Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

WHY HPC AND EXASCALE COMPUTING FOR NUMERICAL MODELLING: THE CLIMATE AND WEATHER PERSPECTIVE

Mario C. Acosta

17/06/2025

SHERIC25, Barcelona

Earth Sciences Department (BSC)

Develop and implement **environmental modelling** and **forecasting** using process-based and AI models, focusing on **weather**, **climate**, and **air quality**

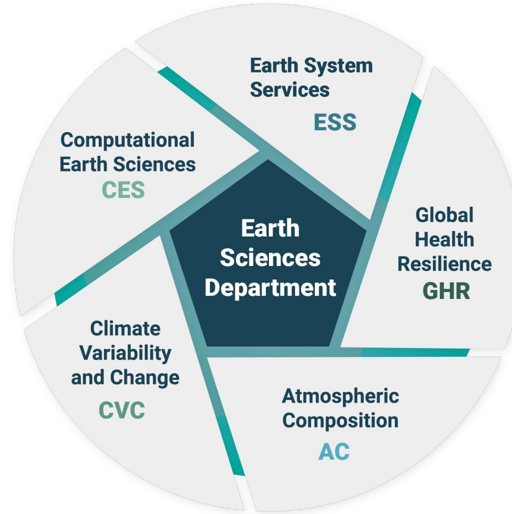

Transferring solutions to support the main **societal & environmental challenges** through service development.

- **210 people**
- **Funding from Horizon programmes, Copernicus and DestinE, private sector, ESA, Spanish, and regional governments**
- **Close link to local universities**

Ocean and Atmosphere 10 km simulation with the EC-Earth model performed by Oriol Tintó (BS

CES-Performance Team & ES Department

- Knowledge about the mathematical and computational side of Earth System Applications
- Knowledge about the specific needs in HPC of the Earth System Applications
- Researching about HPC methods specifically used for Earth System Applications

Who we are

- The refactoring of Earth System Models
  - Computational performance analysis and new optimizations for our numerical models.
  - Studying new algorithms for the new generation of high performance platforms (path to exascale).
- We are collaborating with several institutions on different projects working together in the same direction

Weather vs Climate predictions

Weather prediction — temporal horizon → **around 10 days**

Climate prediction — temporal horizon → **from weeks to decades**

Why?

- Chaotic nature of the atmosphere limits the weather forecast horizon
- Some elements of the climate system change slower than the atmosphere

ocean

sea ice

soil moisture

The ocean is still predictable even when the atmosphere is not!

We can predict how the atmosphere might respond to the oceans.

Weather vs Climate predictions

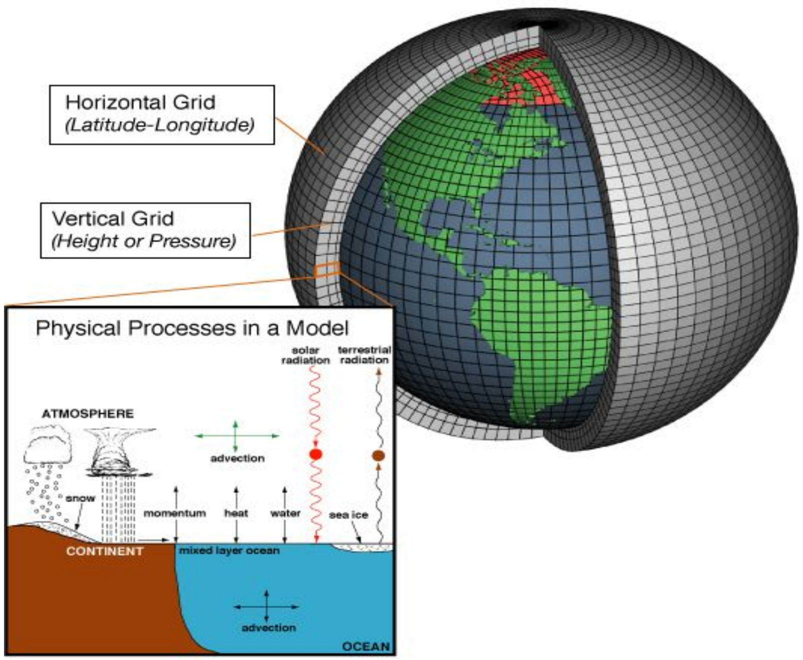

Coupled climate model

(atmosphere + land + ocean + sea ice)

initialised with current observations

Predictions for the next few week/ season

- ***What we can NOT expect from climate predictions***

  The temperature in Albacete on 27th February

- ✔ ***What we can expect from climate predictions***

  How likely next winter is going to be colder/warmer than normal

Temporal Scales

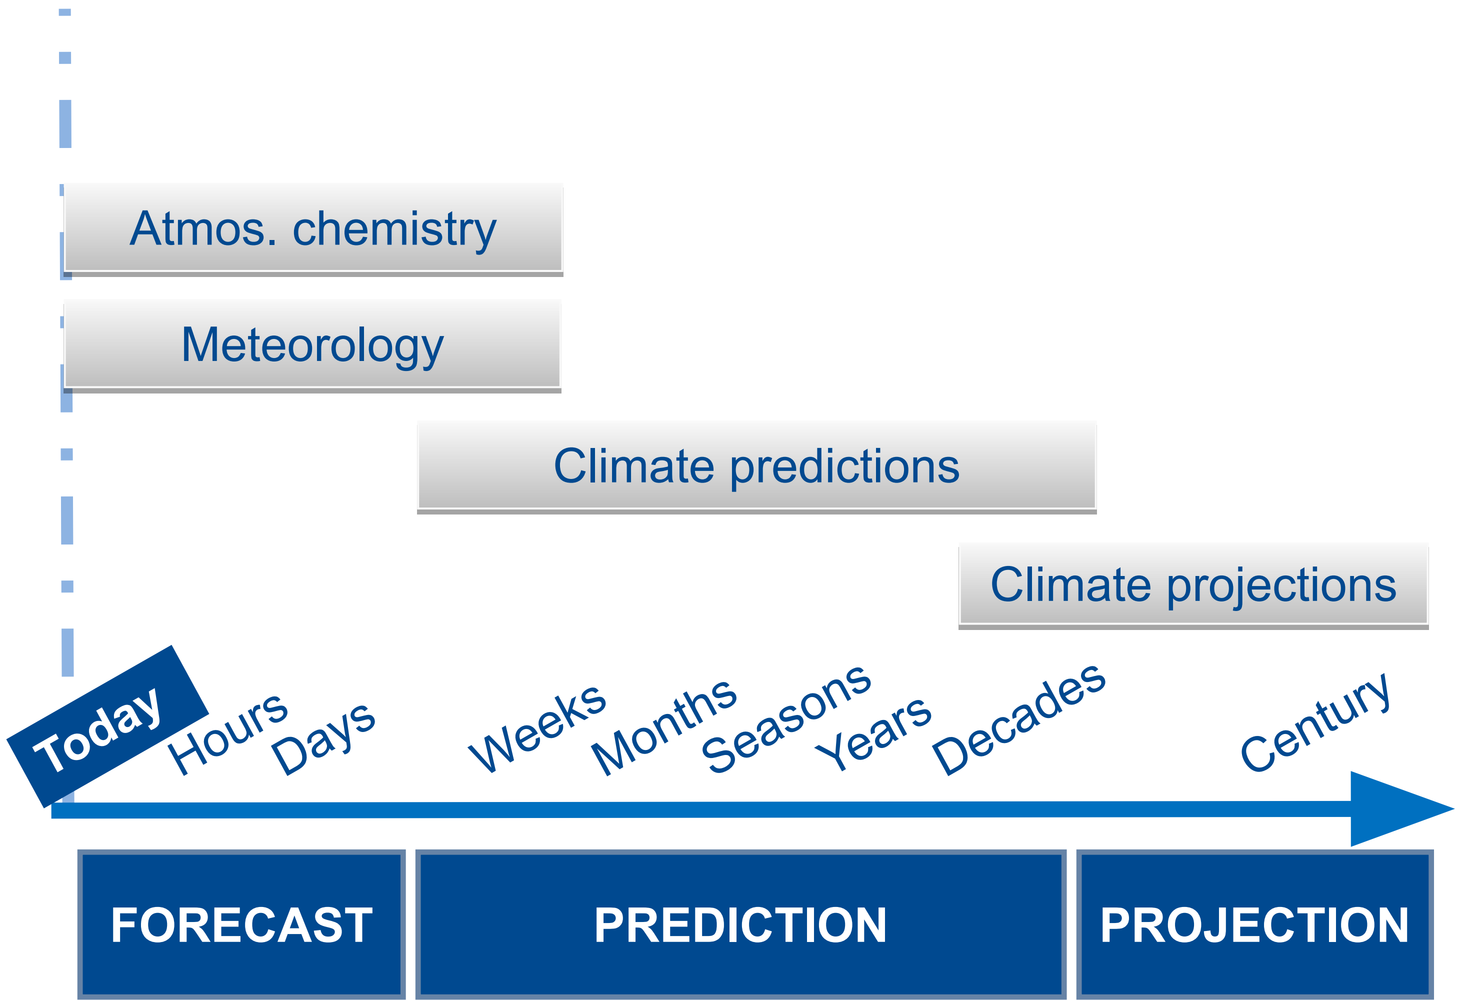

Introduction: Climate overview

- Projections
- Impact analysis
- Adaptation to climate change.

Currently, **only computational models have the potential** to provide geographically and physically consistent estimates.

Introduction: Weather overview

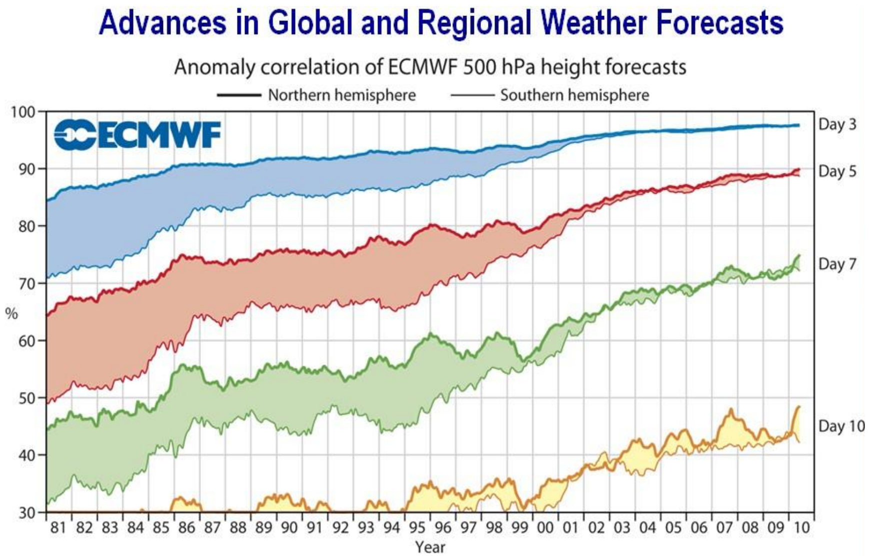

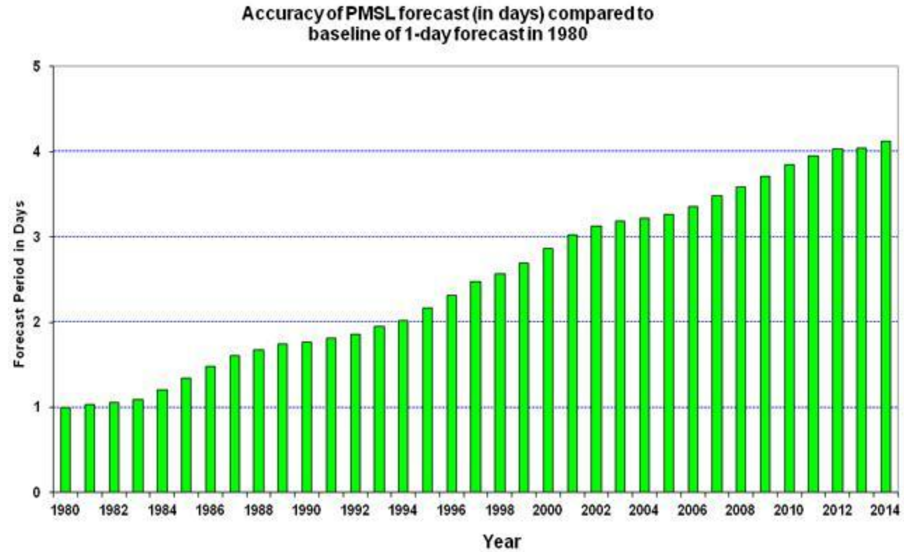

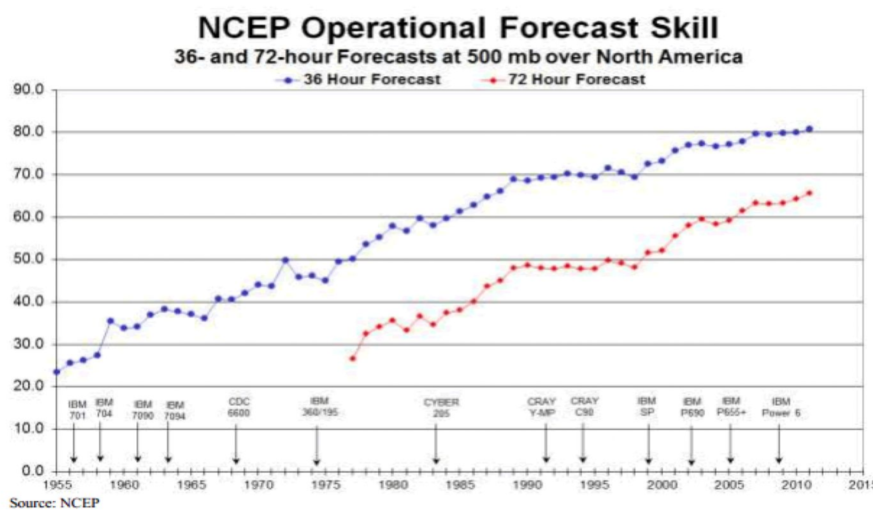

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Introduction

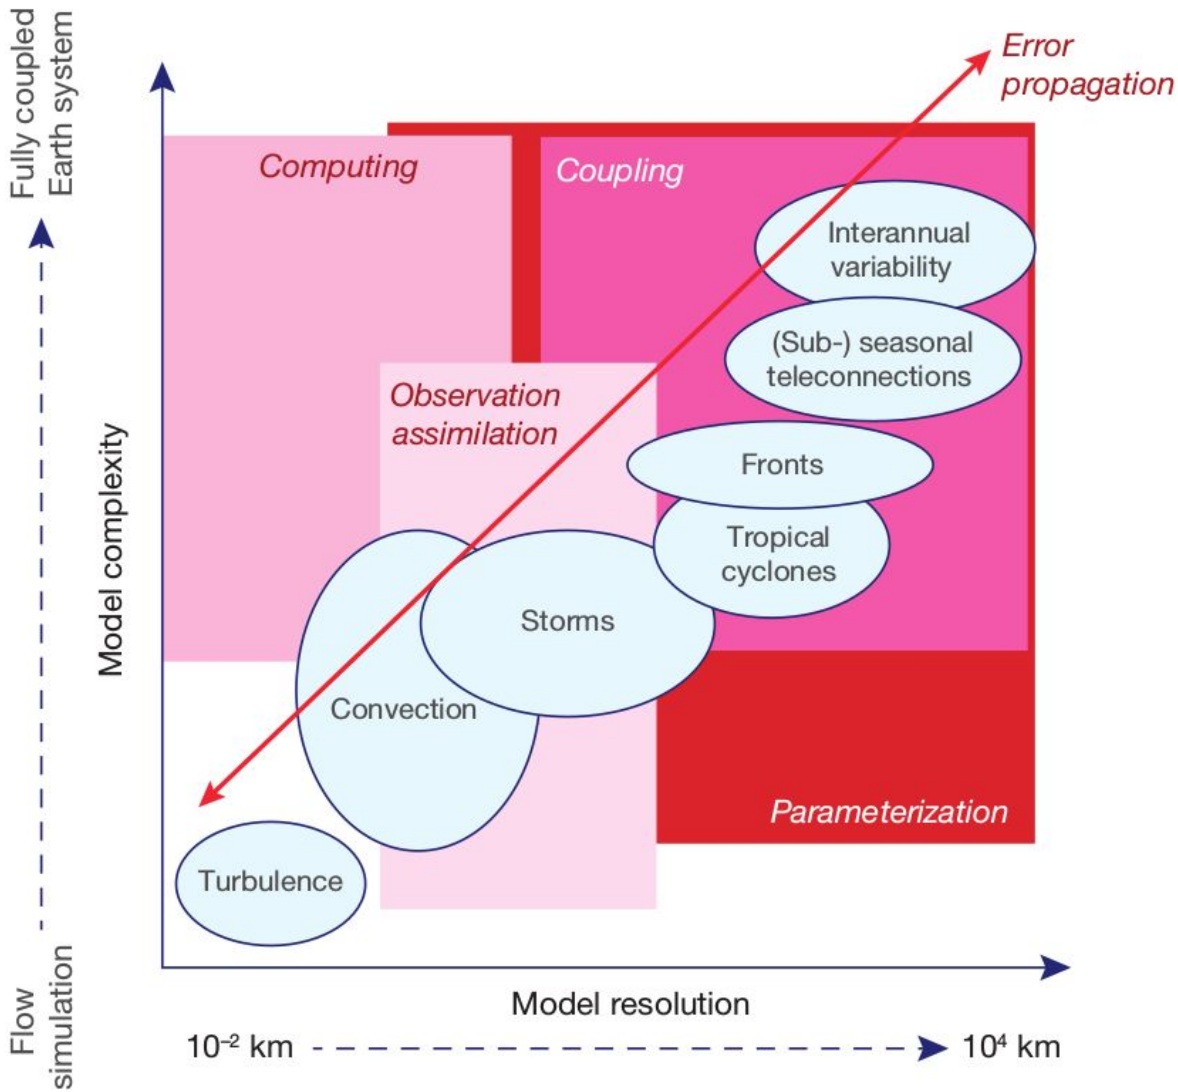

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Why HPC and supercomputers in ocean and climate modelling?

Why HPC and supercomputers

Increasing complexity of ocean and climate models

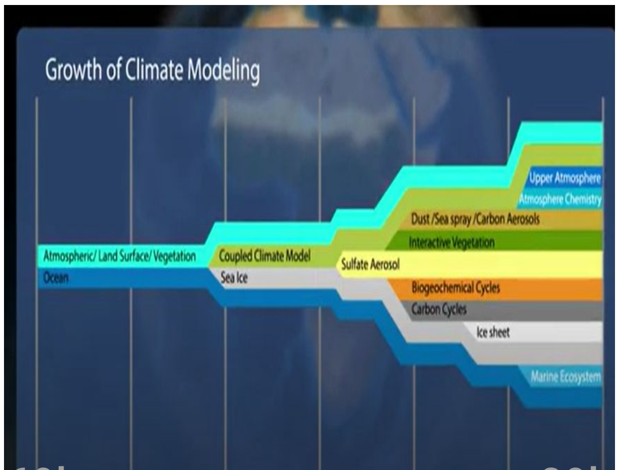

60's 20's

Why HPC and supercomputers

Increasing complexity of ocean and climate models

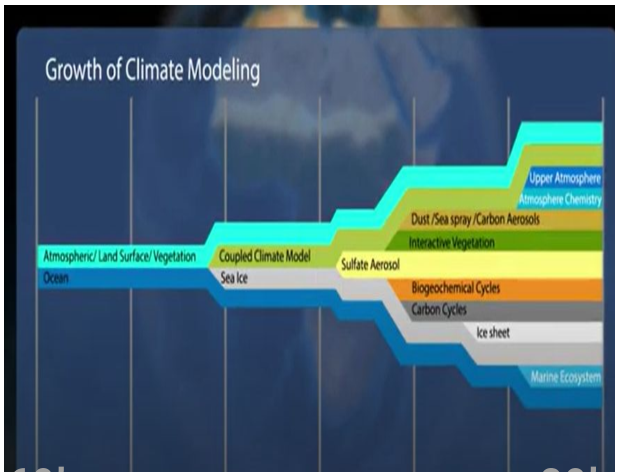

60's 20's

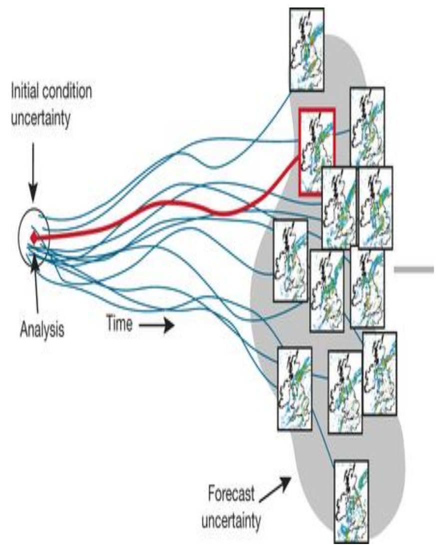

Why HPC and supercomputers

Increasing complexity of ocean and climate models

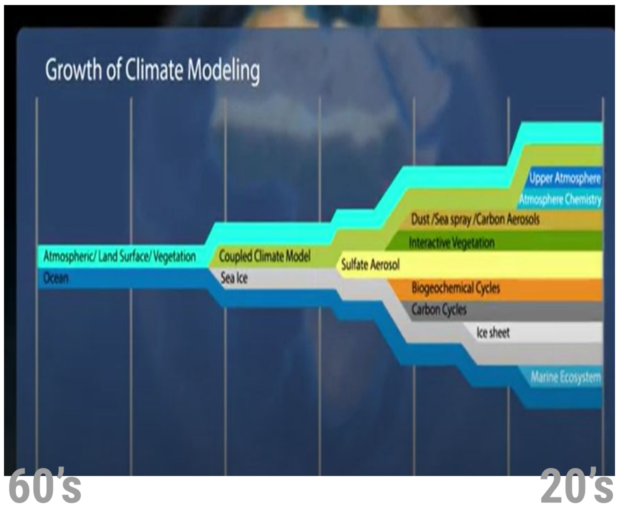

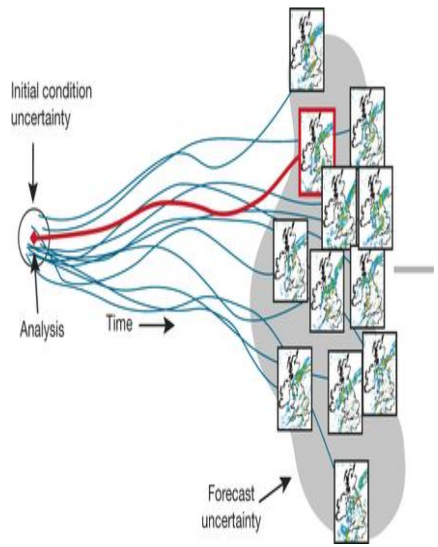

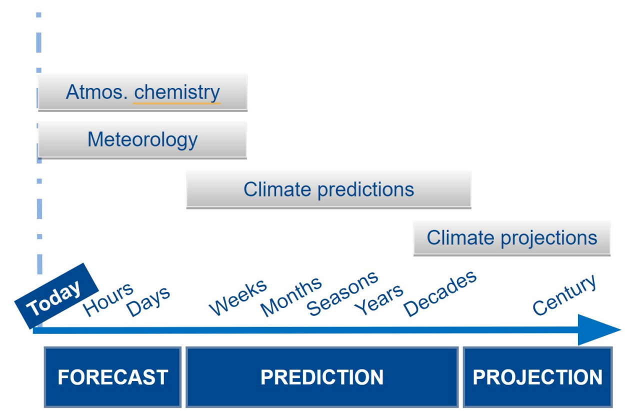

Why HPC and supercomputers

Ocean currents from a 1.25 km simulation with the NEMO model on the LUMI supercomputer

Breaking barriers for decision making

Some limitations for the use of environmental data in different socio-economic sectors.

- Lack of awareness
- Difficult interpretation
- Lack of expert synthesis

Breaking barriers for decision making

Some limitations for the use of environmental data in different socio-economic sectors.

Lack of awareness

Difficult interpretation

Lack of expert synthesis

Possible solution: to distil the information from existing sources to be integrated in decision-making.

Environmental services

- **Goal:** the development and incorporation of environmental data for planning, policy-making and practice at the global, regional and national scale.

  Implementation method: co-production and co-design.
-

BSC Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Computational Performance ESMs

- Especially critical in Earth science models.
- Simulations use a **huge amount of computational resources**.
- Future simulations will need much more resources.
- Mostly of Earth science models use parallelization based on spatial domain decomposition
- Mostly calculations are independent but some interaction is needed among subdomains
- Take advantage of the specific hardware is mandatory
- Optimizations techniques are needed to increase the performance of these models, this is know as High Performance Computing

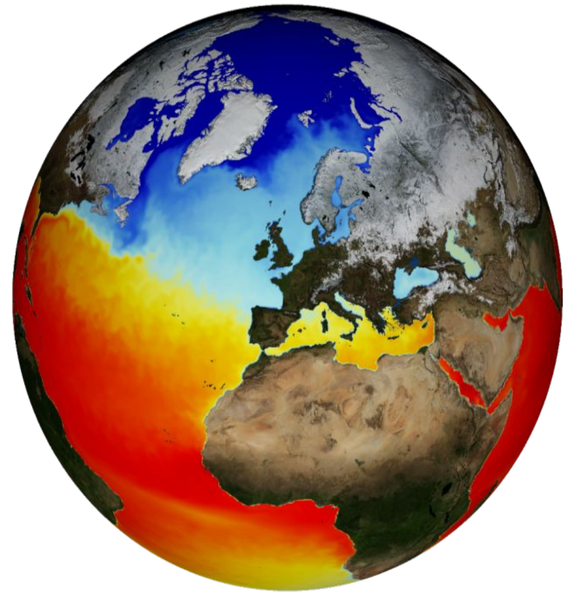

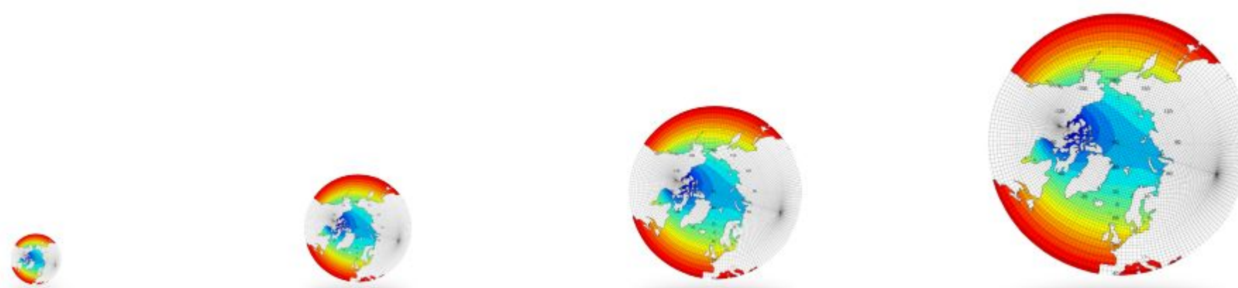

	ORCA 2	ORCA ¼	ORCA 1/12	ORCA 1/36
Grid resolution →	ORCA 2	ORCA ¼	ORCA 1/12	ORCA 1/36
Mem needed →	550 MB of memory	47 Gigabytes of memory	414 Gigabytes of memory	> 1 Terabytes of memory
Simulation time →	8 CPU hours	3500 CPU hours	90 000 CPU hours	~4 000 000 CPU hours
Outputs size →	10 Gigabytes of output (daily)	120 Gigabytes of output (daily)	1 Terabyte of output (daily)	> 5 Terabytes of output (daily)

Computational Performance ESMs

- Especially critical in Earth science models.
- Simulations use a **huge amount of computational resources**.
- Future simulations will need much more resources.
- Mostly of Earth science models use parallelization based on spatial domain decomposition
- Mostly calculations are independent but some interaction is needed among subdomains
- Earth Science Models are very complex
- Take advantage of the specific hardware is mandatory
- Optimizations techniques are needed to increase the performance of these models, this is know as High Performance Computing

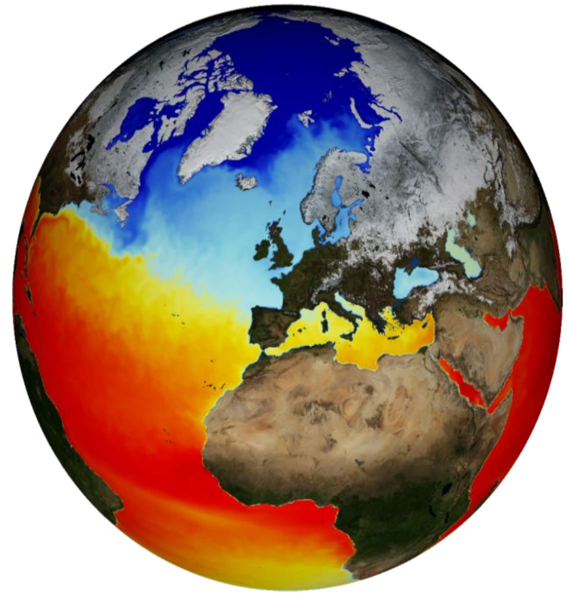

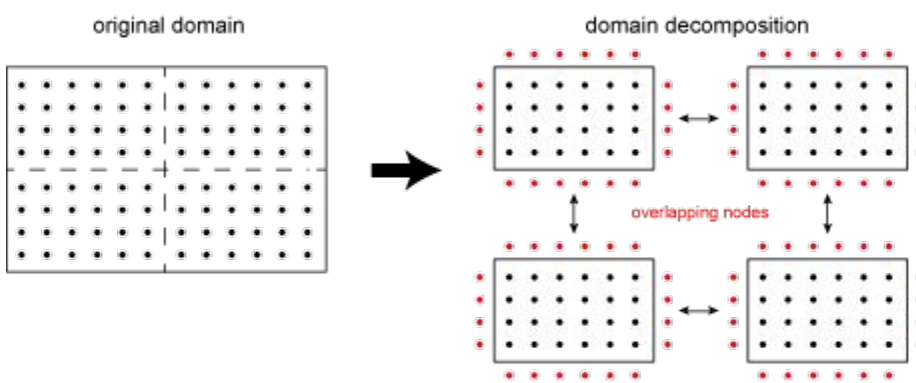

Computational Performance ESMs

- Especially critical in Earth science models.
- Simulations use a **huge amount of computational resources**.
- Future simulations will need much more resources.
- Mostly of Earth science models use parallelization based on spatial domain decomposition
- Mostly calculations are independent but some interaction is needed among subdomains
- Earth Science Models are very complex
- Take advantage of the specific hardware is mandatory
- Optimizations techniques are needed to increase the performance of these models, this is know as High Performance Computing

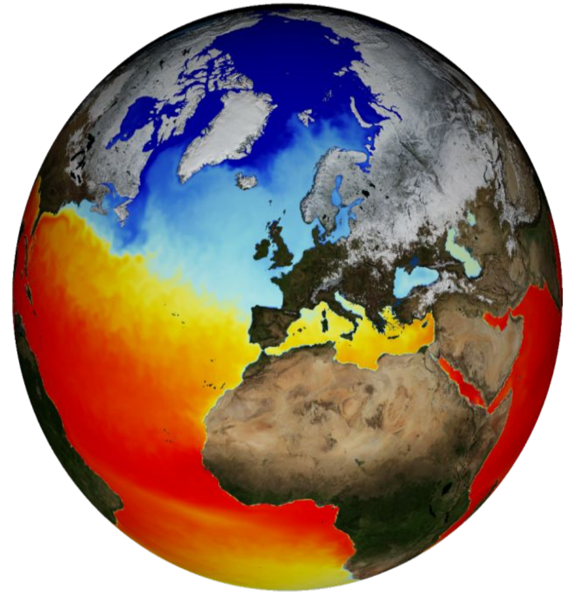

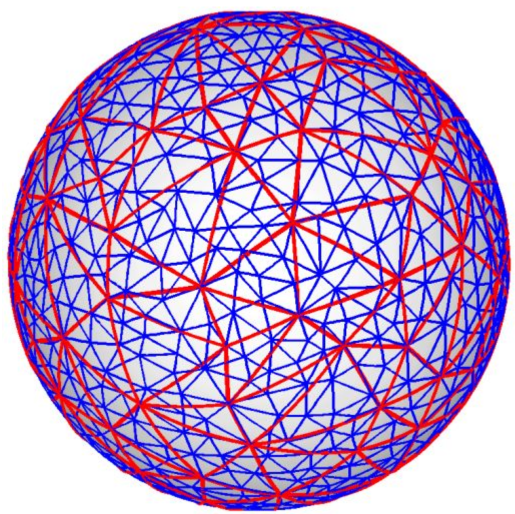

Introduction

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Introduction

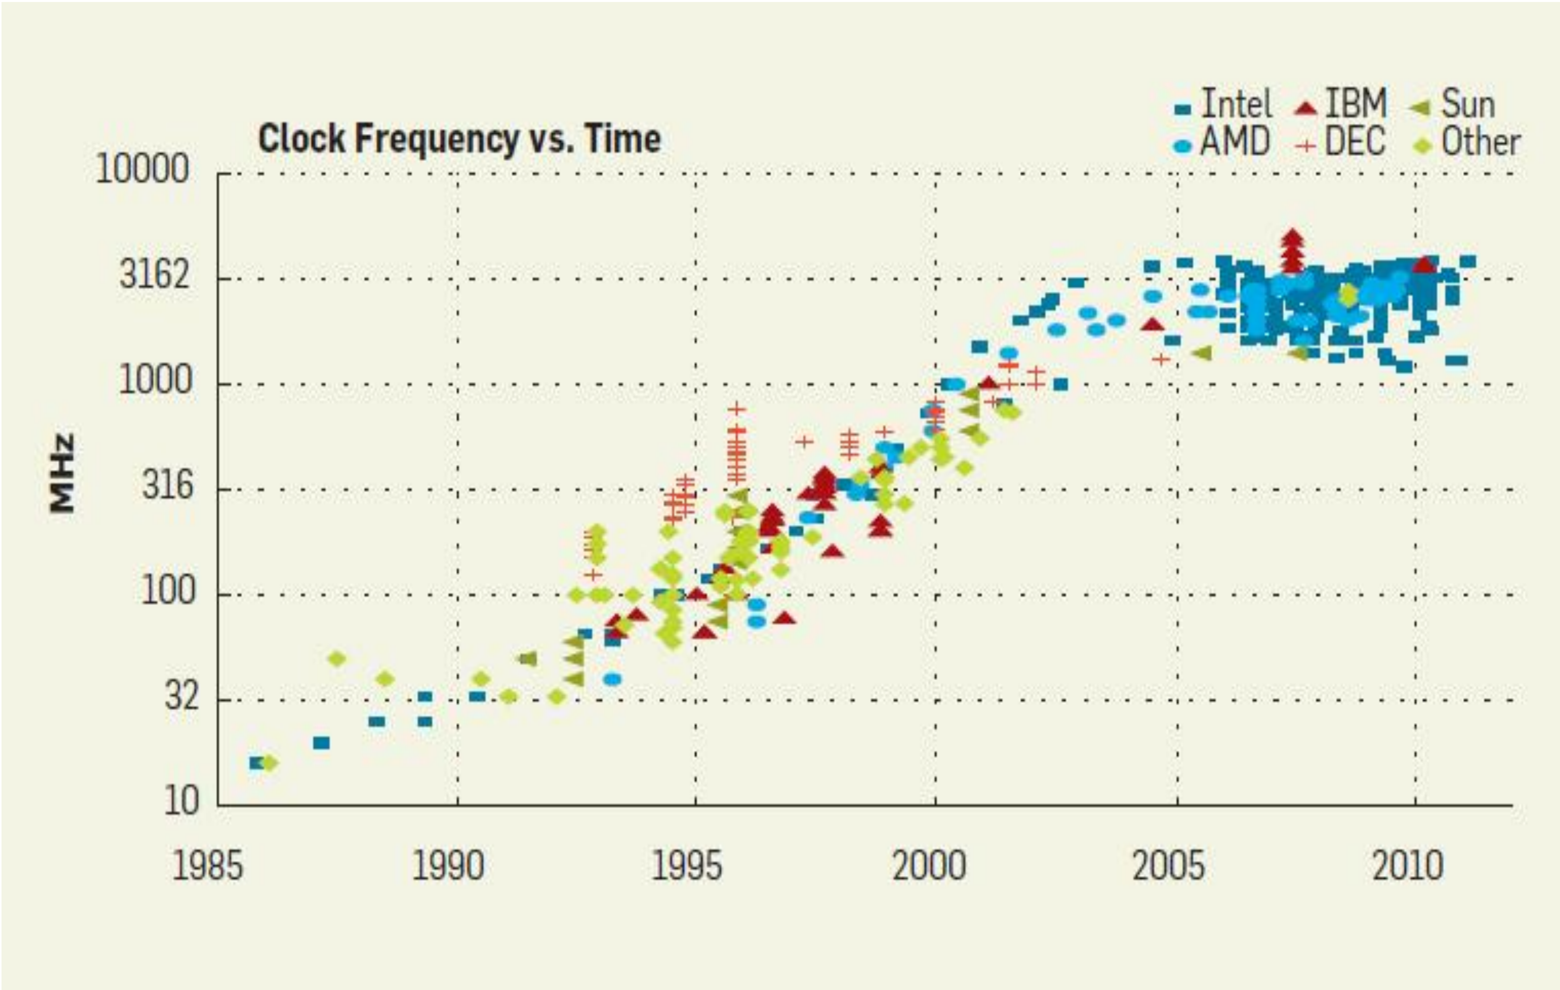

HPC history

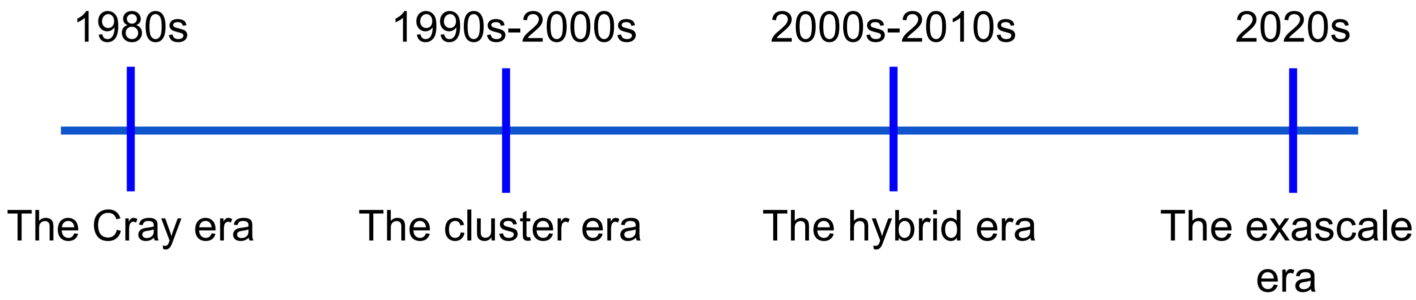

- High performance computing has always pushed the limits and designs of computing technology.
- HPC architectures have gone through rapid changes to facilitate the increasing computational demand of scientific applications in many areas such as Earth Science.

HPC history

1980s	1990s-2000s	2000s-2010s	2020s
The Cray era	The cluster era	The hybrid era	The exascale era

- Supercomputers composed of vector processors such as CRAY-1 and -2 dominated computing fields
- These computations focused on floating point operations per second (FLOPS)
- The supercomputers had central processing units (CPU) that implemented an instruction set designed to operate efficiently and effectively on large one-dimensional arrays of data called vectors.

HPC history

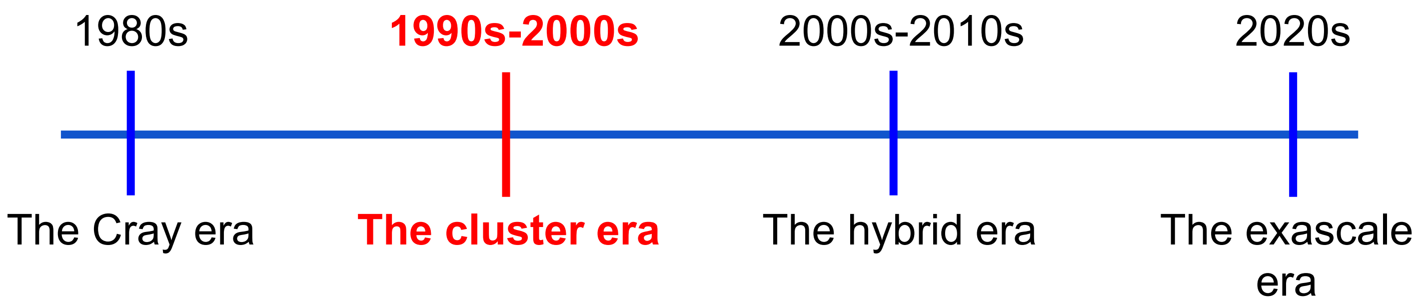

- Multi- and many-core CPUs approach.
- Break the technological limitation of a higher operating clock frequency and a higher density of large integrated circuits to align with Moore's Law.
- Solve other issues as the gap between processor and memory speed or the instruction level parallelism.
- This led to clusters with a large number of processors, each with a small number of core sharing RAM and some cache space.

HPC history

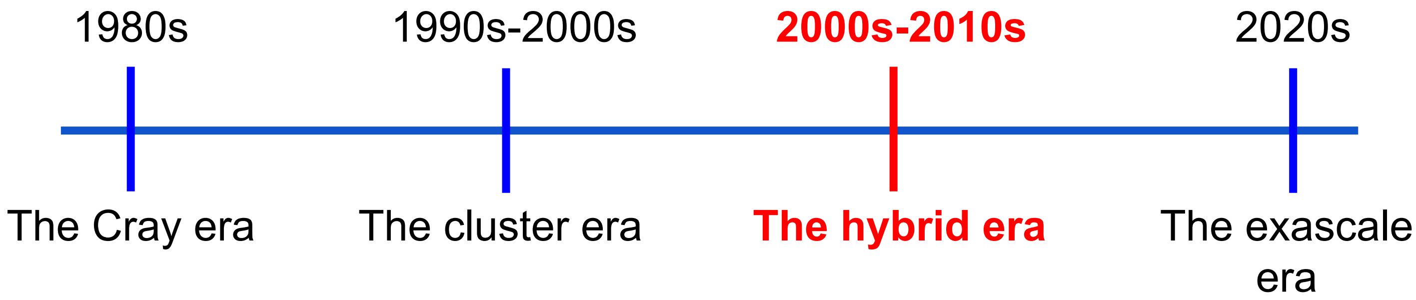

- As the numbers of cores increased, Graphics Processing Units were integrated into HPC clusters to accelerate the performance for many applications.
- These systems contain large numbers (hundreds to thousands) of small efficient cores (many-cores) that worked in unison.
- This approach essentially “offloaded” certain types of operations from the CPU to the GPU.

HPC history

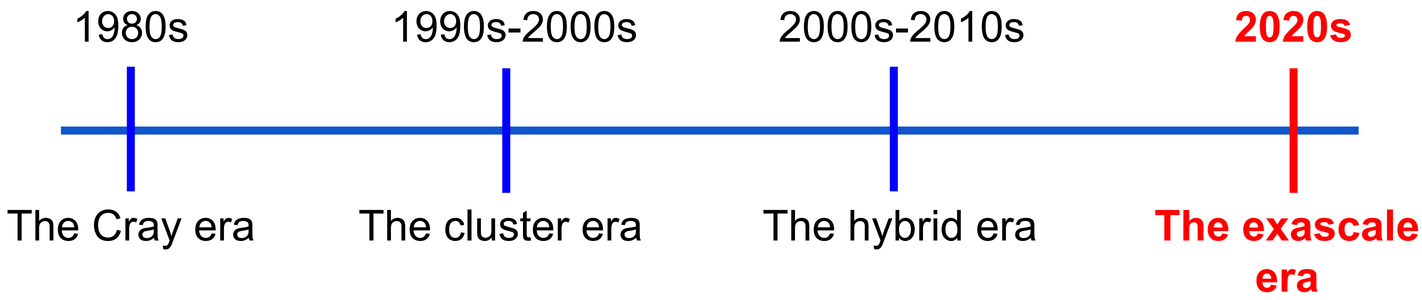

- Exascale machines are more purpose-built and rated using a specific application performance rather than the general Top500 High Performance Linpack benchmark.
- Extending multi-core/many-core clusters to the Exascale range is hampered by the disconnect between hardware and software.
- Heterogeneous computing and co-desing as solution
  - Fixed Accelerators provide order(s) of magnitude more specialized performance
  - The main problem could be complexity and programmability

The pre- and exascale path

- What we can gain
  - New computing elements (GPUs, FPGAs, AI, Quantum, RISCV)
  - More parallelism (Million threads)
  - Much more computing (Exaflops)
  - Data streaming
  - More data (Exabytes)
  - More complexity in our models (increased resolution, more parameters or components, more ensembles)
  - Larger Datasets

The pre- and exascale path

- What we can gain
  - New computing elements (GPUs, FPGAs, AI, Quantum, RISCV)
  - More paralleism (Million threads)
  - Much more computing (Exaflops)
  - Data streaming
  - More data (Exabytes)
  - More complexity in our models (increased resolution, more parameters or components, more ensembles)
  - Larger Datasets

Modeling and simulation + **Machine Learning** + **Big Data Analytics**

The pre- and exascale path

- What we can gain
  - New computing elements (GPUs, FPGAs, AI, Quantum, RISCV)
  - More paralleism (Million threads)
  - Much more computing (Exaflops)
  - Data streaming
  - More data (Exabytes)
  - More complexity in our models (increased resolution, more parameters or components, more ensembles)
  - Larger Datasets

Modeling and simulation + **Machine Learning** + **Big Data Analytics**

The pre- and exascale path

- In short time
  - New pre-exascale machines (LUMI, LEONARDO, Marenostrum5).
  - High-resolution "Digital Twins" using EuroHPC hardware
  - European Projects pursuing HPC improvements (ESiWACE3)
- In medium/long time
  - RISCV, Quantum, FPGAs
  - Clouding
  - Hardware and software acceleration (IA, GPUs)
  - …

Modeling and simulation

Barcelona
Supercomputing

Generalitat de Catalunya
Departament de Recerca
i Universitats

GOBIERNO DE ESPAÑA
MINISTERIO DE CIENCIA E INNOVACIÓN

UNIVERSITAT POLITÈCNICA DE CATALUNYA
BARCELONATECH

UNIÓN EUROPEA
Fondo Europeo de Desarrollo Regiona

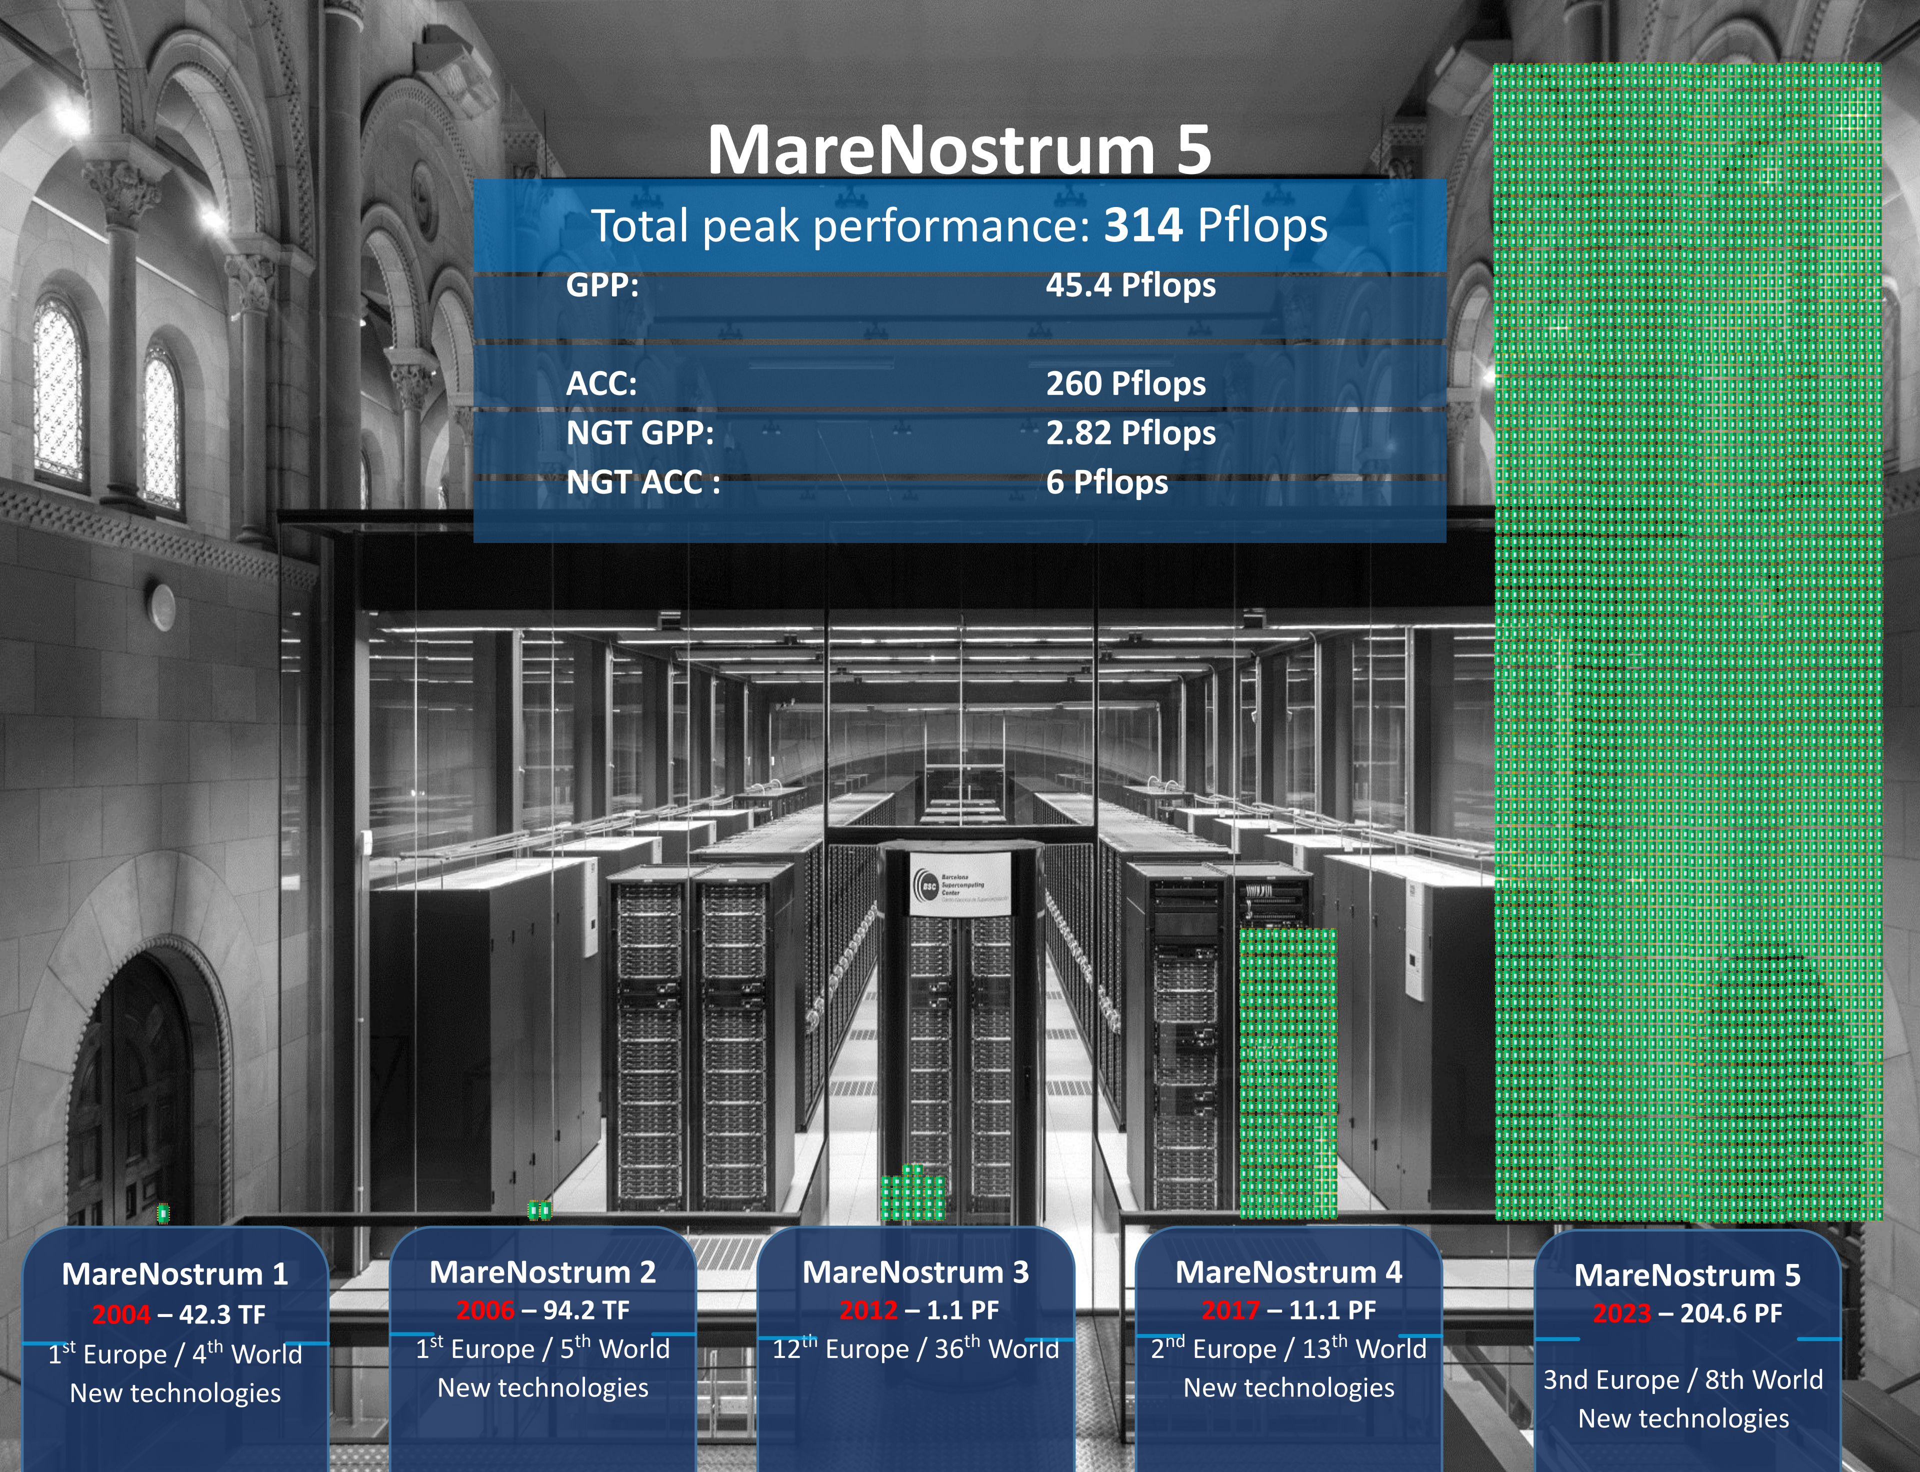
MareNostrum 5
Total peak performance: 314 Pflops
GPP: 45.4 Pflops
ACC: 260 Pflops
NGT GPP: 2.82 Pflops
NGT ACC : 6 Pflops
MareNostrum 1
2004 – 42.3 TF
1st Europe / 4th World
New technologies
MareNostrum 2
2006 – 94.2 TF
1st Europe / 5th World
New technologies
MareNostrum 3
2012 – 1.1 PF
12th Europe / 36th World
MareNostrum 4
2017 – 11.1 PF
2nd Europe / 13th World
New technologies
MareNostrum 5
2023 – 204.6 PF
3nd Europe / 8th World
New technologies

MareNostrum 5

Total net storage capacity: **650** PB

ESS 3500 data: HDD	**248 PB**
ESS 3500 performance NVME	**2.48 PB**
Archive Tape: Tape	**400 PB**

MareNostrum 1
2004 – 236 TB
1st Europe / 4th World
New technologies

MareNostrum 2
2006 – 460 TB
1st Europe / 5th World
New technologies

MareNostrum 3
2012 – 5.7 PB
12th Europe / 36th World

MareNostrum 4
2017 – 149.5 PB
2nd Europe / 13th World
New technologies

MareNostrum 5
2023 – 650 PB
3nd Europe / 8th World
New technologies

MareNostrum5 GPP

Racks	Cooling	Nodes		Processor/Accelerator		Memory	PFlops (HPL)	Local Drive	High-Perf. Network
		Total	per rack						
89	DLC +RDHX	6192	72 (6x6x2)	2x Intel Sapphire R. 8480+	56c @ 2GHz	>2GB/core 256GB DDR5	**40.10**	960GB NVMe	1x NDR200 Shared by 2 nodes
		216				>8GB/core 1024GB DDR5			
1		72		2x Intel Sapphire R. 9480	56c @ 1.9GHz	> 0.5GB HBM/core 128GB HBM + 32GB DDR5	**0.34**		

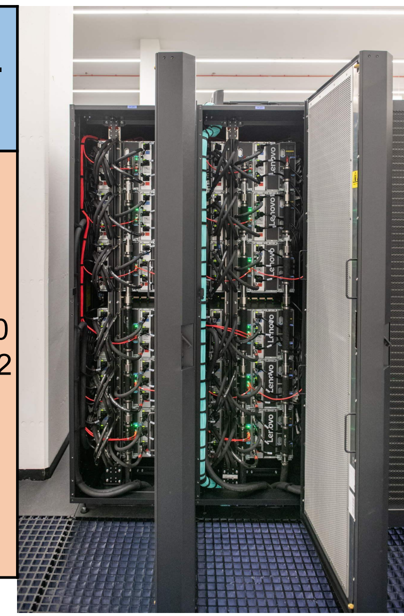

November 2023
HPL: #19 , #1 x86
HPCG: #24
Green500: #81
5.7 MW under HPL

HPL: 40.10 PFlops
HPCG: 484.36 TFlops
Green500: 6.97 Gflops/watt

MareNostrum5 ACC

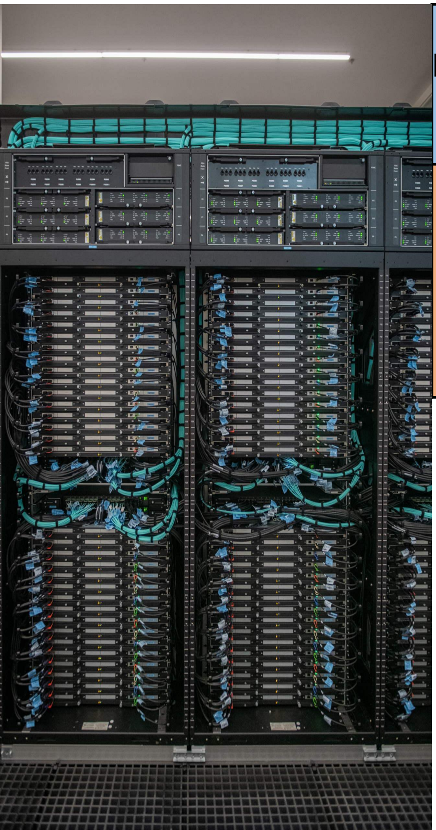

Racks	Cooling	Nodes		Processor/Accelerator		Memory	PFlops (HPL)	Local Drive	High-Perf. Network
		Total	per rack						
35	DLC	1120	32	2x Intel Sapphire R. 8460Y+	40c @ 2GHz	512GB	***175.3***	480GB NVMe	4x NDR200
				4x Nvidia Hopper 64GB HBM					

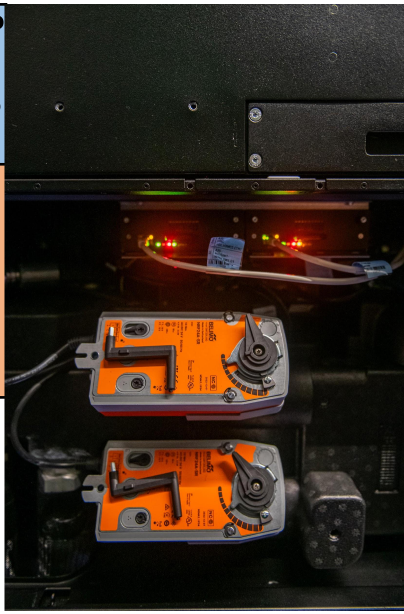

November 2023
HPL: #8
Green500: #5
2.5 MW under HPL

HPL: 175.30 Pflops
HPCG: 1.146 PFlops
Green500: 42.15 Gflops/watt
Graph500: 15.73e+12 GTEPS (BFS)
4.1 MW under HPL (175 PF/s)

MareNostrum 5 – High Speed network IB NDR200

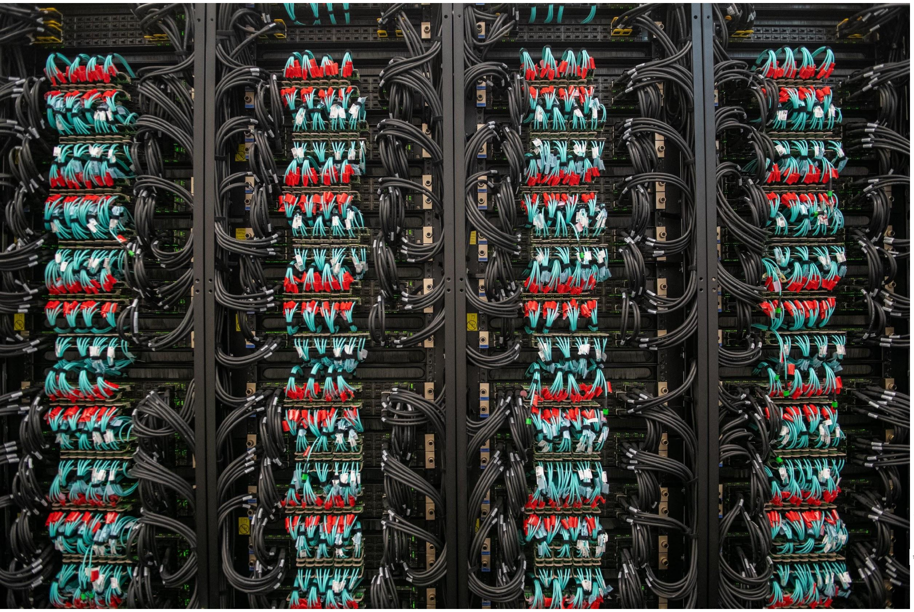

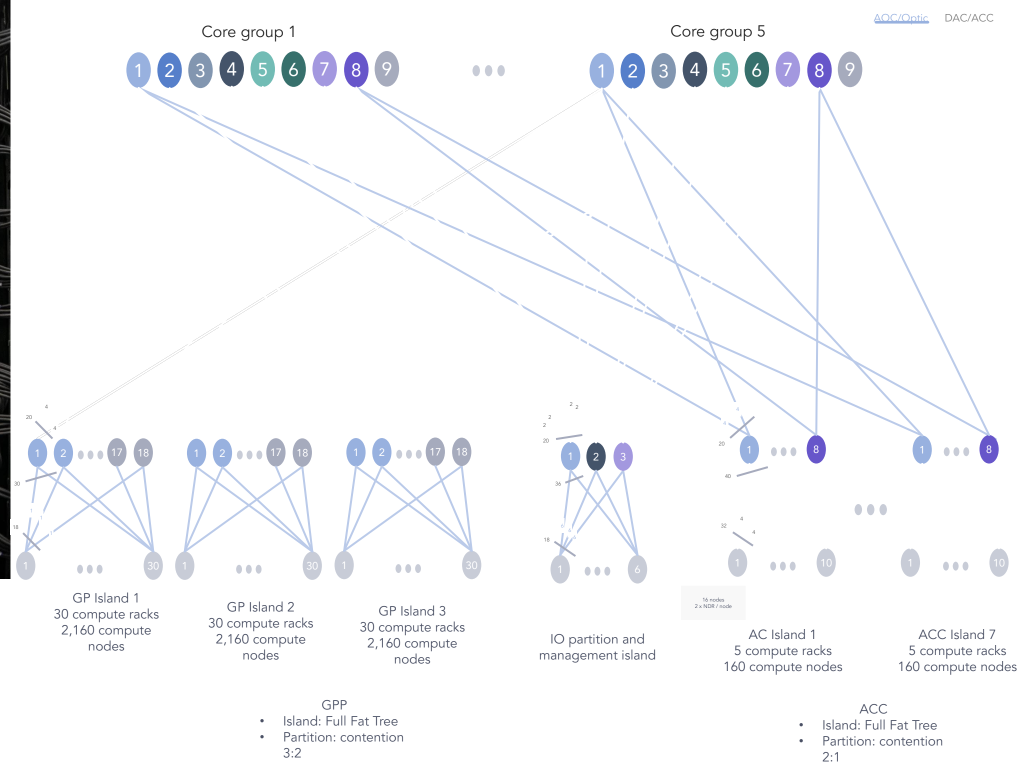

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Next Generation Compute

- General purpose
  - 408 compute nodes
  - NVIDIA Grace processor
  - Air-cooled chassis + RDHX
  - Some immersion cooling pods

- Accelerated

Pending to be confirmed

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Modeling and simulation

Digital Twins

- High-resolution "Digital Twins" using EuroHPC hardware
  - Provide tools, computational analysis and optimizations during the development of the new weather and climate models.
  - Ensure that the novel EuroHPC machines are used efficiently to reach the desired throughput.
  - The creation of approaches and tests which should provide measurable indicators about the quality of the solution: reproducibility, computational efficiency and code quality

  - Accelerate the new models through kernels and data structures suitable for the new GPU partitions.
    - Computationally intensive parts will be optimised for running on GPUs
    - Optimize from the computational point the new kernels to take into account the particularities of the new EuroHPC machines.

Modeling and simulation

Digital Twins

- High-resolution "Digital Twins" using EuroHPC hardware
  - Provide tools, computational analysis and optimizations during the development of the new weather and climate models.
  - Ensure that the novel EuroHPC machines are used efficiently to reach the desired throughput.
  - The creation of approaches and tests which should provide measurable indicators about the quality of the solution: reproducibility, computational efficiency and code quality

  - Accelerate the new models through kernels and data structures suitable for the new GPU partitions.
    - Computationally intensive parts will be optimised for running on GPUs
    - Optimize from the computational point the new kernels to take into account the particularities of the new EuroHPC machines.

Modeling and simulation

Digital Twins

- High-resolution “Digital Twins” using EuroHPC hardware
  - Provide tools, computational analysis and optimizations during the development of the new weather and climate models.
  - Ensure that the novel EuroHPC machines are used efficiently to reach the desired throughput.
  - The creation of approaches and tests which should provide measurable indicators about the quality of the solution: reproducibility, computational efficiency and code quality

  - Accelerate the new models through kernels and data structures suitable for the new GPU partitions.
    - Computationally intensive parts will be optimised for running on GPUs
    - Optimize from the computational point the new kernels to take into account the particularities of the new EuroHPC machines.

ESiWACE3

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

Destination Earth

Destination Earth | EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS

EDITO

EDITO European Digital Twin Ocean

ESiWACE3 services

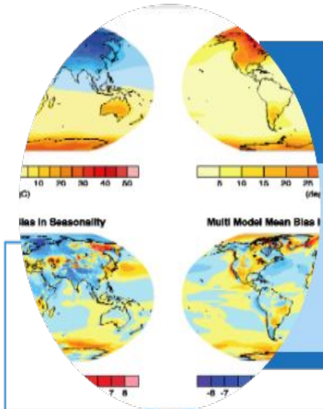

PERFORMANCE IMPROVEMENT

ESMValTool: Earth System Evaluation Tool

Performance improvement and parallelization using Dask

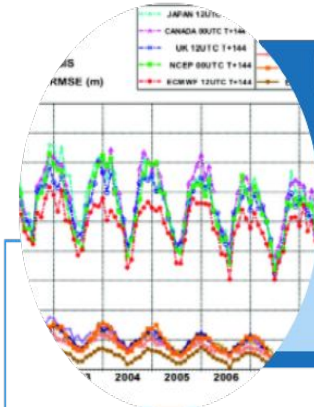

PROFILING & COMPUTATIONAL OPTIMIZATION

GLOBO: Atmospheric General Circulation Model

Profiling, optimizing MPI communications, Parallel I/O and Mixed/lower precision

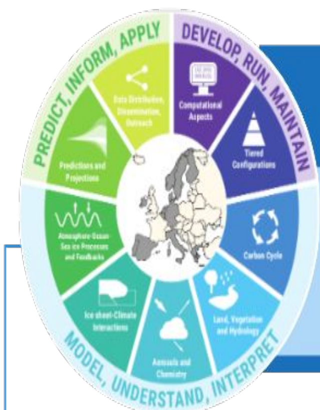

REDUCED PRECISION FOR ESM

EC-EARTH 4: European Community Earth System Model

Porting M7 aerosols in OpenIFS 48r1 to single precision to gain performance without degrading quality

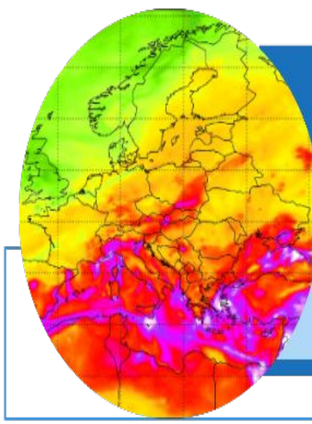

GPU PORTING & OPTIMIZATION

CAMP: Chemistry Across Multiple Phases

Revising and optimizing GPU implementation of CAMP for new HPC systems, focusing on MareNostrum5

ESiWACE3 trainings

- **<u>Training events</u>**
  - **1st ESiWACE3 Hackathon - October 2023**
  - **2nd ESiWACE3 Hackathon - June 2024**
  - **1st ESiWACE3-WarmWorld Summer School - end summer 2024**
    - The participants attended lectures and practical sessions, guiding them through the fundamentals of performing weather and climate simulations on modern HPC systems and managing and analysing Earth system model data.


esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE
SAVE THE DATE
JÜLICH
Forschungszentrum
3RD ESiWACE3 HACKATHON
Porting Earth System Models to the First EuroHPC Exascale System (JUPITER)
Jülich Supercomputing Centre Jülich (Germany)
Hibrid format
21-23 October 2025

ESiWACE3 trainings

- **Training events**
  - **1st ESiWACE3 Hackathon - October 2023**
  - **2nd ESiWACE3 Hackathon - June 2024**
  - **1st ESiWACE3-WarmWorld Summer School - end summer 2024**
  - **Online training courses & Training materials**

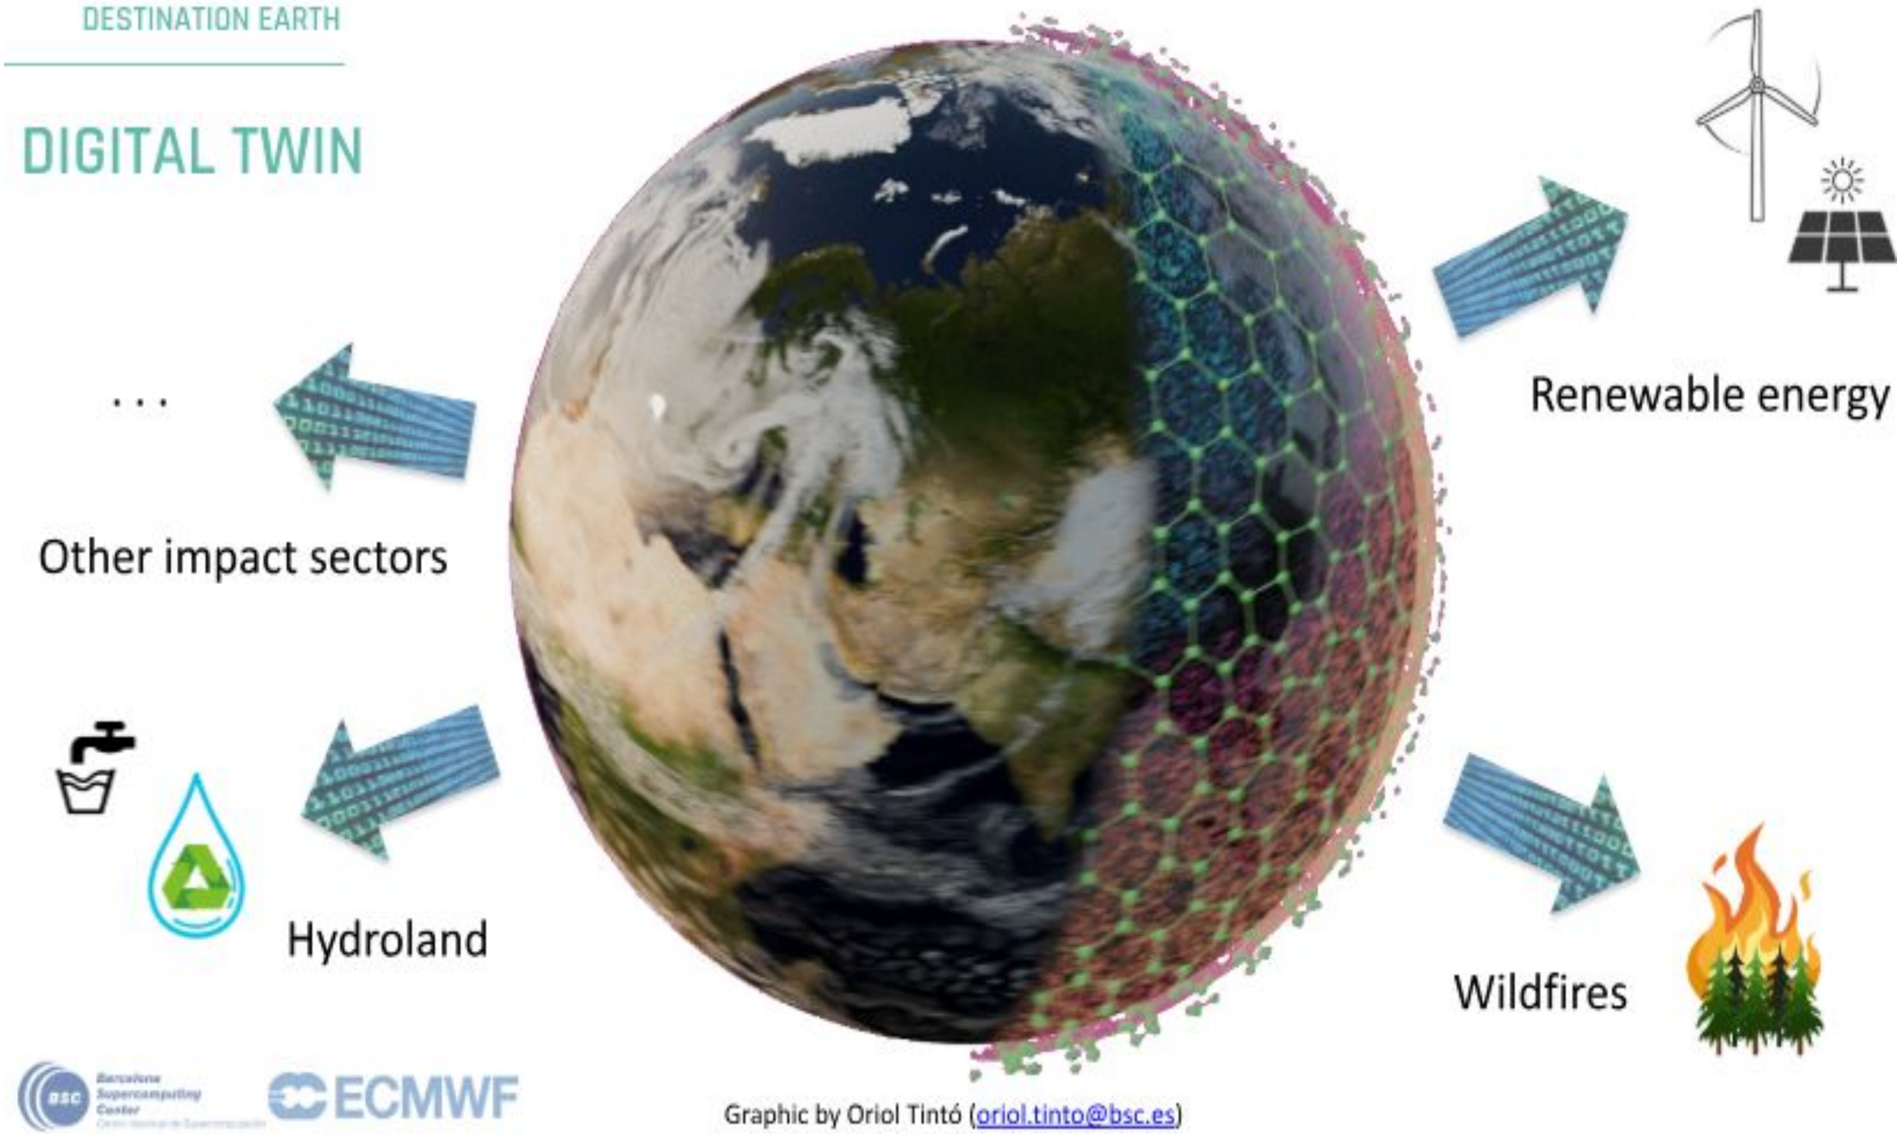
DESTINATION EARTH
DIGITAL TWIN
Renewable energy
. . .
Other impact sectors
Hydroland
Wildfires
Barcelona Supercomputing Center
ECMWF
Graphic by Oriol Tintó (oriol.tinto@bsc.es)

Increasing resolution - ClimateDT IFS-NEMO - Global 5kms

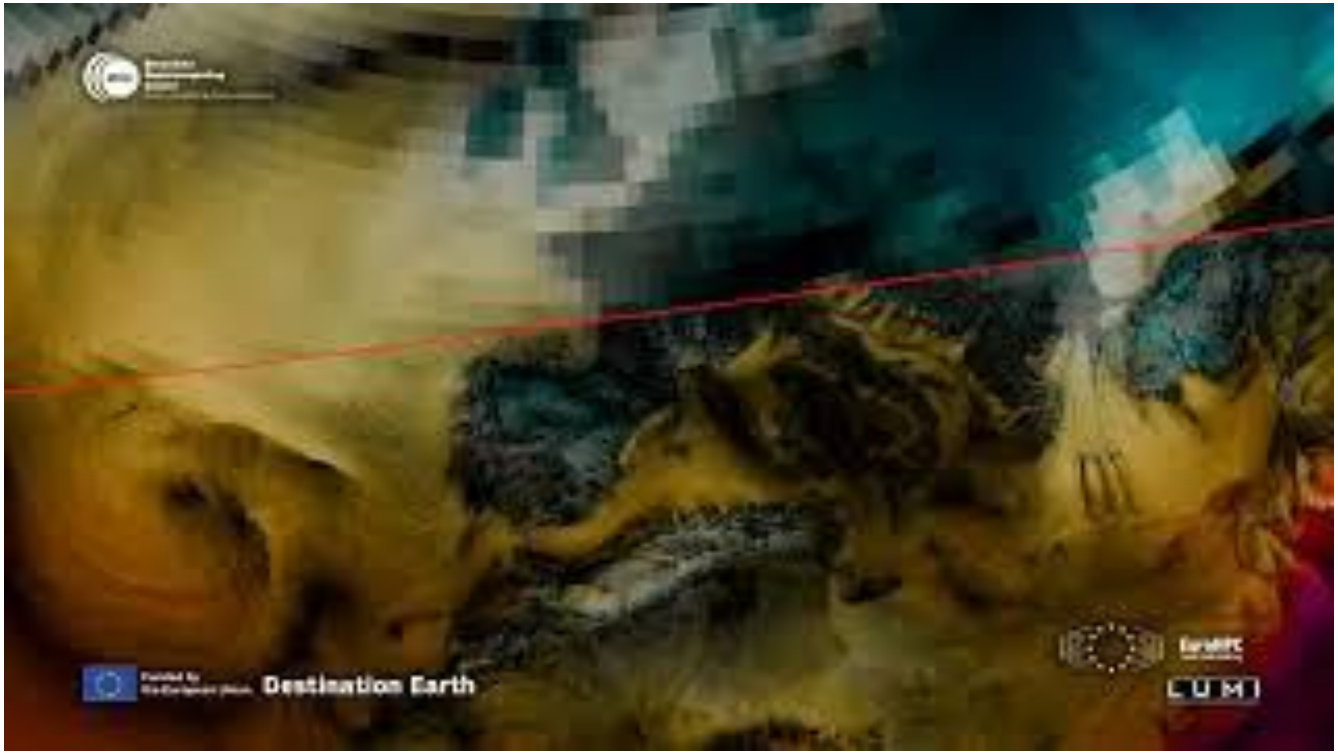

video by Oriol Tintó

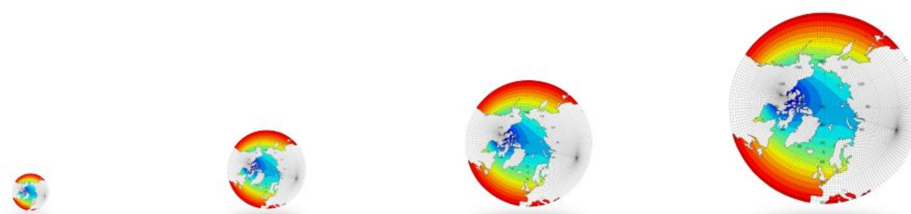

	ORCA 2	ORCA ¼	ORCA 1/12	ORCA 1/36
Grid resolution →	ORCA 2	ORCA ¼	ORCA 1/12	ORCA 1/36
Mem needed →	550 MB of memory	47 Gigabytes of memory	414 Gigabytes of memory	> 1 Terabytes of memory
Simulation time →	8 CPU hours	3500 CPU hours	90 000 CPU hours	~4 000 000 CPU hours
Outputs size →	10 Gigabytes of output (daily)	120 Gigabytes of output (daily)	1 Terabyte of output (daily)	> 5 Terabytes of output (daily)

The pre- and exascale path

- In short time
  - New pre-exascale machines (LUMI, LEONARDO, Marenostrum5).
  - High-resolution “Digital Twins” using EuroHPC hardware
  - European Projects pursuing HPC improvements (ESiWACE3)
- In medium/long time
  - RISCV, Quantum, FPGAs
  - Clouding
  - Hardware and software acceleration (IA, GPUs)
  - …

Modeling and simulation

Software and hardware acceleration

Data Driven

GPU

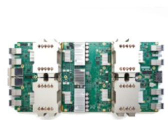

TPU

- A deep learning model trained with a large amount of data.
- Developed with less effort than a physical-based model.

Physical model accelerated

CPU

GPU

- Traditional approach accelerating intensive parts
- GPU porting from CPU
- Progress in science "ensured" through better physical processes simulated

Hardware acceleration

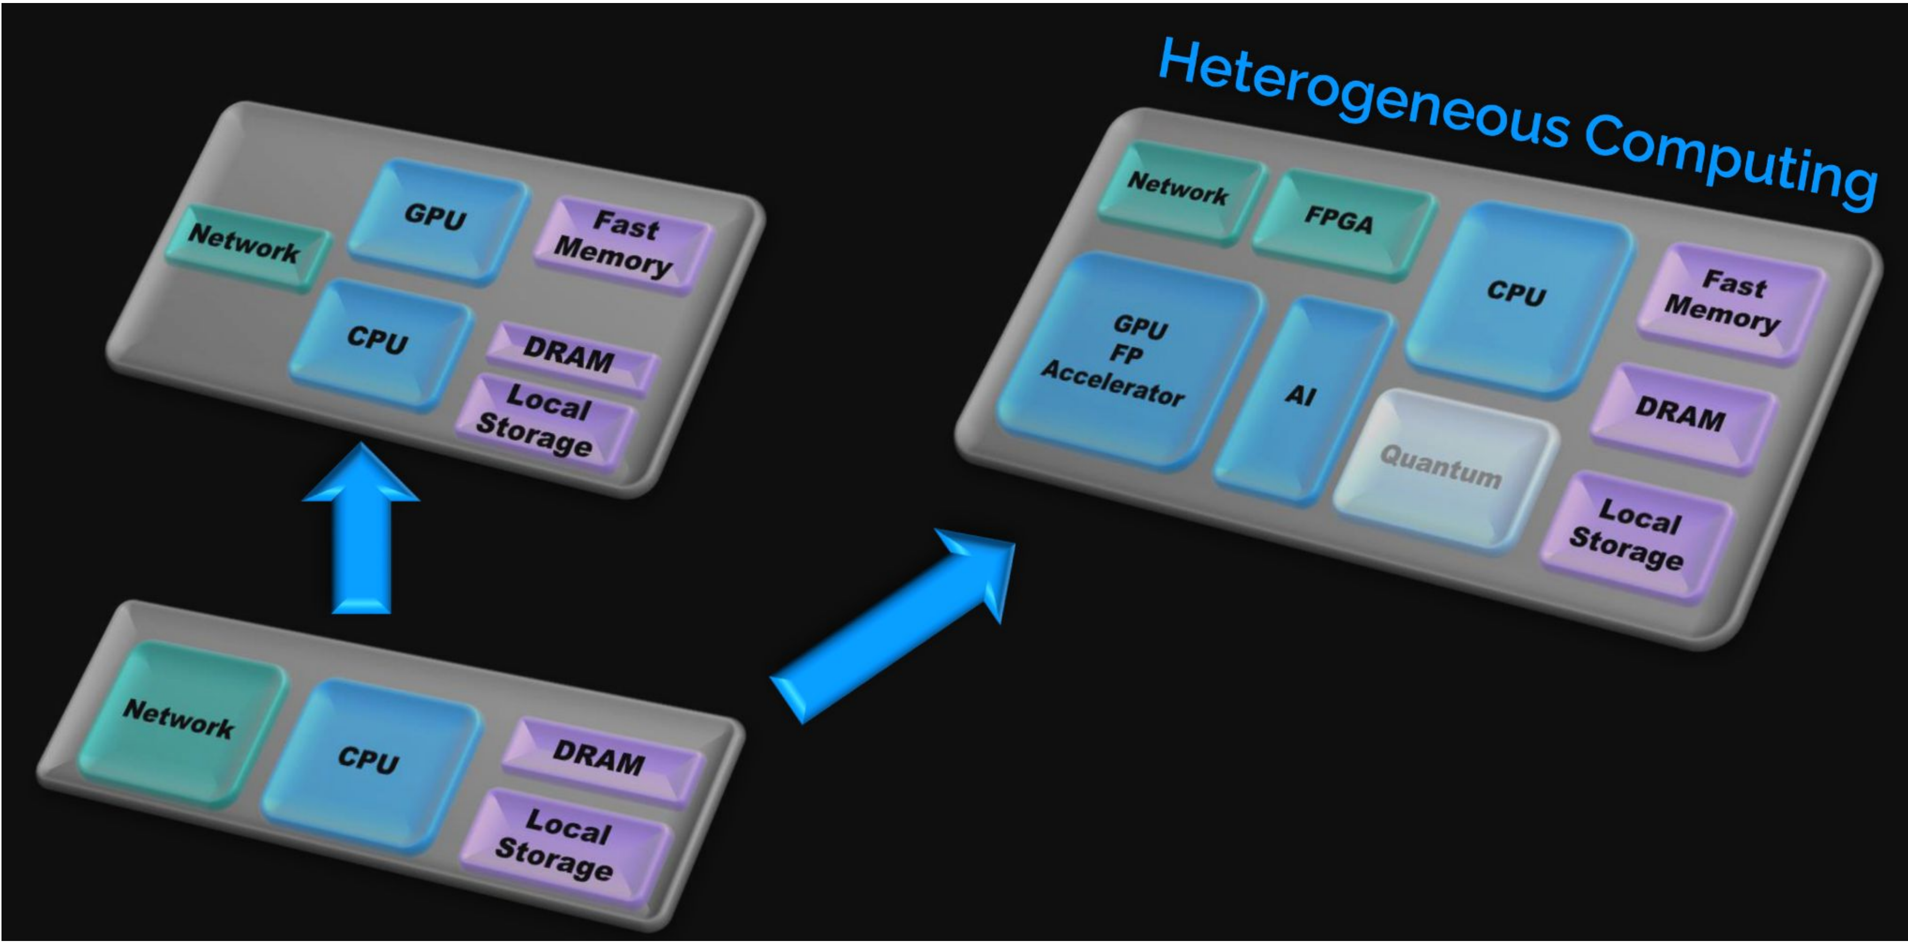

An approach to architectural design for exascale systems, 19th Workshop on HPC in meteorology, 2021, ATOS.

EPAC within EPI

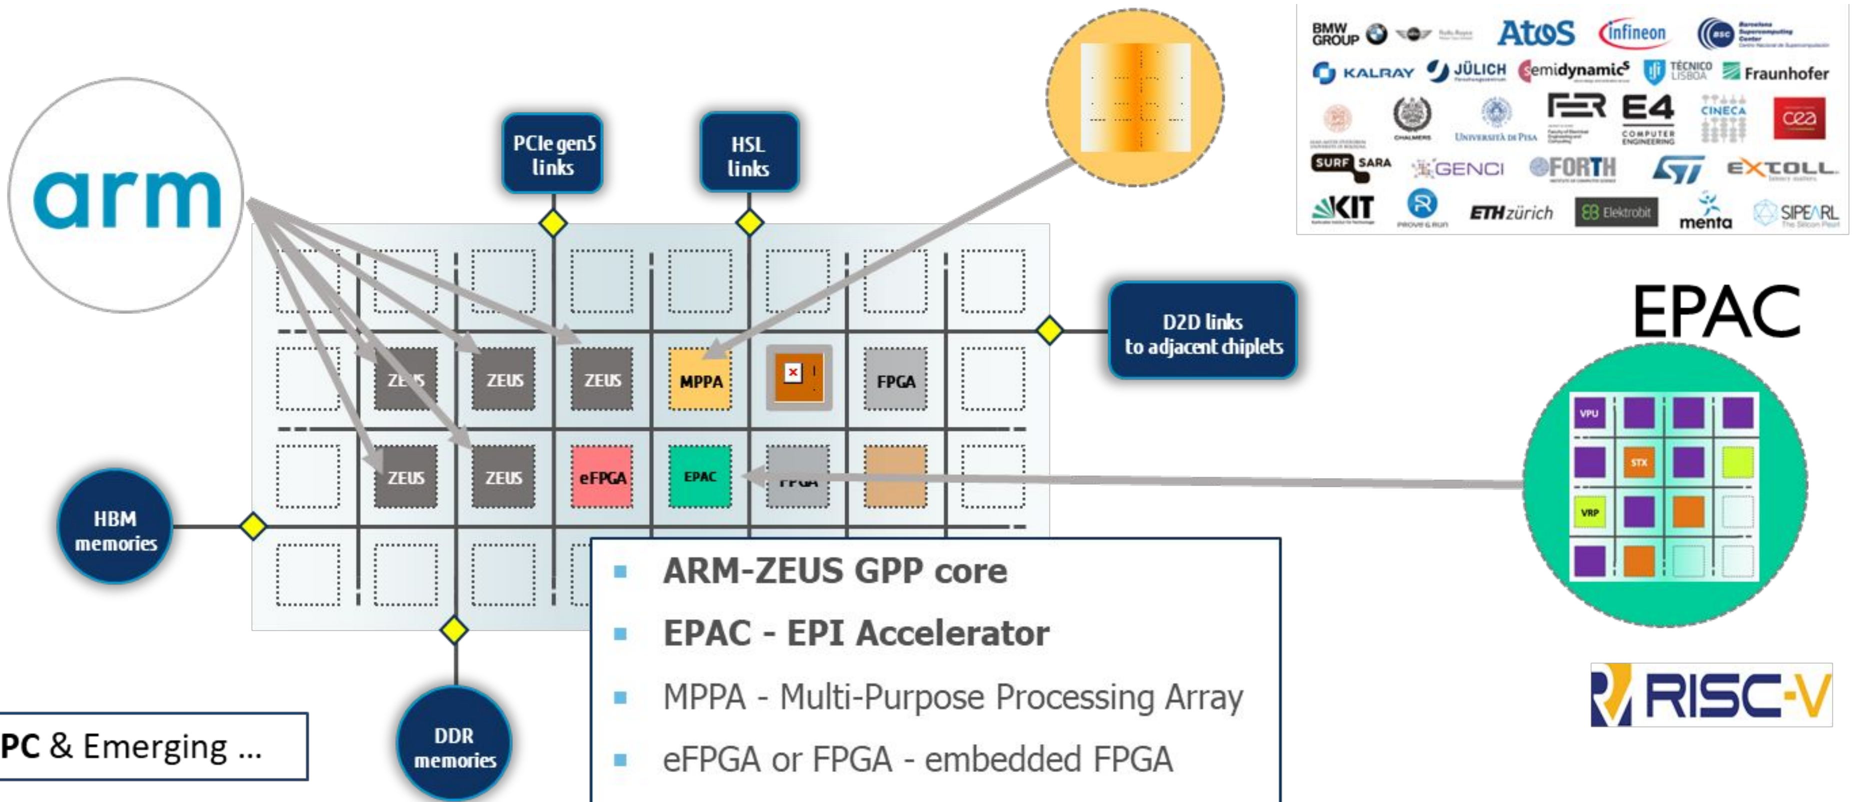

- **ARM-ZEUS GPP core**
- **EPAC - EPI Accelerator**
- MPPA - Multi-Purpose Processing Array
- eFPGA or FPGA - embedded FPGA
- Cryptographic ASIC (EU Sovereignty)

HPC & Emerging …

The RISC-V "accelerator" in EPI, Jesús Labarta (BSC), ACM Summer Schools

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Thank you

mario.acosta@bsc.es