Anticipating the computational performance of Earth System Models for pre-exascale systems

Xavier Yepes-Arbós[1], Miguel Castrillo[1], Mario C. Acosta[1], Kim Serradell[1]

[1]Barcelona Supercomputing Center (BSC), Spain

Session AS4.2 - EGU21-7888

1. Introduction

Overview

The forecast capability of Earth System Models (ESMs) is strongly linked to the amount of computing power, given that the spatial resolution used for global climate experimentation is a limiting factor to correctly reproduce climate mean state and variability. However, this higher spatial resolution requires to improve the computational efficiency of ESMs to take advantage of new High Performance Computing (HPC) platforms. In this context, the porting of a new ultra-high resolution configuration into a new and more powerful HPC cluster involves technical expertise to efficiently deploy and get a good computational performance of such a novel configuration.

EC-Earth 4: OpenIFS 43R3 and NEMO 4

In order to take advantage of this new upcoming exascale era, we are developing the new EC-Earth 4 climate model by coupling the two main components:

- OpenIFS 43R3 as the atmospheric component.
- NEMO 4 as the ocean component.

These two components include different computational optimizations to improve the efficiency, e.g. extending the asynchronous I/O capabilities of the XIOS server to OpenIFS.

2. Benchmarking and setup

Benchmarking

We consider it is important to anticipate the computational behaviour of EC-Earth 4 for new pre-exascale machines, such as the upcoming MareNostrum 5 of the BSC. OpenIFS and NEMO are therefore benchmarked on the MareNostrum 4 (MN4) petascale machine to detect potential computational issues from:

- Developments of new features.
- Investigate if previous known limitations are solved.

We believe this type of analysis can be used to efficiently set up new ultra-high resolutions from a computational point of view, not only for EC-Earth, but also for other ESMs.

In particular, our benchmarking consists of large strong scaling tests by running different output configurations, such as changing multiple XIOS parameters and number of 2D and 3D fields. These tests use up to 2,595 nodes, the 75% of the MN4 supercomputer.

Models' configuration

The models are evaluated on MN4 with the following parameters:

- OpenIFS uses a 9 km global horizontal resolution (Tco1279) and three different I/O configurations: no output, CMIP6-based fields and huge output volume (8.8 TB) to stress the I/O part.
- NEMO uses a 3 km global horizontal resolution (ORCA36) with and without the sea-ice module, and different I/O configurations: no output, only enabling 2D fields, or either producing 3D variables on an hourly basis.

3. OpenIFS results

OpenIFS scalability (Tco1279L137)

No output

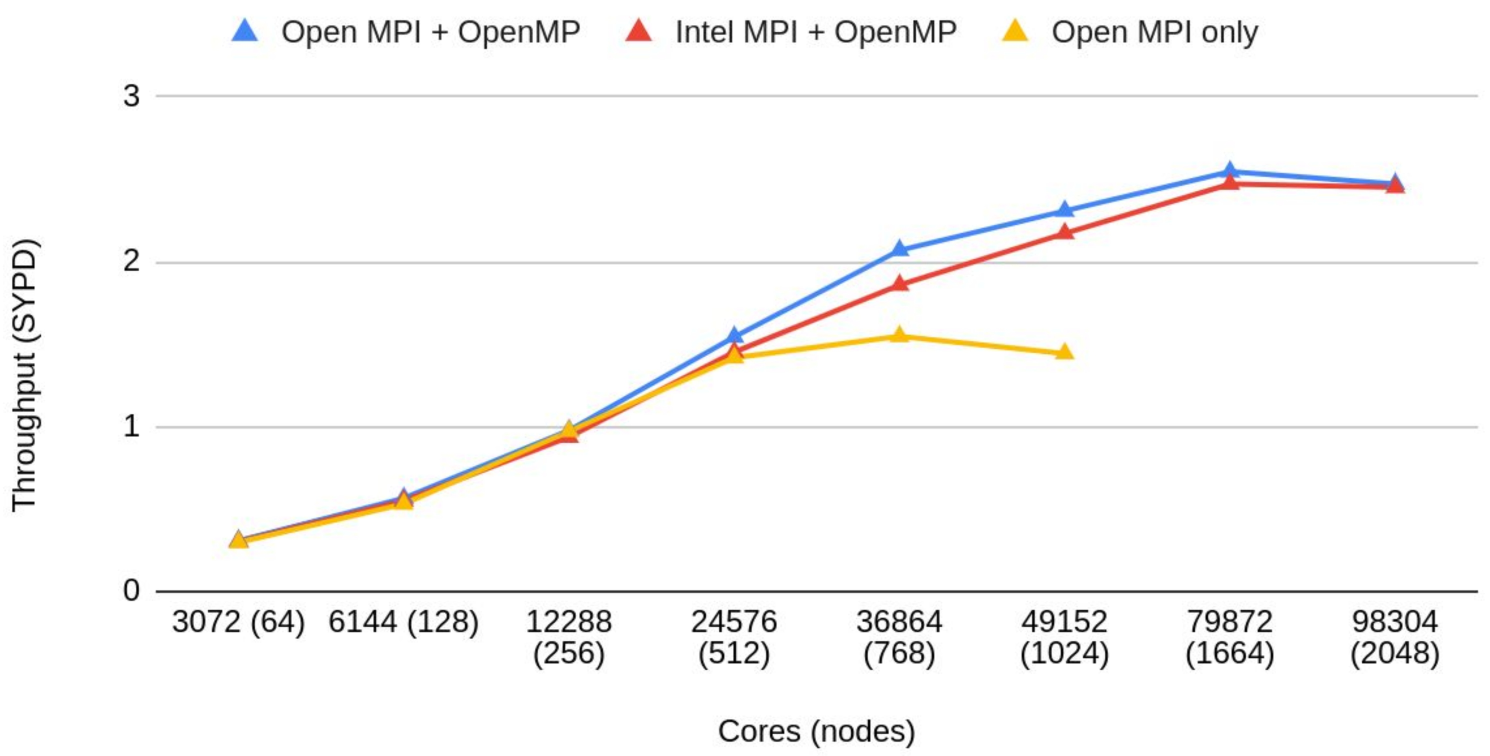

OpenIFS-XIOS scalability for HighResMIP (Tco1279L137)

multiple_file mode, 885 GB output

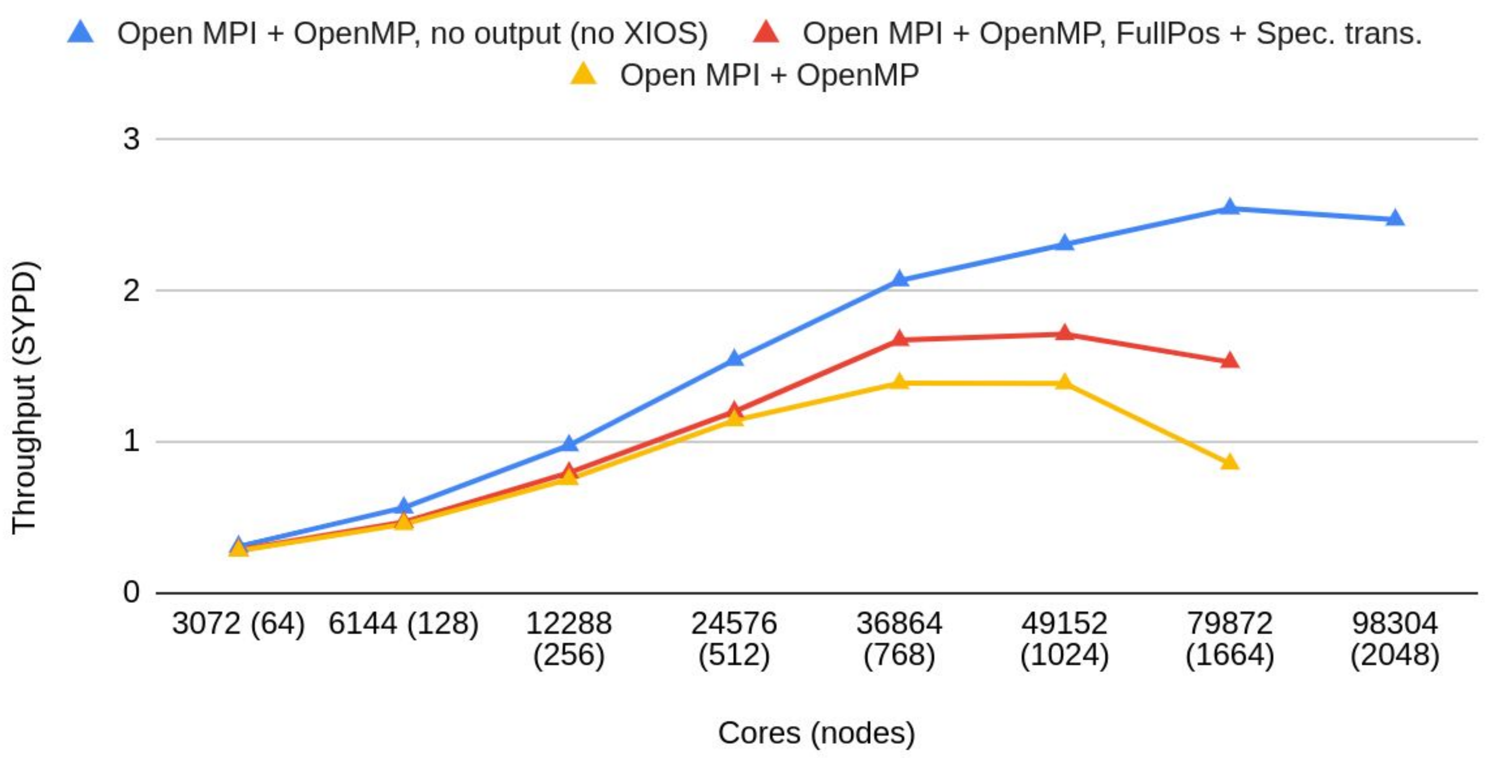

OpenIFS-XIOS scalability for theoretical (Tco1279L137)

multiple_file mode, 8.8 TB output

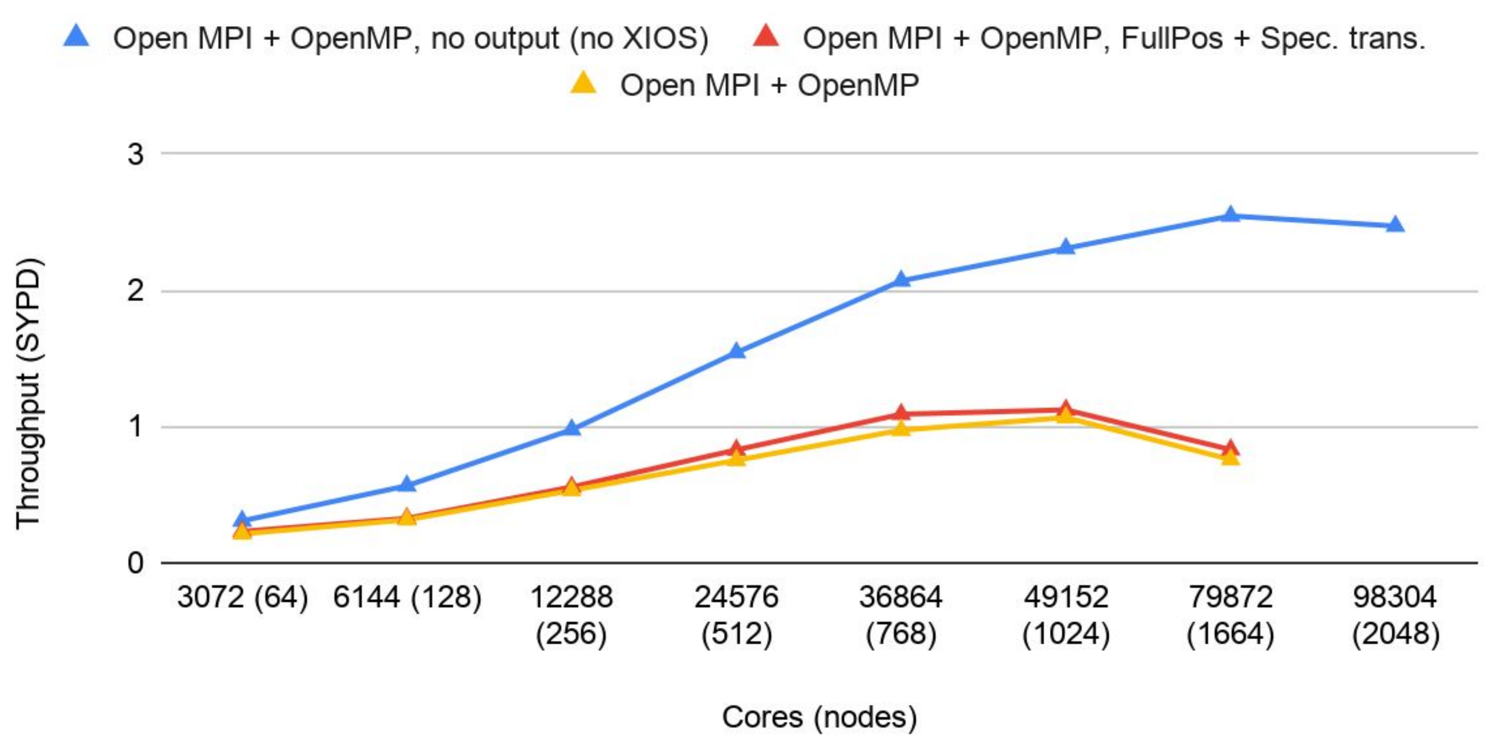

4. NEMO results

NEMO scalability (ORCA36)

No output

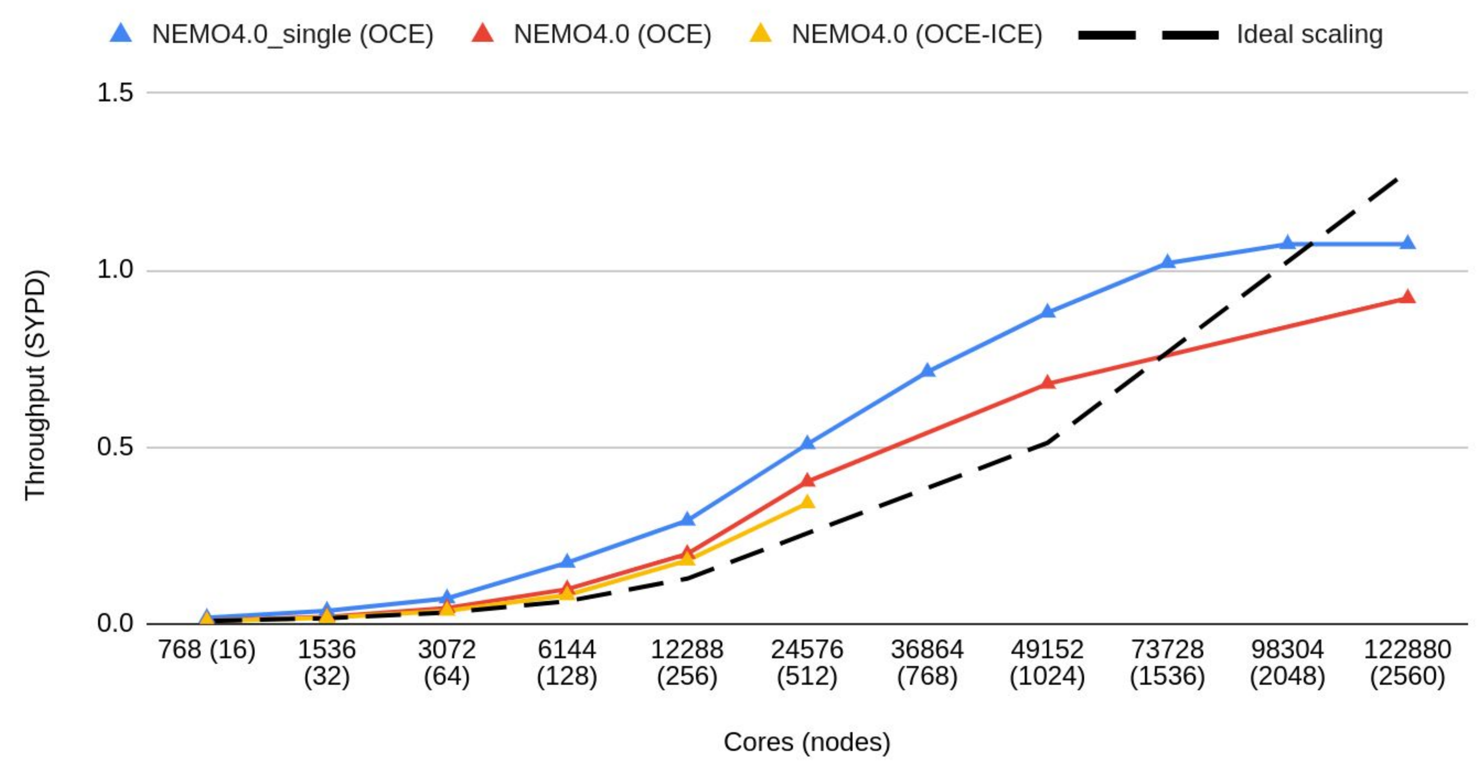

NEMO-XIOS time step scalability

512 nodes for NEMO, multiple_file mode, 5-hour 3D output (340 GB per hour)

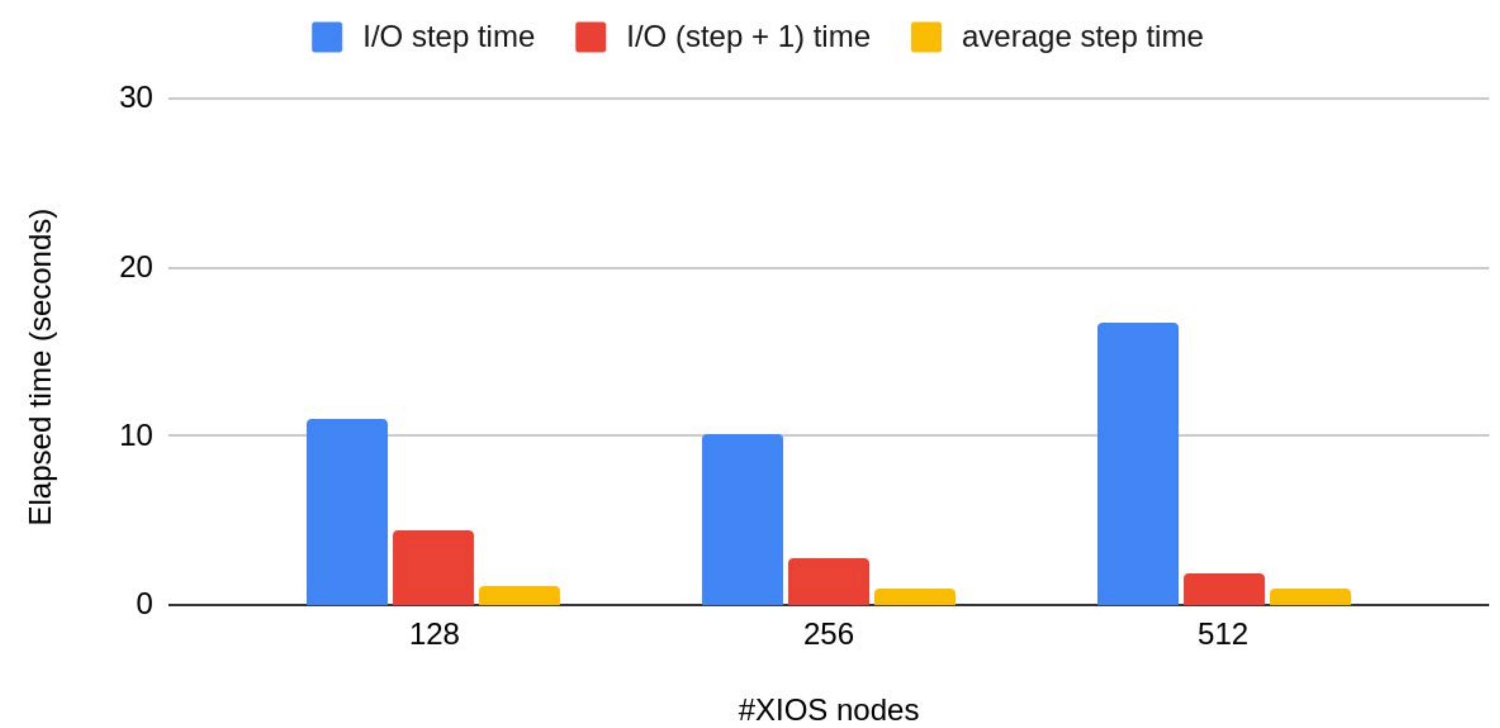

5. Conclusions

OpenIFS 43R3:

- It has a decent scalability until it reaches 1664 nodes.
- The MPI+OpenMP hybridisation is much better than MPI only.
- Open MPI is slightly better than Intel MPI.
- Production throughput depends on different factors such as I/O frequency and size, but it is strongly linked to the FullPos post-processing as well as the spectral transformations. It is necessary to further investigate the weight of each one.
- The XIOS scalability has to be improved, especially the memory management.

NEMO 4:

- It has a good scalability until it reaches 2,546 nodes.
- Using the best XIOS configuration the I/O overhead is moderate, but it uses a lot of dedicated resources (bad memory management).
- Production throughput depends on different factors: time step size, I/O frequency, I/O size, diagnostics computation, etc.
- XIOS does not scale.

The research leading to these results has received funding from the EU H2020 Framework Programme under grant agreement no. 675191 (ESiWACE).
This material reflects only the author's view and the Commission is not responsible for any use that may be made of the information it contains.
This poster can be downloaded at **https://earth.bsc.es/wiki/doku.php?id=library:external:posters**
Corresponding author: **xavier.yepes@bsc.es**

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

Barcelona Supercomputing Center
Centro Nacional de Supercomputación