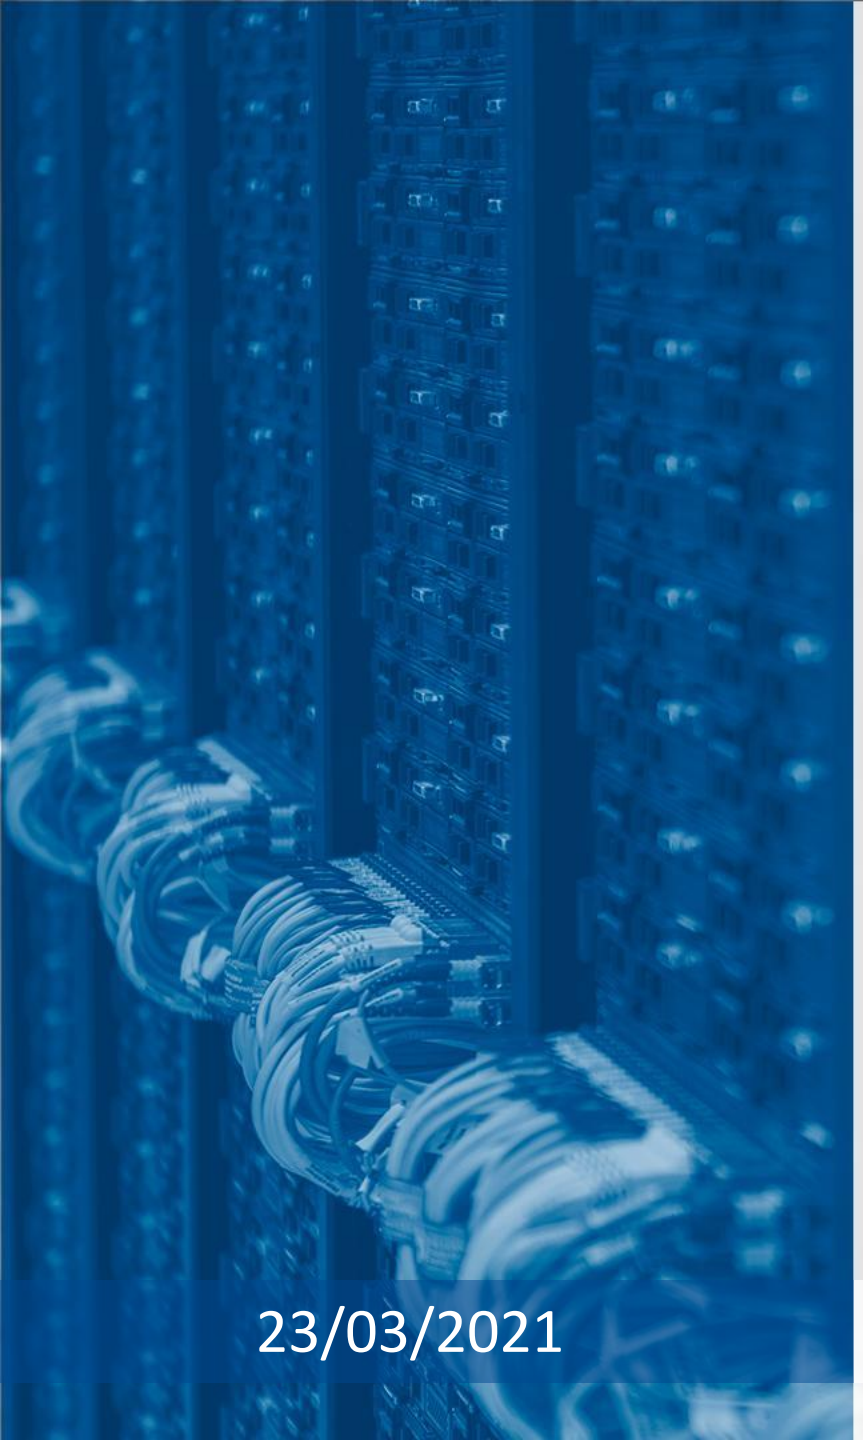

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

The new very-high-resolution coupled global configuration for EC-Earth 4

M. Castrillo, M. Acosta, T. Arsouze, I. Ayan, V. Lapin, G. Montané, S. Palomas, S. Paronuzzi, K. Serradell, O. Tintó, X. Yepes, C. Bricaud

23/03/2021

EuroHPC Summit Week 2022

Outline

- The **EC-Earth** GCM model
- The **coupled** pre-Exascale **demonstrator**
- The **new pre-Exascale VHR-GCM** configuration
- **Future:** Towards **higher resolutions**

EC-Earth

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

The EC-Earth (3) GCM

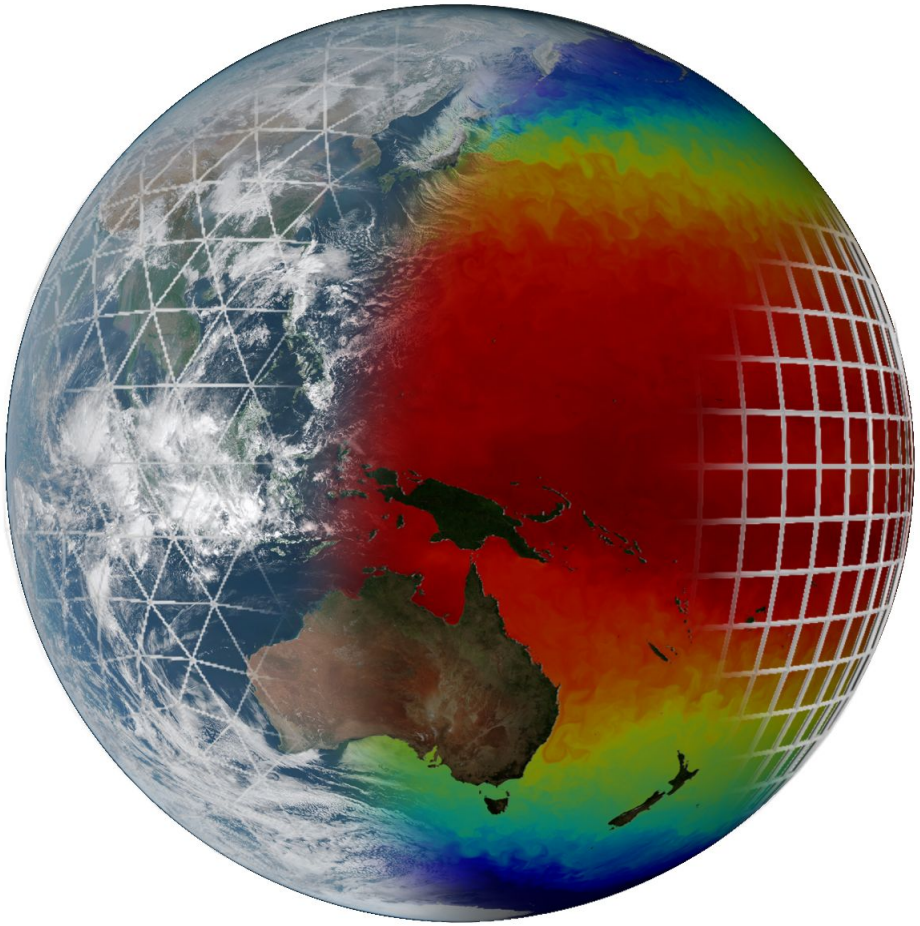

Atmosphere:

IFS

ECMWF

Ocean - ICE:

NEMO - LIM

Coupler:

Oasis

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Resolution in climate models

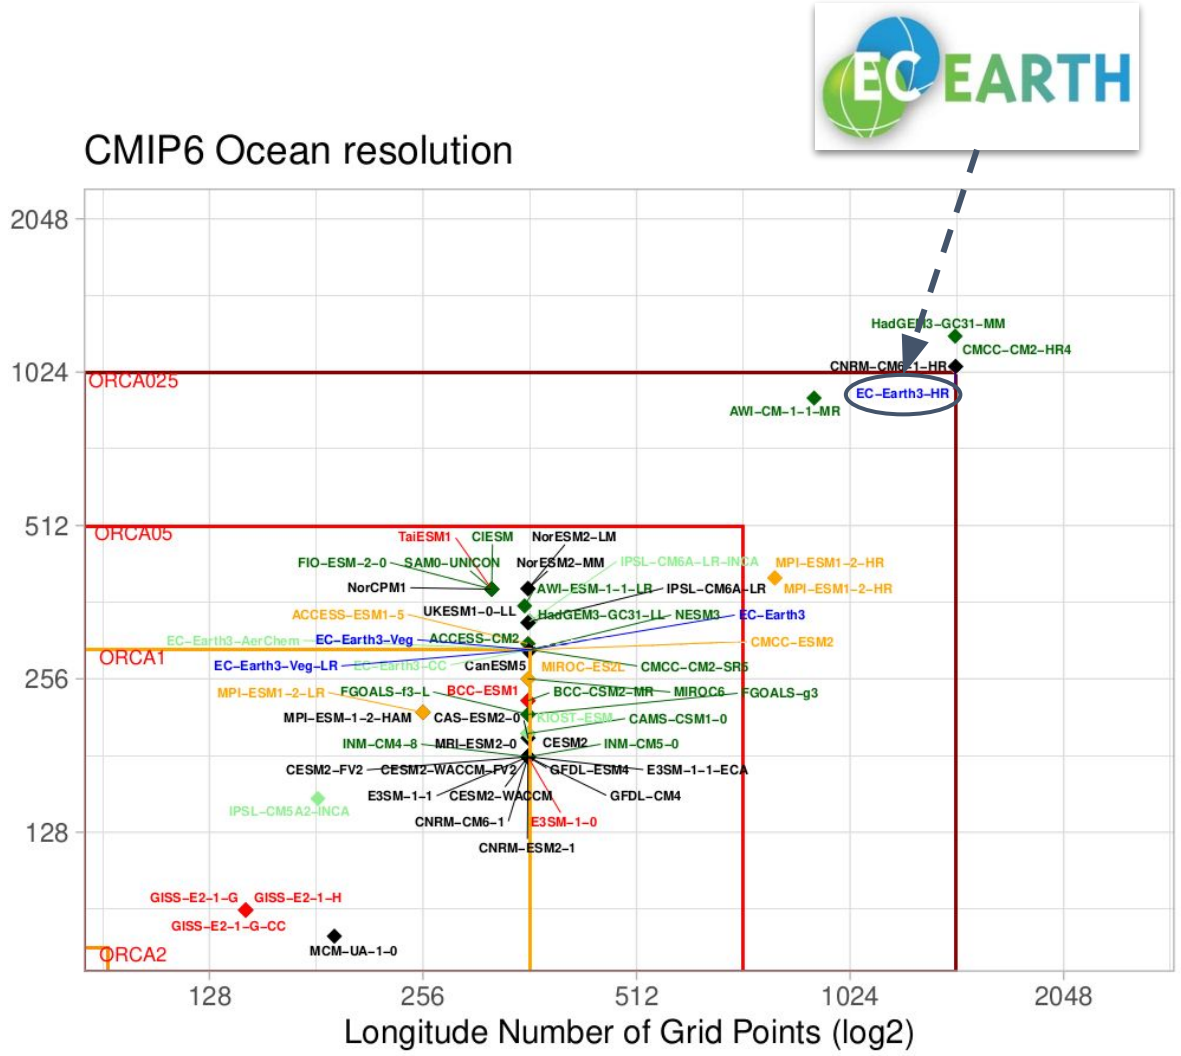

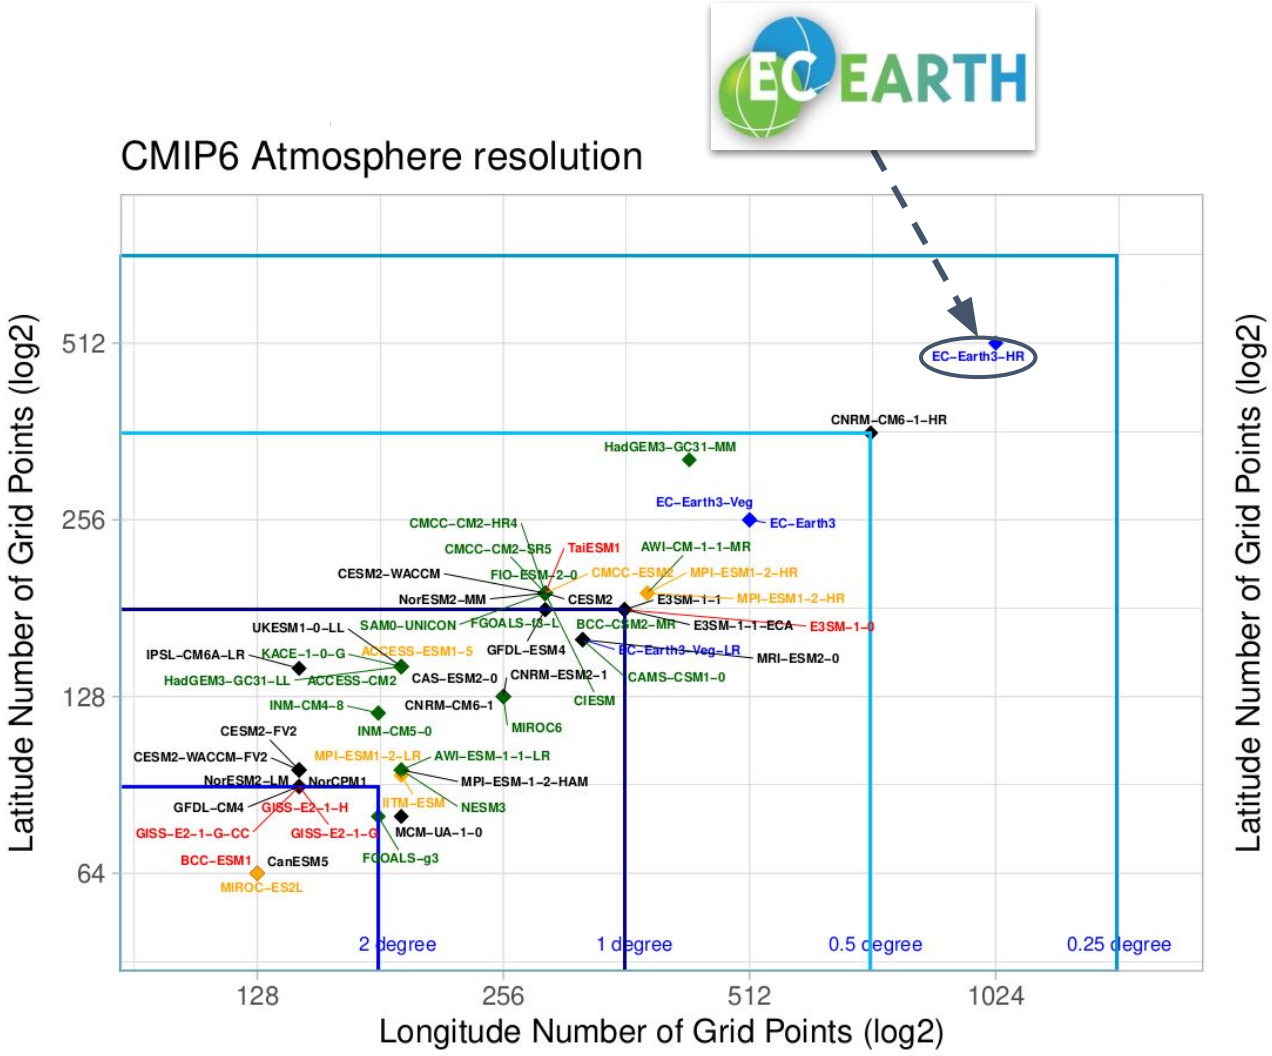

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

P. Davini (CNR)

The coupled pre-exascale demonstrator

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EC-Earth 3 coupled ~10 km

ESiWACE: EC-Earth ~10km coupled demonstrator

- **IFS** cycle 36r4 for atmosphere
  - T1279L91: **~16 km** grid point distance, **2.1 M** grid points
- **NEMO-LIM3** v3.6 for ocean & sea-ice
  - ORCA12L75: **~9 km** grid point distance, **13.2 M** grid points*
- Total 3D space points: **1,181kM vertices**

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

* Including land points (up to 1/3 of the total)

EC-Earth 3 - T1279-ORCA12 in MareNostrum4

Operational global, coupled ~10 km simulations:

- **EC-Earth 3.2** (IFS36r4 + NEMO 3.6 + OASIS3-MCT)
- **4,512 MPI tasks** - 0.44 SYPD, 160 SDPD

100 year exp
~25M computing hours,
227 days

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EC-Earth 3 - T1279-ORCA12: production runs

- **PRIMAVERA** is a **Horizon 2020** project which aims to develop a **new generation of advanced and well-evaluated high-resolution global climate models**, capable of simulating and predicting regional climate with **unprecedented fidelity**, for the **benefit** of governments, business and society in general.

- The **High Resolution Model Intercomparison Project (HighResMIP)** is a **CMIP6** endorsed MIP that applies, for the **first time**, a **multi-model approach** to the systematic investigation of the **impact of horizontal resolution**.

H2020: European HPC & science integration case

Research infrastructure

esiwace

CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

HPC applications (CoEs)

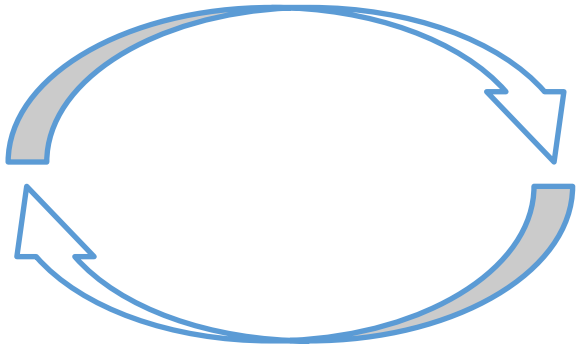

Community

PRIMAVERA

Climate science and HPC

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EC-Earth 3 - T1279-ORCA12: production runs

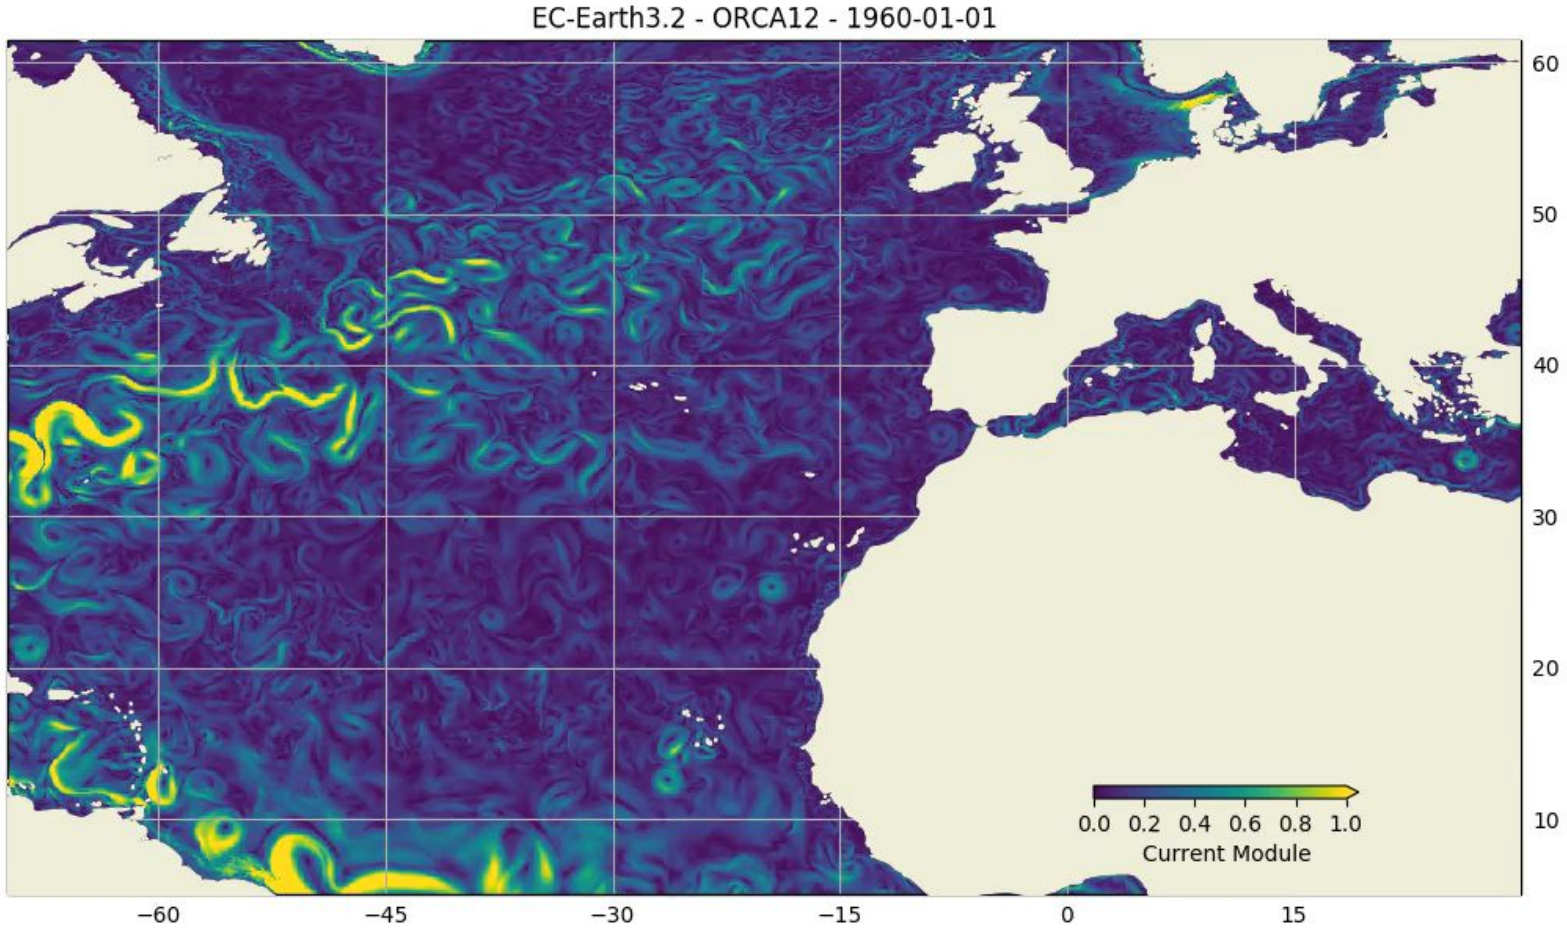

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Courtesy of: Thomas Arsouze

EC-Earth 3 - T1279-ORCA12 in MareNostrum 4

Benchmark scaling

Reduced output: ocean monthly means & 6-hourly atmospheric variables

Timestep: 6 min. (atm. & oce.)

Coupling freq.: 12 min (atm-oce) and (oce-ice).

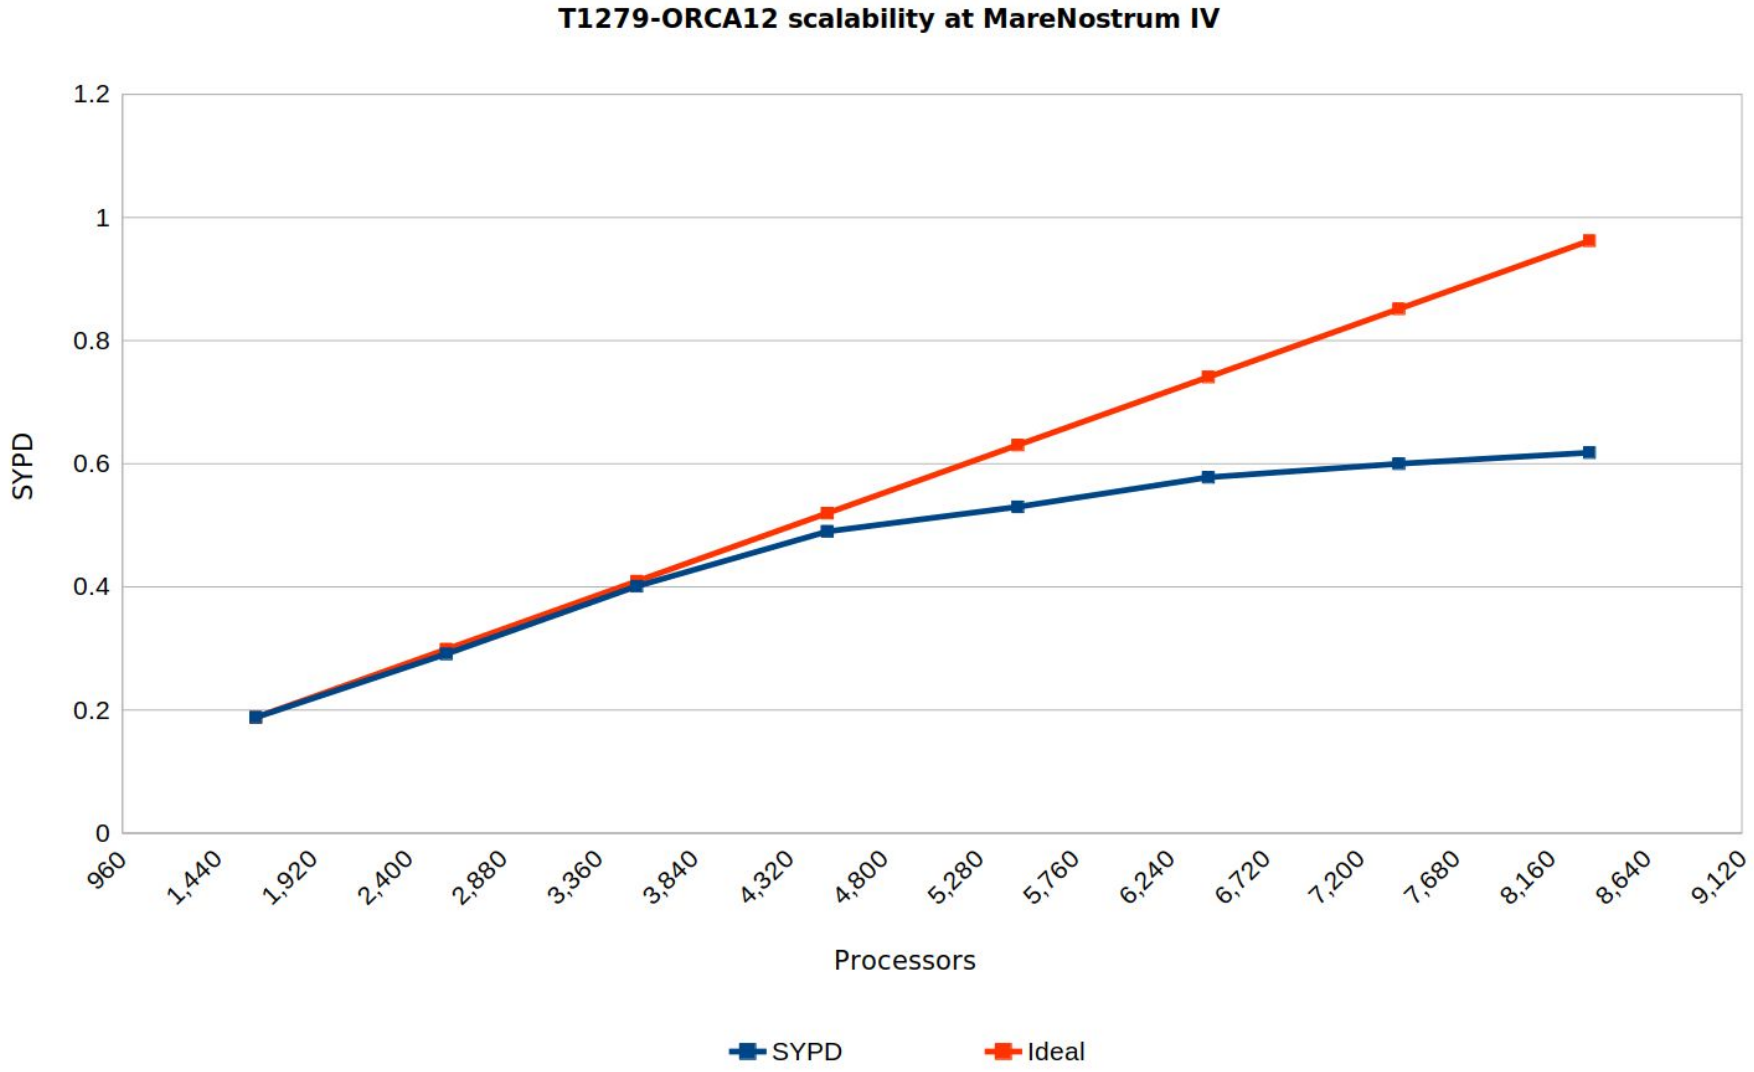

EC-Earth - T1279-ORCA12: Performance analysis

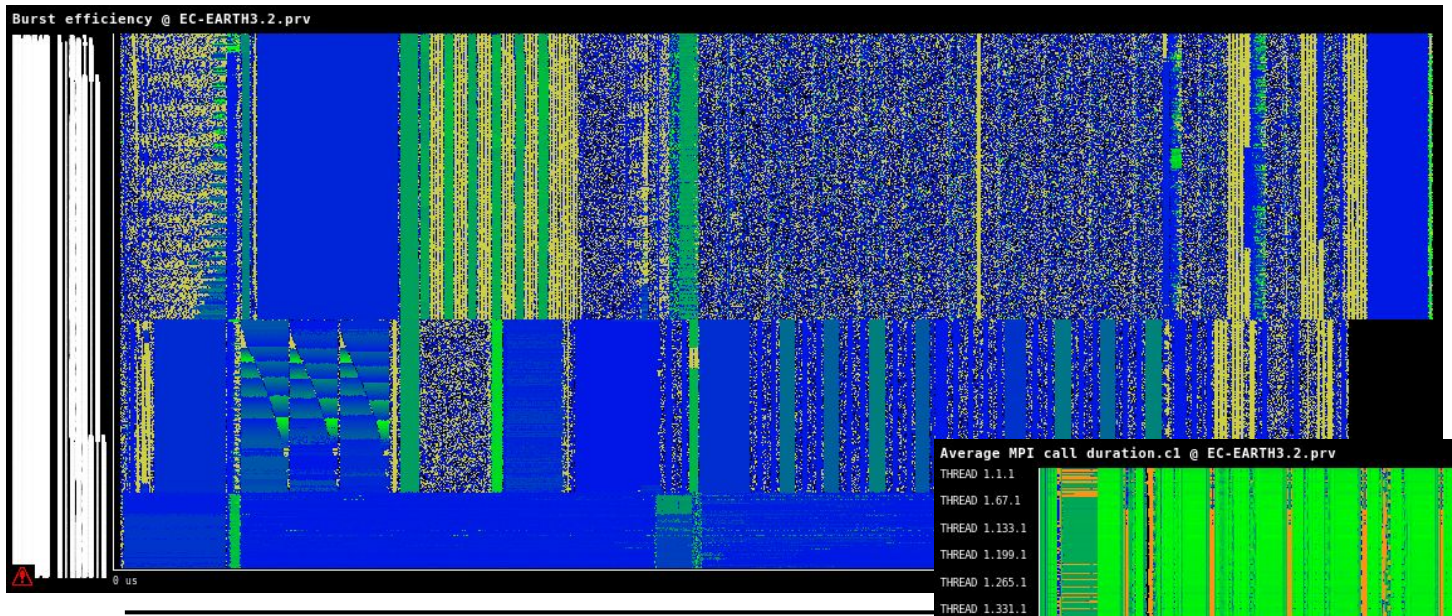

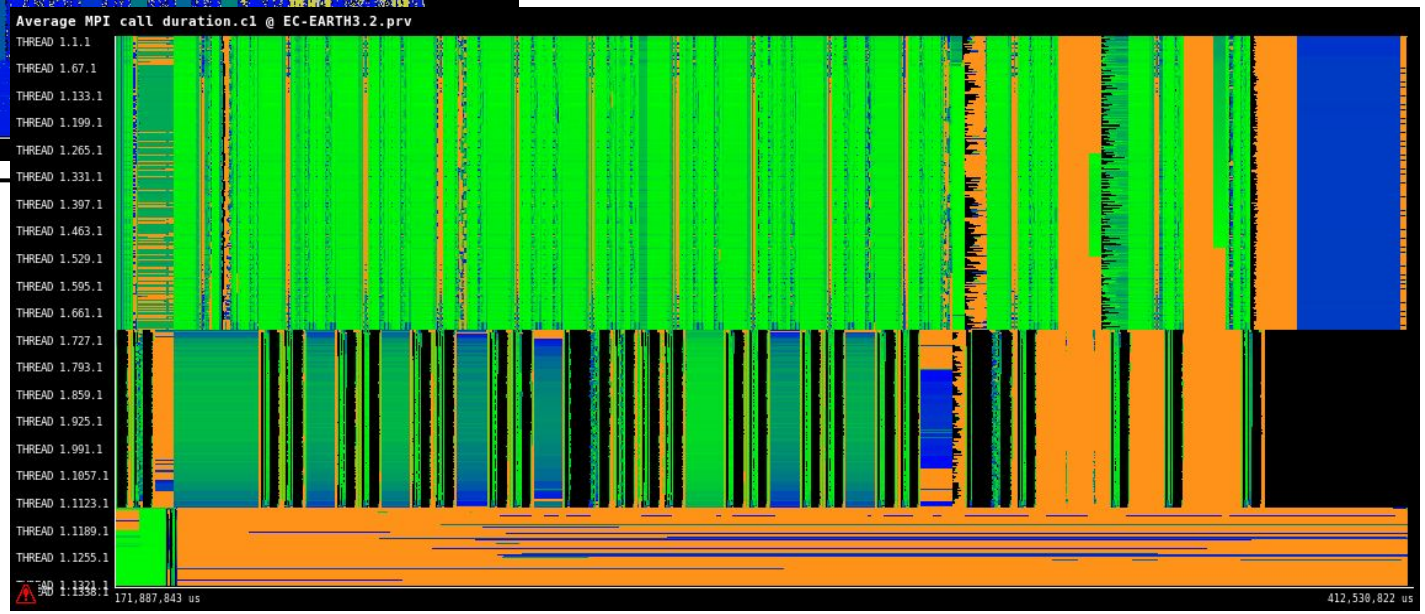

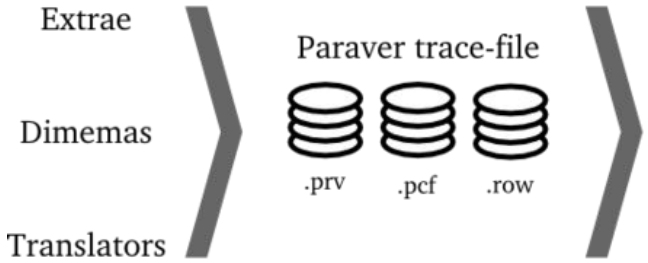

https://tools.bsc.es/

ESiWACE 2: EC-Earth coupled ~10 km production-mode

- **I/O overhead** ⟶ Implement asynchronous I/O
- **Sea-ice scalability** ⟶ Reduce **global** communications
- **Legacy atmospheric model (2010)** ⟶ Update **IFS** to newest cycle, using octahedral grid

esiwace2
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

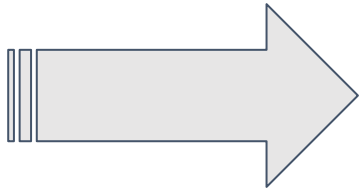

> 1 SYPD

~~**0.45 SYPD**~~

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

The new pre-exascale VHR configuration

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

ESiWACE 2: EC-Earth coupled ~10 km production-mode

- Develop **infrastructure** for production-mode configurations

ESiWACE 2: EC-Earth coupled ~10 km production-mode

- Develop **infrastructure** for production-mode configurations
  - **Coupling** infrastructure **(OASIS)**
  - Improvement of **I/O (XIOS)**
  - **NEMO** for high-resolution
  - Infrastructure for **high-resolution data**

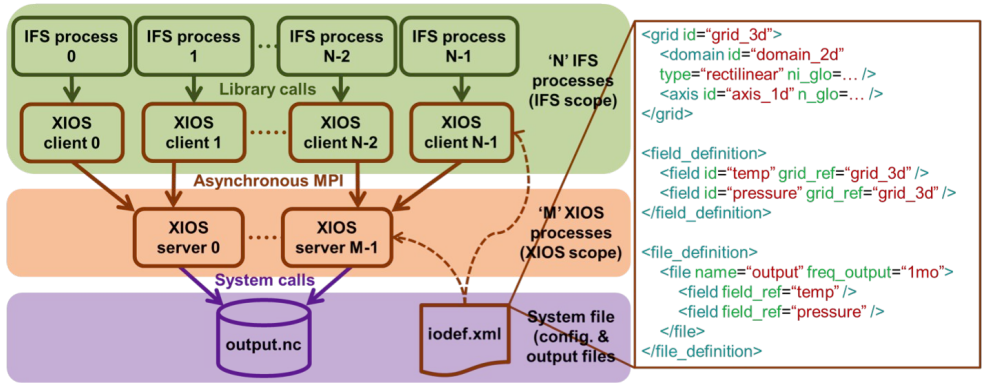

XIOS integration into OpenIFS (X. Yepes)

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

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

ESiWACE 2: EC-Earth coupled ~10 km production-mode

- Develop **infrastructure** for production-mode configurations
- Develop production-mode **configurations** 
- Port models to **pre-exascale** EuroHPC **systems**

EC-Earth coupled ~10 km production-mode

Main assets

- Brand **new ESM: EC-Earth 4**
  - **OpenIFS cycle 43r3** (2020)
  - **NEMO v4.0.2** (2019) (incl. **SI3** sea-ice model)
  - **OASIS3-MCT 4**
- **Common** asynchronous **I/O server** (XIOS v2.5)
- **Octahedral reduced** gaussian grid for OpenIFS
- *Possibility to **switch** numeric **precision***

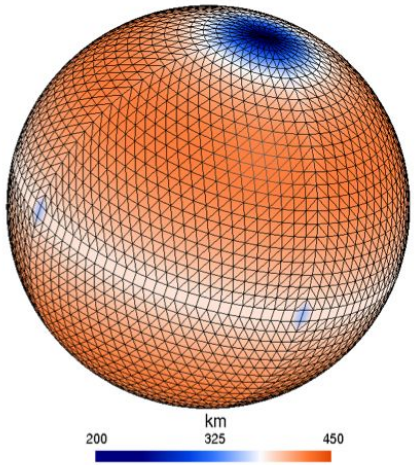

N24 octahedral Gaussian grid

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

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

EC-Earth coupled ~10 km production-mode

ESiWACE2: EC-EARTH 4 VHR coupled demonstrator

- **OpenIFS** cycle 43r3 for atmosphere
  - Tco639L91: ~16 km grid point distance, **1.66 M** grid points
- **NEMO-SI3** v4 & SI3 for ocean & sea-ice
  - ORCA12L75: ~9 km grid point distance, **13.2 M** grid points*
- Total 3D space points: **1,141kM vertices**

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

* Including land points (up to 1/3 of the total)

Tco639-ORCA12 configuration development

ORCA12 inspired from EC-Earth 3 T1279-ORCA12. Namelist "tuning" for NEMO4.

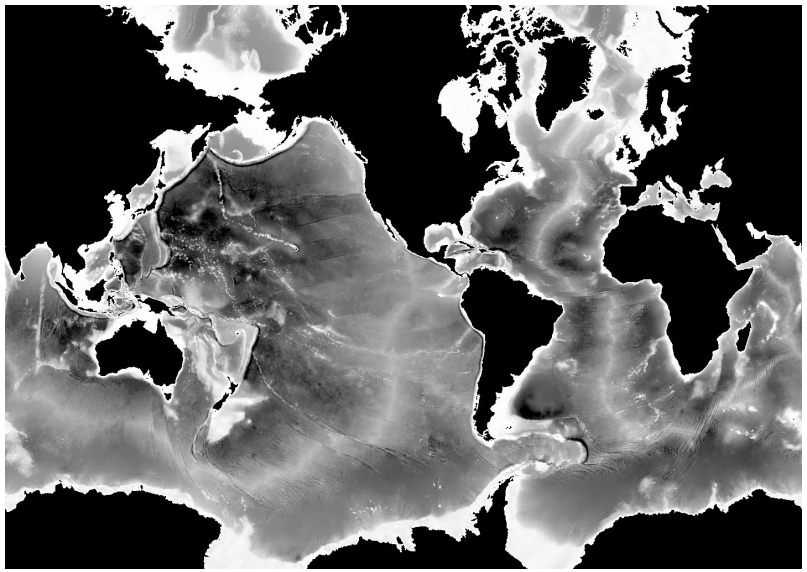

ORCA12 bathymetry

Tco639 configuration and **initial conditions.**

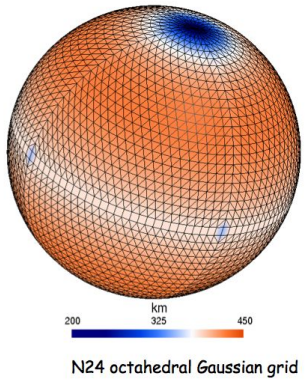

N24 octahedral Gaussian grid

OASIS coupler grids, masks and areas information using the **OCP**[1] tool.

OASIS remapping weights generated in parallel (OpenMP)[2].

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

[1] https://github.com/JanStreffing/ocp-tool
[2] OASIS3-MCT4 new feature

Tco639-ORCA12 in MareNostrum 4

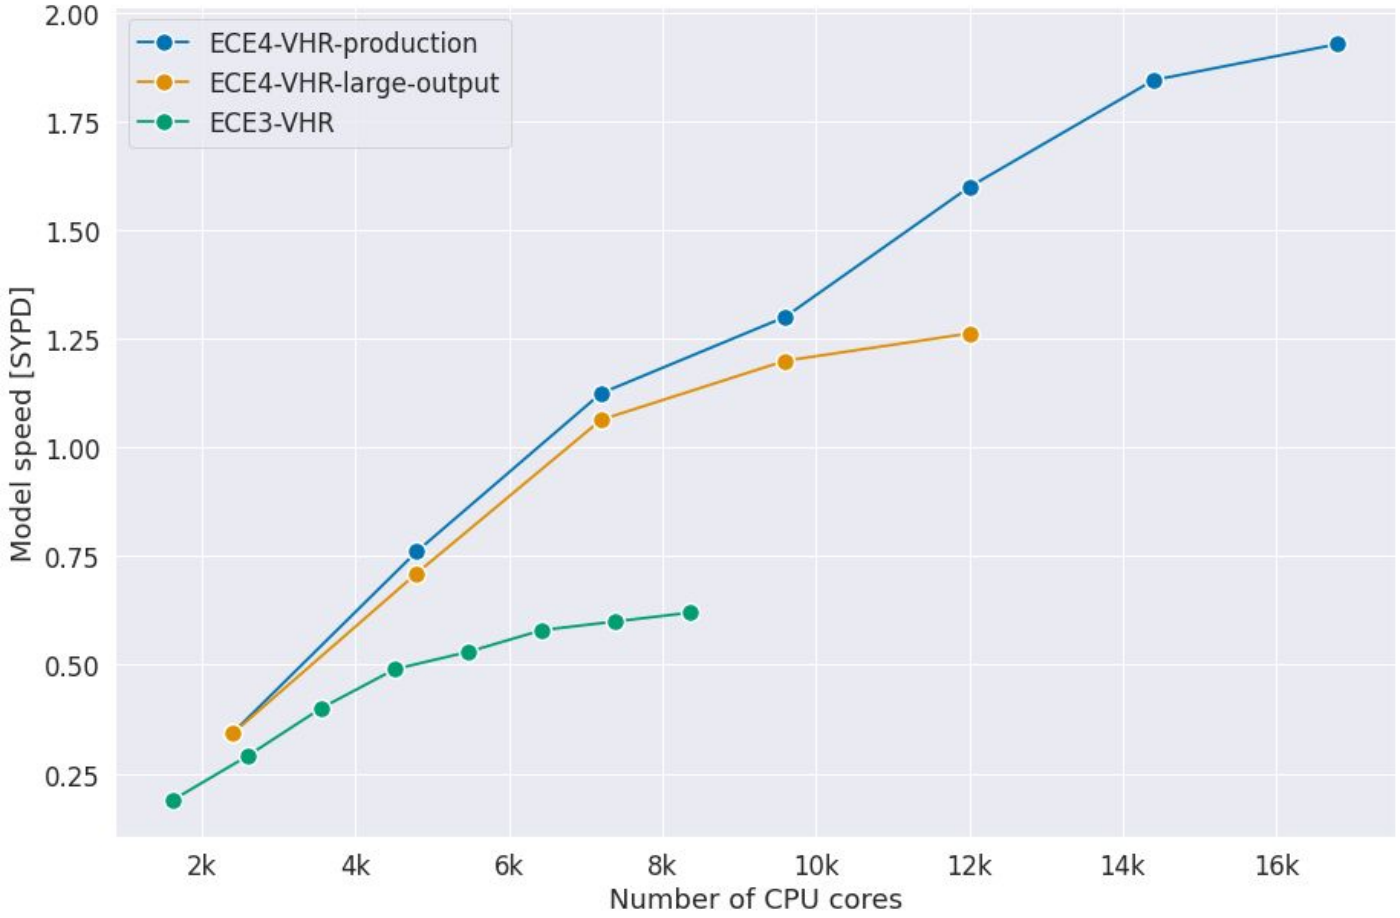

ECE3: TL1279-ORCA12

ATM: 360s, **OCE: 360s**, ICE: 720s, **CPL: 720s**

ECE4: Tco639-ORCA12

ATM: 360s, **OCE: 240s,** ICE: 720s, **CPL: 3600s**

- **Production: Monthly** output (6-hourly and daily averages)
- **Large output: 3-hourly** output

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

100 year experiment, **~16M** computing hours & **68** days

EC-Earth - Tco639-ORCA12 performance analysis

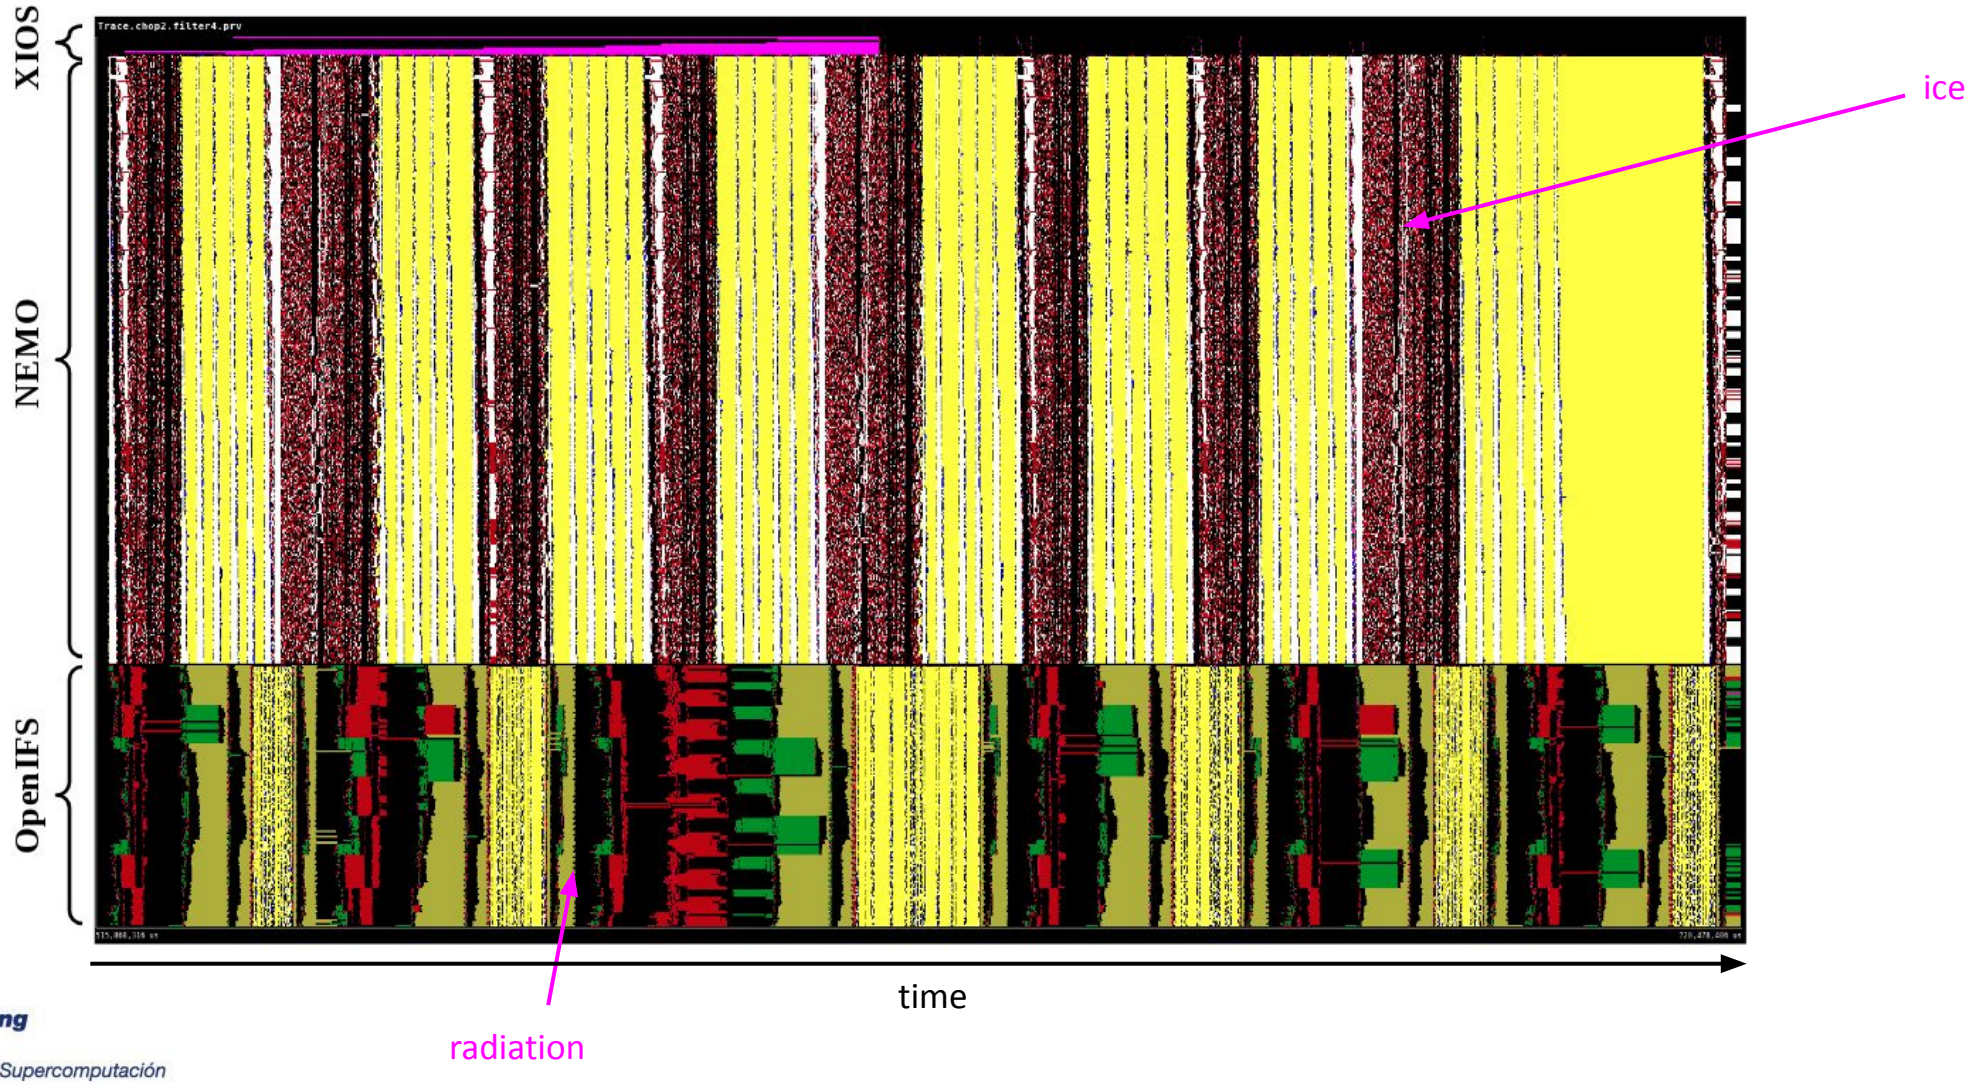

Towards higher resolutions

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Global 1/36° (ORCA36): context

- Model configuration for future CMEMS/MOI global forecasting and reanalysis systems
- Based on NEMO 4

IMMERSE (EU H2020):

Demonstrator for developments in NEMO 4 (HPC developments)

ESIWACE2 (EU H2020):

Demonstrator for « production runs at unprecedented resolution on pre-exascale supercomputers »

CMEMS contract with BSC:

« 87-GLOBAL-CMEMS-NEMO: EVOLUTION AND OPTIMISATION OF THE NEMO CODE USED FOR THE MFC-GLO IN CMEMS » :

NEMO HPC performance, global 1/36°

Clement Bricaud (MOI)

Scaling ORCA36 in MN4

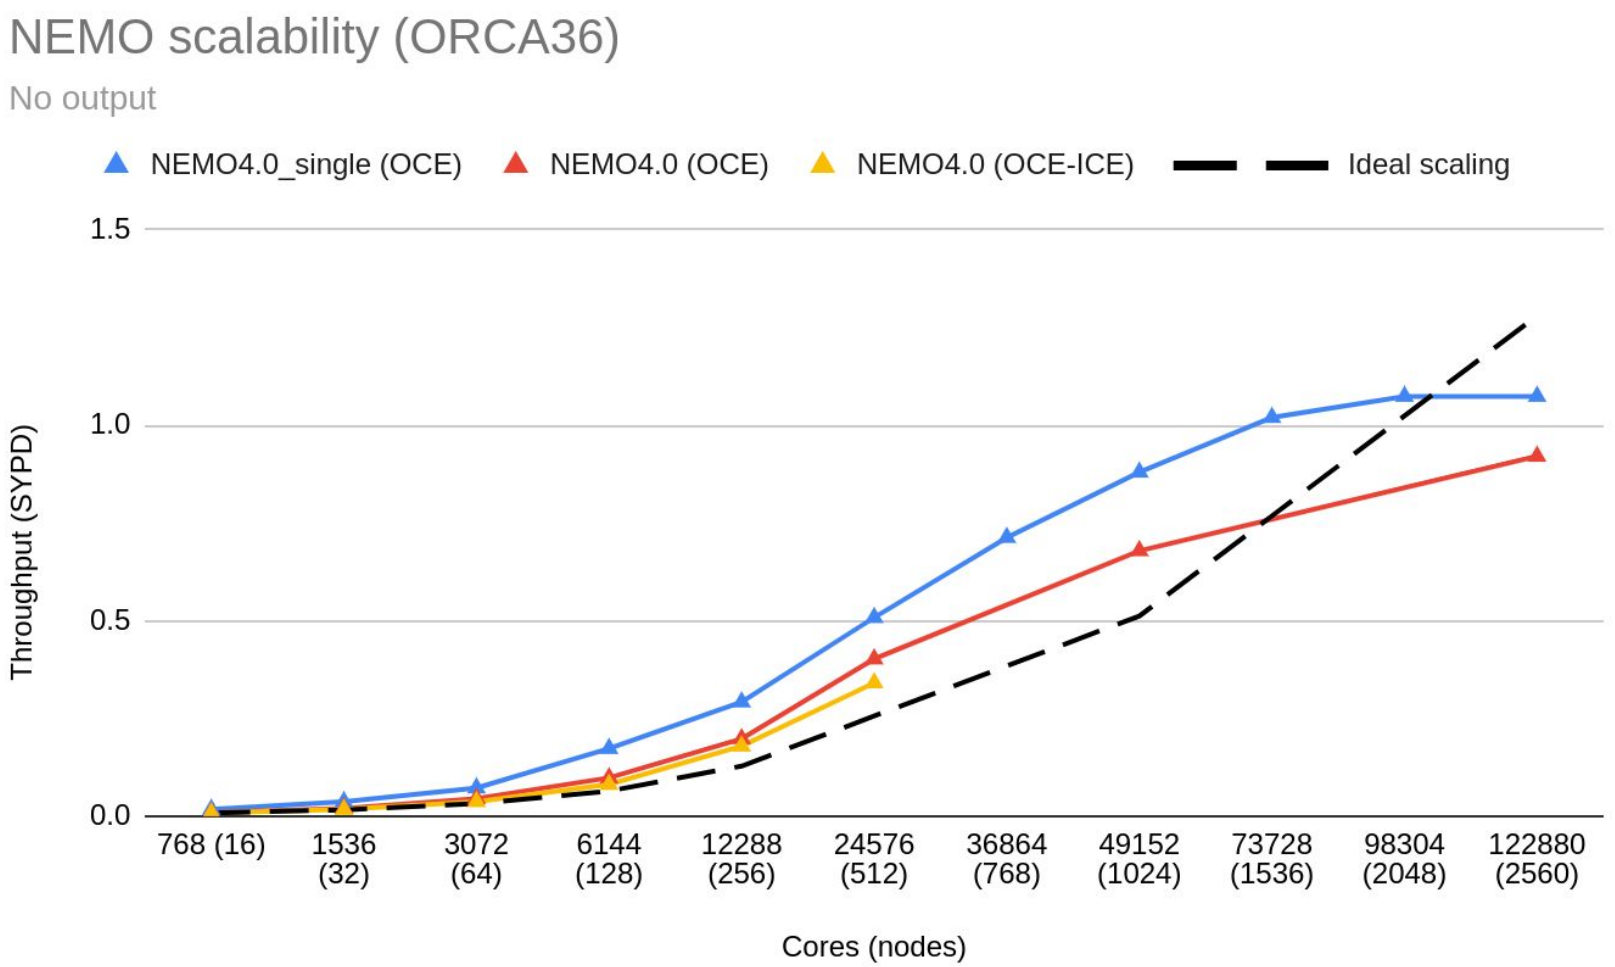

MN4 total = 165,888 cores!!

MareNostrum 5. A European pre-exascale superc

- **200 Petaflops** peak performance (200 x 10^{15})
- **Experimental platform** to create supercomputing technologies “made in Europe”
- **223 M€** of investment

EuroHPC Joint Undertaking

Hosting Consortium:

Spain **Portugal** **Turkey** **Croatia**

European Commission

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Generalitat de Catalunya
Departament d'Empresa i Coneixement

Gobierno de España
Ministerio de Ciencia, Innovación y Universidades

UPC Universitat Politècnica de Catalunya BarcelonaTech

Conclusions

- **First coupled 10km** configuration developed within ESiWACE:
  - Shared with the **EC-Earth consortium** partners
  - Used in **production** for **different** projects
  - Used to investigate **very-high resolution scalability** for coupled systems

- **10 km production-mode** configuration developed within ESiWACE2:
  - Solves the most important **bottlenecks**. Uses **updated** model components
  - Will be deployed and analysed in **pre-Exascale** EuroHPC systems
  - Will allow running VHR simulations with a **production throughput**

- **Bases** established for **increased resolution** and **parallelism**

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Thank you!

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

immerse
IMPROVING OCEAN MODELS FOR THE COPERNICUS PROGRAMME

The ESiWACE2 project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 823988
The IMMERSE project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 821926

miguel.castrillo@bsc.es

VHR-GCM: Technical take-home messages

- **Parallelizing** the **remapping weights** creation **saves** a lot of **time**.
  - **Independent** process for each pair of source-target grid.
- MareNostrum4 (OmniPath): **Open MPI** more **robust** than **Intel MPI**.
- XIOS: **parallel I/O** has a significant **overhead**. Multiple-file fairly efficient.
- OASIS: **Initialization** time increases with the number of processes (need to **tune** the remapping search method)
- NEMO: Using EN4 T&S. Need of a **spin-up** to **increase the timestep**.
- Properly **distributing** resources to handle **memory** needs.

Tco639-ORCA12 in MareNostrum 4

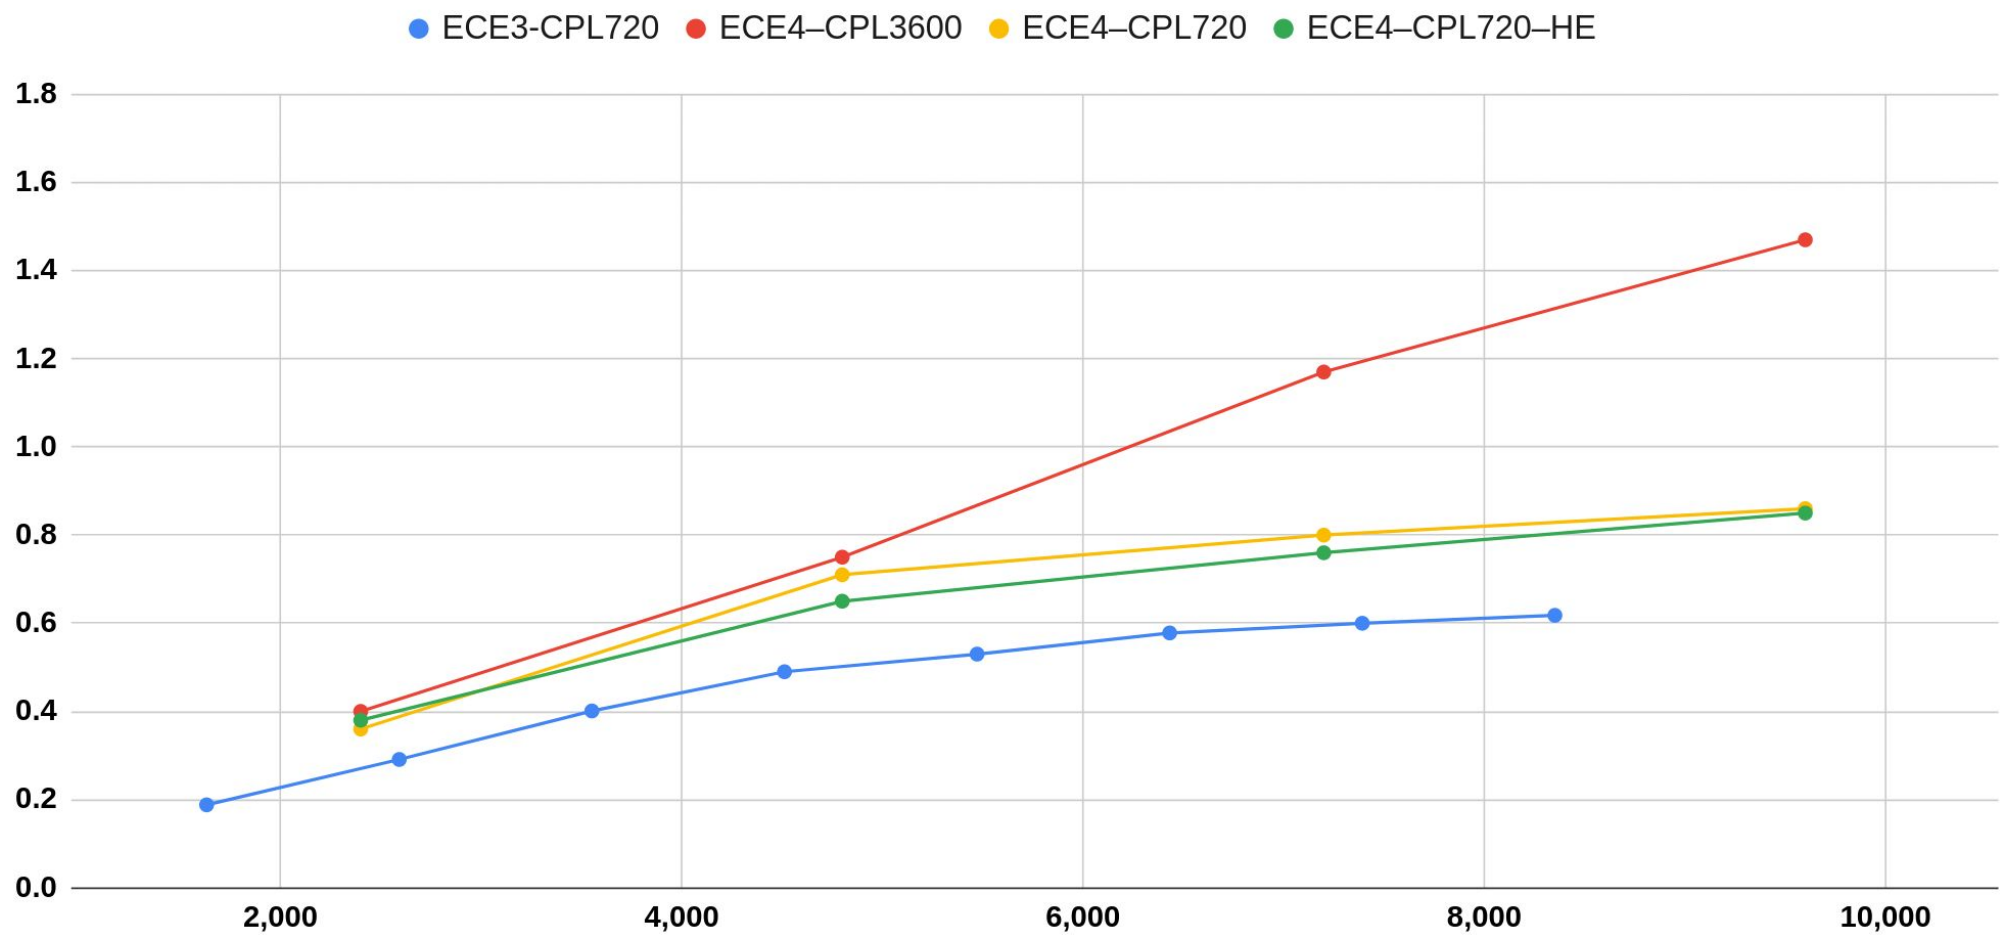

Scientific objectives

- Develop and prepare a **new generation** of **global high-resolution** climate models
- Evaluating global high-resolution climate models at a **process level**
- Focus on **air-sea interactions** at oceanic mesoscale:
  - Thermal feedback
  - Evaluate the role of the mechanical interactions between oceanic surface currents and atmospheric winds ("**current-feedback**")

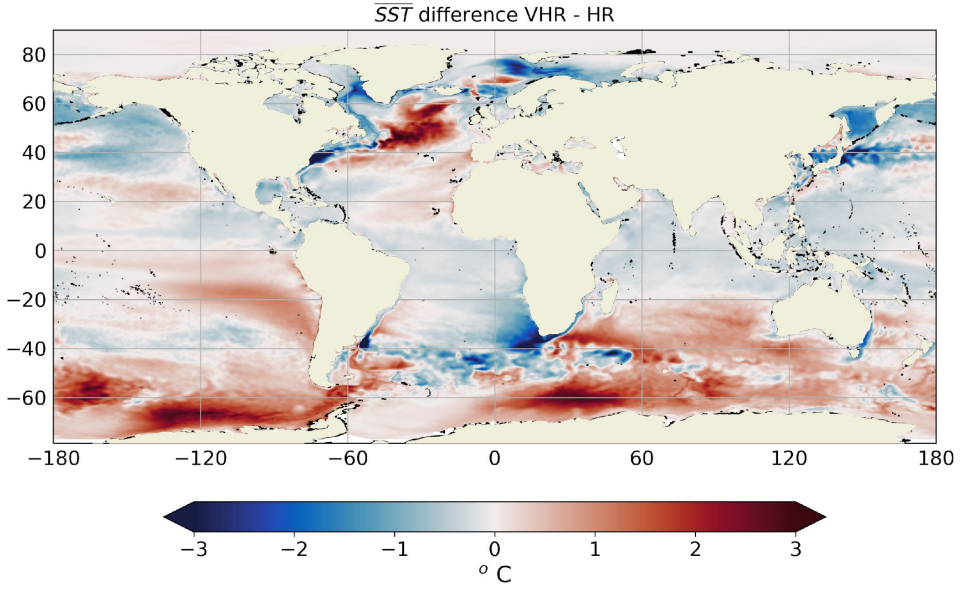

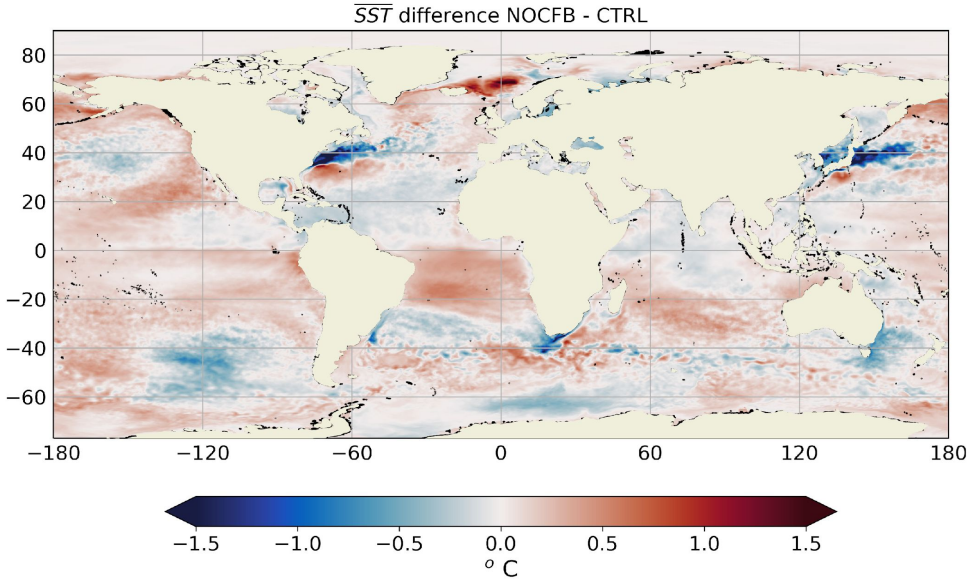

Courtesy of: Thomas Arsouze

EC-Earth 3 - T1279-ORCA12: production runs

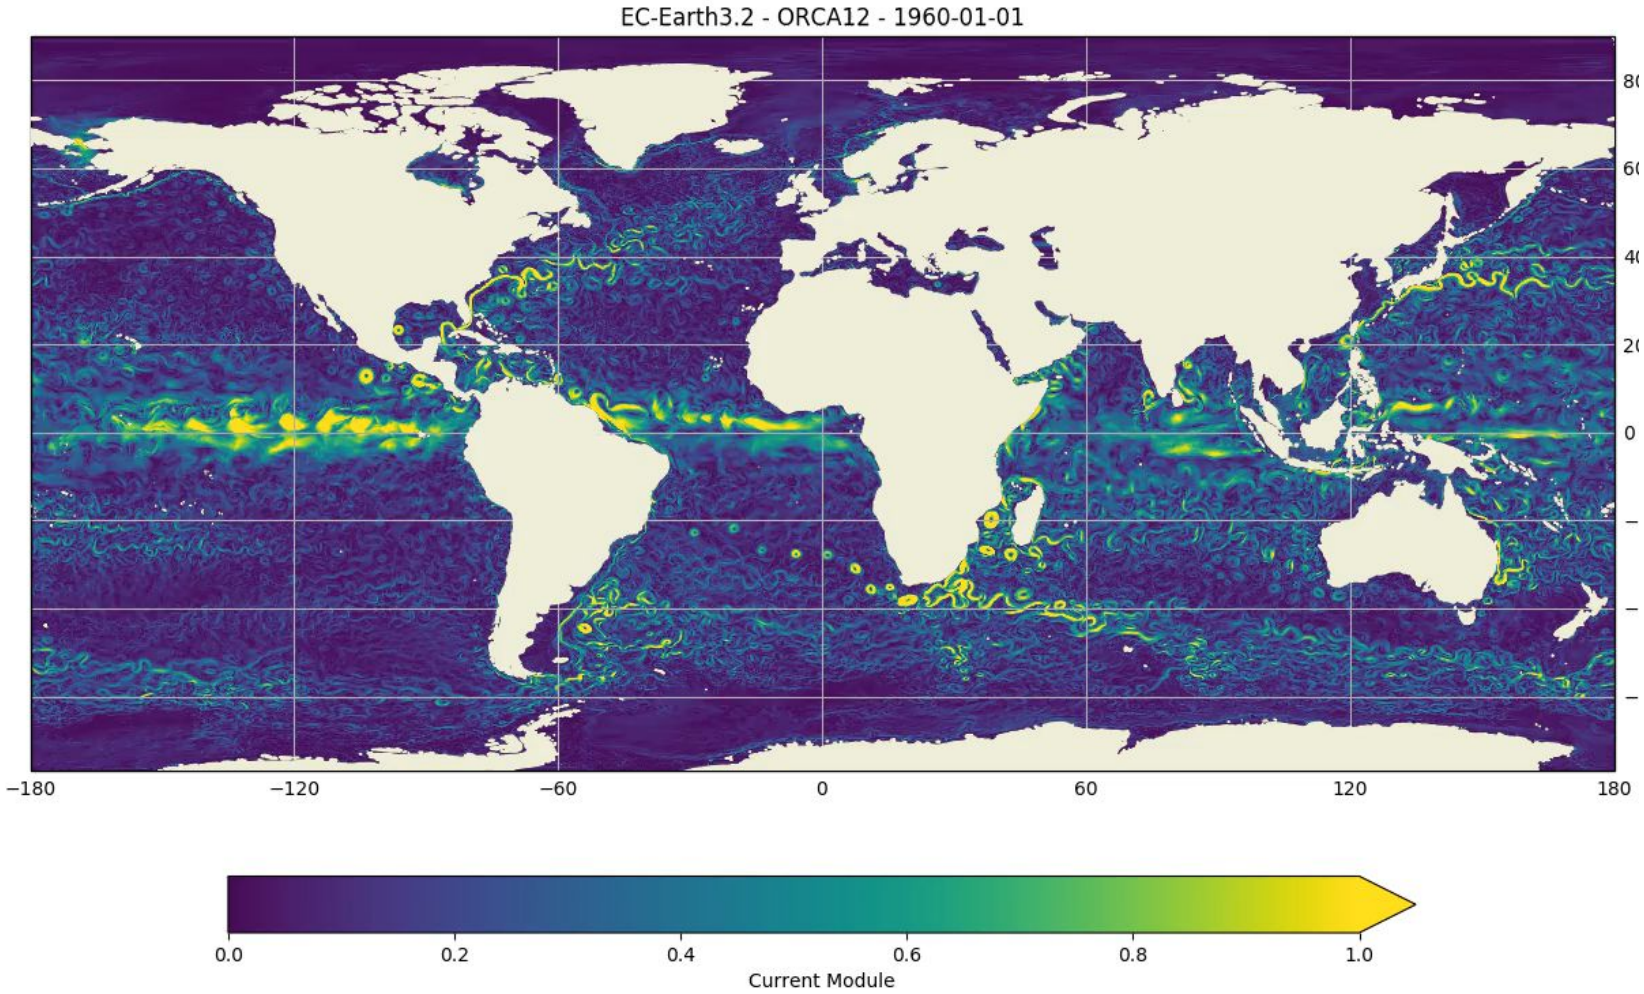

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Courtesy of: Thomas Arsouze

Mixed precision calculations in NEMO

In a nutshell

- Different computations might require a **different** level of **precision**.
- Customized precision has emerged as a **promising approach** to improve power / performance trade-offs.
- We developed a **method** useful to **determine** the **precision needed** for the different variables in a model.
- Branch in NEMO 2020 WP ⟶ Aiming for merge in **NEMO 4.2**.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

2019: "How to use mixed precision in ocean models: exploring a potential reduction of numerical precision in NEMO 4.0 and ROMS 3.6."
Tintó Prims, O., Acosta, M.C., Moore, A.M., Castrillo,M., K. Serradell, A. Cortés and F.J. Doblas-Reyes

IMMERSE: Porting diagnostics to GPUs

The diagnostics dia_hsb kernel

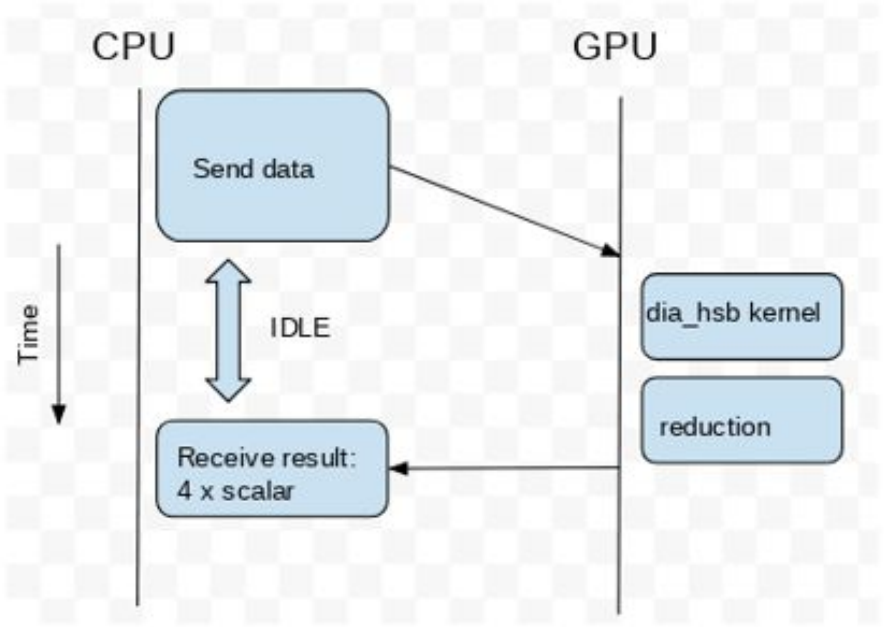

DIA_HSB diagnostics time, using one Power9 node