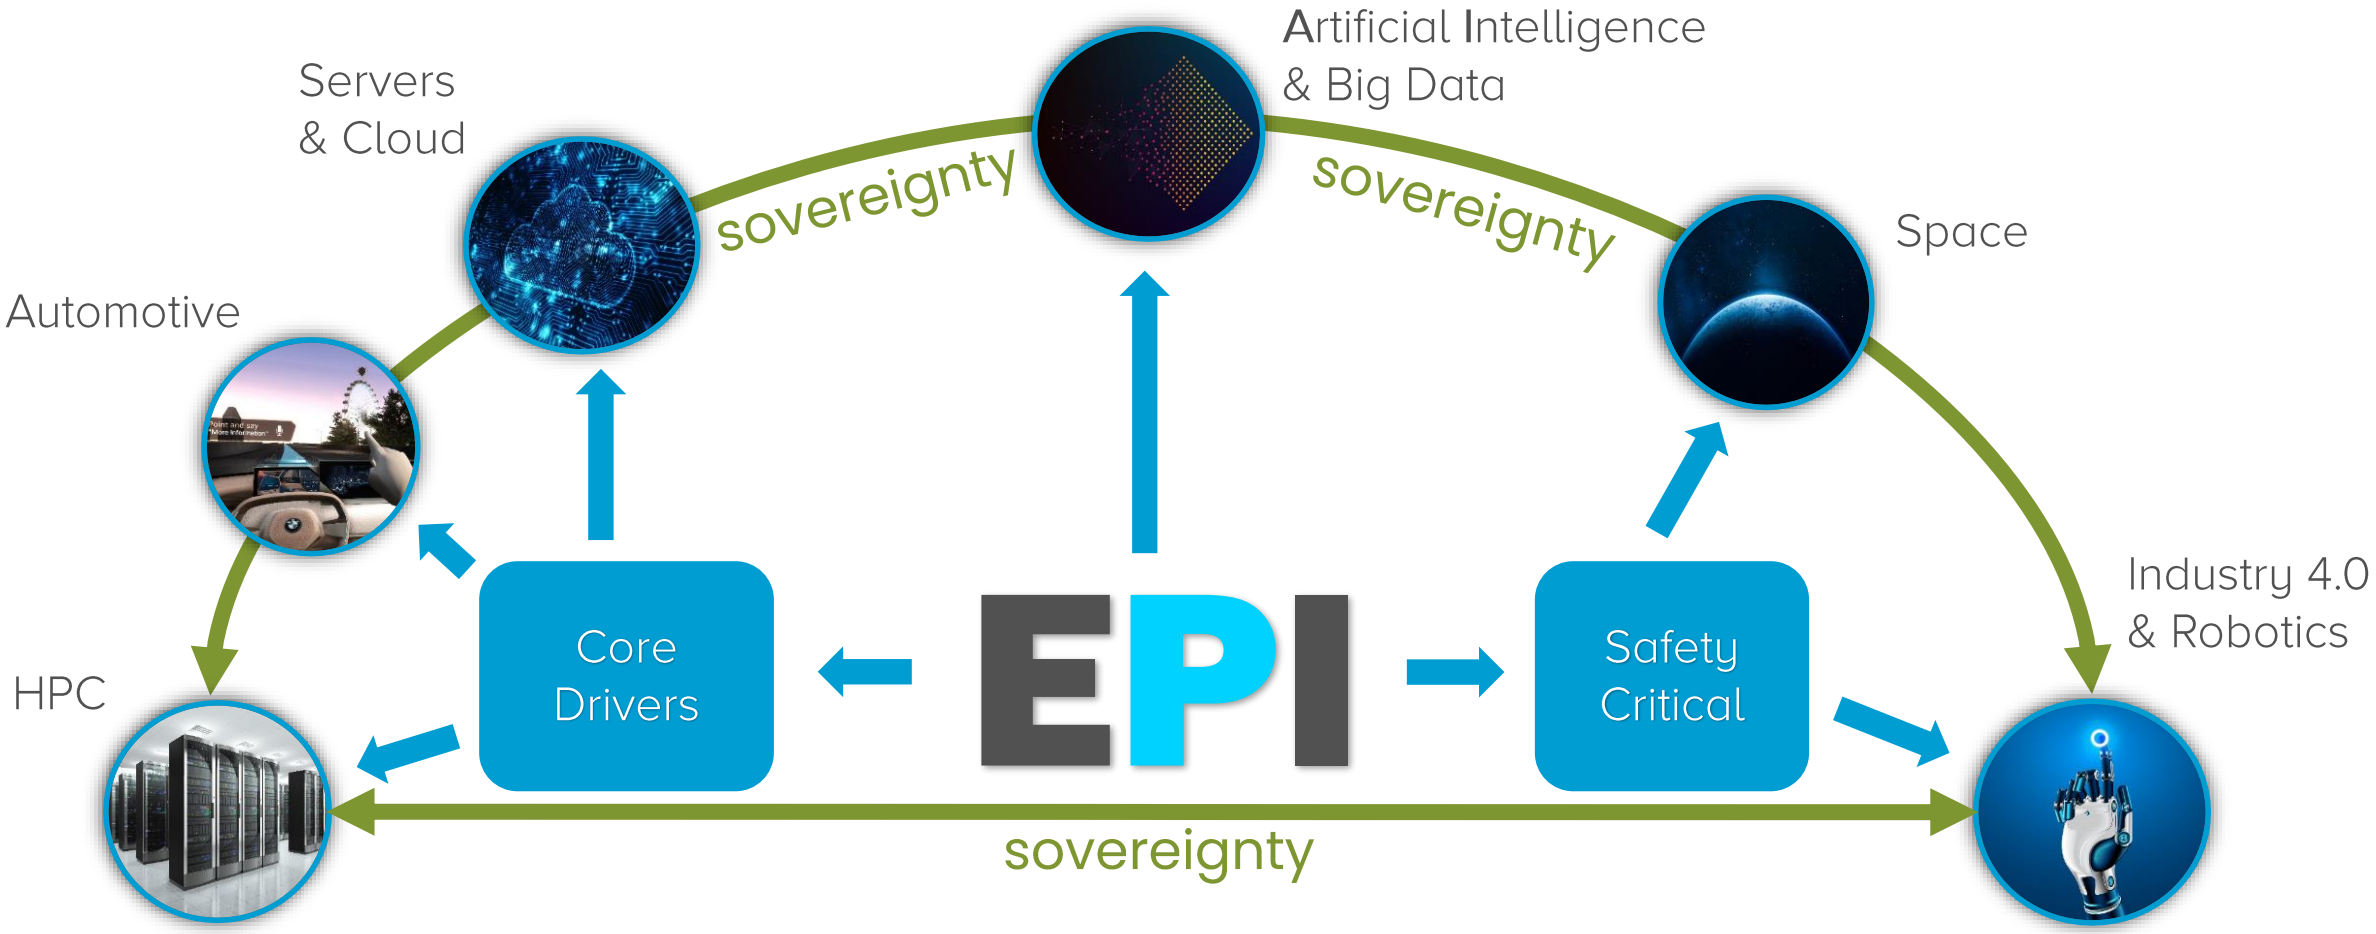

CONCLUSION

- 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

PROCESSOR AND GENERAL ARCHITECTURE

DENIS DUTOIT – YINGCHIH YANG – ANDREA BARTOLINI – PATRICE HAMEAU

GENERAL ARCHITECTURE

- Memory-coherent NoC connects
  - Array of computing units (CU)
  - Memory and I/O controllers
  - Bridge to links
- High speed links
  - D2D links to connect on-package dies
  - HSL links to connect on-board packages
- Top level infrastructures
  - Power management & controller
  - Security

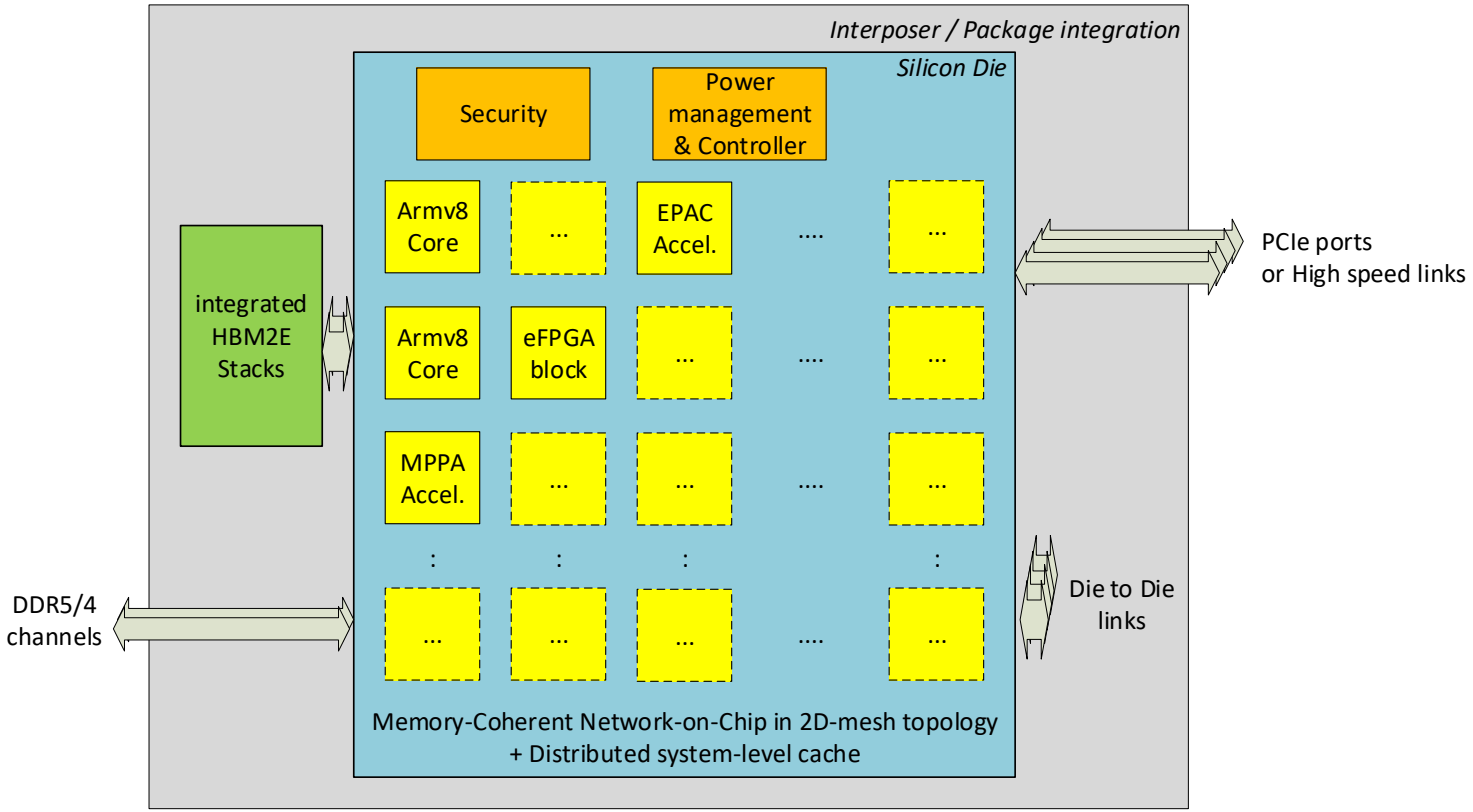

NoC: network on chip
HSL: High speed links (with memory coherent support)

GENERAL ARCHITECTURE (#2)

- Interfaces to connect acceleration functions to the NoC
  - Data access and sharing throught 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

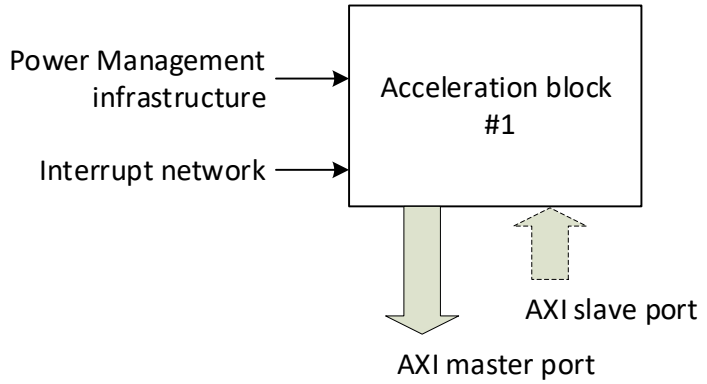

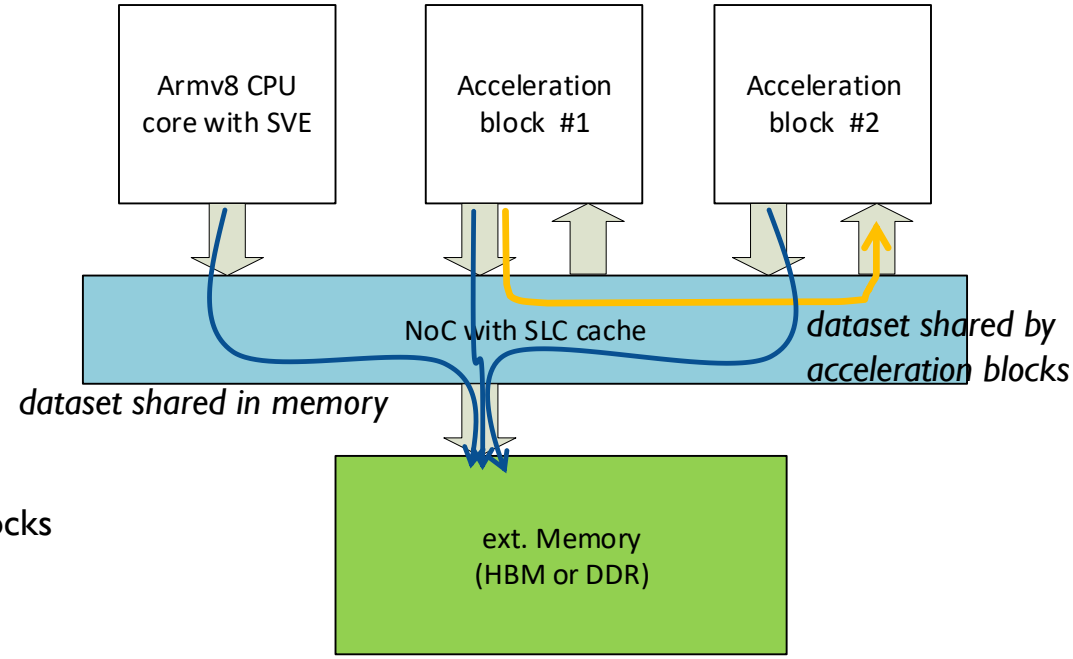

CU: Computing Unit; either Armv8 core with SVE or the EPAC/MPPA acceleration blocks
SLC: System Level Cache; a last-level cache before ext. memories

HETEROGENEOUS INTEGRATION

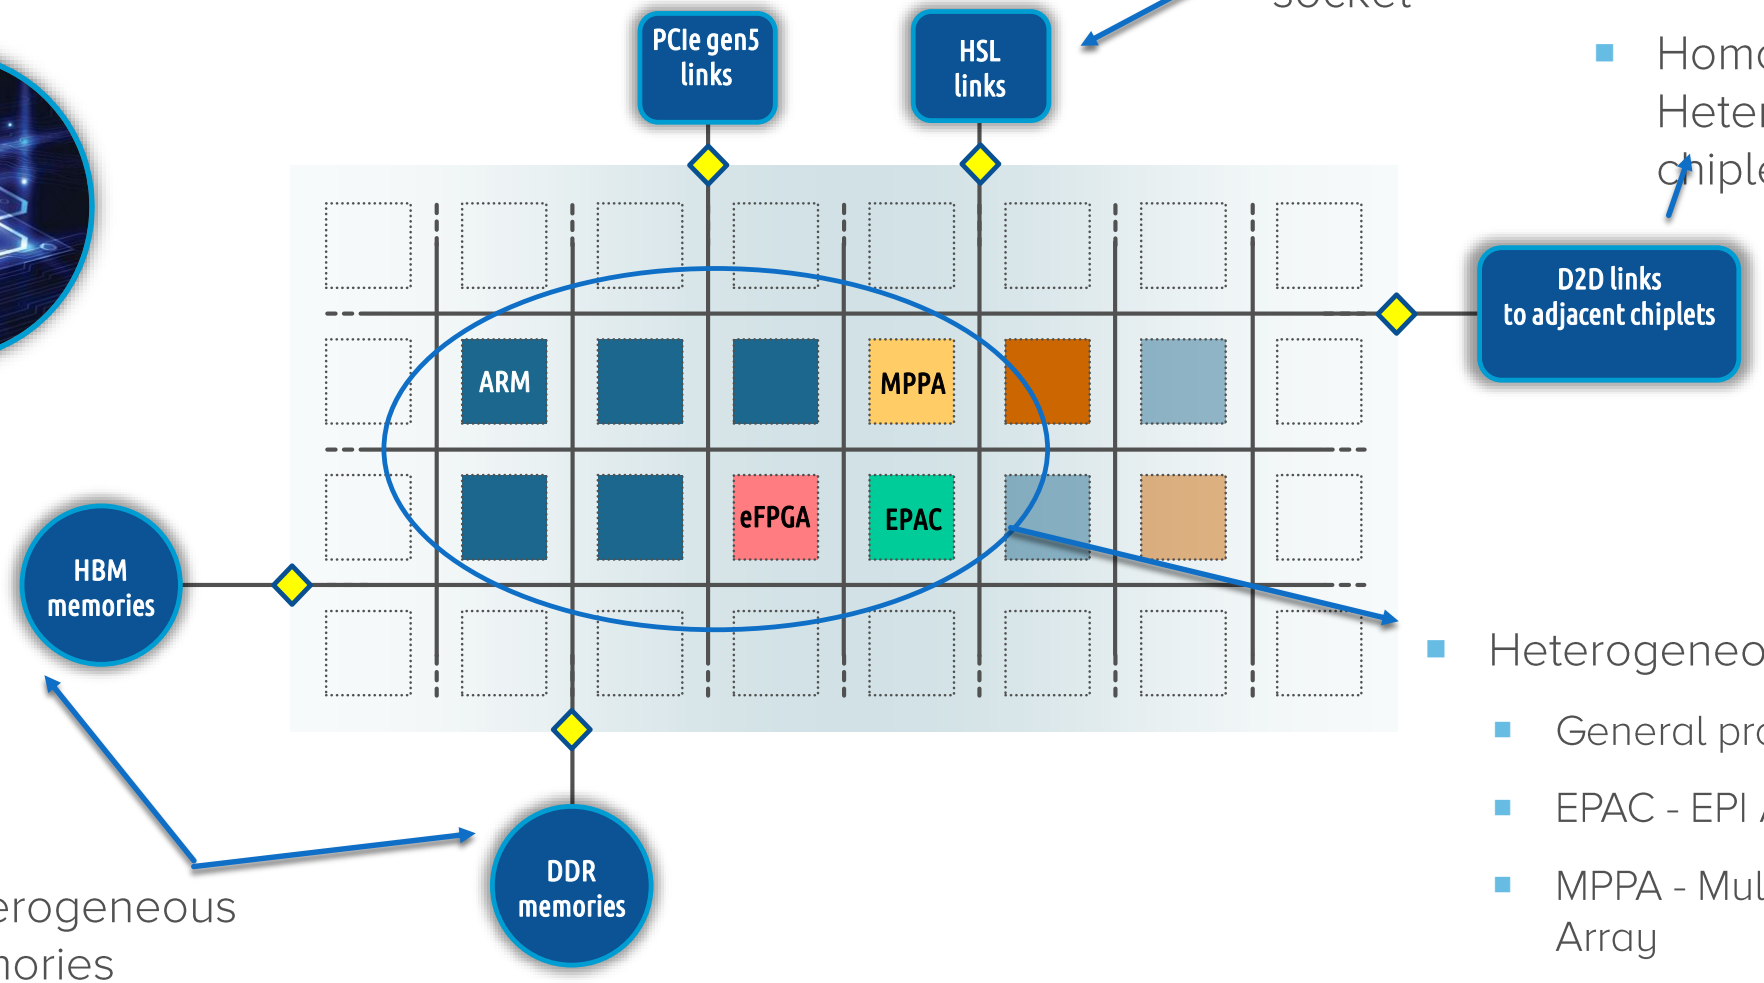

- Heterogeneous socket
- Homogenegous/ Heterogeneous chiplet
- Heterogeneous SoC:
  - General processing core
  - EPAC - EPI Accelerator
  - MPPA - Multi-Purpose Processing Array
  - eFPGA - embedded FPGA
- Heterogeneous memories

WP3 – POWER MANAGEMENT & CONTROLLER

- Out-of-band – zero overhead telemetry
- Node Pcap – Max perf @ Pnode<Pmax
- RAS – error and conditions reporting
  - cpufreq/ cpuidle
  - Based on O.S. metrics
  - Slow & often unused

System Management / RM

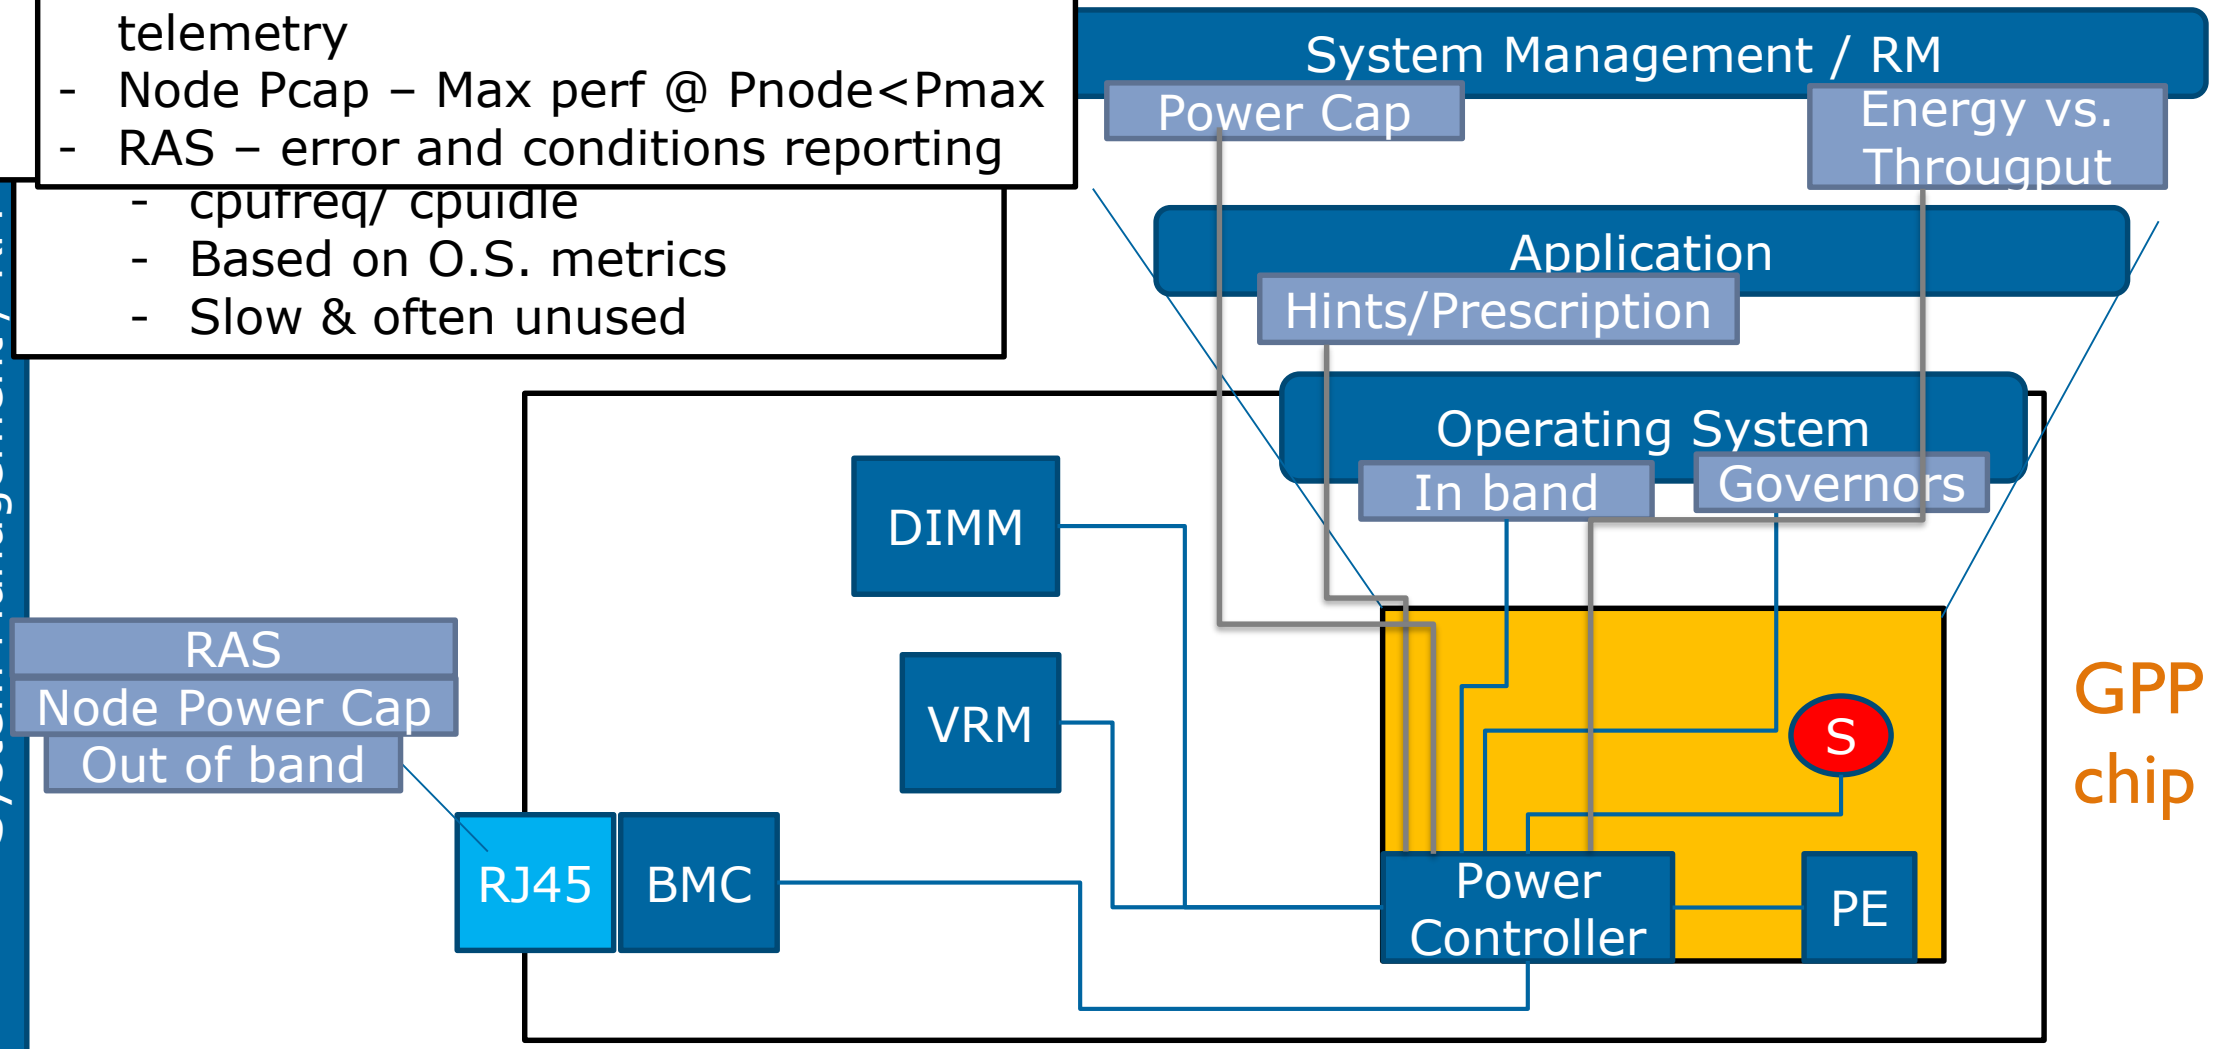

POWER MANAGEMENT SOA & REQUIREMENTS

	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 is powered UNIBO and targets:

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

SECURITY REQUIREMENTS

- Security aspects taken into account since the start of architecture definition:
  - Strong Root of Trust (EAL4+ embedded Secure Element) always active during all chip lifecycle.
  - Global security policy and chip monitoring can be managed independently of applications.
  - Modular security levels to adapt to different applicative needs.
  - Level can be raised to be compatible with automotive needs for security (including safety constraints)
  - Future proof considerations for cryptographic assets (including post-quantic)
  - Core security services made available to applications through API (mailbox) to ease independent and quick security update.

SECURITY IN EPI PROCESSOR

- Root of trust
- Secure Boot
- Security domains
- Security services isolated
- Advanced cryptographic functions

- Various monitors for fault-injection, physical intrusion and other conditions

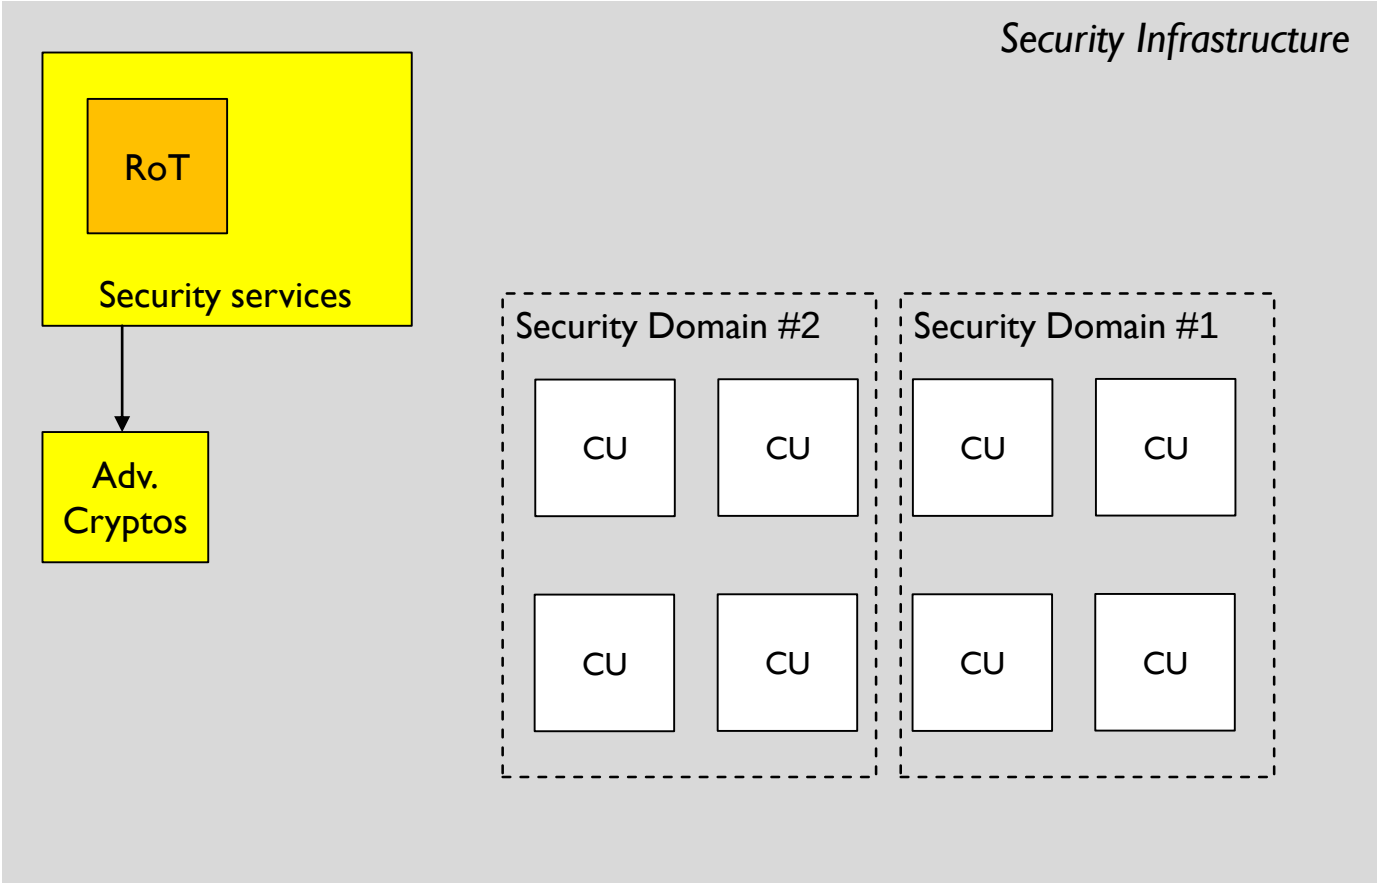

CO-DESIGN PROCESS AND MODELING

DENIS DUTOIT – ESTELA SUAREZ – NICOLAS VENTROUX

EPI CO-DESIGN

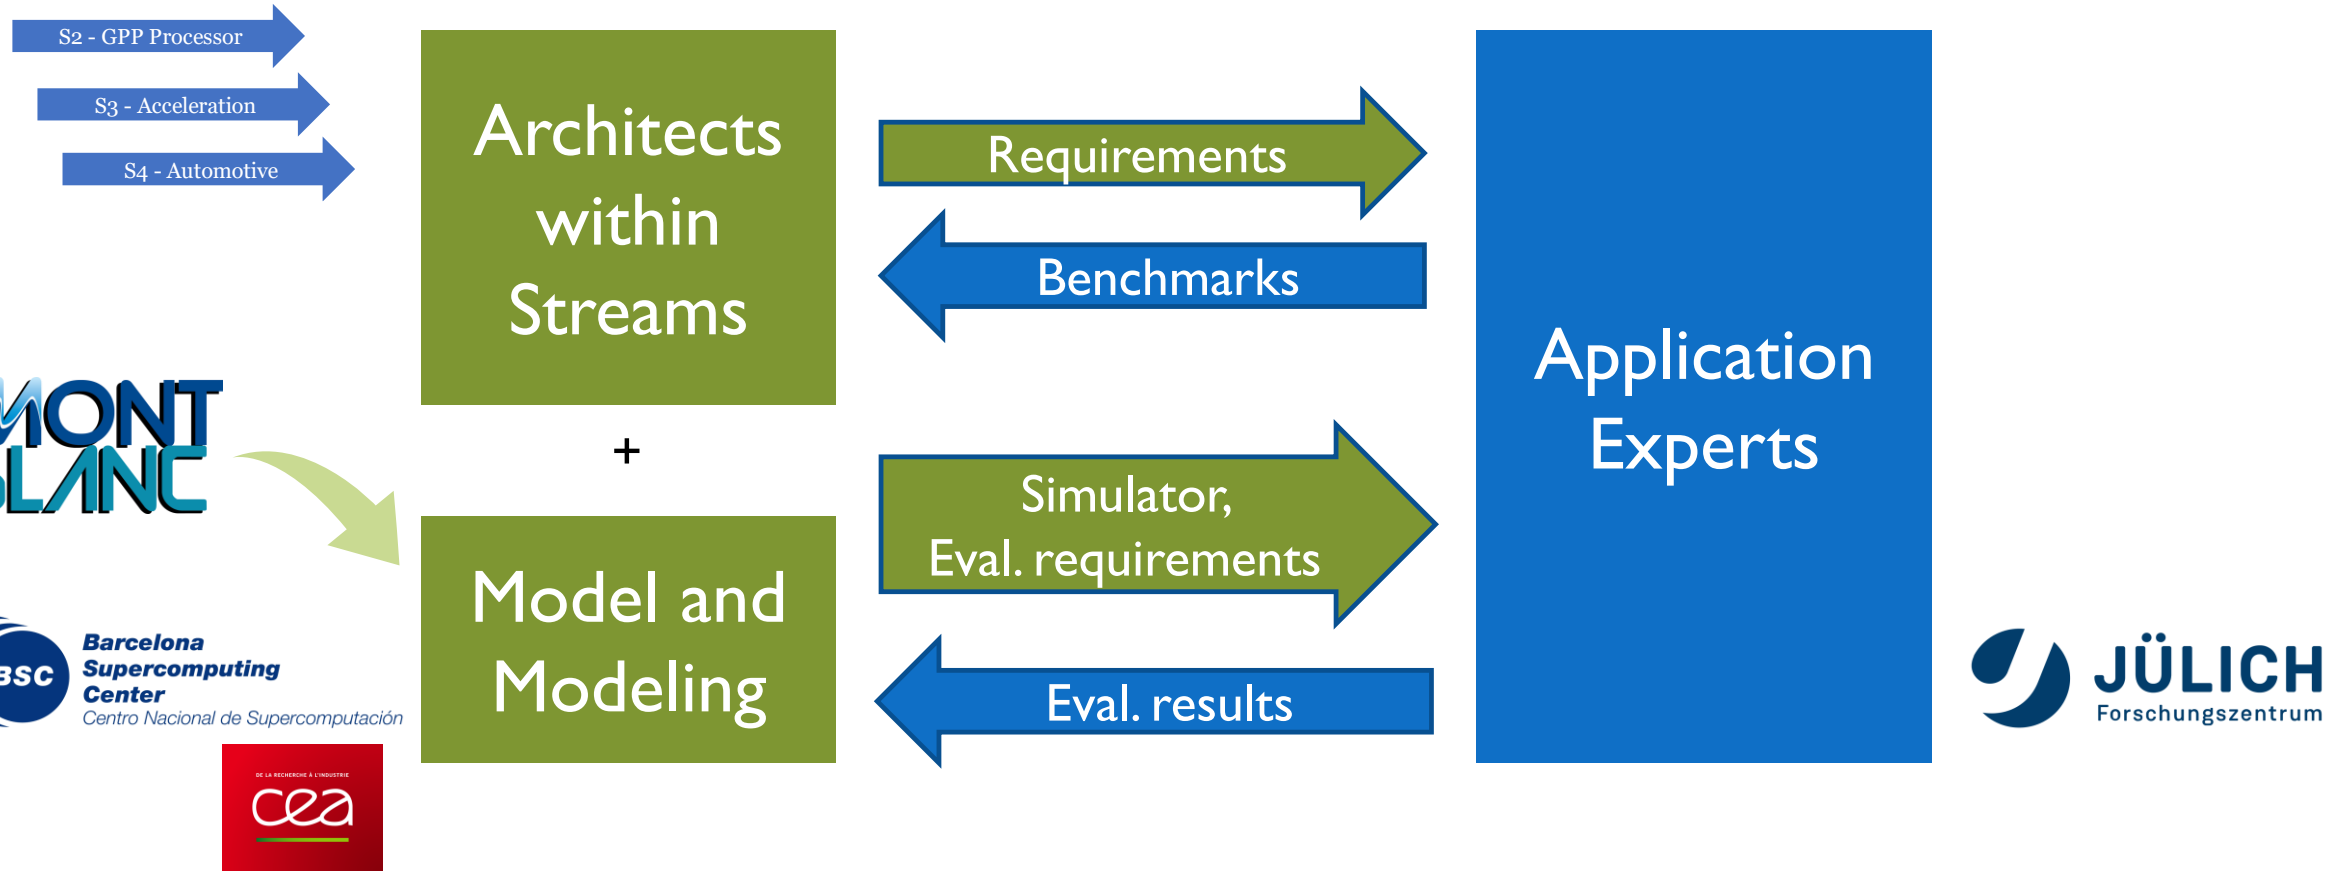

CO-DESIGN PROCESS

- Bi-directional and iterative interaction process between:
  - application experts and
  - hardware (HW) and system-software (SW) developers
- Multi-level suite of benchmarks
  - from very low-level synthetic benchmarks to high-level applications
- Methodology with multi-level models & simulators
  - 1). analytical models, high level
  - 2). simulation based (e.g. gem5 simulation engine)
  - 3). reference platform (e.g. Marvell ThunderX2)
- Node-level co-design parameters (e.g.):
  - GPP: SVE length, number of SVE pipelines per core
  - Accel.: vector registers length, ratio accelerator-vs-GPP cores
  - Memory (size and BW): Cache, HBM, DDR

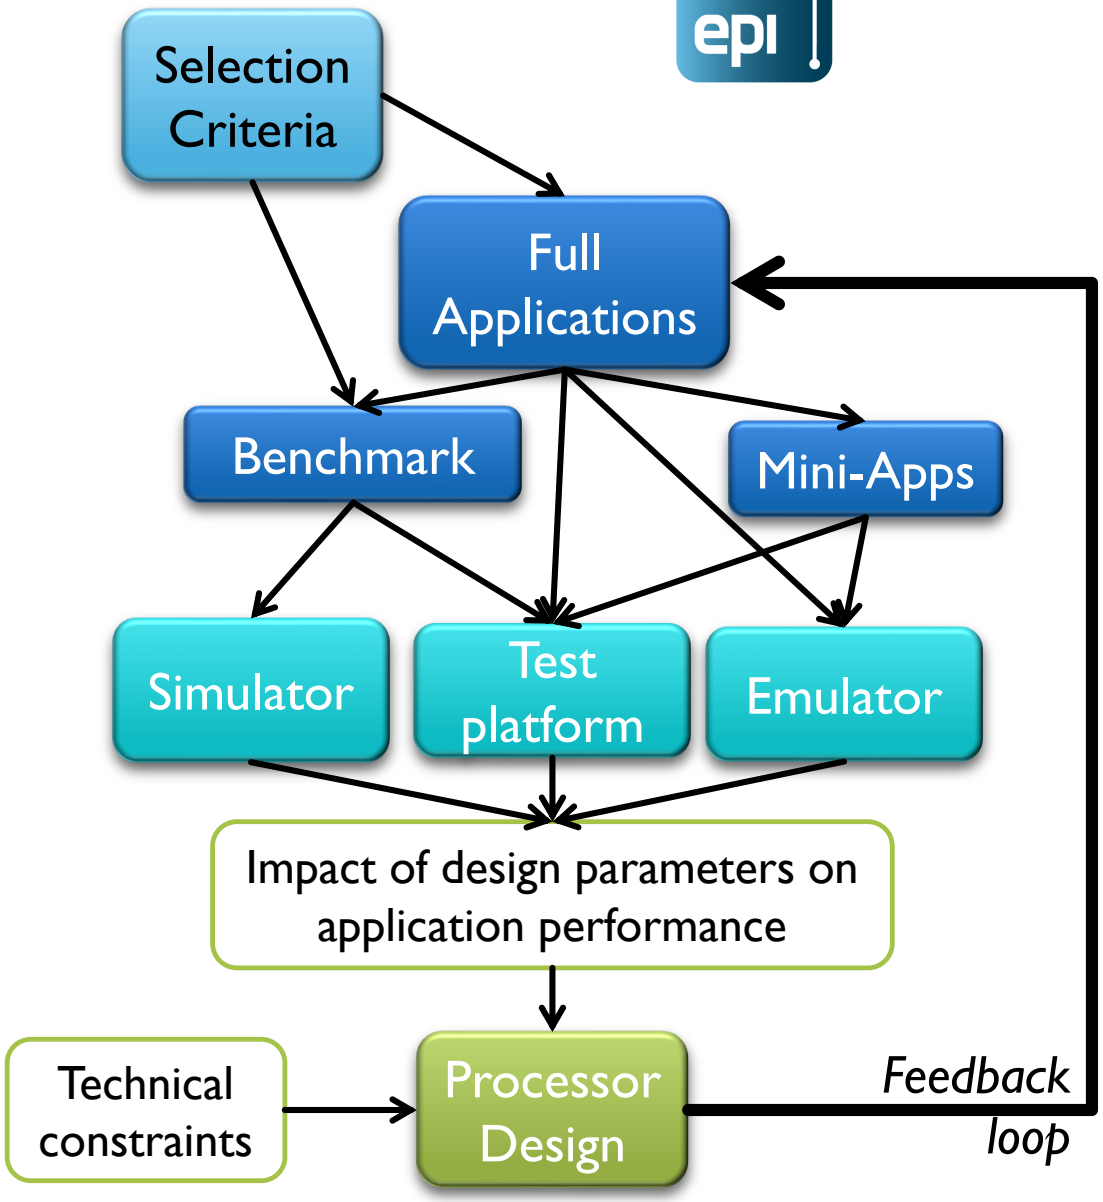

APPLICATION SELECTION

Selection Criteria

C1	Relevant (now or in 5-years) markets
C2	Its requirements covers architectural components
C3	Represent a family/class of applications
C4	Close relation to code developers
C5	Licence allows development of mini-apps/derived benchmarks (preferable OpenSource)
C6	Reference data available from other platforms
C7	Application uses/covers a software component/ programming model relevant for an EPI market
C8	Application features relatively simple kernel
C9	High societal impact
C10	Part of an existing benchmark suite, or widely known
C11	Mini-app or kernels already available

Application Fields

- Biophysics
- Biology/Medicine
- Earth Sciences/Climate
- HEP & Fusion
- Material Sciences
- CFD
- Hydrodynamics
- PDE
- Image / Media
- Automotive
- Cryptography
- HPDA
- Machine Learning
- Deep learning
- Cloud
- Data Base
- Reference benchmarks (HPL, HPCG, Stream, DGEMM…)

EXAMPLE: DGEMM

L2 cache exploration with BLIS microkernel

CREDITS:

-P.Petrakis, V. Papaefstathiou et al. (FORTH): simulation execution an analysis
-B.Brank, S.Nassyr (FZJ): BLIS micro-kernel
-A.Portero (FZJ): Gem5 simulator setup

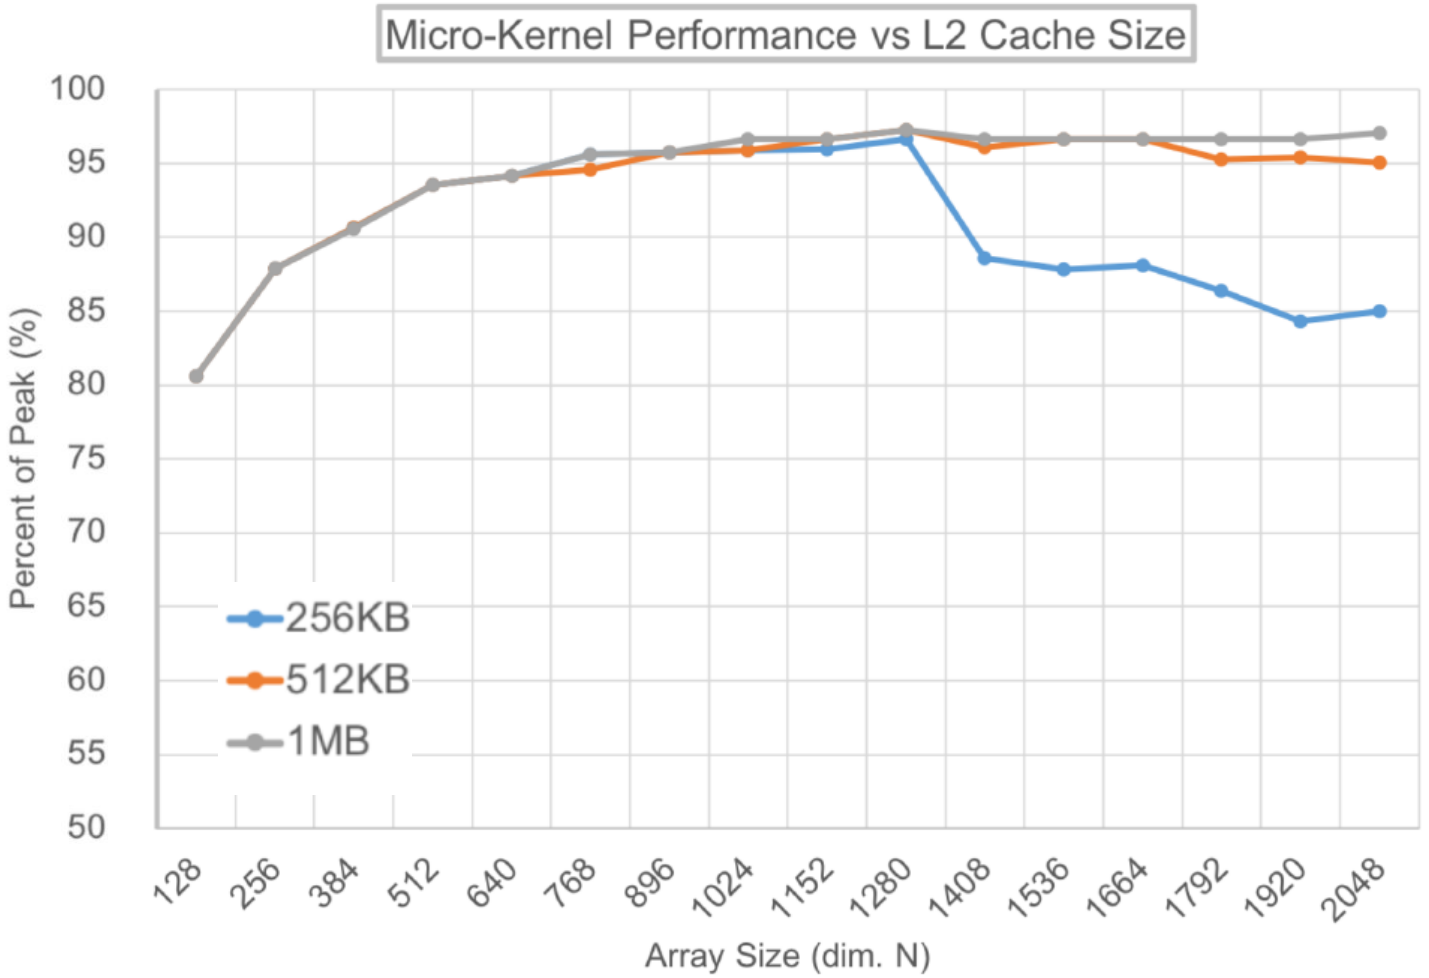

SIMULATION ACTIVITIES IN EPI

- Why simulation tasks in EPI?
  - To make available high-level simulation tools to partners during the project
  - Develop advanced innovative solutions for simulation
- Two technologies will be tuned for the EPI
  - Early stage DSE with MUSA simulation (BSC)
    - Early-stage performance prediction
    - Trace-based simulators
  - Virtual prototypes design with SESAM (CEA)
    - Fast and parallel High-level SystemC/TLM simulation of manycore chips
    - DSE and early SW development and validation

SESAM/VPSIM

- Virtual prototyping framework
- Large and flexible IP portfolio
  - Supports models from native library & ext. providers (QEMU 3.1, ARM Fast Models, NoCs…)
- SystemC / TLM 2.0 simulations
  - Co-simulation & co-emulation for RTL valid.
  - Co-simulation with 3rd-party tools (FMI standard)
- Designed to run full software stacks
  - Firmware, bootloader, hypervisor, OS kernel, user applications
  - Develop, debug, profile using std. tools

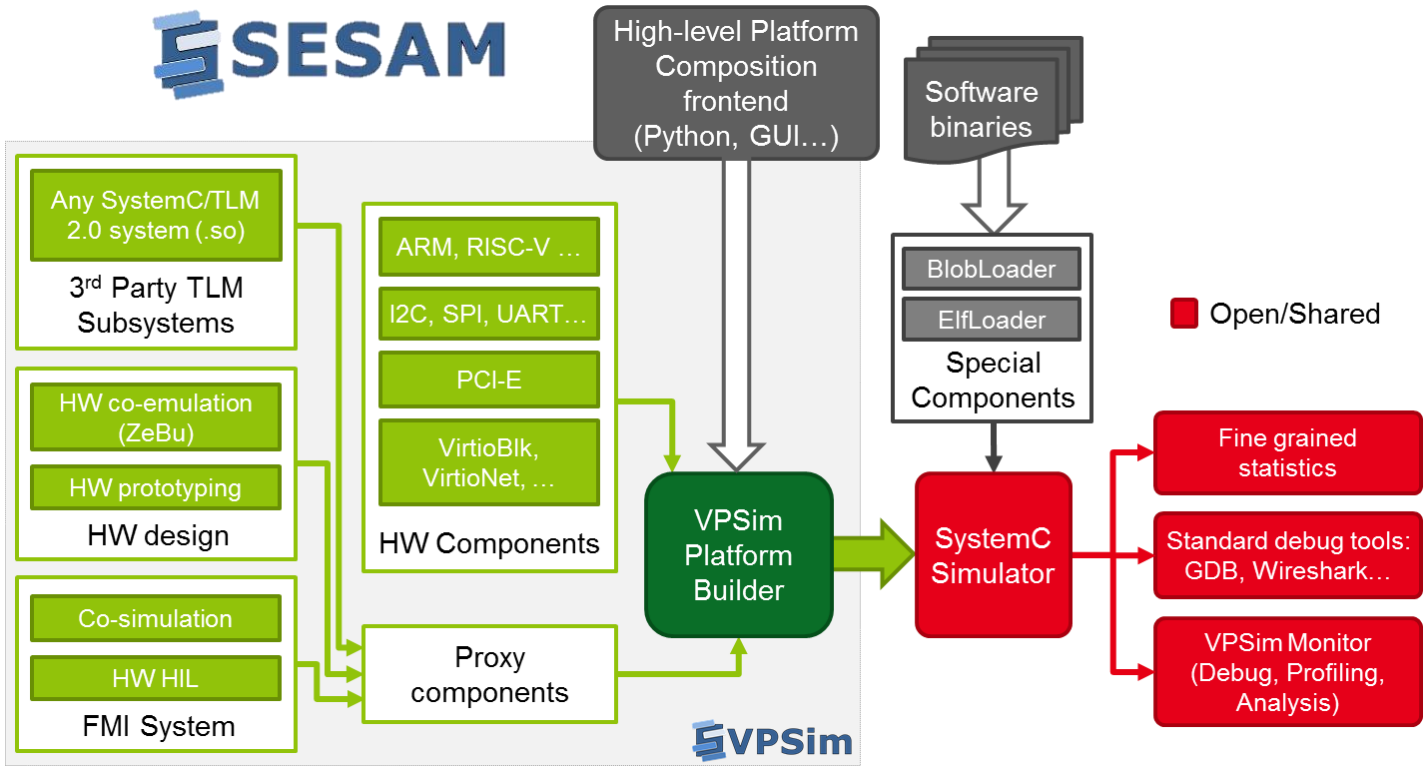

MUSA

- Usage/objectives
  - Combine detailed trace driven simulation with sampling strategies
  - Exploration of HPC architectural parameters affecting performance at scale
- A multi-level approach
  - First level: trace generation
    - OpenMP runtime system plugin
    - MPI call instrumentation
    - Pintool / DynamoRIO
  - Second level: simulation
    - A network simulator plays the obtained traces
    - Multi-core detailed simulation of a set of computation phases
    - Extrapolation of results for large scale HPC
  - Support different modes
    - Burst, detailed, sampling) trading accuracy for speed