
BSC
Barcelona
Supercomputing
Center
Centro Nacional de Supercomputación

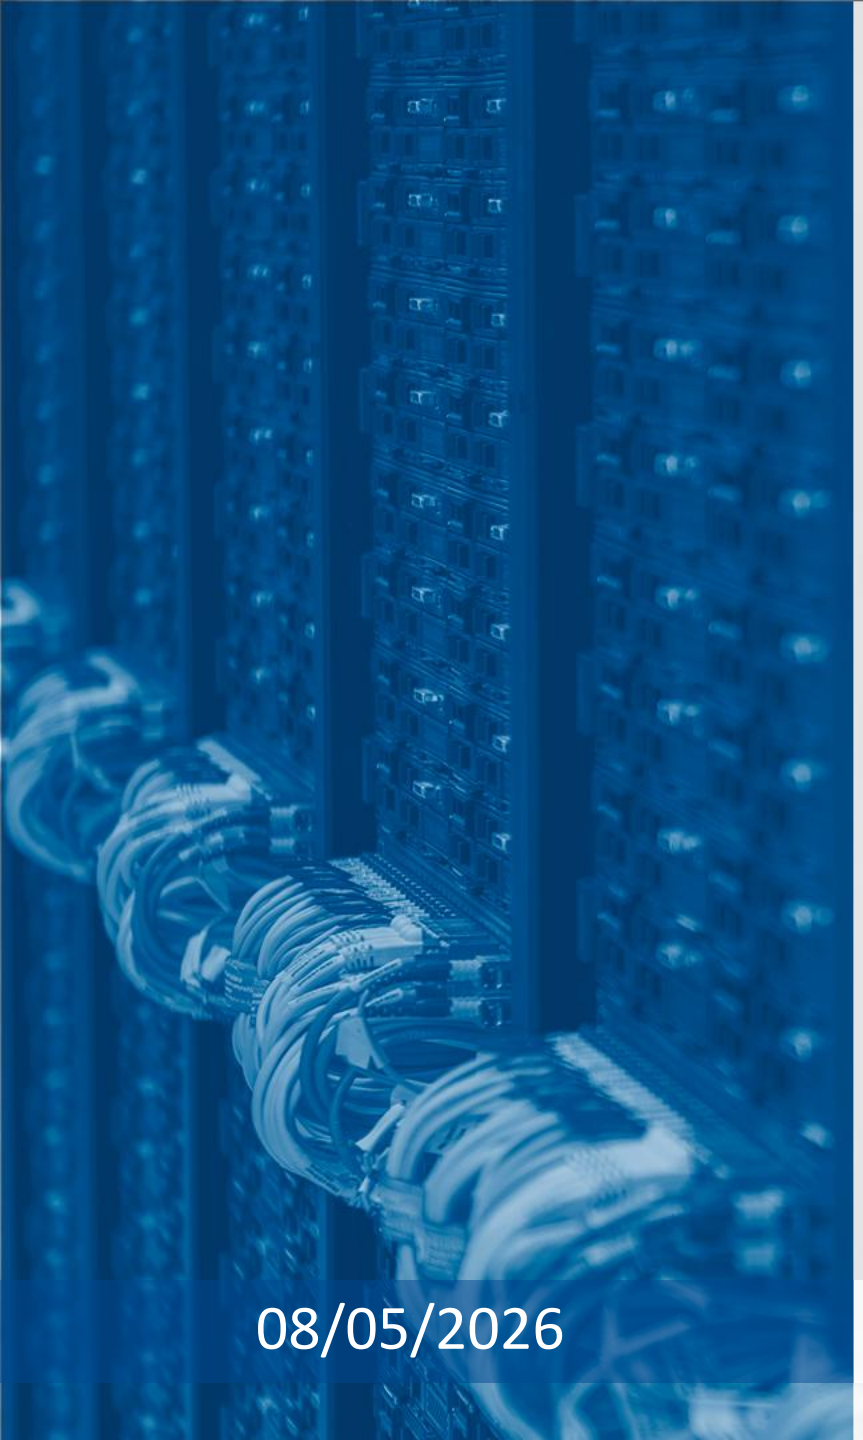

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Enhancing Earth system models efficiency: Leveraging the Automatic Performance Profiling framework

Roc Salvador Andreazini
Xavier Yepes Arbós, Oriol Tintó Prims, Stella Paronuzzi Ticco, and Mario Acosta Cobos

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

Destination Earth

08/05/2026

EGU General Assembly 2026, Vienna, Austria

Why?

The Challenge: Manual regression testing is important for every model update to verify performance.

Our Approach: Automatic Performance Profiling (APP)

- **Automates Repetition:** Eliminates manual testing cycles, instantly flagging performance regressions and pinpointing root causes.
- **Universal Accessibility:** Delivers actionable insights for all users:
  - General Users: High-level metrics (e.g. SYPD).
  - Performance Experts: Deep-dive diagnostics (e.g. PAPI, MPI).

Tool overview

AUTOMATIC PERFORMANCE PROFILING

AUTOSUBMIT

PROFILING TOOLS

BSC Performance Tools

APP PARSERS AND POSTPROCESSING

Tool overview

Modular design to add new models and platforms easily

SUPPORTED MODELS

+ Add new models here

SUPPORTED HPC PLATFORMS

- BSC MARENOSTRUM5
- ECMWF HPC2020

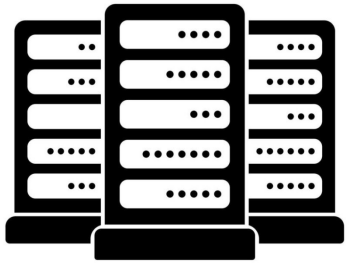

+ Add new platforms here

Tool overview

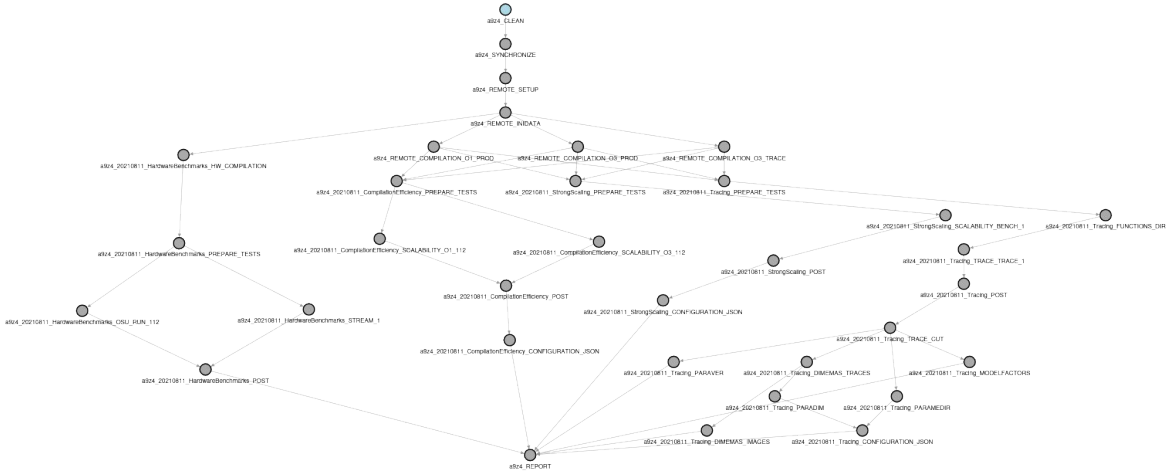

HARDWARE BENCHMARKS

STRONG SCALING

WEAK SCALING

COMPILATION EFFICIENCY

TRACING

Update the configuration

COMPILATION EFFICIENCY

Hardware Benchmarks

Runs multiple hardware benchmarks to characterize the performance of the platforms where the model is run.

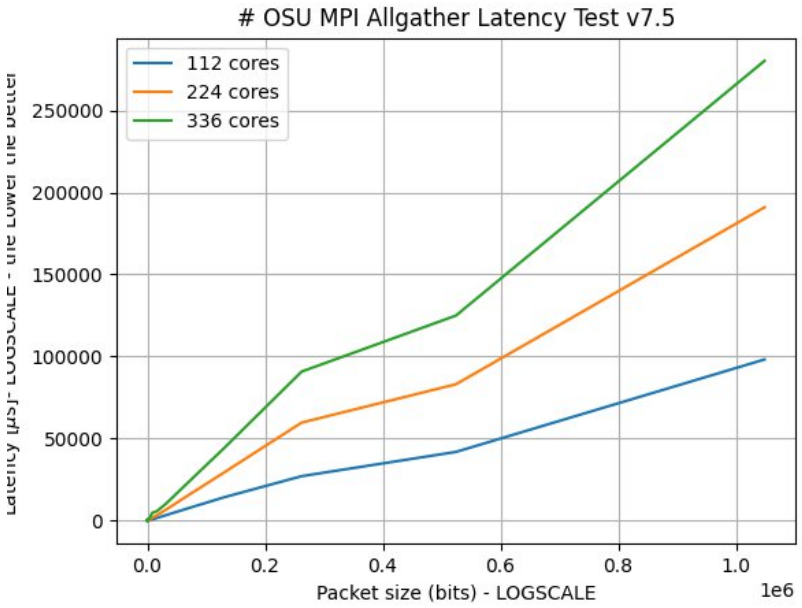

OSU MPI Allgather benchmark.

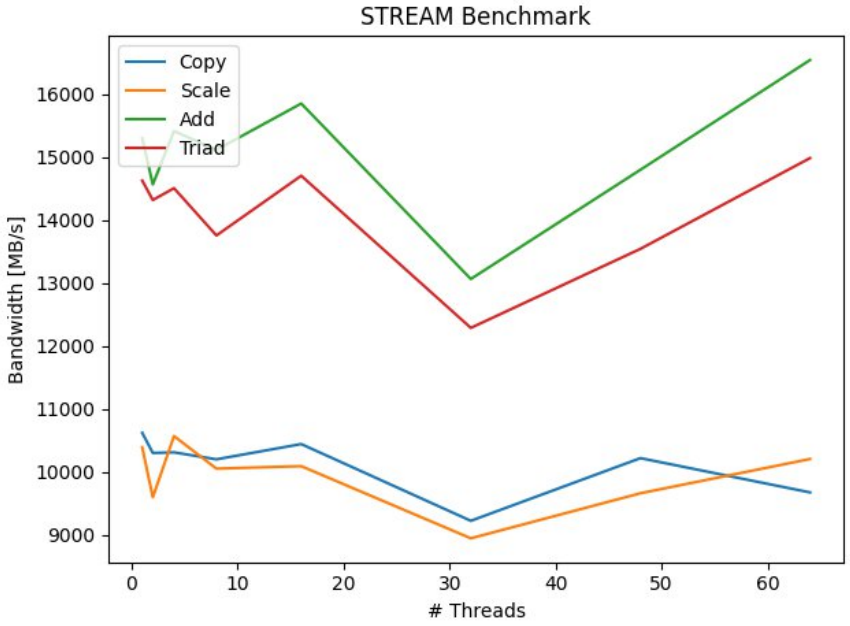

STREAM memory bandwidth benchmark.

Strong Scaling

Increase processor count while keeping the same problem size (resolution).

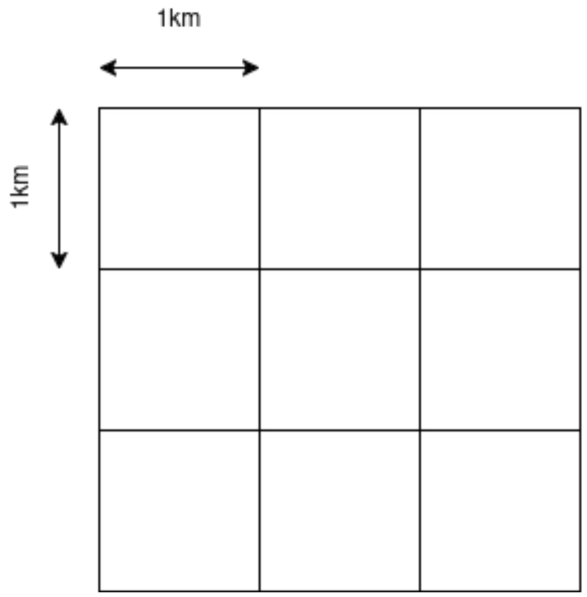

9 processors

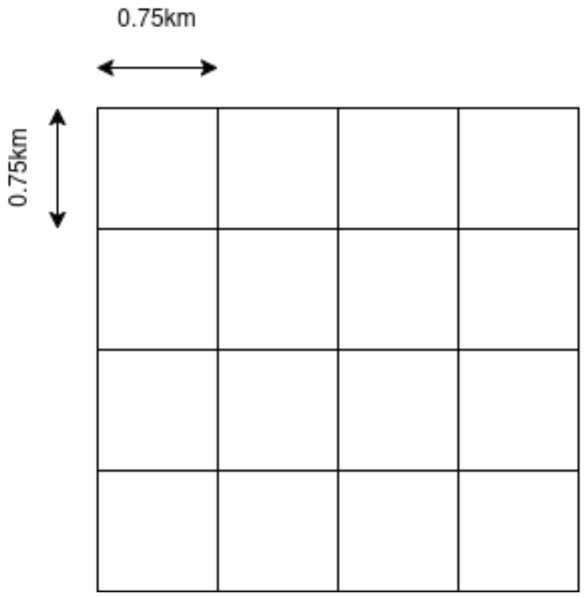

16 processors

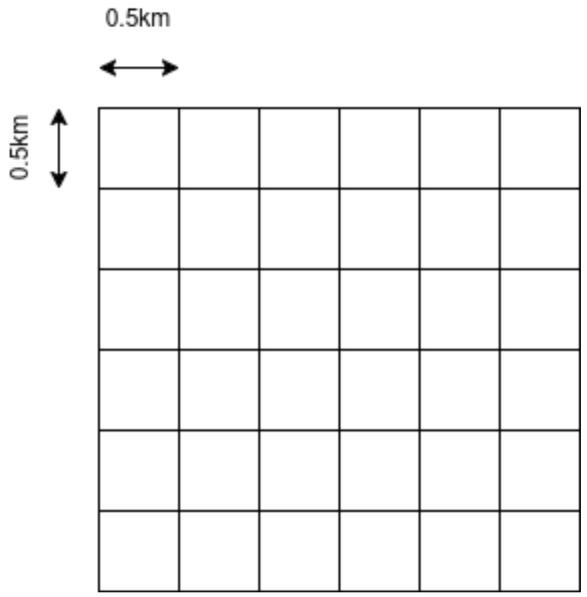

36 processors

Strong Scaling

Ideally we would expect the model speed to scale proportionally to the processors added.

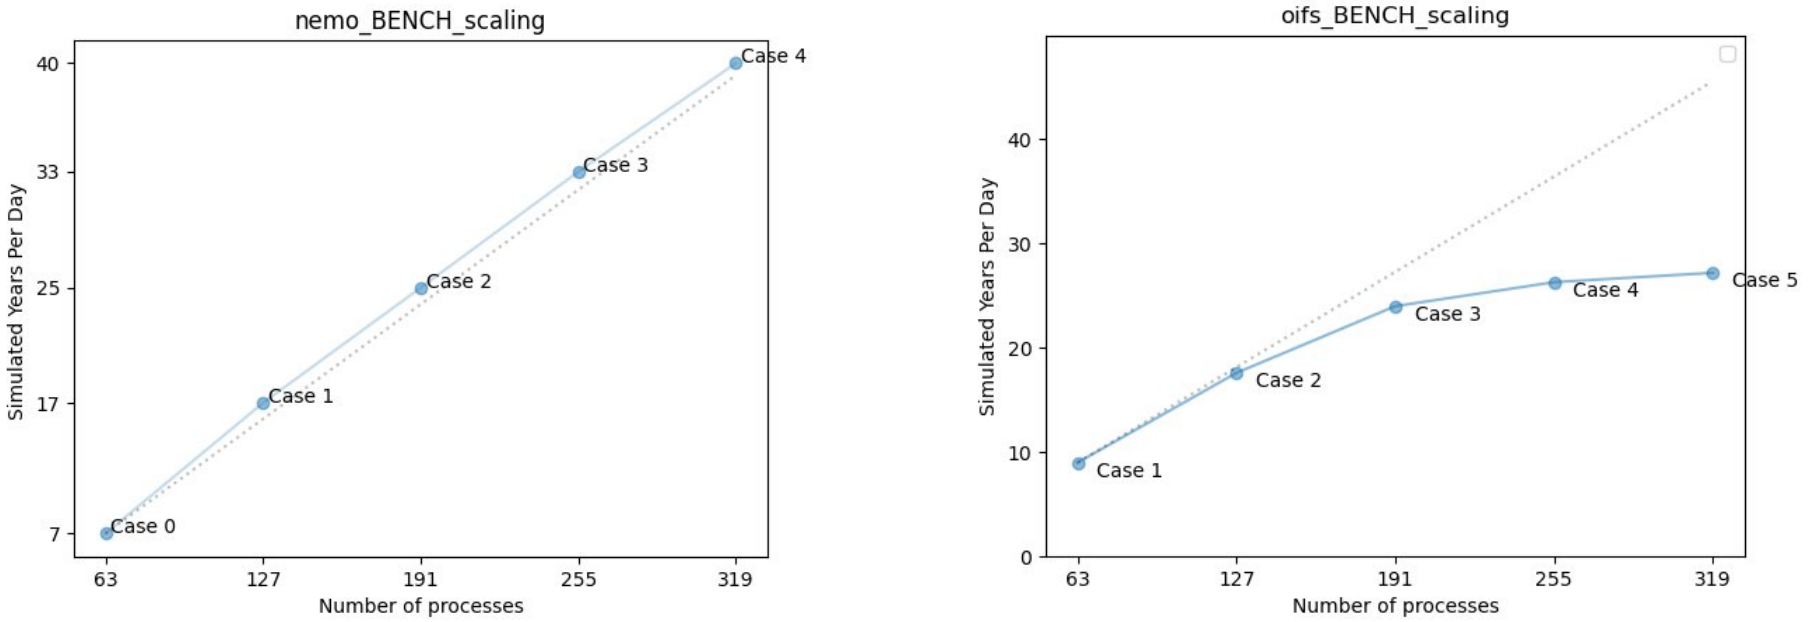

NEMO (left) and OIFS (right) scalability plots on HPC2020 using ECE4 coupled.

Weak Scaling

Increase processor count while increasing problem size proportionally.

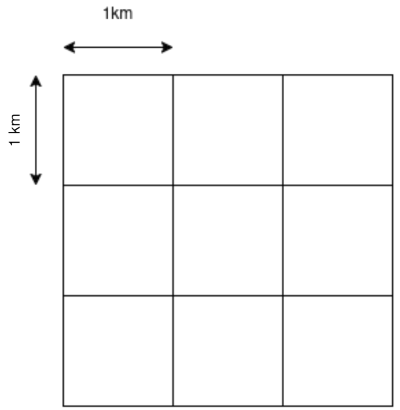

9 processors

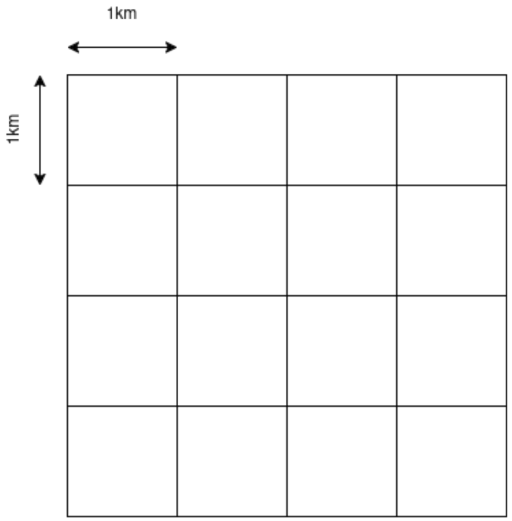

16 processors

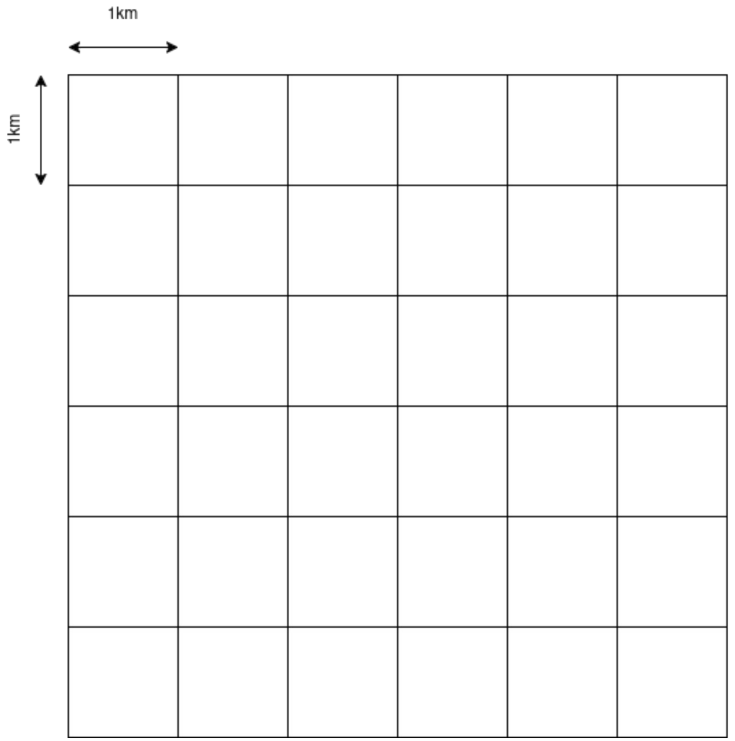

36 processors

Weak Scaling

Ideally the model speed should remain the same.

Config:	Tile Size	Ncores	SYPD
orca1	1000	24	4.5
orca025	16000	384	4.3
orca12	144000	3456	3.6

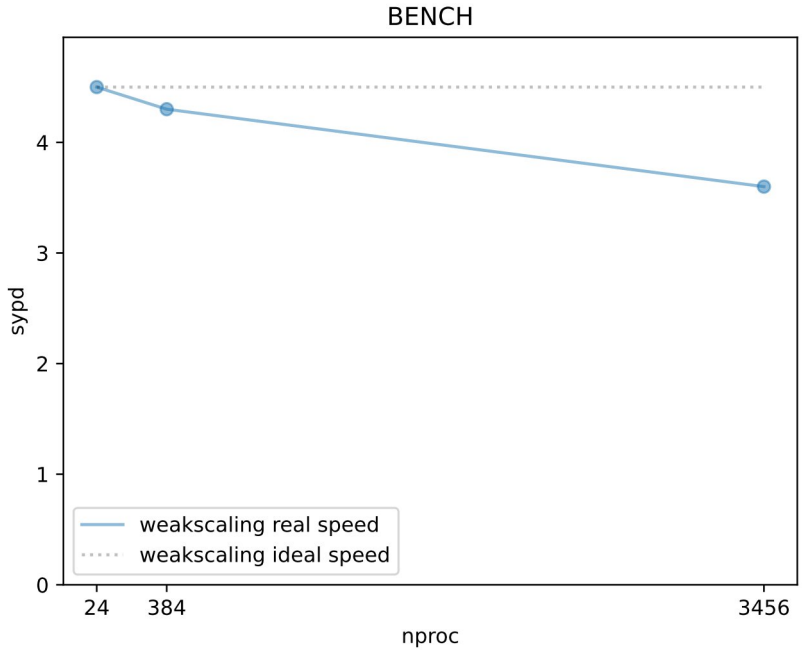

Weak scaling results for NEMO 4.2.1.

Compilation Efficiency

Compares different compiled versions of the model using a strong scaling test. This allows to compare compiler flags/compilers/libraries.

Nproc	O1-prod (SYPD)	O3-prod (SYPD)
112	1.51	1.62
224	2.82	2.89
336	3.83	4.66
448	4.87	6.18
560	6.33	6.38

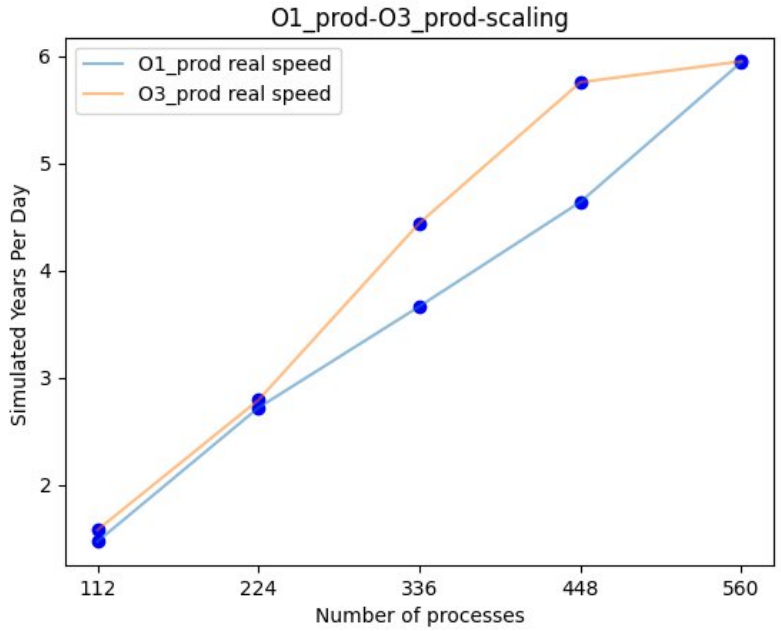

HPCW NEMO 5 scaling for different compiled versions.

Tracing

Short simulation strong scaling test that is traced using Extrae to extract advanced performance metrics.

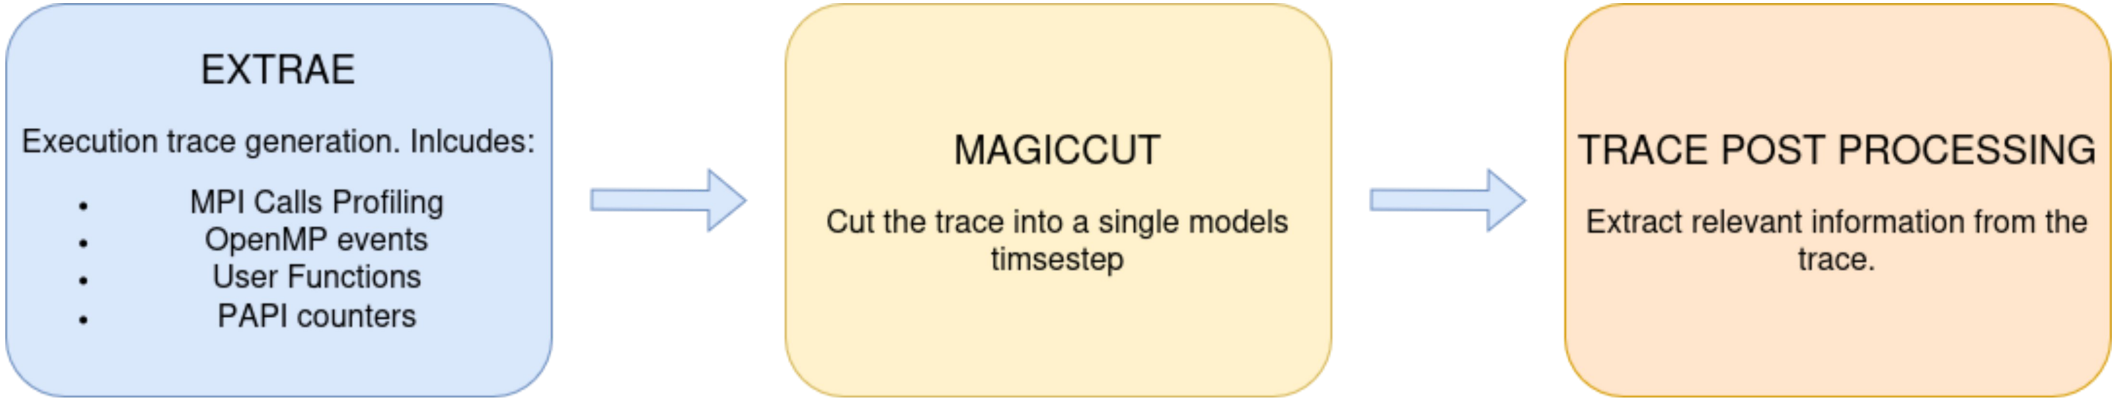

Tracing

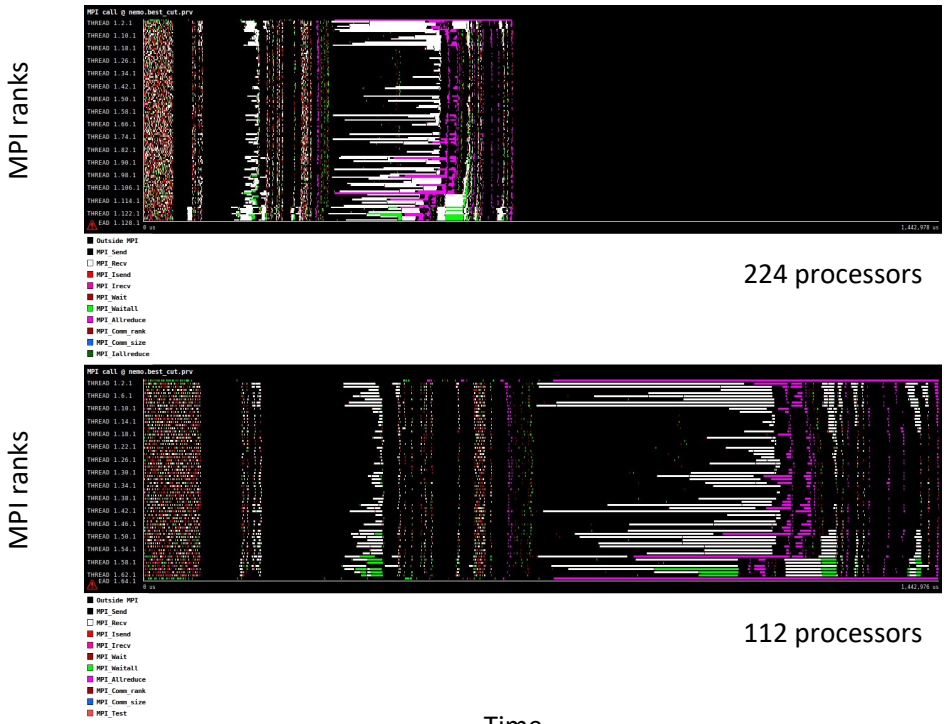

MPI Call Paraver view for two different cases (112 vs 224 procs) using the same timescale.

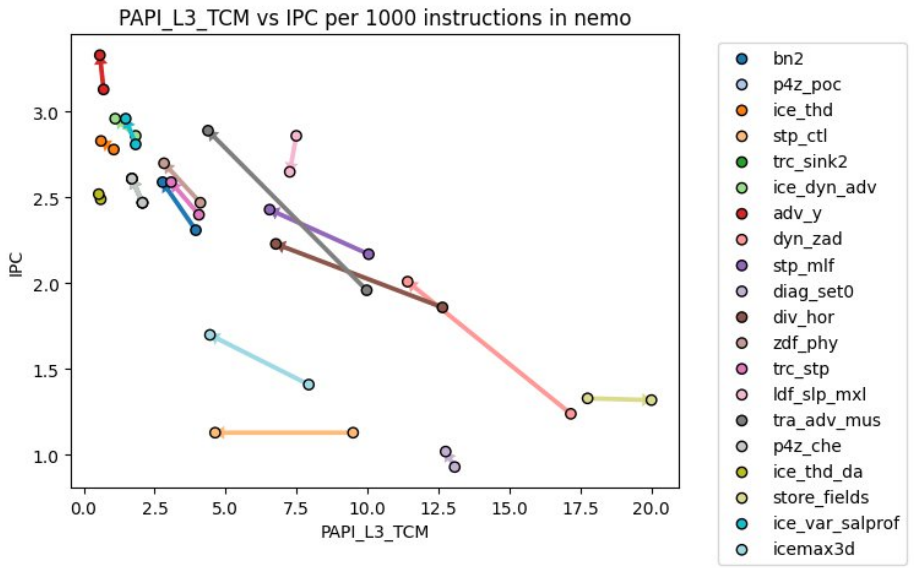

IPC vs L3 cache misses evolution for NEMO 4.2.1 functions when increasing resources

Tracing

-	OutsideMPI	MPIRecv	MPIIsend	MPIIrecv	MPIWait	MPIWaitall
average	92.54	6.52	0.60	0.10	0.25	0.08
maximum	96.52	24.56	0.98	0.12	7.34	0.30
minimum	73.59	2.70	0.41	0.08	0.05	0.01
stdev	5.27	4.51	0.09	0.02	1.00	0.07

MPI Call profile for 112 processors for NEMO.

nemo_ice_dyn_adv

resources	112.0	224.0
speedup	1.0	1.58
efficiency	1.0	0.79

Table 13: nemo_ice_dyn_adv

nemo_ice_thd

resources	112.0	224.0
speedup	1.0	1.86
efficiency	1.0	0.93

Table 14: nemo_ice_thd

nemo_ice_thd_da

resources	112.0	224.0
speedup	1.0	1.8
efficiency	1.0	0.9

Table 15: nemo_ice_thd_da

NEMO User Functions scalability.

The final report

The final results can be inspected in a LaTeX report or an interactive dashboard (still experimental).

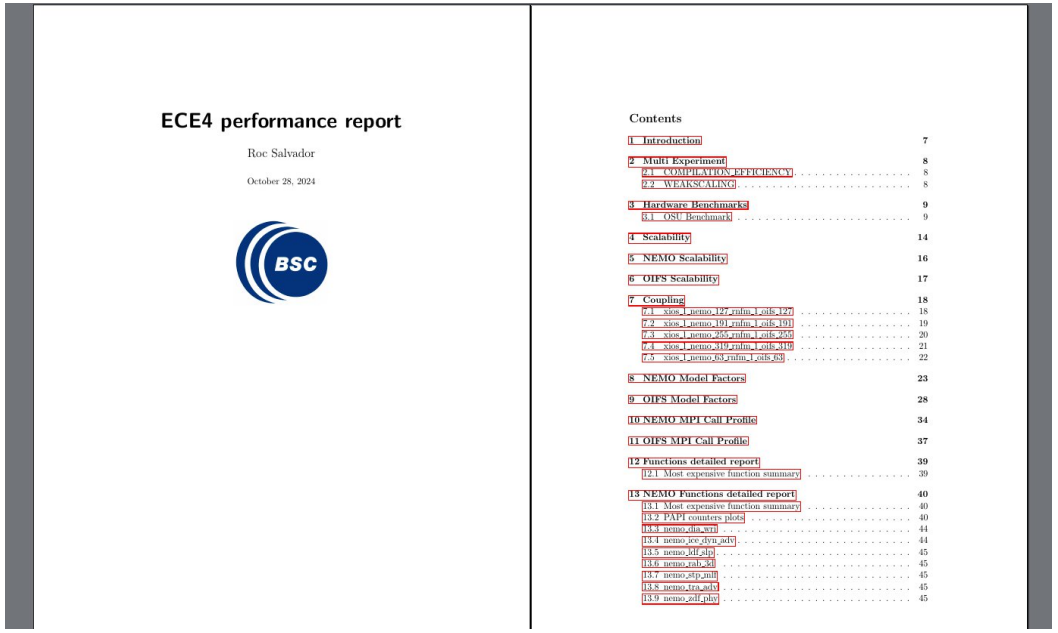

ECE4 performance report

Roc Salvador

October 28, 2024

Contents

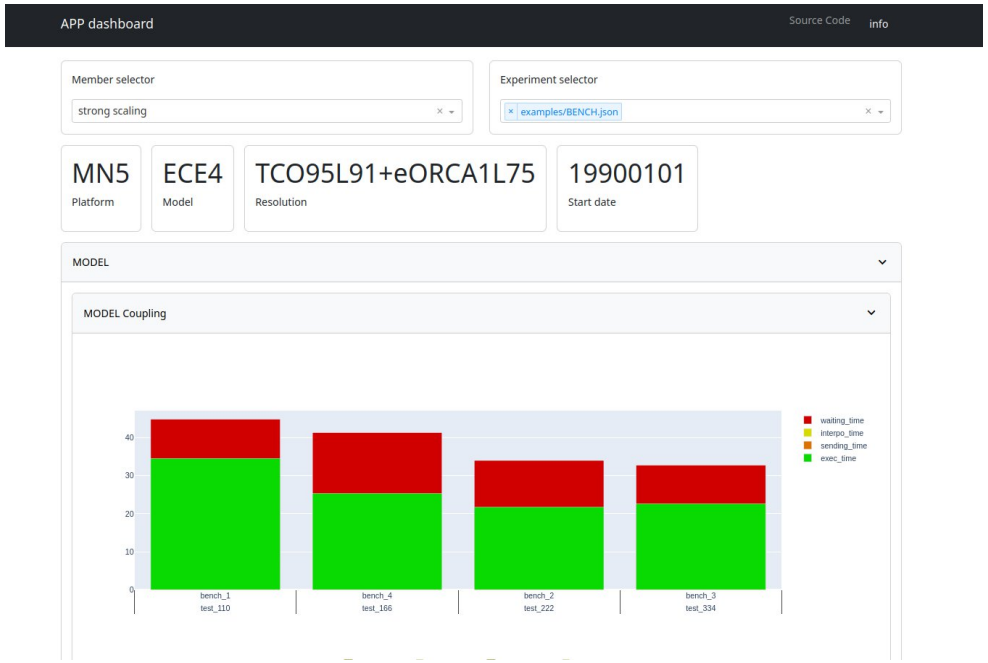

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

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

Destination Earth

Summary

The APP tool automates performance profiling, generating detailed and accessible reports without manual intervention, significantly reducing the workload for users.

The tool currently supports several models, including EC-Earth4 and NEMO, and is compatible with HPC systems such as Marenostrum 5 and ECMWF's supercomputers.

We are currently working on the support for the Auto-Monarch model. We plan to add support of HPC platforms widely used by the community, and improve the tool robustness and long-term sustainability.

Links

Read The Docs documentation: https://automatic-performance-profiling.readthedocs.io/en/latest/

Tool repository: https://gitlab.earth.bsc.es/ces/hpc-for-es-team/automatic_performance_profiling

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Thank you

This work was supported by Destination Earth, a European Union initiative aimed at creating a digital replica of the Earth system by 2030. We acknowledge the joint implementation efforts of ECMWF (Digital Twins & Engine), ESA (Core Service Platform), and EUMETSAT (Data Lake).

Funded by the European Union. This work has received funding from the European High Performance Computing Joint Undertaking (JU) under grant agreement No 101093054.

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

Destination Earth

roc.salvador@bsc.es