THE EUROPEAN APPROACH FOR EXASCALE AGES

HOW DOES THE FUTURE PROCESSORS LOOK LIKE?

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

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**

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

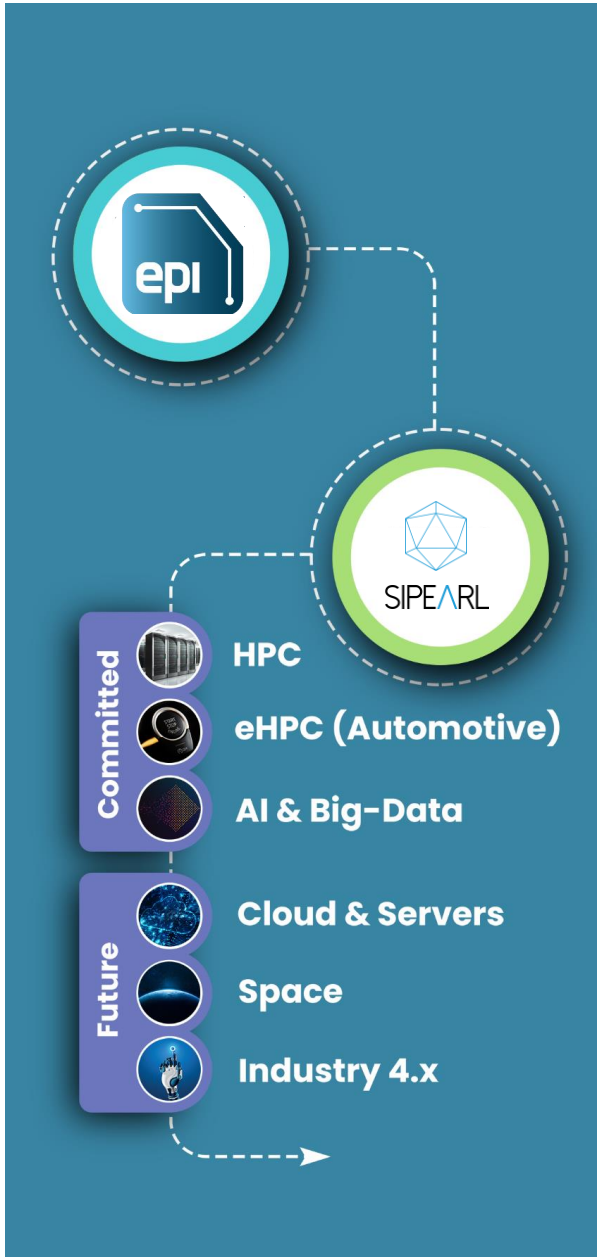

FROM OBJECTIVES TO ROADMAP, FROM ROADMAP TO PRODUCTS

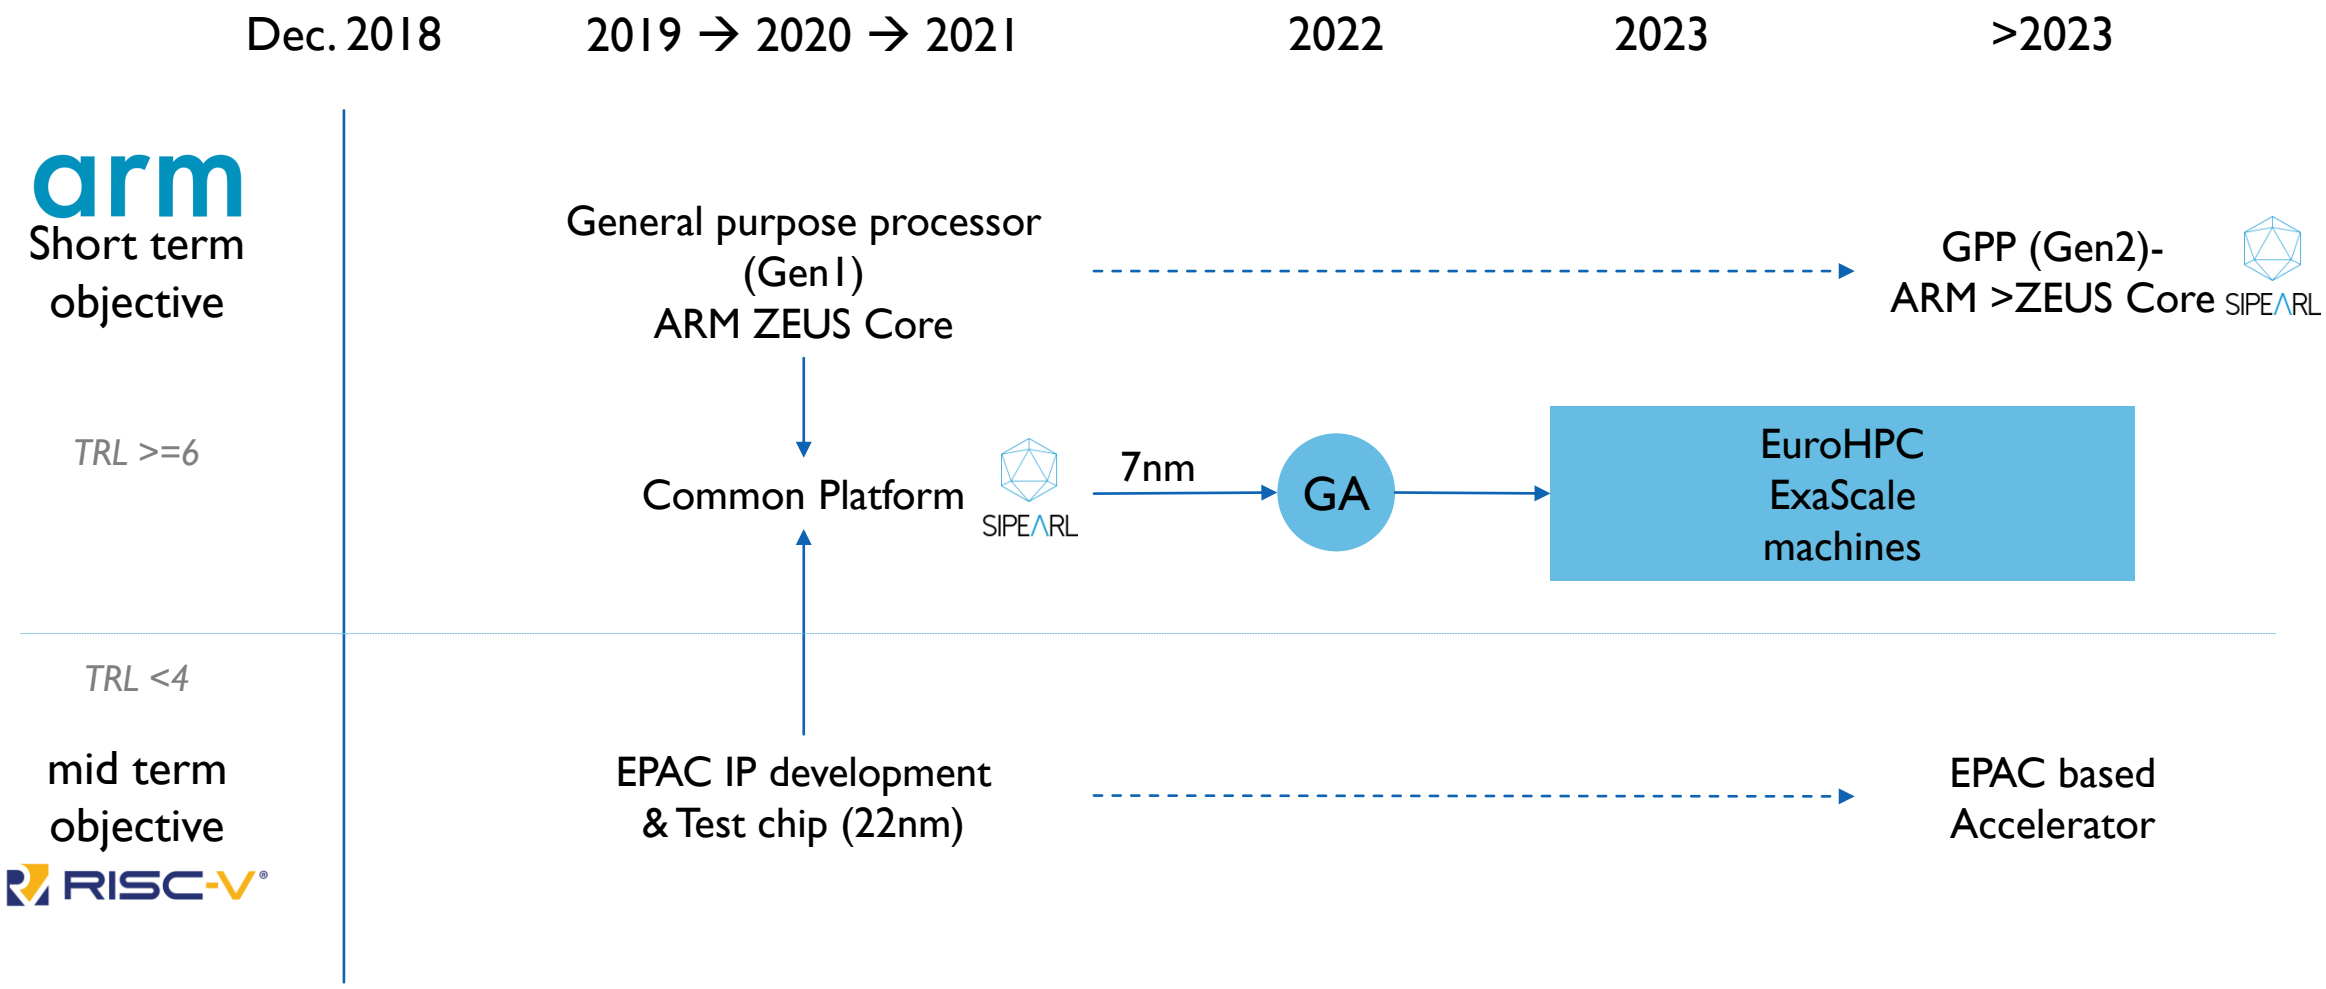

(POST)-EXASCALE SUPERCOMPUTERS OVERALL SPECIFICATIONS

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(4,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)
```

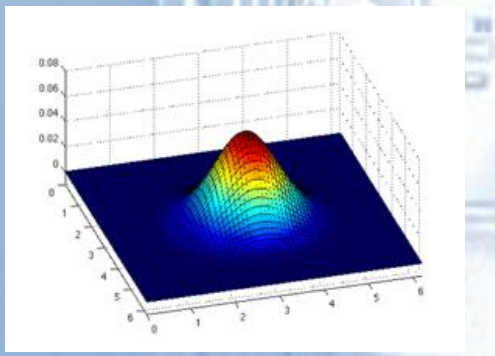

HPC WITH ARTIFICIAL INTELLIGENCE

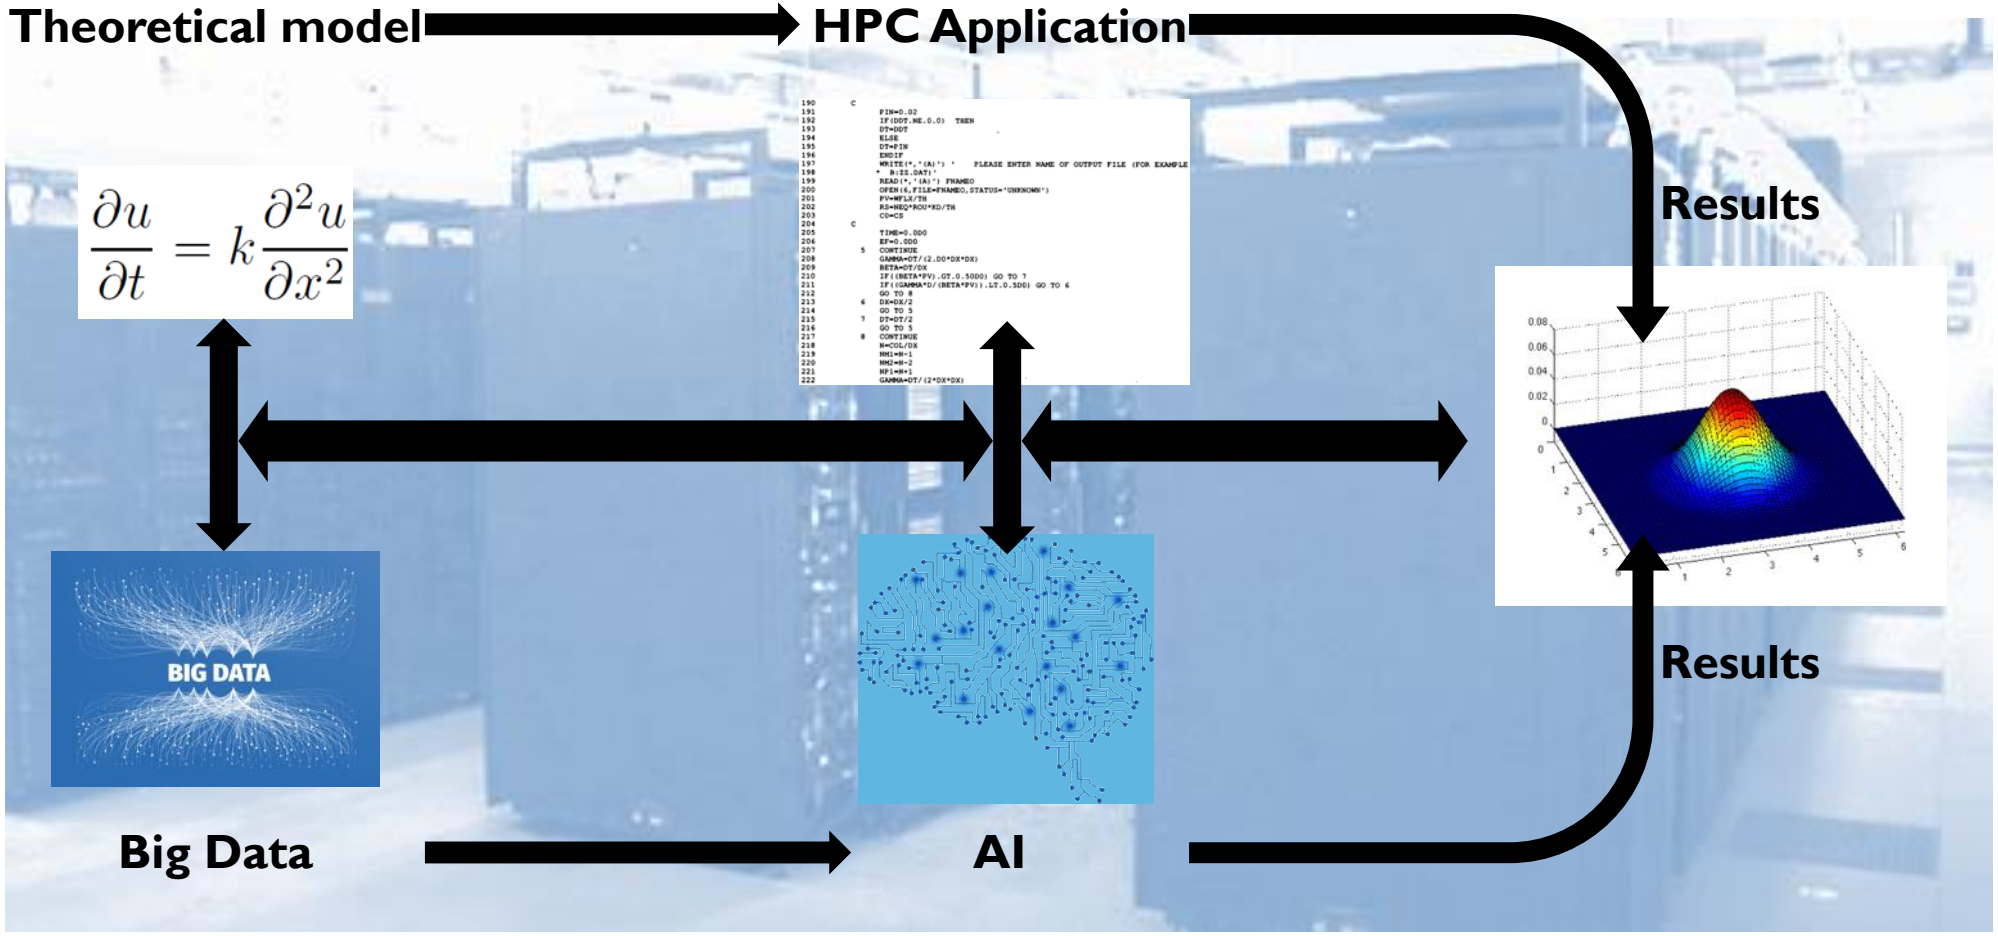

HPC & AI AT EXASCALE: IT'S ALL ABOUT WORKFLOWS (1/3)

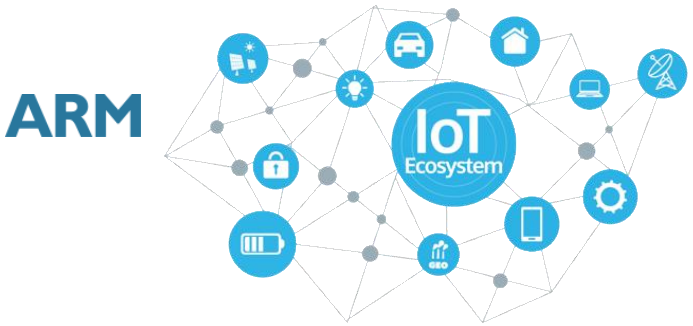

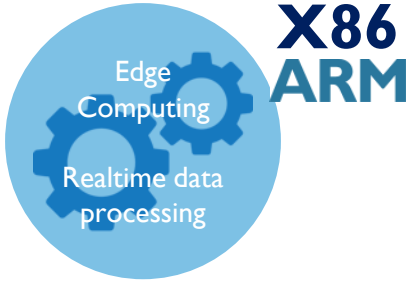

SUMMIT

sunway taihulight

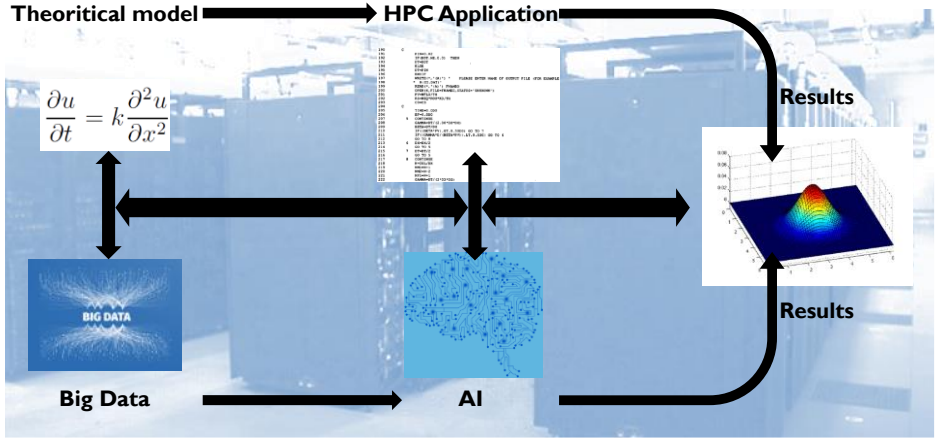

FUGAKU

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

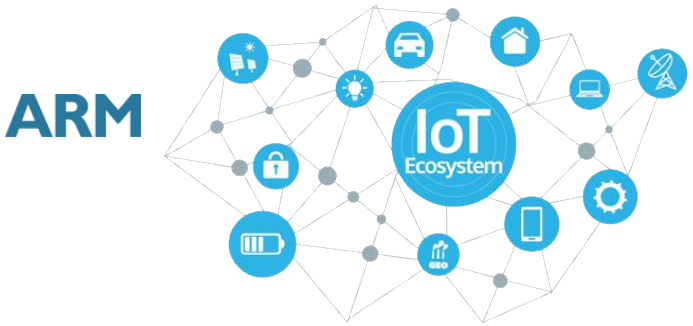

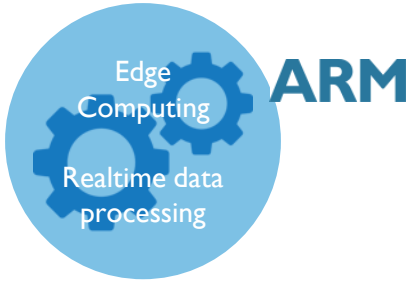

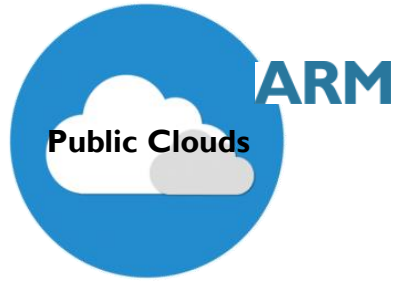

European approach

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

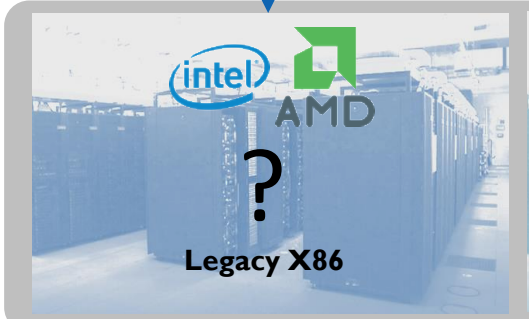

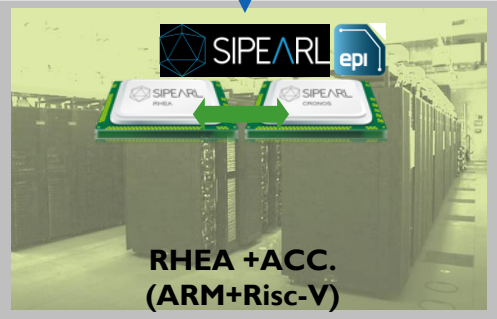

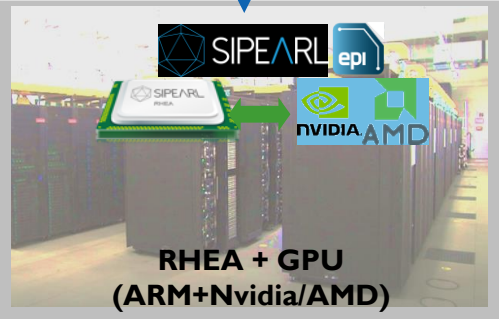

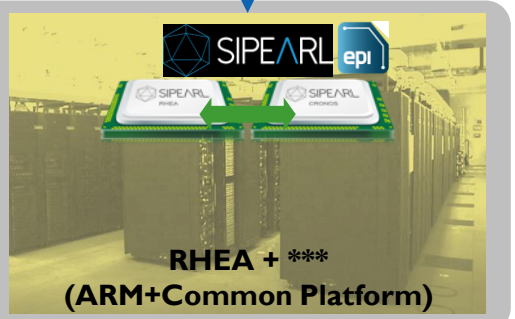

HPC & AI AT EXASCALE: IT'S ALL ABOUT WORKFLOWS (3/3)

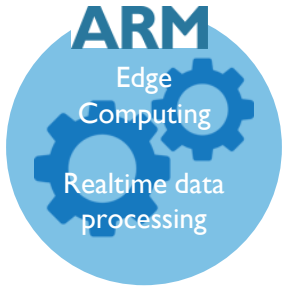

European approach

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

Unified environment: Compilers / programming languages / libraries

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 + [illegible] (ARM) — ACC. E

TAKEWAY #1

- The HPC Datacenter will “disappear”: users need a continuum between sensors (IoT) and data analysis (supercomputers) through edge layer.
- Hybrid computing (in-house & Cloud) is a must
- Workflows from IoT to Supercomputers need seamless dataflows
  ➔ develop once, run on many

TAKEWAY #2

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 #3

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

SPECIFICATIONS FOR AN EXASCALE CLASS GENERAL PURPOSE PROCESSOR

FROM TAKE AWAYS #1 TO #3

General Purpose Processors have to be (much) more open

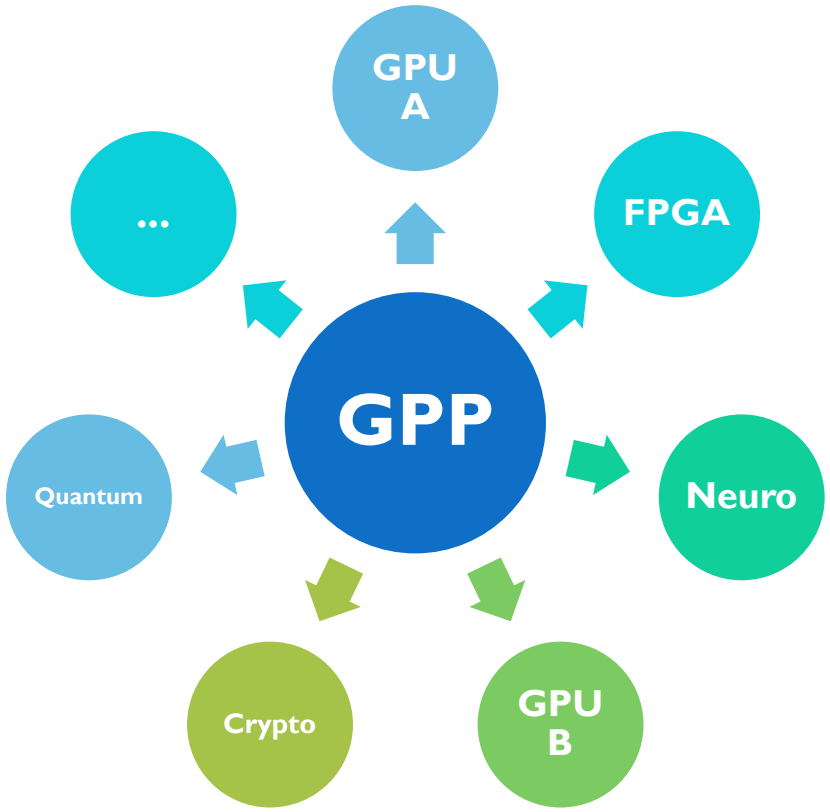

The race to FLOPS is now in the accelerators area

TAKEWAY #4
SPECIFICATION FOR EXASCALE GENERAL PURPOSE PROCESSORS (PART 1)

- Need extreme flexibility and performances on external links
  - 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

EXASCALE CLASS GENERAL PURPOSE PROCESSORS UNDER THE HOOD

THE R_PEAK CHALLENGE – CPU VS. GPU

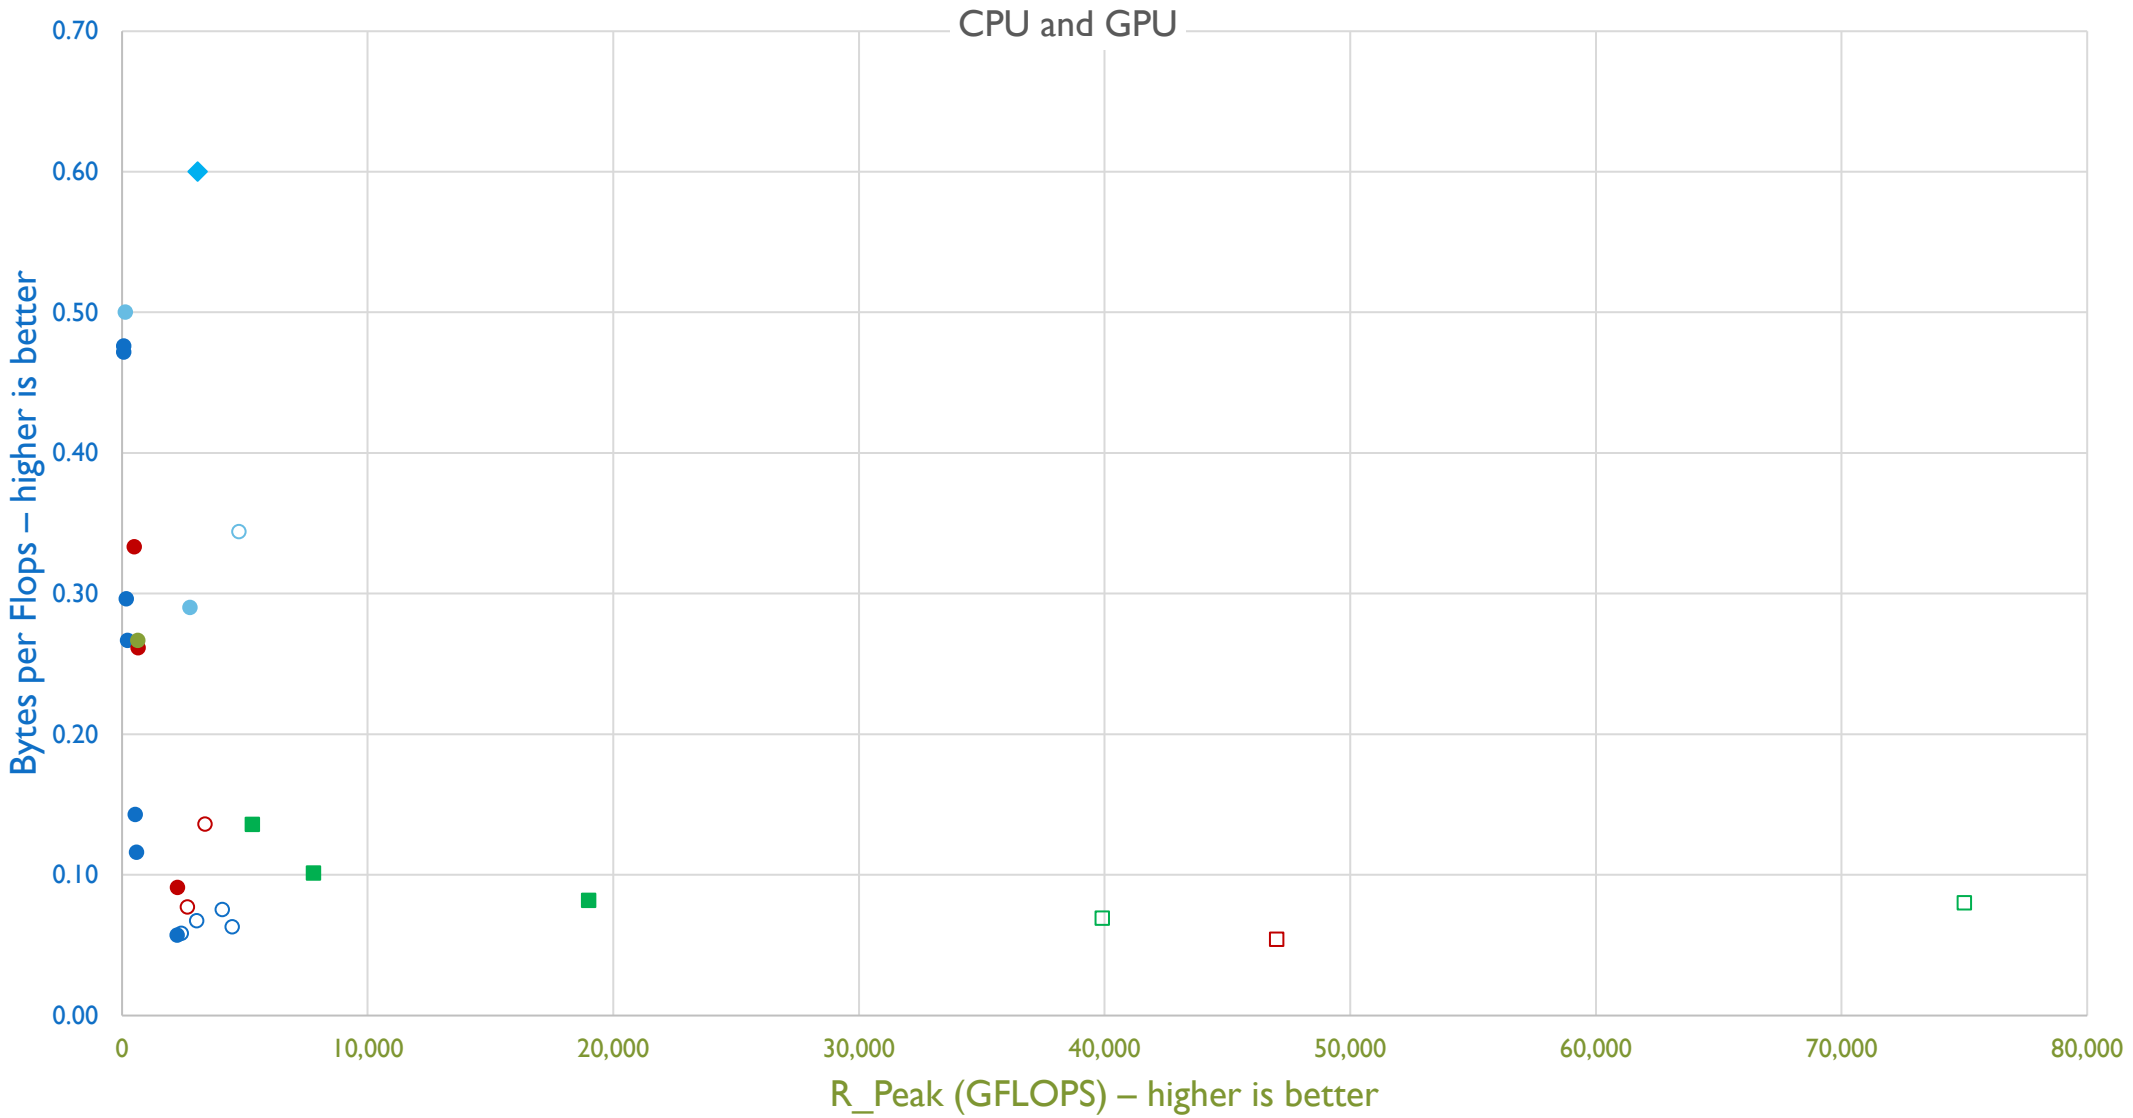

THE R_PEAK CHALLENGE – CPU VS. GPU

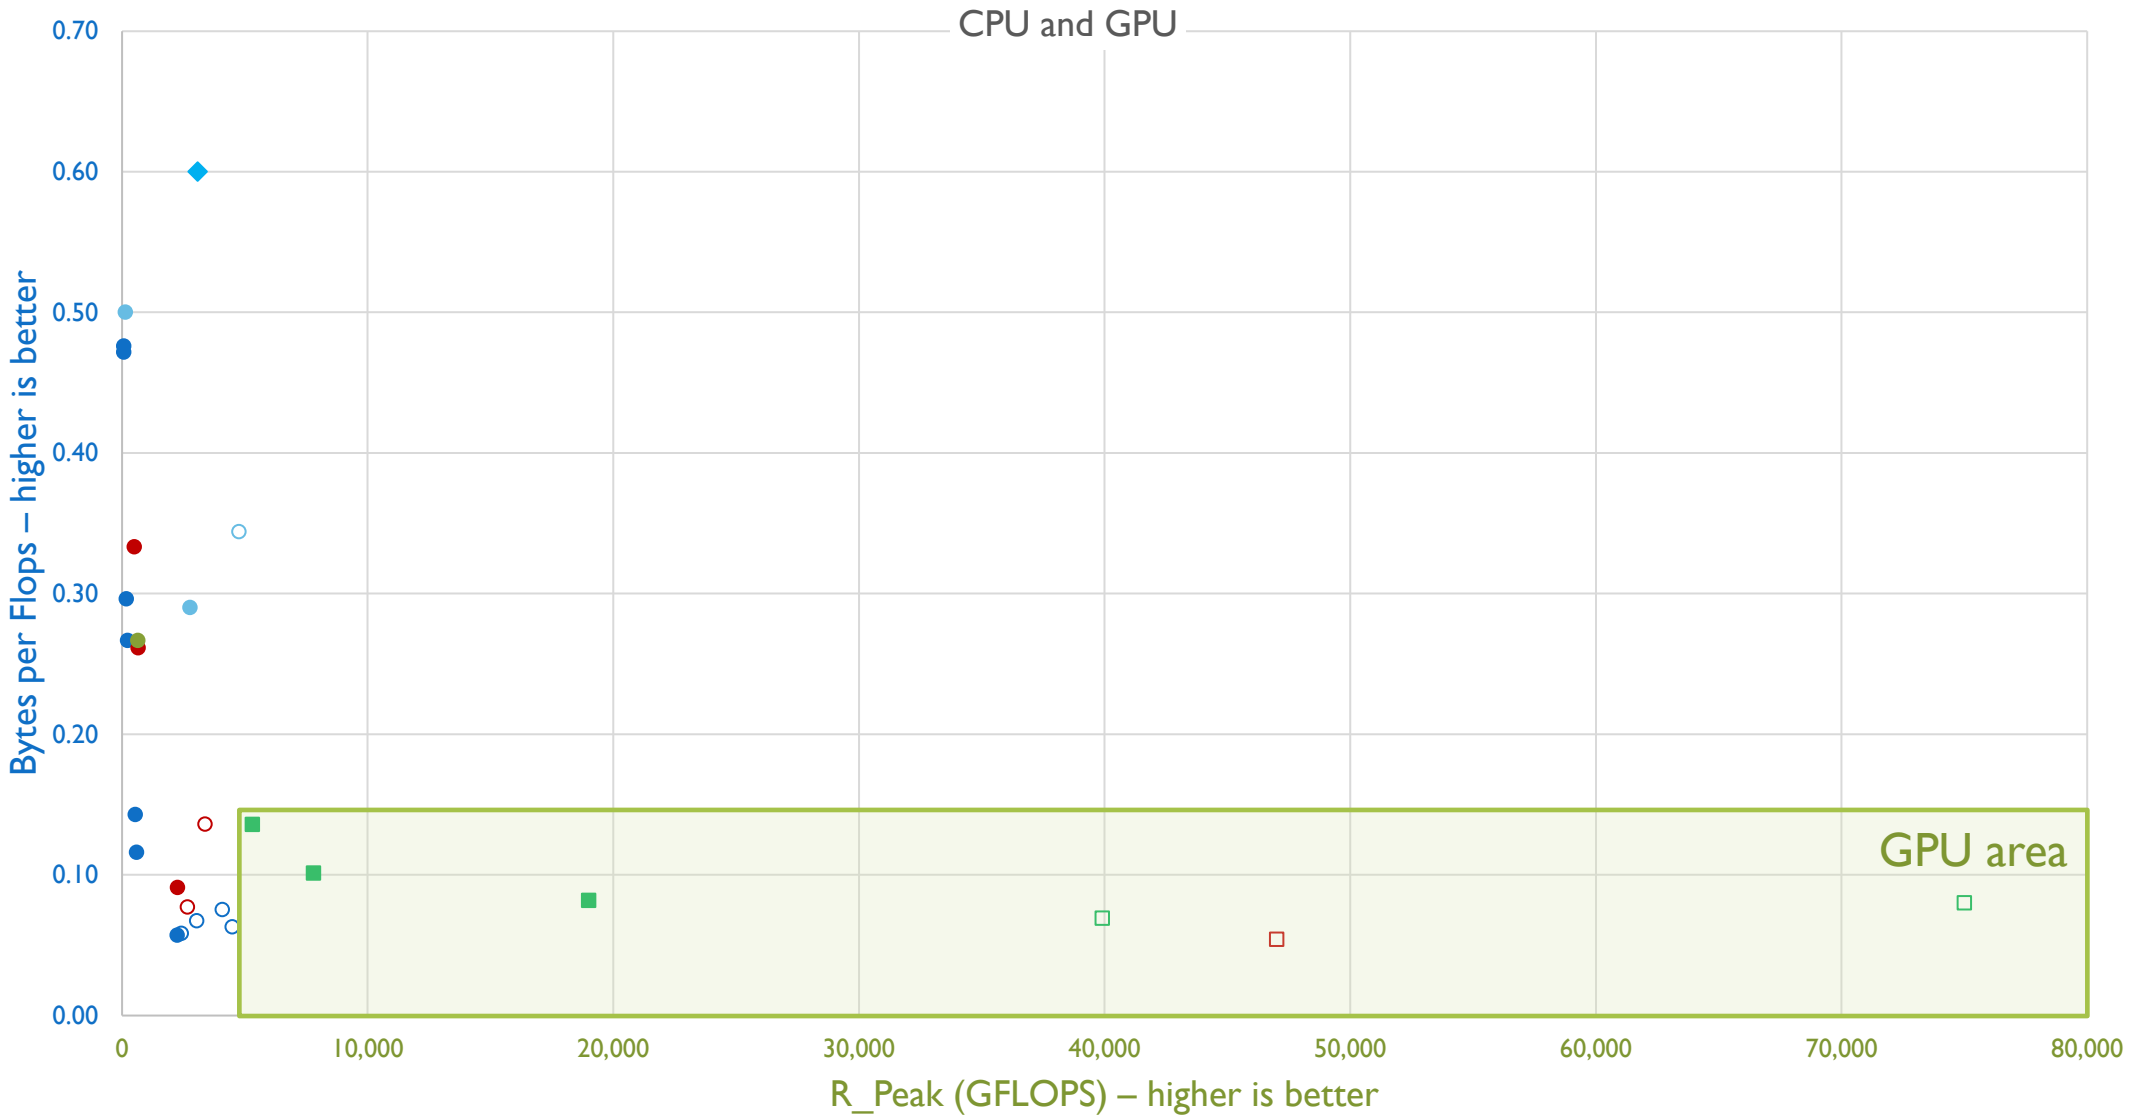

THE R_PEAK CHALLENGE – CPU VS. GPU

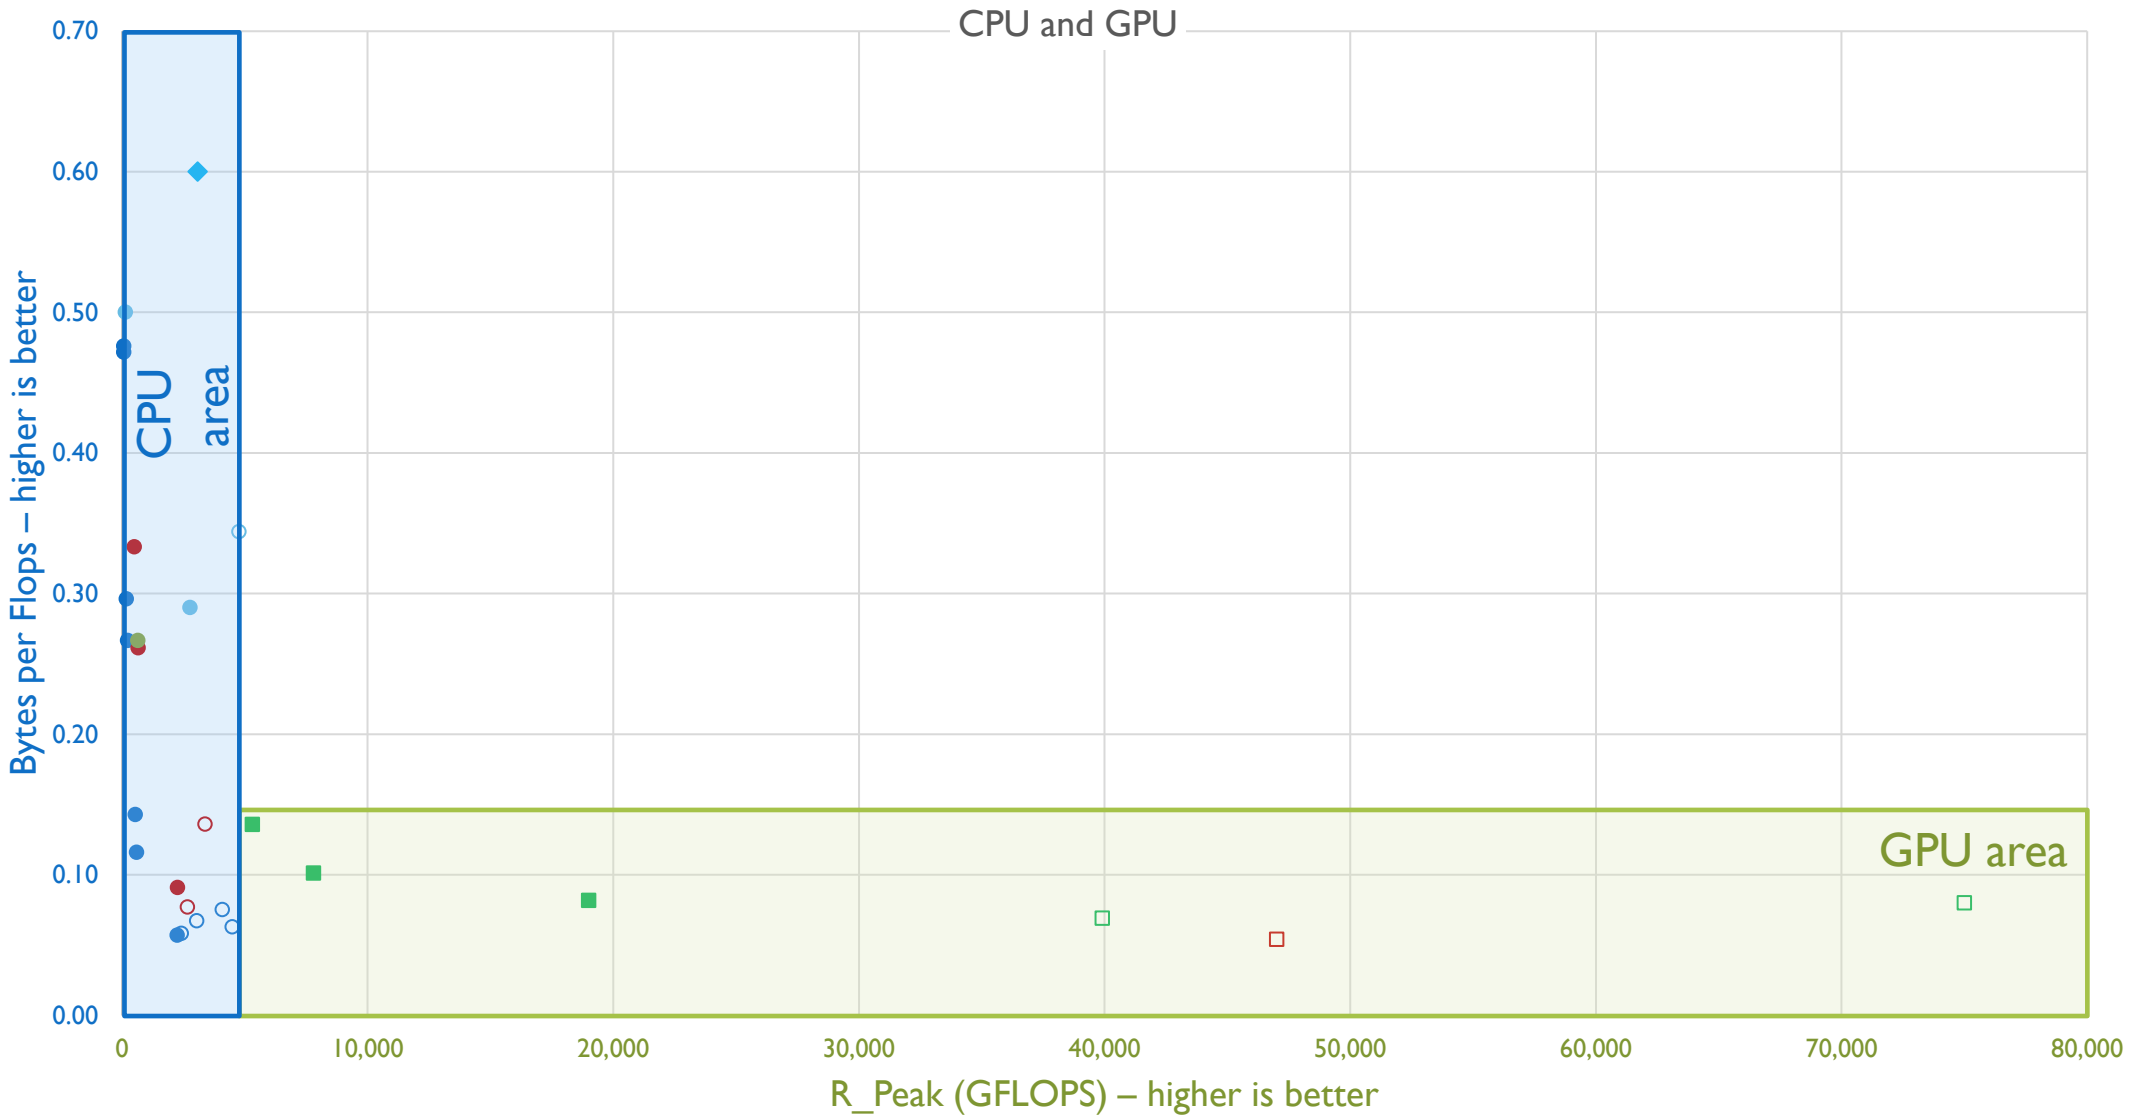

THE R_PEAK CHALLENGE – CPU VS. GPU

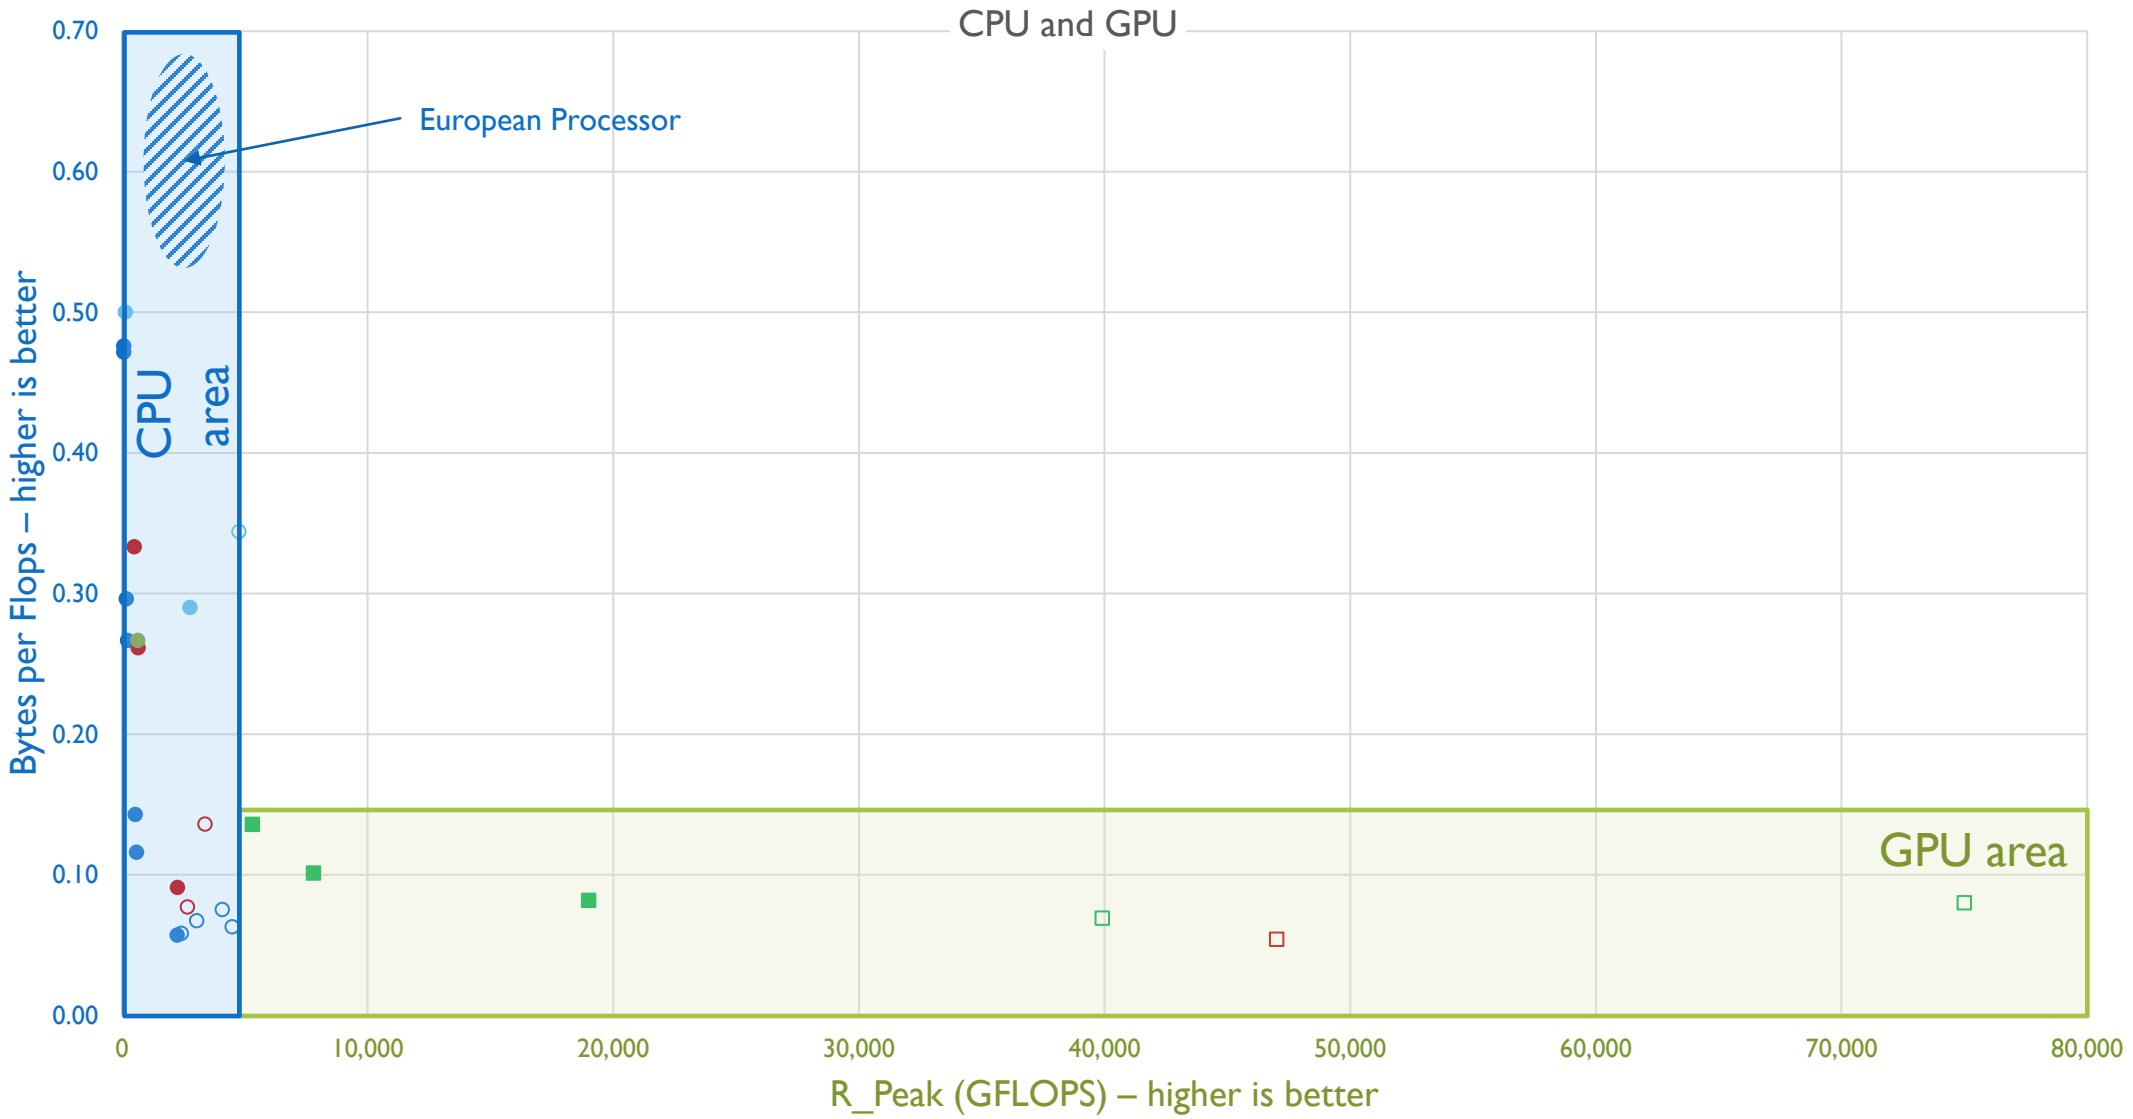

THE R_PEAK CHALLENGE – CPU VS. GPU

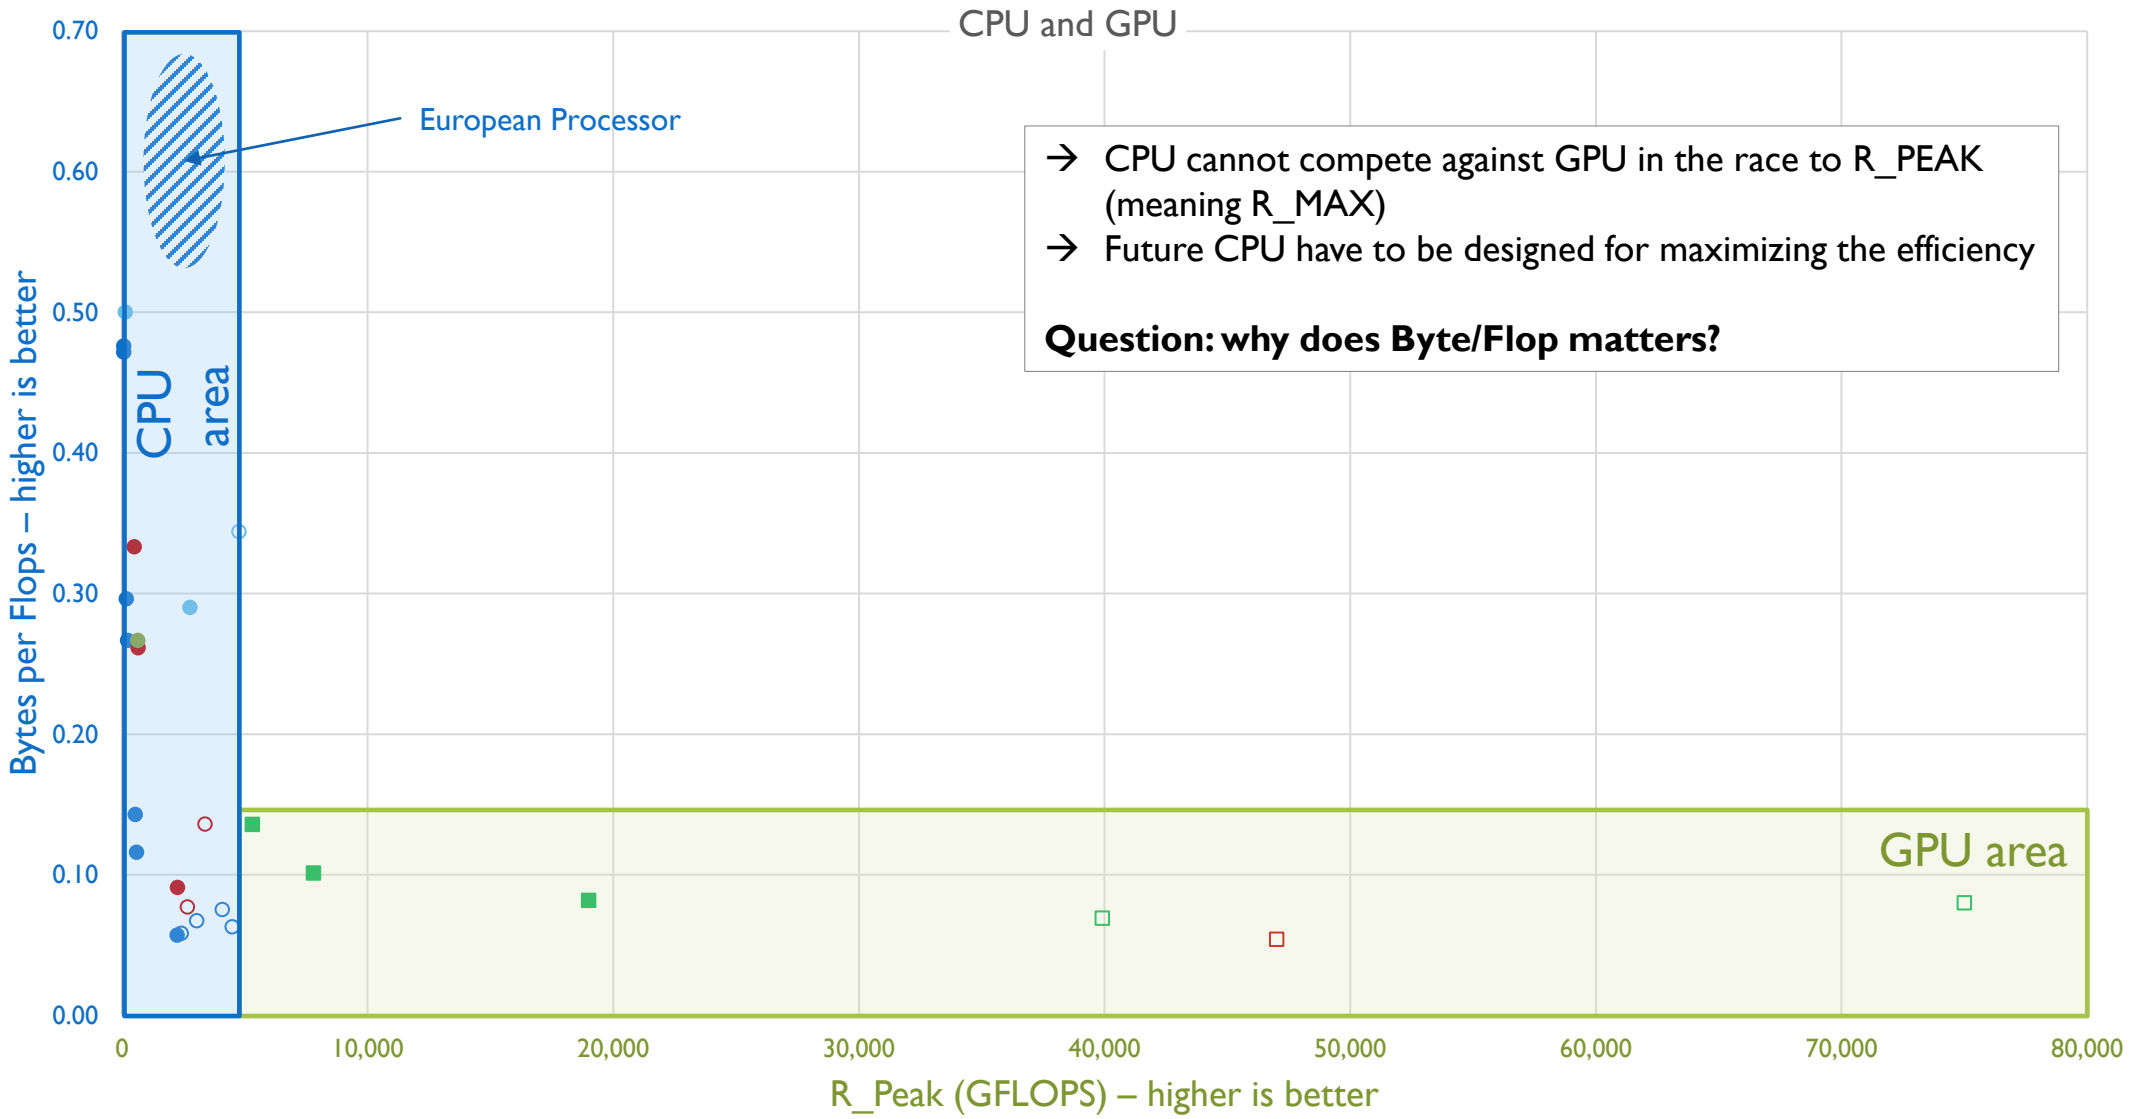

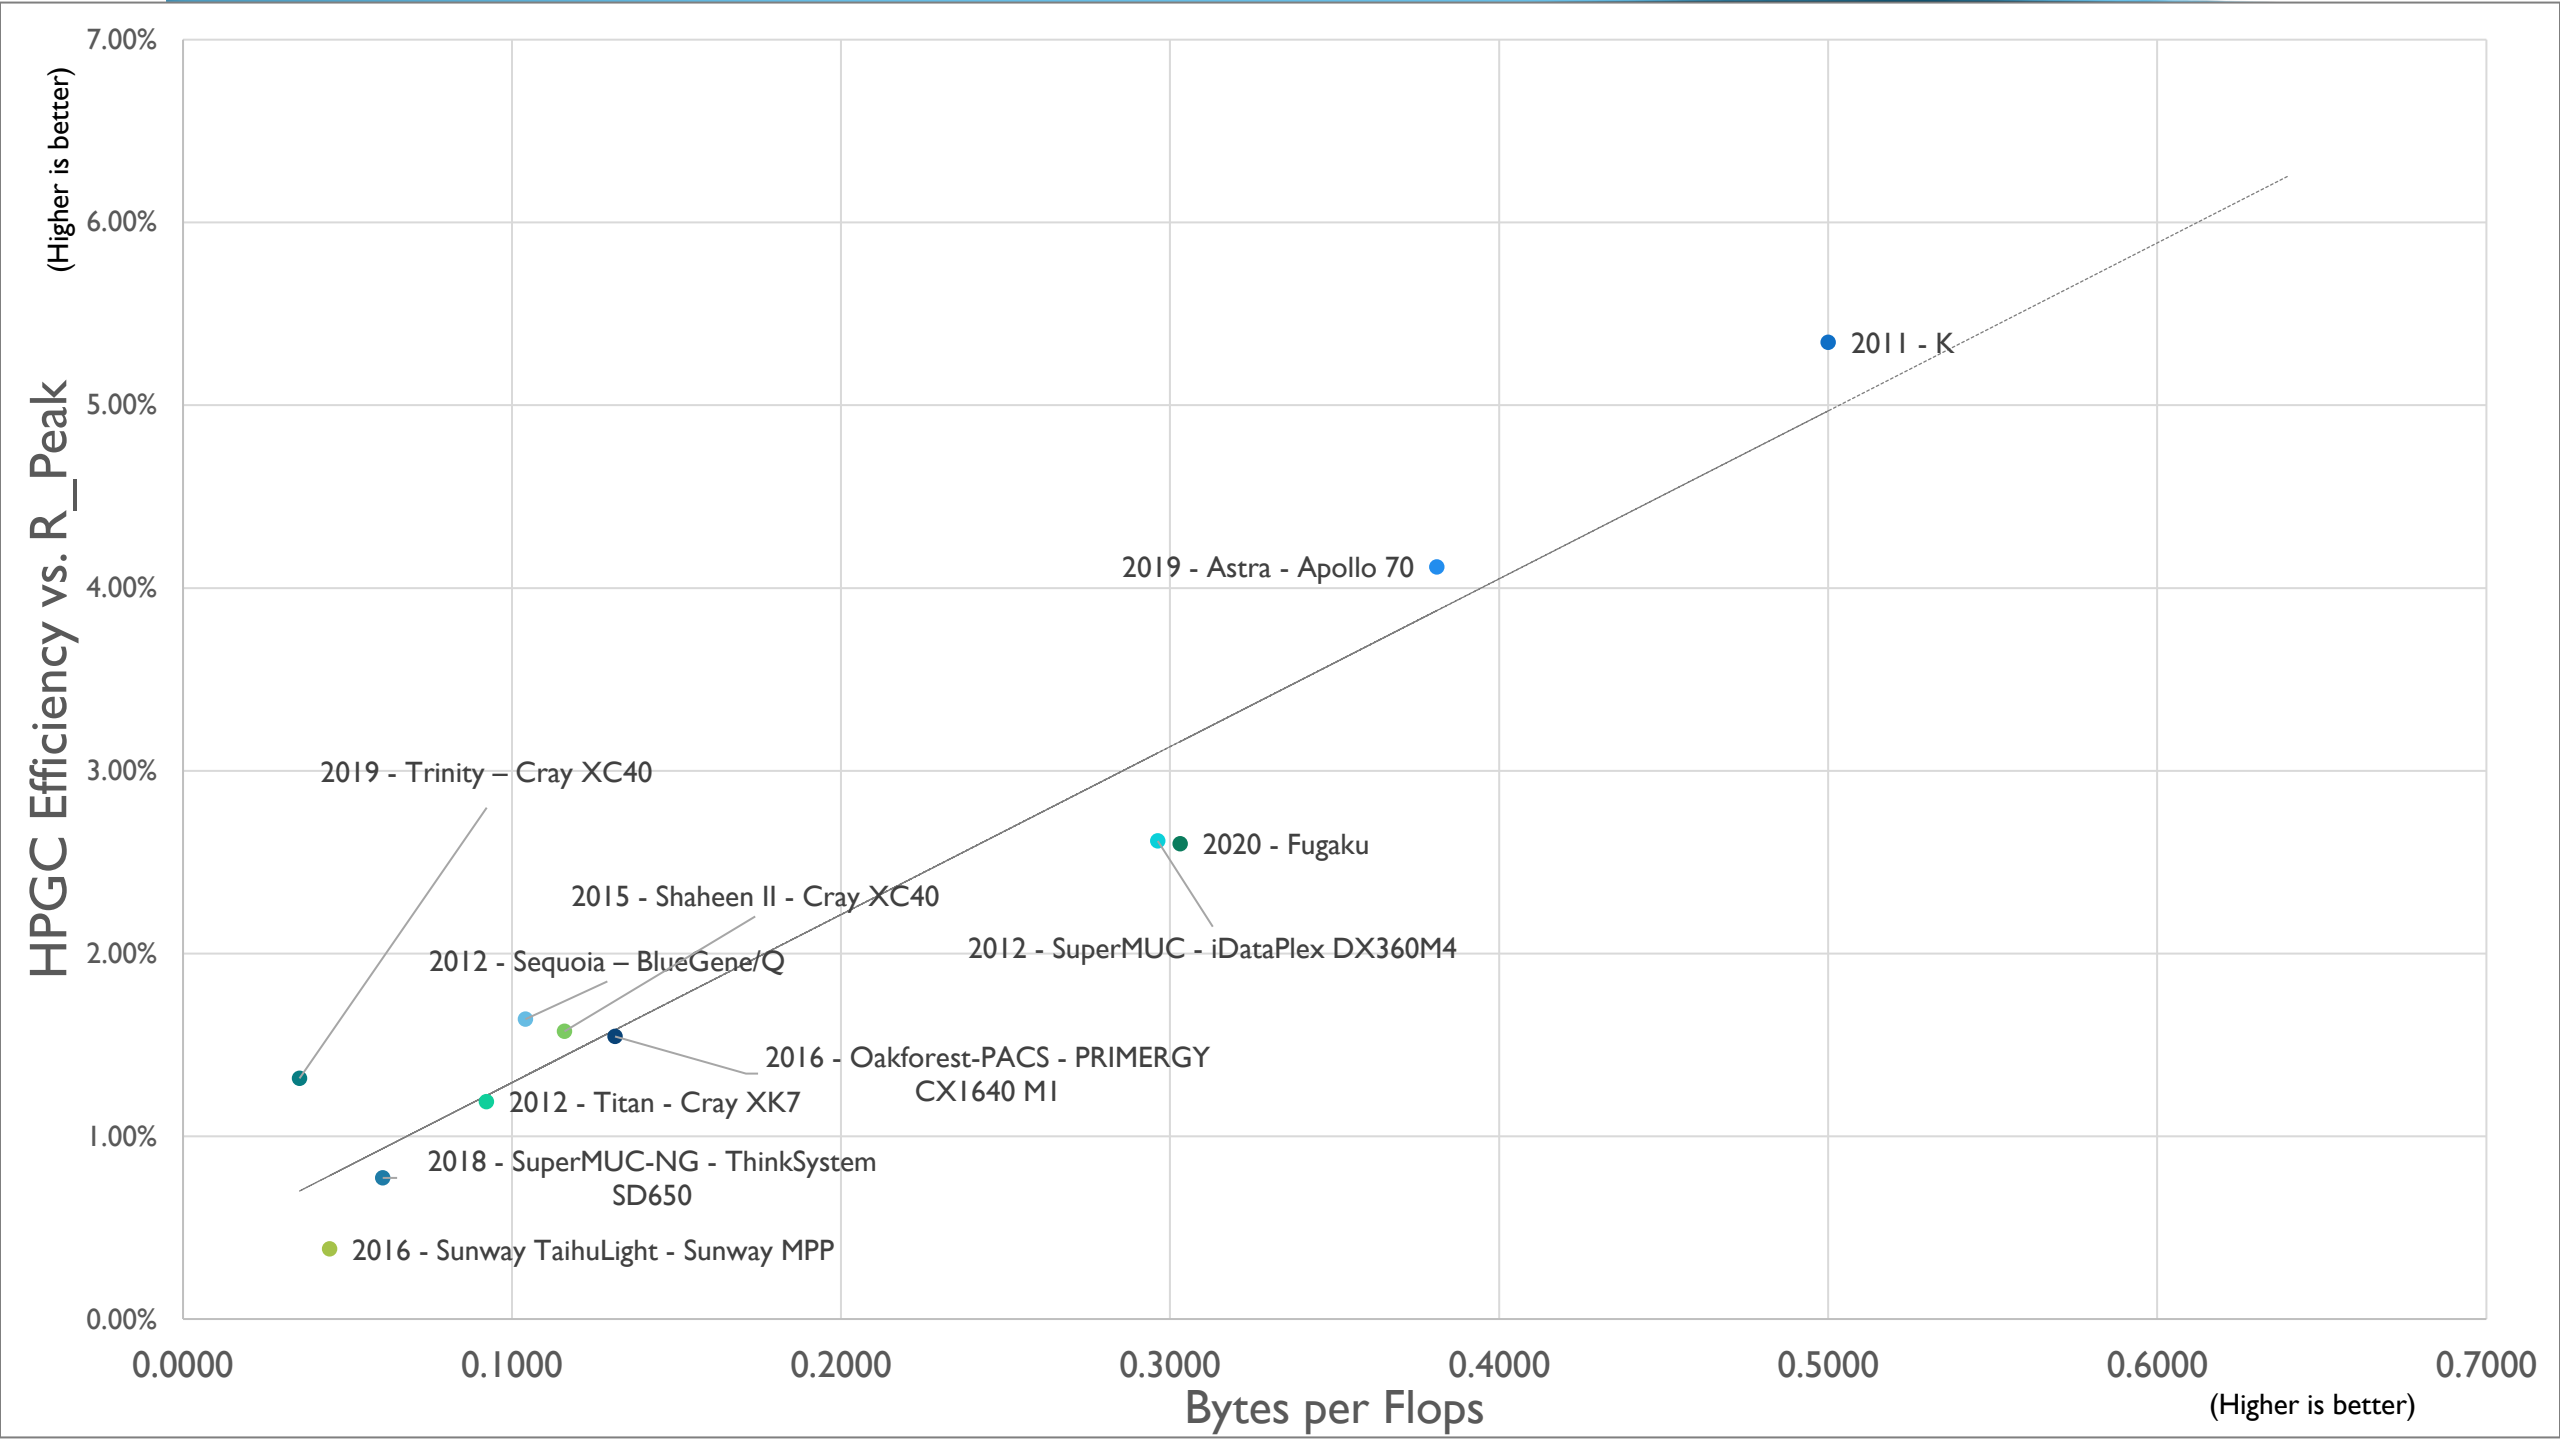
HPGC Efficiency vs. R_Peak
(Higher is better)
7.00%
6.00%
5.00%
4.00%
3.00%
2.00%
1.00%
0.00%
0.0000
0.1000
0.2000
0.3000
0.4000
0.5000
0.6000
0.7000
Bytes per Flops
(Higher is better)
2011 - K
2019 - Astra - Apollo 70
2020 - Fugaku
2012 - SuperMUC - iDataPlex DX360M4
2019 - Trinity – Cray XC40
2015 - 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

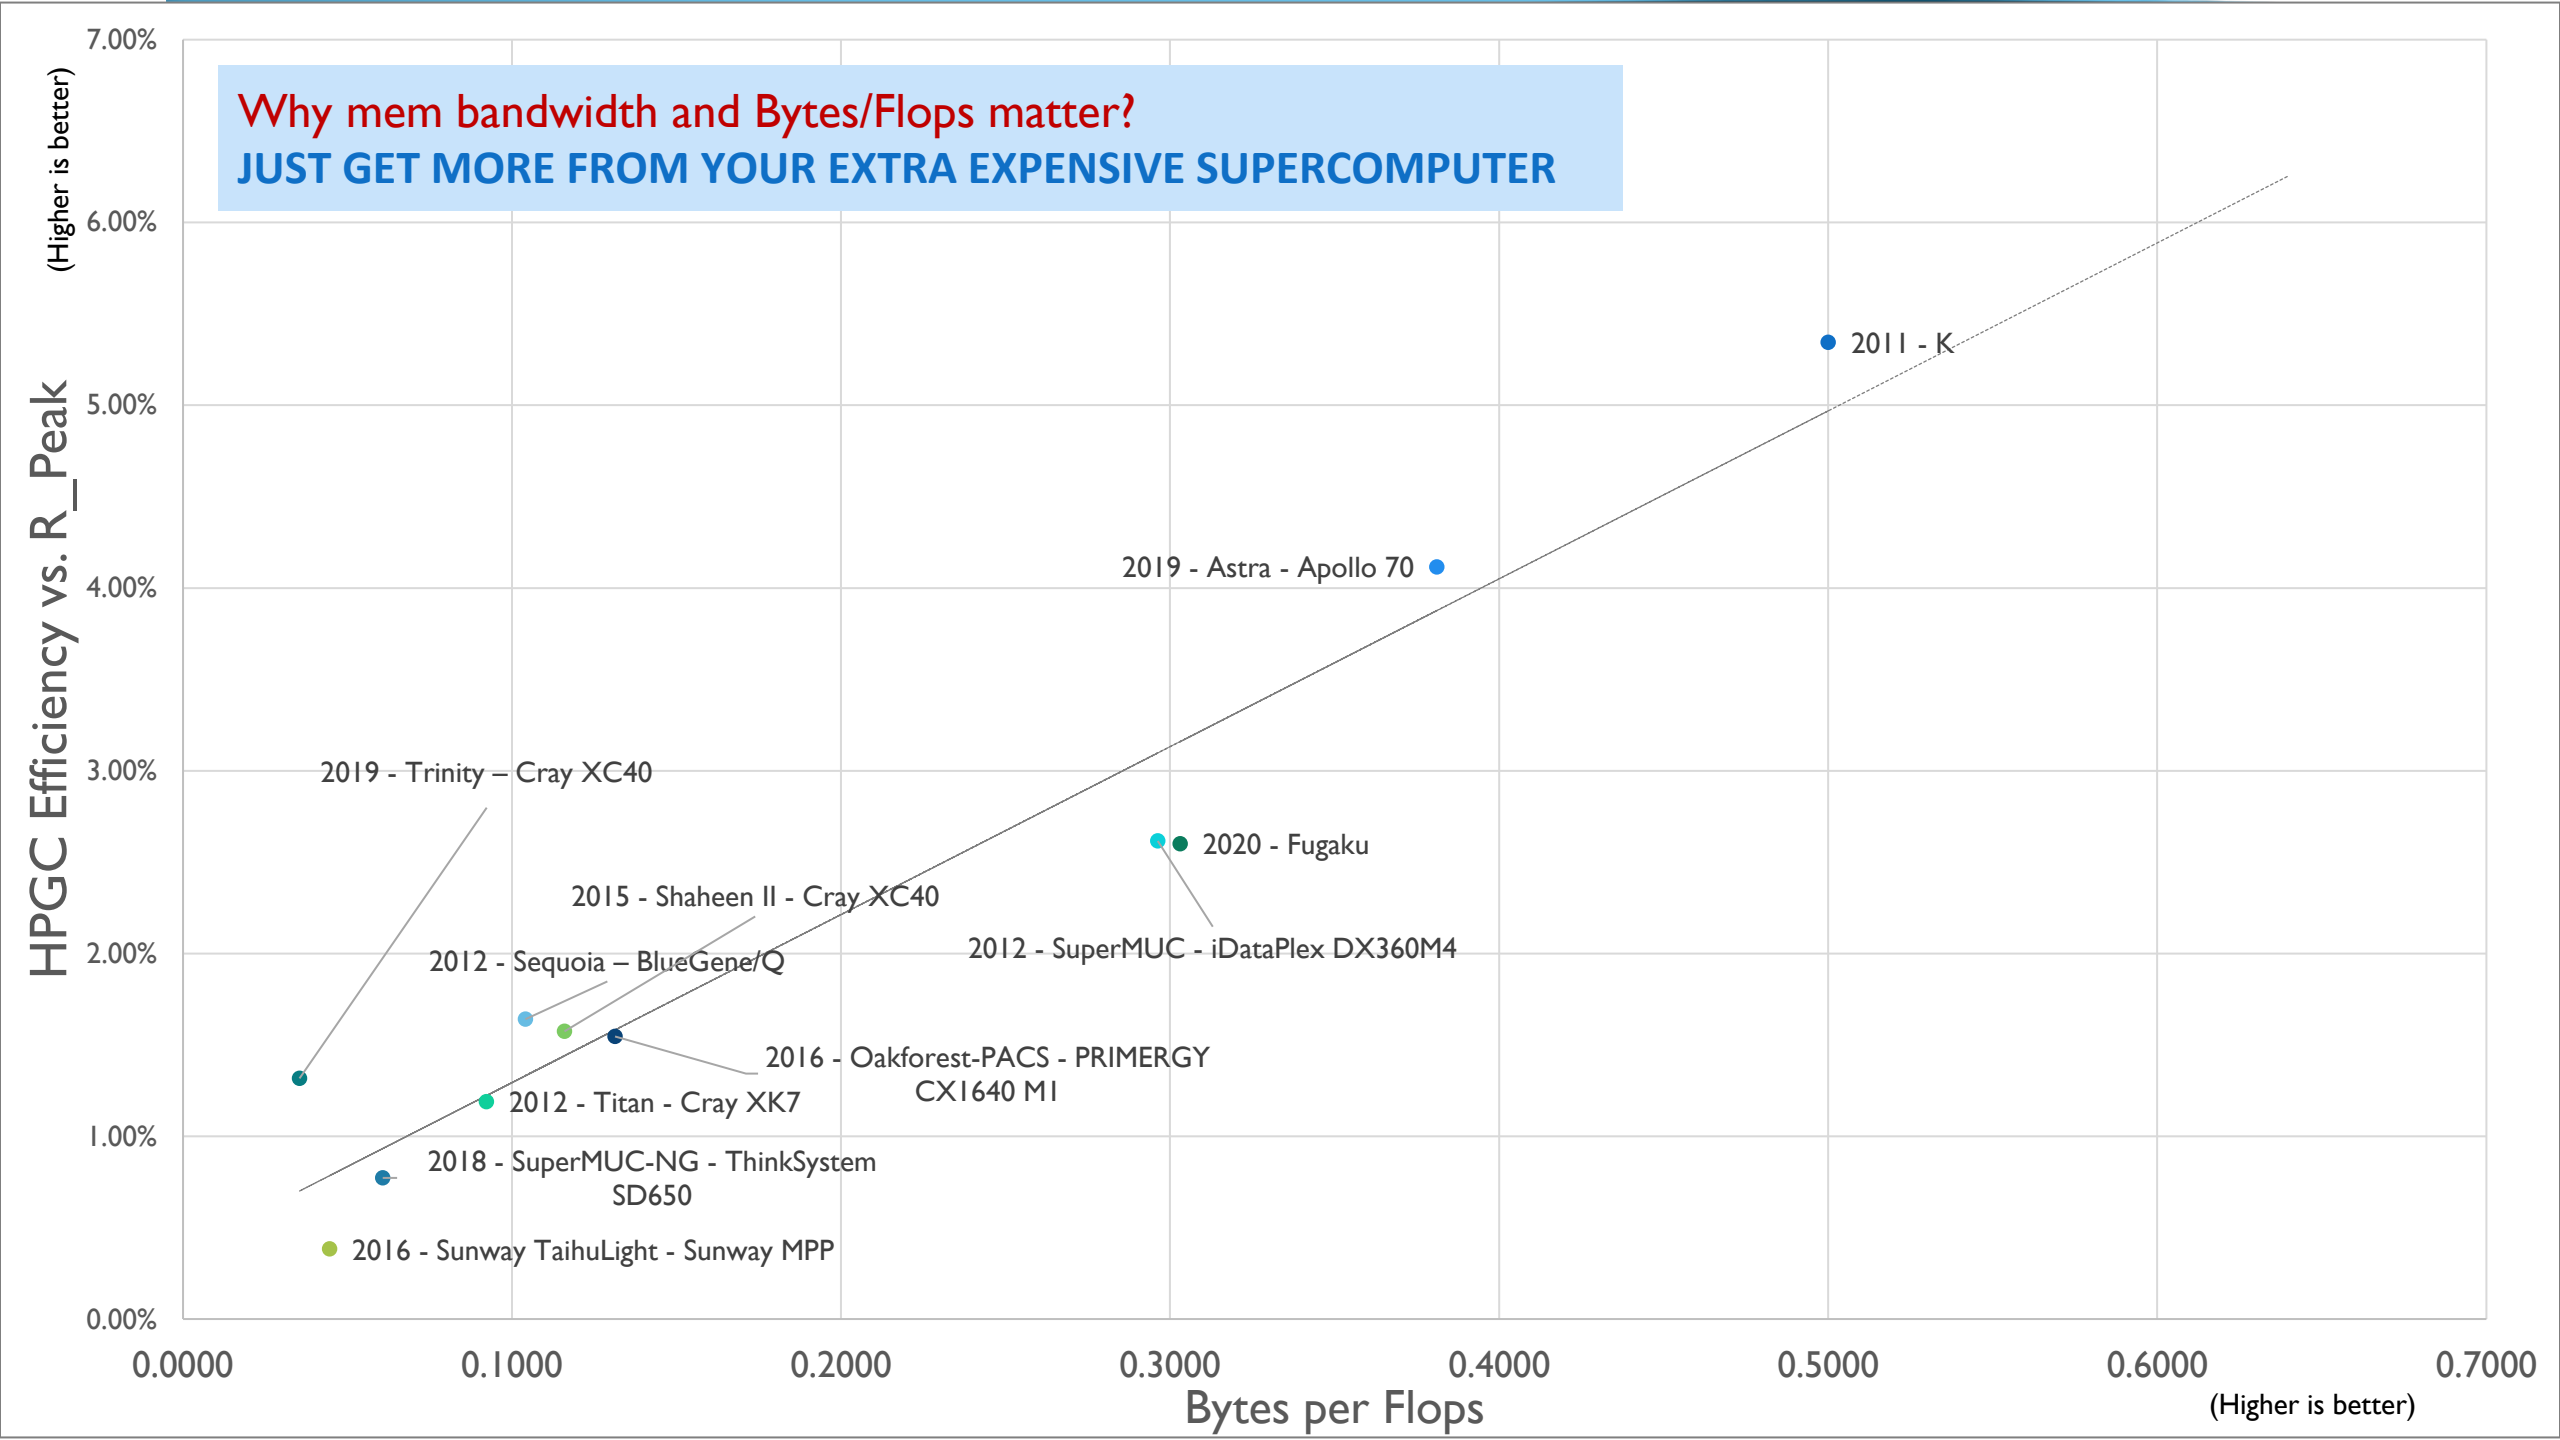
Why mem bandwidth and Bytes/Flops matter?
JUST GET MORE FROM YOUR EXTRA EXPENSIVE SUPERCOMPUTER
HPGC Efficiency vs. R_Peak
(Higher is better)
7.00%
6.00%
5.00%
4.00%
3.00%
2.00%
1.00%
0.00%
0.0000
0.1000
0.2000
0.3000
0.4000
0.5000
0.6000
0.7000
Bytes per Flops
(Higher is better)
2011 - K
2019 - Astra - 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

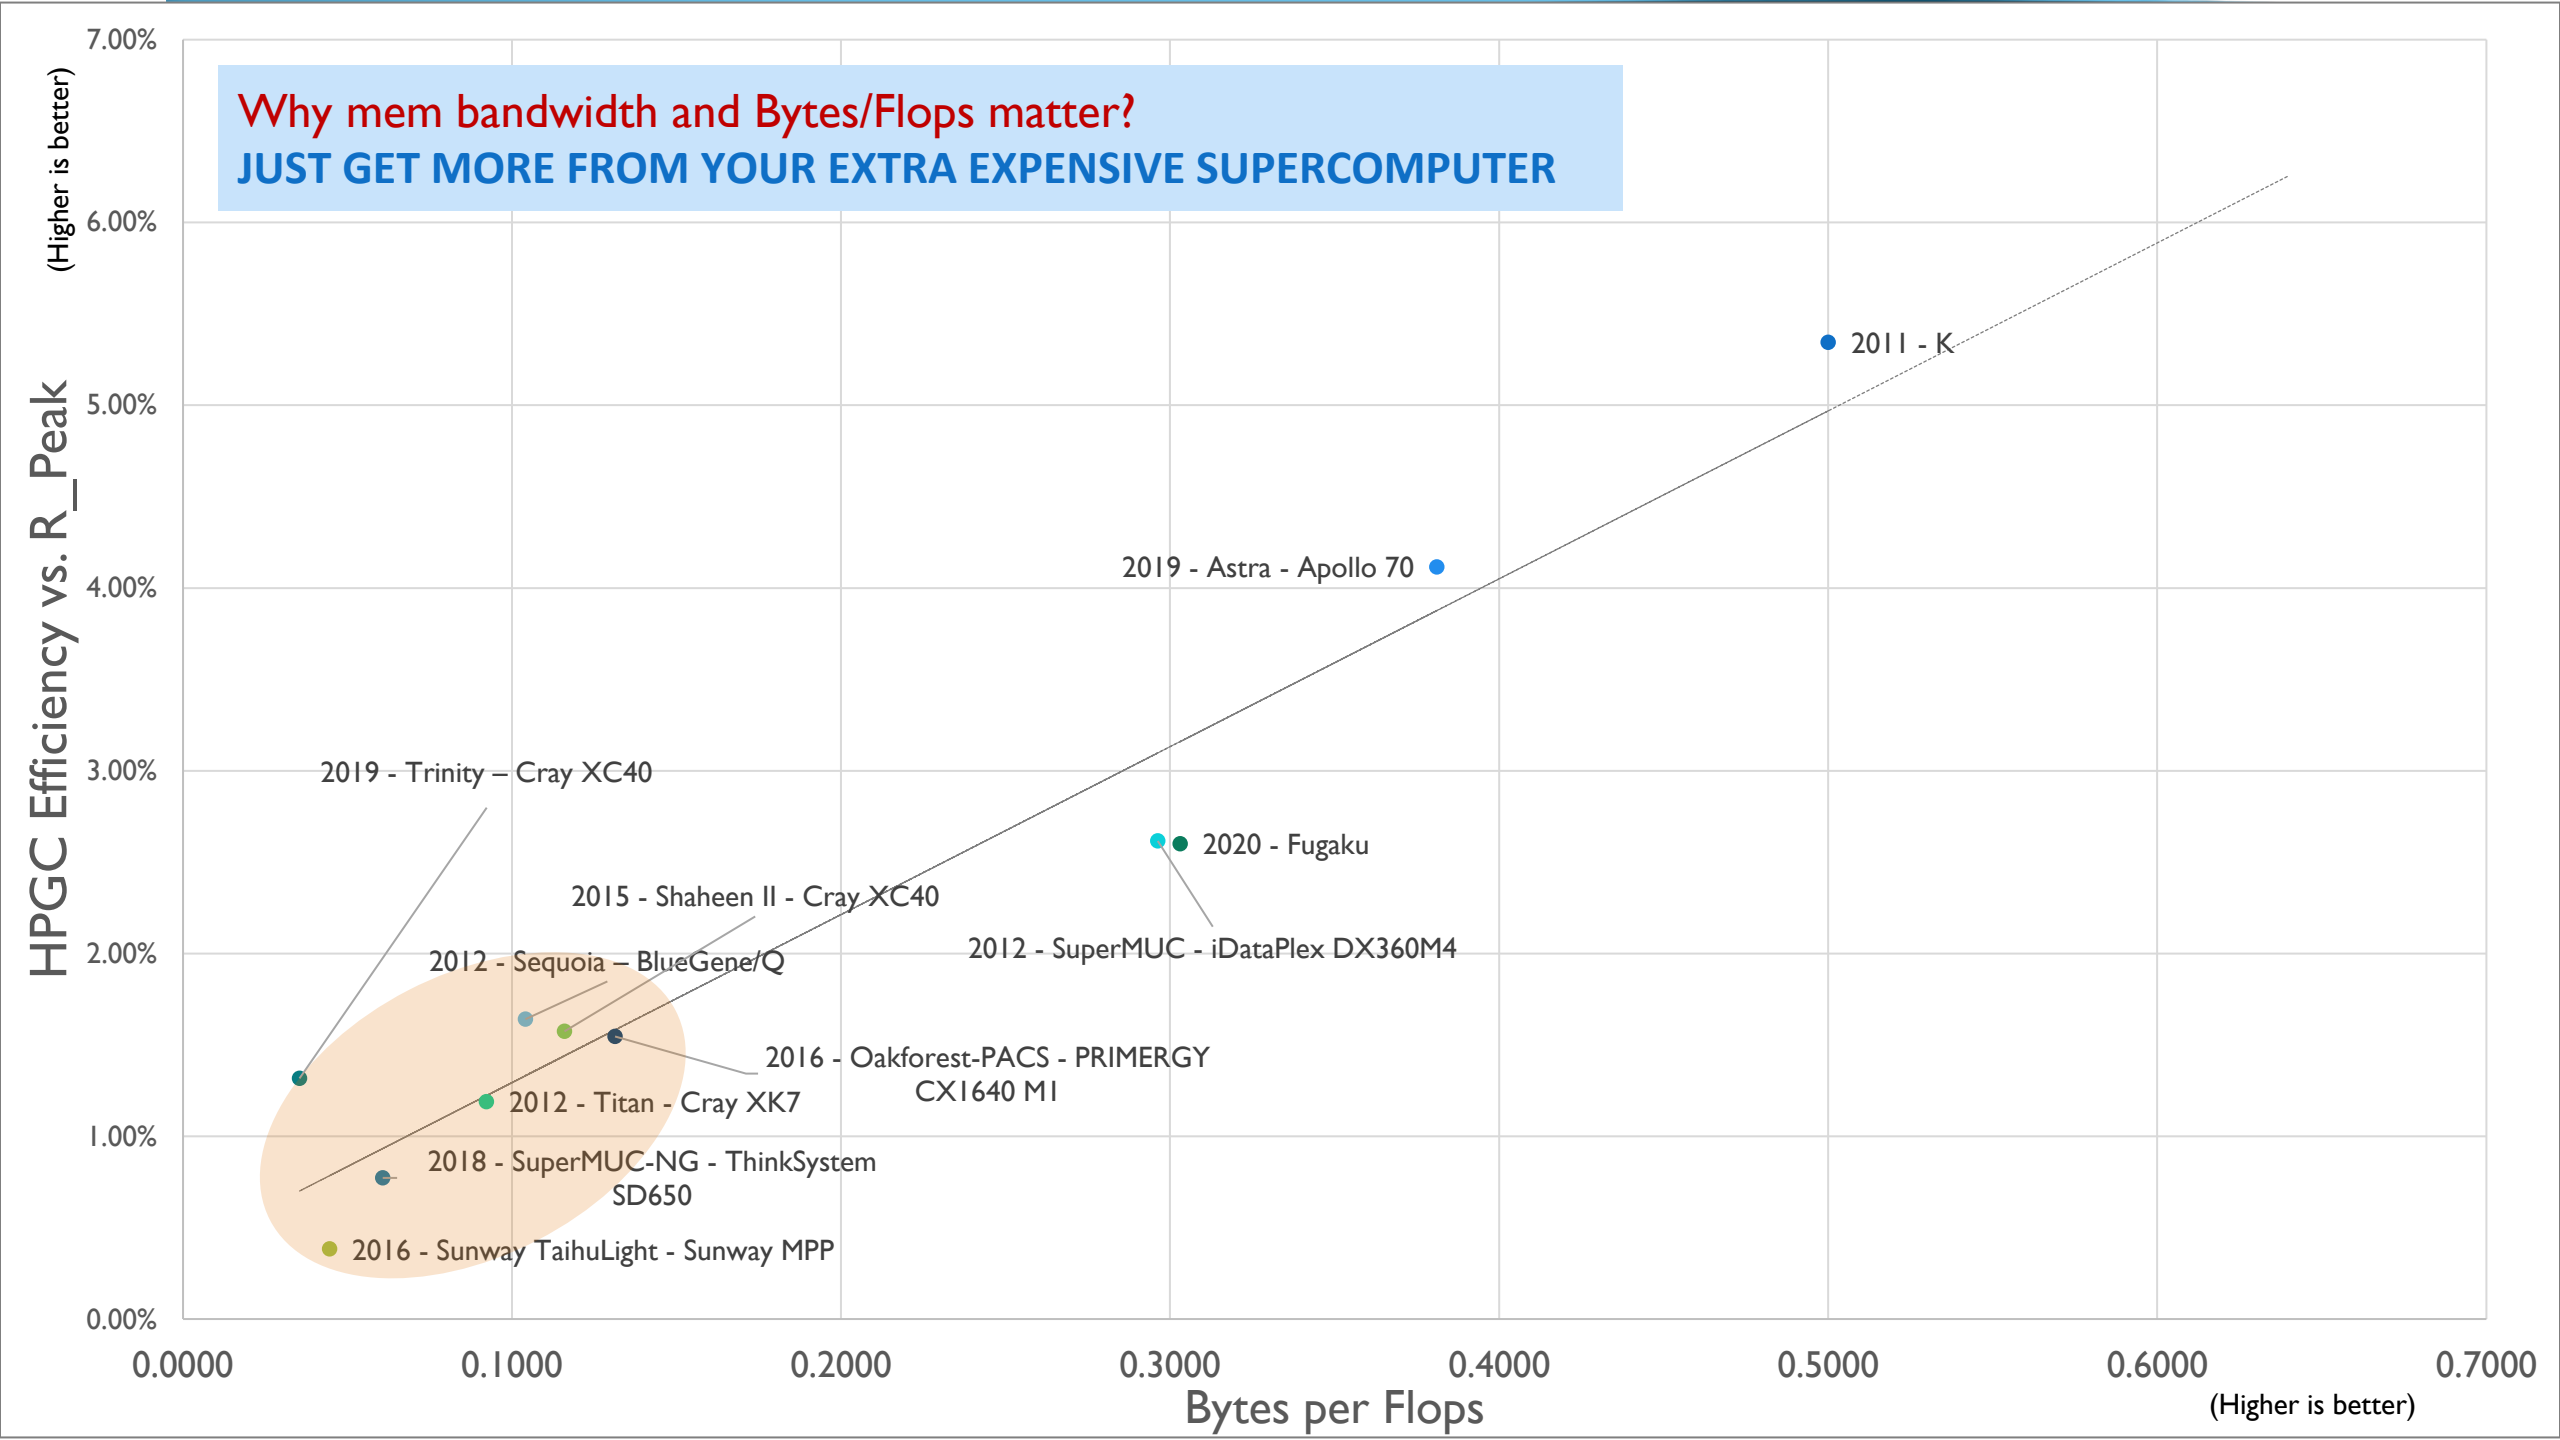
Why mem bandwidth and Bytes/Flops matter?
JUST GET MORE FROM YOUR EXTRA EXPENSIVE SUPERCOMPUTER
HPGC Efficiency vs. R_Peak
(Higher is better)
7.00%
6.00%
5.00%
4.00%
3.00%
2.00%
1.00%
0.00%
0.0000
0.1000
0.2000
0.3000
0.4000
0.5000
0.6000
0.7000
Bytes per Flops
(Higher is better)
2011 - K
2019 - Astra - 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

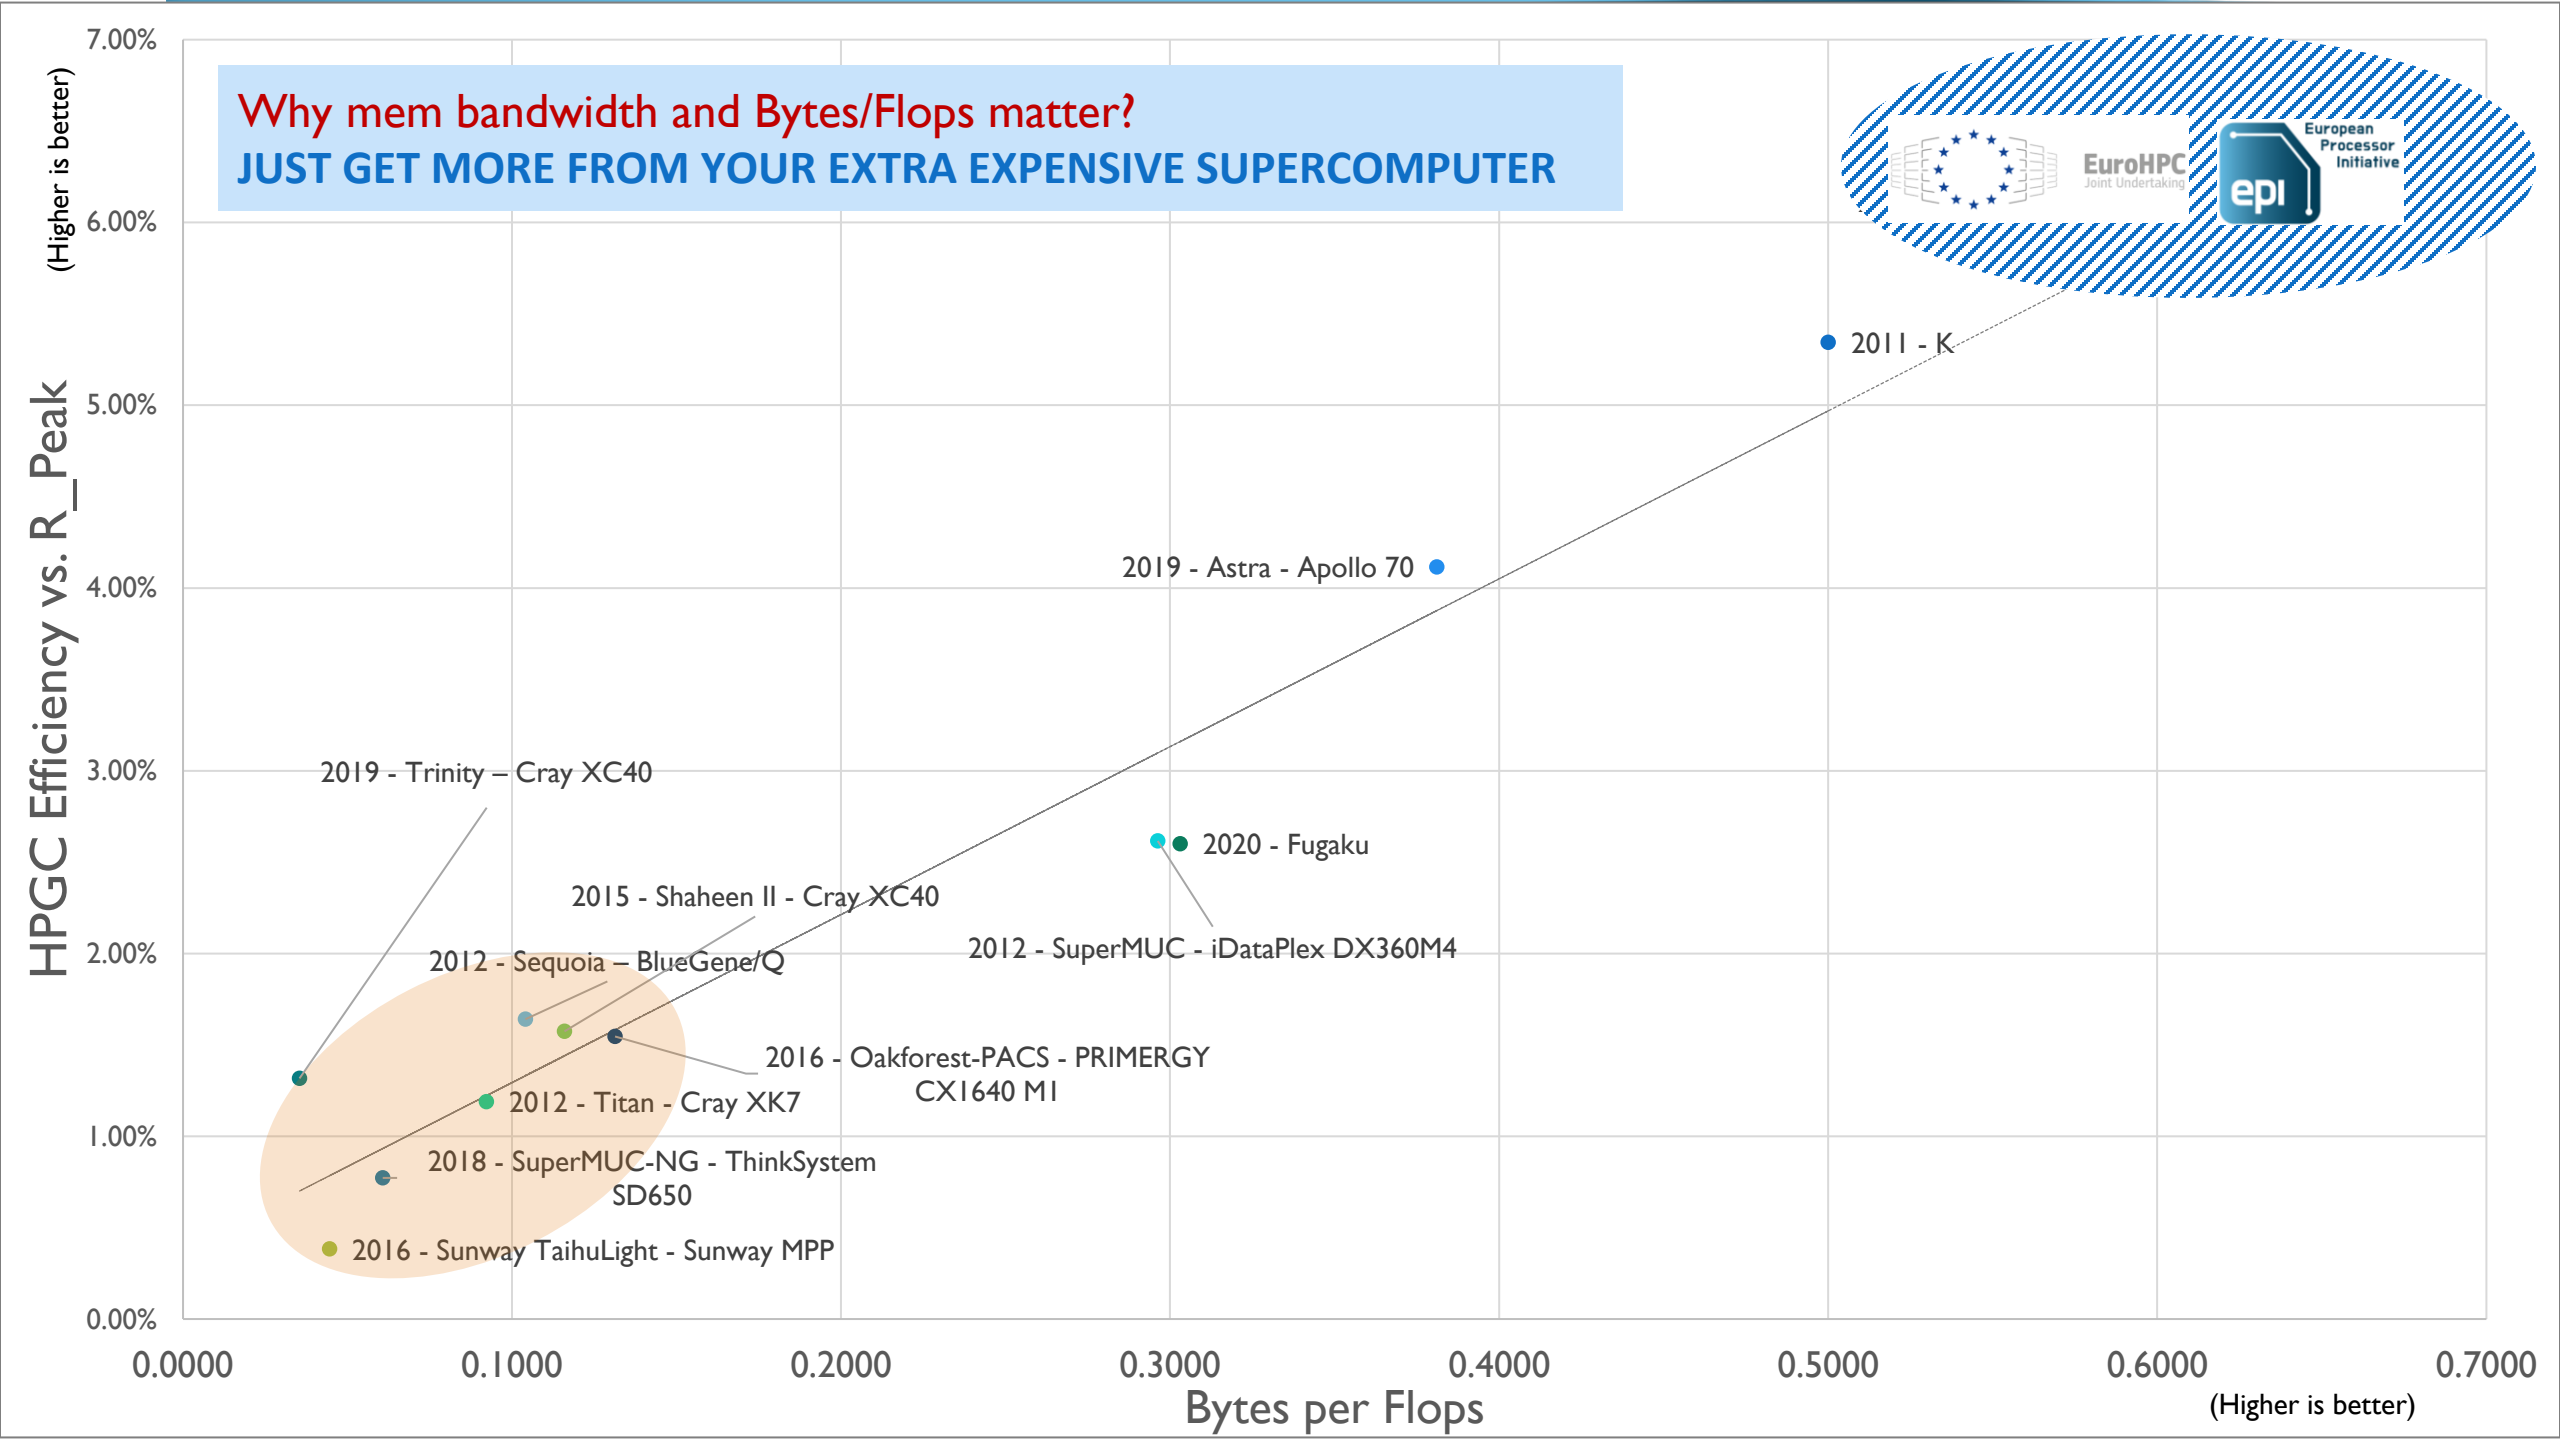
Why mem bandwidth and Bytes/Flops matter?
JUST GET MORE FROM YOUR EXTRA EXPENSIVE SUPERCOMPUTER
EuroHPC Joint Undertaking
European Processor Initiative
epi
HPGC Efficiency vs. R_Peak
(Higher is better)
7.00%
6.00%
5.00%
4.00%
3.00%
2.00%
1.00%
0.00%
2011 - K
2019 - Astra - Apollo 70
2020 - Fugaku
2012 - 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

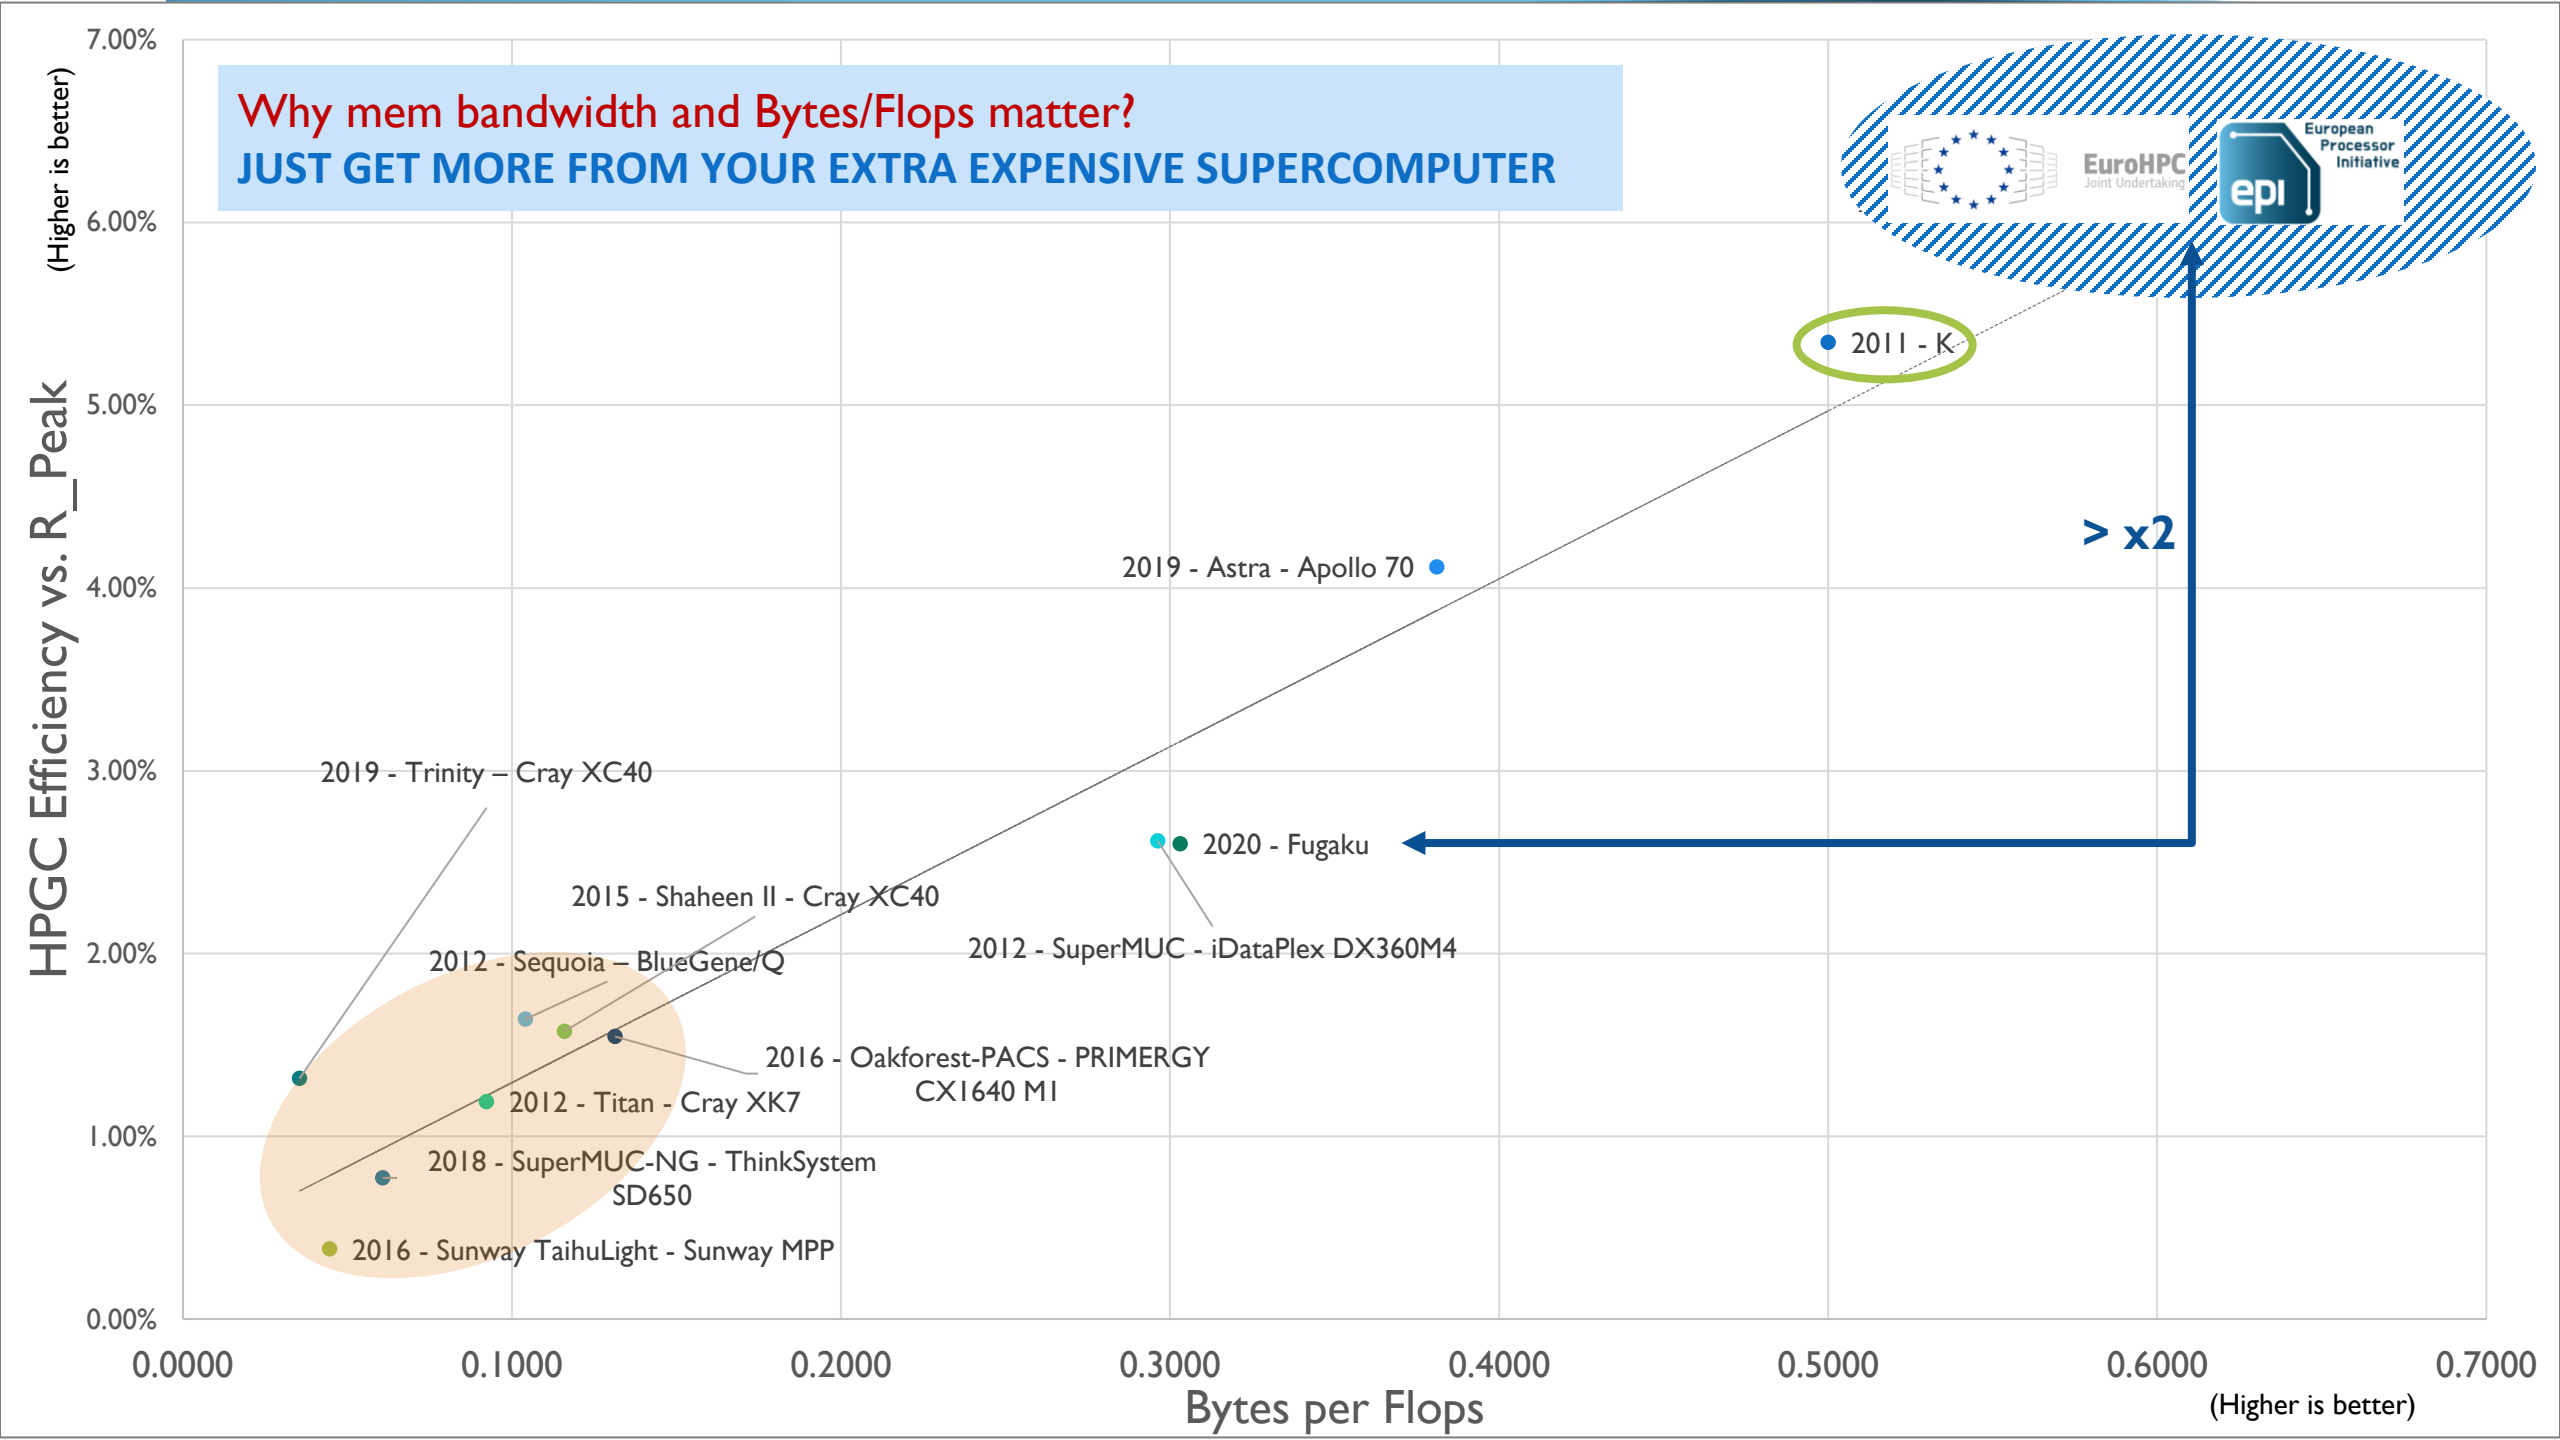

TAKEWAY #5
SPECIFICATION FOR EXASCALE GENERAL PURPOSE PROCESSORS (PART 2)

- No need to compete with specialized devices like GPUs
  ➔ Need “good enough” but excellent FP64 performances

- Improve overall efficiency
  ➔ Need much better byte/Flop ratio than today

SO… WHAT ARE THE OVERALL SPECIFICATIONS FOR THE EUROPEAN PROCESSOR INITIATIVE?

- ARM Instruction set (IoT to HPC)
- Strong memory subsystem with HBM and DDR
- Extra large PCIe subsystem (maximize PCIe lanes)
- HPC standard interface (CXL)
- Maximum interoperability with accelerators
- Focus on ease of use
  - +++ byte per flop
  - +++ open source SW tools

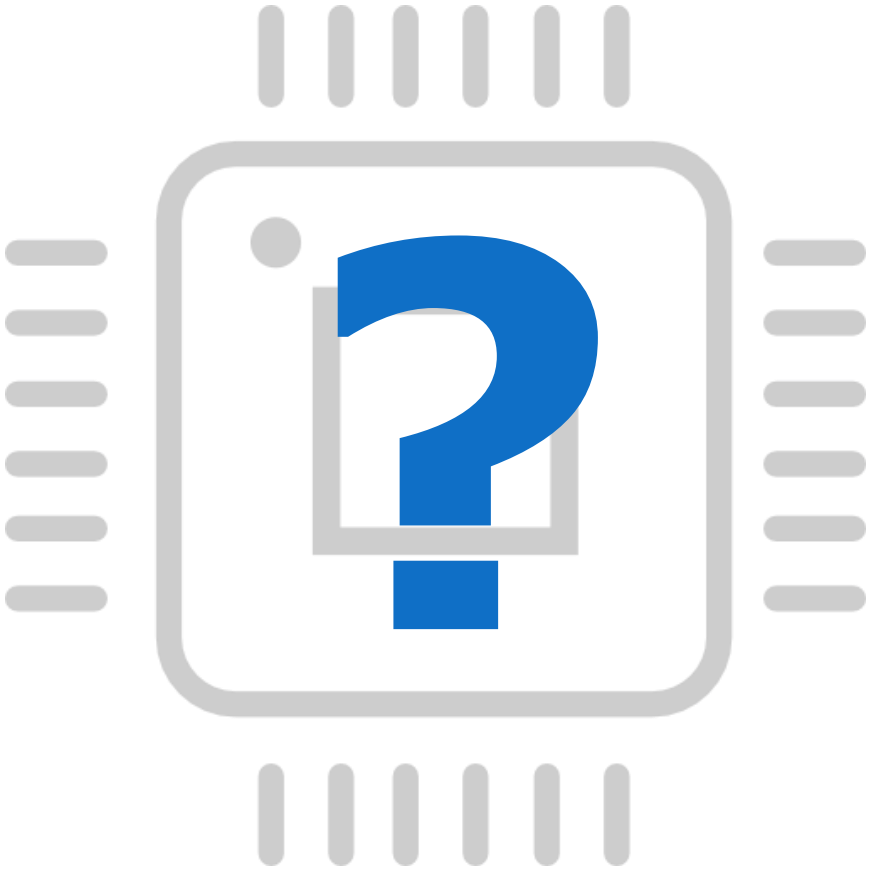

GENERAL PURPOSE PROCESSOR (GPP) AND COMMON OPEN ARCHITECTURE

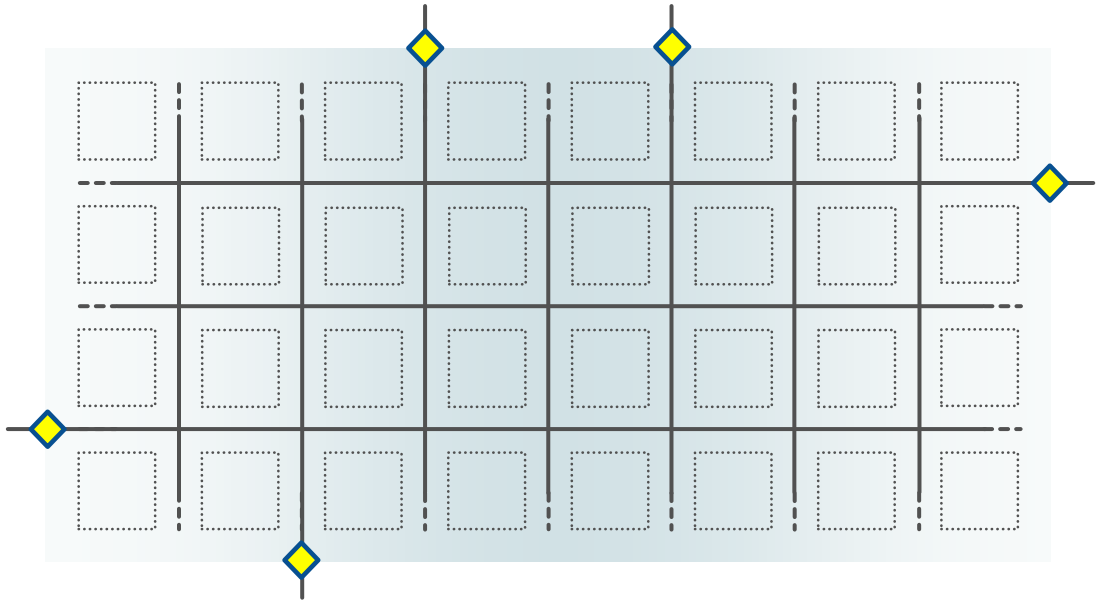

GENERAL PURPOSE PROCESSOR (GPP) AND COMMON OPEN ARCHITECTURE

- Multi-Die

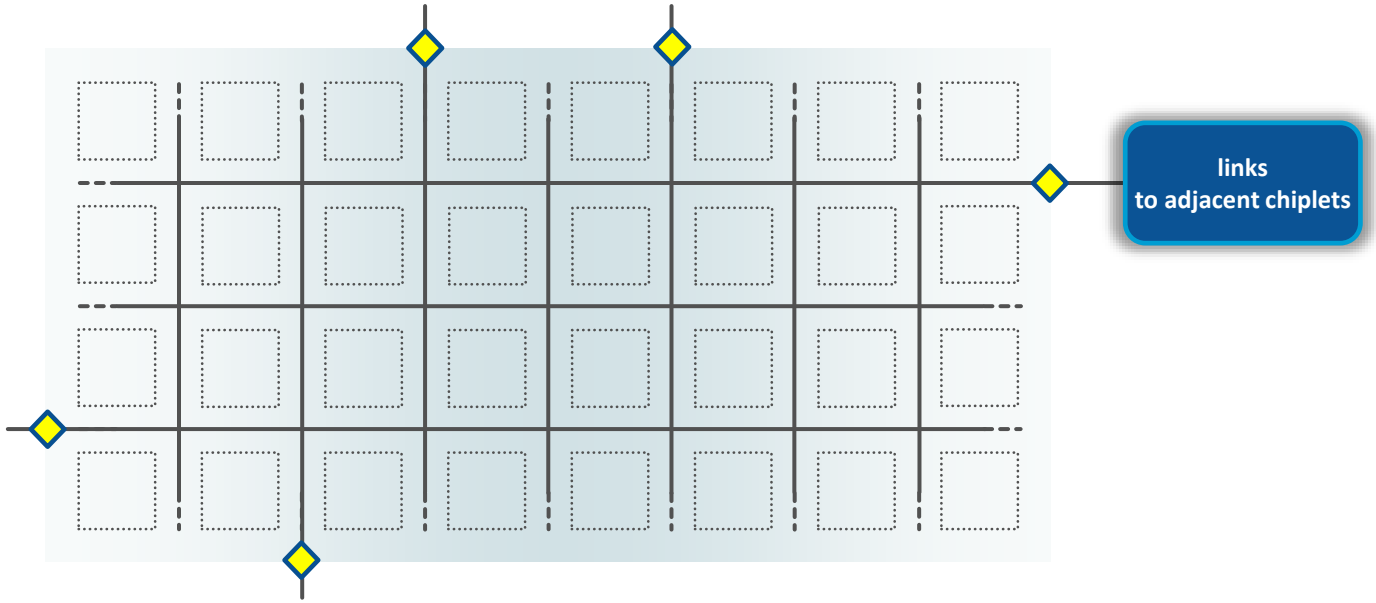

GENERAL PURPOSE PROCESSOR (GPP) AND COMMON OPEN ARCHITECTURE

- Multi-Die
- HBM
- DDR

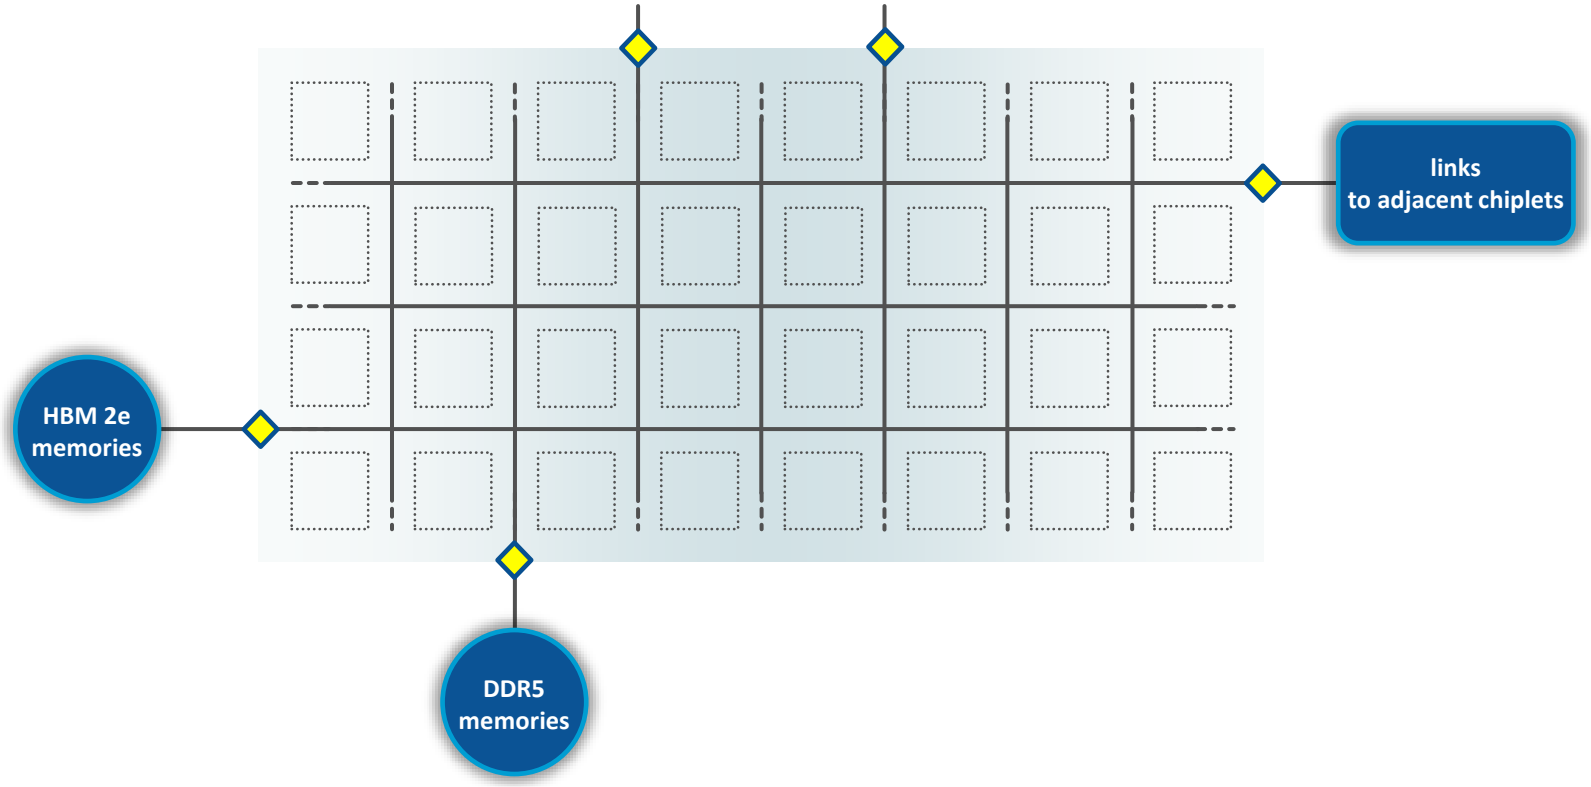

GENERAL PURPOSE PROCESSOR (GPP) AND COMMON OPEN ARCHITECTURE

- Multi-Die
- HBM
- DDR

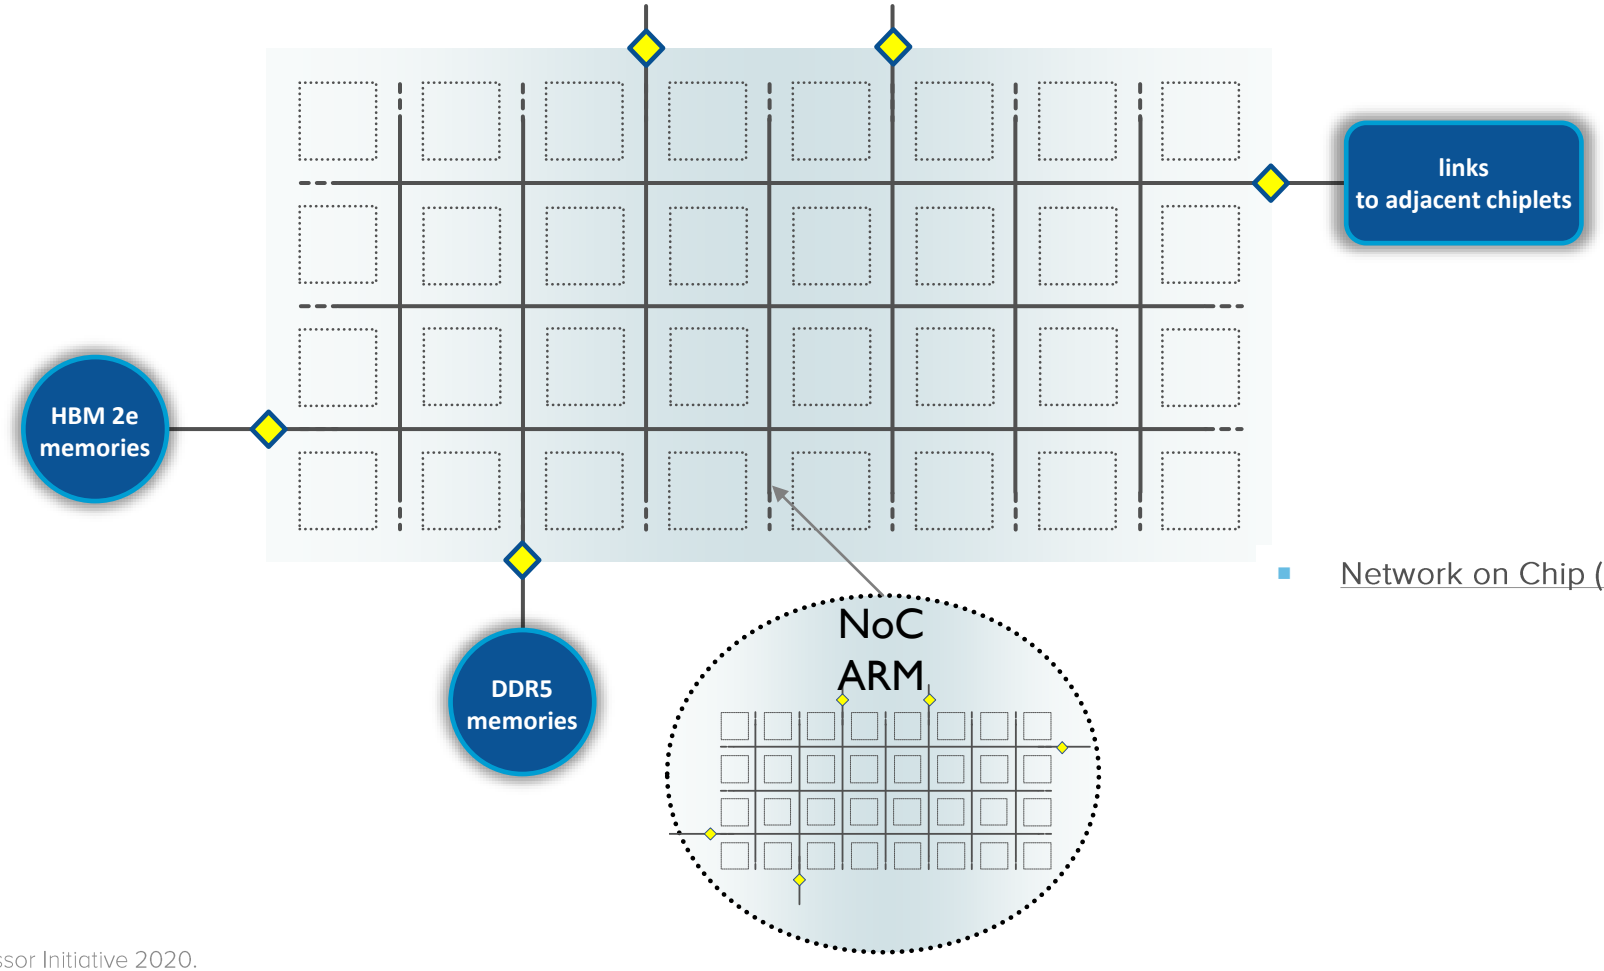

- Network on Chip (NoC) – ARM

GENERAL PURPOSE PROCESSOR (GPP) AND COMMON OPEN ARCHITECTURE

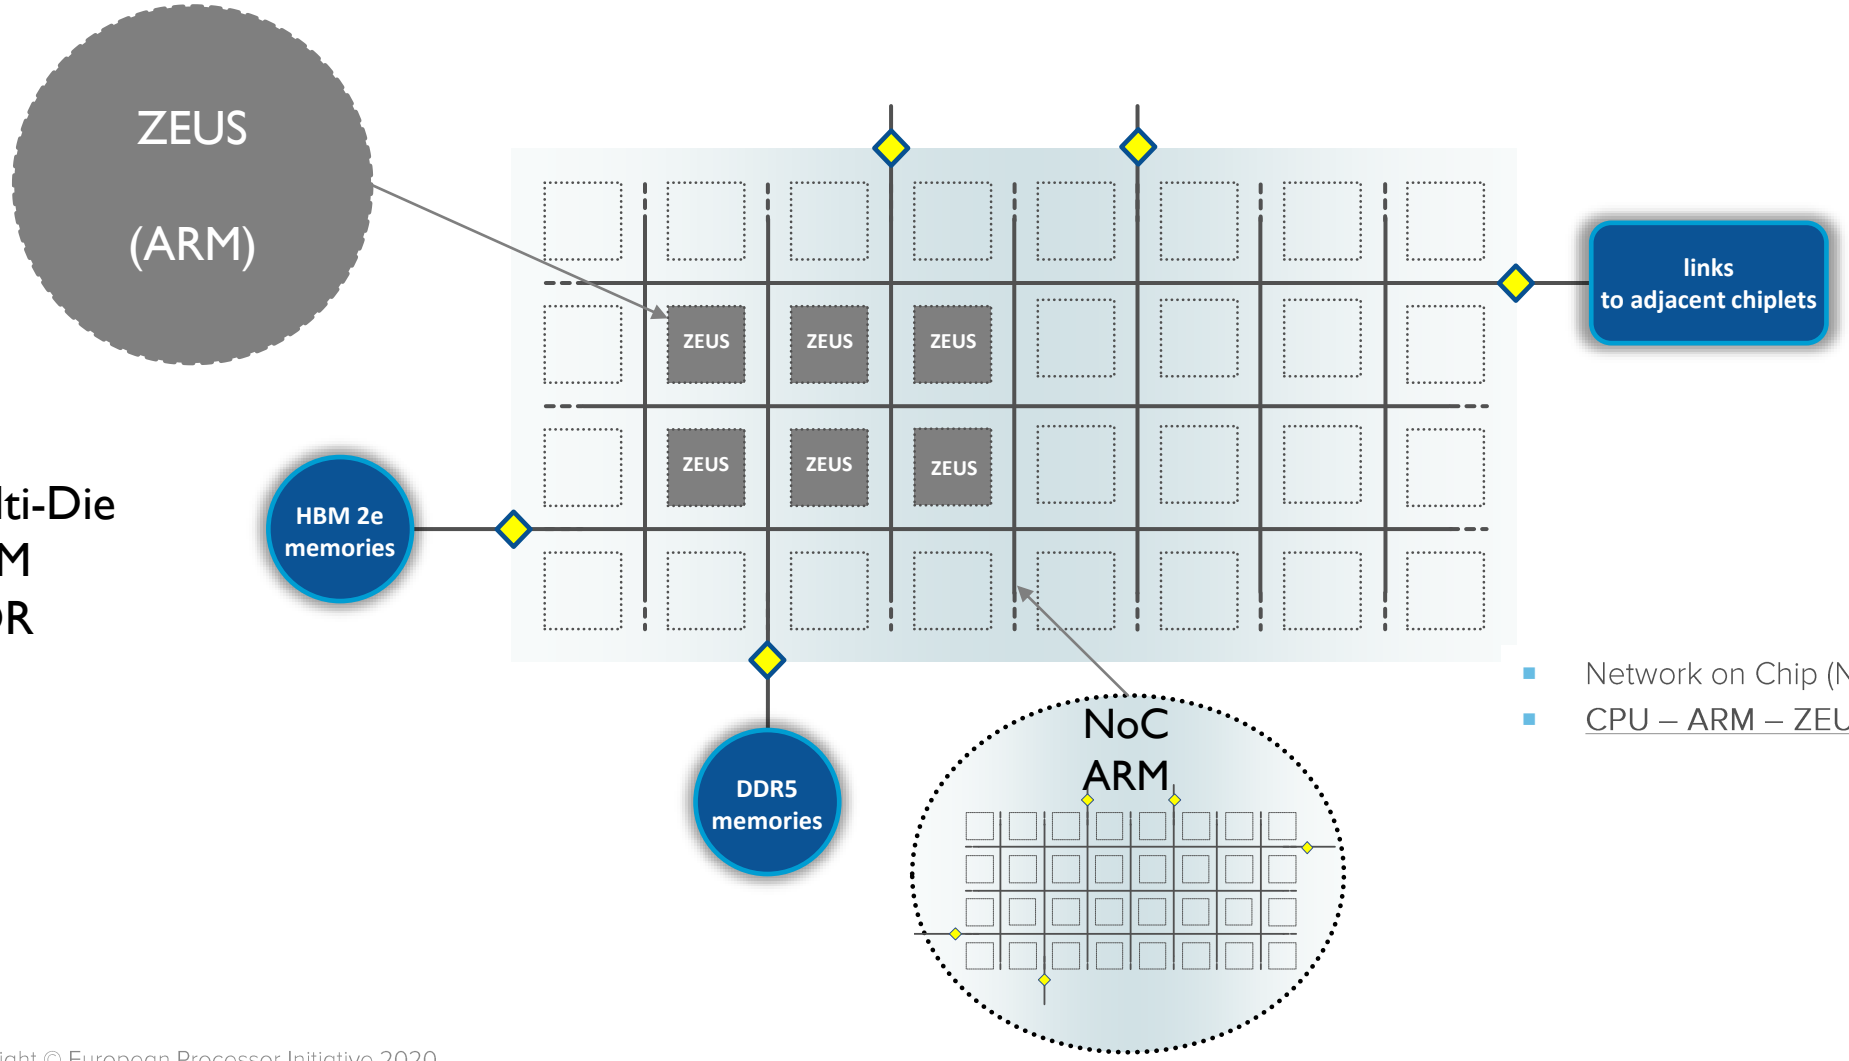

- Multi-Die
- HBM
- DDR

- Network on Chip (NoC) – ARM
- CPU – ARM – ZEUS

GENERAL PURPOSE PROCESSOR (GPP) AND COMMON OPEN ARCHITECTURE

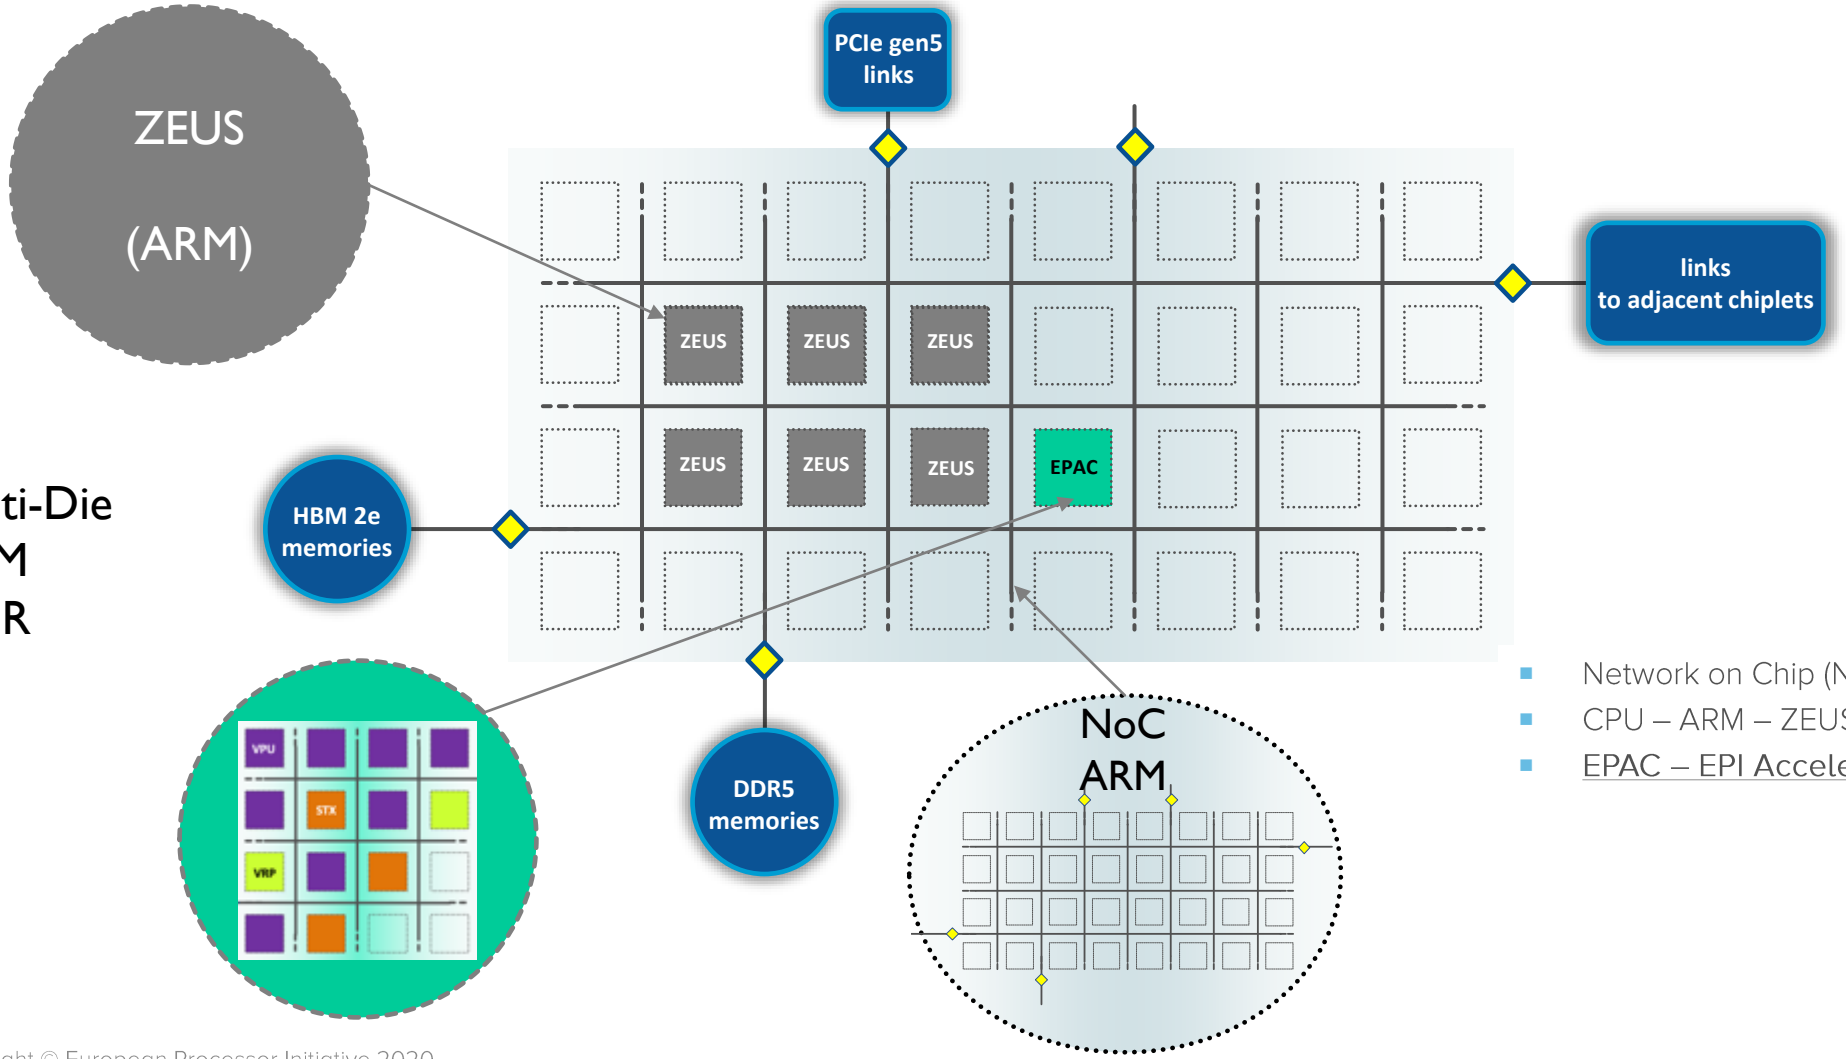

- Multi-Die
- HBM
- DDR

- Network on Chip (NoC) – ARM
- CPU – ARM – ZEUS
- EPAC – EPI Accelerator (TITAN)

GENERAL PURPOSE PROCESSOR (GPP) AND COMMON OPEN ARCHITECTURE

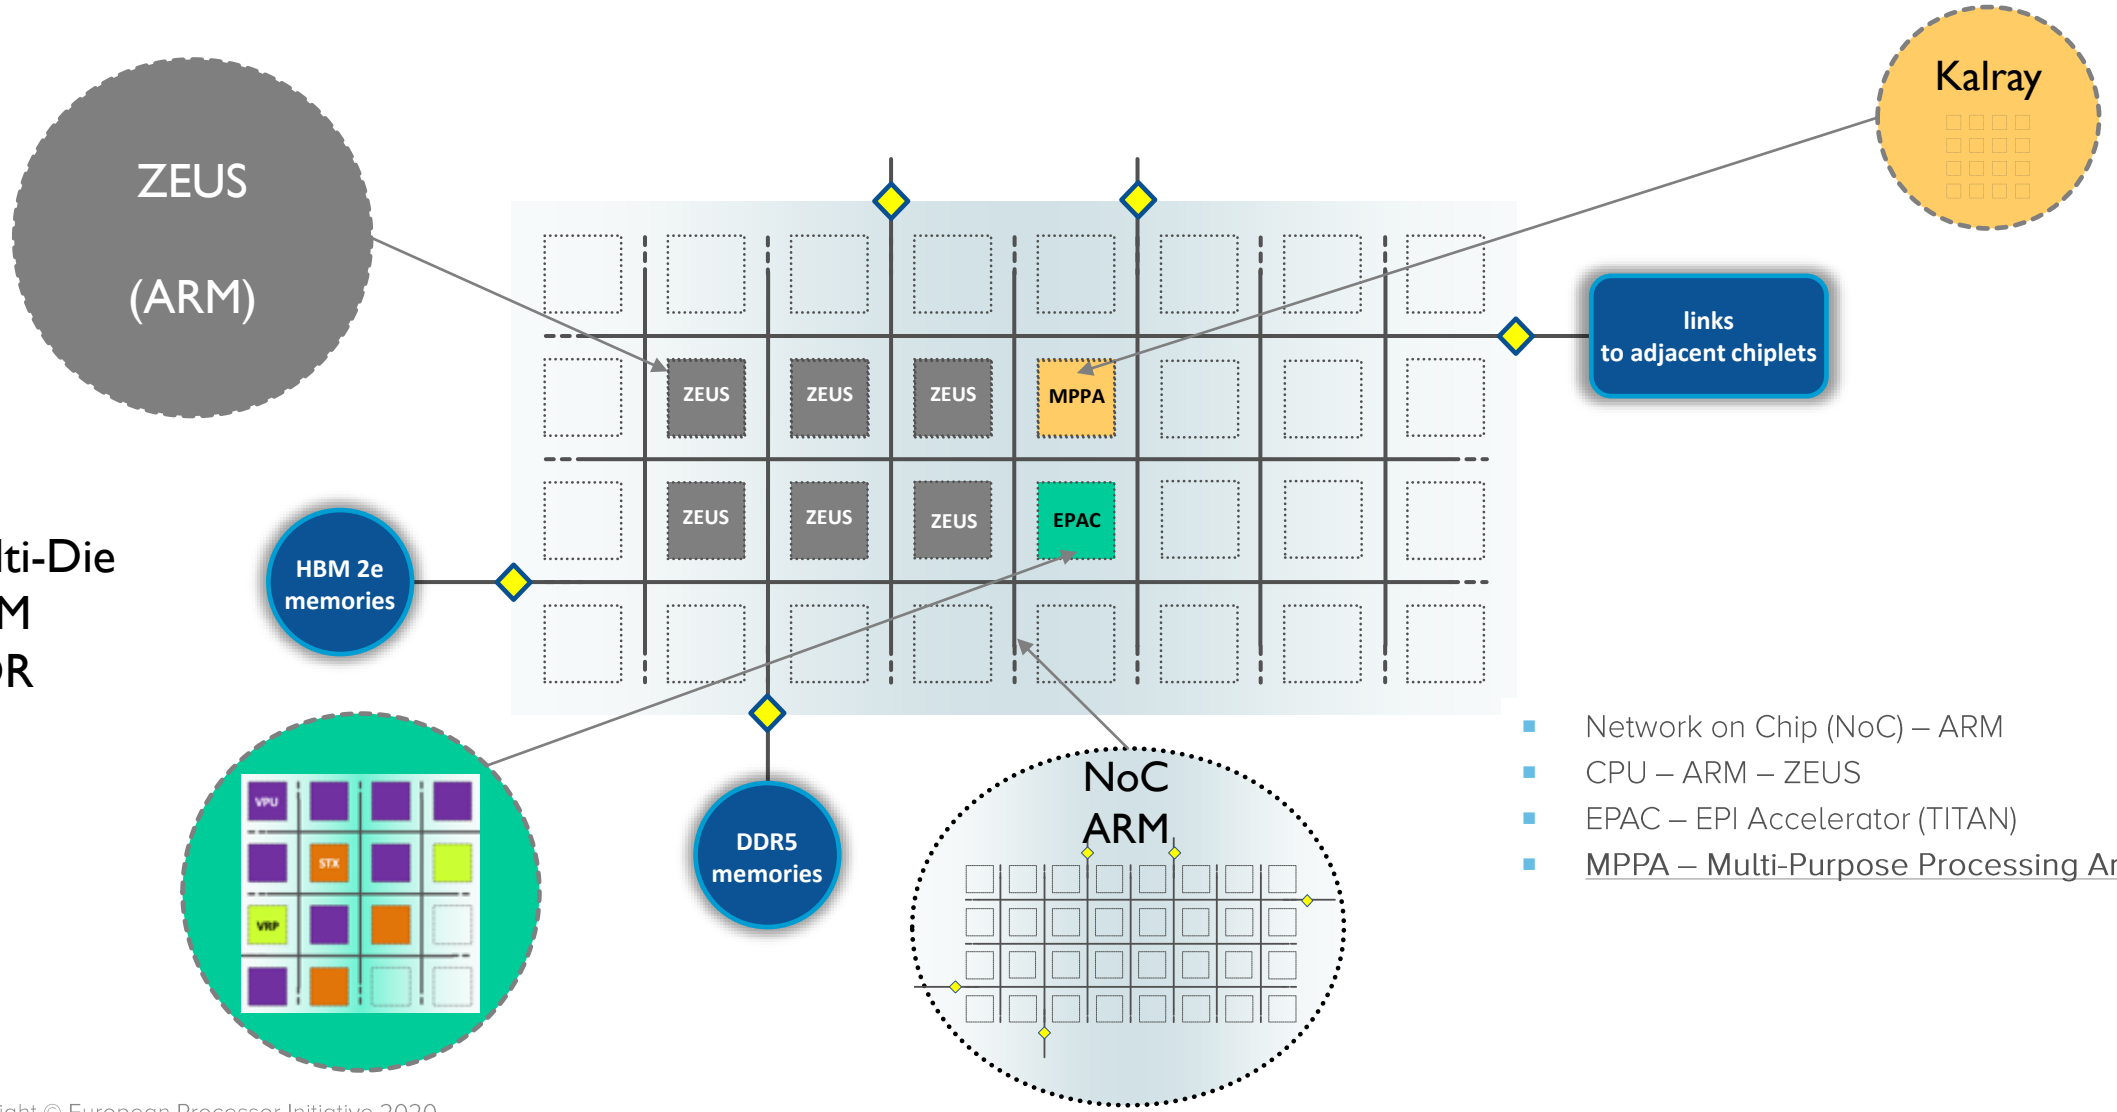

- Multi-Die
- HBM
- DDR

- Network on Chip (NoC) – ARM
- CPU – ARM – ZEUS
- EPAC – EPI Accelerator (TITAN)
- MPPA – Multi-Purpose Processing Array

GENERAL PURPOSE PROCESSOR (GPP) AND COMMON OPEN ARCHITECTURE

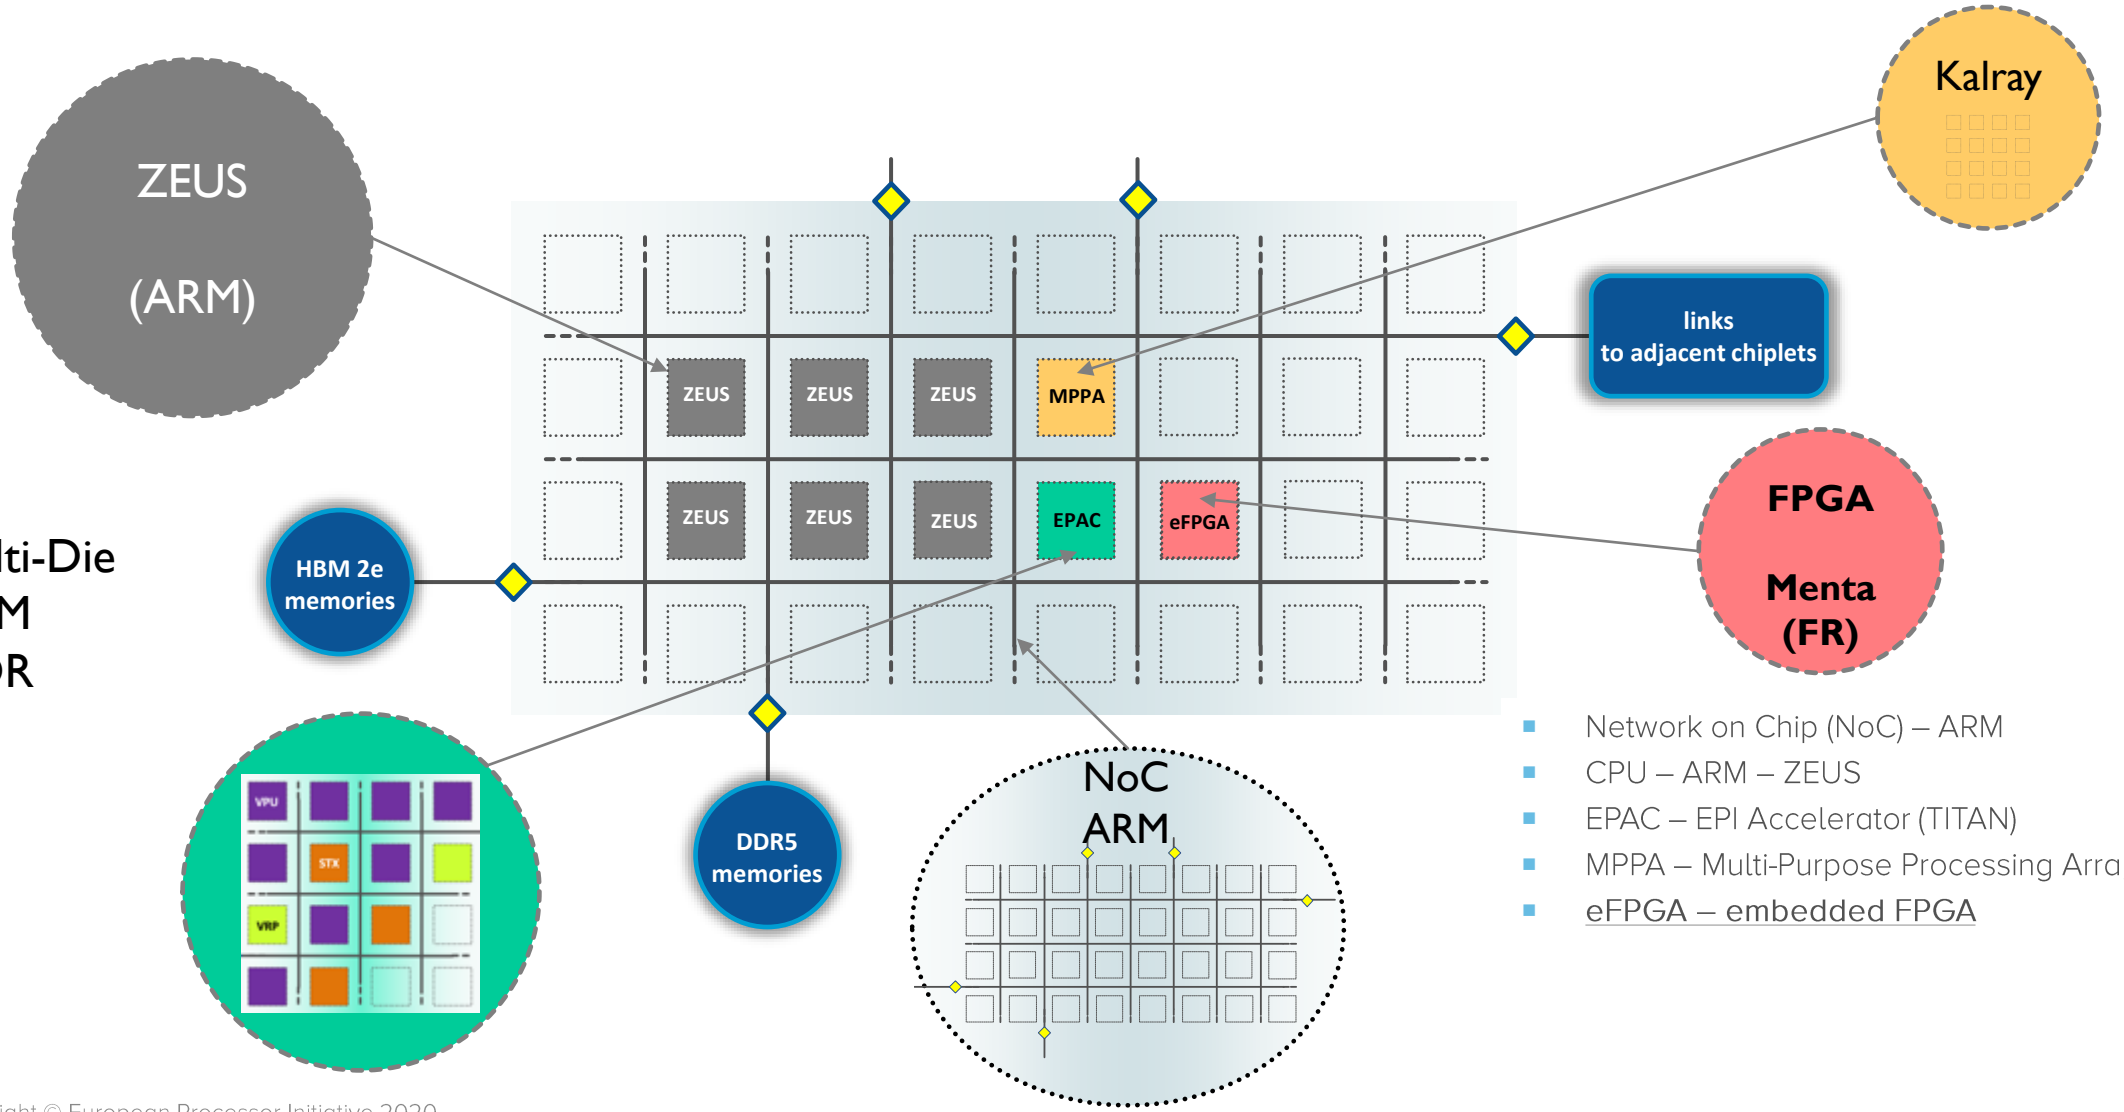

- Multi-Die
- HBM
- DDR

- Network on Chip (NoC) – ARM
- CPU – ARM – ZEUS
- EPAC – EPI Accelerator (TITAN)
- MPPA – Multi-Purpose Processing Array
- eFPGA – embedded FPGA

GENERAL PURPOSE PROCESSOR (GPP) AND COMMON OPEN ARCHITECTURE

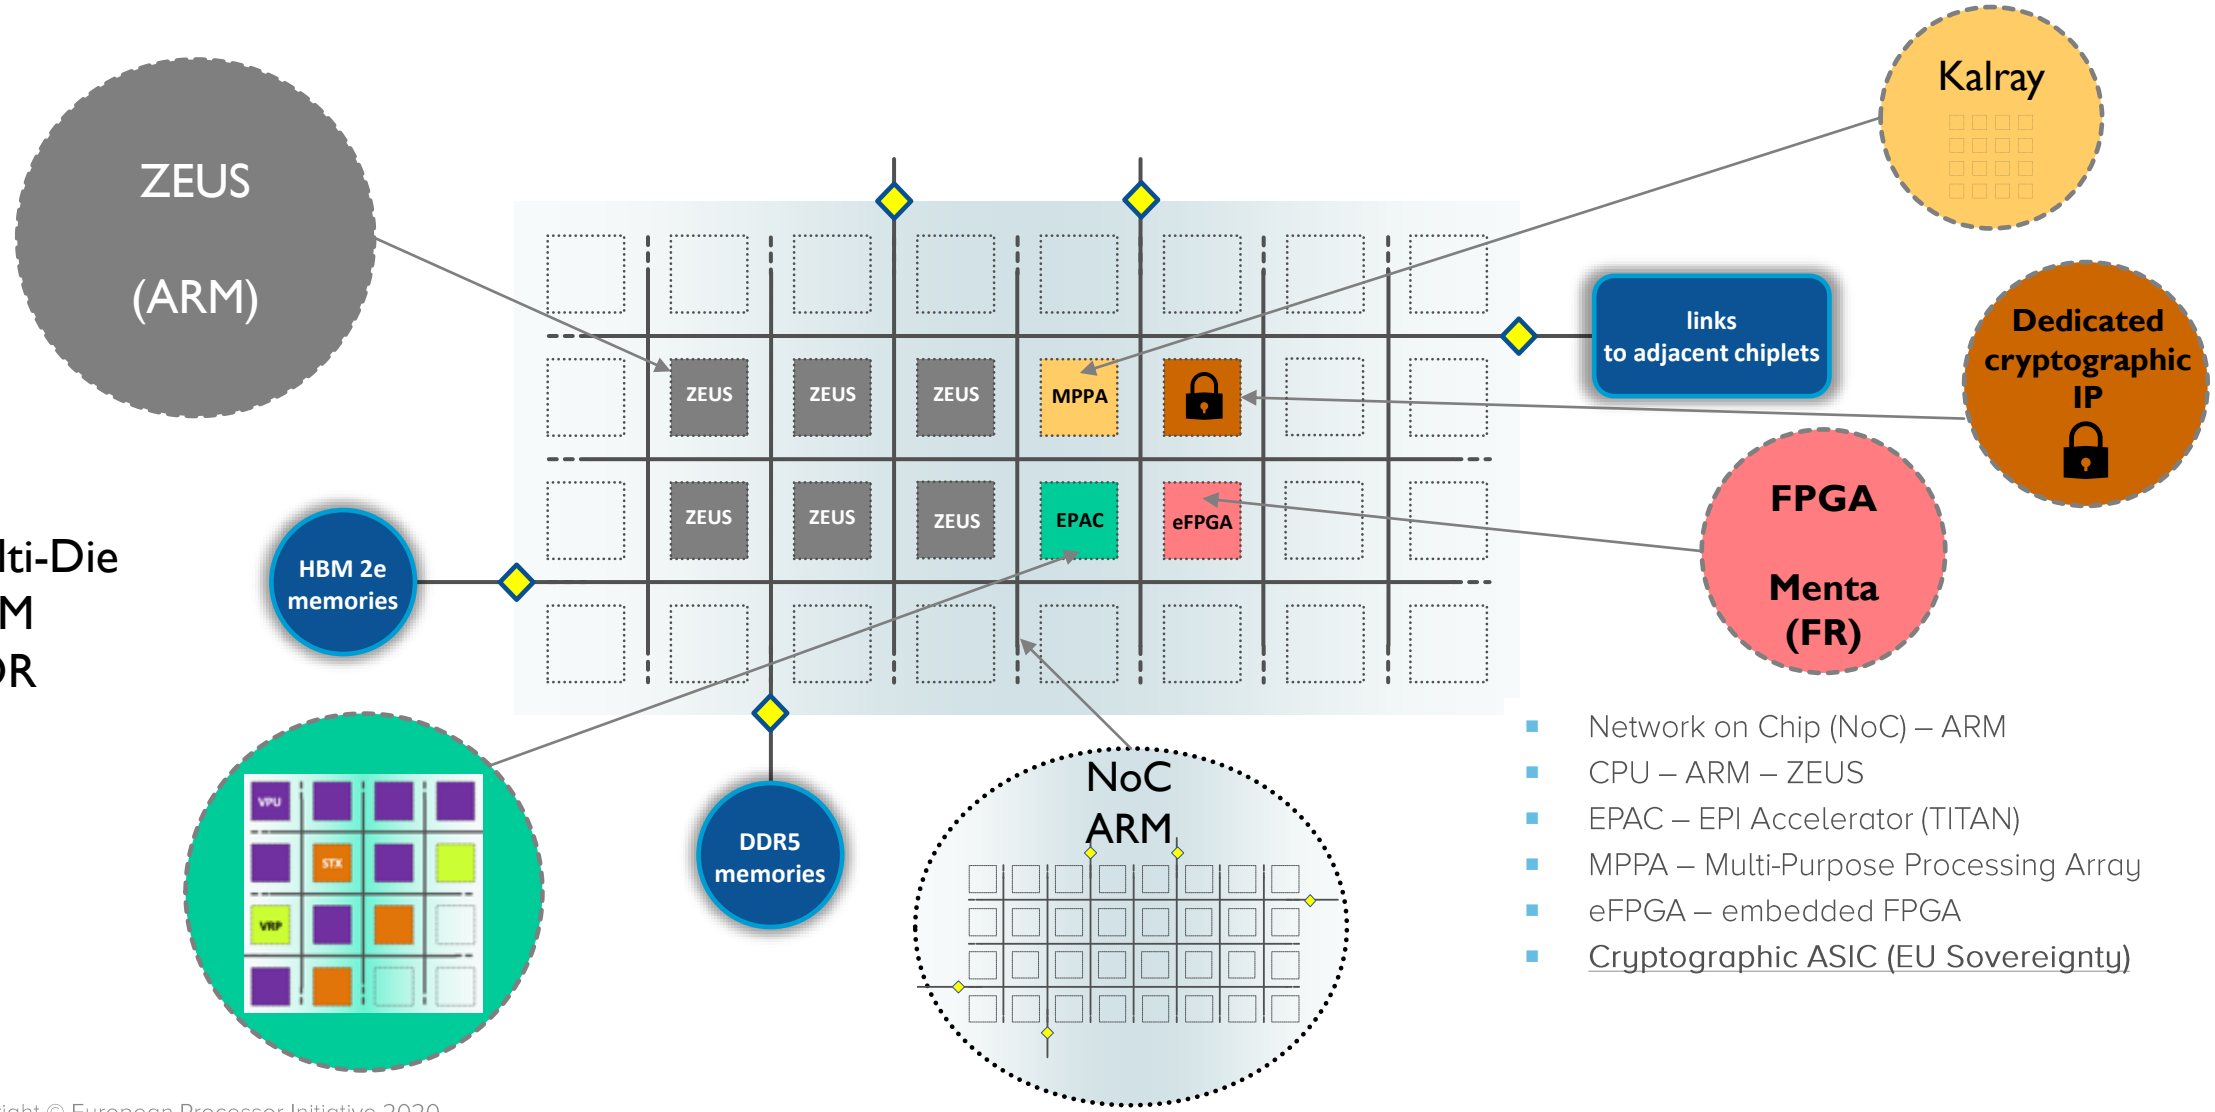

- Multi-Die
- HBM
- DDR

- Network on Chip (NoC) – ARM
- CPU – ARM – ZEUS
- EPAC – EPI Accelerator (TITAN)
- MPPA – Multi-Purpose Processing Array
- eFPGA – embedded FPGA
- Cryptographic ASIC (EU Sovereignty)

GENERAL PURPOSE PROCESSOR (GPP) AND COMMON OPEN ARCHITECTURE

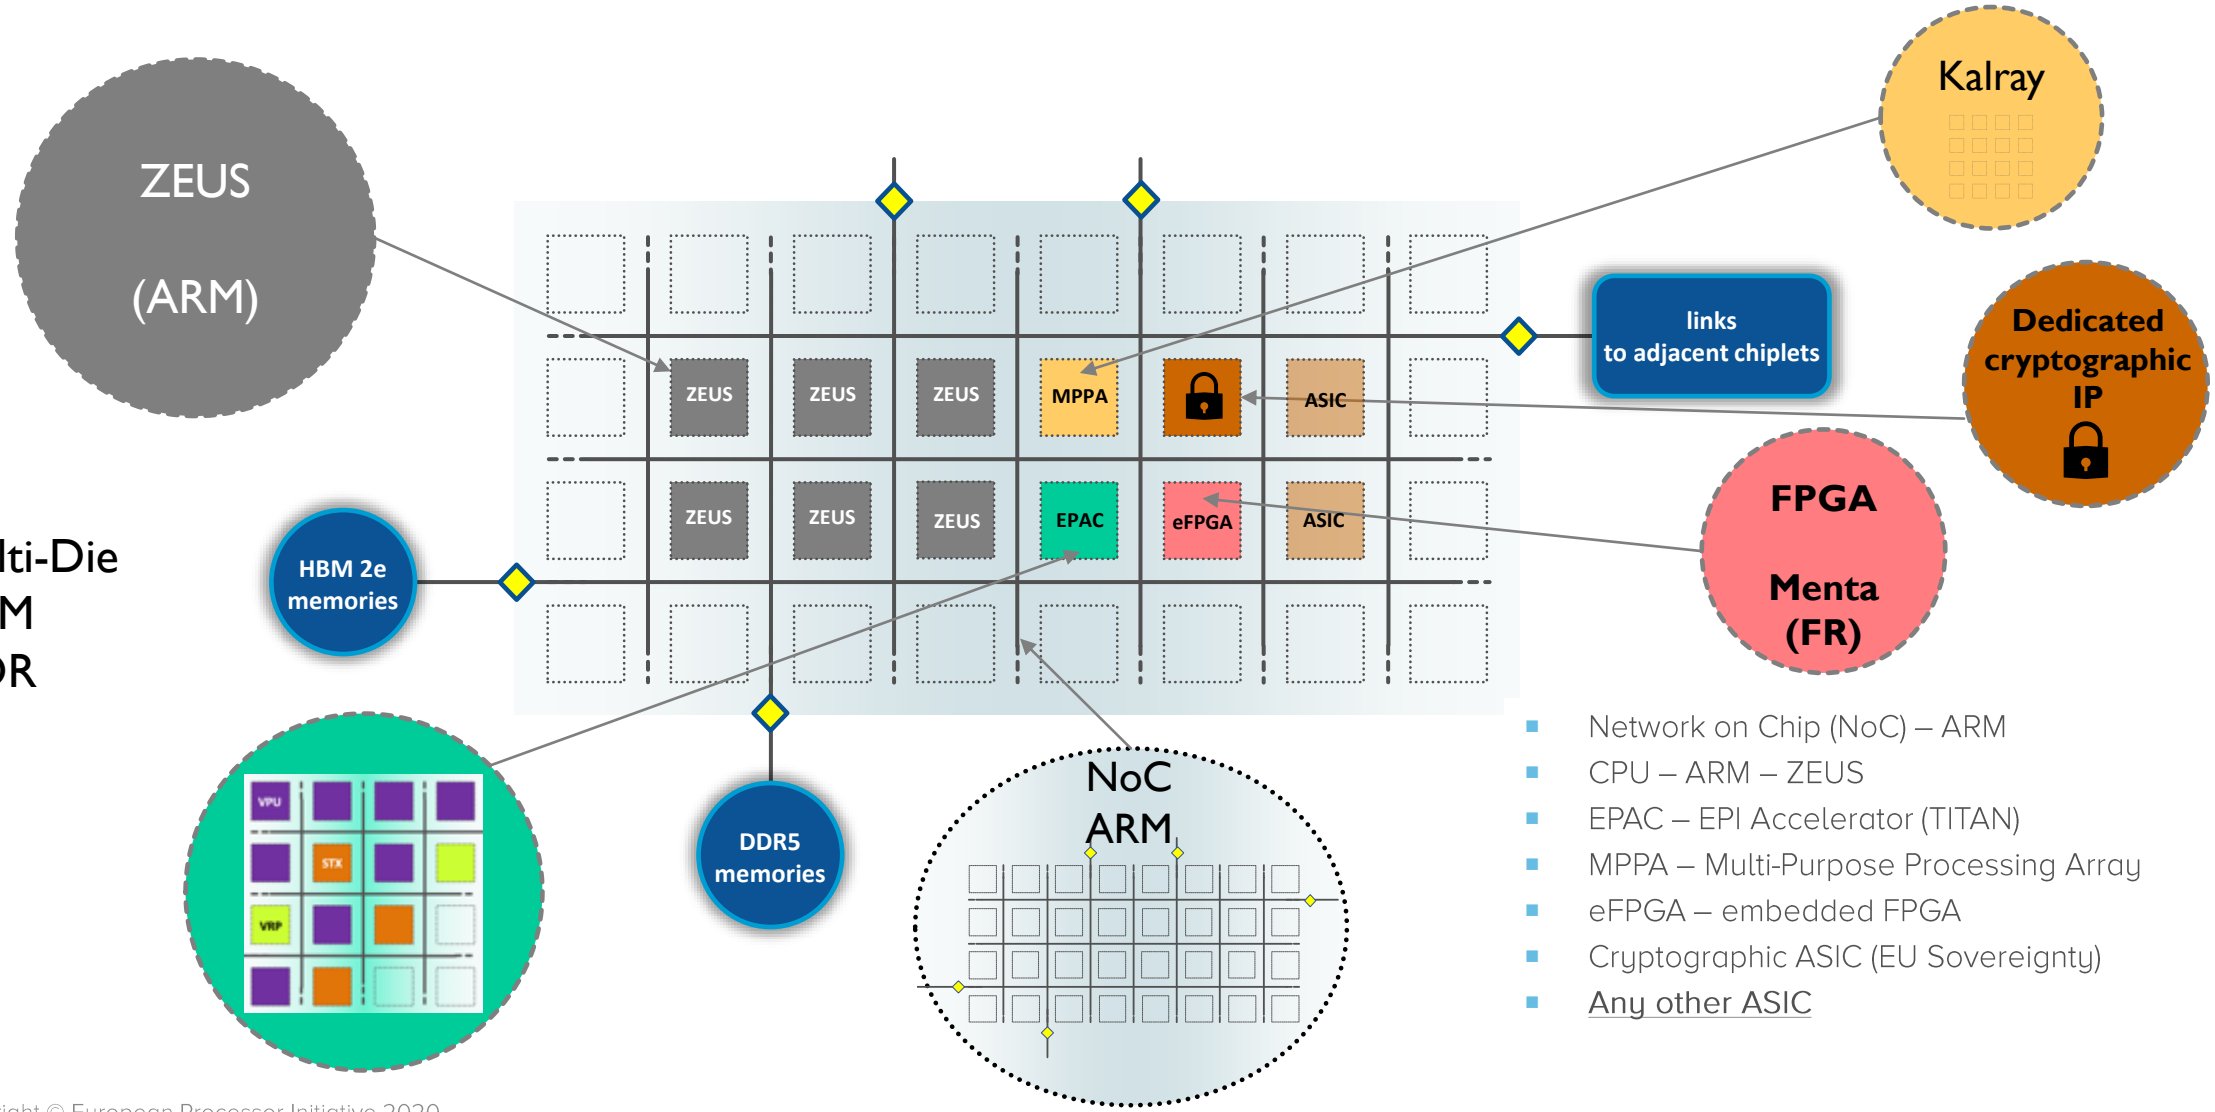

- Multi-Die
- HBM
- DDR

- Network on Chip (NoC) – ARM
- CPU – ARM – ZEUS
- EPAC – EPI Accelerator (TITAN)
- MPPA – Multi-Purpose Processing Array
- eFPGA – embedded FPGA
- Cryptographic ASIC (EU Sovereignty)
- Any other ASIC

GENERAL PURPOSE PROCESSOR (GPP) AND COMMON OPEN ARCHITECTURE

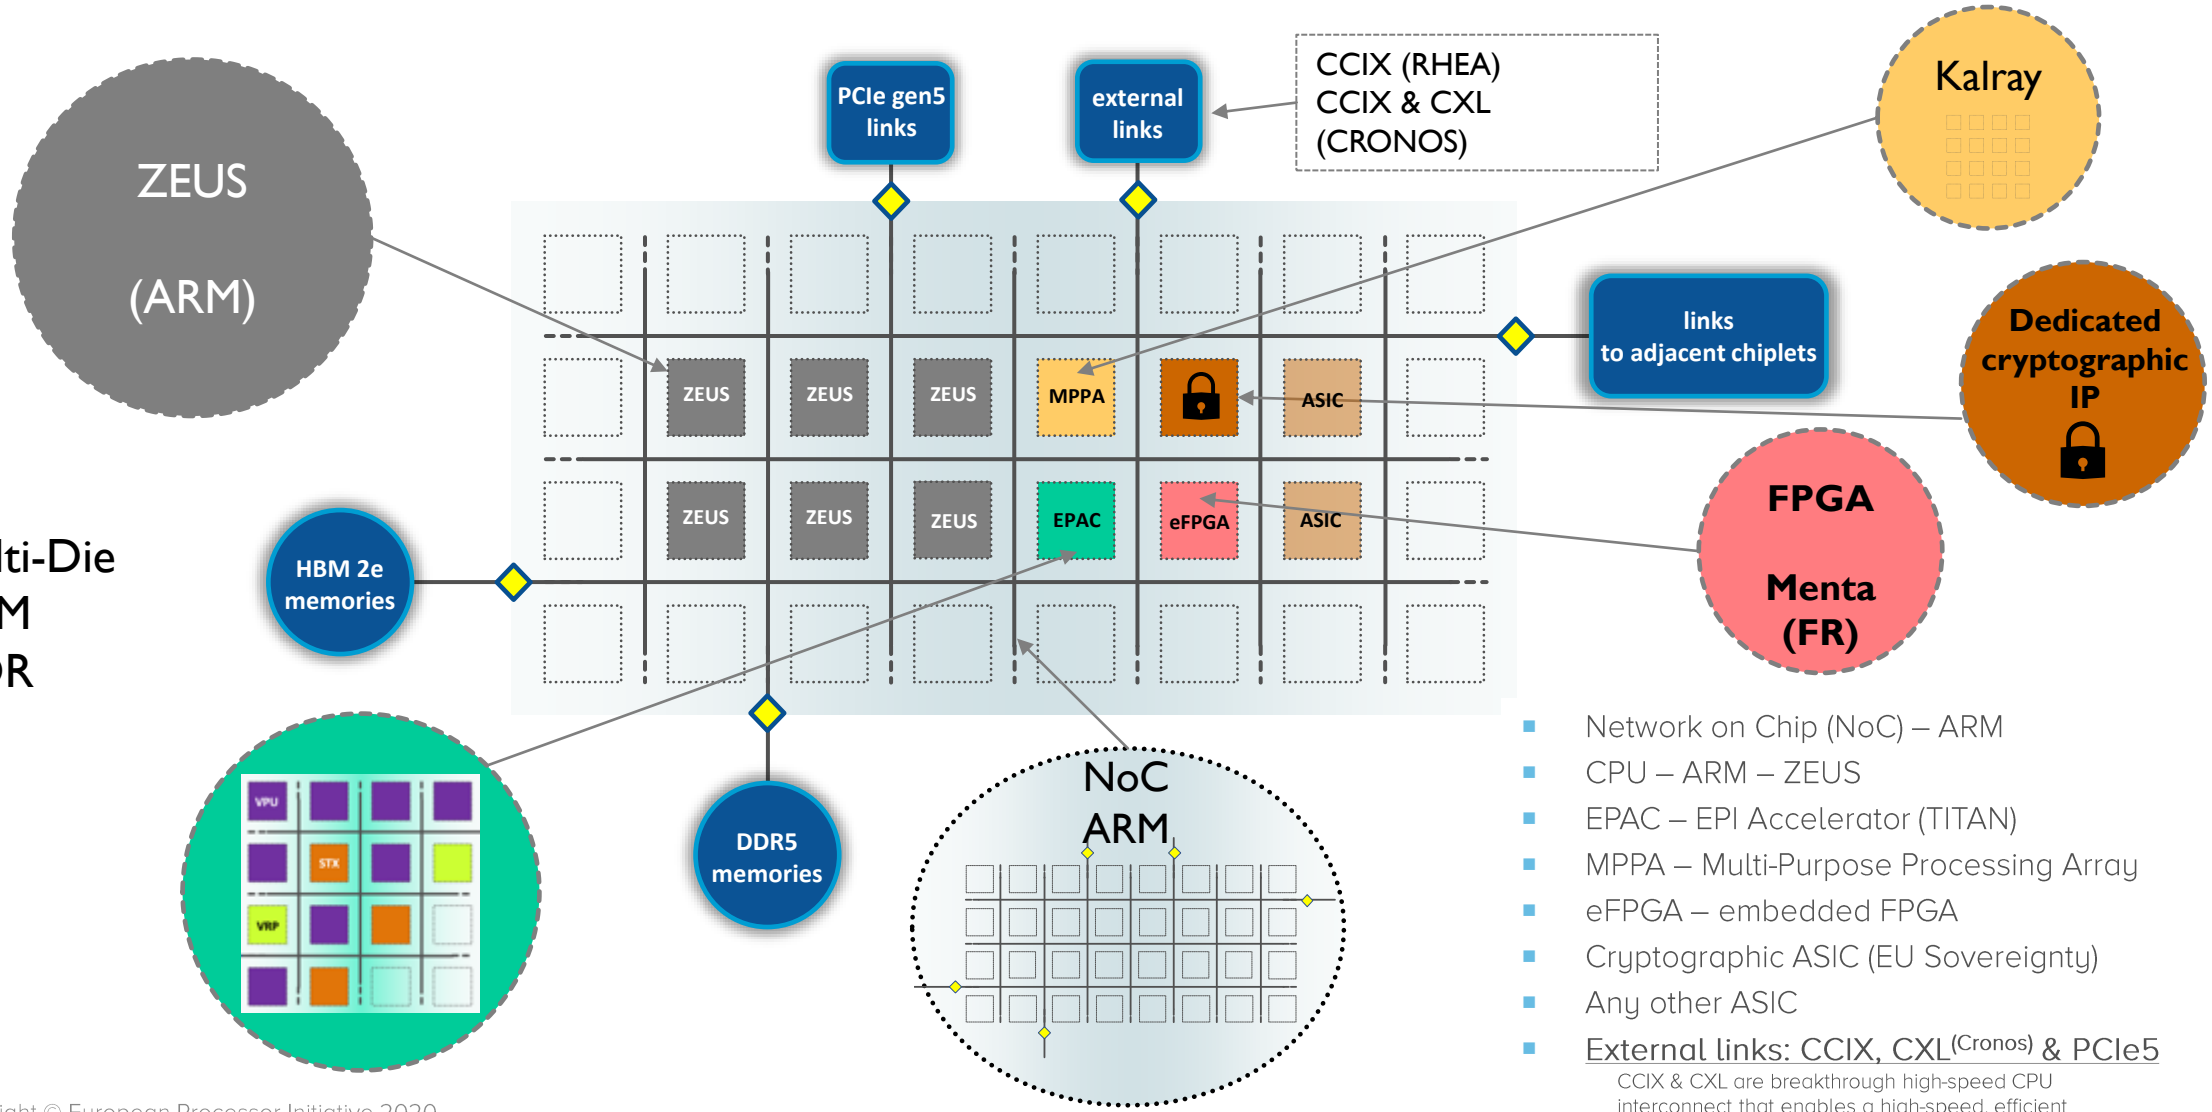

- Multi-Die
- HBM
- DDR

- Network on Chip (NoC) – ARM
- CPU – ARM – ZEUS
- EPAC – EPI Accelerator (TITAN)
- MPPA – Multi-Purpose Processing Array
- eFPGA – embedded FPGA
- Cryptographic ASIC (EU Sovereignty)
- Any other ASIC
- External links: CCIX, CXL(Cronos) & PCIe5

  CCIX & CXL are breakthrough high-speed CPU interconnect that enables a high-speed, efficient performance between the CPU and platform enhancements and workload accelerators

GENERAL PURPOSE PROCESSOR (GPP) AND COMMON OPEN ARCHITECTURE

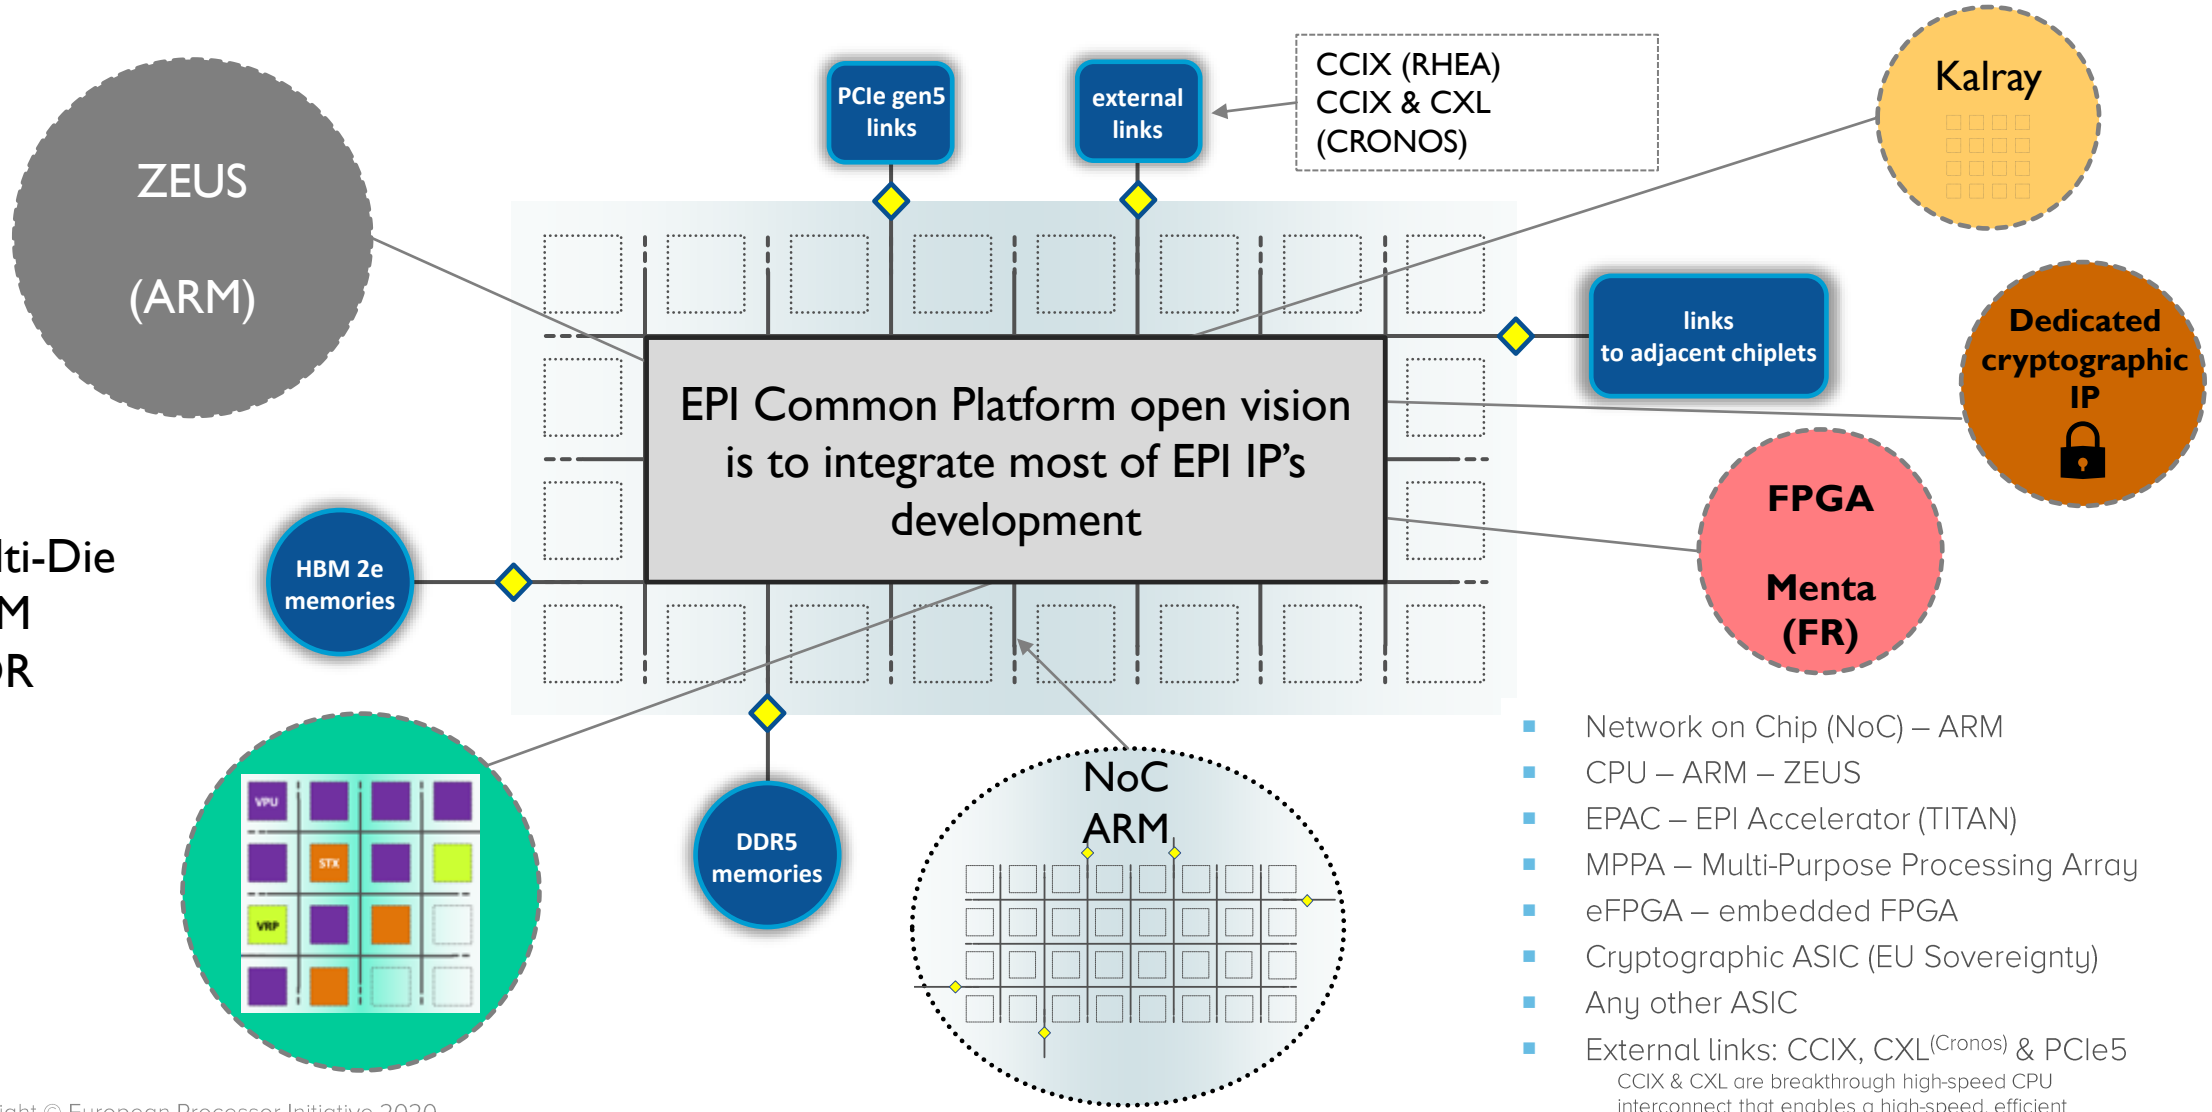

- Multi-Die
- HBM
- DDR

- Network on Chip (NoC) – ARM
- CPU – ARM – ZEUS
- EPAC – EPI Accelerator (TITAN)
- MPPA – Multi-Purpose Processing Array
- eFPGA – embedded FPGA
- Cryptographic ASIC (EU Sovereignty)
- Any other ASIC
- External links: CCIX, CXL(Cronos) & PCIe5

  CCIX & CXL are breakthrough high-speed CPU interconnect that enables a high-speed, efficient performance between the CPU and platform enhancements and workload accelerators

GENERAL PURPOSE PROCESSOR (GPP) AND COMMON OPEN ARCHITECTURE

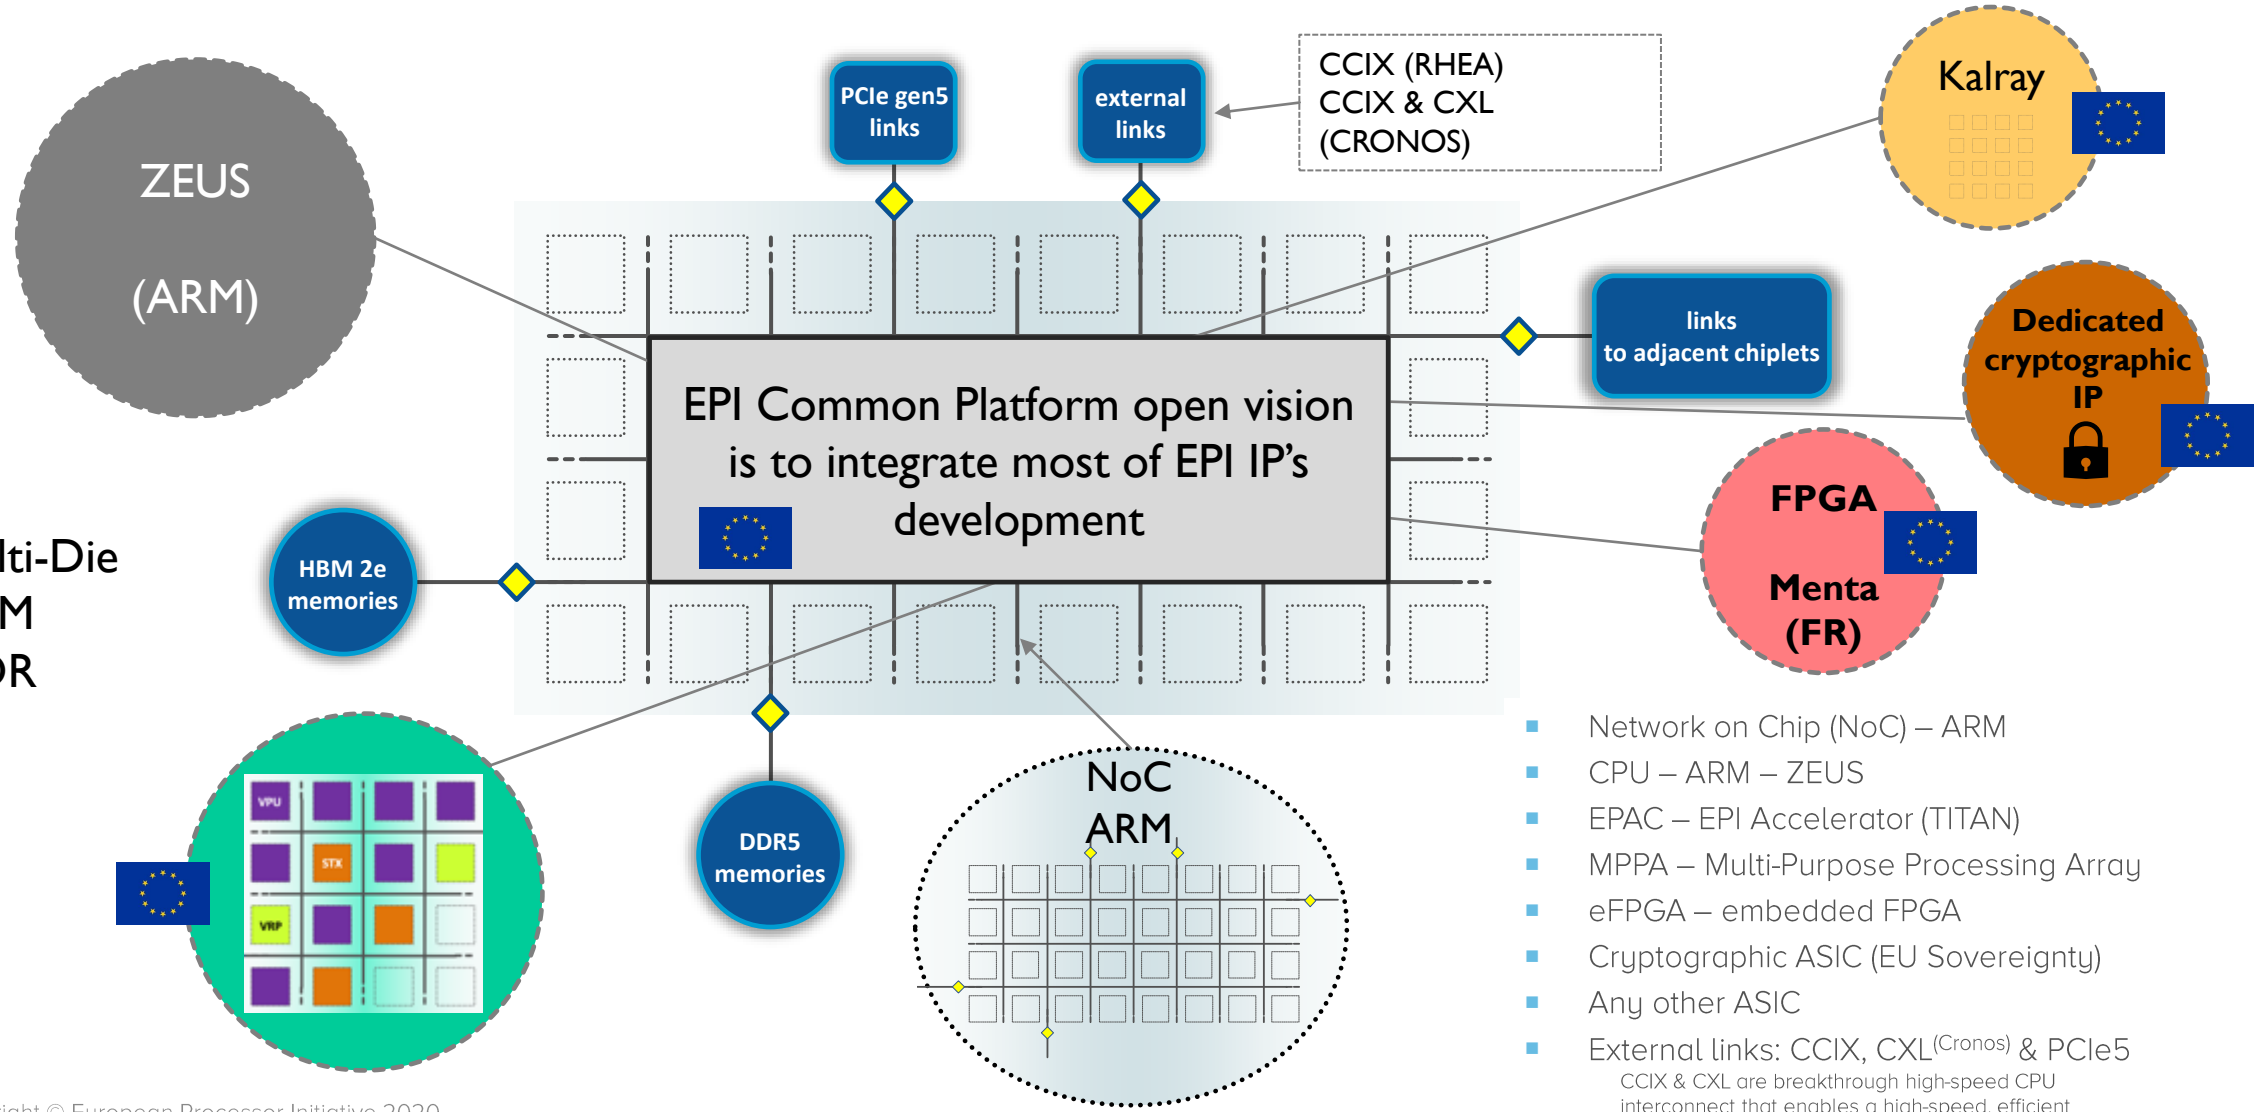

- Multi-Die
- HBM
- DDR

- Network on Chip (NoC) – ARM
- CPU – ARM – ZEUS
- EPAC – EPI Accelerator (TITAN)
- MPPA – Multi-Purpose Processing Array
- eFPGA – embedded FPGA
- Cryptographic ASIC (EU Sovereignty)
- Any other ASIC
- External links: CCIX, CXL(Cronos) & PCIe5

  CCIX & CXL are breakthrough high-speed CPU interconnect that enables a high-speed, efficient performance between the CPU and platform enhancements and workload accelerators

PERSPECTIVES AND CHALLENGES

EPI DELIVERS!

- The expertise for developing high-end and complex processing units in Europe, after decades of dis-investment
- A General Purpose Processor for HPC machines can be developed in EU by a EU Company (SiPearl)
- We'll be ready to move to the next step: engage on the development of a 100% EU IP general purpose processor.

CHALLENGES?

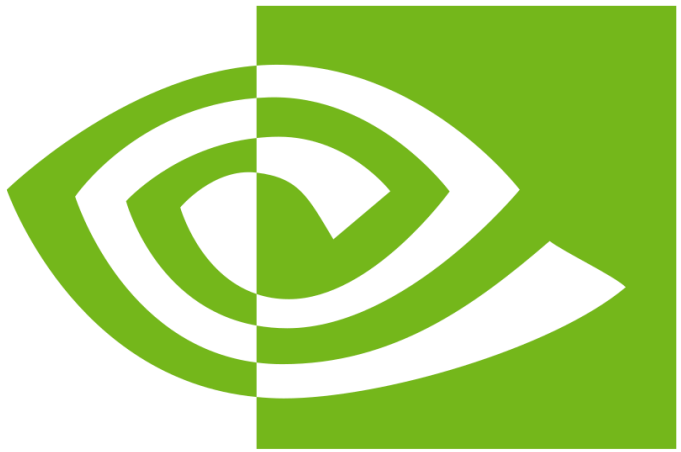

NVIDIA®

CHALLENGES?

THANK YOU FOR YOUR ATTENTION

European Processor Initiative

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