

THE STRATEGIC TOUCH BETWEEN SUPERCOMPUTERS AND EMBEDDED SYSTEMS IN THE ROADMAP TOWARDS EXASCALE: THE EUROPEAN PROCESSOR INITIATIVE

HIPEAC ACACES SUMMER SCHOOL, FIUGGI, ITALY

14 JULY 2019



MAURO OLIVIERI

VISITING RESEARCHER, BARCELONA SUPERCOMPUTING CENTER

PROFESSOR, SAPIENZA UNIVERSITY OF ROME



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

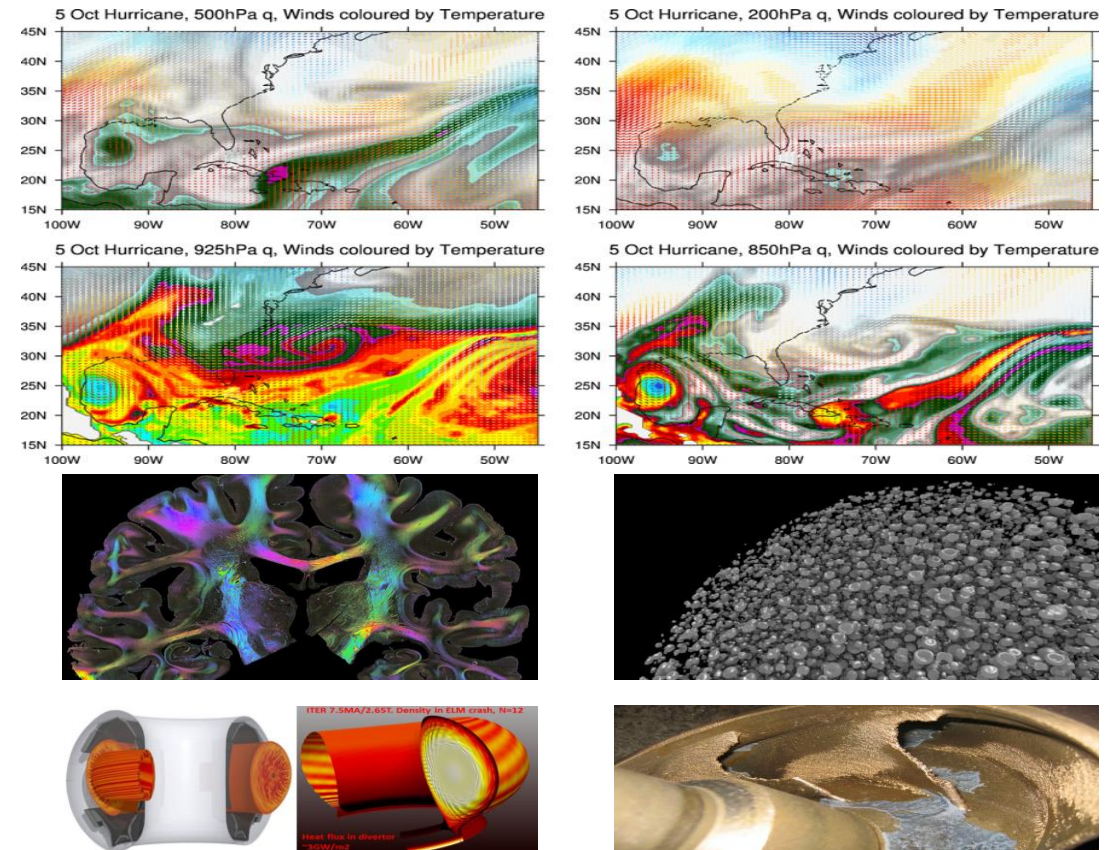


**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación




WHY IS HPC NEEDED?

- Will save billions by helping us to adapt to climate change
- Will improve human health by enabling personalized medicine
- Will improve fuel efficiency of aircraft & help design better wind turbines
- Will help us to understand how the human brain works
- But not only this....




Images courtesy of The PRACE Scientific Steering Committee, "The Scientific Case for Computing in Europe 2018-2026"

MICROPROCESSOR TRENDS



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2015 by K. Rupp

SO WHAT'S HAPPENED IN SUPERCOMPUTING...?



[1]	ASCI Red (Intel Teraflops)
[2]	Red(memory upgrade)
[3]	ASCI Blue Pacific (IBM RS/6000)
[4]	ASCI White
[5]	Cray XT3
[6]	IBM Roadrunner
[7]	IBM Sequoia (BLUGENE/q)
[8]	Tianhe-1A
[9]	Nebulae
[10]	PRIMEHPC FX10
[11]	Eurora
[12]	Fujitsu PRIMEHPC FX100 (SPARC64 Xifx)
[13]	Sunway TaihuLight
[14]	IBM Summit

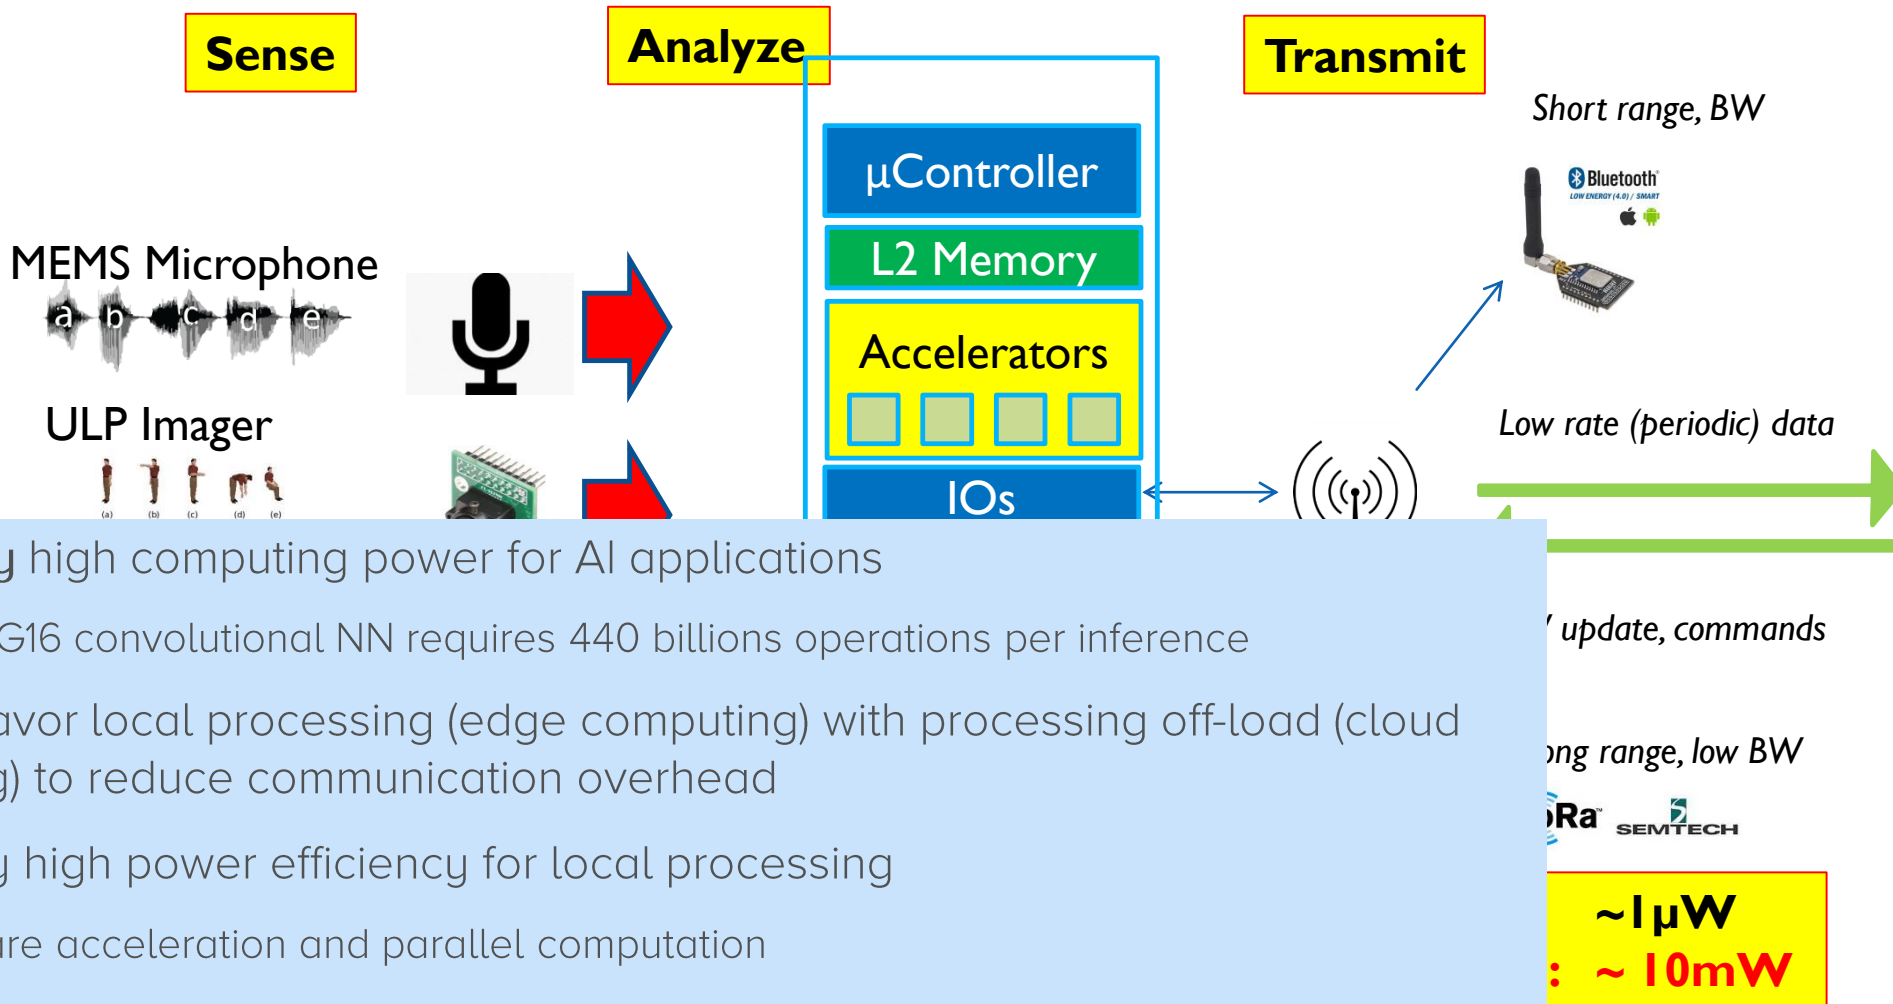
Power efficiency exponential growth is expected to reach between 33 GFLOPS and 50 GFLOPS per Watt in 7-5 nm technologies, through extensive use of hardware accelerators

(source: Gagliardi, F., Moreto, M., Olivieri, M., Valero, M., CCF Transactions on HPC, 2019)

SO WHAT'S HAPPENED IN SUPERCOMPUTING...?

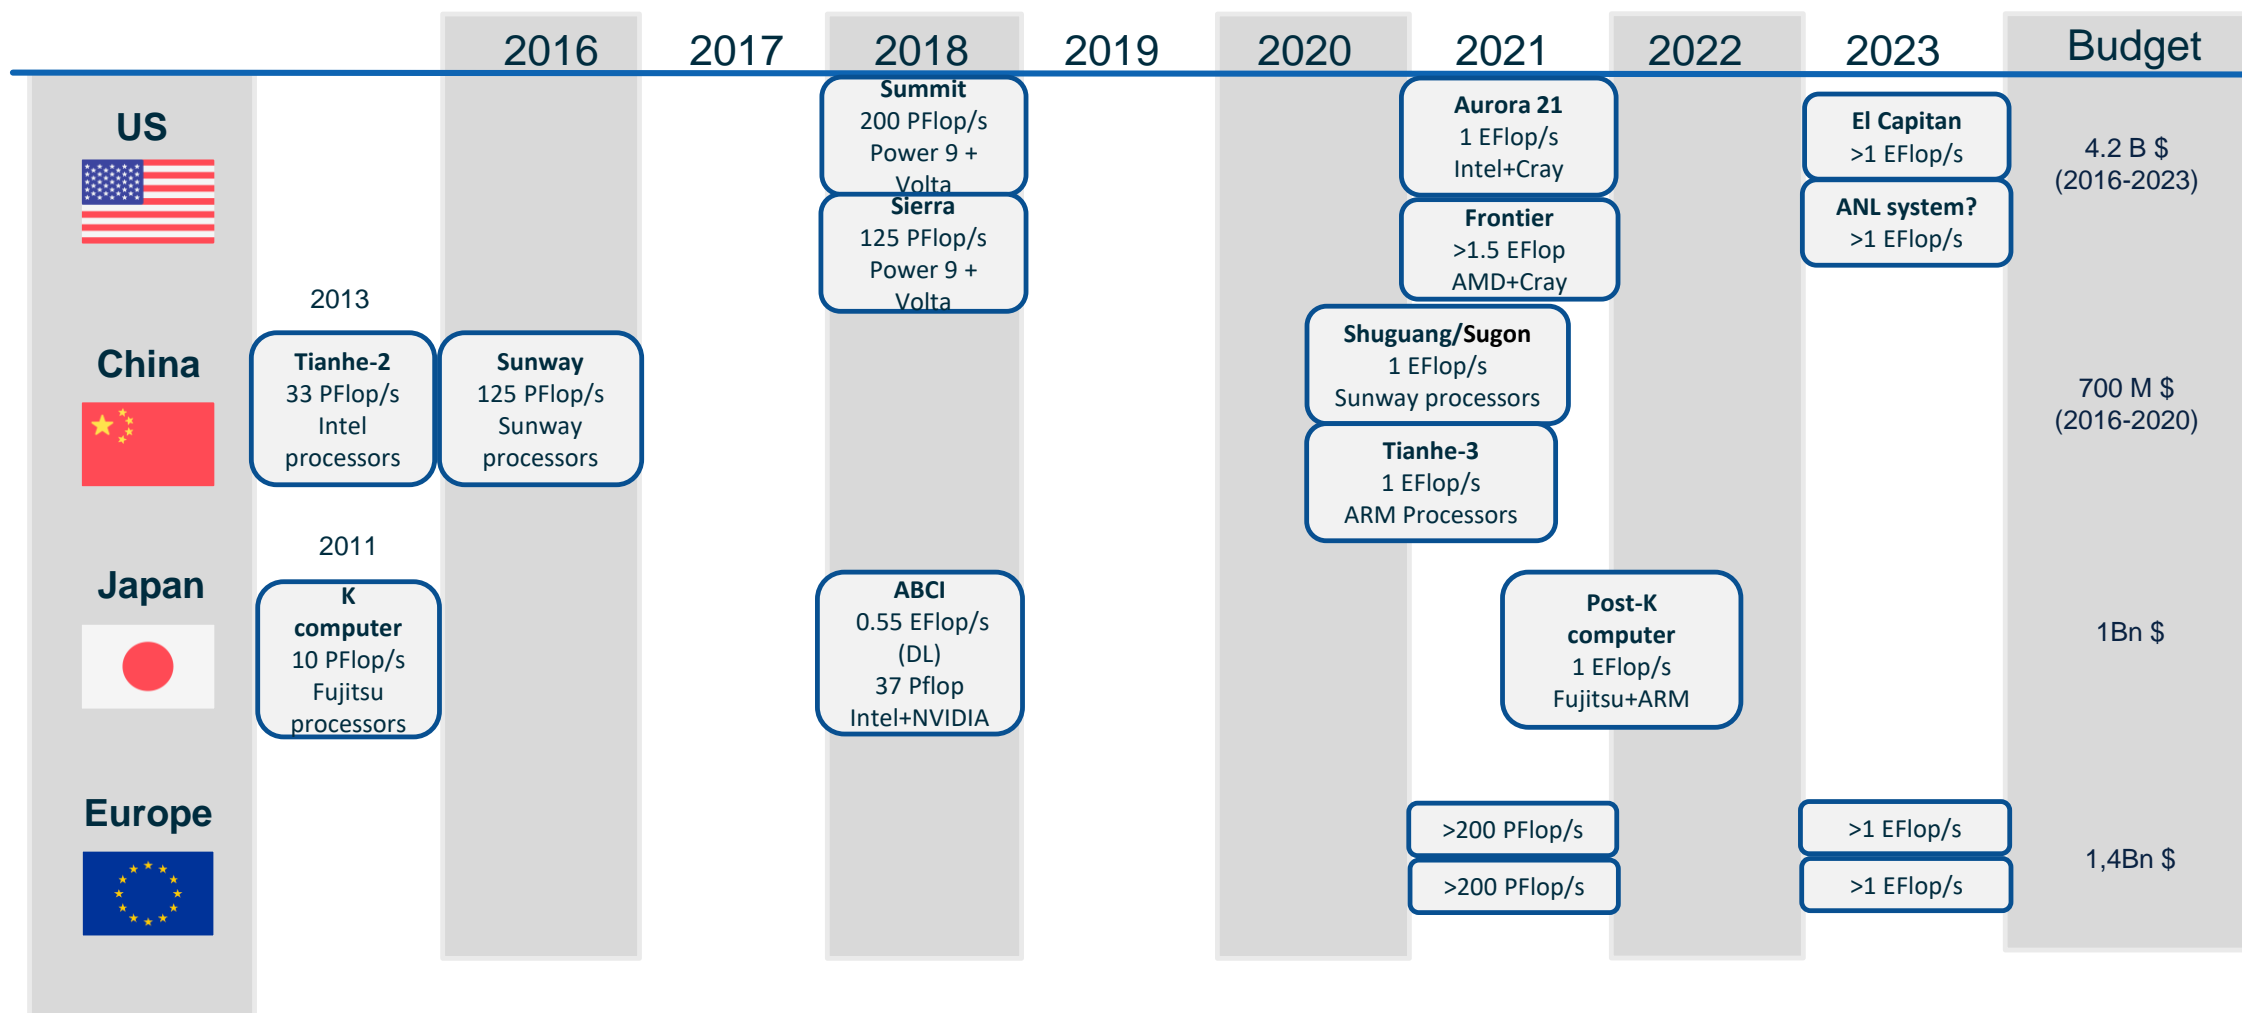
- **Mainframe Era (circa 1953 - circa 1972)**
 - Memory capacity is the main limit
 - All fundamental computer architecture techniques are invented
- **First rise of HW acceleration: vector computers (1974-1993)**
 - Processing speed on matrix algebra is the main limit
 - SIMD processing, domain specific architectures
- **Rise of massive homogeneous parallelism (1994-2007)**
 - Memory bandwidth is the main limit (memory wall)
 - Moore's law boosts clock speed and scale of integration in HIGH VOLUME MARKET processors
 - Parallel architectures with commodity CPUs overcome vector processors (the killer processor effect)
- **The renaissance of acceleration units (2008 - ...?)**
 - Power consumption is the main limit – Hit the power wall
 - Hardware specialization allows better power efficiency (FLOPS/W)
 - The first example are GPUs because they come from HIGH VOLUME MARKET

...AND WHAT'S HAPPENED IN EMBEDDED SYSTEMS FOR IOT?



- Need **very** high computing power for AI applications
 - E.g. VGG16 convolutional NN requires 440 billions operations per inference
- Need to favor local processing (edge computing) with processing off-load (cloud computing) to reduce communication overhead
- Need very high power efficiency for local processing
 - Hardware acceleration and parallel computation
- **Need a supercomputer chip on the sensor node**

THE ROADMAP TOWARDS EXASCALE



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 NRCP
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

WHERE EUROPE NEEDS TO BE STRONGER

- Only 1 of the 10 most powerful HPC systems is in the EU
- HPC codes must be upgraded
- Vital HPC hardware elements are missing: General Purpose Processor and Accelerators



EU needs its own source of as many of the system elements as possible

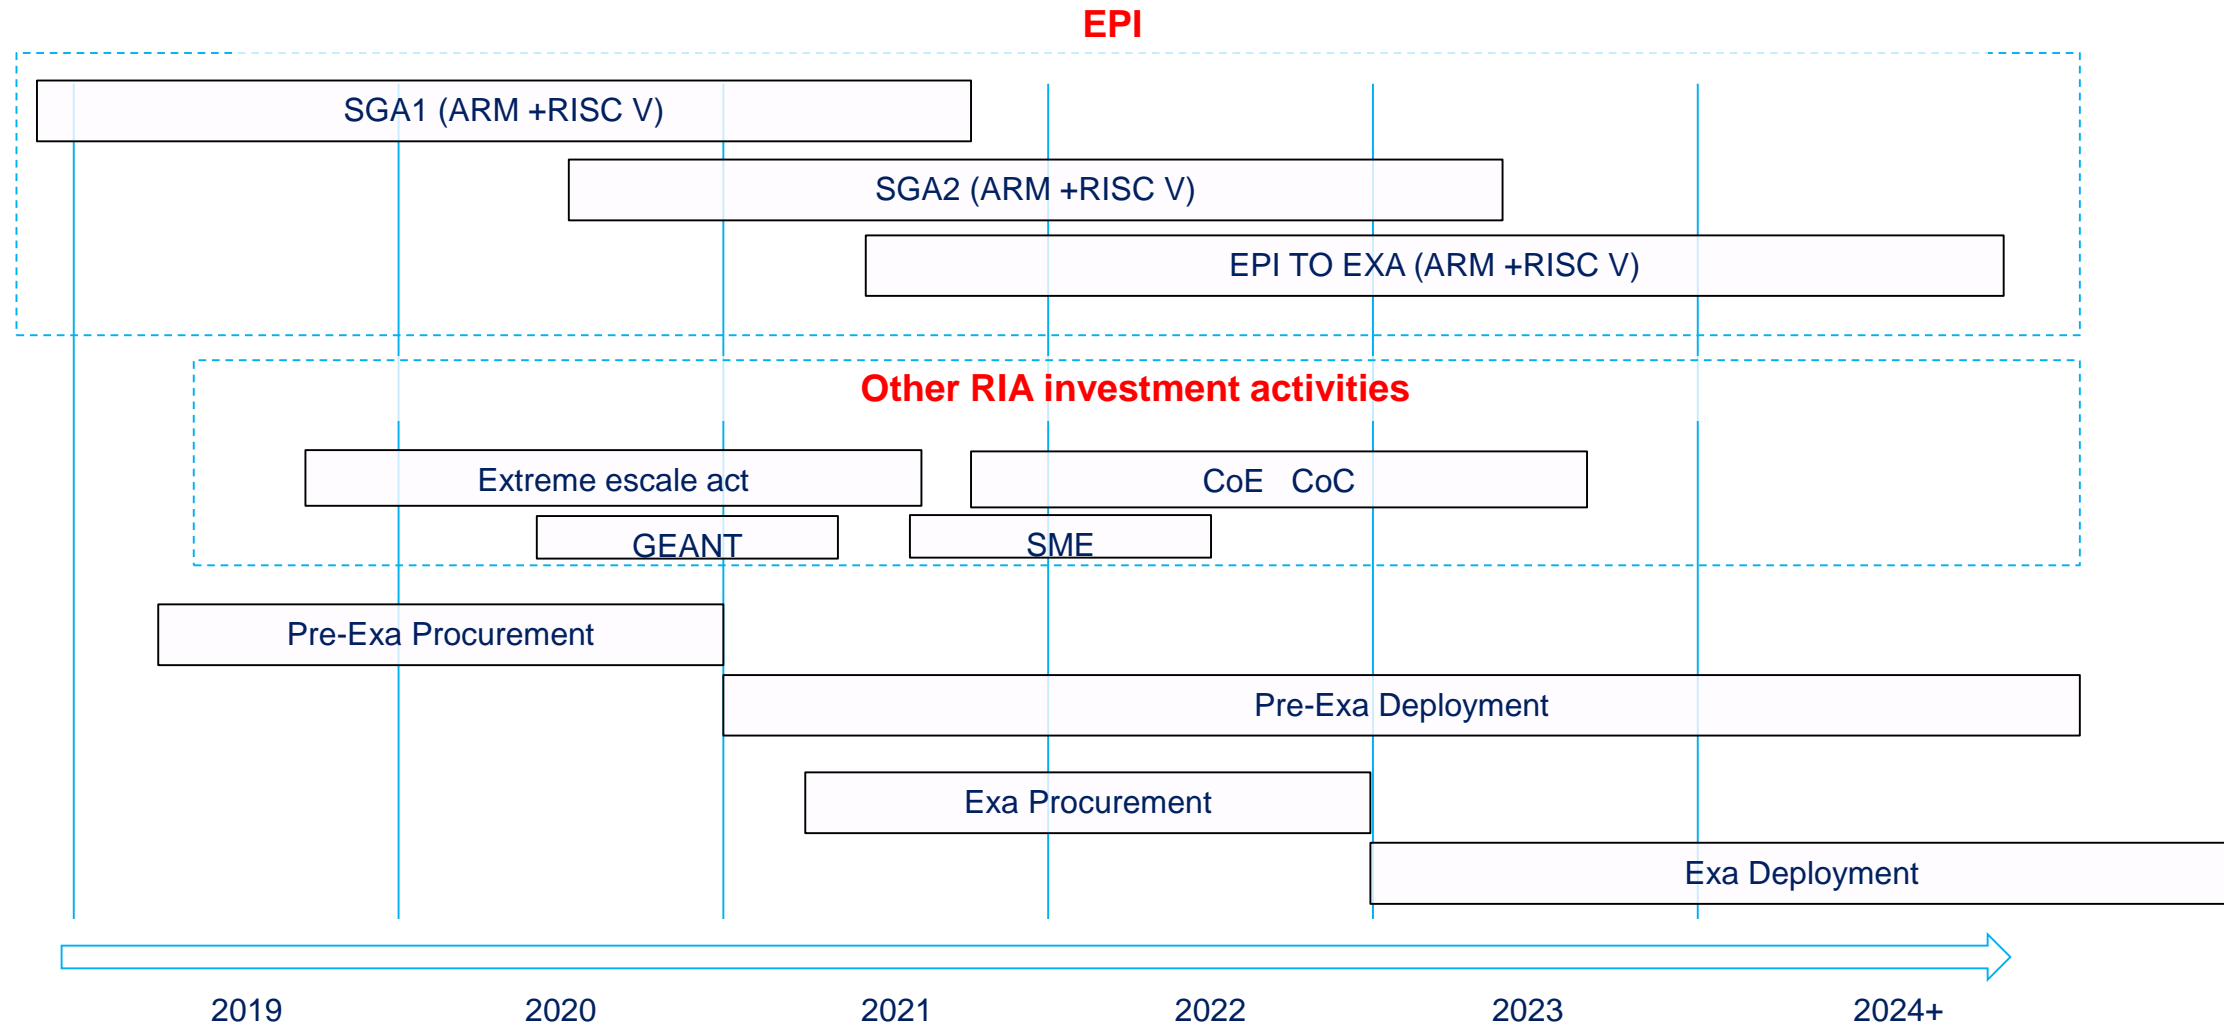
WHY EUROPE NEEDS ITS OWN PROCESSORS

- Processors now control almost every aspect of people's 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



Images courtesy of Reuters.com and European Processor Initiative

EuroHPC ROADMAP



Sign up for our newsletter and get the latest HPC news and analysis.

Email Address

[Home](#) » [HPC Hardware](#) » Future MareNostrum 5 Supercomputer at BSC to host Made-in-Europe Technologies

Future MareNostrum 5 Supercomputer at BSC to host Made-in-Europe Technologies

June 10, 2019 by [staff](#) [Leave a Comment](#) 

Today the [Barcelona Supercomputing Center](#) announced plans to host EuroHPC processor technologies within its future Mare Norstrom 5 supercomputer. The [EuroHPC Joint Undertaking](#) has chosen BSC to be one of the main centers to host pre-exascale computers co-funded by the European Union.



MareNostrum 5

circa 220 Million Euros budget

Pre-Exascale system

Including new technologies

MareNostrum 1

2004 – 42,3 Tflops

1st Europe / 4th World

New technologies

MareNostrum 2

2006 – 94,2 Tflops

1st Europe / 5th World

New technologies

MareNostrum 3

2011 – 1,1 Pflops

1st Europe / 36th World

New technologies

MareNostrum 4

2017 – 11,1 Pflops

2nd Europe / 13th World

New technologies

- 3 “
- 5 P

European Commis
Digital Single I
supercompute

Luxembourg, 7 June

EUROPEAN **PROCESSOR** INITIATIVE

PROJECT PILLARS:

- 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

EPI OBJECTIVES

- Architect 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)
- Implement a Real-time acceleration PoC based on the first EPI GPP Processor (MPPA)
- Interfacing with the Automotive MCU
- Development efficient power management 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 KEY PERFORMANCE INDICATORS

- 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**

- 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 SGA I, focus on **programming models** to include accelerations.

EPI STREAMS

S1 - Common Stream

Codesign, Architecture, System software and key technologies for the Common Platform

S2 - GPP Processor

Design and implement the processor chip(s) and PoC system

S3 - Acceleration

Foster acceleration technologies and create building blocks

S4 - Automotive

Address automotive market needs and create a pilot eHPC system


S5 - Administration

Manage and support activities

EUROPEAN PROCESSOR INITIATIVE



www.european-processor-initiative.eu



PROJECT PILLARS

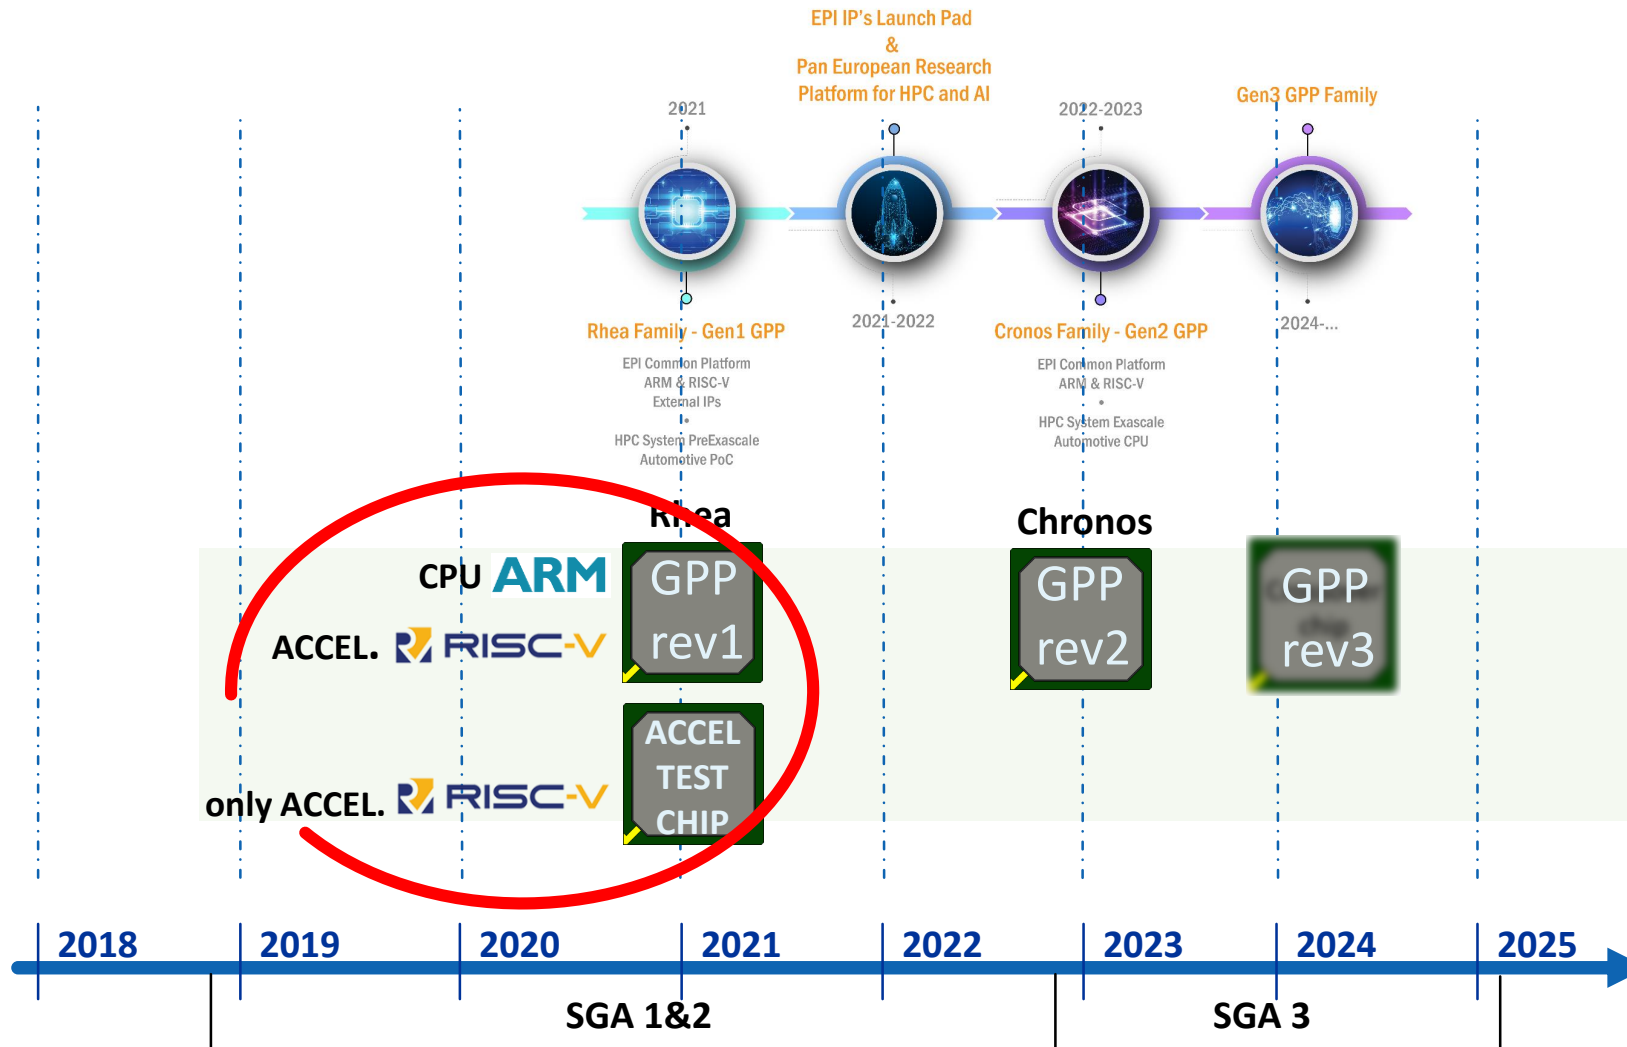
- Common platform and global architecture stream
- HPC general purpose processor stream
- Accelerator stream
- Automotive platform stream



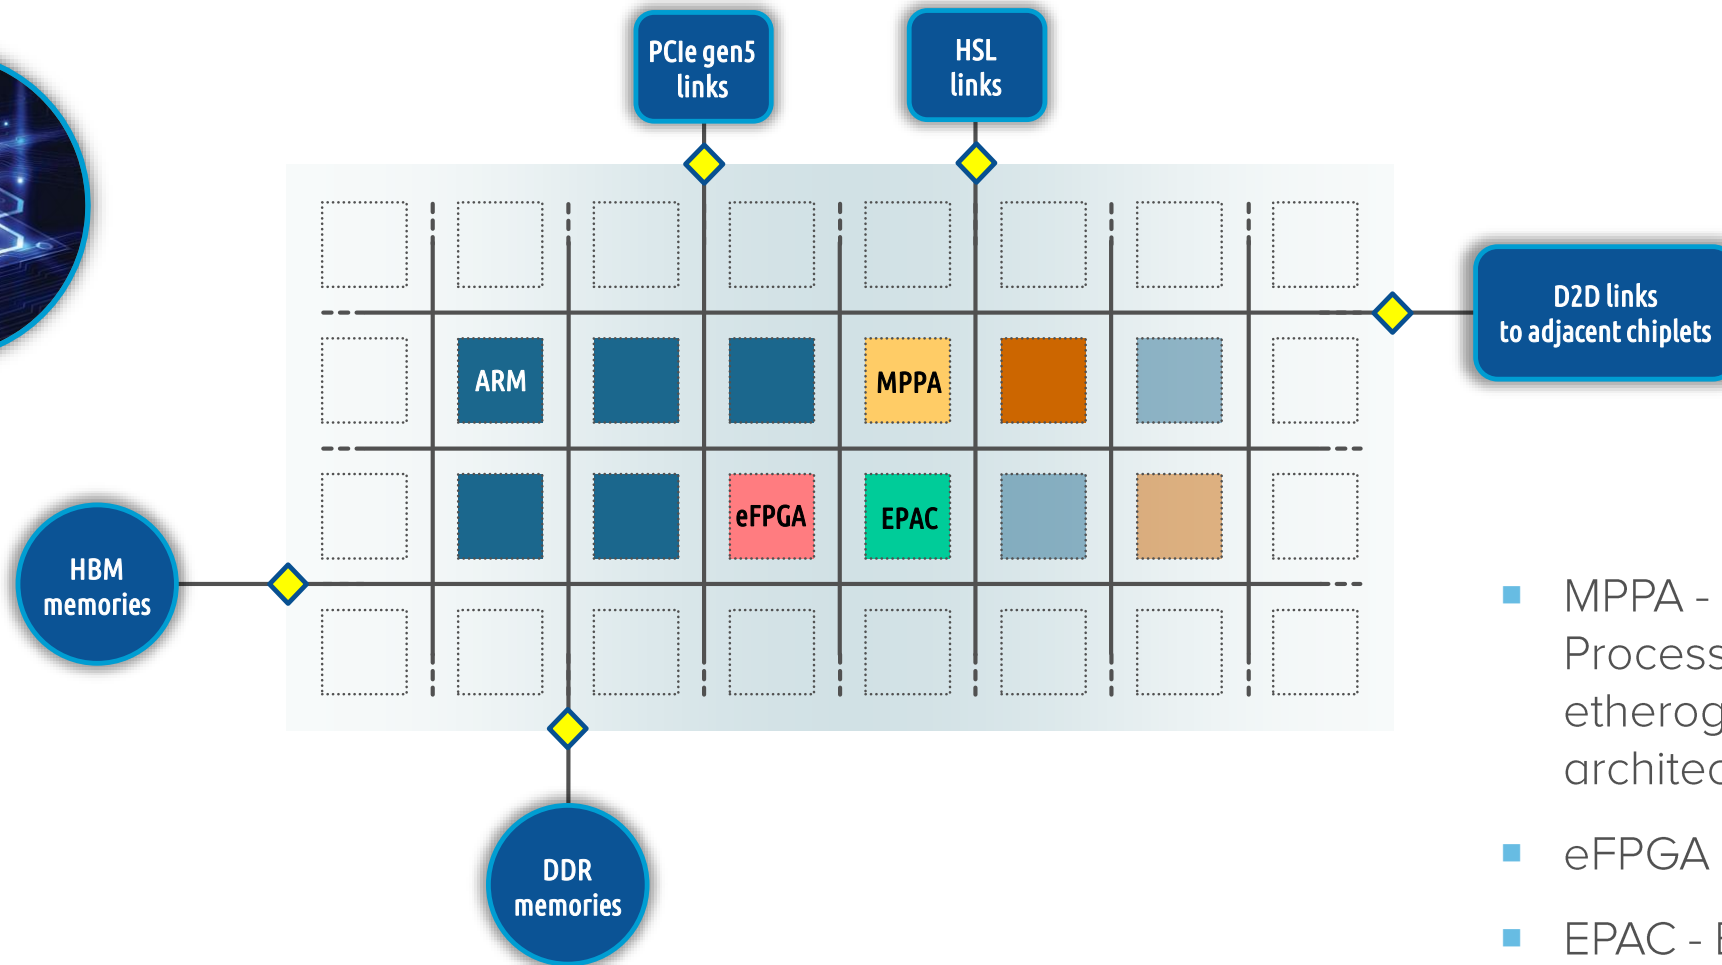
This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 826647



EPI ROADMAP

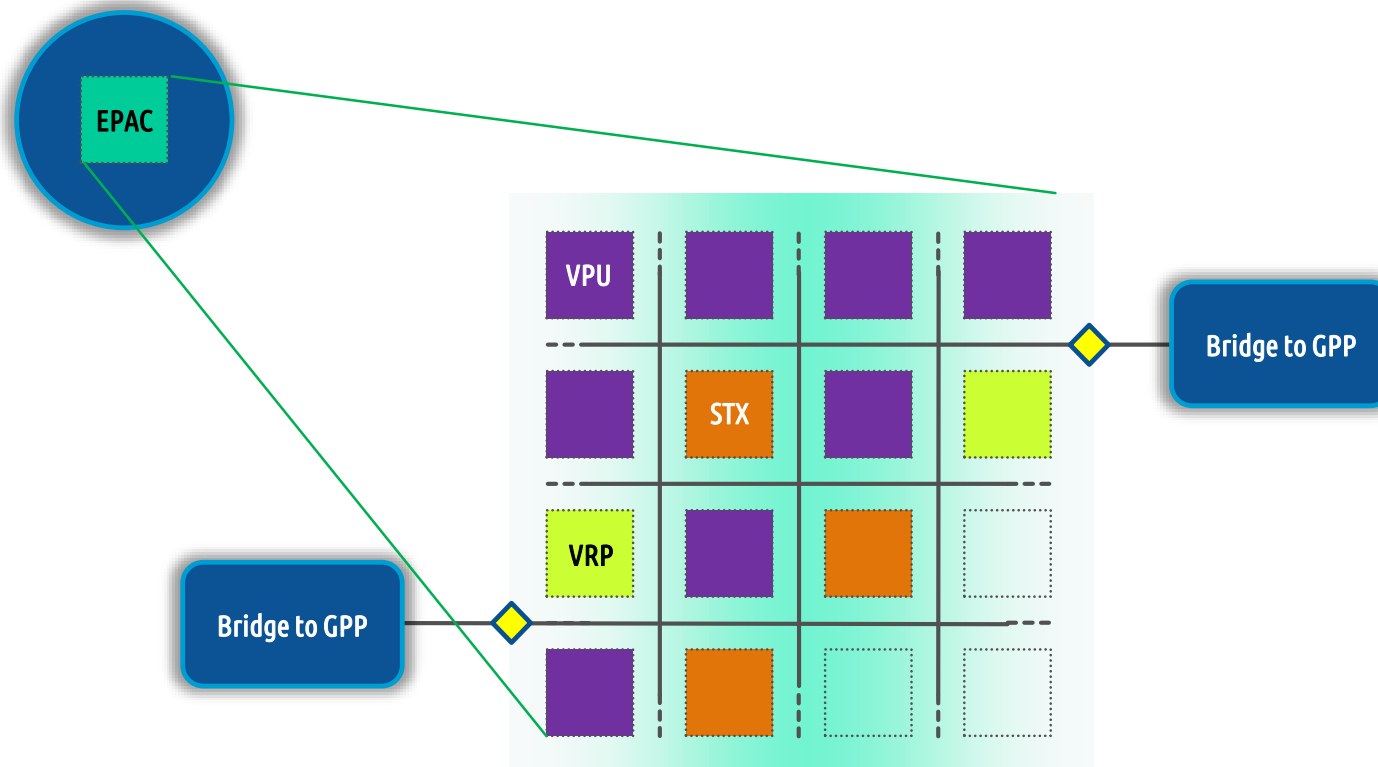


GPP AND COMMON ARCHITECTURE




- MPPA - Multi-Purpose Processing Array
heterogeneous architecture
- 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

1ST GENERATION EPI CHIPS



General Purpose Processor (GPP) chip

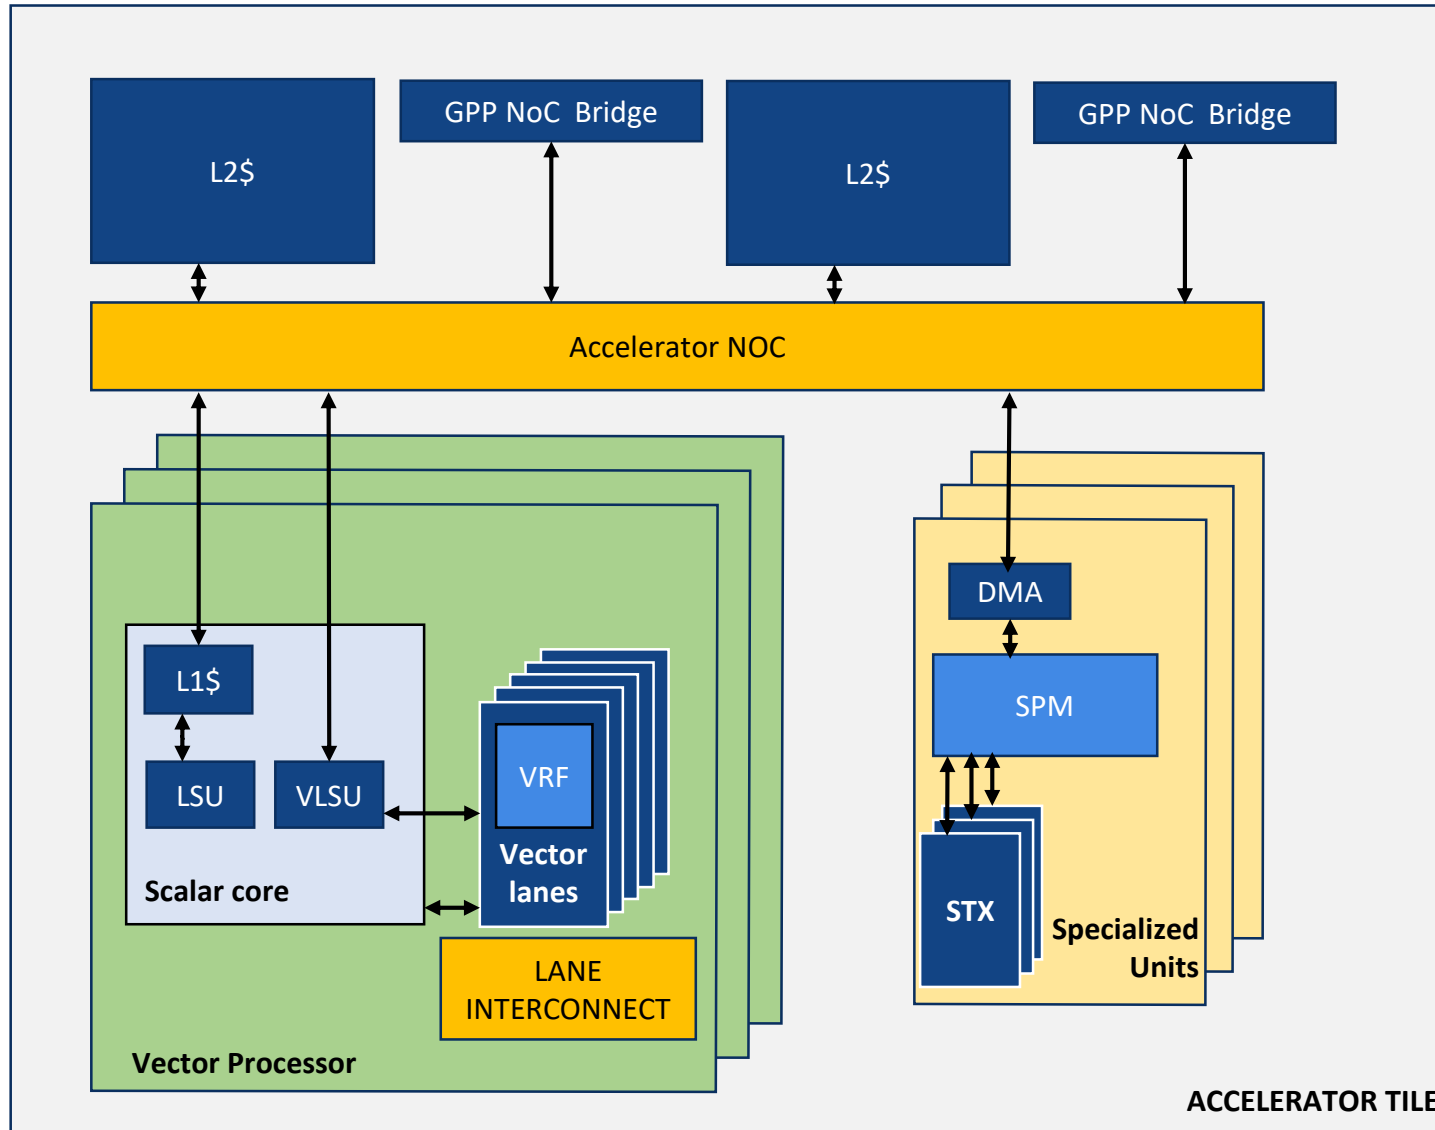
- 7 nm, chip-let technology
- ARM-SVE tiles
- EPAC RISC-V vector+AI accelerator tiles
- L1, L2, L3 cache subsystem + HBM + DDR



RISC-V Accelerator Demonstrator Test Chip

- 22 nm FDSOI
- Only one RISC-V accelerator tile
- On-chip L1, L2 + off-chip HBM + DDR PHY
- Targets 128 DP GFLOPS (vector processor only)

EPAC ARCHITECTURE VIEW



- Up to 8 vector processors per tile
- The Vector Lanes act as tightly coupled (ISA mapped) acceleration units to the scalar core in the vector processor
- Heavily pipelined
- RISC-V vector extension compliant
- Up to 8 Specialized Units per tile
- The STX Units act as loosely coupled, memory mapped acceleration units to the scalar cores
- Fast single-cycle MACs in parallel
- Shared L2 cache banks
- Cache coherent NoC


LESSONS LEARNED AND HINTS

- Never forget the ETERNAL LAWS

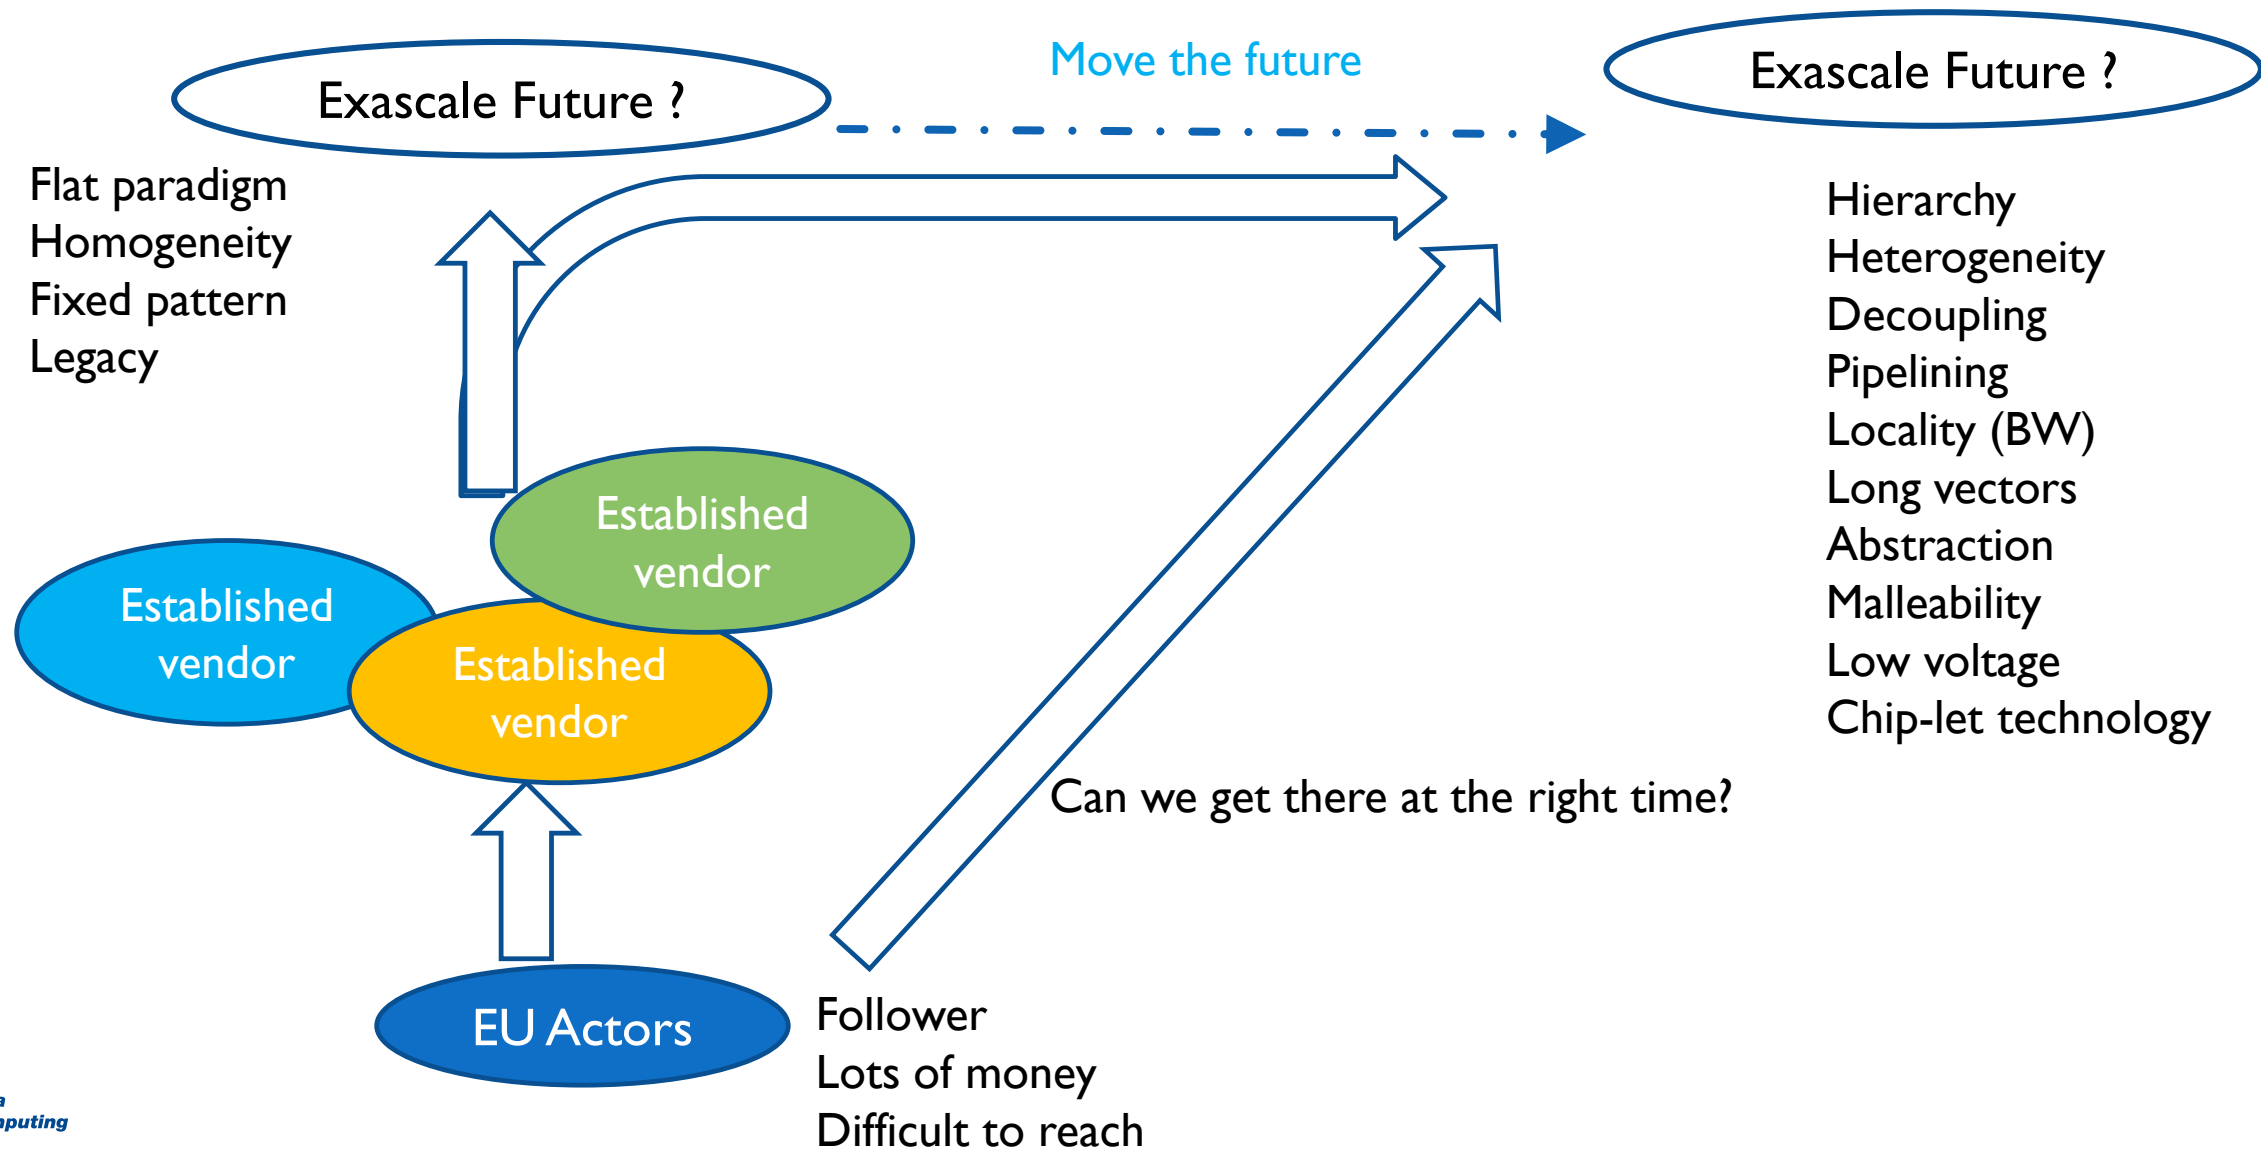
Execution time: $T_{EXEC} = N \cdot CPI \cdot T_{CK}$

Power (dynamic): $P = C_{eff} \cdot V^2 \cdot (1/T_{CK})$

- Forget you can design application, runtime environment, architecture and micro-architecture as independent things
- Future **sure** challenges:
 - Power efficiency
 - Usability/adaptability
 - Fault tolerance
 - Security




BSC'S VISION OF THE EXASCALE TRANSITION



ABOUT THE INSTRUCTION SET FOR THE ACCELERATOR

- In 2015, Mateo Valero said he believed a European Supercomputer based on ARM was possible (Mont-Blanc).
- Even though ARM is no longer European, it can (must) form part of the short-term solution
- The fastest-growing movement in computing at the moment is Open-Source and is called RISC-V
- The future is Open and RISC-V is democratizing chip-design
- Very-high-performance general purpose RISC-V CPUs are not available yet. But the time is right to develop powerful RISC-V accelerator processors.





THE ADVANTAGES....

- RISC-V has no legacy constraints
- RISC-V has (standard and non-standard) extensions
- RISC-V extensions facilitate decoupling from hardware (avoid low level accelerator control, e.g. memory mapped)
- RISC-V is contributed by committees of experts doing a great job

THE LIMITATIONS....

- Many people's contributions = many different requirements
- Some features may result less important to some members
- The eco-system issue (but smaller than – for example – in consumer market)
- The political issue question

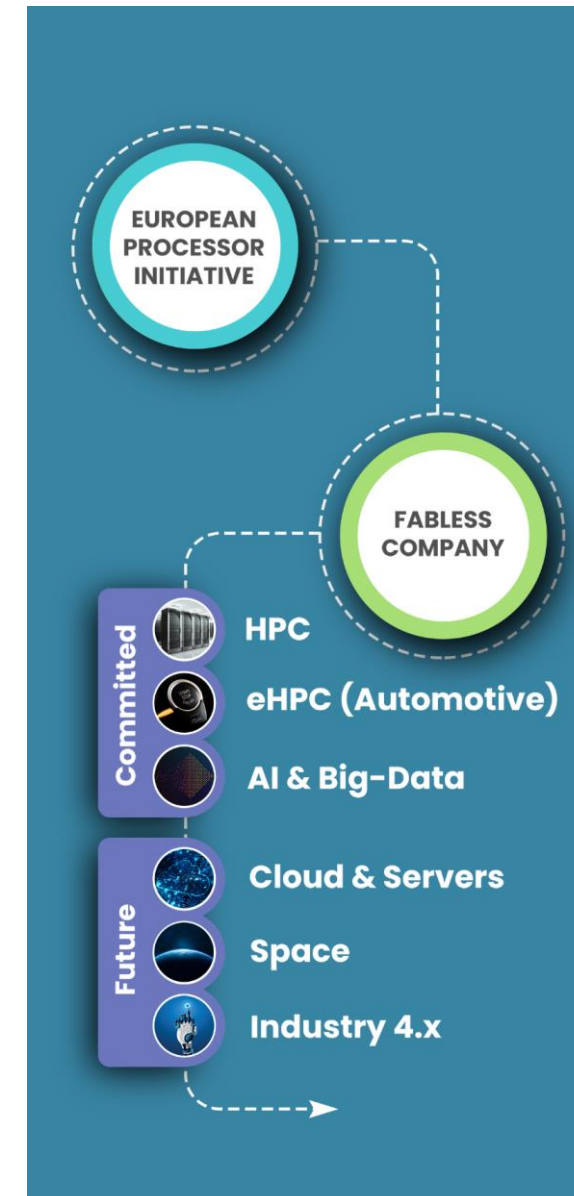
EPI AUTOMOTIVE

- 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
- Other issues to be faced: safety standards

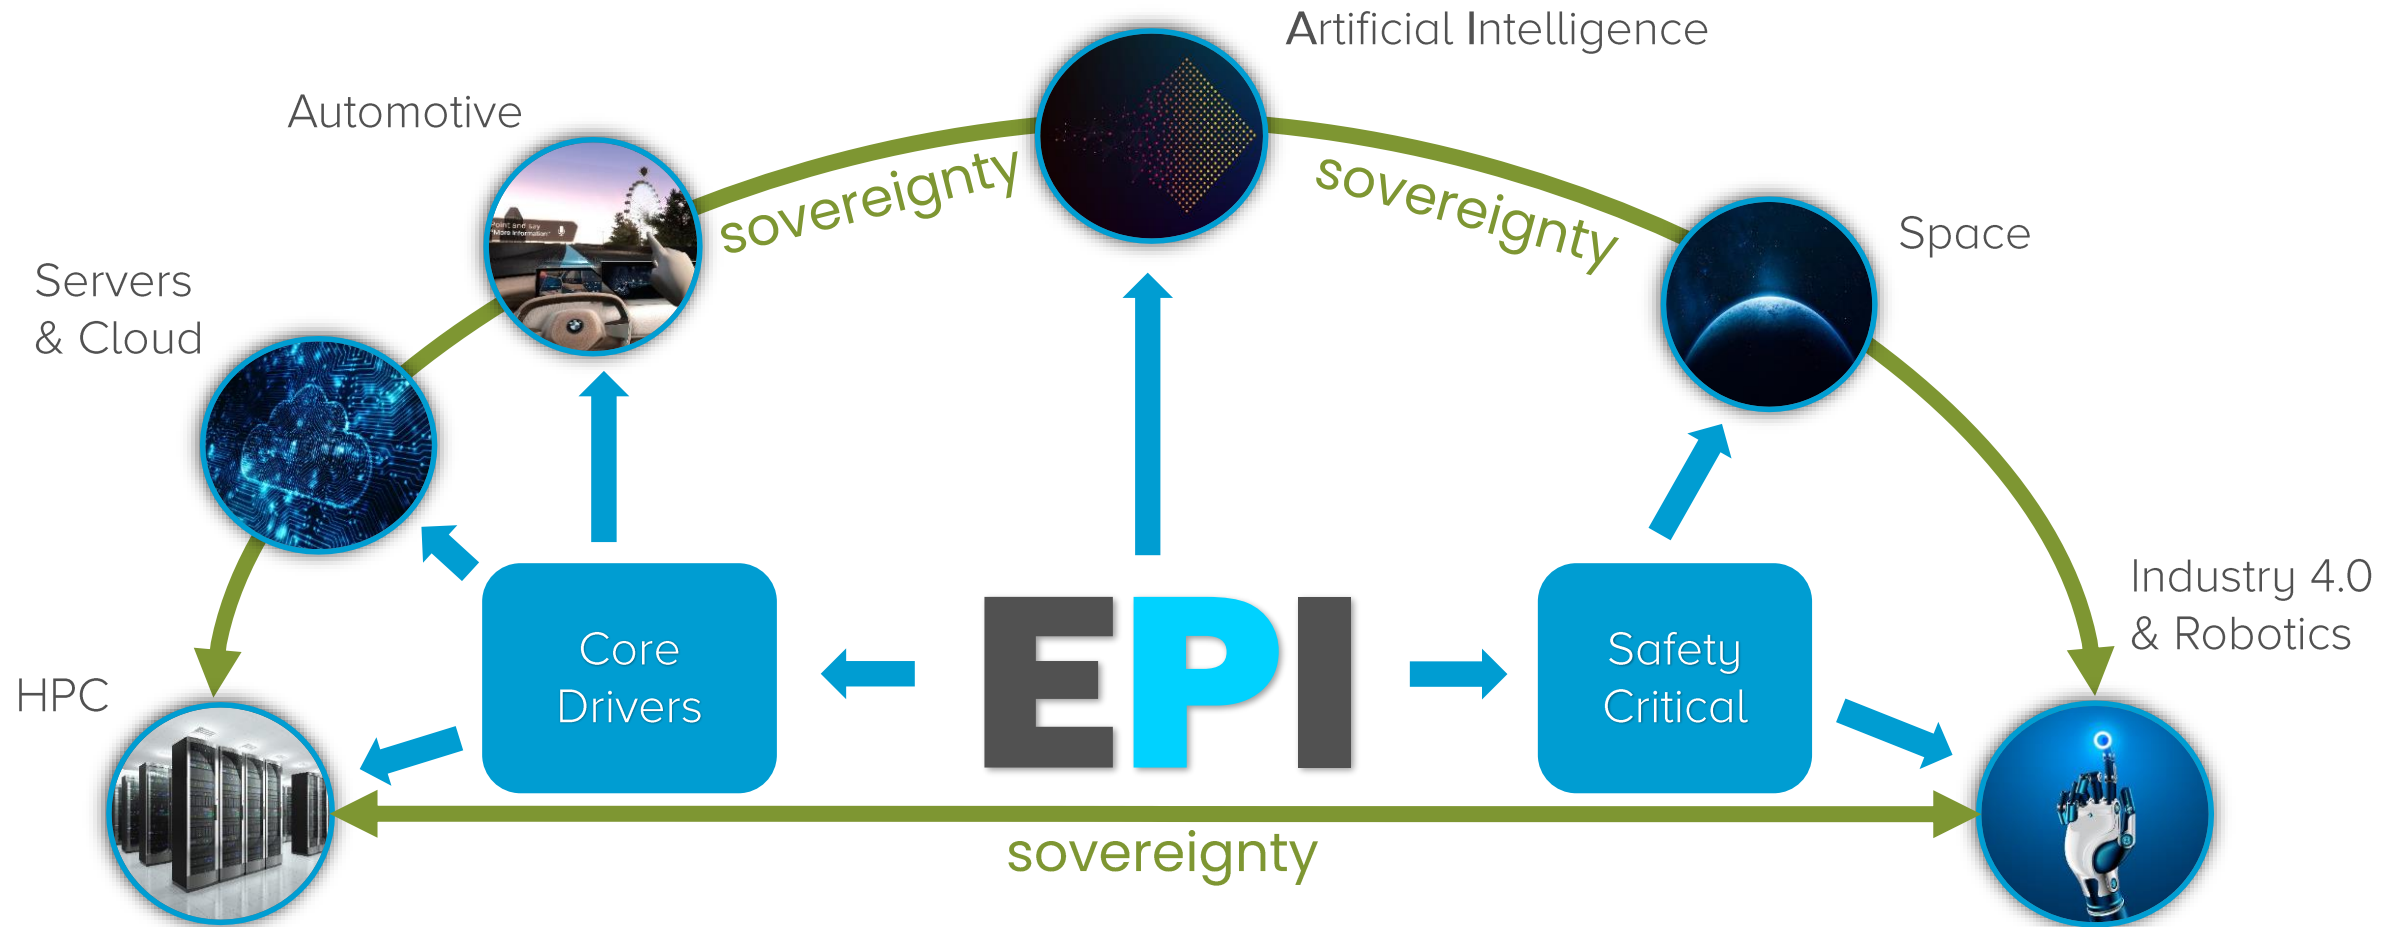


EPI FABLESS COMPANY

- 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

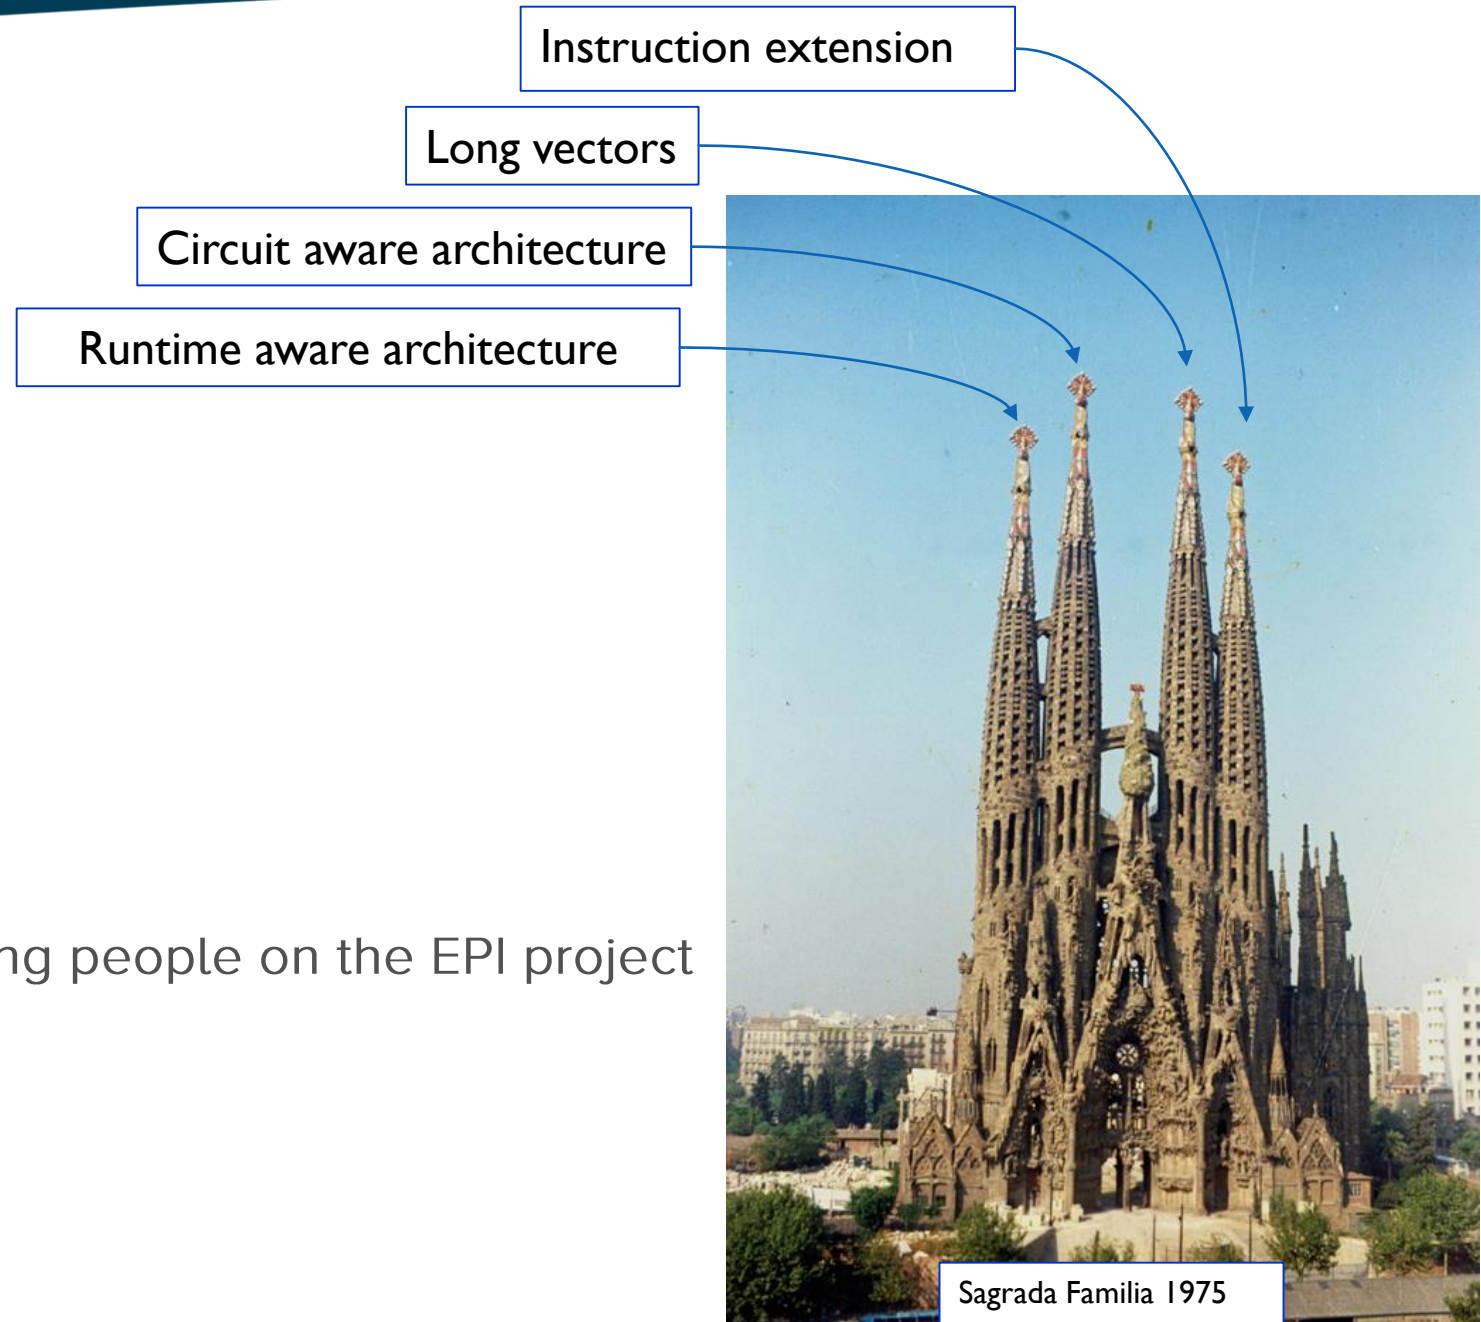


SCALABILITY ALLOWS WIDE MARKET POTENTIAL COVERAGE



EUROPEAN MICROPROCESSOR RESEARCH AND THE *SAGRADA FAMILIA*

- Do something special ...
- ... showing the way ...
- ... sustaining the effort
- BSC is hiring talented young people on the EPI project



Special thanks to

- Mateo Valero
- Jesus Labarta
- Adrian Cristal Kestelman
- Osman Unsal
- Peter Hsu
- Luca Benini
- Ying Chih Yang
- Philippe Notton
- Denis Dutoit
- Eugene Griffiths
- Fabrizio Gagliardi

Thank you for attending

mauro.olivieri@bsc.es

mauro.olivieri@uniroma1.it