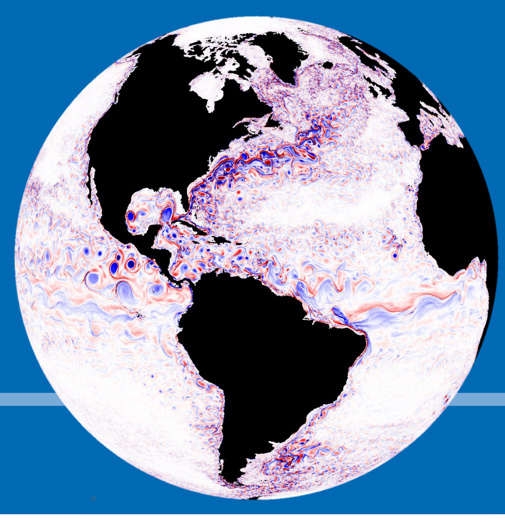

Performance study of Earth System Models

Mario C. Acosta, Xavier Yepes, Oriol Tintó, Miguel Castrillo, Alicia Sánchez and Kim Serradell

Earth Sciences Department, Barcelona Supercomputing Center (BSC), Spain

Motivation

Global **ESMs are complex** systems made of heterogeneous components, that require a substantial amount of effort to work correctly and use the resources wisely. The goal of running these models using grids on the order of 1 km will be impossible to achieve without using adequate optimization techniques.

Earth System Models (ESMs) consume billions of computing hours every year. Our motivation is to optimize these models to **save both time and energy**, and **allow further increases in resolution and complexity** to improve their capabilities of simulating small-scale features and reduce the impact of the parametrizations.

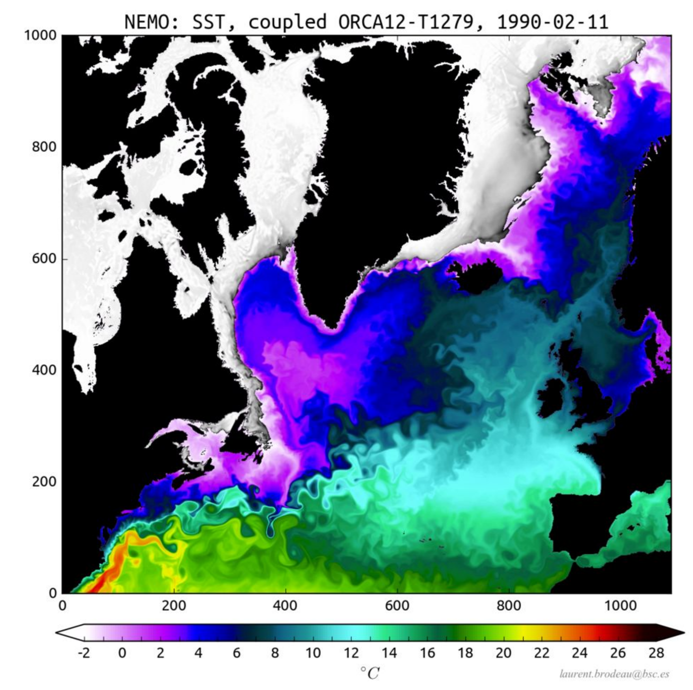

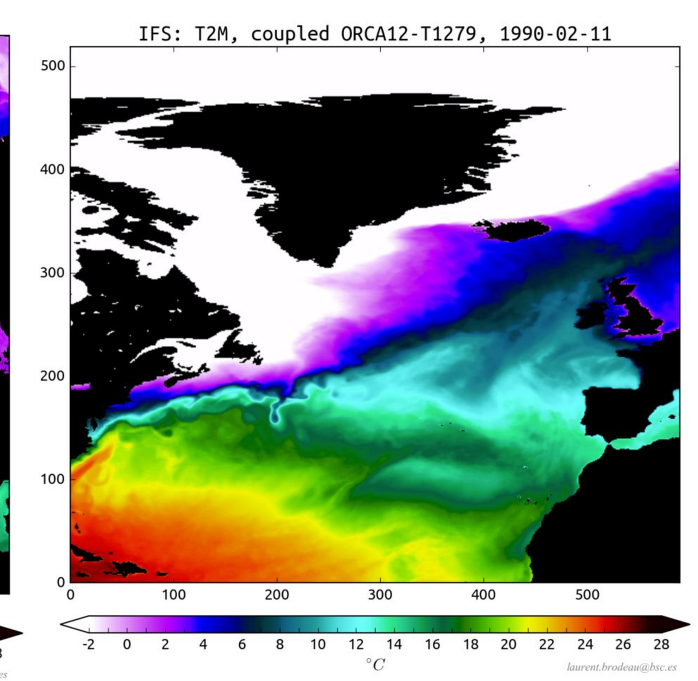

Daily mean of: sea surface temperature (NEMO, left) and 2m air temperature (IFS, right) from a ORCA12-T1279 EC-Earth coupled simulation. The influence of the sea surface temperature (cold eddy detaching from the Gulf Stream) on lowest layer of the atmosphere can be clearly observed.

Methodology to evaluate Earth System Models

Complex models require complex studies. Several techniques are needed to optimize an ESM:

Mathematical and Computational study

Understand the algorithms used and the overhead or drawbacks introduced by the parallel implementation.

Scalability Study

The evaluation of an ESM requires deep knowledge about the complex workflows and experiments done [1].

Profiling Study

General and oriented profiling analysis for ESMs [2].

Introducing optimizations

Computational improvements, optimizing or introducing new approaches [3].

Reproducibility study

Keeping in mind the chaotic nature of climate models, evaluate the impact in accuracy and reproducibility of the model after applying the proposed optimizations [4].

Scalability Study: How efficient is your model?

Climate predictions have complex workflows. **New metrics are needed to evaluate the computational efficiency.**

Metric	used to evaluate ...
Simulation Year Per Day (SYPD)	how efficient is your sim job per each year of the simulation
Core-hours Per Year (CHPY)	how efficient is your sim job with respect to the number of parallel resources used
Queue Time	how much time are all the jobs waiting in the queue
Actual SYPD	how affect queue time to the complete experiment, from the first to the last sim job
Energy Cost Per Year (JCPY)	how much energy is needed per each year of simulation
Memory Bloat	how much memory is needed for the execution of the sim job
Data Output Cost	how much time is used for the output post-processing, only done for specific time steps
Coupling Cost	how much time is wasted in waiting time, produced by irregular executions

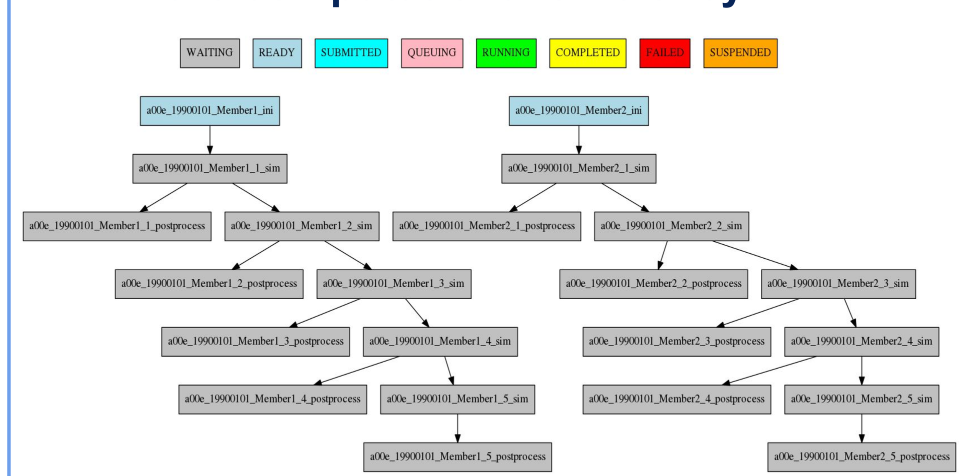

- The simulation does **not only involve the execution of a model during a sequence of time steps** (represented by the sim job).
- The experiment **adds complexity in the horizontal for ensembles** (members with perturbations in the initial conditions).
- The experiment **adds complexity in the vertical, running long simulations** divided into chunks and including **pre- and post-processing.**

Why do we need to evaluate Output and Coupling cost?

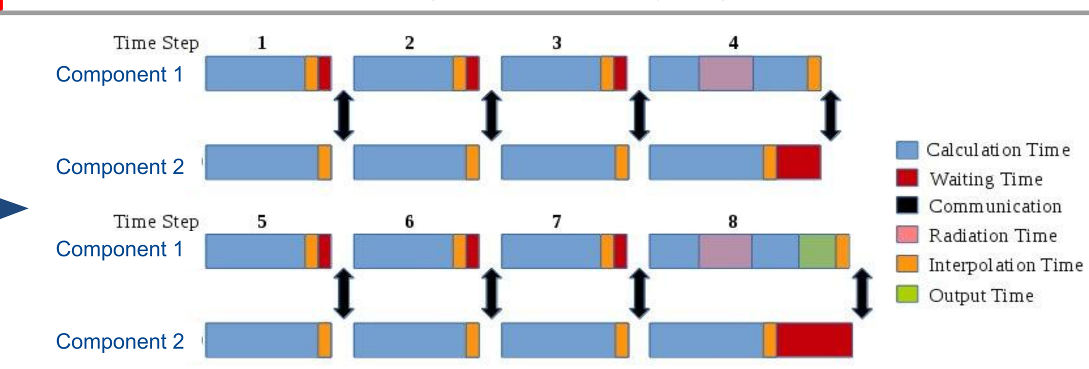

- The execution of a coupled model is complex. **Different components run in parallel**, exchanging some information and adding some extra overhead (Communication and Interpolation Time). **Load imbalances produce extra waiting time as well**.
- The execution of **each component could involve some irregular extra overhead**, such as output processing or calculations occasionally done, such as radiation.

Profiling Study: What can you see in a trace?

Performance tools [5] are essential to study the behavior of ESMs:

- **Extrae**: is a package used to instrument the code. It generates trace-files with hardware counters, MPI messages and other information.
- **Paraver**: is a browser used to analyze both visually and analytically trace-files.
- **Dimemas, Clustering, Folding...**: Other tools to evaluate ideal conditions, gather performance data...

-The view of a trace consists of **threads or processes** in the Y axis and the **timeline** on the X axis.

-The base trace (figure on the right) shows **MPI functions**, where each type of call is identified by a color. Blue color represents a computation area.

-Other traces can contain **PAPI events to collect information regarding the microprocessor performance.** Colors represent in this case a gradient between maximum and minimum values.

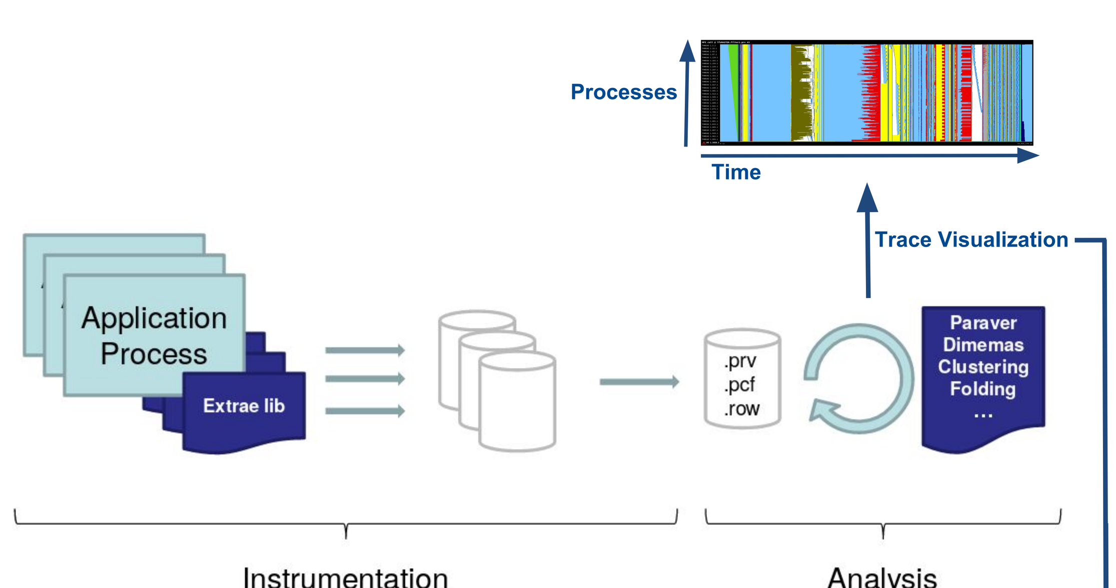

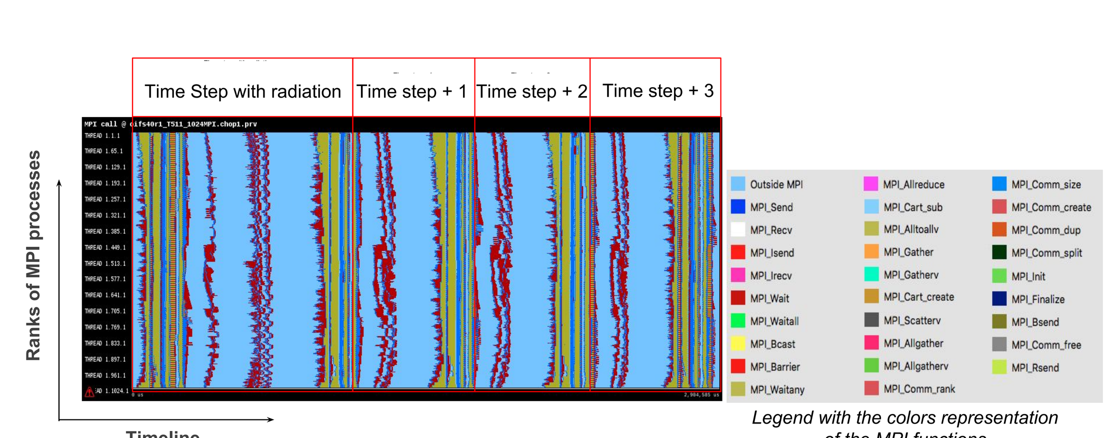

Base trace with MPI events for four time steps of IFS using 512 processes.

Some results that you can obtain from the analysis

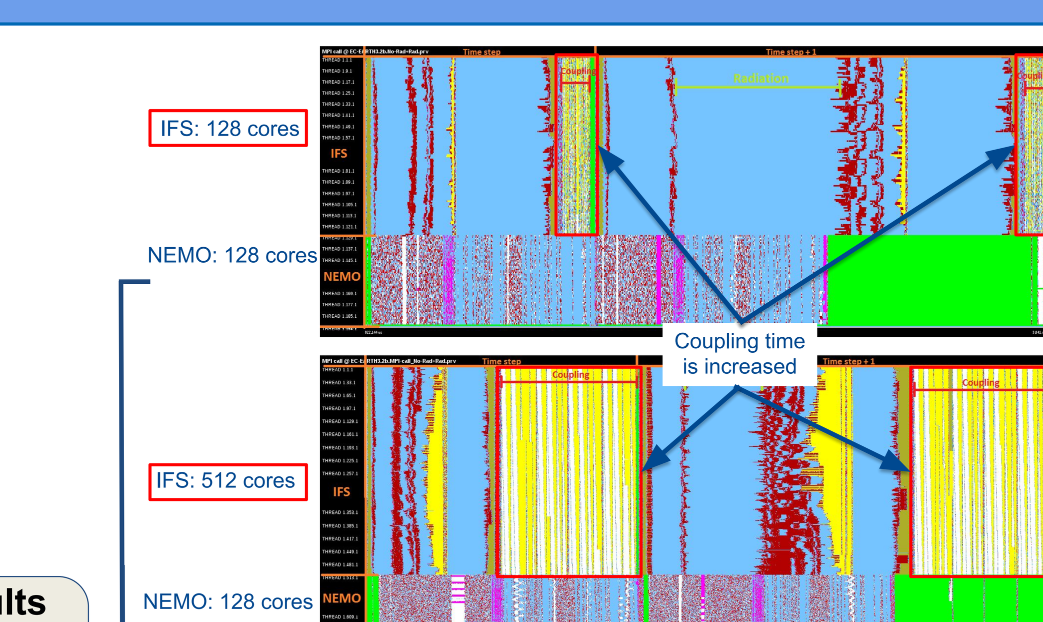

- Coupling time dramatically increases when more parallel resources are used.
- This means that there is a bottleneck in the coupling process due to sequential work in the master process.
- This analysis proves that this configuration is not scalable.

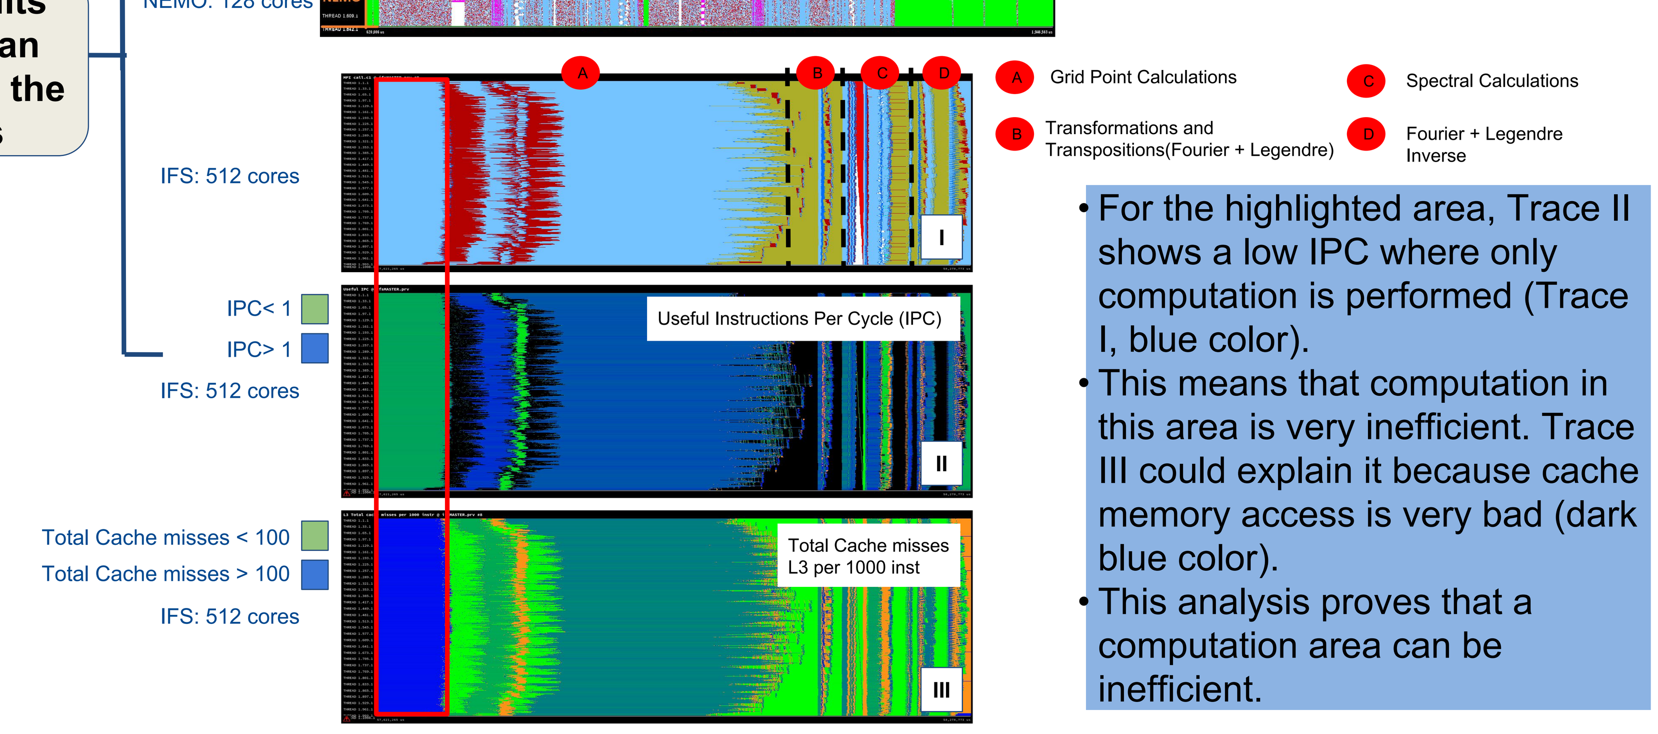

- For the highlighted area, Trace II shows a low IPC where only computation is performed (Trace I, blue color).
- This means that computation in this area is very inefficient. Trace III could explain it because cache memory access is very bad (dark blue color).
- This analysis proves that a computation area can be inefficient.

About us

The **Performance Team** is part of the **Computational Earth Sciences** Group (CES), at the **BSC Earth Sciences Department**.

The mission of **CES** is to ensure the efficient use of the computational resources by the Earth scientists, to develop Earth Sciences related HPC user-friendly software and to provide computational services to the department and the Earth Sciences community.

Kim Serradell | Mario C. Acosta | Xavier Yepes | Miguel Castrillo | Oriol Tintó

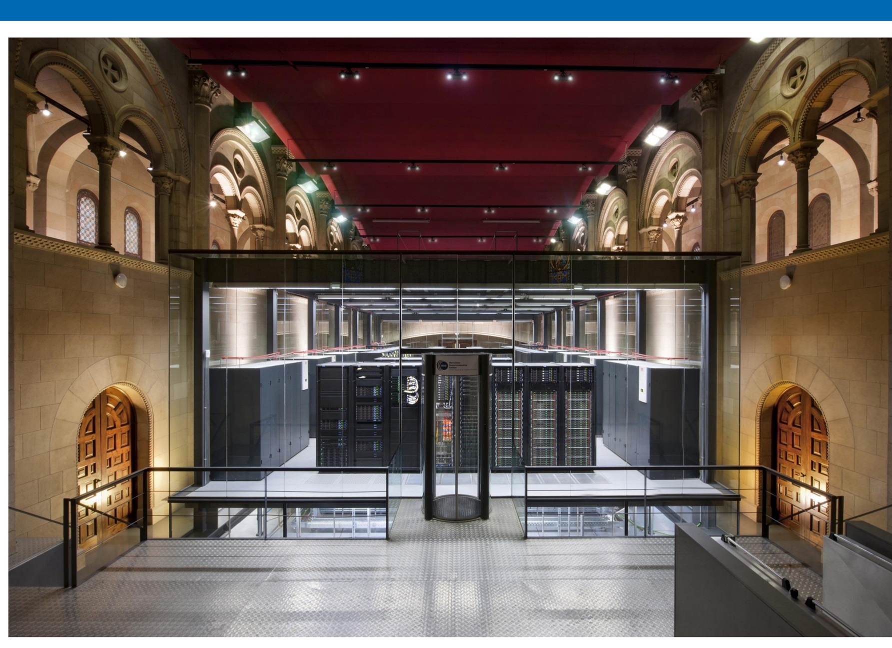

Marenostrum IV Supercomputer

References

[1] Yepes-Arbós, X. et al., 2016: "Scalability and performance analysis of EC-Earth 3.2.0 using a new metric approach (Part II)". Barcelona Supercomputing Center

[2] Acosta, M.C. et al., 2016: "Performance analysis of EC-Earth 3.2: Coupling". Barcelona Supercomputing Center

[3] Tintó Prims, O., et al. 2015: "Optimization of an ocean model using performance tools". Barcelona Supercomputing Center

[4] Acosta, M.C. at al., 2016: "Reproducibility of Earth System Models". Barcelona Supercomputing Center

[5] Barcelona Supercomputing Center. BSC performance tools, 2016: https://tools.bsc.es/

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

The research leading to these results has received funding from the EU H2020 Framework Programme under grant agreements no 675191 (ESiWACE) and no 641727 (PRIMAVERA).

This poster can be downloaded at
https://earth.bsc.es/wiki/doku.php?id=library:external:posters

Corresponding authors:
mario.acosta@bsc.es

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

PRIMAVERA