



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación

HPC adaptation and developments for EC-Earth4 at BSC

Xavier Yepes-Arbós, Mario Acosta, Marta Alerany, Marcel Cases, Maitane Fariñas, Kai Keller, Sergi Palomas, Roc Salvador, Oriol Tintó

24 February 2026

GANANA C-DAC Colloquium, Pune, India

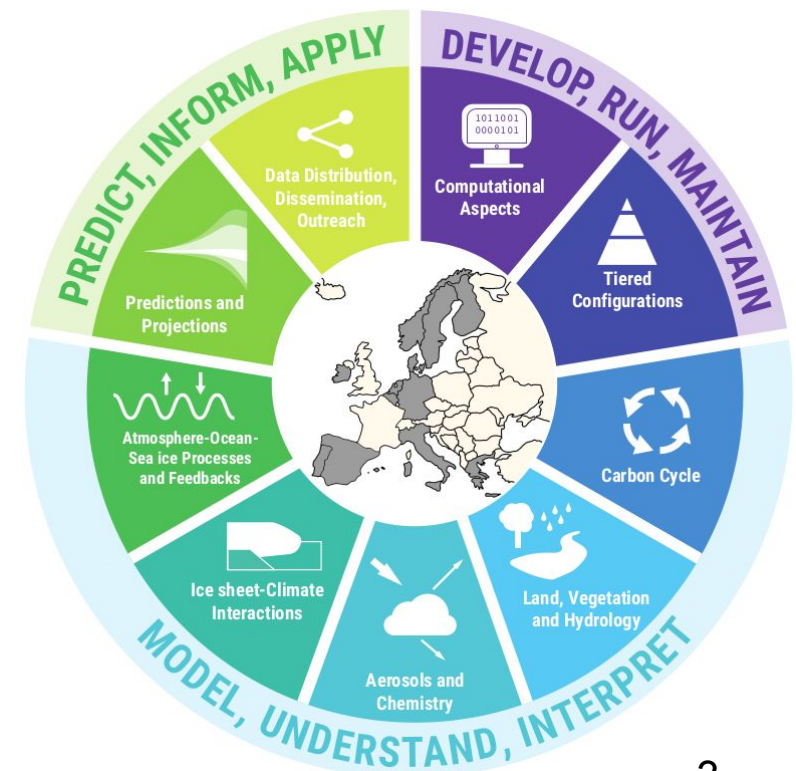
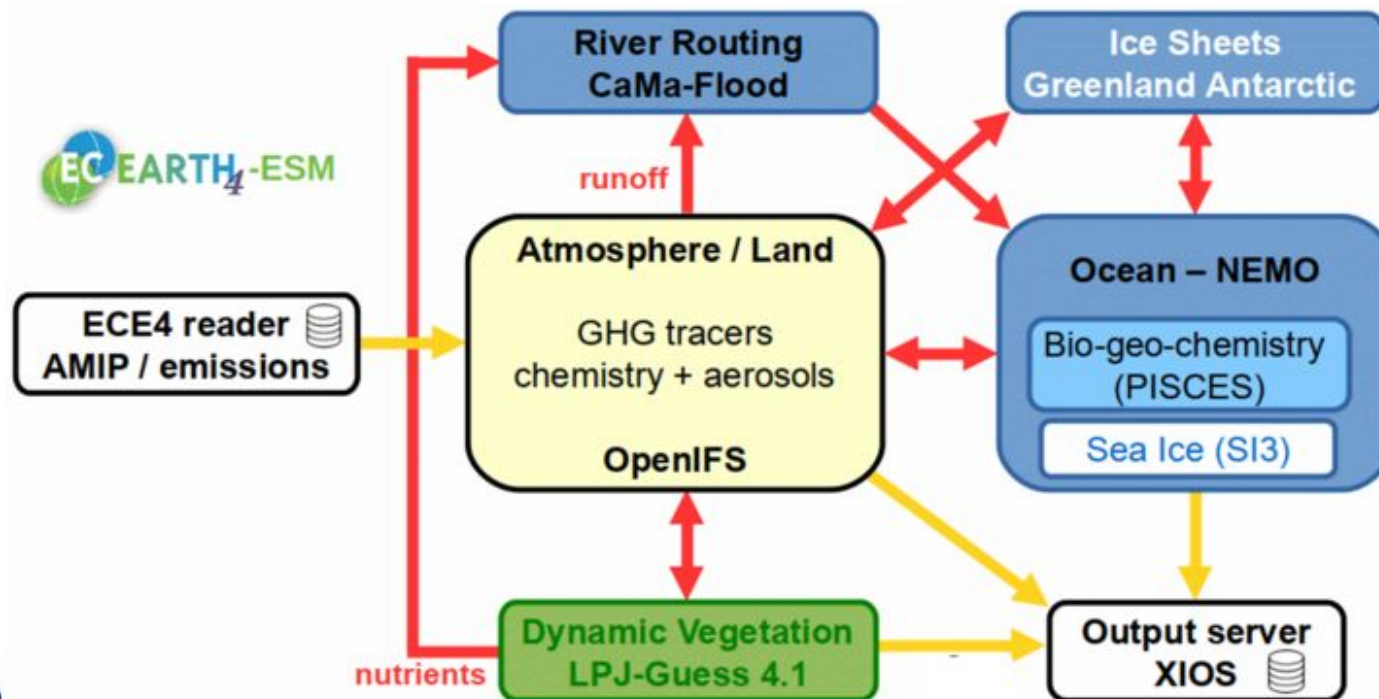
HPC for Earth Sciences Team at BSC

- Knowledge about:
 - Mathematical and computational aspects of Earth system models (ESMs)
 - HPC adaptation and development
- Support on computational performance analysis and optimization
- Research on different HPC topics, including GPUs, reproducibility, RISC-V co-design, etc.



EC-Earth 4

- EC-Earth 4 is a next-generation European Earth-system model that integrates atmosphere, ocean, land, and biogeochemical components to produce more realistic and flexible climate simulations.
- The purpose is to simulate the key physical, chemical, and biological processes of the Earth system in order to better understand climate variability and improve predictions of future climate change.



The workflow manager



Autosubmit is a lightweight workflow manager designed to meet climate research necessities. Unlike other workflow solutions in the domain, it integrates the capabilities of an experiment manager, workflow orchestrator and monitor in a self-contained application.

Automatization

Meta-scheduling

Task dependencies

High-level config

Automatic retries

Interoperability

Multi-platform

Single point of access

Python

Web GUI

Efficiency

Custom granularity

Dyn. task aggregation

Performance metrics



Robustness

Scalable database

Manage multiple hosts

Auto-recovery

Traceability

Workflow database

Manage multiple hosts

Auto-recovery

Monitoring

Real-time status

Workflow statistics

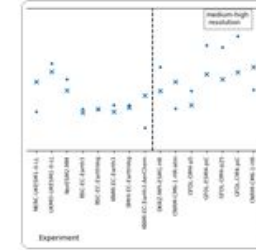
Task logs



CPMIP: The impact of EC-Earth

The computational and energy cost of simulation and storage for climate science: lessons from CMIP6

Mario C. Acosta ✉, Sergi Palomas ✉, Stella V. Paronuzzi Ticco, Gladys Utrera, Joachim Biercamp, Pierre-Antoine Bretonniere, Reinhard Budich, Miguel Castrillo, Arnaud Caubel, Francisco Doblas-Reyes, Italo Epicoco, Uwe Fladrich, Sylvie Joussaume, Alok Kumar Gupta, Bryan Lawrence, Philippe Le Sager, Grenville Lister, Marie-Pierre Moine, Jean-Christophe Rioual, Sophie Valcke, Niki Zadeh, and Venkatramani Balaji



Scan for DOI!



SMHI

Years Simulated: **38 854**

Core-hours: **46,36 * 10⁶**

Energy: **1,24 * 10¹² J**

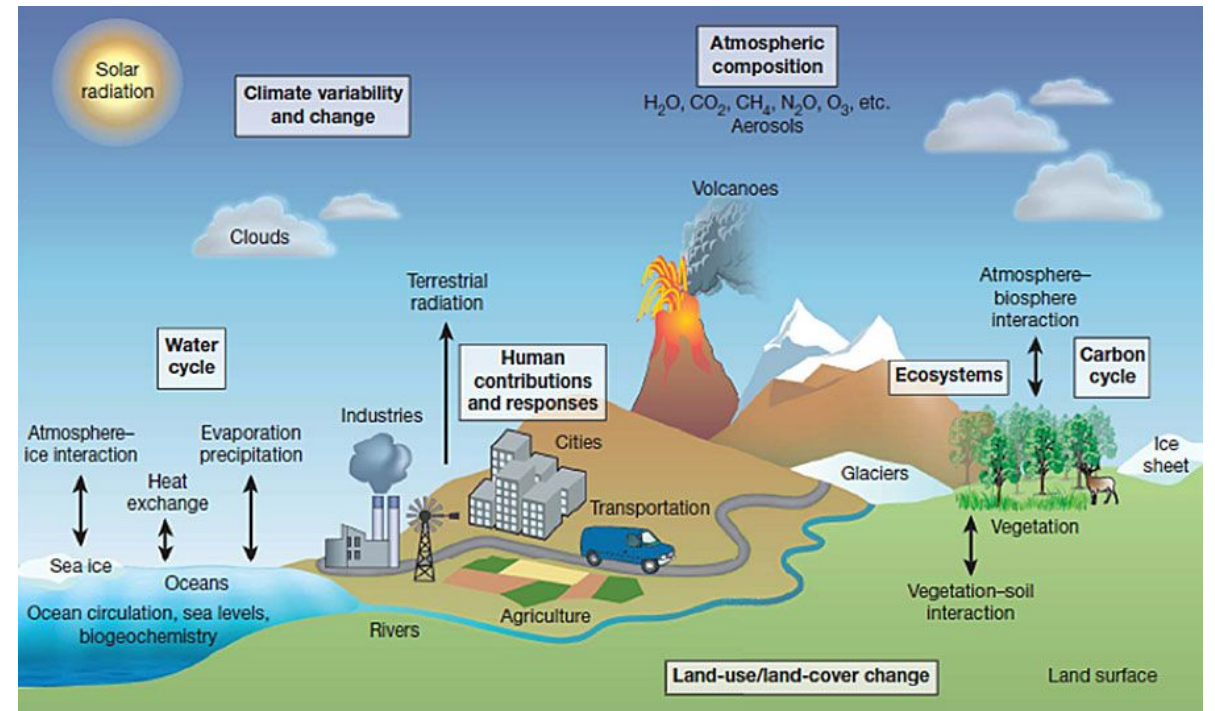
Carbon Footprint*: **165 tons of CO₂**

CPMIP: What is it?

Provide **standardized metrics** to assess the computational performance of climate models (ESMs).

Identifying efficiency, energy costs of ESMs.

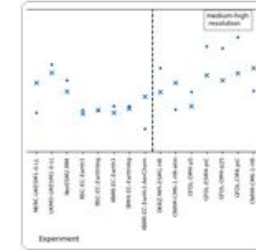
Make **climate science** more **efficient, sustainable, and globally comparable**.



CPMIP: Why automate?

The computational and energy cost of simulation and storage for climate science: lessons from CMIP6

Mario C. Acosta ✉, Sergi Palomas ✉, Stella V. Paronuzzi Ticco, Gladys Utrera, Joachim Biercamp, Pierre-Antoine Bretonniere, Reinhard Budich, Miguel Castrillo, Arnaud Caubel, Francisco Doblaz-Reyes, Italo Epicoco, Uwe Fladrich, Sylvie Joussaume, Alok Kumar Gupta, Bryan Lawrence, Philippe Le Sager, Grenville Lister, Marie-Pierre Moine, Jean-Christophe Rioual, Sophie Valcke, Niki Zadeh, and Venkatramani Balaji



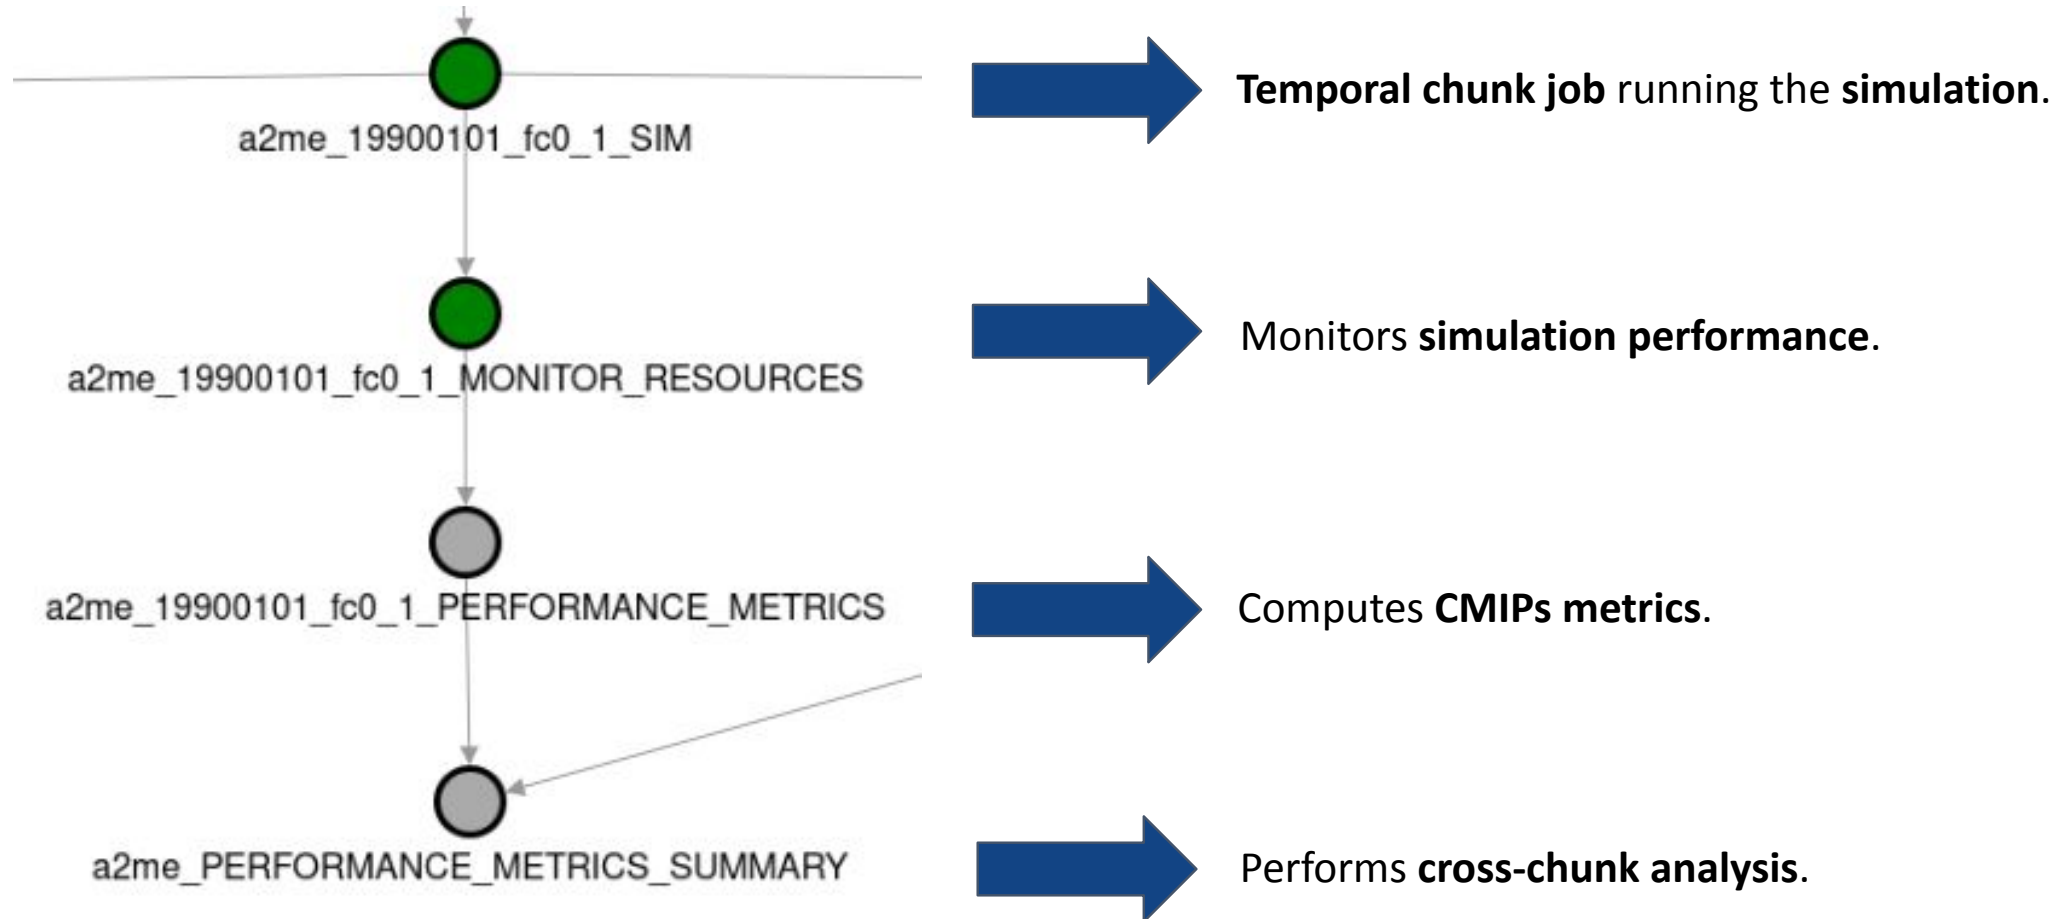
Drawbacks

Lack of resources and time to collect metrics after CMIP experiments.

Non-tech people have to learn the definition of the metrics.


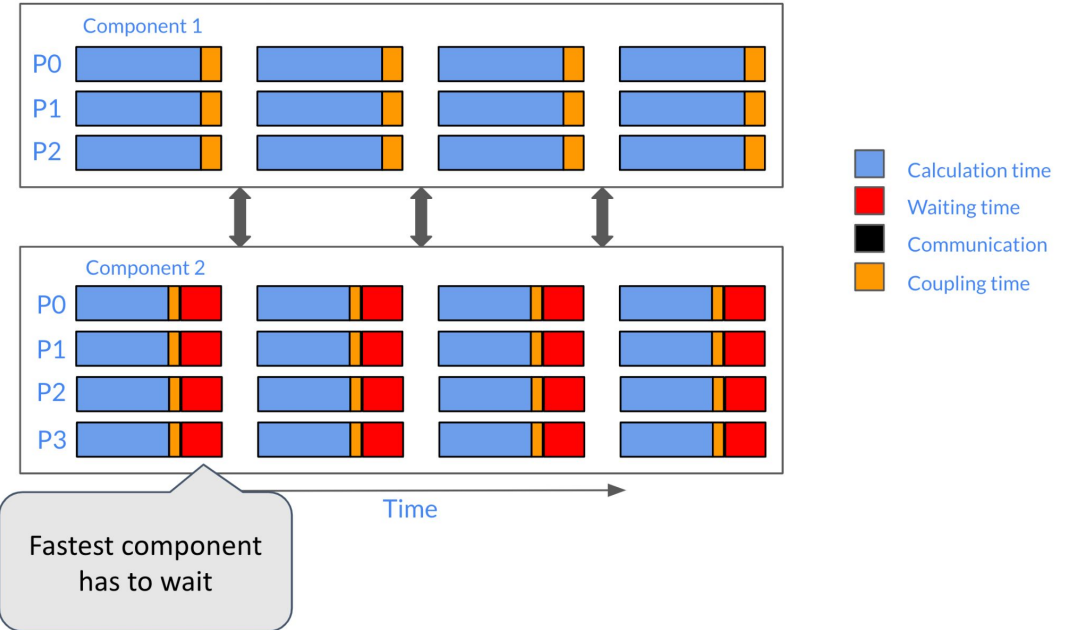
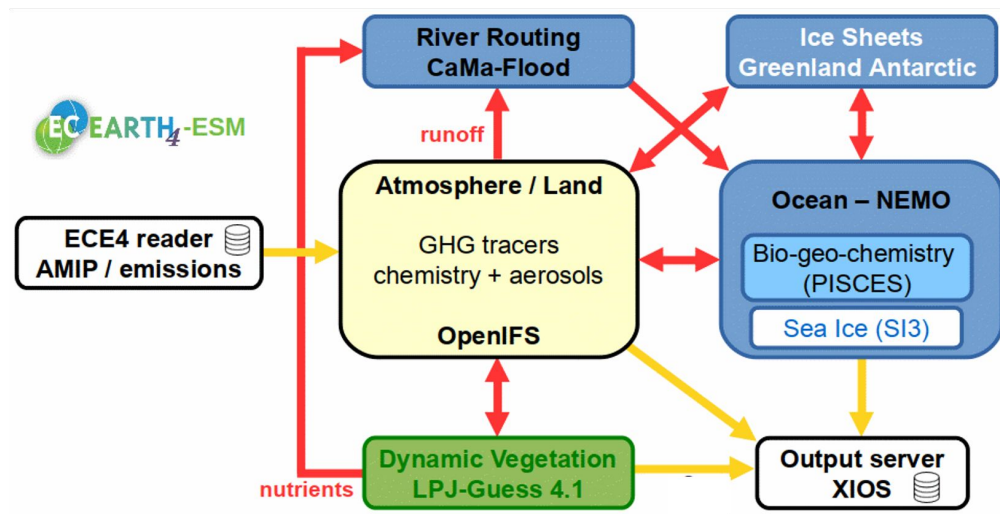
Errors and differences in the collection of the performance metrics.

CPMIP: How to Automate?



Load-balance for Coupled ESMs

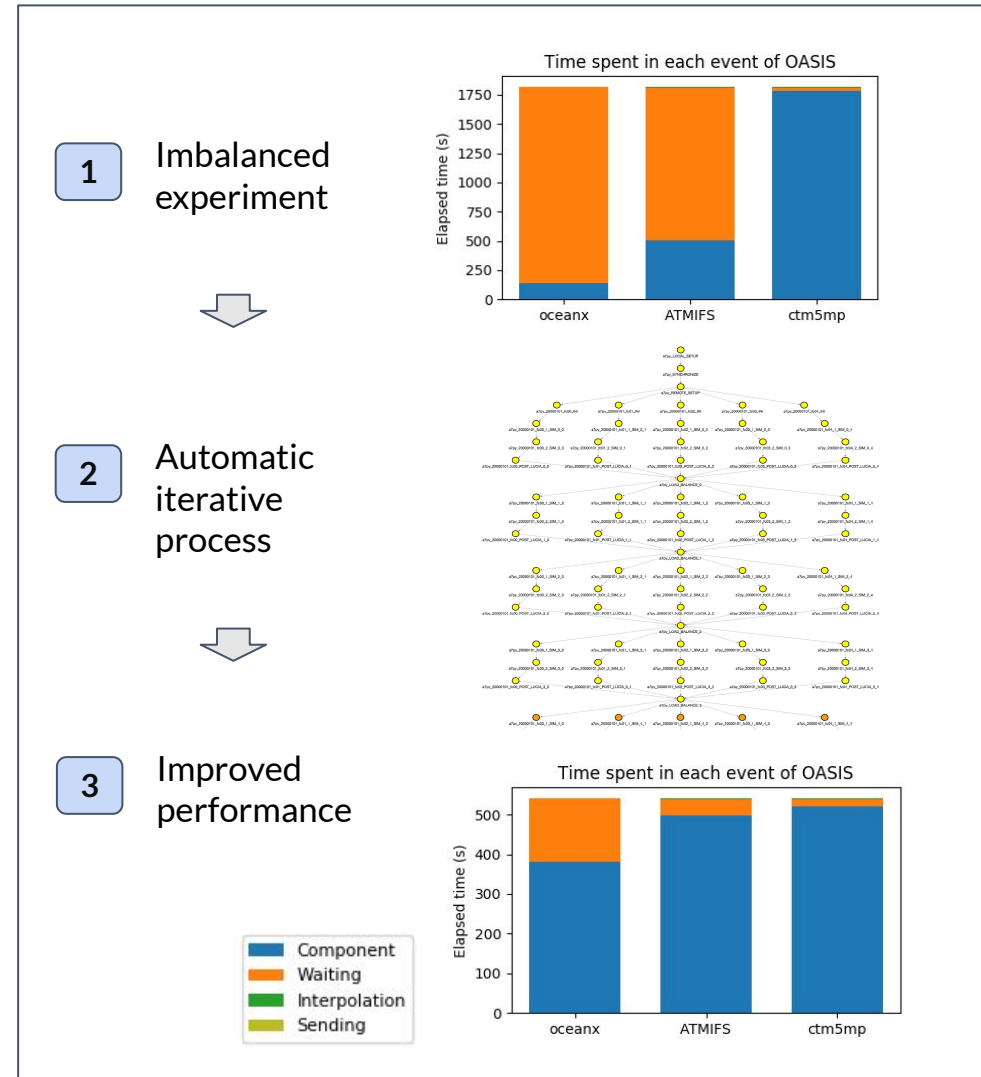
Load-balance is one key limiting factors of coupled ESMs performance.
Average CMIP6 12.8% ([Acosta et. al. 2024](#))



Load-balance for Coupled ESMs

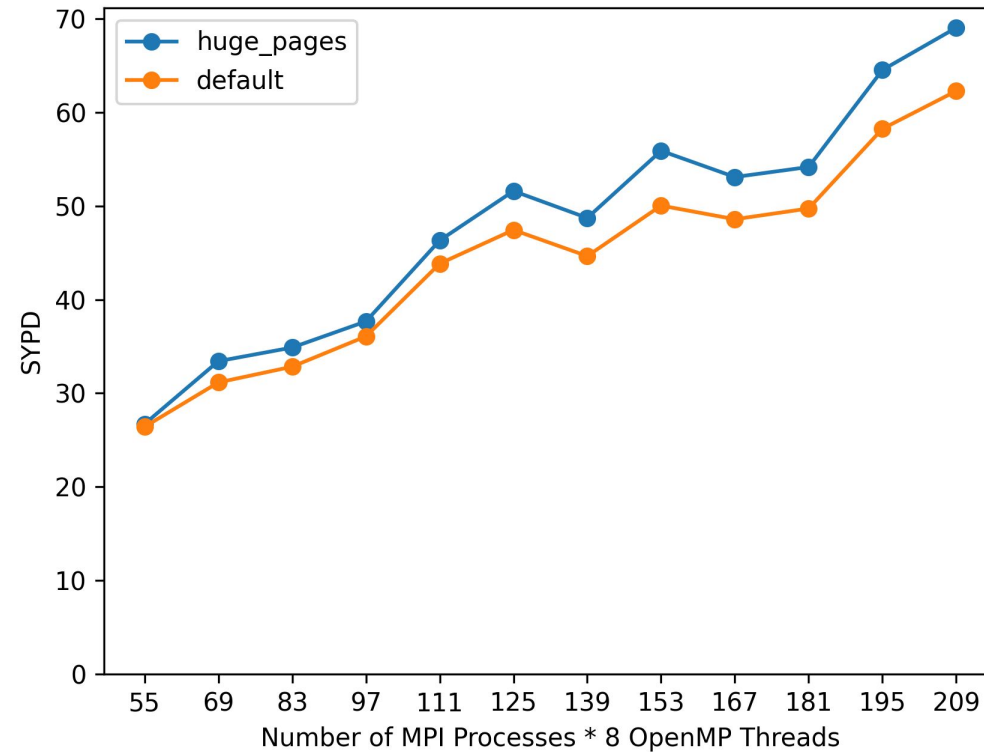
Auto Load-Balance Workflow

- A tool to automatically minimize the performance loss on the coupling for ESMs by:
 - Minimizing synchronization cost
 - Finding a **good resource configuration** for the components based on specified Energy-to-Solution / Time-to-Solution criteria
- Easy to adopt by the climate community → integrated into climate workflows, directly compatible with OASIS Coupler
- Tested **portability** to HPC platforms (HPC2020, MeluXina, MareNostrum 4, MareNostrum 5)



Profiling and optimization: TBB Huge Pages on MN5

Average performance improvement using Huge Pages: **+7,09%**

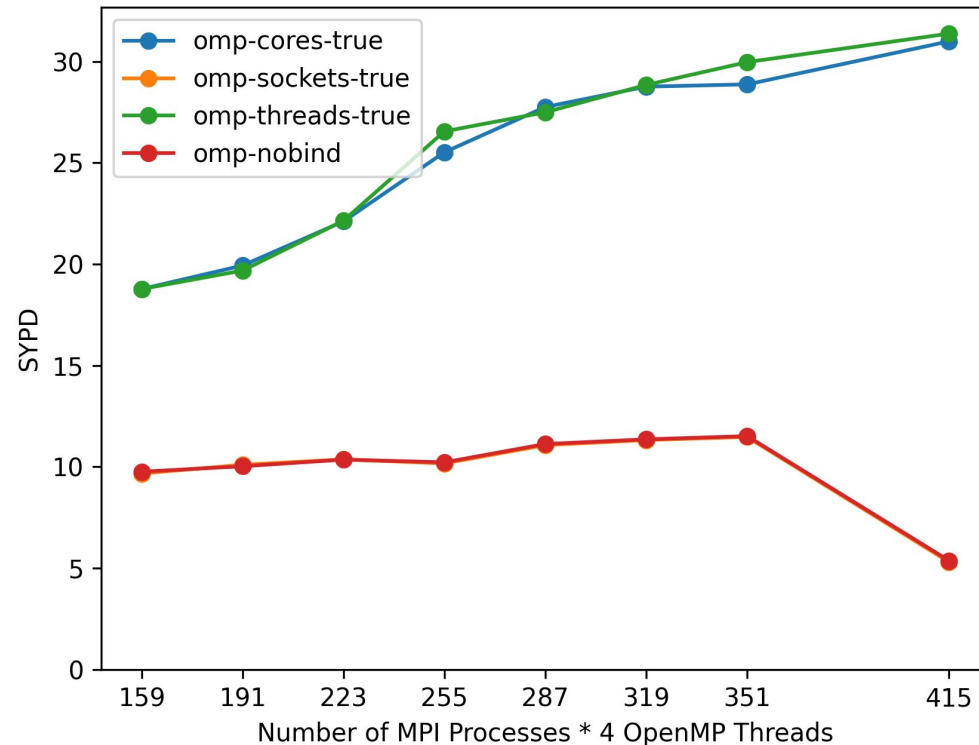


OIFS TL255 scalability in MN5 using HP export vs not using it

Profiling and optimization: OpenMP affinity on HPC2020

Degradation of performance on some OMP binding settings.


OpenMP Threads binding is already the default in EC-Earth4 and provides best in performance together with the Cores binding.




OIFS TL255 scalability in HPC2020 using different OMP binding settings

Profiling and optimization: Tracing

Preliminary profiling of EC-Earth 4 on MN5 running TL255L91 and eORCA1.



- Outside MPI
- MPI_Send
- MPI_Recv
- MPI_Isend
- MPI_Irecv
- MPI_Wait
- MPI_Waitall
- MPI_Bcast
- MPI_Barrier
- MPI_Allreduce
- MPI_Alltoallv
- MPI_Gatherv
- MPI_Comm_rank
- MPI_Comm_size
- MPI_Bsend
- MPI_Test
- MPI_Waitany
- MPI_Iallreduce



APP: Why?

Each time there is a change to a model, we have to run the same tests to check its performance.

The **Automatic Performance Profiling (APP)** aims to automate this repetitive tests, allowing developers to know what is the performance and if there is any issue know where to look.

It is meant for any user since it reports basic performance metrics (like SYDP) and also advanced metrics (like PAPI, MPI...) for performance experts.

APP: Overview

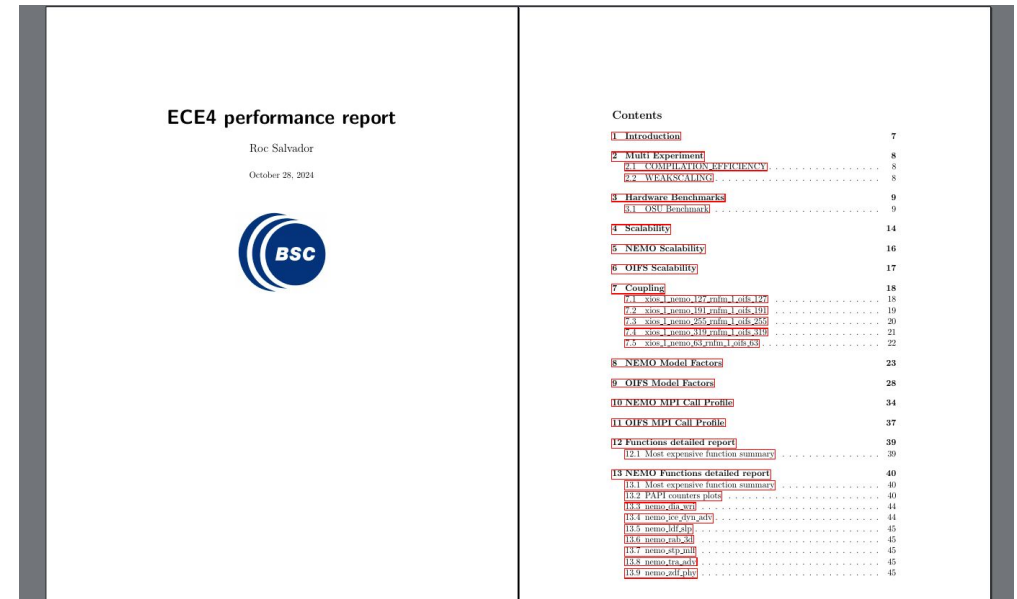
APP is implemented using Autosubmit 4 as a workflow manager to be user-friendly since it has features like GUI support.

It supports multiple models:

- EC-Earth 4
- NEMO 4.2.1
- HPCW 3.0: NEMO 5 benchmark
- MONARCH (in development)

Portable across HPC platforms, currently:

- MN5
- HPC2020



The image shows a preview of a performance report. The left side displays the cover page for 'ECE4 performance report' by Roe Salvador, dated October 28, 2024, featuring the BSC logo. The right side shows a detailed table of contents with page numbers.

Contents	
1 Introduction	7
2 Multi Experiment	8
2.1 COMPUTATION EFFICIENCY	8
2.2 WFASSCAN/CC	8
3 Hardware Benchmarks	9
3.1 CSI Benchmark	9
4 Scalability	14
5 NEMO Scalability	16
6 OIFS Scalability	17
7 Coupling	18
7.1 _sso_ _nemo_02_min_1_016_147	18
7.2 _sso_ _nemo_01_min_1_016_101	19
7.3 _sso_ _nemo_05_min_1_016_205	19
7.4 _sso_ _nemo_01_min_1_016_310	21
7.5 _sso_ _nemo_03_min_1_016_03	22
8 NEMO Model Factors	23
9 OIFS Model Factors	28
10 NEMO MPI Call Profile	34
11 OIFS MPI Call Profile	37
12 Functions detailed report	39
12.1 Most expensive function summary	39
13 NEMO Functions detailed report	40
13.1 Most expensive function summary	40
13.2 PAPI counters list	40
13.3 nemo_01a_w3	44
13.4 nemo_02_01v_015	44
13.5 nemo_01a_w3	45
13.6 nemo_01a_w3	45
13.7 nemo_01a_w3	45
13.8 nemo_01a_w3	45
13.9 nemo_01a_w3	45

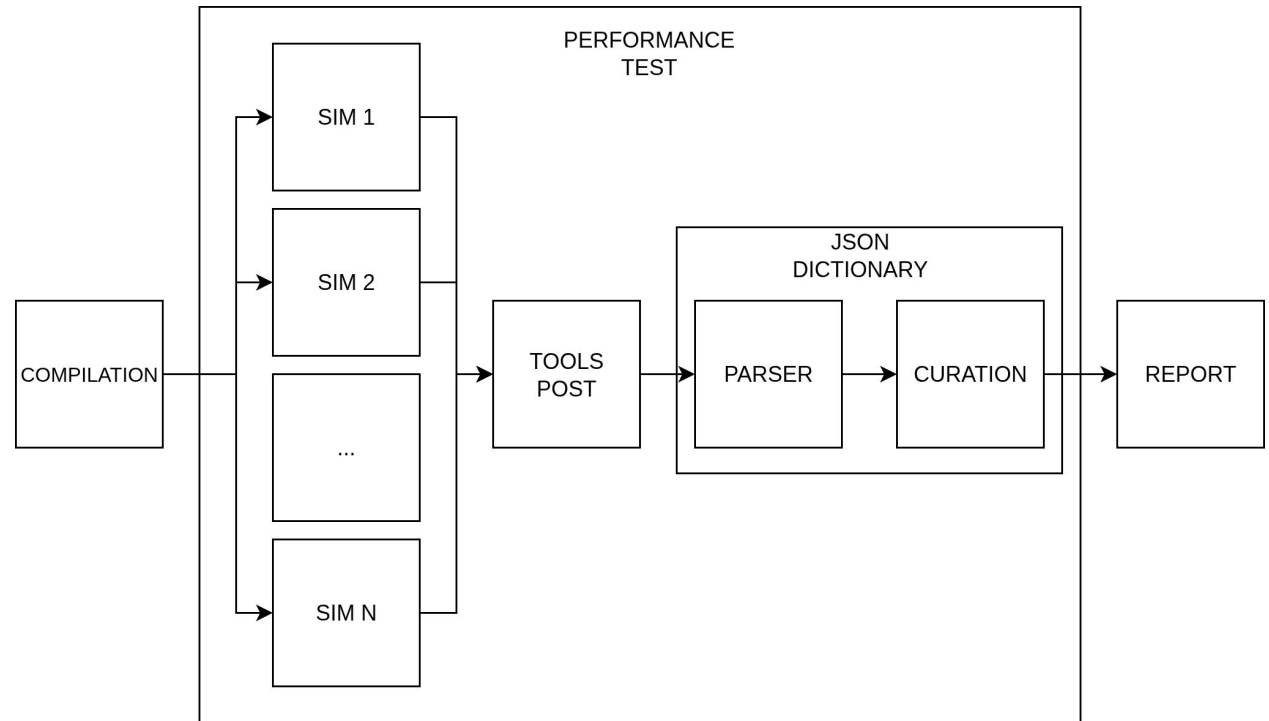
Check example reports [here](#)

APP: Performance test types

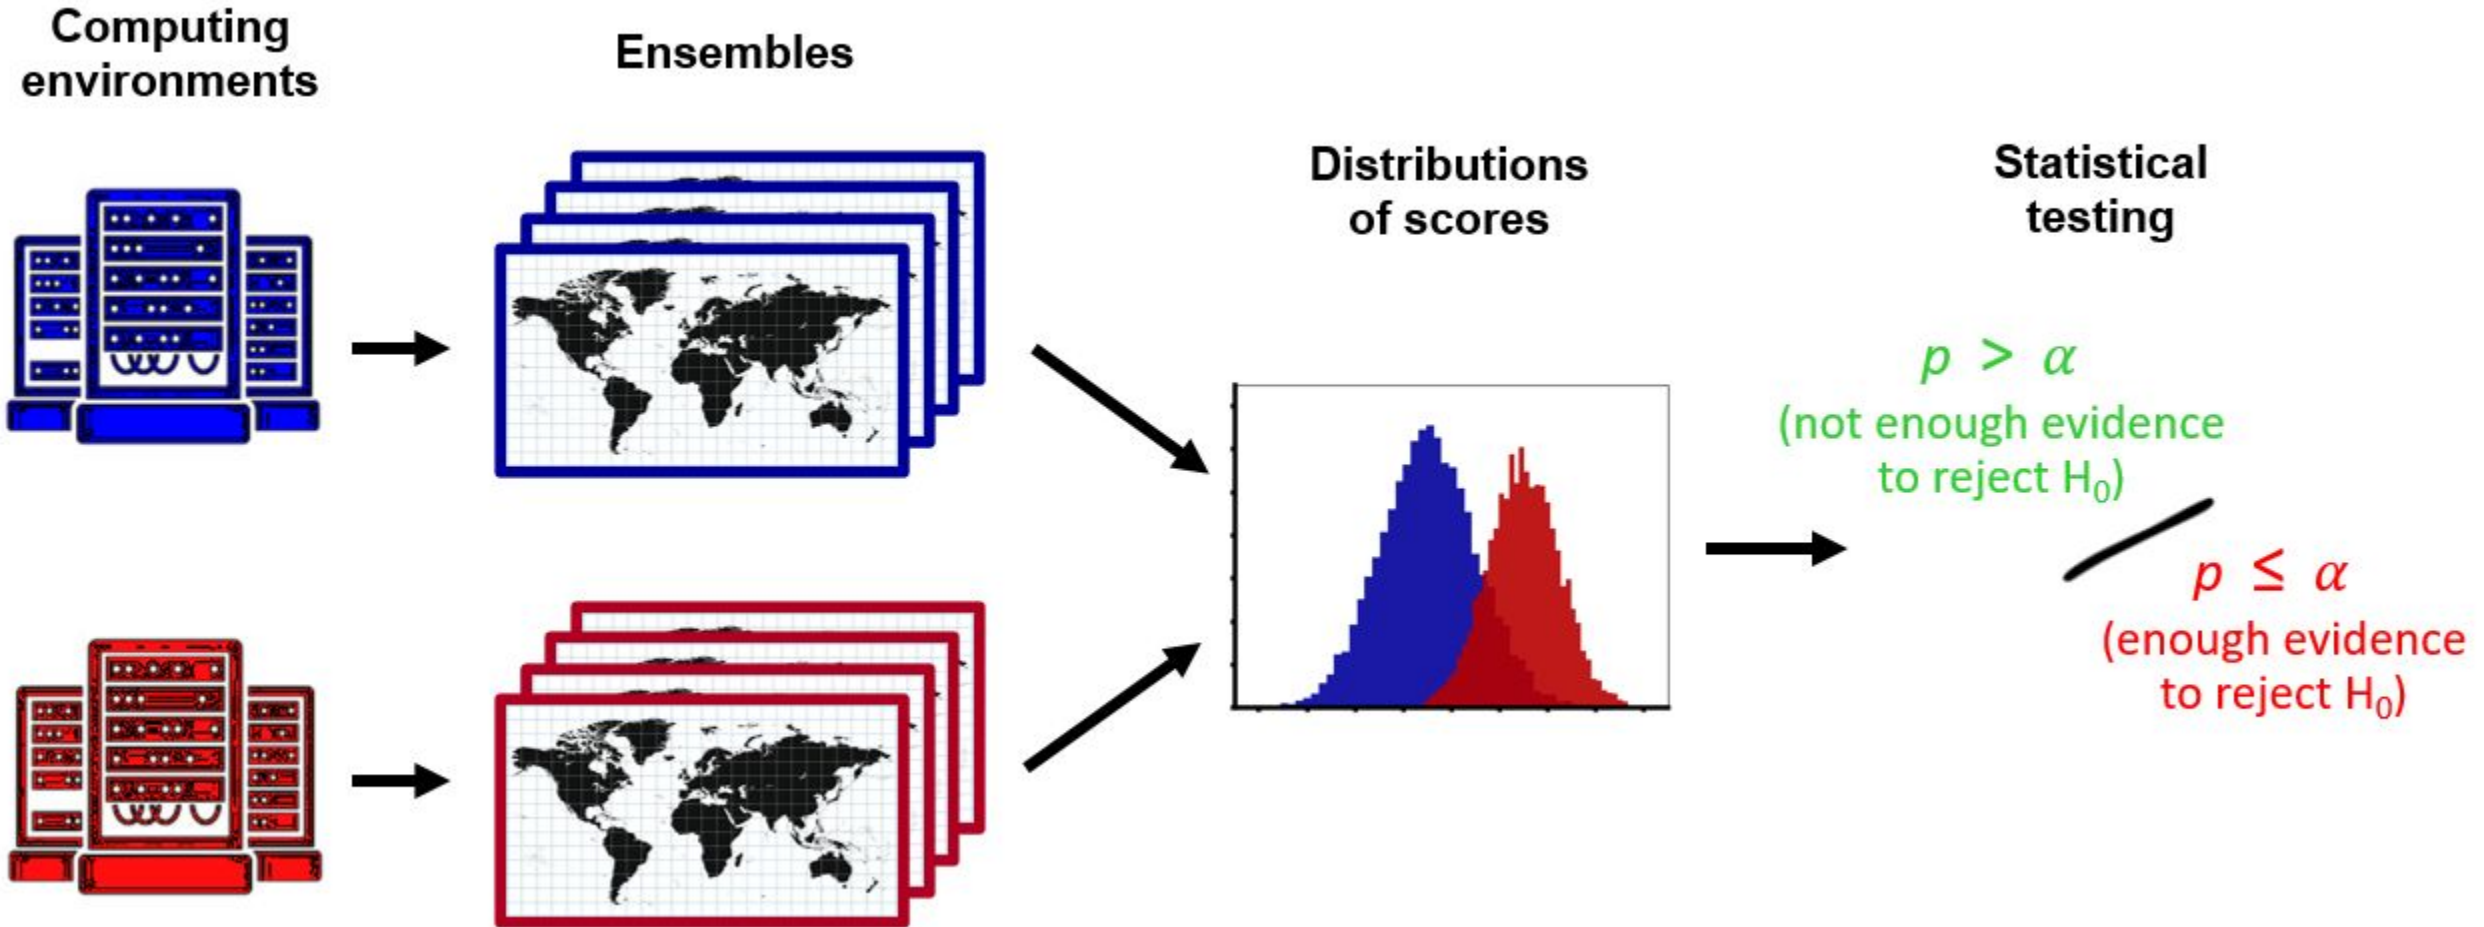
The APP's workflow is divided in 5 different tests, that provide different insights on the model's performance.

These tests are:

- Hardware Benchmarks
- Compilation Efficiency
- Weak Scaling
- Strong Scaling
- Tracing



Climate Model Replicability



Climate Model Replicability

Improved methodology, based on Massonnet et al., 2022 (<https://doi.org/10.5194/gmd-13-1165-2020>)

Metrics

- **RK08 score**

$$eRK08_m^X = \sqrt{\sum_{i=1}^N \tilde{\omega}_i \cdot \exp \left[- \left(\frac{x_{mi} - \bar{y}_i}{\sigma_i^y} \right)^2 \right]}.$$

- **Bias score**

$$eBIAS_m^X = \sum_{i=1}^N \tilde{\omega}_i \cdot e^{-|\text{bias}_{mi}^X|/\sigma_i^y}.$$

- **RMSE score**

$$eRMSE_m^X = \sum_{i=1}^N \tilde{\omega}_i \cdot e^{-\text{crmse}_{mi}^X/\sigma_i^y}$$

- **combined score**

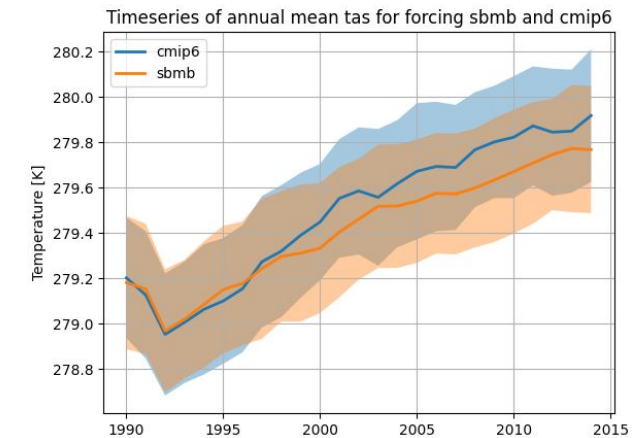
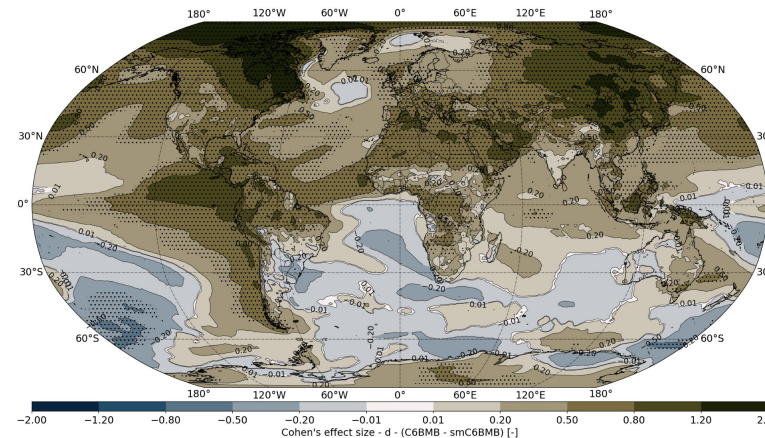
$$\text{COMBINED} = \frac{eRK08_m^X + eBIAS_m^X + eRMSE_m^X}{3}$$

Tests


- **T-test (Welsh's t-test)**
- **U-test (Mann–Whitney test)**
- **Bootstrap test**
- **Kolmogorov-Smirnov test**

Reports

- **Time series plots (ensemble mean and spread for both ensembles)**
- **Spatial differences plots**
- **Matrix plot (see next slide)**



Climate Model Replicability: Matrix Plot





**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación

Thank you



The GANANA project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement No. 101196247. The views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the granting authority can be held responsible for them.

xavier.yepes@bsc.es