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

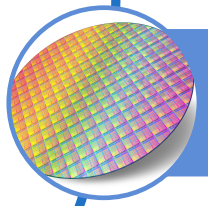

(Post) Exascale processor specifications

Wrap-Up

EPI IN EUROPEAN HPC STRATEGY

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

23 members from 10 EU countries

→ General Purpose processor in 2022

→ Accelerator IP

Brussels, 1 Dec 2017

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) – CONTEXT

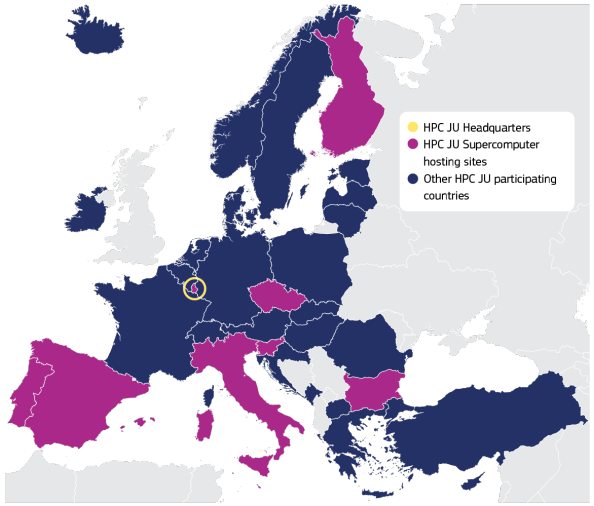

- The European High Performance Computing Joint Undertaking (EuroHPC JU) is pooling European resources to buy and deploy top-of-the-range supercomputers and develop innovative exascale supercomputing technologies and applications.

- 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
TBD	EFlops

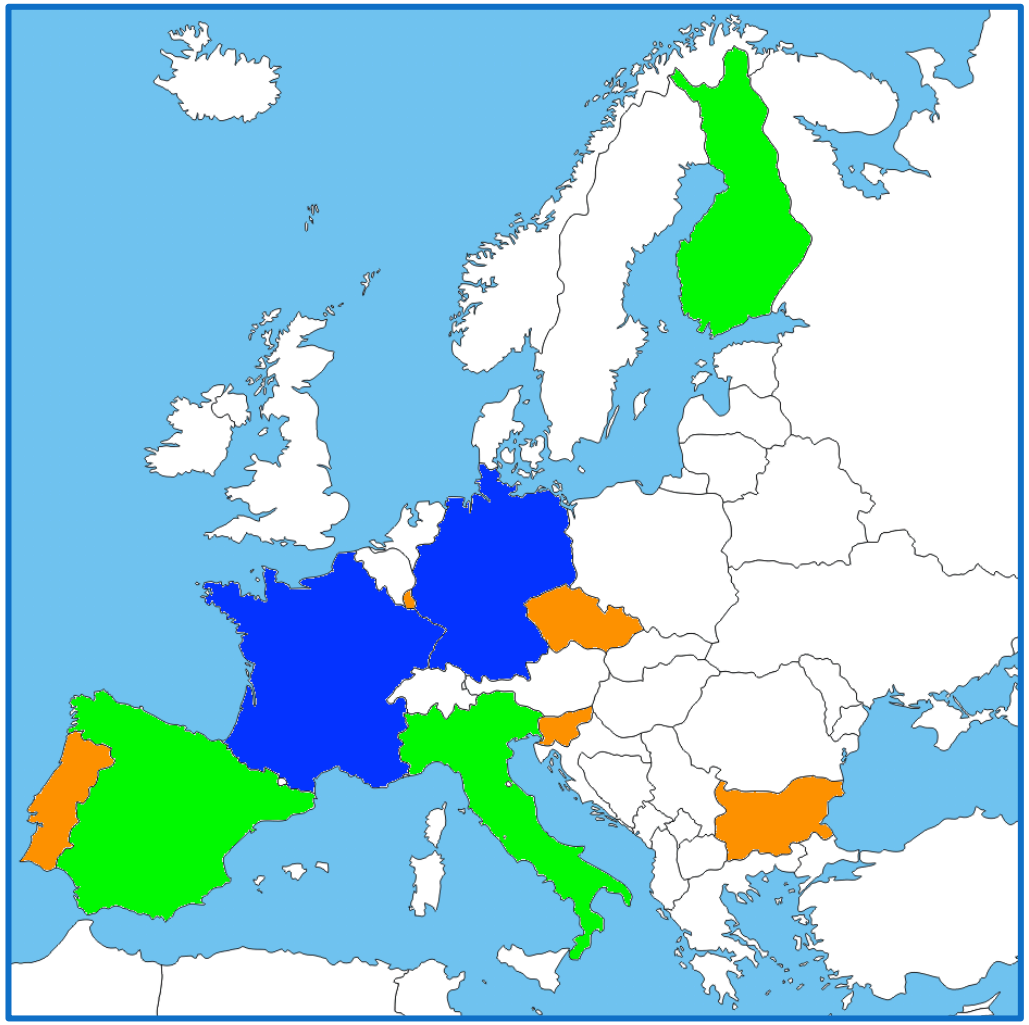

Country	Machine	Supplier	PFLOPS	Year
Finland	LUMI	HPE	550	2021
Italy	Leonardo	ATOS	248	2021
Spain(*)	*MareNostrum5*	*TBD*	*>200*	*2021*
Luxembourg	MeluXina	ATOS	10	2021
Portugal	Deucalion	ATOS Fujitsu	10	2021
CZ Rep	IT4I (name tbd)	HPE	15,2	2021
Bulgaria	NCSA	ATOS	6	2021
Slovenia	Vega	ATOS	6,8	2021
TBD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

(*) announced on Dec. 15th at the earliest
(**) with EU processor based on EPI developments)

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 1-3 are for long-term objectives (EU IP)

  and
  - TRL 4-9 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 (~100M€)
- generate revenue from both the HPC, IA, server and eHPC markets
- integrate, market, support & sell the chip
- work on the next generations

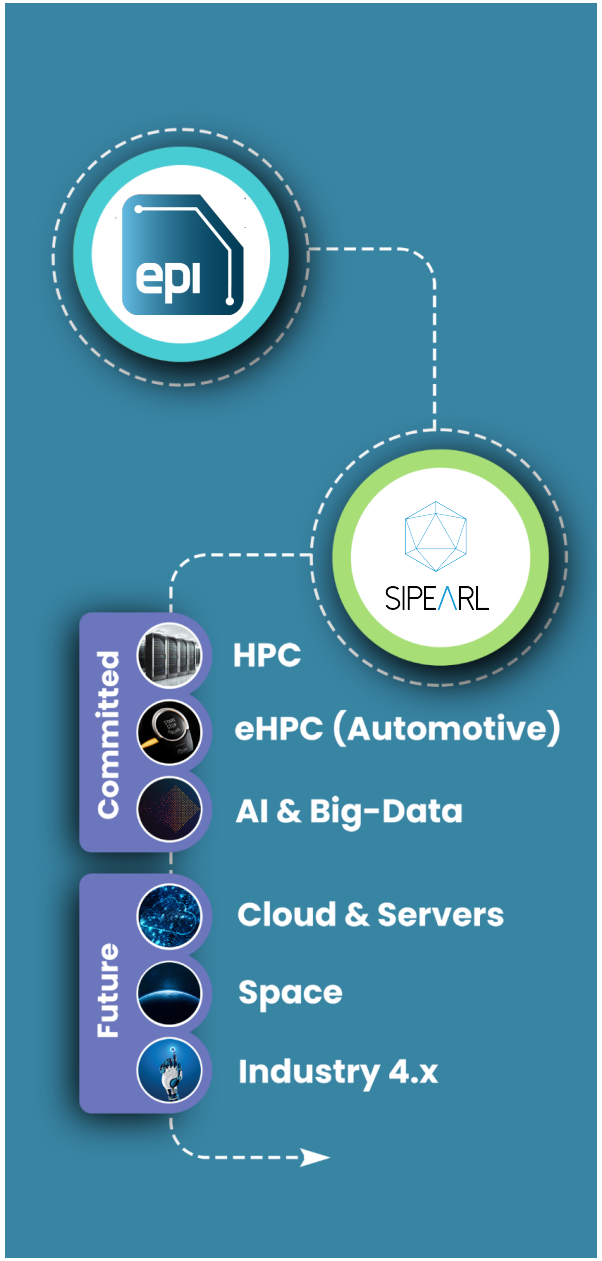

OVERALL ROADMAP

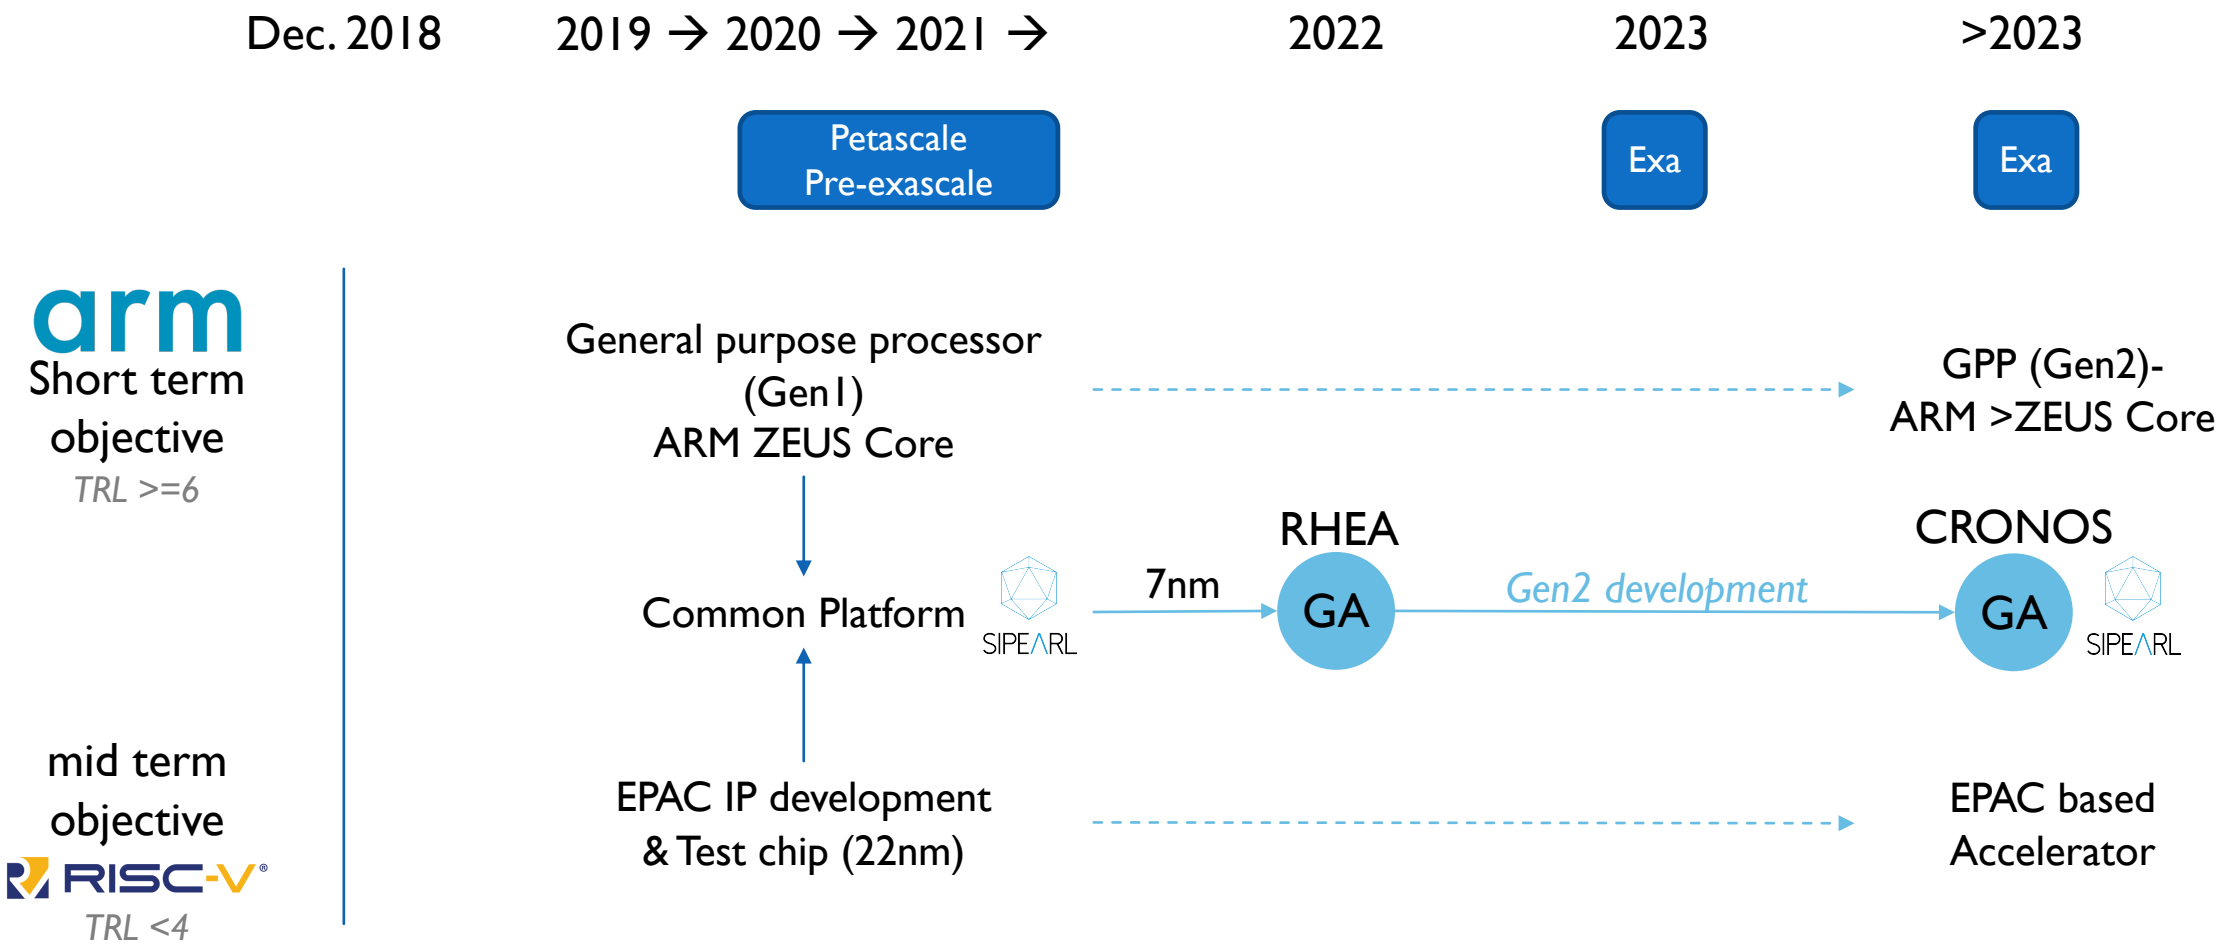

(POST)-EXASCALE SUPERCOMPUTERS SPECIFICATIONS

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:IS.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.0.5D0) 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)
```

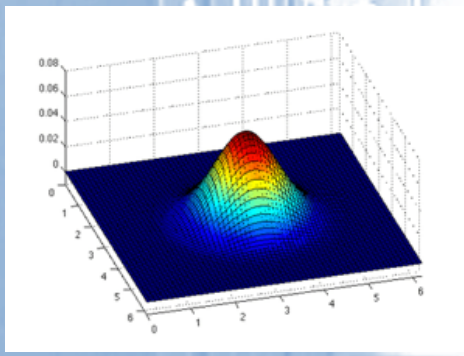

HPC WITH ARTIFICIAL INTELLIGENCE

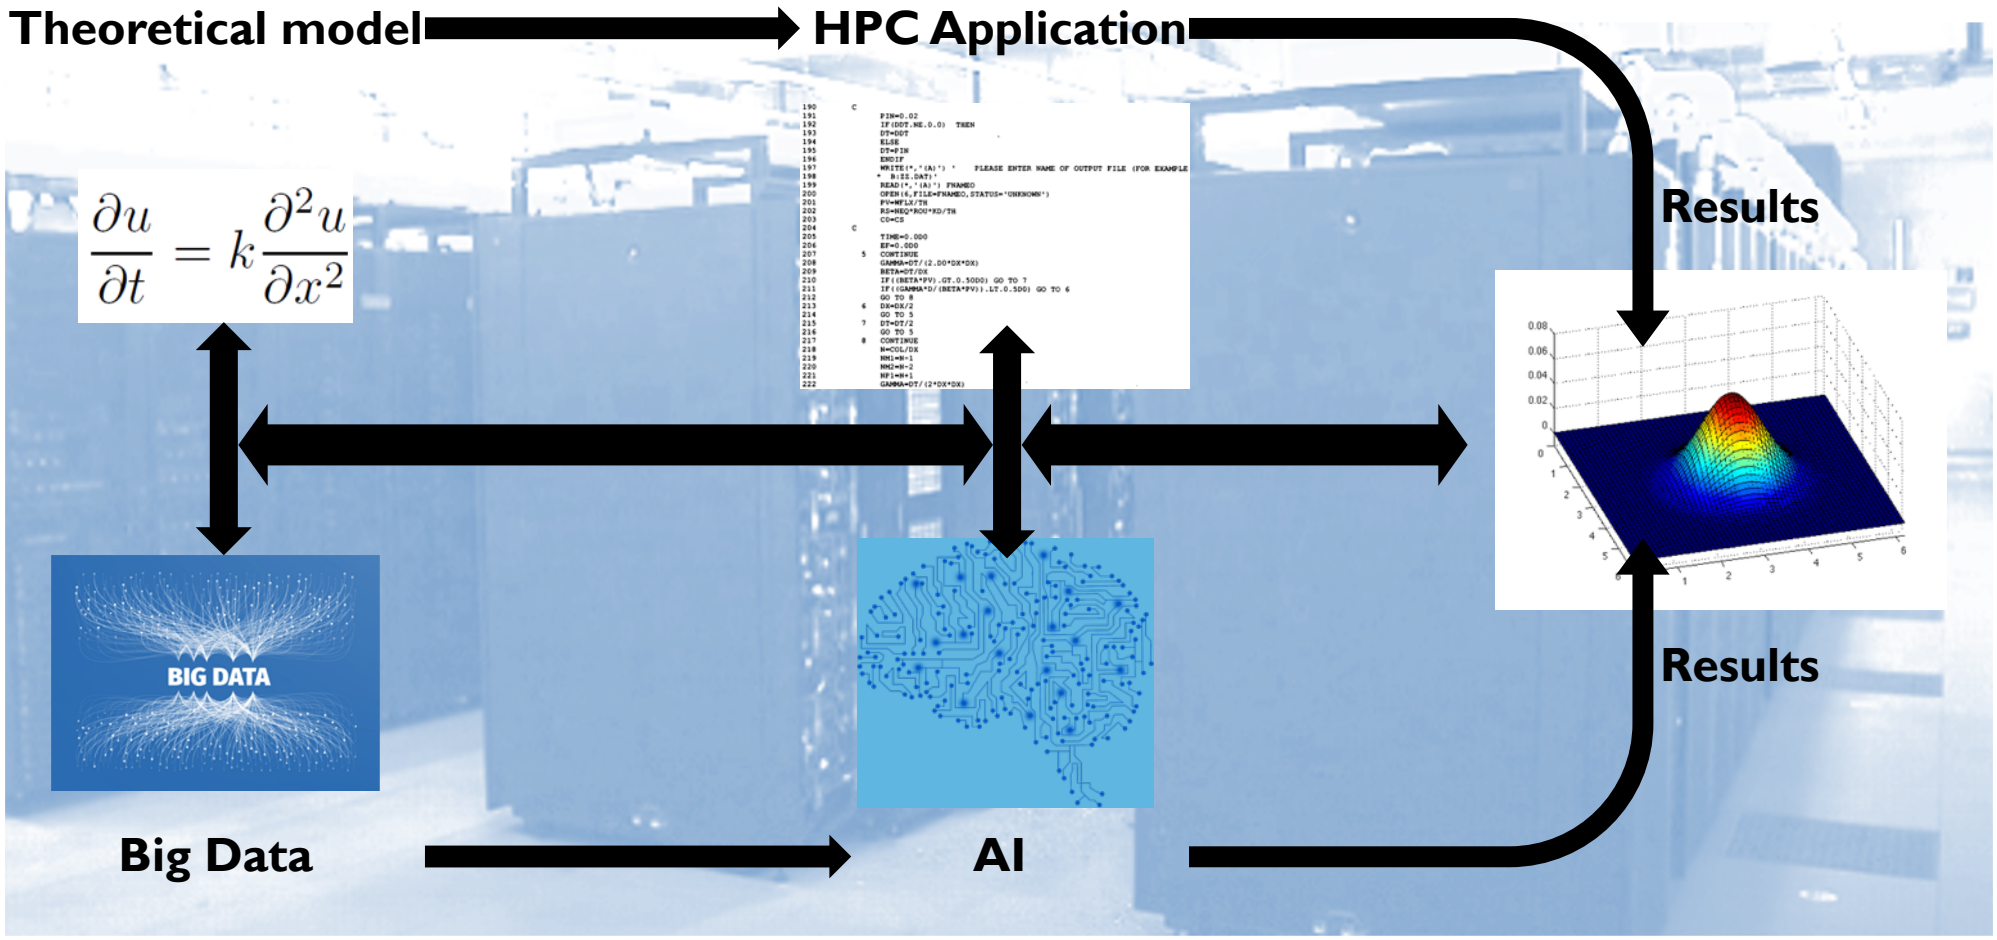

HPC & AI AT EXASCALE: IT'S ALL ABOUT WORKFLOWS (1/3)

ARM

X86
ARM

Edge Computing

Realtime data processing

X86
ARM

Public Clouds

SUMMIT

sunway taihulight

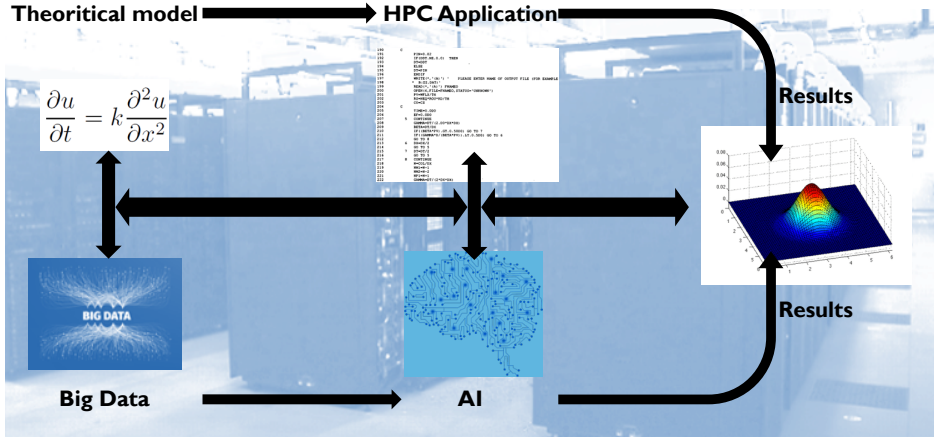

FUGAKU

HPC & AI AT EXASCALE: IT'S ALL ABOUT WORKFLOWS (2/3)

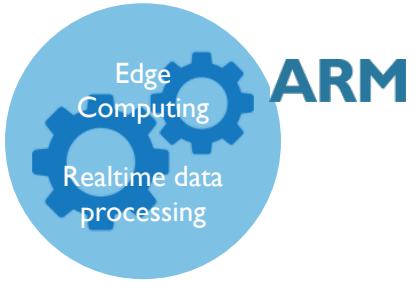

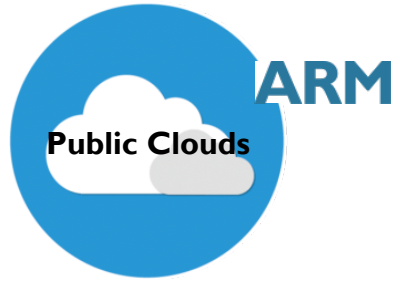

European approach

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

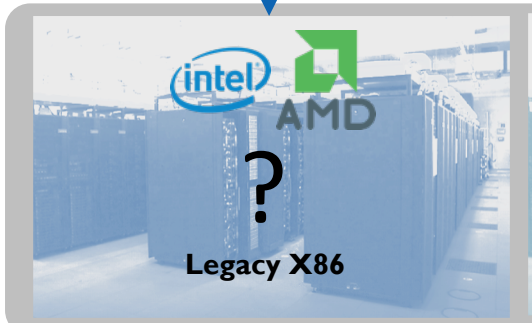

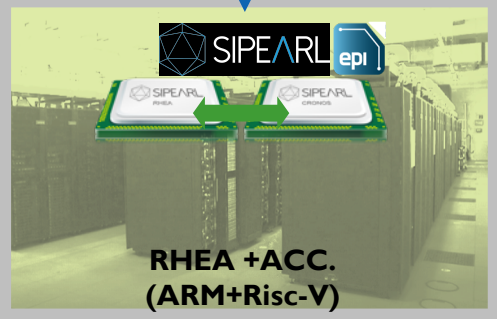

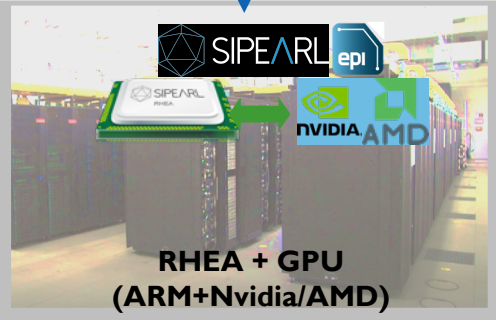

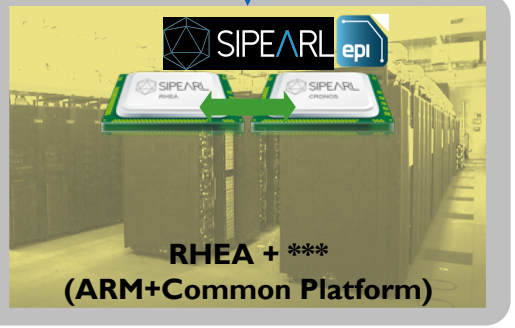

HPC & AI AT EXASCALE: IT'S ALL ABOUT WORKFLOWS (3/3)

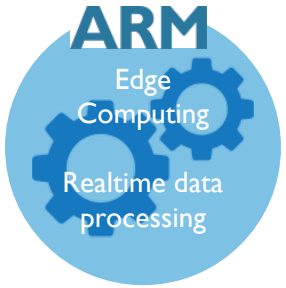

European approach

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

Unified environment: Compilers / programming languages / libraries

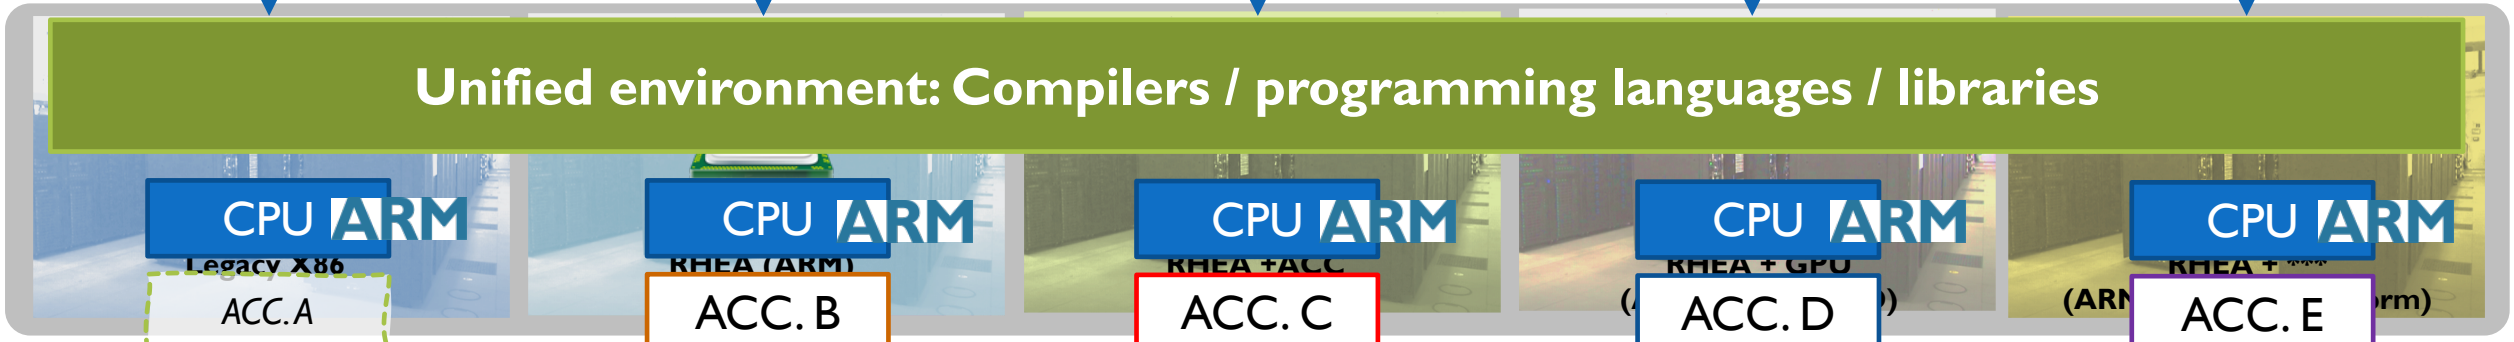

TAKEWAY #1

Exascale and post-Exascale superpercomputers will modular.

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

TAKEWAY #2

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

(POST) EXASCALE PROCESSOR SPECIFICATIONS

TAKE AWAYS #1 & #2 CONSEQUENCES

General Purpose Processors have to be (much) more open

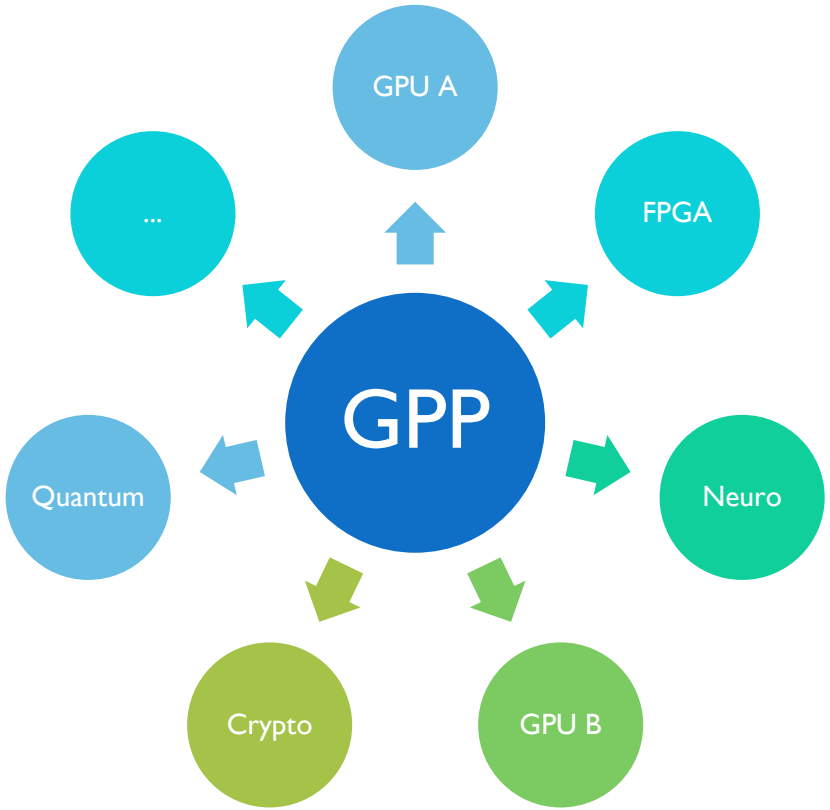

The race to FLOPS is now in the accelerators area

TAKEWAY #3
SPECIFICATION FOR EXASCALE GENERAL PURPOSE PROCESSORS (PART 1)

- 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
  - *and* PCIe G5
  - *and* CXL
- Transparent integration in end-to-end dataflow : IoT ←→ Edge ←→ Datacenter ←→ Cloud
  - Easy to port / optimize
  - Opensource tools
  - Unified development tools

THE R_PEAK CHALLENGE – CPU VS. GPU

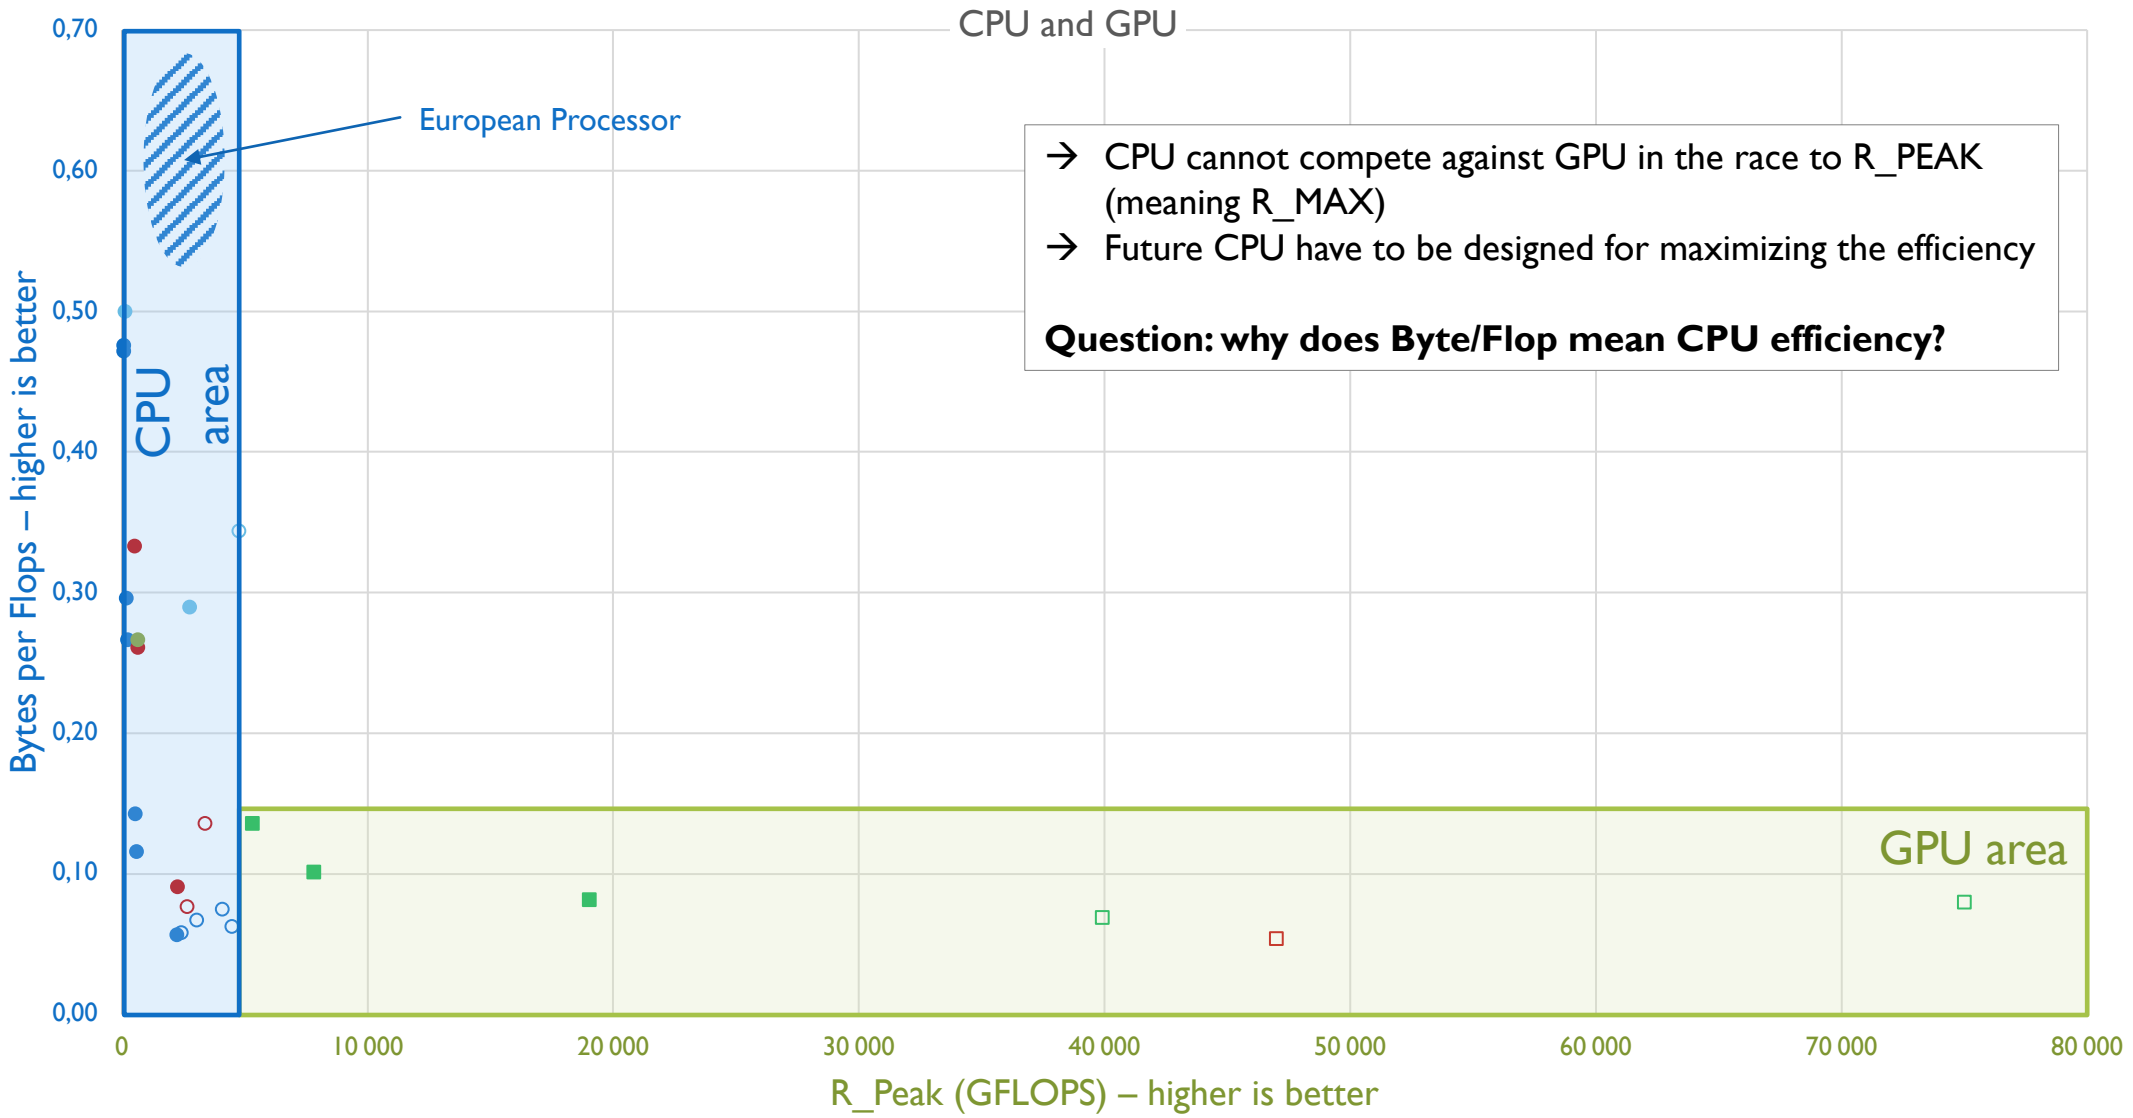

THE R_PEAK CHALLENGE – FOCUS ON CPU

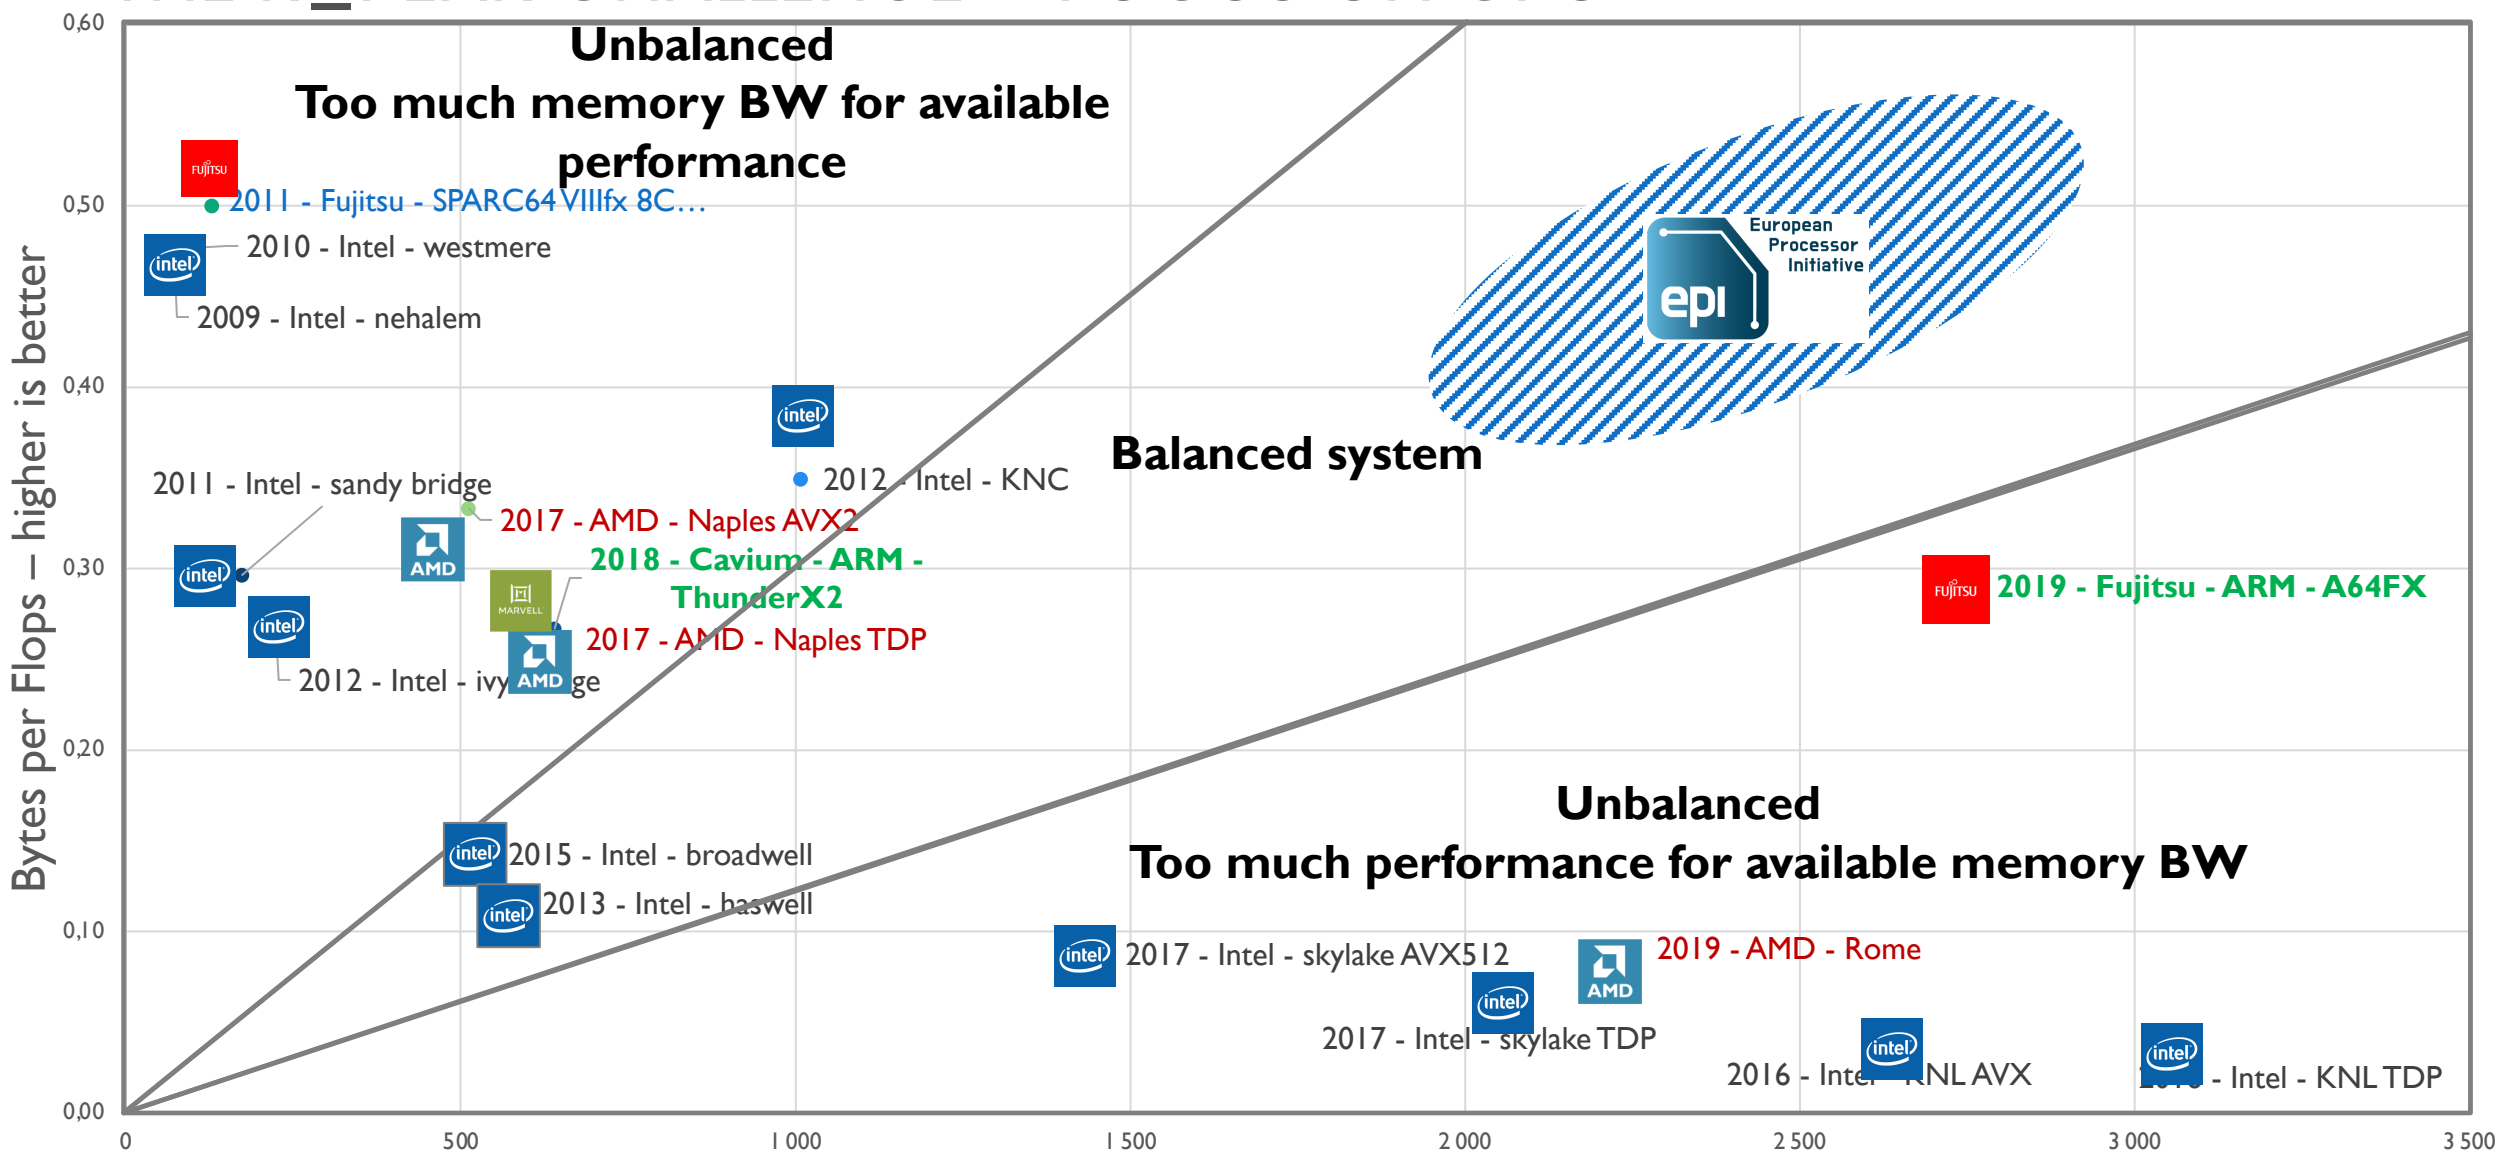

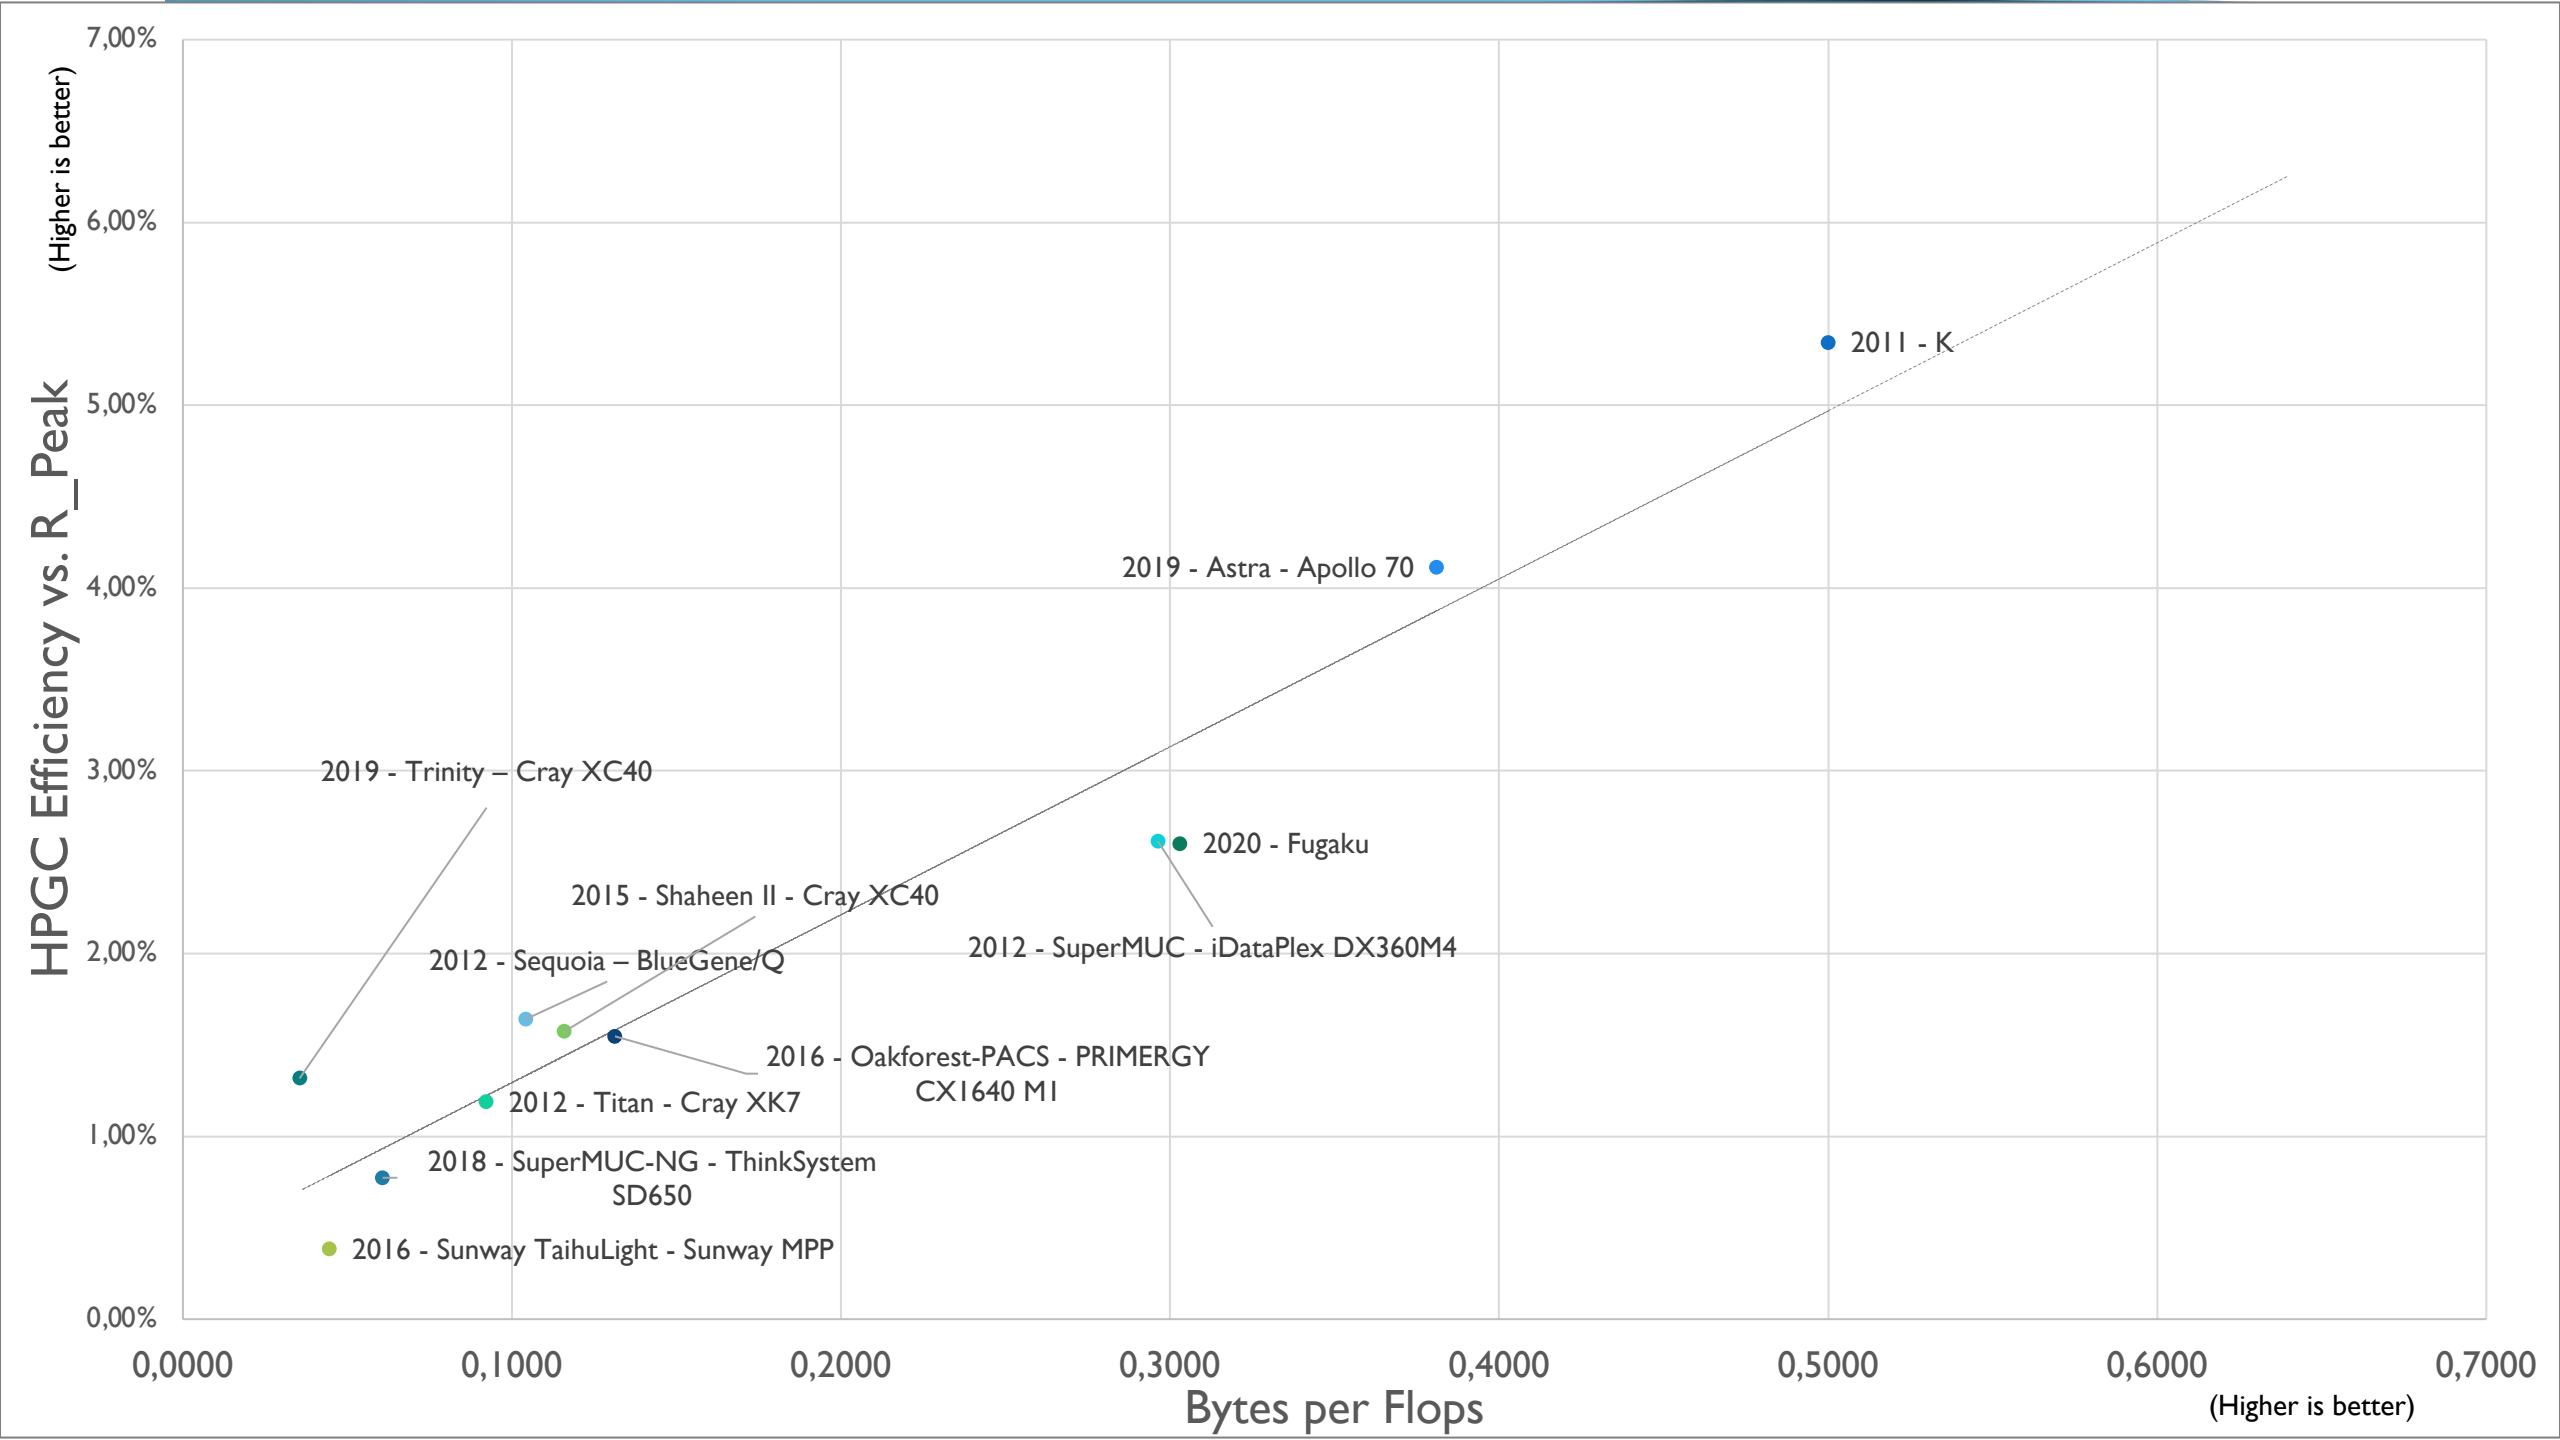
HPGC Efficiency vs. R_Peak
(Higher is better)
7,00%
6,00%
5,00%
4,00%
3,00%
2,00%
1,00%
0,00%
0,0000
0,1000
0,2000
0,3000
0,4000
0,5000
0,6000
0,7000
Bytes per Flops
(Higher is better)
2011 - K
2019 - Astra - Apollo 70
2020 - Fugaku
2012 - SuperMUC - iDataPlex DX360M4
2019 - Trinity – Cray XC40
2015 - Shaheen II - Cray XC40
2012 - Sequoia – BlueGene/Q
2016 - Oakforest-PACS - PRIMERGY CX1640 M1
2012 - Titan - Cray XK7
2018 - SuperMUC-NG - ThinkSystem SD650
2016 - Sunway TaihuLight - Sunway MPP

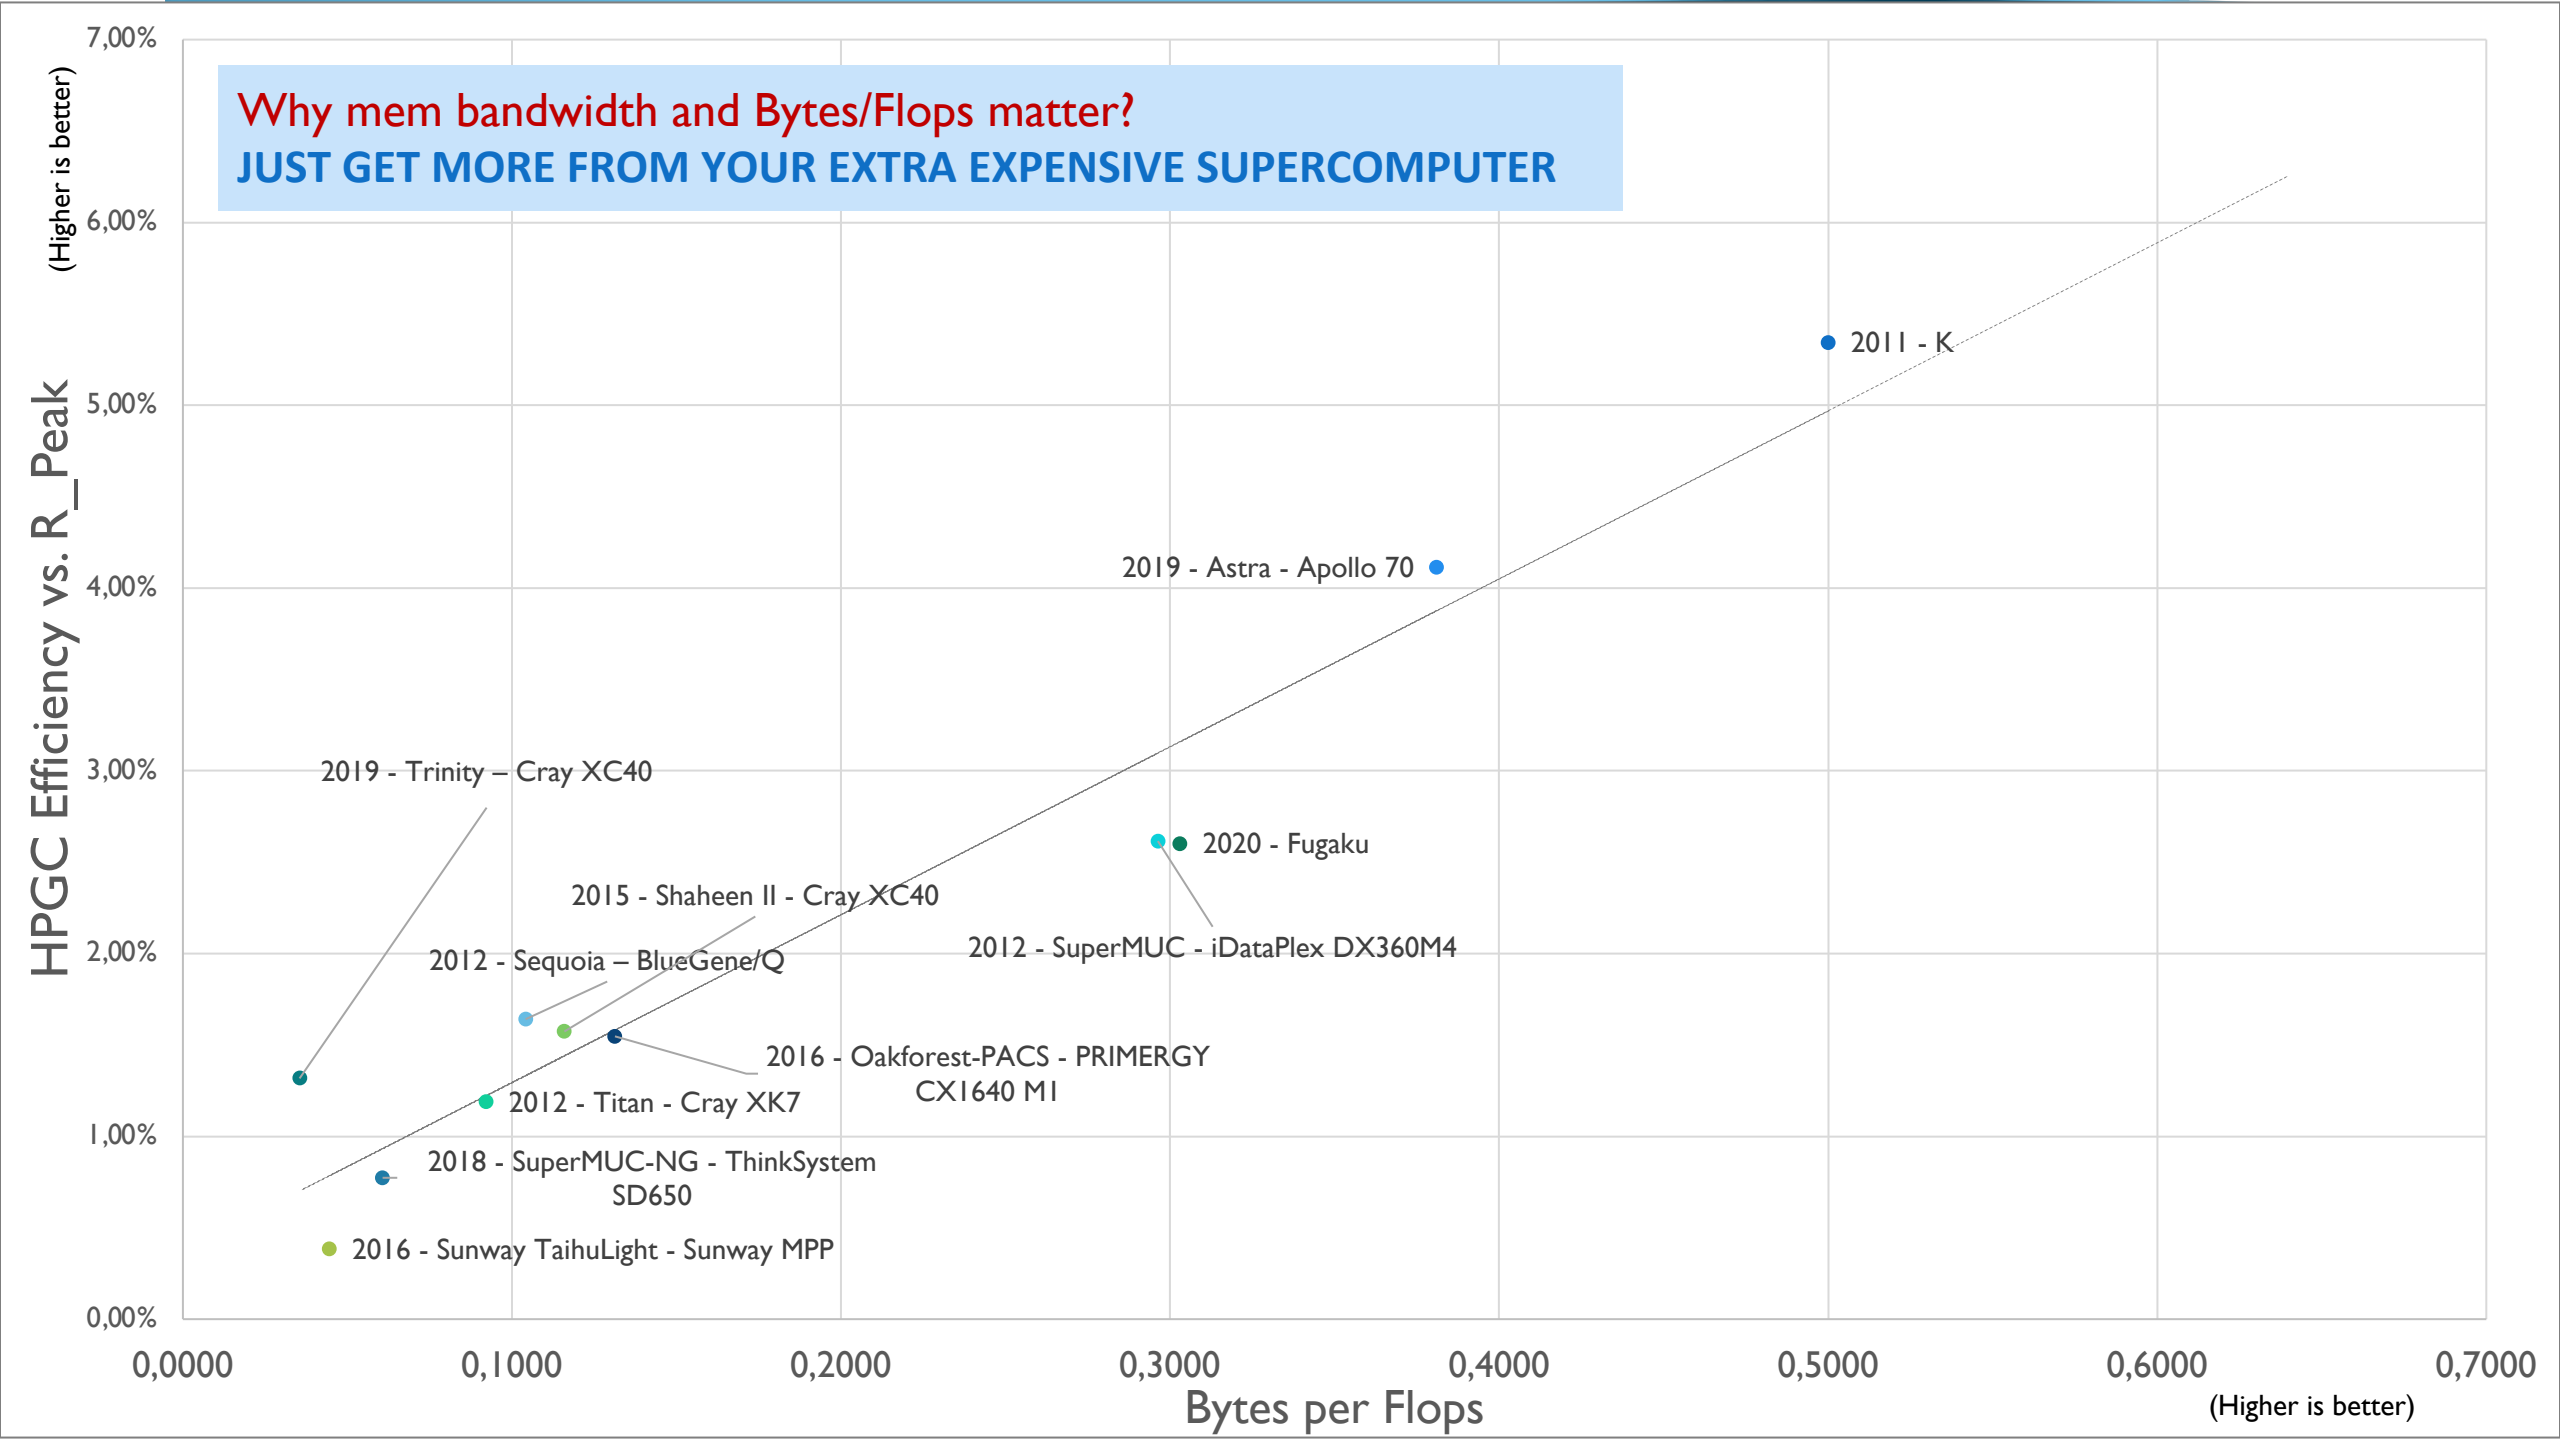
Why mem bandwidth and Bytes/Flops matter?
JUST GET MORE FROM YOUR EXTRA EXPENSIVE SUPERCOMPUTER
HPGC Efficiency vs. R_Peak
(Higher is better)
7,00%
6,00%
5,00%
4,00%
3,00%
2,00%
1,00%
0,00%
2011 - K
2019 - Astra - Apollo 70
2020 - Fugaku
2012 - SuperMUC - iDataPlex DX360M4
2019 - Trinity – Cray XC40
2015 - Shaheen II - Cray XC40
2012 - Sequoia – BlueGene/Q
2016 - Oakforest-PACS - PRIMERGY CX1640 M1
2012 - Titan - Cray XK7
2018 - SuperMUC-NG - ThinkSystem SD650
2016 - Sunway TaihuLight - Sunway MPP
0,0000
0,1000
0,2000
0,3000
0,4000
0,5000
0,6000
0,7000
Bytes per Flops
(Higher is better)

Why mem bandwidth and Bytes/Flops matter?

JUST GET MORE FROM YOUR EXTRA EXPENSIVE SUPERCOMPUTER

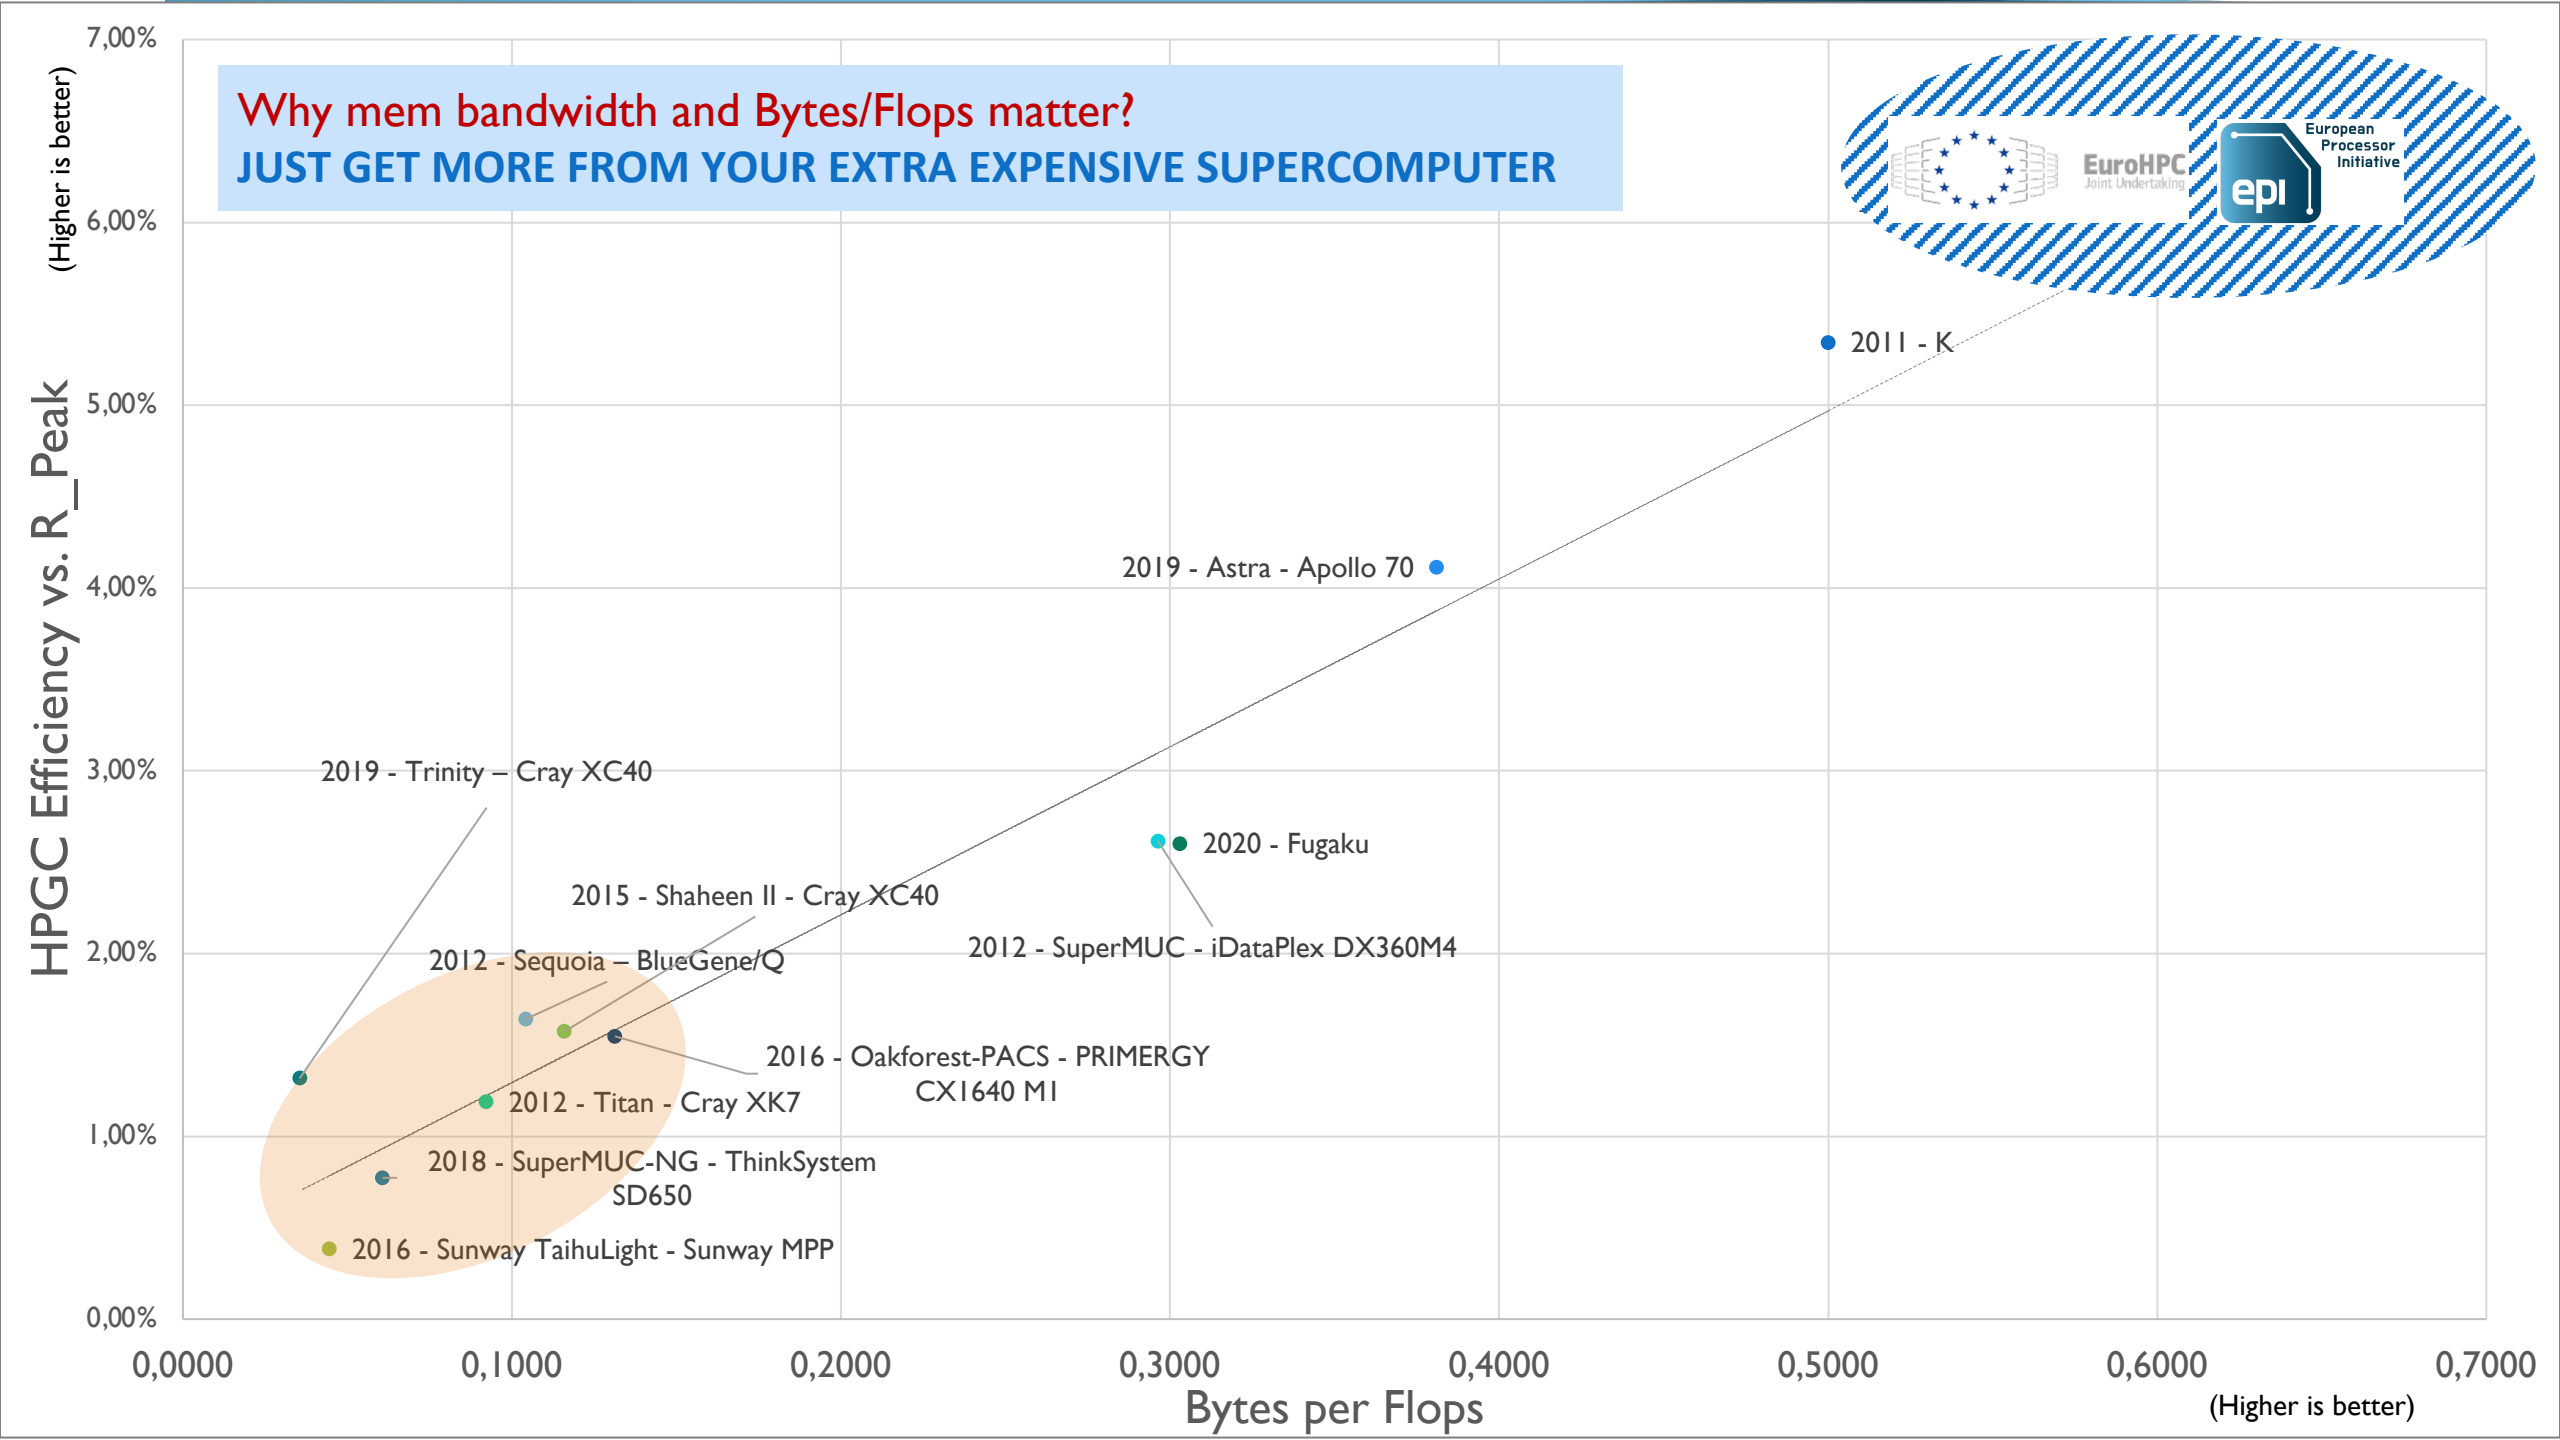

SO… WHAT TO EXPECT FROM AN EXASCALE GENERAL PURPOSE PROCESSOR? (PART 2)

- 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
  - *and* PCIe G5
  - *and* CXL
- Transparent integration in end-to-end dataflow : IoT ←→ Edge ←→ Datacenter ←→ Cloud
  - Easy to port / optimize
  - Opensource tools
  - Unified development tools
- **Need “good enough” FP64 performances.**
  No need to compete with specialized devices like GPUs
- **Change the metric. Need much better byte/Flop ratio than today**

GPP ARCHITECTURE CONCEPT

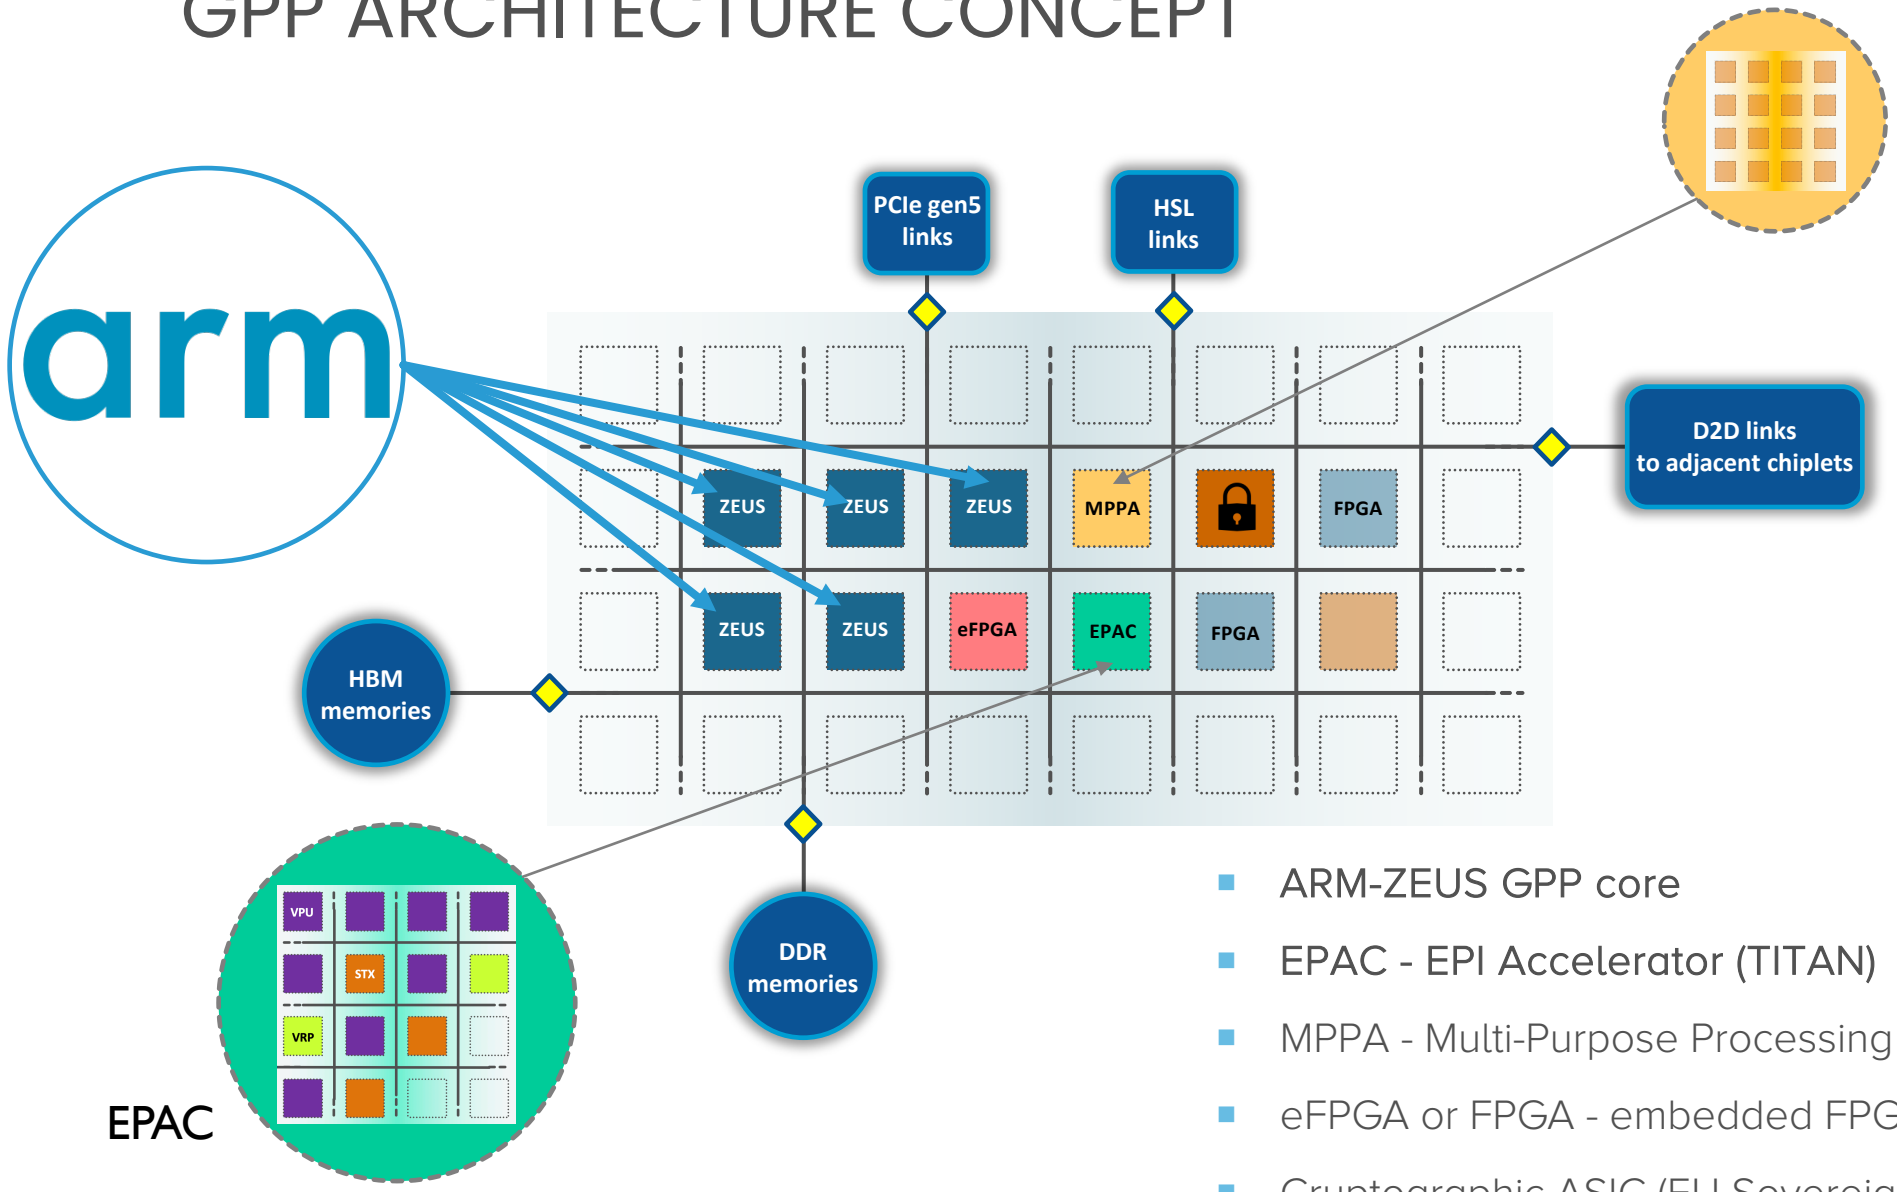

- ARM-ZEUS GPP core
- EPAC - EPI Accelerator (TITAN)
- MPPA - Multi-Purpose Processing Array
- eFPGA or FPGA - embedded FPGA
- Cryptographic ASIC (EU Sovereignty)

CONCEPT OF COMMON PLATFORM : INTERPOSER & MULTI-CHIPLET

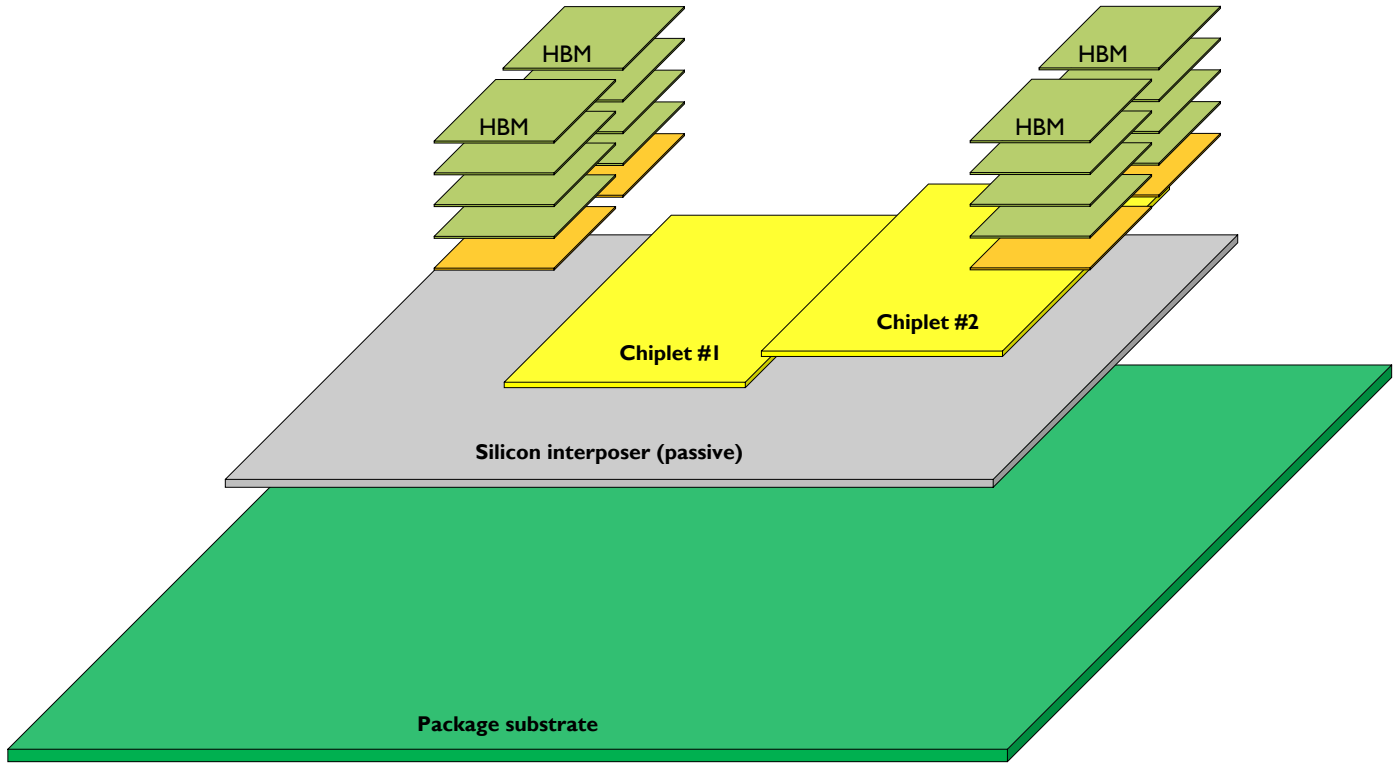

CONCEPT OF COMMON PLATFORM : INTERPOSER & MULTI-CHIPLET

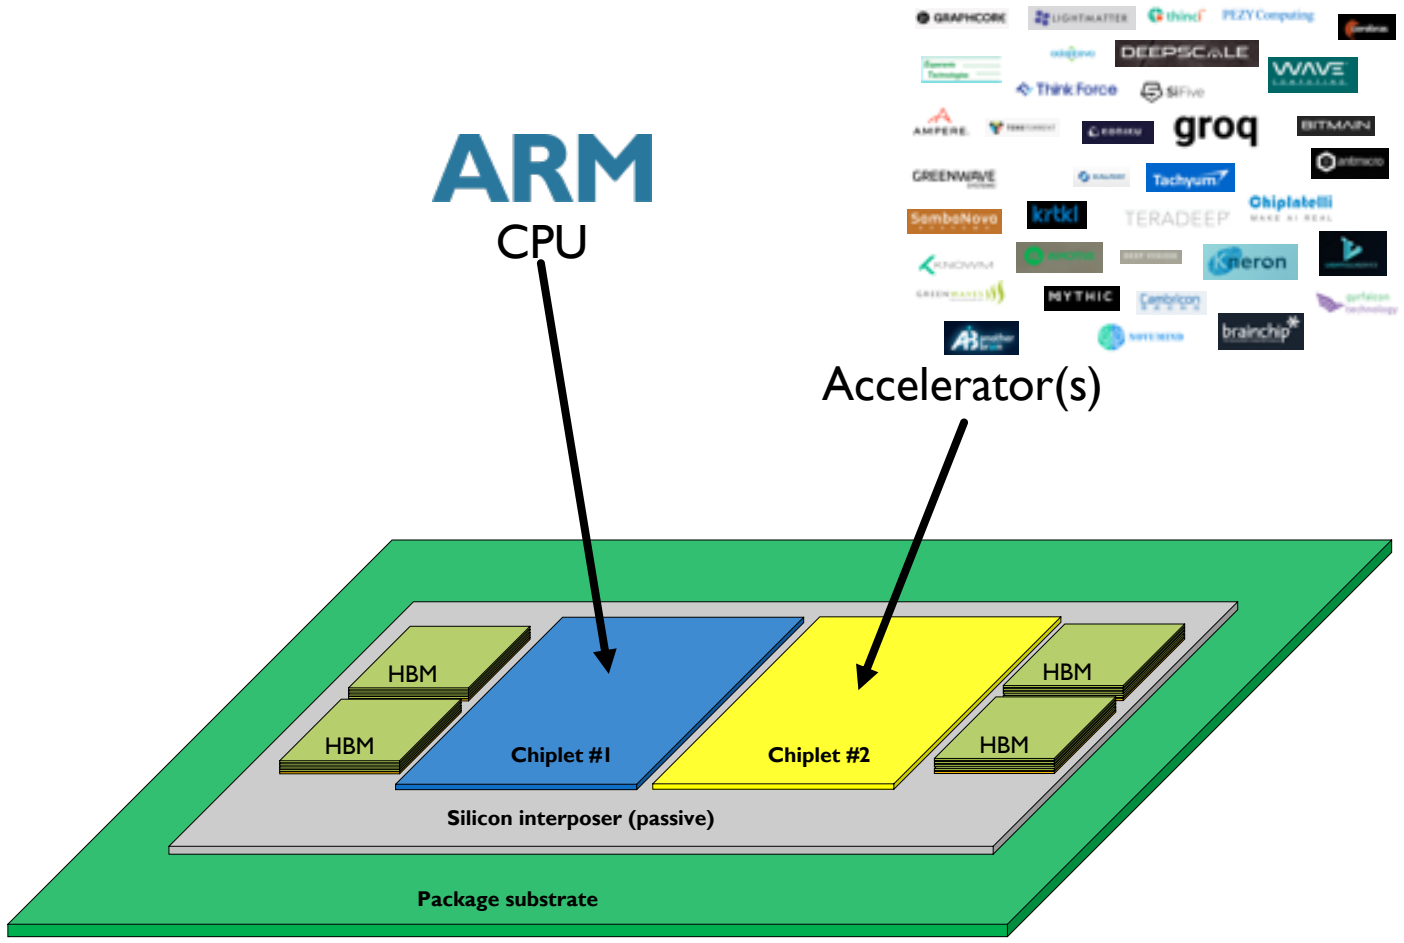

HETEROGENEOUS INTEGRATION

- Integrating customized functions at different levels
- EPI accelerator IPs today are integrated in Rhea design

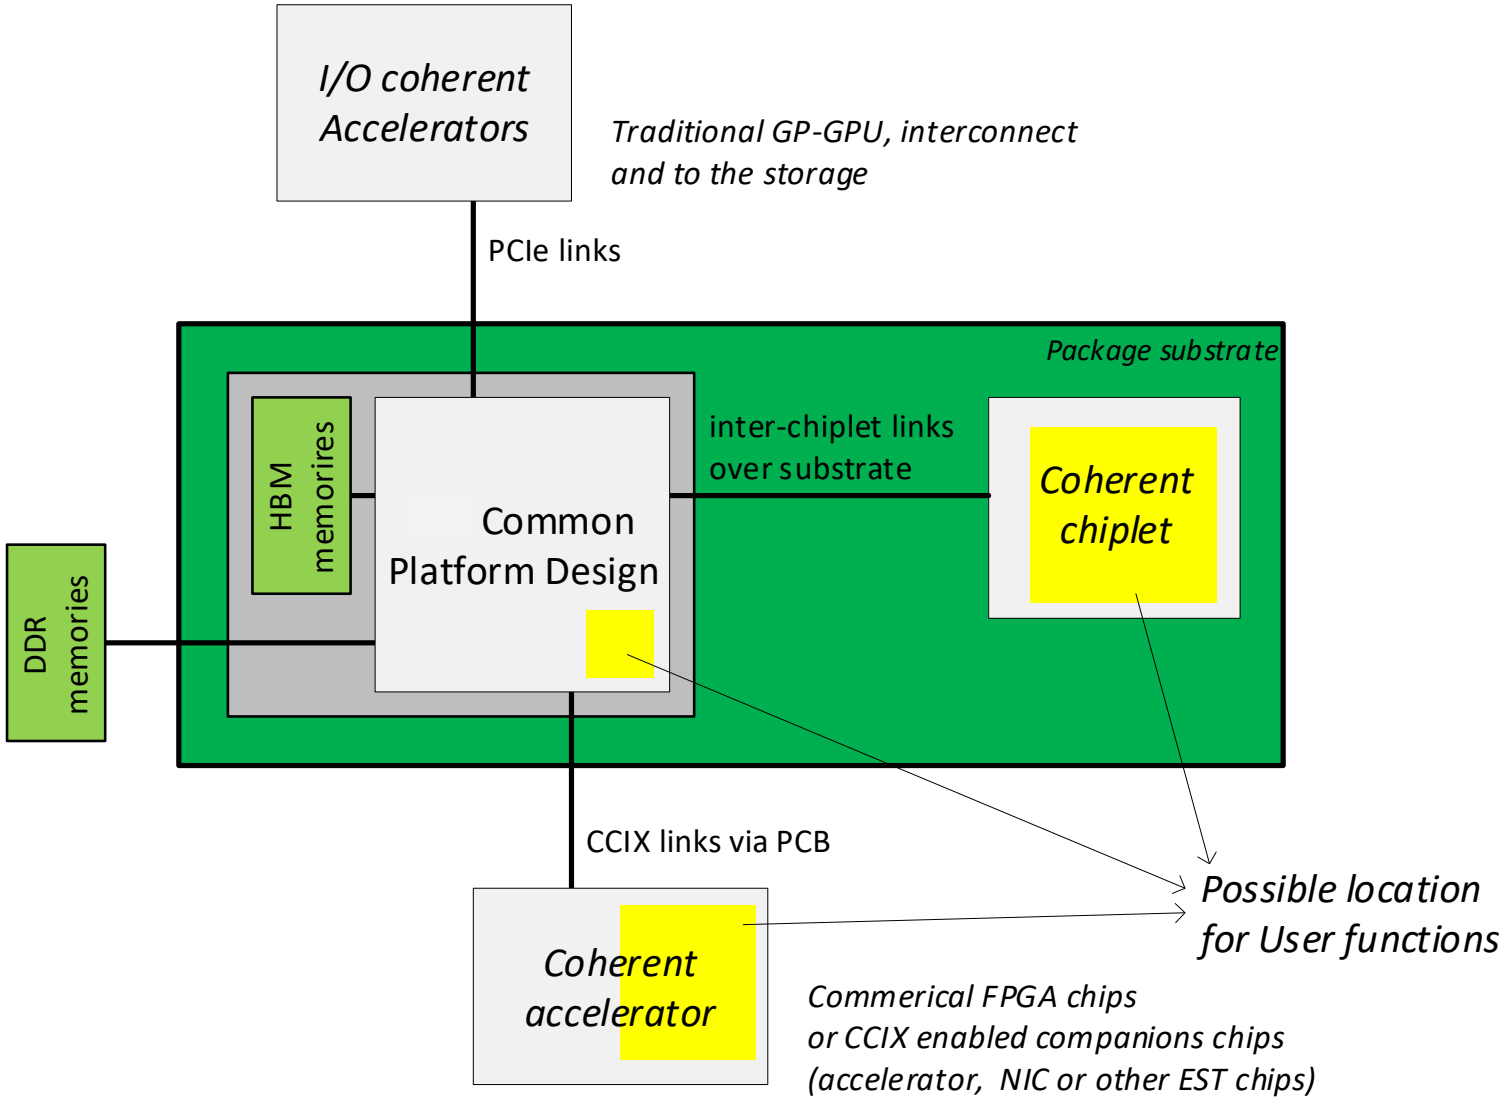

SO… WHAT TO EXPECT FROM AN EXASCALE GENERAL PURPOSE PROCESSOR? (PART 3)

- 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
  - *and* PCIe G5
  - *and* CXL
- Transparent integration in end-to-end dataflow : IoT ←→ Edge ←→ Datacenter ←→ Cloud
  - Easy to port / optimize
  - Opensource tools
  - Unified development tools
- Need “good enough” but excellent FP64 performances.
  No need to compete with specialized devices like GPUs
- Need much better byte/Flop ratio than today
- **Based on multi-die and IO-DIE building blocks, combined at package level**

WRAP-UP

FINAL TAKEAWAYS

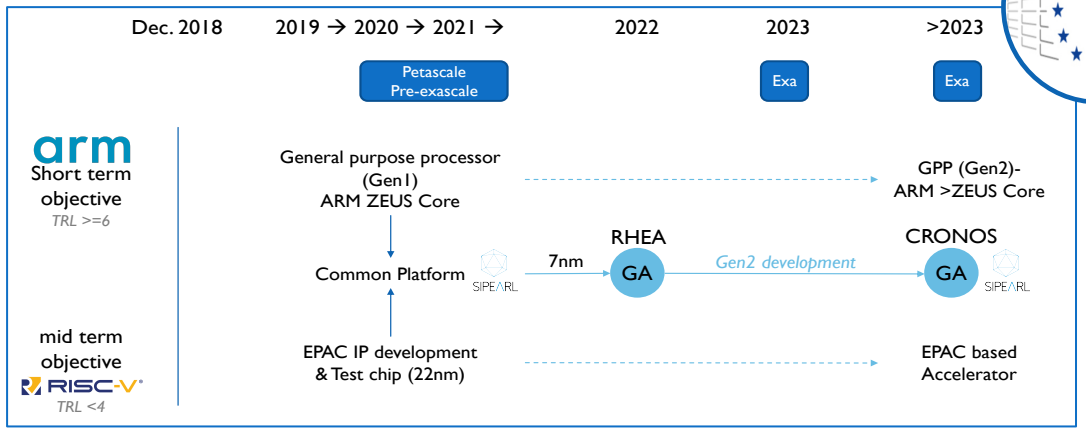

Exascale and post-Exascale supercomputers will modular.

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

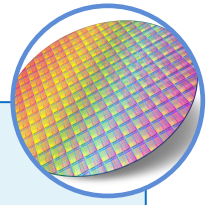

- World class manufacturing process (7nm or better)
- Need extreme flexibility and performances on external links
  - Composable architectures, from 1 to >=4 sockets (CCIX and/or CXL) HBM 2e / 3
  - *and* DDR 5
  - *and* PCIe G5
  - *and* CXL
- Transparent integration in end-to-end dataflow
  IoT ←→ Edge ←→ Datacenter ←→ Cloud
  - Easy to port / optimize
  - Opensource tools
  - Unified development tools
- Need "good enough" but excellent FP64 performances.
  No need to compete with specialized devices like GPUs
- Need much better byte/Flop ratio than today
- Based on multi-die and IO-DIE building blocks, combined at package level

THANK YOU FOR YOUR ATTENTION

European Processor Initiative

- www.european-processor-initiative.eu
- @EuProcessor
- European Processor Initiative
- European Processor Initiative