Barcelona Supercomputing Center
Centro Nacional de Supercomputación

is-enes
INFRASTRUCTURE FOR THE EUROPEAN NETWORK FOR EARTH SYSTEM MODELLING

CPMIP: Computational evaluation of Earth System Models. Multi-model results from CMIP6 and challenges for the exascale computing.

Mario Acosta, Venkatramani Balaji, Sergi Palomas and Stella Paronuzzi

25/05/2022

EGU General Assembly, Vienna, 23-27 May

Introduction

- Traditional metrics of computational efficiency such as performance counters and scaling curves do not provide us enough about real sustained performance from climate models on different machines.

- Most applications targeting exascale machines require some degree of rewriting to expose more parallelism. **Understand the performance of our models is critical.**

- CPMIP (Balaji et al. 2017): a set of metrics that can be used for the study of computational performance of Earth System Models (ESMs)

The IS-ENES3 project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 824084

Introduction

- CMIP6: Coordinated experiments designed to understand specific aspects of the model response.
  - The ideal context to create a performance data-base from a multi-model context using different configurations, resolutions, platforms...
- **IS-ENES3 provided a unique opportunity to exploit this new set of metrics**, performing for the first time a complete computational/energy analysis and the creation of a novel data-base based on CMIP6 experiments, using the different models and platforms available all across Europe.
- The outcome of this work was published in the D4.3 deliverable including:
  - The possibilities for collaboration with other groups (ES-DOC, HPC-TF and Carbon Footprint G.)
  - The analysis illustrating some practical examples, and proving the usefulness of the metrics to the community.
  - Main difficulties encountered in the coordination of the collection, including general recommendations on how to solve these problems for future collections and analyzes.

The IS-ENES3 project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 824084

CPMIP metrics (1, 2)

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
Complexity	the number of prognostic variables per component
Actual SYPD	how affect queue time and interruptions to the complete experiment
Parallelization	total number of cores allocated for the run
Energy Cost Per Year (JCPY)	how much energy is needed per each year of simulation
Memory Bloat	the ratio between actual and ideal memory size
Data Output Cost	how much time and resources are used performing I/O
Data Intensity	the amount of data produced per compute-hour
Coupling Cost	how much time and resources are used in the cost of the coupling algorithm as well as load imbalance

(1) Balaji et al. 2017, CPMIP: measurements of real computational performance of Earth system models in CMIP6.
(2) Mario Acosta et al. 2021, ISENES3 D4.3: CPMIP performance metrics evaluation for CMIP6 and community advice. https://doi.org/10.5281/zenodo.6394049

The IS-ENES3 project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 824084

CPMIP metrics

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
Complexity	the number of prognostic variables per component
Actual SYPD	how affect queue time and interruptions to the complete experiment
Parallelization	total number of cores allocated for the run
Energy Cost Per Year (JCPY)	how much energy is needed per each year of simulation
Memory Bloat	the ratio between actual and ideal memory size
Data Output Cost	how much time and resources are used performing I/O
Data Intensity	the amount of data produced per compute-hour
Coupling Cost	how much time and resources are used in the cost of the coupling algorithm as well as load imbalance

The IS-ENES3 project has received funding from the European Union's Horizon 2020
research and innovation programme under grant agreement No 824084

CPMIP: Community List

Model / Institution	People Involved
CNRM-CM6	Sophie Valcke, Marie Pierre Moine
IPSL-CM	Arnaud Caubel
EC-Earth	Mario Acosta, Uwe Fladrich, Philippe Le Sager
MetO	Harry Shepherd, JC Rioual
CMCC	Italo Epicoco, Silvia Mocavero
MPI-M-DKRZ	Maria Moreno, Reinhard Budich, Joachim B.
U. Read	Grenville Lister, Bryan Lawrence
Nor-ESM	Alok Kumar Gupta
TOPAZ/MOM5	Paulo Nobre
GFDL	Niki Zadeh

Include 11 models with 32 CMIP6 configurations (AMIP, OCE, Coupled, different resolutions…)

The IS-ENES3 project has received funding from the European Union's Horizon 2020
research and innovation programme under grant agreement No 824084

Main outcome

- The outcome of this work was published in the D4.3 deliverable including:
  - **The possibilities for collaboration with other groups (ES-DOC, HPC-TF and Carbon Footprint G.)**
  - The analysis illustrating some practical examples, and proving the usefulness of the metrics to the community.
  - Main difficulties encountered in the coordination of the collection, including general recommendations on how to solve these problems for future collections and analyzes.

The IS-ENES3 project has received funding from the European Union's Horizon 2020
research and innovation programme under grant agreement No 824084

Carbon footprint collaboration

CMIP6 Experiments: Institutions/Models	Useful SY	Useful CH (Mh)	Total Energy Cost (Joules)	PUE	CF (g CO2/kWh)	Total Carbon Footprint (CO2)
EC-Earth	28,105	31.3	1.24E+12	1.35	357	165t
CNRM-CERFACS	47,000	160	6.18E+12	1.43	40	97t
IPSL	75,000	150	8.72E+12	1.43	50	172t
CMCC	965	1.99	1.61E+12	1.84	408	329t
UKMO	37.237	683	2.67E+13	1.35	87	868t
DKRZ	1,276	5.52	4.09E+11	1.19	184	24t
NCC-NORESM2	23,096	27.23	1.69E+12			
NERC	640	55.497	2.17E+12	1.10	0*	
MPI	24,175	16.32	7.10E+11	1.19	184	42t

*Green tariff according to NERC

Carbon Footprint = Total Energy Cost (MWh) * Conversion Factor (CF) * Power Usage Effectiveness (PUE)

The IS-ENES3 project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 824084

Main outcome

- The outcome of this work was published in the D4.3 deliverable including:
    - The possibilities for collaboration with other groups (ES-DOC, HPC-TF and Carbon Footprint G.)
    - **The analysis illustrating some practical examples, and proving the usefulness of the metrics to the community.**
    - Main difficulties encountered in the coordination of the collection, including general recommendations on how to solve these problems for future collections and analyzes.

The IS-ENES3 project has received funding from the European Union's Horizon 2020
research and innovation programme under grant agreement No 824084

Analysis and results

- The analysis illustrated some practical examples, and proving the usefulness of the metrics to the community.
  - Resolution impact
  - Complexity impact
  - Data output impact
  - ASYPD: Queue time and interruptions impact
  - Coupling impact

- In a previous work, we also studied a specific model (EC-Earth) to evaluate the computational efficiency on different machines or configurations
  - Complexity Impact: Identify which component is the bottleneck of the coupled version
  - ASYPD: Queue time could differ between machines, due to the different set-up of the queue systems
  - Comparison through machines: Detect bottlenecks according to the limitations of each hardware

The IS-ENES3 project has received funding from the European Union's Horizon 2020
research and innovation programme under grant agreement No 824084

Analysis and results

- The analysis illustrated some practical examples, and proving the usefulness of the metrics to the community.
  - **Resolution impact**
  - Complexity impact
  - Data output impact
  - **ASYPD: Queue time and interruptions impact**
  - Coupling impact

- In a previous work, we also studied a specific model (EC-Earth) to evaluate the computational efficiency on different machines or configurations
  - Complexity Impact: Identify which component is the bottleneck of the coupled version
  - ASYPD: Queue time could differ between machines, due to the different set-up of the queue systems
  - Comparison through machines: Detect bottlenecks according to the limitations of each hardware

The IS-ENES3 project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 824084

CPMIP Results

- Grouping experiments with similarities

Resol ↑ → **Paralel** ↑ **CHSY** ↑ **Mem. B.** ↑

- Mem. B. results support the idea that our models are memory bound

---

- If we group these experiments per complexity

Resol ↑ → **SYPD** ↓ **CHSY** ↑ **Paralel** ↑

- No model using the highest resolution is scaling ideally
  - These configurations are very demanding
  - This could be due to hardware restrictions, the high cost of other phases (CO, DO), the overhead introduced for a higher number of parallel resources and the memory consumption… or a combination of some of them

The IS-ENES3 project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 824084

CPMIP Results

- **If we consider ASYPD_OH as the percentage decreased compared to SYPD**
- ASYPD_OH is around 10-50% for the most of the cases.
- ASYPD_OH can be classified in two clear groups
  - Queue time represents around 10-15% and adding interruptions and workflow management around 40-50%

- BSC results using the same configuration prove that the percentage could change between two different platforms. A finer granularity could be needed.

- **The typical SYPD that we see reported for the community could be half!**

The IS-ENES3 project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 824084

Main outcome

- The outcome of this work was published in the D4.3 deliverable including:
  - The possibilities for collaboration with other groups (ES-DOC, HPC-TF and Carbon Footprint G.)
  - The analysis illustrating some practical examples, and proving the usefulness of the metrics to the community.
  - **Main difficulties encountered in the coordination of the collection, including general recommendations on how to solve these problems for future collections and analyzes.**

The IS-ENES3 project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 824084

Conclusion 3

- Main difficulties encountered in the coordination of the collection, including general recommendations on how to solve these problems for future collections and analyzes.
  - Performance metrics collection is secondary. Facilitate the collection should help to the institutions.
  - A coordinated collection has been proved useful to ensure to get the metrics and solve possible inconsistencies or gaps during the process.
  - Create a finer granularity for some of the metrics could be important to improve the analysis in the future.

The IS-ENES3 project has received funding from the European Union’s Horizon 2020
research and innovation programme under grant agreement No 824084

The Consortium

Coordinated by **CNRS-IPSL**, the IS-ENES3 project gathers **22 partners** in **11 countries**

with

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement N°824084

Our website
https://is.enes.org/

Follow us on Twitter !
@ISENES_RI

Contact us at
is-enes@ipsl.fr

Follow our channel
IS-ENES3 H2020

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Thank you

mario.acosta@bsc.es