THE EUROPEAN APPROACH FOR EXASCALE AGES

THE ROAD TOWARD SOVEREIGNTY

Jean-marc.denis@European-processor-initiative.eu
Chairman of the Board

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

AGENDA

EPI in European HPC strategy

(post)-Exascale supercomputers specifications

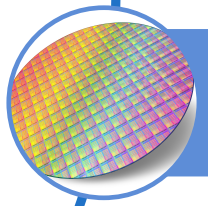

Consequences on compute units architecture

Wrap-Up

SECTION 1
EPI IN EUROPEAN HPC STRATEGY

EUROPEAN CONTEXT

European Commission President Jean-Claude Juncker

Paris, 27 October 2015

« Our ambition is for Europe to become one of the top 3 world leaders in high-performance computing by 2020 »

Creation of the European processor Initiative

Brussels, 1 Dec 2017

23 members from 10 EU countries

- General Purpose processor in 2022
- Accelerator IP

Thierry Breton, Commissioner for Internal Market

Brussels, December, 7th, 2020

- « Europe has all it takes to diversify and reduce critical dependencies, while remaining open. We will therefore need to set ambitious plans, from design of chips to advanced manufacturing progressing towards 2nm nodes […]
- **Budget 2021-2023 up to 145B€**

EuroHPC Joint Undertaking

Vice President Andrus Ansip

« I encourage even more EU countries to engage in this ambitious endeavour »

Digital Day Rome, 23 March 2017

Ministers from seven MS (France, Germany, Italy, Luxembourg, Netherlands, Portugal and Spain) sign a declaration to support the next generation of computing and data infrastructures

Ursula Von Der Lyen State of the Union

Brussels – September, 16th, 2020

- Investment **of 8 billion euros in the next generation of supercomputers** - cutting-edge technology made in Europe.
- **The European industry will develop our own next-generation microprocessor**

EUROHPC JOINT UNDERTAKING (JU)

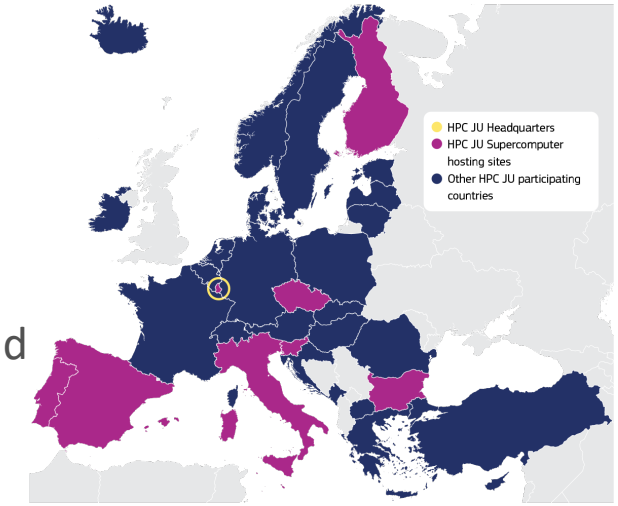

- The European High Performance Computing Joint Undertaking (EuroHPC JU) is pooling European resources to buy and deploy top-of-the-range supercomputers and develop innovative exascale supercomputing technologies and applications.

- It aims to improve quality of life, advance science, boost industrial competitiveness, and **ensure Europe's technological autonomy**.
- The JU is currently supporting two main activities (2020-2021):
  - **Developing a pan-European supercomputing infrastructure**: 3 pre-exascale supercomputers (aim to be among the top 5 WW), and 5 petascale supercomputers. Benefit European private and public users, working in academia and industry, everywhere in Europe.
  - **Supporting research and innovation activities**: developing a European supercomputing ecosystem, stimulating a technology supply industry (from low-power processors to software and middleware, and their integration into supercomputing systems), and making supercomputing resources in many application areas available to a large number of public and private users, including small and medium-sized enterprises.

RECENT NEWS FROM EUROHPC

>1,000 Pflops
38 Pflops
>2 EFlops

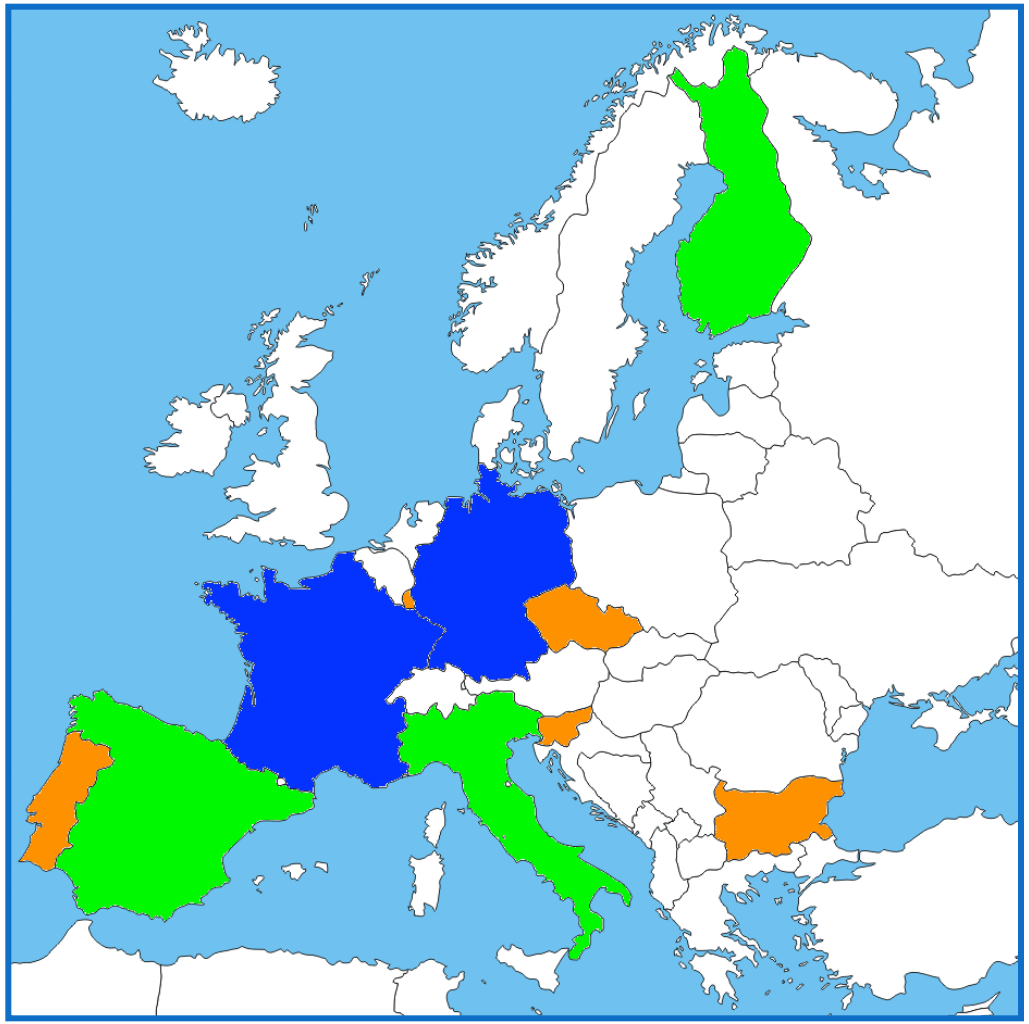

Country	Machine	Supplier	PFLOPS	Year
Finland	LUMI	HPE	550	2021/22
Italy	Leonardo	ATOS	248	2021
Spain(*)	*MareNostrum5*	*TBD*	*>200*	*2022*
Luxembourg	MeluXina	ATOS	10	2021
Portugal	Deucalion	Fujitsu ATOS	10	2021
CZ Rep	IT4I (name tbd)	HPE	15,2	2021
Bulgaria	NCSA	ATOS	6	2021
Slovenia	Vega	ATOS	6,8	2021
TBD (DE?)	*TBD(**)*	*TBD*	*>1,000*	*2023*
TBD (FR?)	*TBD(**)*	*TBD*	*>1,000*	*2024*

(*) announced in 2021/Q1
(**) with EU processor based on EPI delivrables

EPI OBJECTIVES

- **Overall: Develop a complete EU designed high-end microprocessor, addressing Supercomputing and edge-HPC segments.**
- Short-term objective
  - supply the EU-designed microprocessor to empower the EU Exascale machines
- Long-term objective
  - Europe needs a sovereign (=not at risk of limitation or embargo by non-EU countries) access to high-performance, low-power microprocessors, from IP to products
- EPI has been set to fulfil this objective
- EPI has to cover all Technical Readiness levels (TRL)
  - TRL 1-5 are for long-term objectives (EU IP)

    and
  - TRL 6-9 are for short to mid-term objectives (decade) with products designed in EU

27 PARTNERS FROM 10 EU COUNTRIES

FROM **IPR** TO **PRODUCTS**
FROM **EPI** TO **SIPEARL**

- SIPEARL is
  - Incorporated in EU (France)
  - the industrial and business ‘hand’ of EPI
  - the Fabless company
- licence of IPs from the partners
- develop own IPs around it
- licence the missing components from the market
- Raise in equity the missing budget
- generate revenue from both the HPC, IA, server and eHPC markets
- integrate, market, support & sell the chip
- work on the next generations

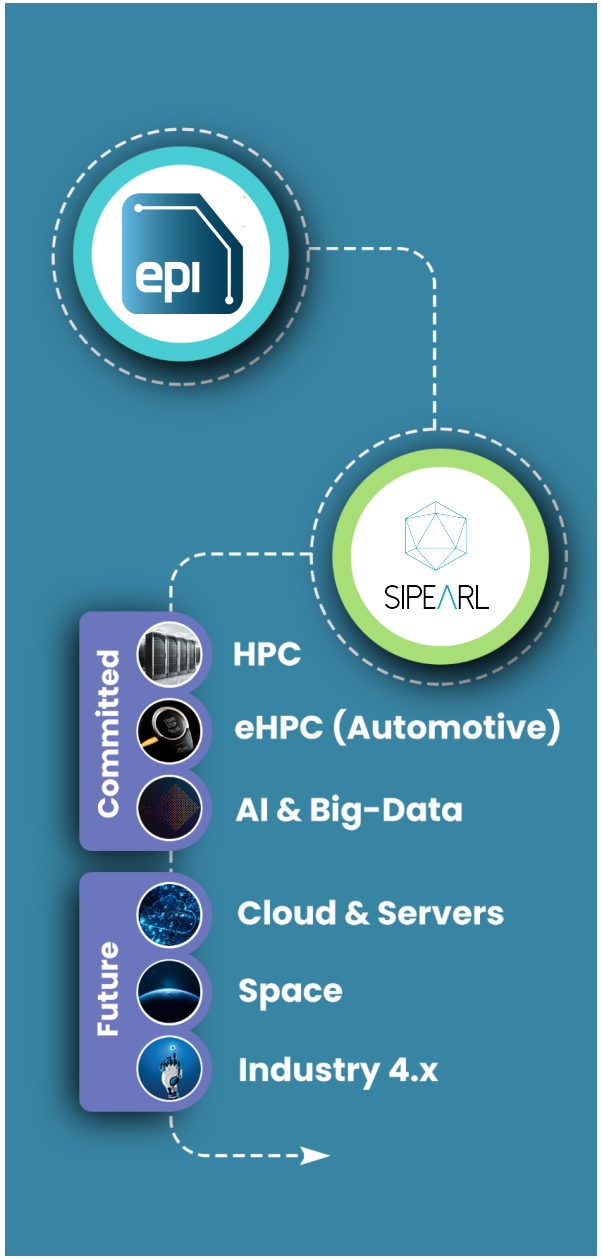

OVERALL ROADMAP

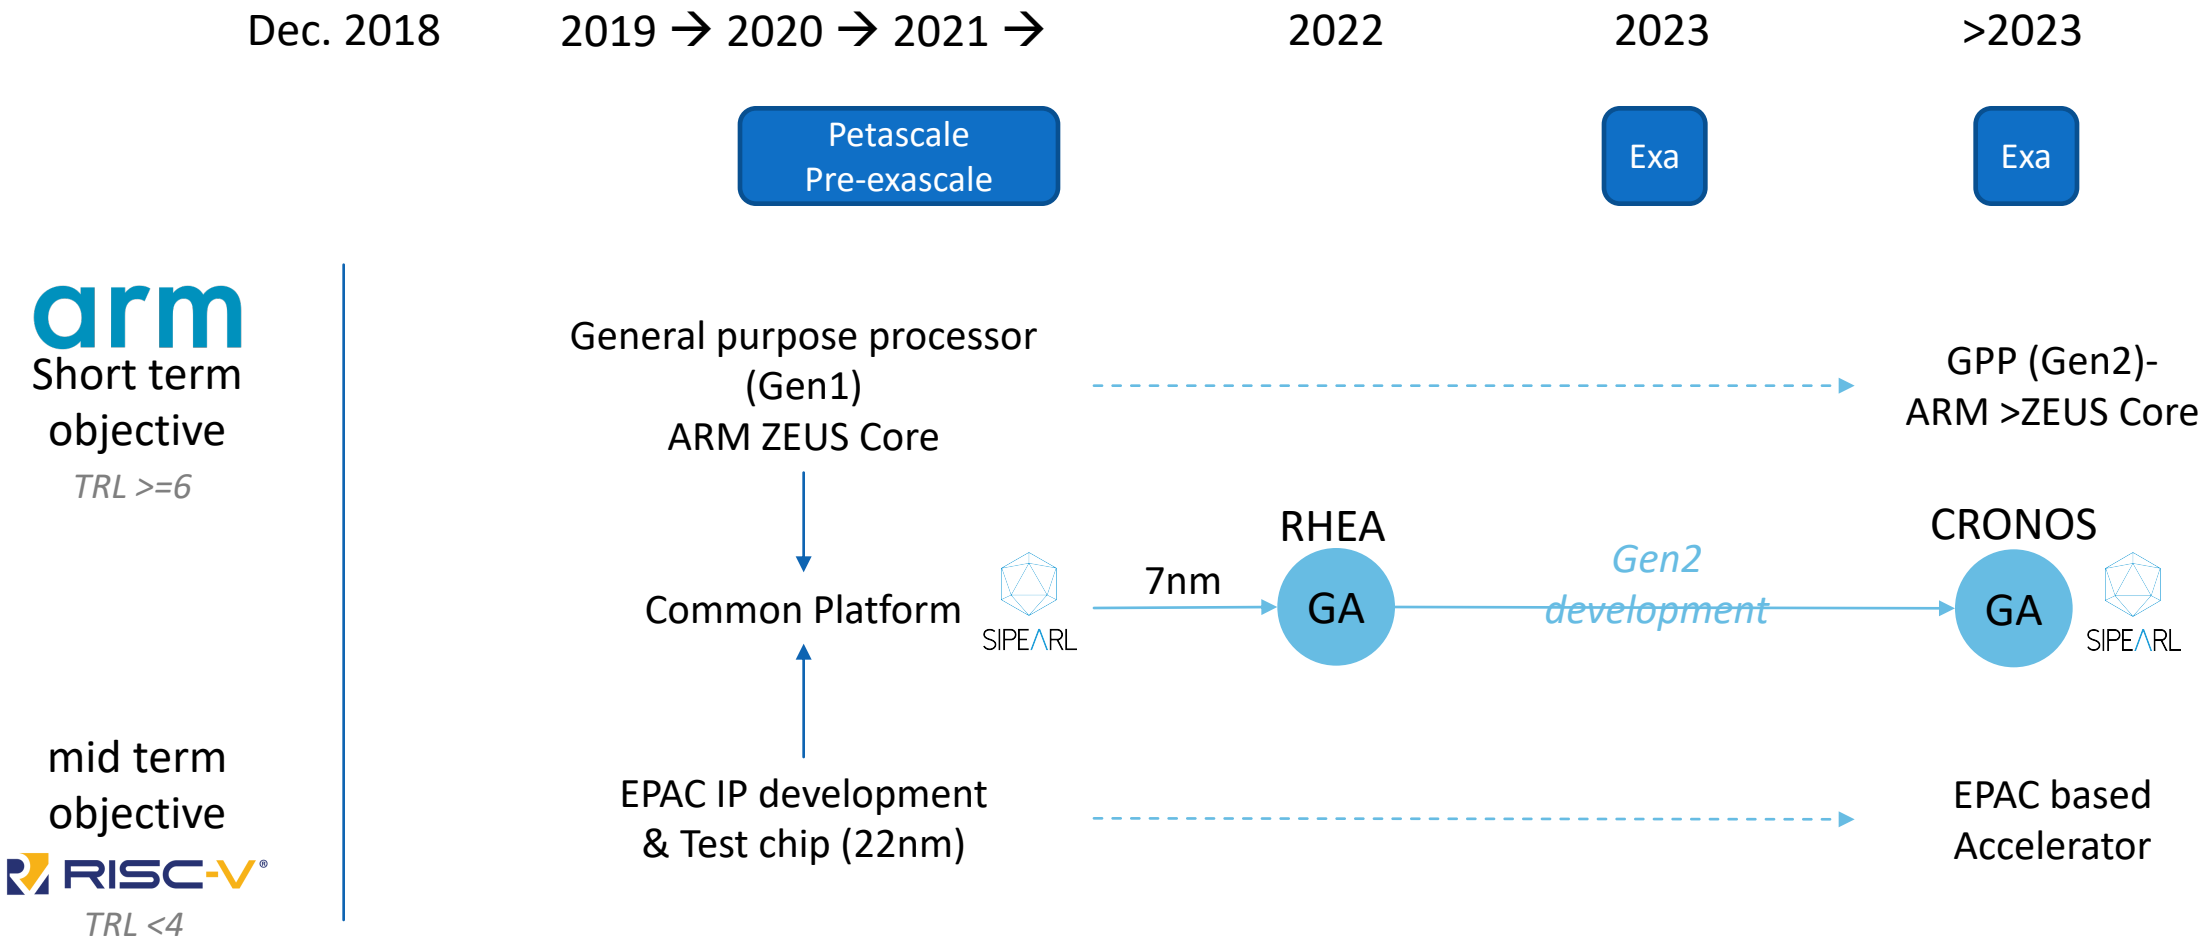

SECTION 2
(POST)-EXASCALE SUPERCOMPUTERS SPECIFICATIONS

HPC BEFORE ARTIFICIAL INTELLIGENCE

Theoretical model → **HPC Application** → **Results**

$$\frac{\partial u}{\partial t} = k \frac{\partial^2 u}{\partial x^2}$$

```
      C
            PIN=0.02
            IF(DDT.NE.0.0)  THEN
            DT=DDT
            ELSE
            DT=PIN
            ENDIF
            WRITE(*,'(A)') '    PLEASE ENTER NAME OF OUTPUT FILE (FOR EXAMPLE
     *  B:ZZ.DAT)'
            READ(*,'(A)') FNAMEO
            OPEN(6,FILE=FNAMEO,STATUS='UNKNOWN')
            PV=WFLX/TH
            RS=NEQ*ROU*KD/TH
            CO=CS
      C
            TIME=0.0D0
            EF=0.0D0
      5     CONTINUE
            GAMMA=DT/(2.D0*DX*DX)
            BETA=DT/DX
            IF((BETA*PV).GT.0.50D0) GO TO 7
            IF((GAMMA*D/(BETA*PV)).LT.0.5D0) GO TO 6
            GO TO 8
      6     DX=DX/2
            GO TO 5
      7     DT=DT/2
            GO TO 5
      8     CONTINUE
            N=COL/DX
            NM1=N-1
            NM2=N-2
            NP1=N+1
            GAMMA=DT/(2*DX*DX)
```

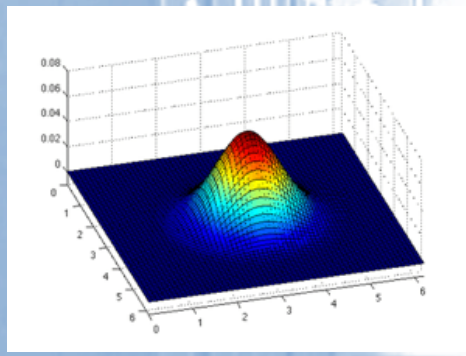

HPC WITH ARTIFICIAL INTELLIGENCE

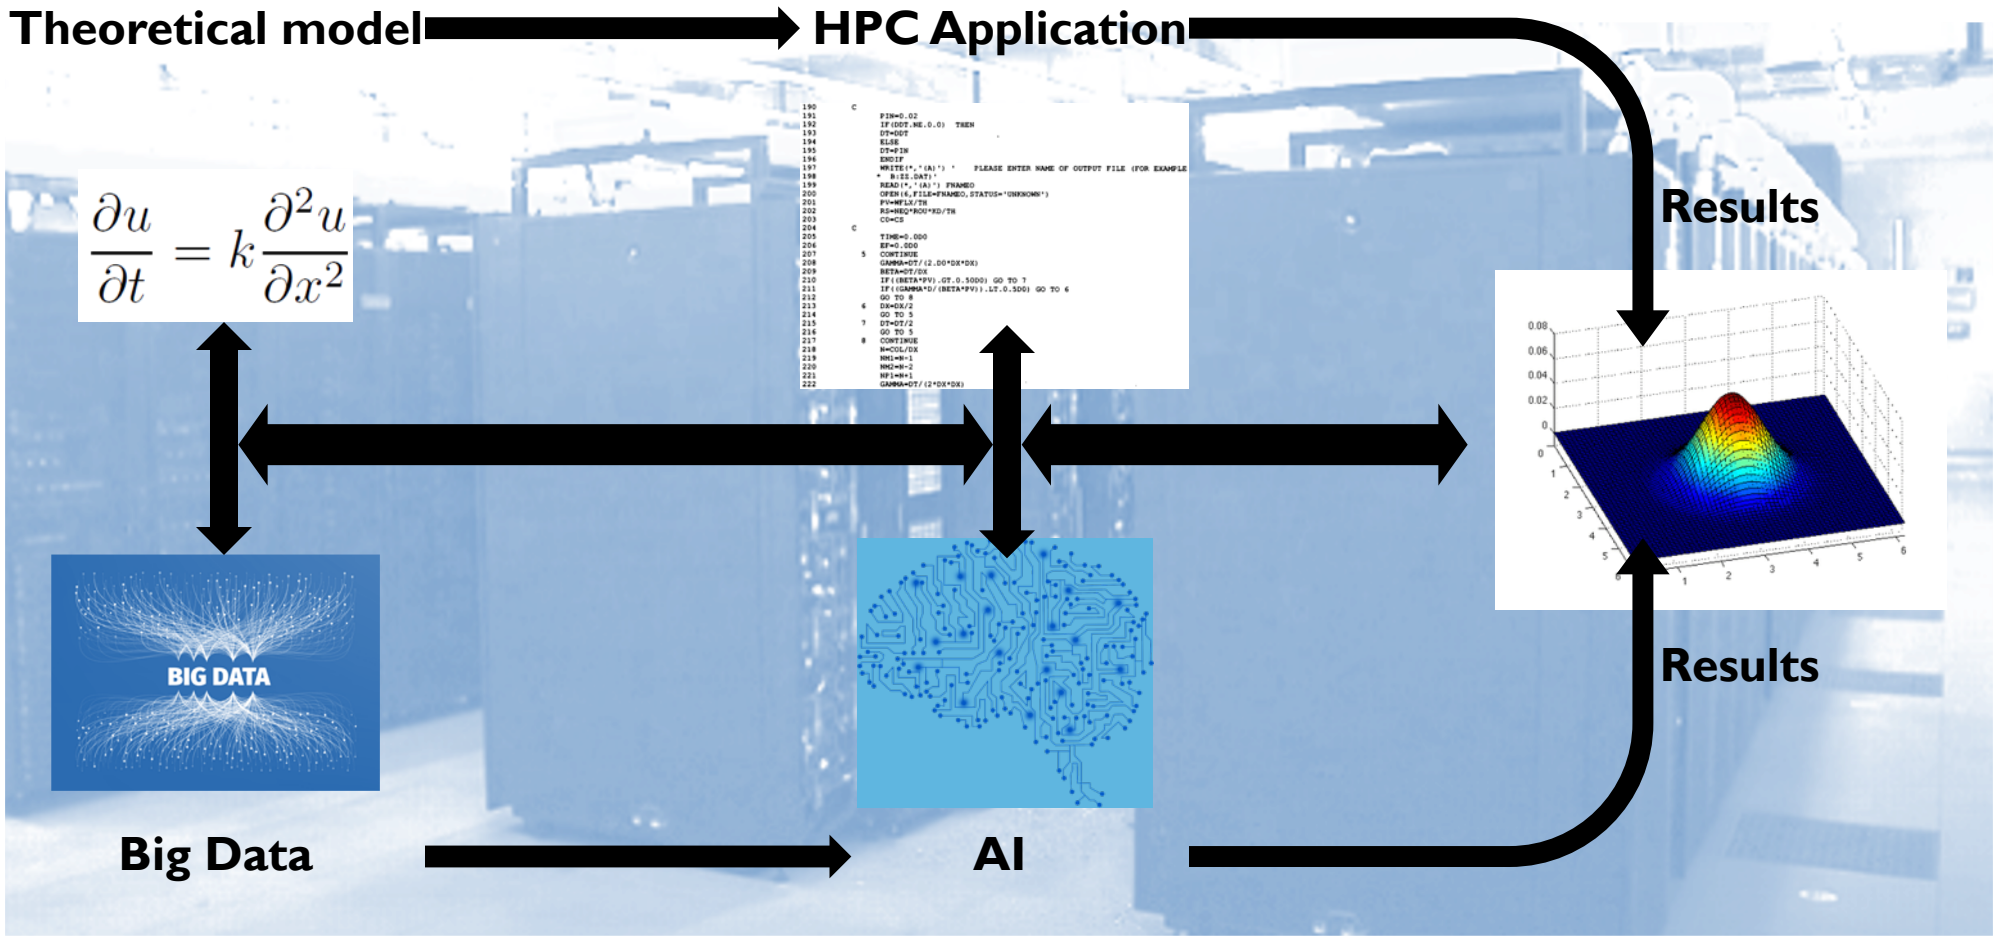

EXAMPLE: EXTREME PREDICTIONS WORKFLOW

Source: The extremes prediction use case, by Peter Bauer, Marc Duranton and Michael Malms, © HiPEAC VISION 2021

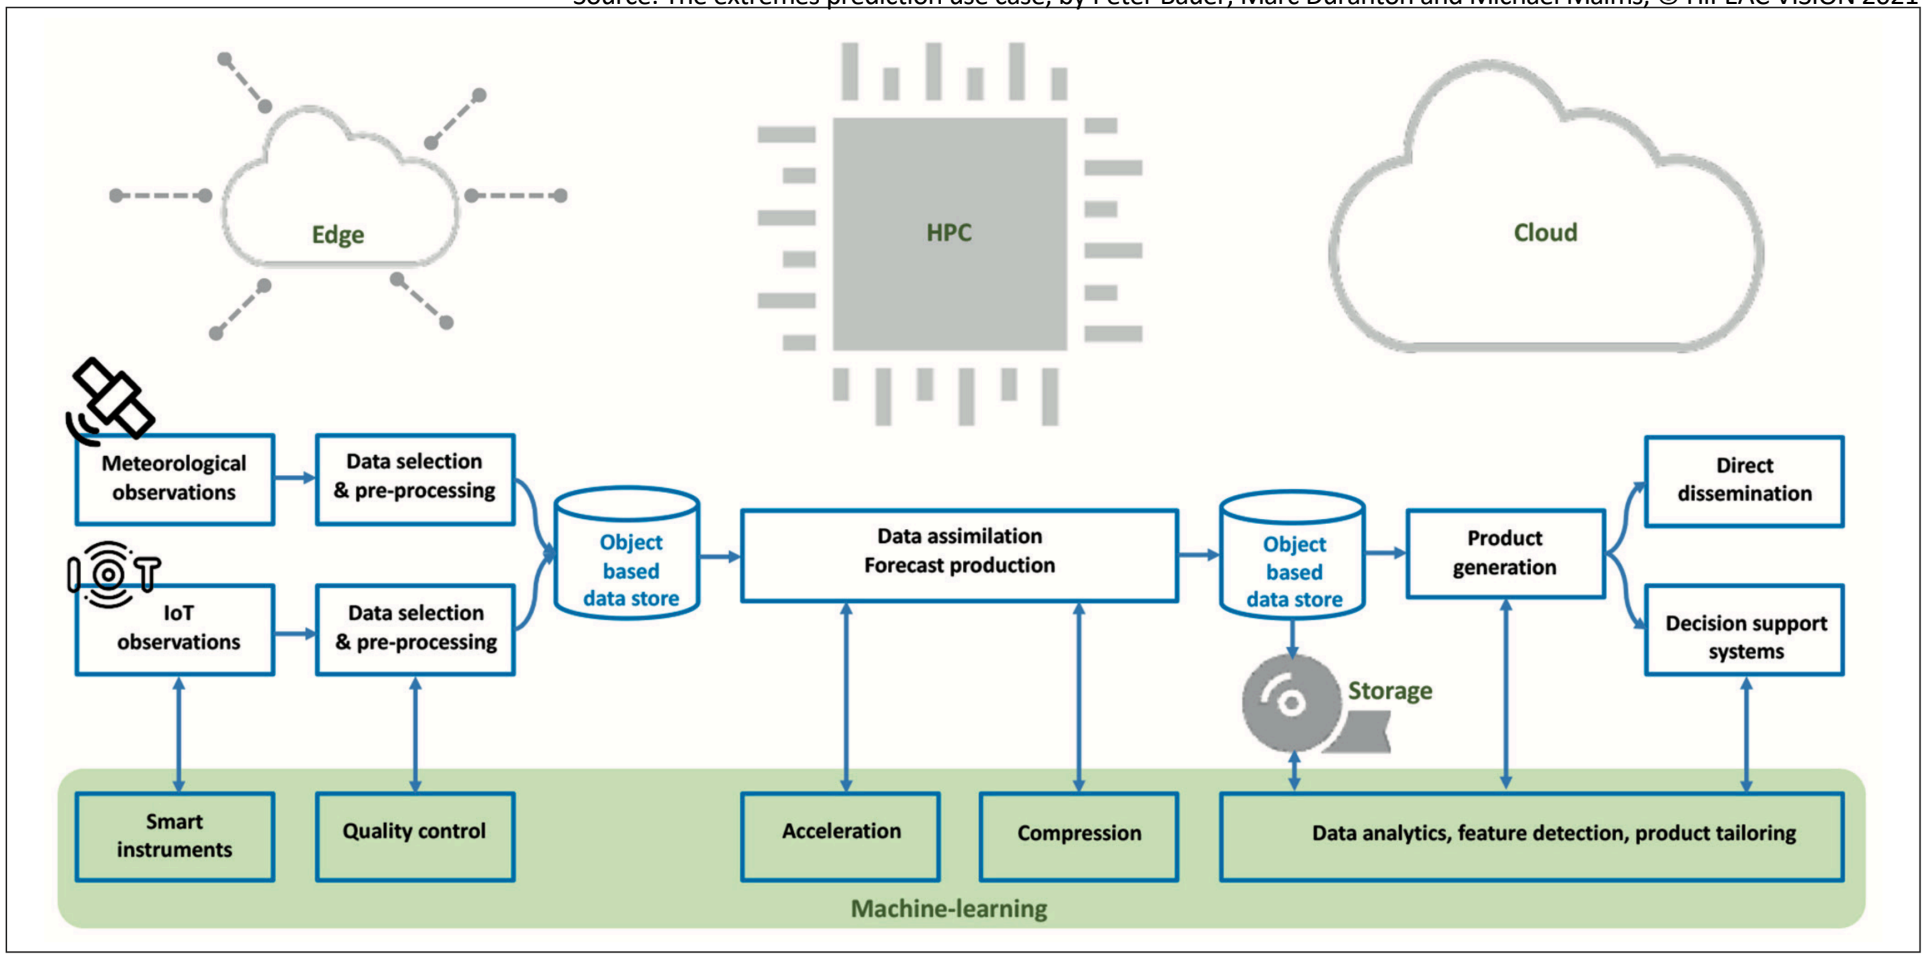

Main components of extremes prediction workflow, their alignment with edge, high-performance and cloud computing elements of the TransContinuum, and key contributions of machine learning to workflow enhancements.

HPC & AI AT EXASCALE: IT'S ALL ABOUT WORKFLOWS ➔ ONE DOES NOT FIT ALL

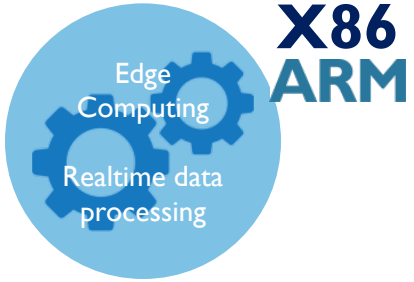

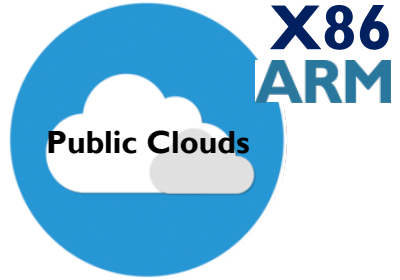

SUMMIT

sunway taihulight

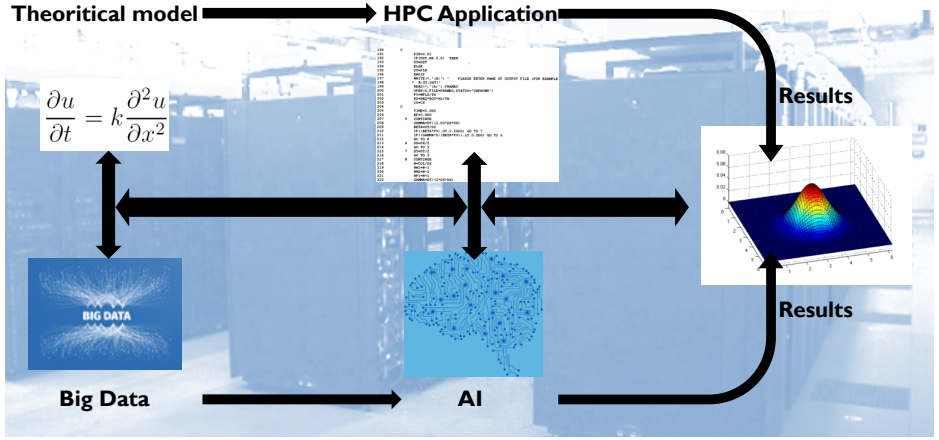

FUGAKU

HPC & AI AT EXASCALE: IT'S ALL ABOUT WORKFLOWS ➔ MODULAR ARCHITECTURE

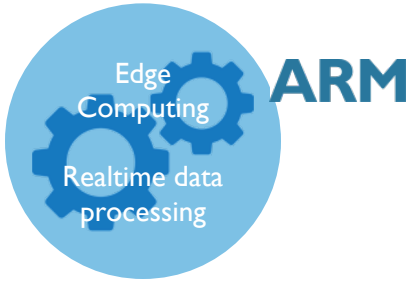

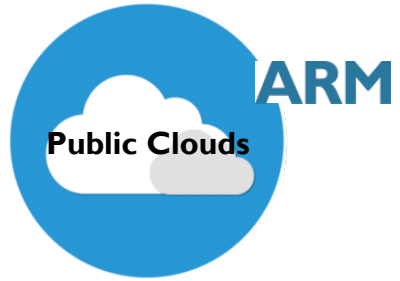

European approach

- ✓ EU Architecture
- ✓ EU HW
- ✓ EU SW
- ✓ Handle all complexity the others don't know how to

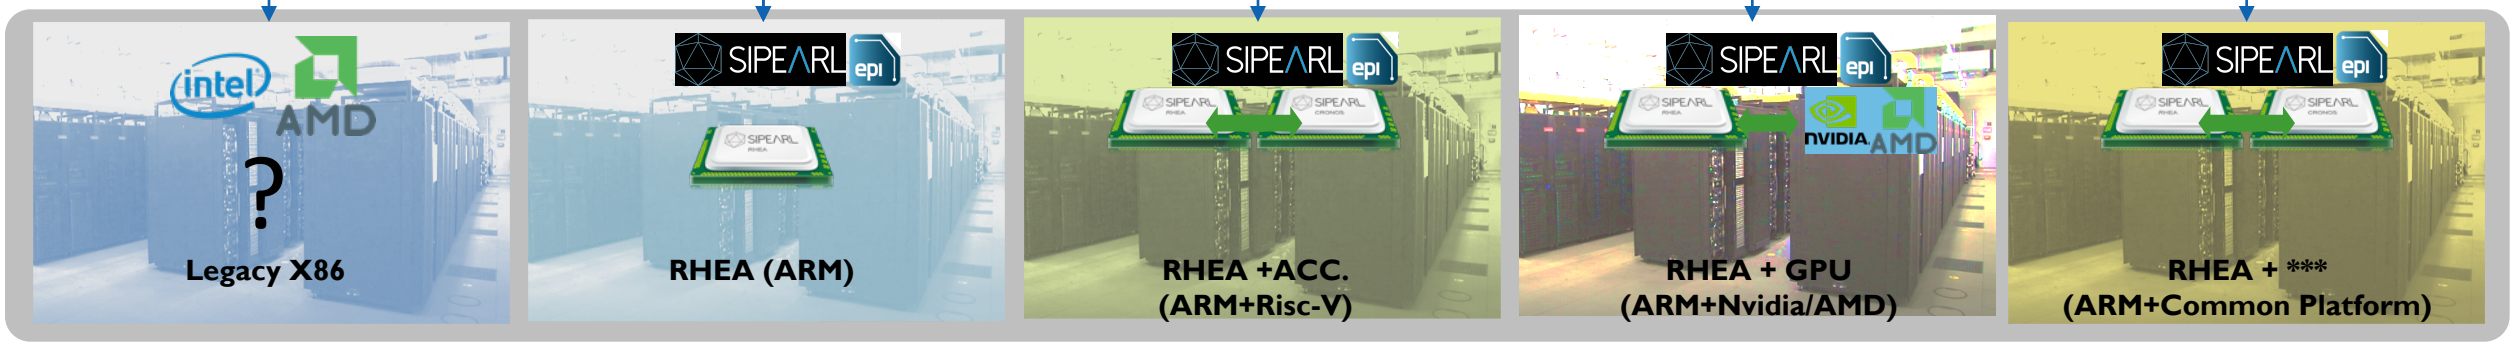

HPC & AI AT EXASCALE: IT'S ALL ABOUT WORKFLOWS ➔ COMPILE ONE, RUN ON MANY

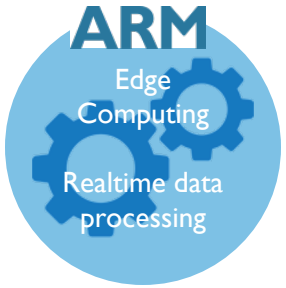

European approach

- ✓ EU Architecture
- ✓ EU HW
- ✓ EU SW
- ✓ Handle all complexity in a unified environment

Unified environment: Compilers / programming languages / libraries

CPU ARM Legacy X86 — ACC. A

CPU ARM RHEA (ARM) — ACC. B

CPU ARM RHEA +ACC — ACC. C

CPU ARM RHEA + GPU — ACC. D

CPU ARM RHEA + *** (ARM) — ACC. E

CONSEQUENCES

Hardware

Future supercomputers will be modular.

They'll have massively non-homogenous architectures, combining one general purpose processor with several different accelerator kinds

- → Processing units must properly handle modularity
  - → Open
  - → Agile
  - → Data driven

Software

Software will play an even more important role as a unification layer between all technologies, between all modules, in and out of the walls

- → Hybrid unified SW layers from IoT, Edge to Supercomputers and Cloud
- → Opensource and standardization are more important than ever
- → Proprietary SW stacks, especially for specialized HW will become problematic

SECTION 3
(POST) EXASCALE PROCESSOR OVERALL SPECIFICATIONS

CONSEQUENCES FROM SECTION 2

General Purpose Processors have to be (much) more open

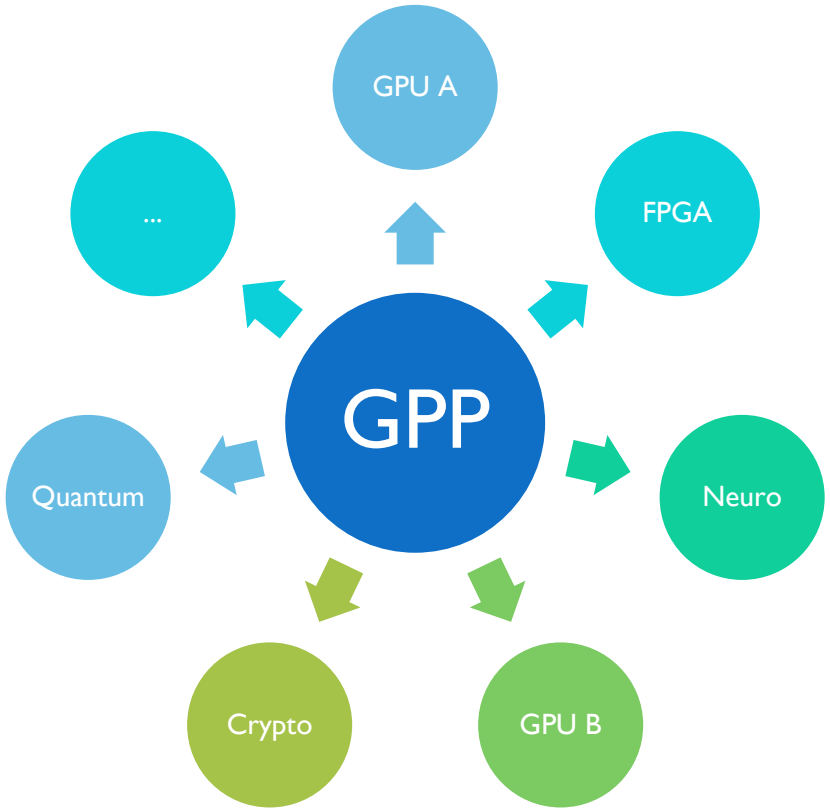

The race to FLOPS is now in the accelerators area

SO… WHAT TO EXPECT FROM AN EXASCALE GENERAL PURPOSE PROCESSOR? (PART 1)

- World class manufacturing process (7nm or better)
- Need extreme flexibility and performances on external links
  - Composable architectures, from 1 to >=4 sockets (CCIX and/or CXL)
  - HBM 2e / 3 and DDR 5/6
  - and PCIe G5/G6
  - and CXL 1 / 2

- Transparent integration in end-to-end dataflow : IoT ←→ Edge ←→ Datacenter ←→ Cloud
  - Easy to port / optimize
  - Opensource tools
  - Unified development tools
  - Compile one, run on many

THE R_PEAK CHALLENGE – CPU VS. GPU

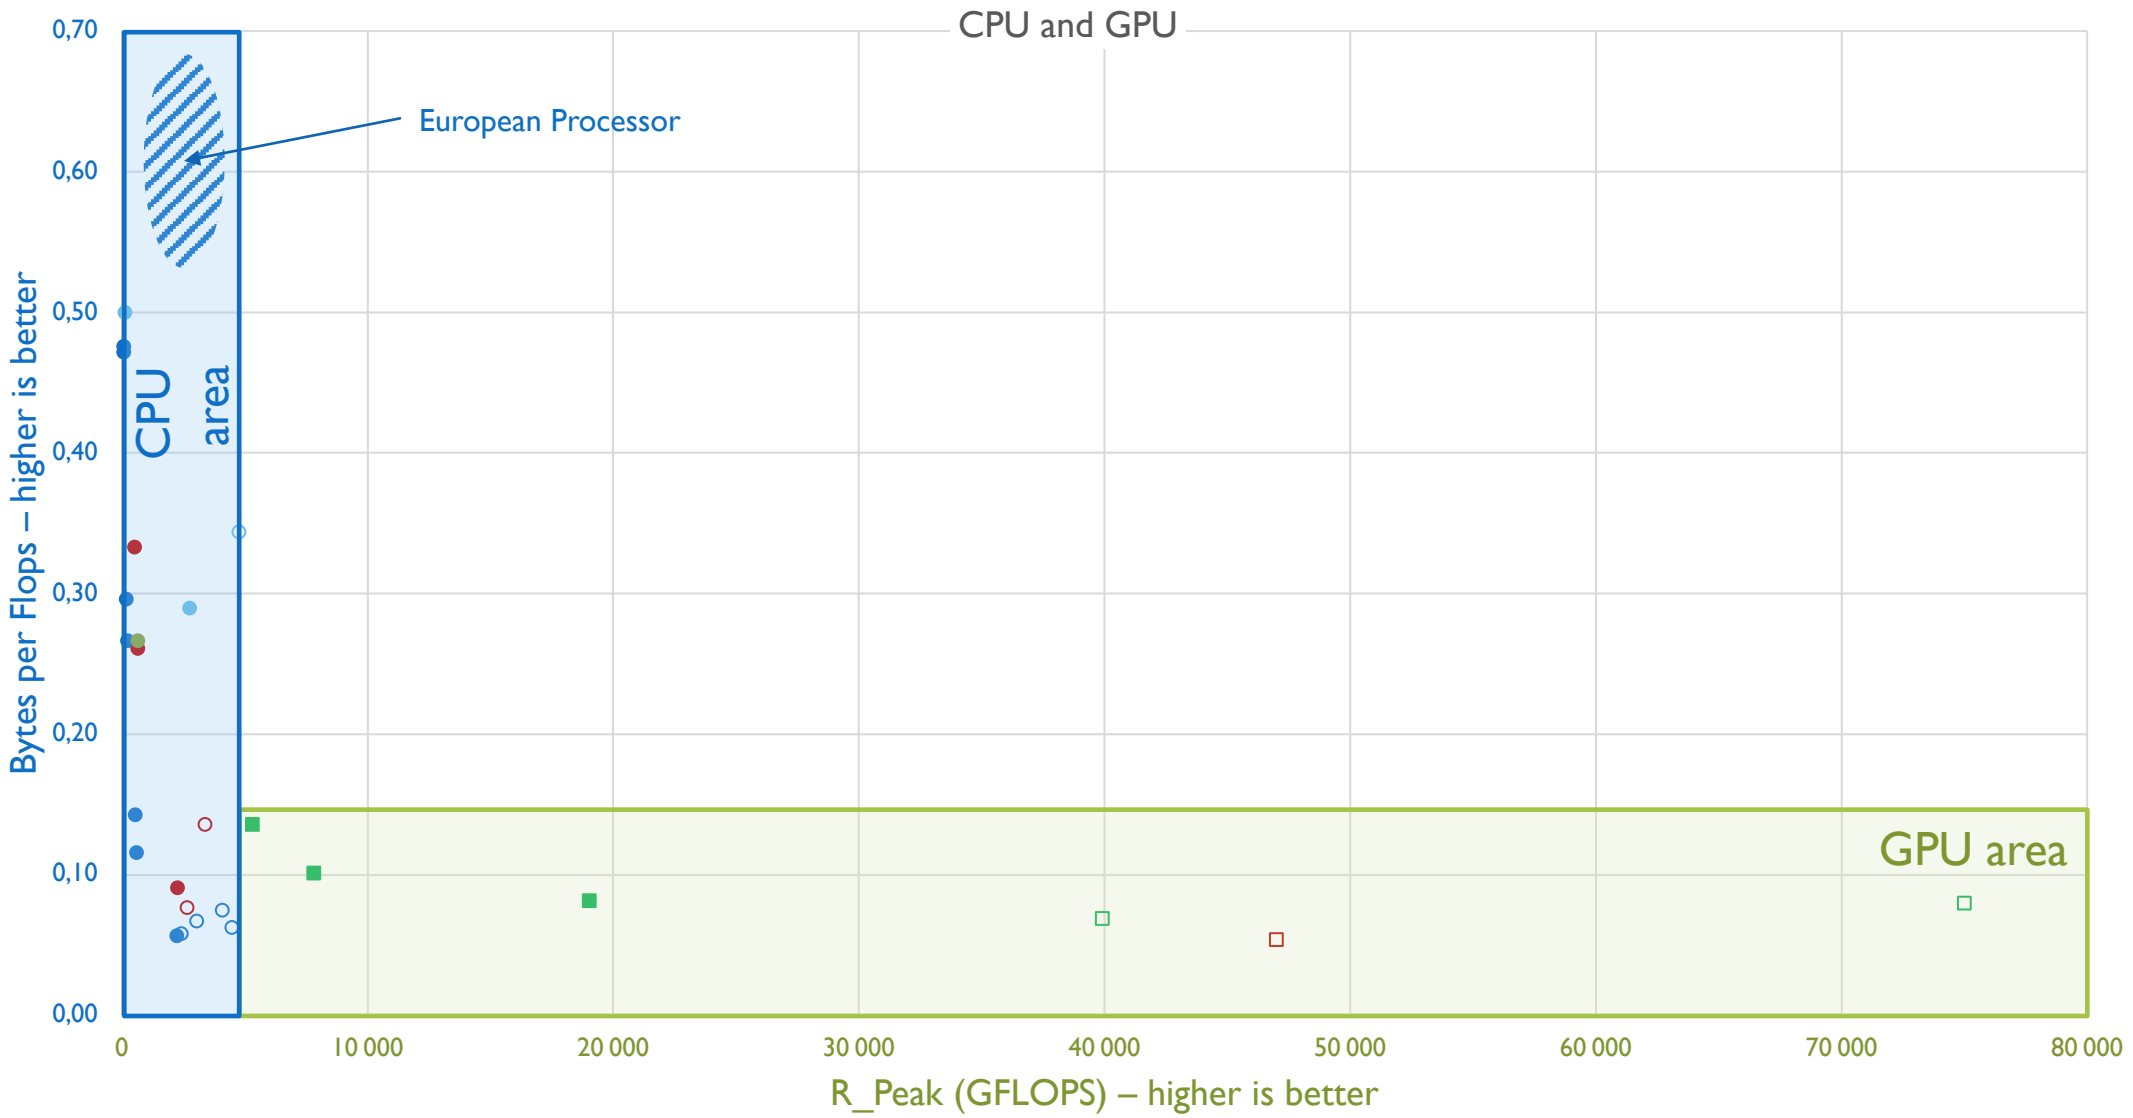

THE R_PEAK CHALLENGE – CPU VS. GPU

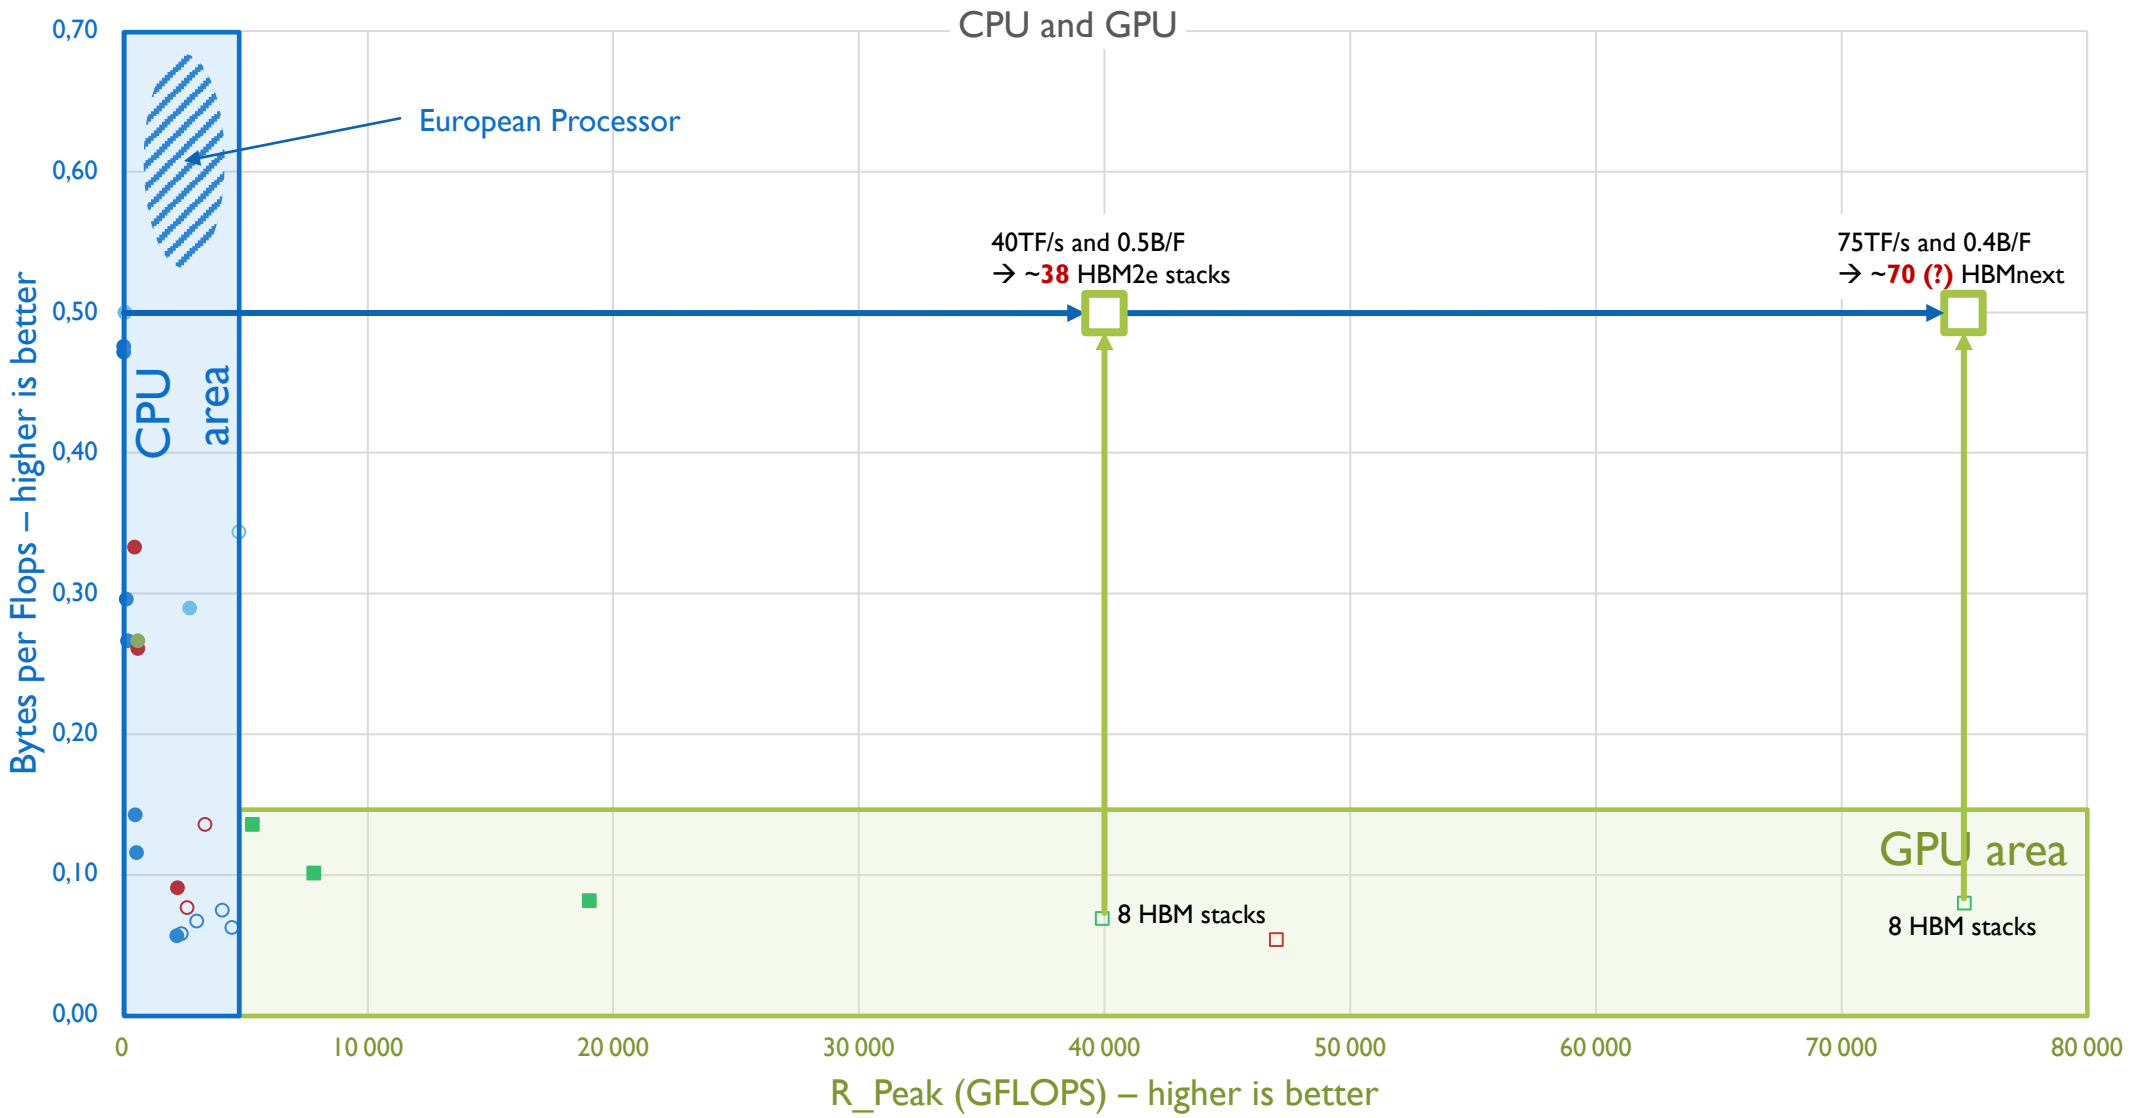

THE R_PEAK CHALLENGE – CPU VS. GPU

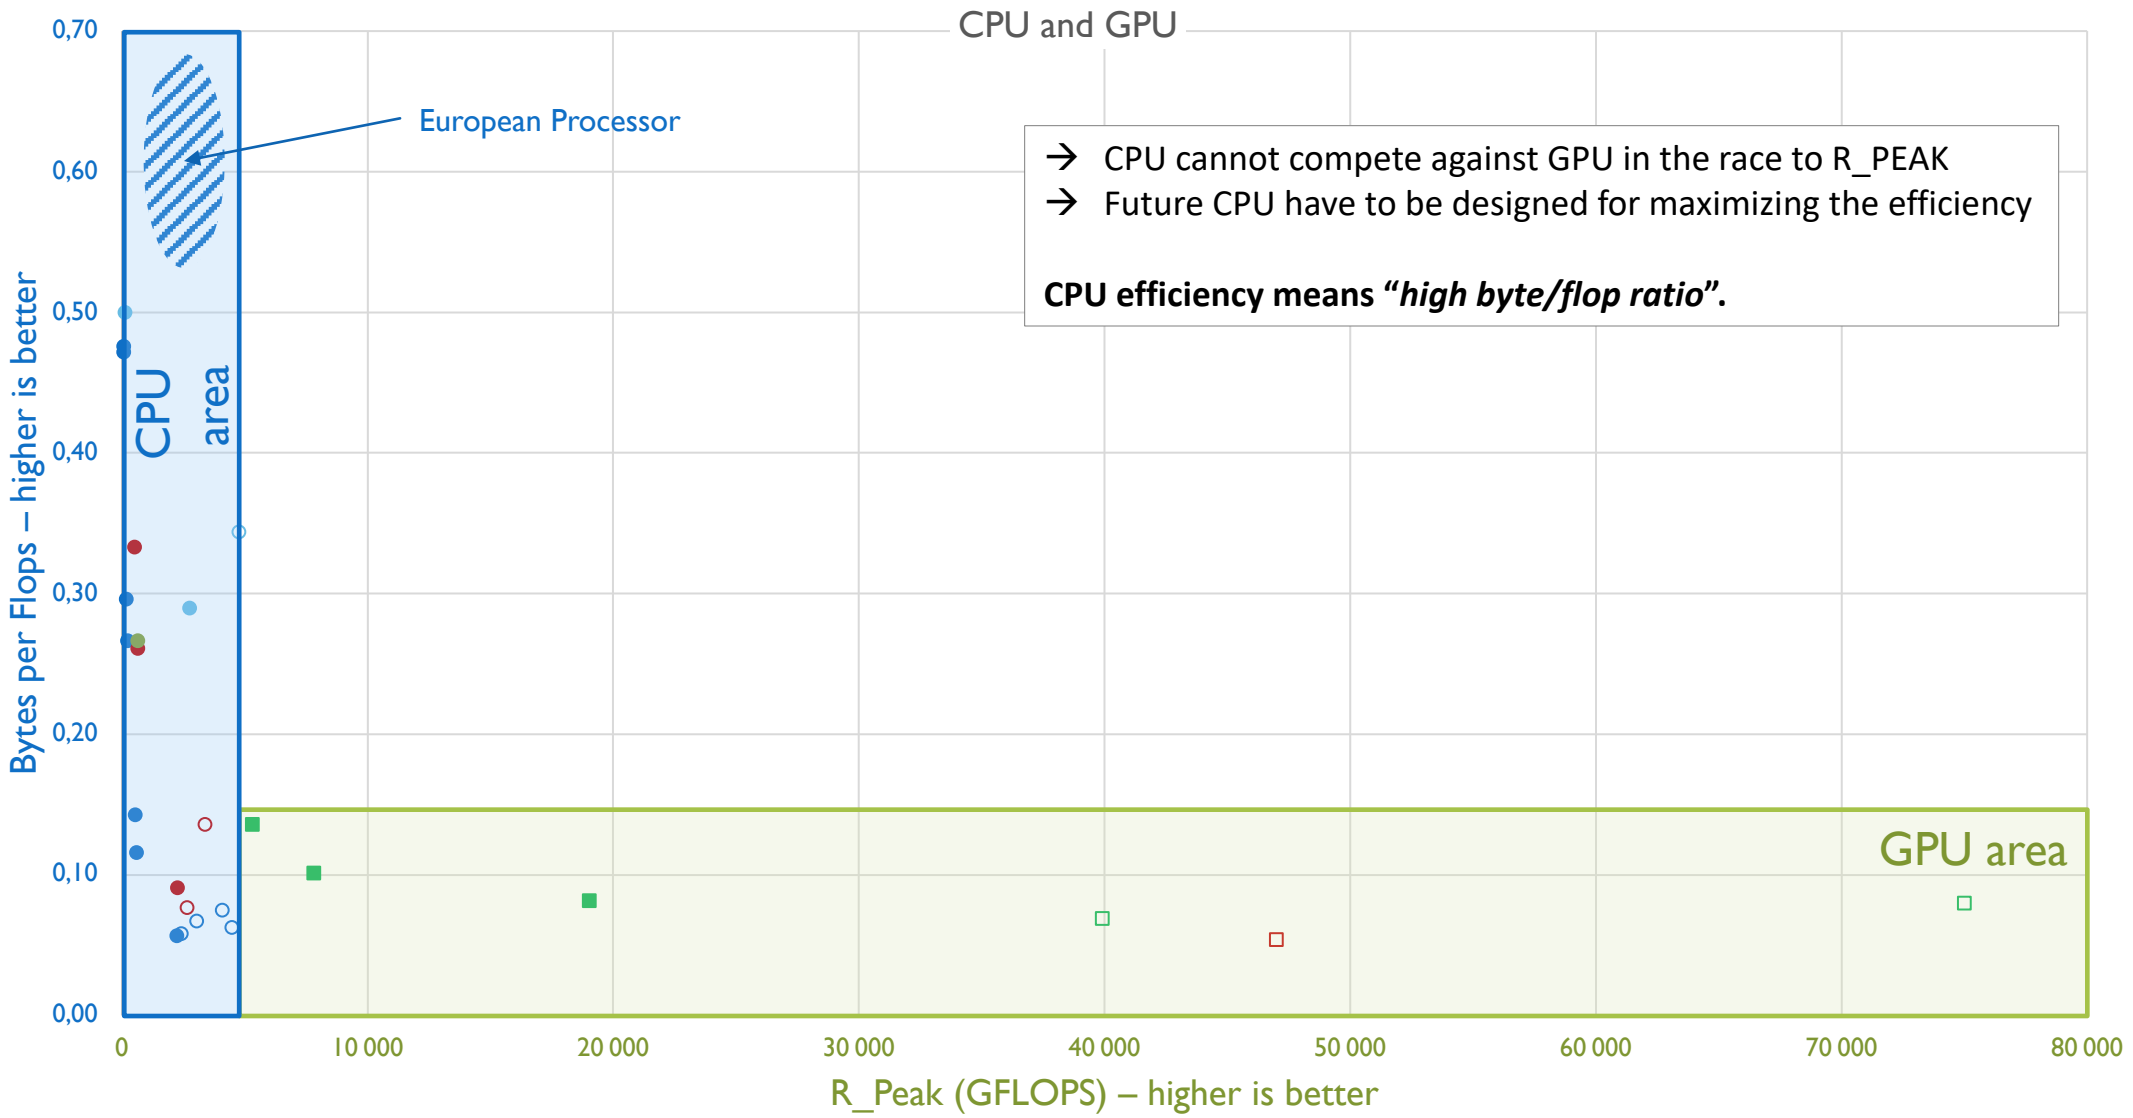

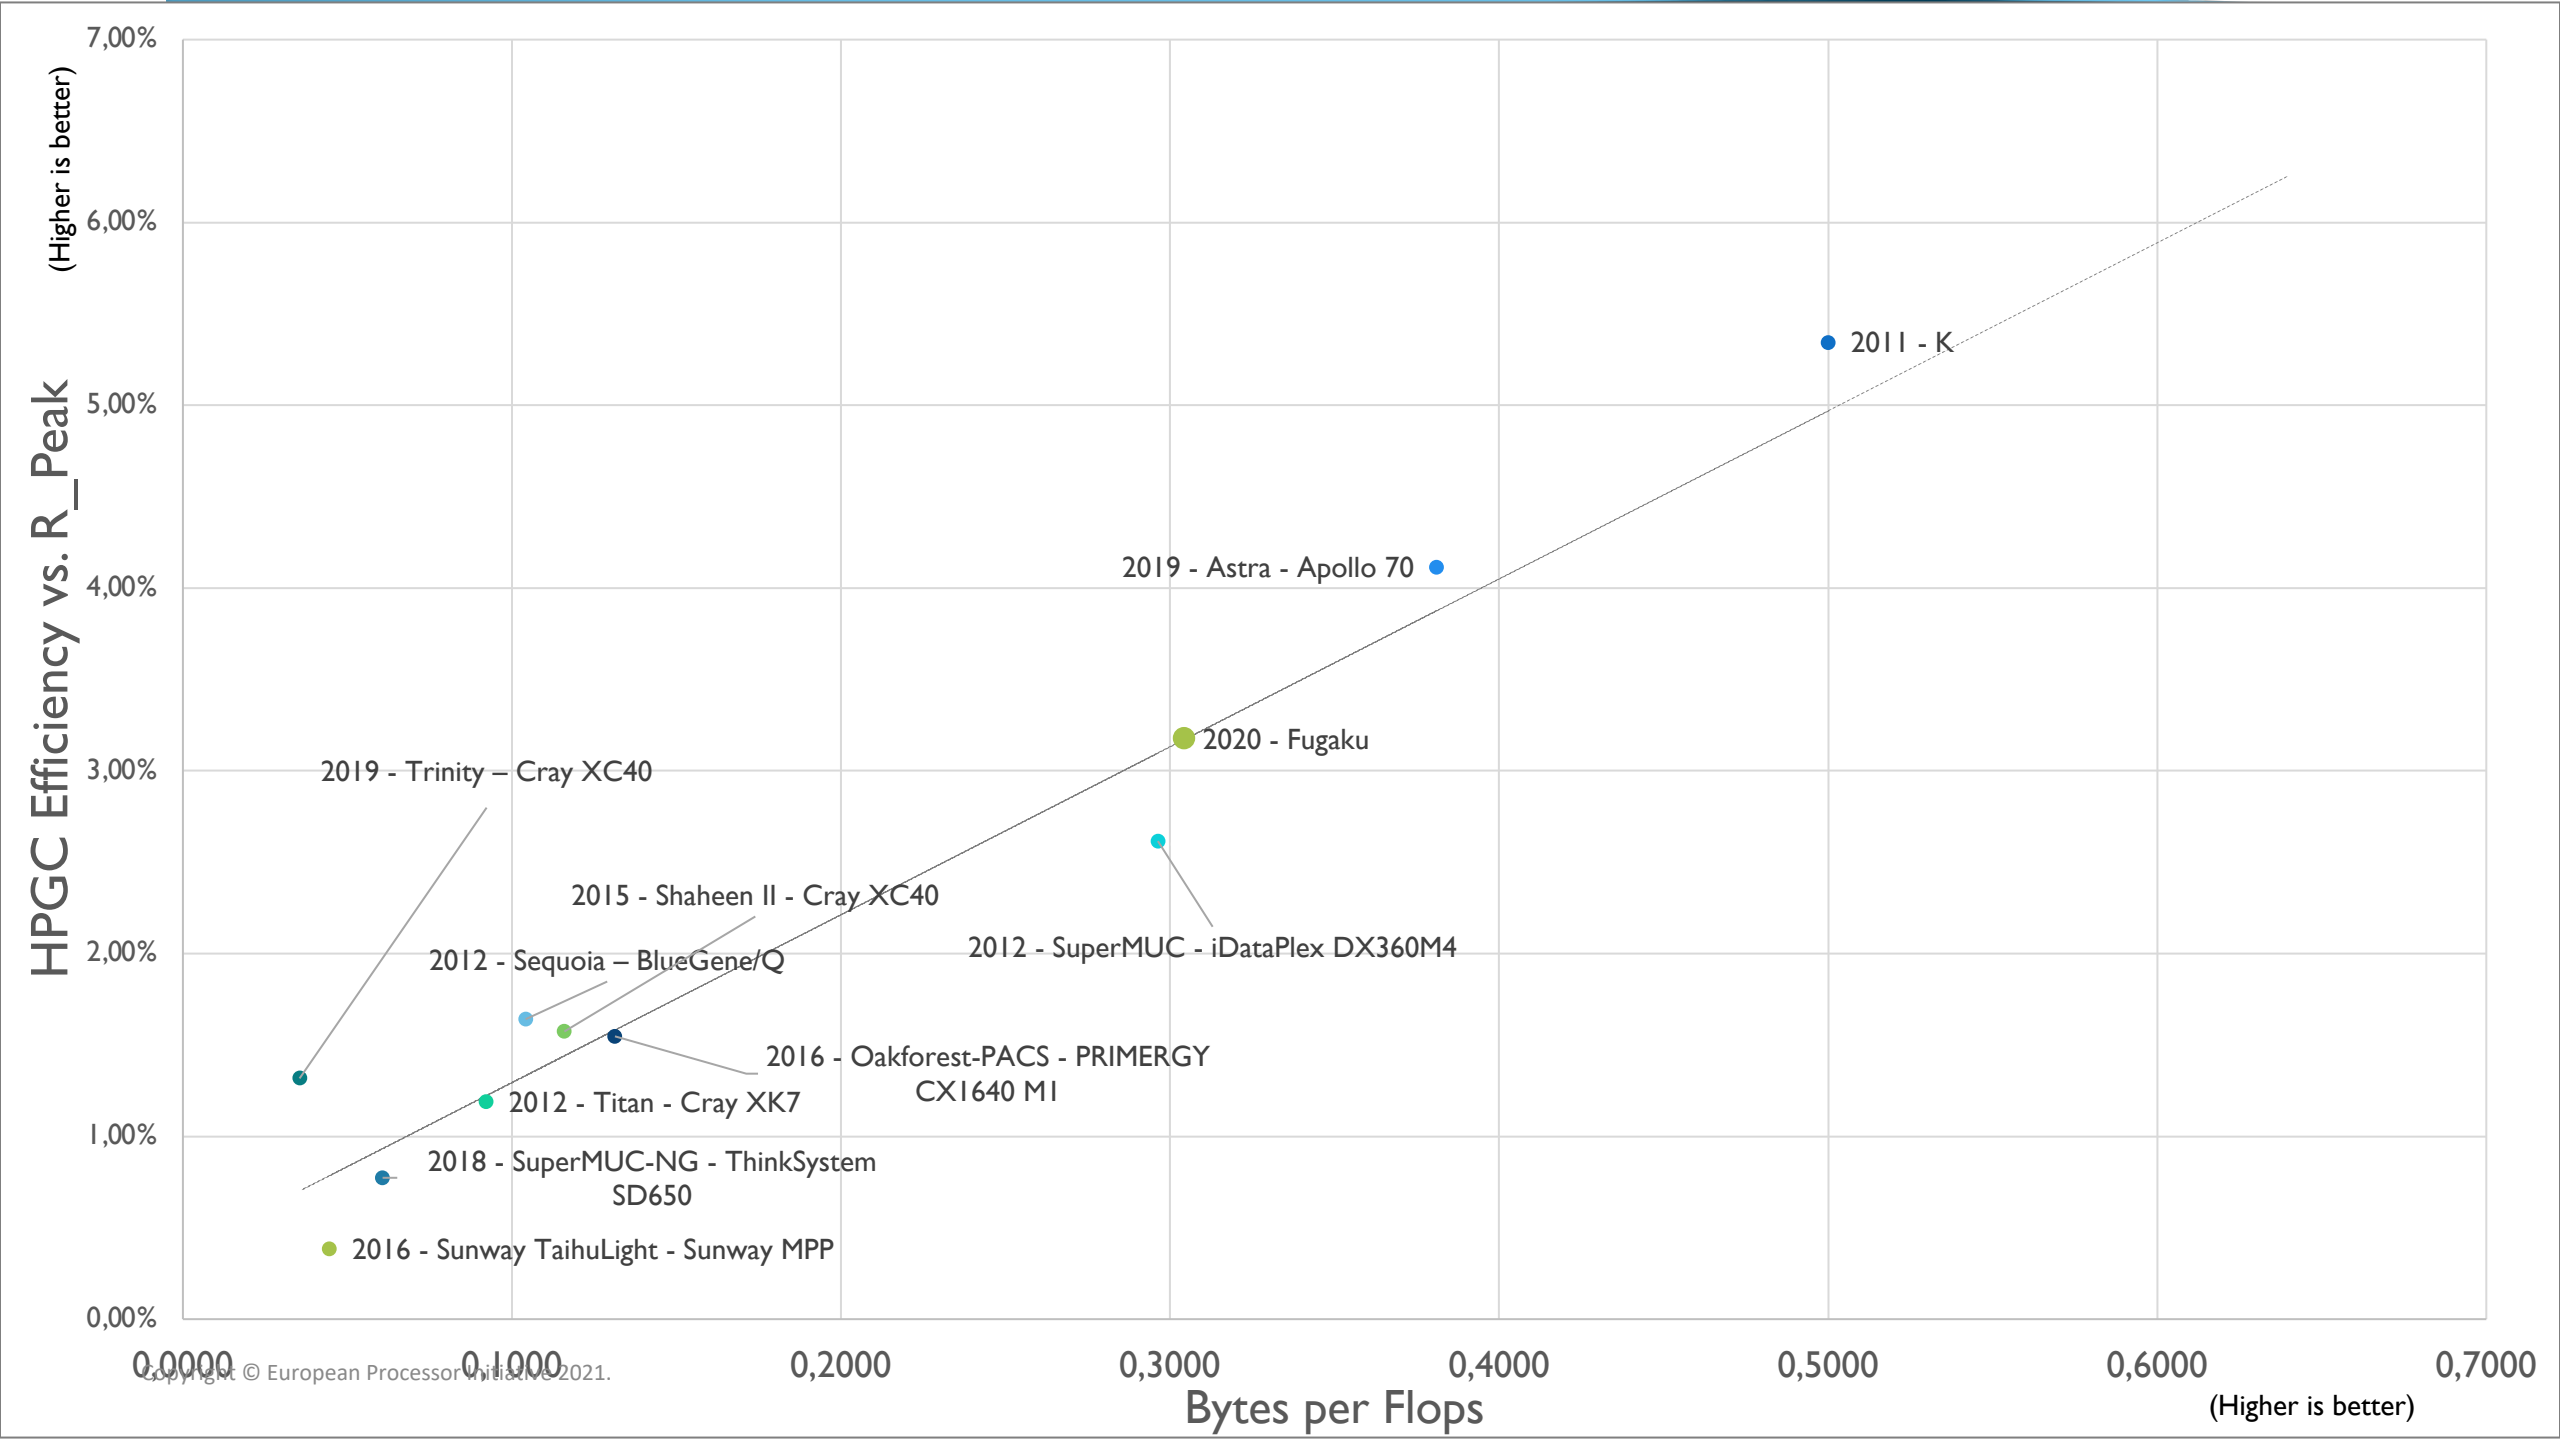

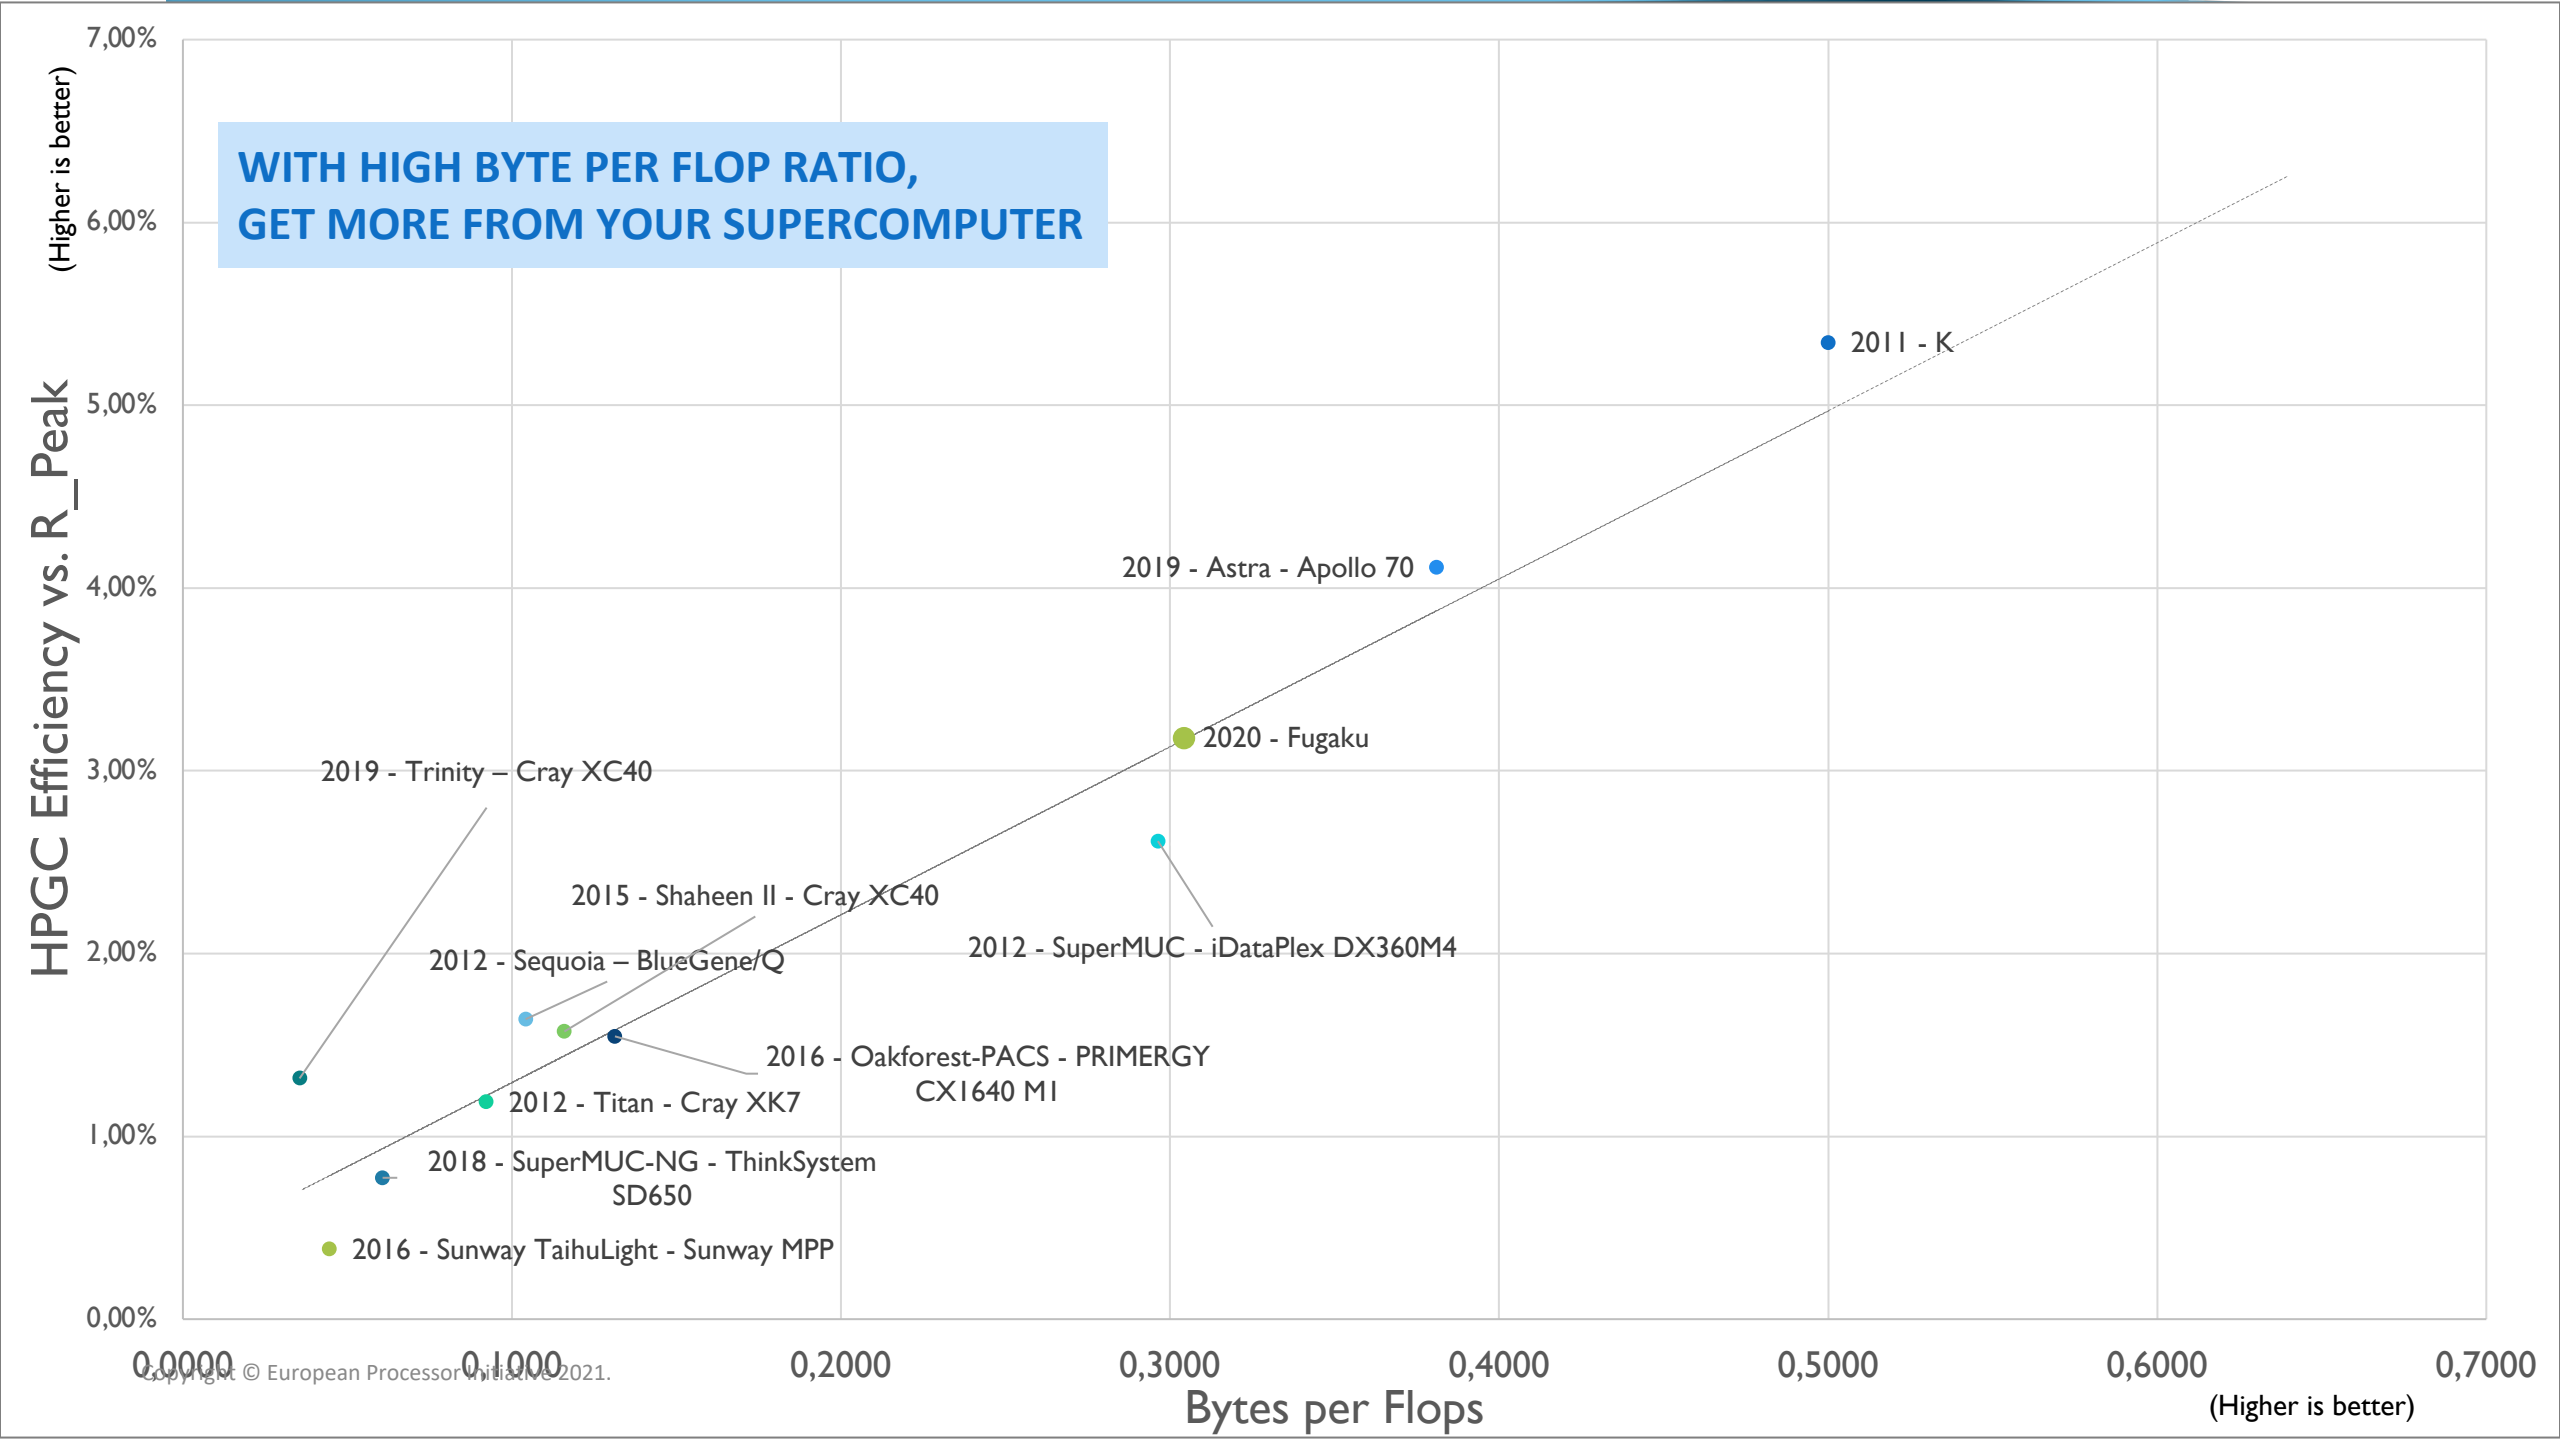

Why mem bandwidth and Bytes/Flops matter?

JUST GET MORE FROM YOUR EXTRA EXPENSIVE SUPERCOMPUTER

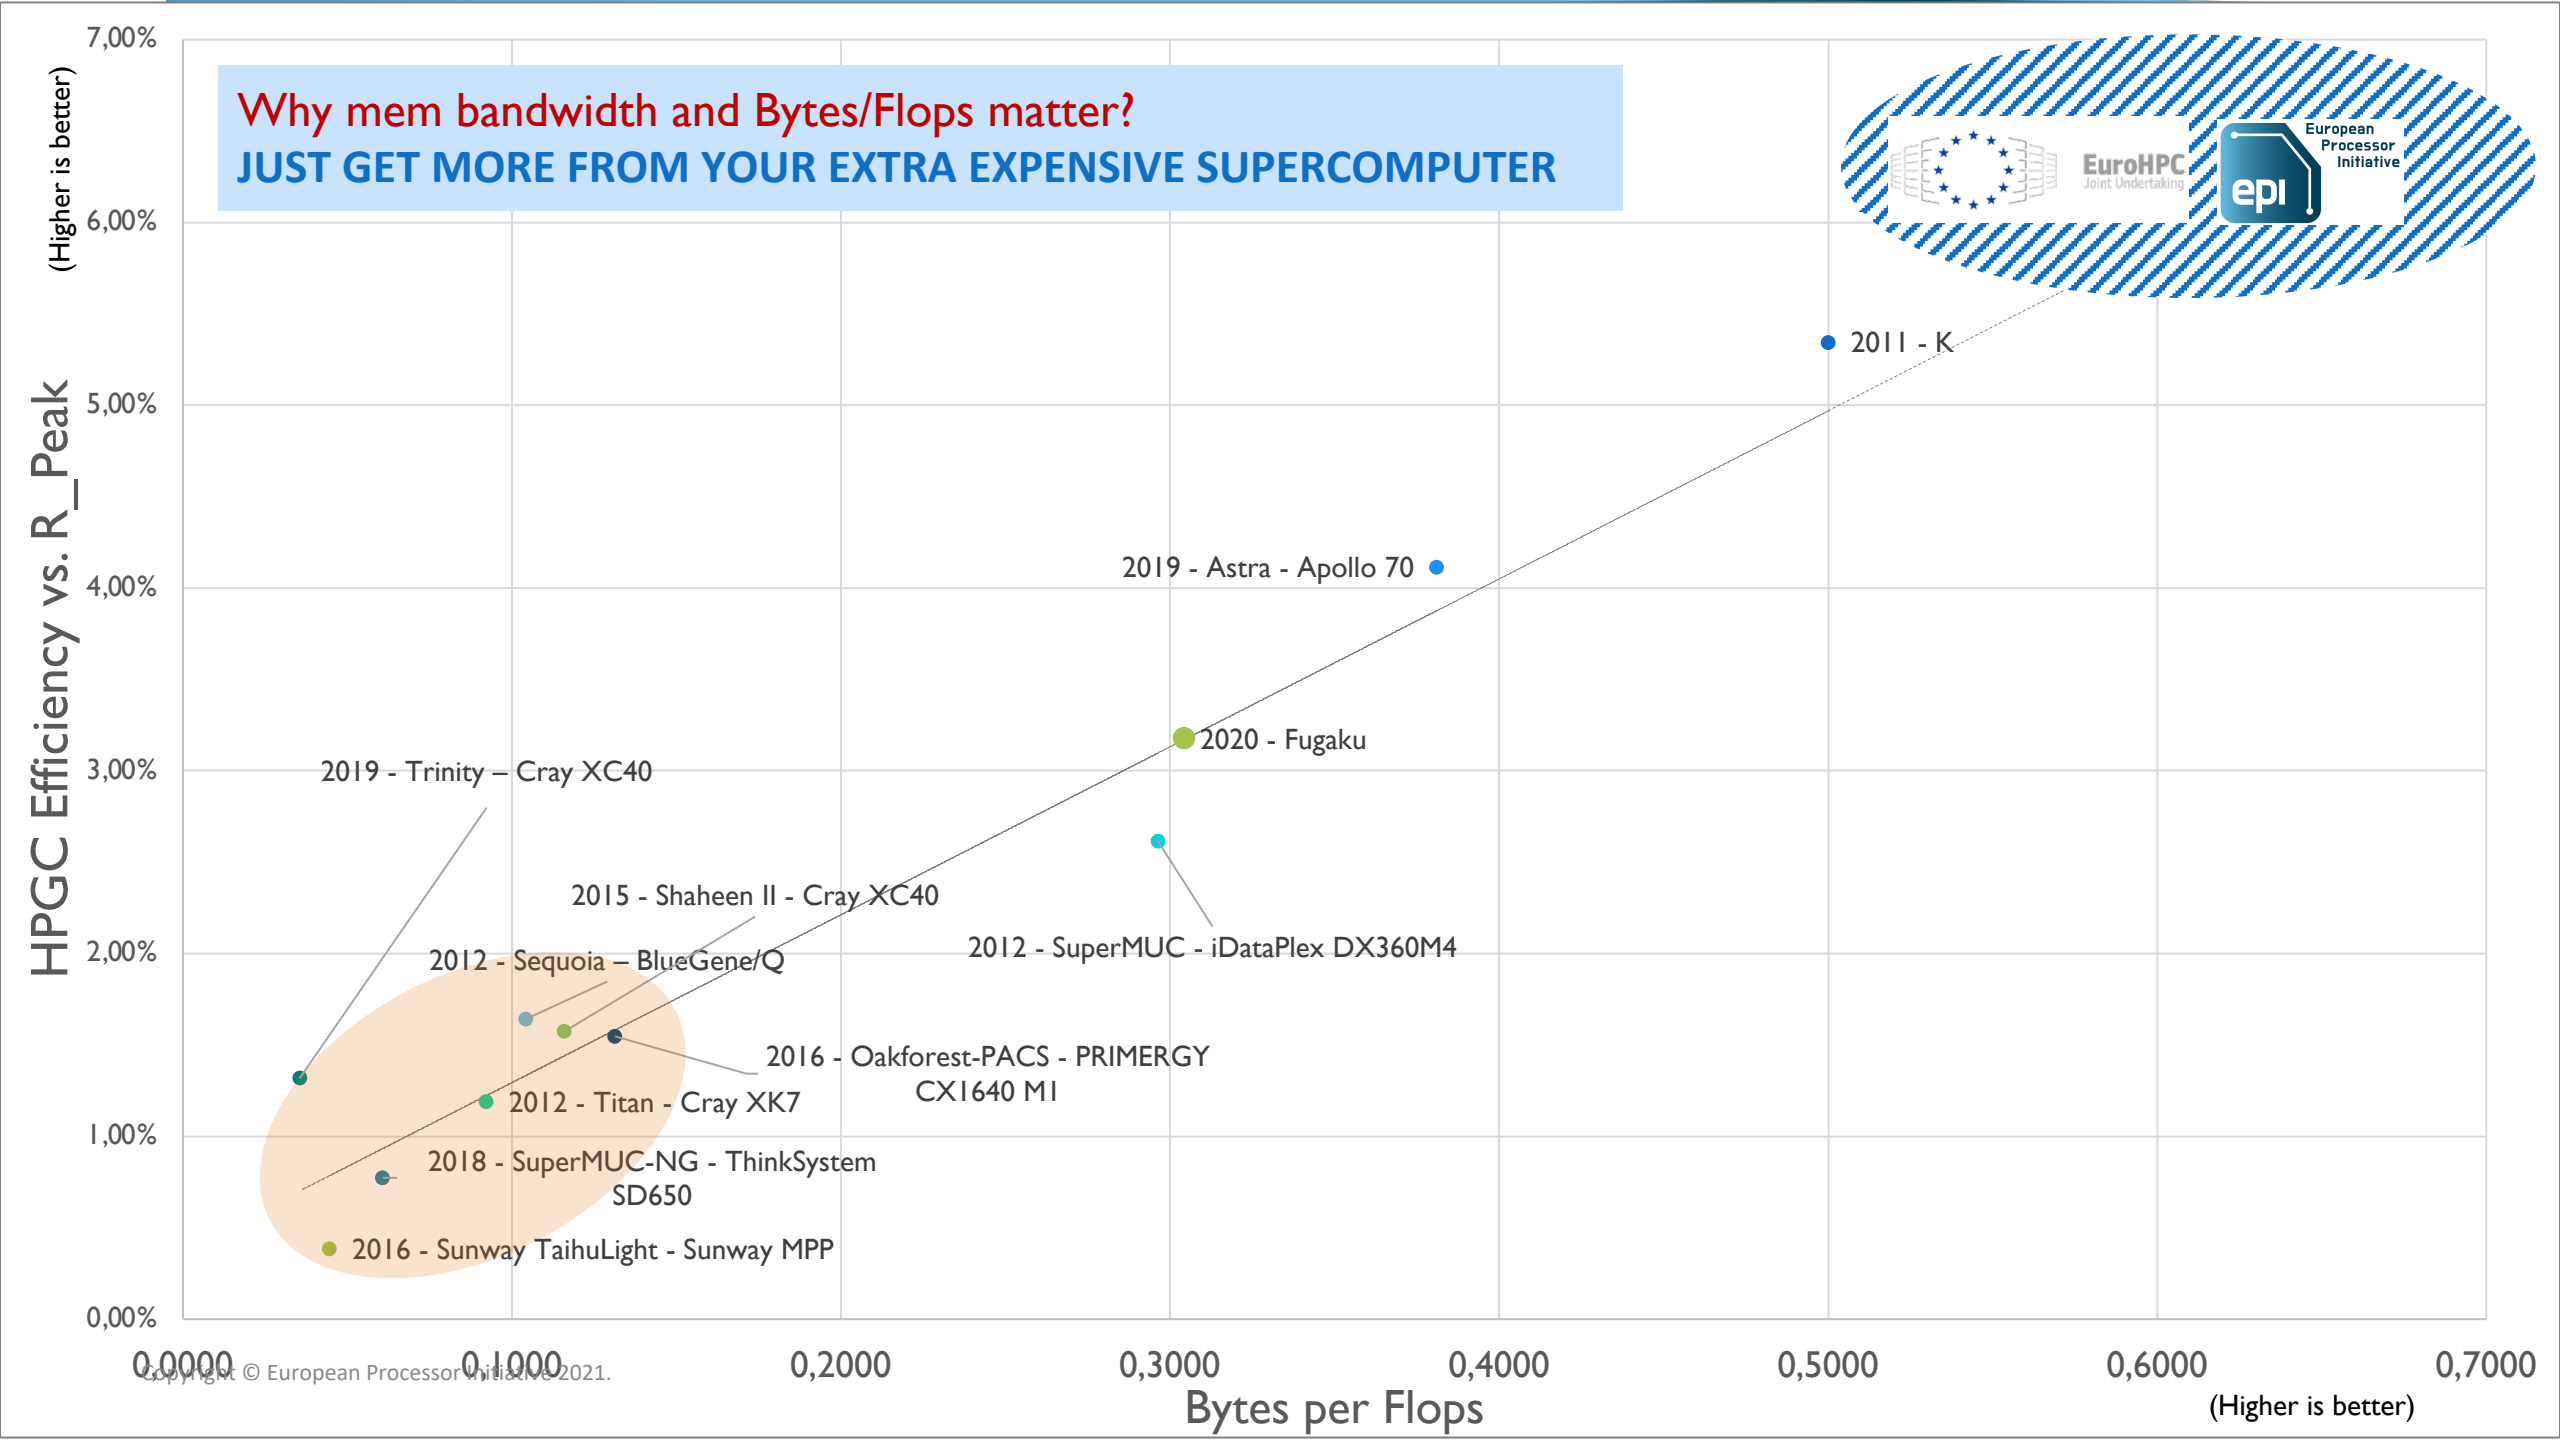

SO… WHAT TO EXPECT FROM AN EXASCALE GENERAL PURPOSE PROCESSOR? (PART 2)

- World class manufacturing process (7nm or better)
- Need extreme flexibility and performances on external links
  - Composable architectures, from 1 to >=4 sockets (CCIX and/or CXL)
  - HBM 2e / 3 and DDR 5/6
  - and PCIe G5/G6
  - and CXL 1 / 2

- Transparent integration in end-to-end dataflow : IoT ←→ Edge ←→ Datacenter ←→ Cloud
  - Easy to port / optimize
  - Opensource tools
  - Unified development tools
  - Compile one, run on many

- No need to compete with specialized devices like GPUs
- Need “good enough” but excellent FP64 performances.
- Need much better byte/Flop ratio than today. Target 0.5+ Byte/Flop ➔ improved efficiency on real workflows

HETEROGENEOUS INTEGRATION ➔ COMMON PLATFORM

- Allows integration of customized functions in chip, in package, on board, or over PCIe or network link
- EPI Accelerators work in I/O coherent mode and share the same memory viewSingle or dual chiplet package for power efficient sizing
- Coherent NoC with system level cache to keep data local
- D2D interface open to EPI (and beyond)

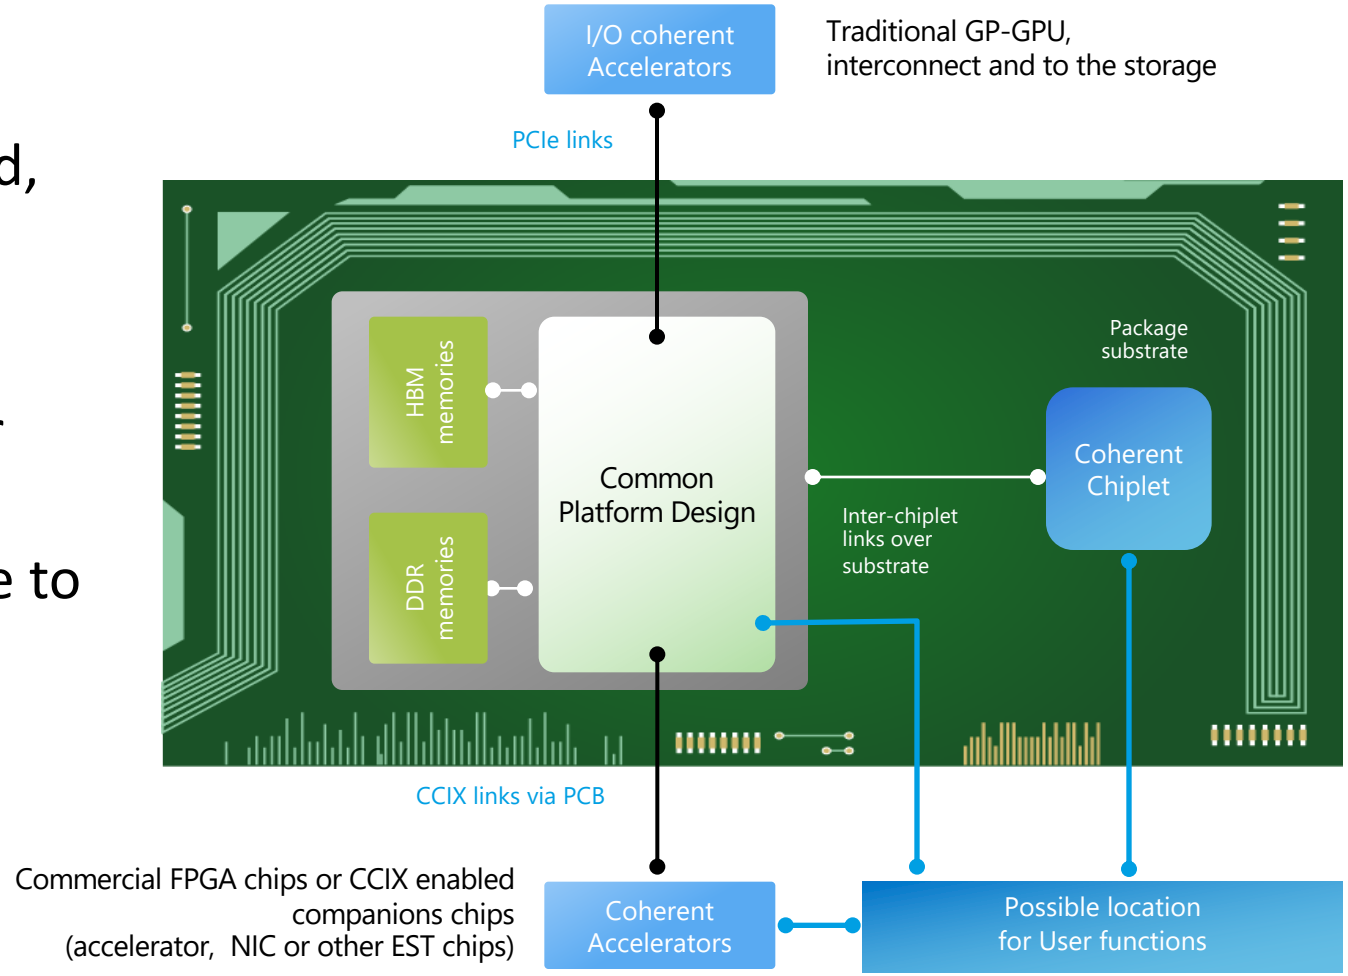

COMMON PLATFORM TO HARMONIZE THE HETEROGENEOUS COMPUTING ENVIRONMENT

Computing Units

- Arm – Scalable Vector Extension
- MPPA - Multi-Purpose Processing Array
- EPAC – RISC-V based Accelerators
- eFPGA - embedded FPGA

METHODOLOGY

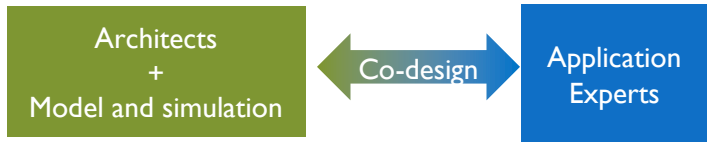

SOFTWARE

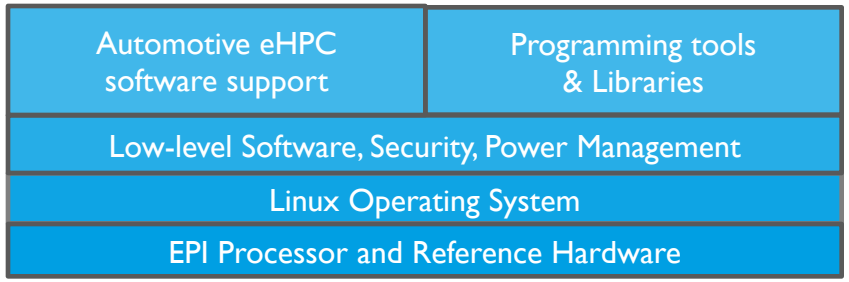

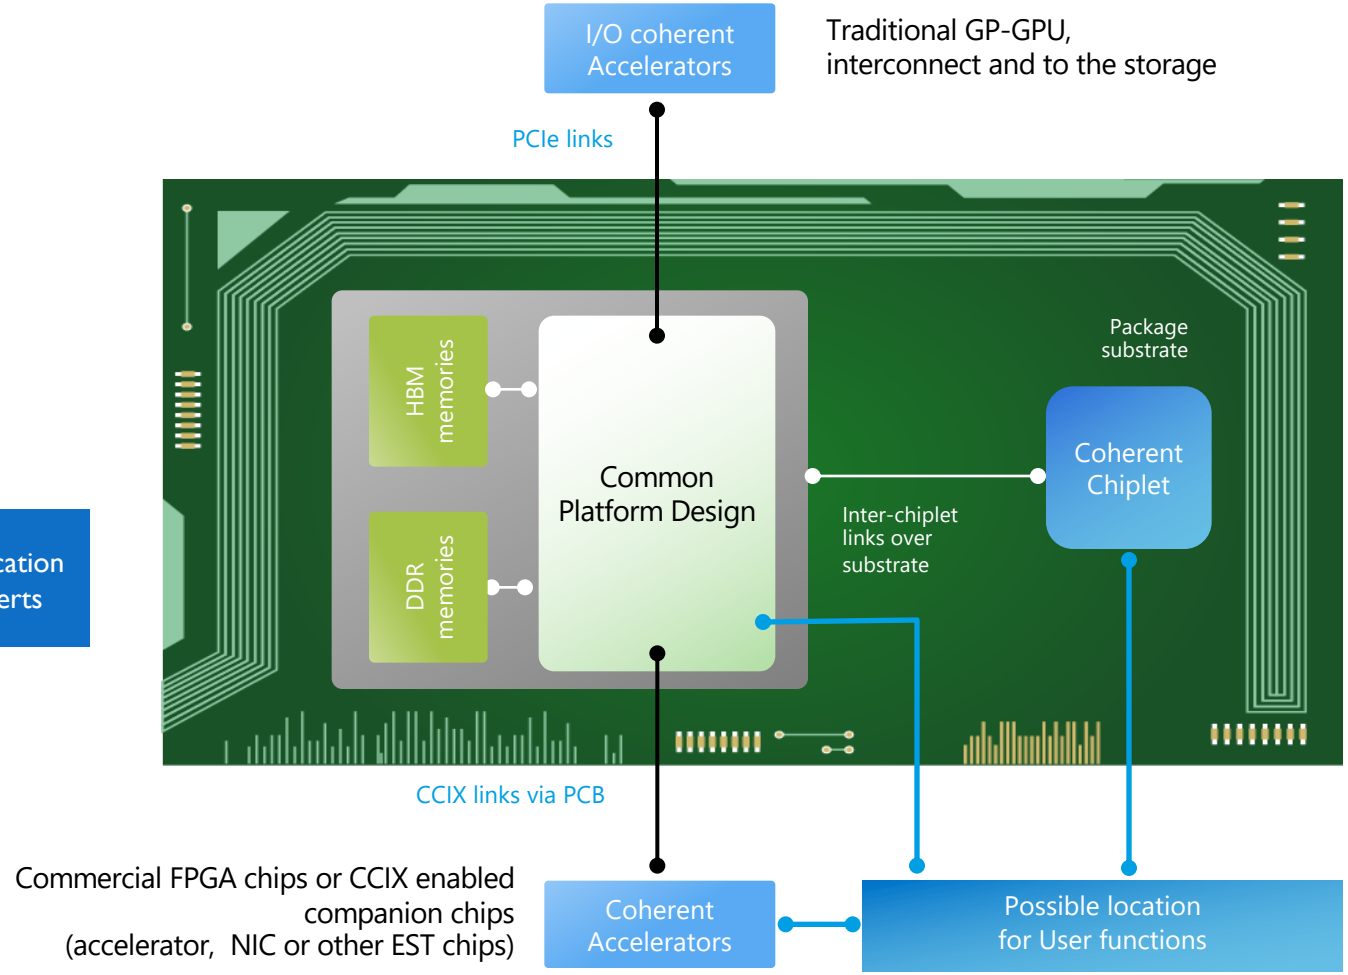

ON-CHIP HETEROGENEOUS INTEGRATION

- 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).

SOFTWARE

Automotive eHPC software support	Programming tools & Libraries
Low-level Software, Security, Power Management	
Linux Operating System	
EPI Processor and Reference Hardware	

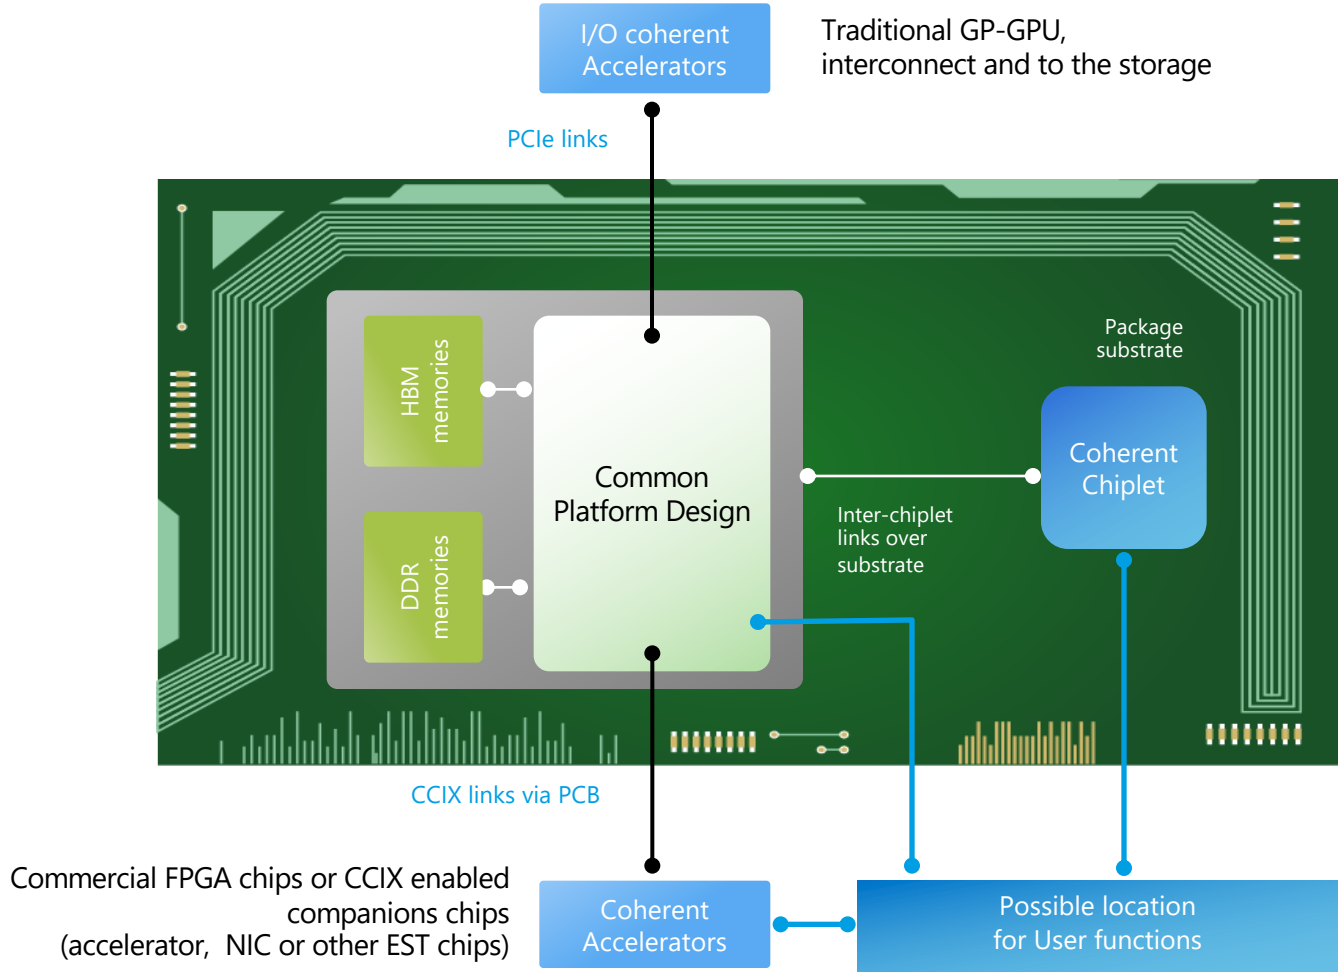

OFF-CHIP HETEROGENEOUS INTEGRATION

- Chiplet
- Socket
- Network

Distance (time)

SOFTWARE

Automotive eHPC software support	Programming tools & Libraries
Low-level Software, Security, Power Management	
Linux Operating System	
EPI Processor and Reference Hardware	

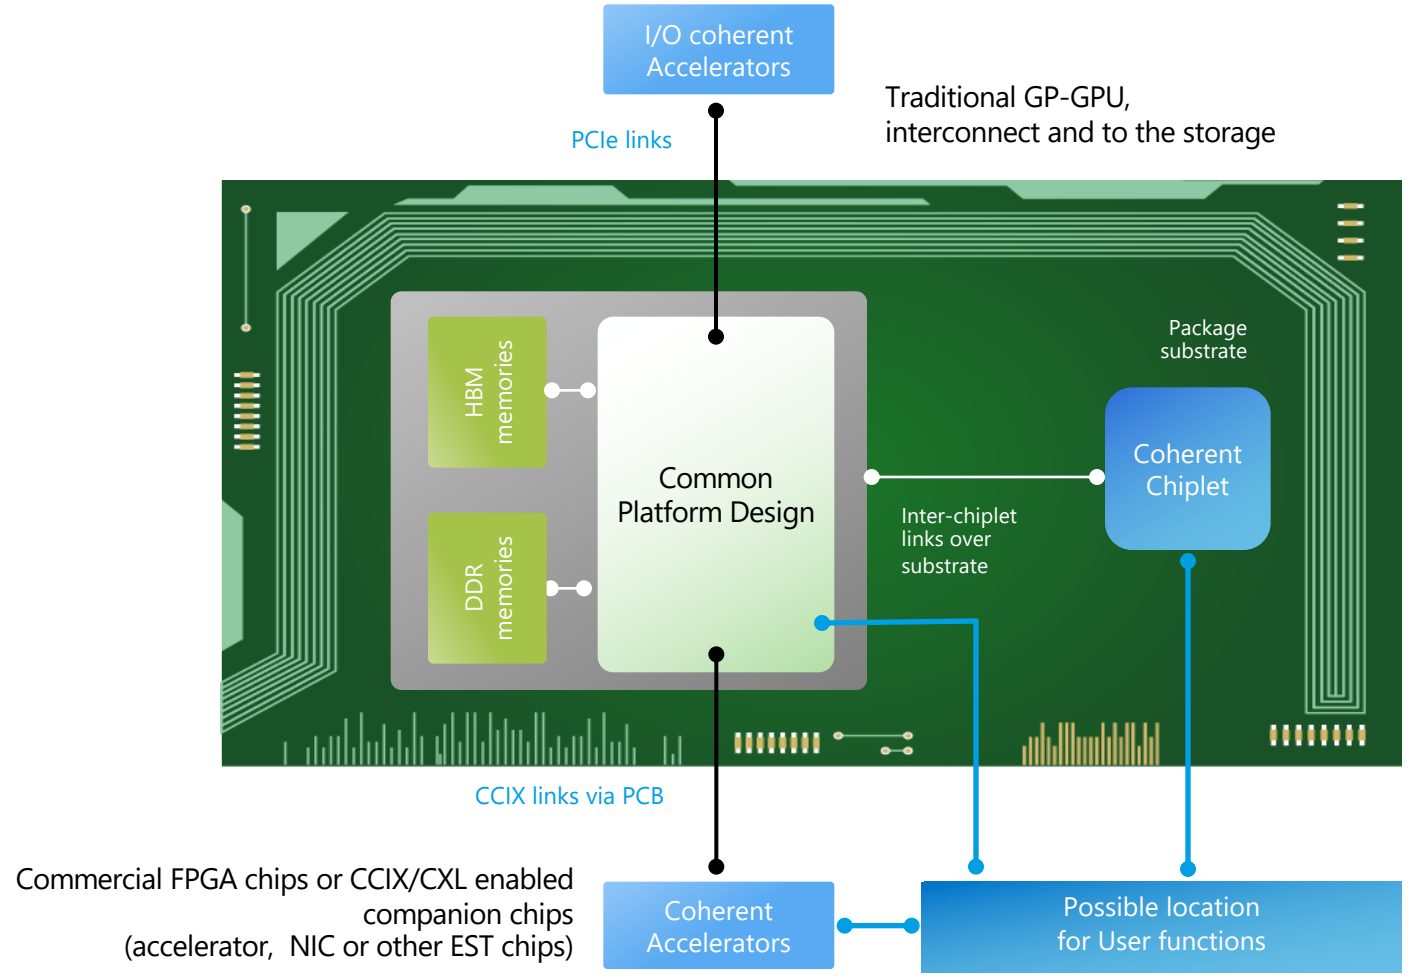

SO… WHAT TO EXPECT FROM AN EXASCALE GENERAL PURPOSE PROCESSOR? (PART 3)

- World class manufacturing process (7nm or better)
- Need extreme flexibility and performances on external links
  - Composable architectures, from 1 to >=4 sockets (CCIX and/or CXL)
  - HBM 2e / 3 and DDR 5/6
  - and PCIe G5/G6
  - and CXL 1 / 2

- Transparent integration in end-to-end dataflow : IoT ←→ Edge ←→ Datacenter ←→ Cloud
  - Easy to port / optimize
  - Opensource tools
  - Unified development tools
  - Compile one, run on many

- No need to compete with specialized devices like GPUs
- Need “good enough” but excellent FP64 performances.
- Need much better byte/Flop ratio than today. Target 0.5+ Byte/Flop ➔ improved efficiency on real workflows

- Based on multi-die and IO-DIE building blocks, combined at package level or out of package

  ➔ common platform

WRAP-UP

TAKEAWAYS

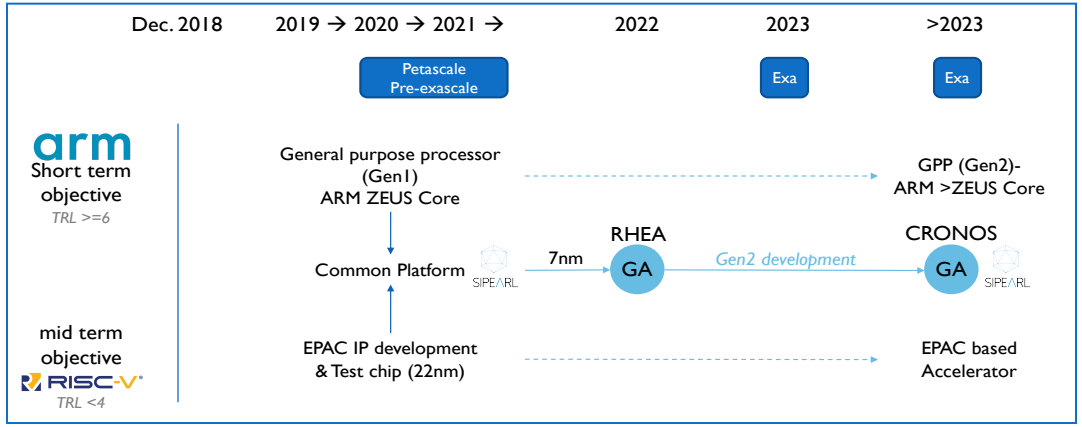

TAKEAWAYS

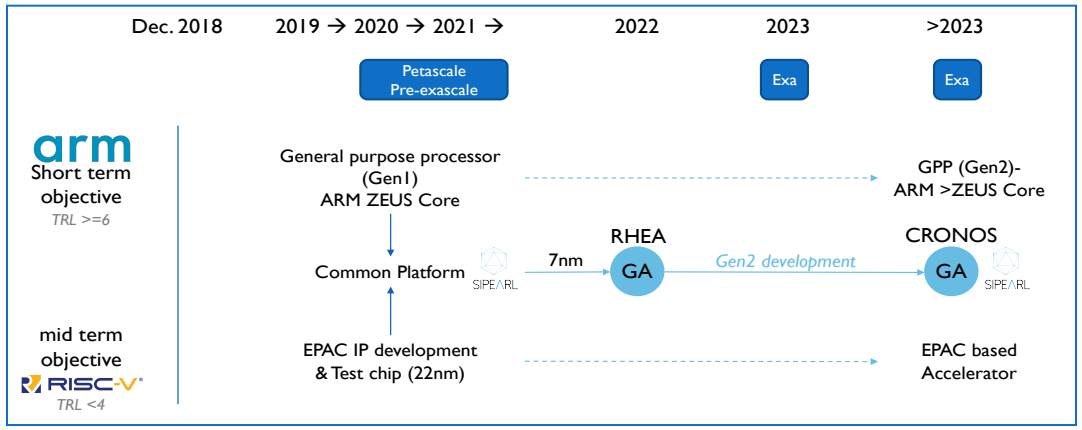

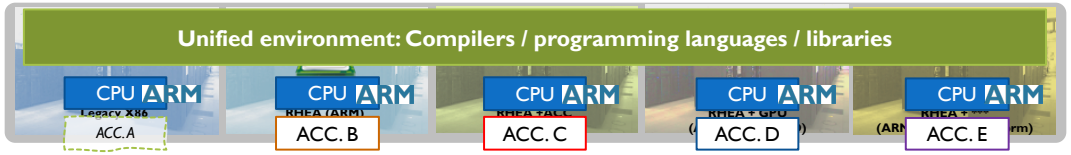

Exascale and post-Exascale supercomputers will modular.

They'll have massively non-homogenous architectures, combining one general purpose processor with several different accelerator kinds

Software will play an even more important role as a unification layer between all technologies, between all modules

→ Opensource and standardization are more important than ever

→ Proprietary SW stacks, especially for specialized HW will become problematic

TAKEAWAYS

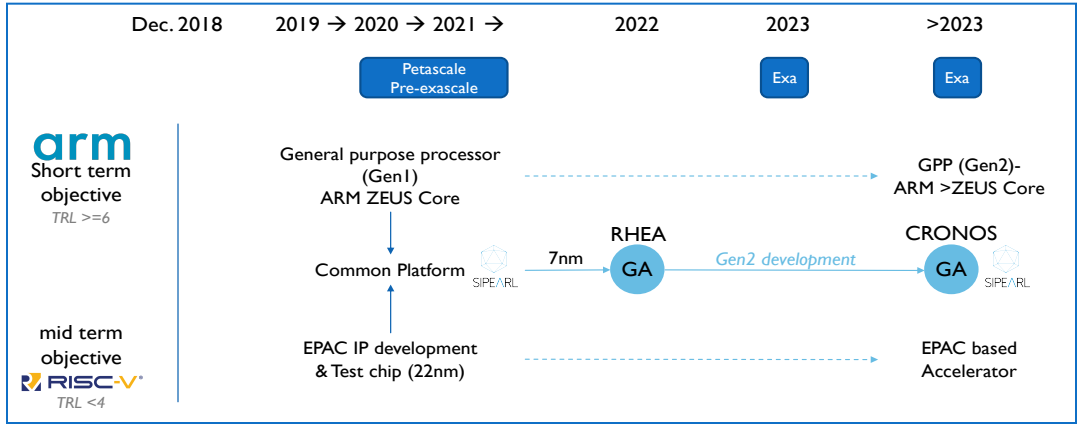

- Need extreme flexibility and performances on external links
  - Composable architectures, from 1 to >=4 sockets (CCIX and/or CXL)
  - HBM 2e / 3 and DDR 5 / 6
  - PCIe G5 / G6 and CXL 1 / 2
- Transparent integration in end-to-end dataflow :
  IoT ←→ Edge ←→ Datacenter ←→ Cloud
  - Easy to port / optimize
  - Opensource tools
  - Unified development tools
  - Compile one, run on many
- Need “good enough” but excellent FP64 performances.
  Need much better byte/Flop ratio than today.
  Target 0.5+ Byte/Flop
- Based on multi-die and IO-DIE building blocks, combined at package level or out of package ➔ common platform

Exascale and post-Exascale supercomputers will modular.

They’ll have massively non-homogenous architectures, combining one general purpose processor with several different accelerator kinds

Software will play an even more important role as a unification layer between all technologies, between all modules

→ Opensource and standardization are more important than ever

→ Proprietary SW stacks, especially for specialized HW will become problematic

THANK YOU FOR YOUR ATTENTION

European Processor Initiative

Jean-marc.denis@European-processor-initiative.eu

www.european-processor-initiative.eu

@EuProcessor

European Processor Initiative

European Processor Initiative