IS-ENES3 Deliverable D4.3

CPMIP performance metrics and community advice

Reporting period: 01/07/2020 – 31/12/2021

Authors: Mario Acosta (BSC), Sergi Palomas (BSC) and Stella Paronuzzi (BSC)
Reviewer(s): Italo Epicoco (CMCC) and Uwe Fladrich (SMHI)

Release date: 22/12/2021

ABSTRACT

This deliverable shows the main results obtained from the CPMIP collection done for CMIP6, including ISENES3 and some external partners. The document provides the list of partners involved and the CPMIP metrics per institution/model, as well as the approach used for the collection and the coordination behind this process. Additionally, a section has been included to explain the collaborations done with other groups (HPC-TF within the ENES network and the carbon footprint group) and the results produced, proving that the possibilities using CPMIP metrics go beyond a simple performance evaluation.

Furthermore, a section has been included to analyze the results. The analysis has been done illustrating some practical examples, and proving the usefulness of the metrics to the community. However, a more detailed analysis will be done in a scientific paper pending for publication. Moreover, it is described the main difficulties encountered in the coordination of the collection, including general recommendations on how to solve these problems for future collections and analyzes.

Revision table			
Version	**Date**	**Name**	**Comments**
V1	19/11/2021	Mario Acosta	Initial version for revision
V2	02/12/2021	Mario Acosta	Final version after revision
V3	17/12/2021	Mario Acosta	Fixed some errors

Dissemination Level		
PU	Public	X
CO	Confidential, only for the partners of the IS-ENES3 project	

Table of contents

Executive summary

WP4/NA3 performed for the first time a complete computational and energy performance analysis of CMIP6 experiments using the CPMIP computational metrics.

This document, while describing the data collection done, will highlight the coordination effort made by the various groups to ensure dissemination of results and to take advantage of the defined metrics for performing new analyses. Here, the complete CPMIP results are presented, grouped per model/institution involved. Additionally, a section has been included to analyse the results. The analysis has been done pursuing the previously commented goals, illustrating some practical examples, and proving the usefulness of the metric to the community. The next examples has been included to evaluate the impact of different issues:

- Resolution impact
- Complexity impact
- Data output impact
- Queue time and interruptions impact
- Coupling impact

From these examples we can conclude the first a main conclusion of this deliverable:

- CPMIP collection is not only a dissemination process. It is a very powerful tool to analyze the computational efficiency of a model across platforms, configurations; find bottlenecks…or even for the multi-model comparison.

An additional section has been included, specifically for this deliverable, where the main difficulties encountered in the coordination of the collection are highlighted. This section also includes general recommendations on how to solve these problems for future collections, including the main conclusions:

- Although the CPMIP collection is important, it is secondary during the CMIP execution. 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.
- Collect CPMIP metrics before or during the CMIP experiment to avoid unexpected issues for a post-collection.
- Develop new approaches to analyze the results and produce more interesting and valuable conclusions in the future.
- Inconsistencies in the way that some metrics are collected and how to avoid this for future collections
- Create a finer granularity for some of the metrics. Filling the gaps of some metrics or simply to improve the analysis in the future.

1 Introduction and objectives

When talking about climate models, traditional metrics, such as performance counters or scaling curves, may fall short in describing their real computational performance on different platforms. In the same way, traditional benchmarks (think about the LINPACK benchmark or a scientific domain test) are not fully representative of the sustained performance of a complete Earth System Model (ESM) running an operational configuration as CMIP6.

For these reasons, the community felt the need to create a new paradigm to measure ESM computational performance: the CPMIP (a computational performance model intercomparison project) set of metrics saw the light in the framework of the IS-ENES2 project.

The following points summarise the philosophy at the base of this new set of metrics:

- the set of metrics has to be universally available from current ESMs
- the chosen measure has to be easy to collect and should not require specialized instrumentation or software
- when evaluating the ESMs performance, the configuration used should be the operational one
- the performance measurements should be taken across the entire lifecycle of modeling and cover both data and computational load

The CPMIP metrics provide a new way to study ESMs from a computational point of view. IS-ENES3 provided a unique opportunity to exploit this new set of metrics to create a novel data-base based on CMIP6 experiments, using the different models and platforms available all across Europe.

CMIP6 represents the last iteration of experiments created to understand past, present, and future climate changes arising from natural, unforced variability or in response to changes in radiative forcing in a multi-model context. It is a set of coordinated experiments, designed to understand specific aspects of the model response, where several institutions contribute with different configurations, resolutions and platforms among others. In summary, it is the ideal environment to create a performance data-base from a multi-model context, where both differences and similarities among the models can be observed on a variety of different hardware. Moreover, the current database could be used for different studies, such as the comparison of different models running similar configurations or the same model and configuration but executed on different platforms. All these possibilities create a unique context that has to be exploited by the community to improve the evaluation of the computational performance of the ESMs, using this information for future optimizations.

For this reason, through the IS-ENES3 project, a collection of measurements was planned. The goal was to create a database that could be used for performance analysis and, eventually, for improving the computational performance evaluation of our ESMs. As an example for the reader, the performance analysis could pursue the following goals:

- Compare the performance of models running similar CMIP6 experiments.
- Compare the performance of models with similar complexity

- Compare the performance of platforms when running the same model
- Evaluate the efficiency of a model when the cost increases (more components, higher resolution…).
- Evaluate the performance of a model from different points of view and find main bottlenecks.

Some of these points will be covered in this document through the analysis section as examples, while others were presented in the dissemination done for the first and second general assembly. A complete analysis covering all the points will be done and published in the Geoscientific Model Development journal.

This document, while describing the data collection done, will highlight the coordination effort made by the various groups to ensure dissemination of results and to take advantage of the defined metrics for performing new analyses. Here, the complete CPMIP results are presented, grouped per model/institution involved. Additionally, a section has been included to analyse the results. The analysis has been done pursuing the previously commented goals, illustrating some practical examples, and proving the usefulness of the metric to the community. However, a more detailed analysis will be done in the scientific paper. The conclusions presented here will be extended in the final paper. Moreover, an additional section has been included, specifically for this deliverable, where the main difficulties encountered in the coordination of the collection are highlighted. This section also includes general recommendations on how to solve these problems for future collections. Finally, future steps and conclusions are presented.

2 Description of the collection

WP4/NA3 performed for the first time a complete computational and energy performance analysis of CMIP6 experiments using the CPMIP computational metrics. The collection included IS-ENES3 and some external partners as it is shown in Table 1.

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
U. Read	Grenville Lister, Bryan Lawrence
Nor-ESM	Alok Kumar Gupta
TOPAZ/MOM5	Paulo Nobre
GFDL	Niki Zadeh

Table 1. Model/institutions involved in the CPMIP collection and people in charge.

The collection process included complete tracking and coordination to get the metric results, including meetings, reporting, and surveys in different moments of the CMIP6 simulations (before, during, and after the simulation runs). All the partners included in Table 1 were invited to participate in the tracking process. The coordination, meetings, and reporting were useful to evaluate the state of the collection from the partners, and support was provided to those institutions that required it during the collection process. The support was provided from additional explanations to understand correctly each metric to advices about how to collect specific metrics for a particular environment. Among the activities done, the reporting through a survey was useful to know details about the collection from each partner, the difficulties they encountered and to evaluate potential problems during the collection. This information will be shared in the general recommendations section.

The collection was carried out using a google spreadsheet [2] ready to be directly updated to ES-DOC project [3]. The collection was divided into two steps: the first comprehends the collection up to March 2020, coinciding with the first general assembly where the first results were presented; the second includes the data collected up to the end of 2020 when all the institutions had finished the CMIP6 runs. Finally, ES-DOC completed the last update to ES-DOC during the middle of 2021, publishing CPMIP along with the other CMIP6 results.

3 Collaboration with other groups and activities

During the life-cycle of the collection process, we collaborated with external groups to ensure the quality of the collection or to improve the usability of the results.

3.1 ES-DOC

The Earth System Documentation (ES-DOC) project aims to nurture an eco-system of tools & services in support of Earth System documentation creation, analysis and dissemination. As it is explained in the previous section, ES-DOC was working in collaboration with this task from the beginning. The google spreadsheet template used for the collection was provided by ES-DOC and it was compatible with the future update of the ES-DOC documents. This process proved to be a complete success and the ES-DOC update was done without extra issues.

Additionally, the collection and the analysis were useful to understand how the ES-DOC infrastructure could be improved in the future, providing a finer granularity for the performance metrics. For this reason and as a future requirement, performance metrics per components of a coupled model could be used in future collections. This requirement will be useful to understand the computational bottlenecks at component level.

3.2 HPC-TF

The ENES High Performance Computing Task Force (HPC-TF) aims to advice the ENES Board on all issues relevant to High Performance Computing for the European climate modelling community. Additional data were collected from the ISENES3 partners in parallel with the CPMIP metrics thanks to a collaboration with the HPC-TF group (Jean-Claude Andre and Sophie Valcke). Although the CPMIP metrics provide deep insight on each CMIP6 configuration, the new data required from the HPC-TF provide a general overview of the CMIP6 execution (Table 2), showing: 1) total and useful simulated years, 2) total and useful output produced and 3) core hours spent. Useful simulated years is the number of years simulated for the whole CMIP6 for the useful runs only. Total simulated years is the number of years run for the whole CMIP6, including tuning and runs that were finally discarded for any reason. On the other hand, useful data produced size is the data output volume produced for the whole CMIP6 by useful runs only. Finally, total data produced size is the total data output volume produced for the whole CMIP6, including tuning and runs that were finally discarded for any reason.

CMIP6 Experiments: Institutions/Models	Useful SY	Total SY	Useful Data Produced (PB)	Total Data Produced (PB)	Useful CH (Mh)	Total CH (Mh)
EC-Earth	28,105	38,854	0.80	1.405	31.3	46.8
CNRM-CERFACS	47,000		1.35	2.48	160	365
IPSL	75,000	165,000	1.8	7.6	150	320
CMCC	965		0.965		1.99	
UKMO	37.237		10.4		683	
DKRZ	1,276	1,321	0.606		5.52	5.90
NCC-NORESM2	23,096		0.596		27.23	80
NERC	640		0.460		55.497	
MPI	24,175	35,000	1.925		968.116	

Table 2. Useful and total simulated years(SY), output produced and core hours (CH) spent. The results for each CMIP6 configuration are grouped per model/institution.

As seen in Table 2) the "Total Data Produced" by some institutions is missing. The corresponding institutions reported that it was impossible to measure them during or after the model's executions.

These data will be used by the HPC-TF to analyze and compare CMIP6 results with the data obtained in CMIP5 and as a base for planning CMIP7 's collection.

3.3 Carbon Footprint Group

The CPMIP metrics can be used not only for computational evaluation but also to provide broader information about the analysis. For this reason, we collaborated with the Carbon footprint group created in ISENES3, led by Sophie Valcke and divided in different actions.

The collaboration included the evaluation of the total energy cost of the CMIP6 experiments and to give a first estimation of the carbon footprint related to these experiments. To make this possible, Mario Acosta has been in charge of the Action 4 in order to ensure that the energy cost of each CMIP6 experiment was collected. The total energy cost was produced using useful simulated years and energy cost for each configuration. Finally, the novel total energy cost produced by each configuration was grouped per model/institution to provide a general overview (Table 3). Additionally, the carbon footprint was also calculated following the equation provided by Action 5 [4] to calculate the level 2 suggested in this action (Eq. 1).

Carbon Footprint = Total Energy Cost (MWh) * CF * PUE (Eq. 1) being CF the conversion factor used to convert the kWh produced from the total energy (converted before from J to MWh) to CO2 kilogram according to the supplier bill or the country energy mix. PUE accounts for other costs sustained by the data center, as cooling.

The PUE and Conversion Factor (CF) were collected from each institution in a survey thanks to the coordination of the CPMIP collection. One institution did not provide PUE and CF values. Additionally, NERC reported that their carbon footprint is zero because they have a green tariff from their power supplier.

CMIP6 Experiments: Institutions/Models	Useful SY	Total Energy Cost (Joules)	PUE	CF (g CO2/kWh)	Total Carbon Footprint (CO2)
EC-Earth	28,105	1.24E+12	1.35	357	165t
CNRM-CERFACS	47,000	6.18E+12	1.43	40	97t
IPSL	75,000	8.72E+12	1.43	50	172t
CMCC	965	1.61E+12	1.84	408	329t
UKMO	37.237	2.67E+13	1.35	87	868t
DKRZ	1,276	4.09E+11	1.19	184	24t
NCC-NORESM2	23,096	1.69E+12			
NERC	640	2.17E+12	1.10	0	
MPI	24,175	6.20E+11	1.19	184	37t

Table 3. Total energy cost and carbon footprint grouped per model/institution.

4 Results

The collection was done for all the CPMIP metrics presented in [1]. Table 2 shows a summary for each metric.

Metric	Description
Simulation Year Per Day (SYPD)	Simulated years per day in a 24 h period, collected by timing a segment of a production run of usually one year.
Core-hours Per year Simulated (CHSY)	Simulated years produced with respect to the number of parallel resources used
Complexity (Cmpx)	number of prognostic variables per component
Actual SYPD (ASYPD)	how queue time and interruptions affect the complete experiment duration
Parallelization (Paral)	total number of cores allocated for the run
Energy Cost Per Year (JPSY)	Energy in Joules needed per year of simulation
Memory Bloat (Mem B)	ratio between actual and ideal memory size
Data Output Cost (DO)	time and resources used for performing I/O. The value is given as the percentage added to the simulation without outputs. For example, 1.05 means that DO is 5%.
Data Intensity (DI)	amount of data produced in GB per compute-hour
Coupling Cost (Coup C)	time and resources used in the execution of the coupling algorithm as well as load imbalance among model components. The value is given as the percentage represented comparing to the simulation of the components without coupling. For example, 0.05 means that Coup. C. is 5%.

Table 4. CPMIP metrics and description according to [1].

Complete results are presented below, classified per institution/model. They can be also found in [2].

Model- Exp-Inst.	Resol	Cmpx	SYPD	ASYPD	CHSY	Paral.	JPSY	Coup. C..	Mem. B.	DO.	DI	UsefulYears
EC-Earth3-BSC	1.60 E+07	31	15.2	9.87	1119	768	4.41 E+07	0.080	11	1.12	0.03	14020
EC-Earth-KNMI	1.60 E+07	31	16.2	16.2	1286	868		0.15		1.12	0.074	1009
EC-EarthAer Chem-KNMI	1.64 E+07		3.03	3.03	3549	448		0.16		1.10	0.09	2391
EC-EarthVeg-SMHI	1.60 E+07	39	12.44	6.653	1676	864					0.028	6337
EC-Earth Veg-BSC	1.60E+07	39	12.36	7.42	1269	768	4.87 E+07	0.10	12	1.13	0.04	252
GFDL-CM4-piC	8.35 E+07	31.00	9.98	8.16	15383	6399	5.88 E+08	0.260	16.09	1.24	89.43	657
GFDL-ESM4-piC	2.45 E+07	140.00	8.65	7.46	13570	4893	5.19 E+08	0.270	40.57	1.10	43.99	1124
GFDL-OM4-p25	8.16 E+07	11.00	11.50	7.05	9746	4671	3.72 E+08	0.130	47.64	1.32	173.96	300
GFDL-OM4-p5	2.18 E+07	13.00	15.90	12.22	1962	1300	7.50 E+07	0.140	33.61	1.11	76.86	300

Model- Exp-Inst.	Resol	Cmpx	SYPD	ASYPD	CHSY	Paral.	JPSY	Coup. C..	Mem. B.	DO.	DI	UsefulYears
CNRM-CM6-1	1.02 E+07	181	8.1	7.3	1920	520	4.80 E+07					
CNRM-CM6-1-atm	2.24 E+06	128	7.3	6.1	1320	393	3.50 E+07					
CNRM-CM6-1-HR	2.79 E+08	181	1.5	1.48	19040	1347	5.28 E+08					
CNRM-CM6-1-HR-atm	1.65 E+08	128	2.2	1.8	8720	781	2.28 E+08					
IPSL-CM6A	1.04 E+07	750	12	11.5	1900	950	1.16 E+08	0.050	10	1.20	0.07	53,000
NorESM2-LM	7.77 E+06		13.84	3.03	1664	960	5.60 E+07	0.035			0.065	5463
NorESM2-MM	9.10 E+06		8.96	6.14	4885	1824	1.65 E+08	0.32			0.06	1021
CMCC-CM2-SR5	8.00 E+06	844	6.68	6.5	2068	576	1.67E+09	0.074	17.8	1.04	0.05	965

Model- Exp-Inst.	Resol	Cmpx	SYPD	ASYPD	CHSY	Paral.	JPSY	Coup. C..	Mem. B.	DO.	DI	UsefulYears
HadGEM3-GC31-MM-UKMO	1.42 E+08	236	1.65	1.32	62836	4320	2.33 E+09	0.105	120	1.02	0.05	2386
HadGEM3-GC31-LL-UKMO	1.41 E+07	228	4	3.55	13392	2232	4.97 E+08	0.061	46	1.03	0.074	5610
UKESM1-0-LL-UKMO	1.41 E+07	372	4.3	3.6	16074	2880	5.97 E+08	0.098	4.6	1.03	0.019	15435
DKRZ-MPI-ESM1-HR	2.00 E+07		13.33	11	4728	2616	3.21 E+08					1864
MPI-ESM1-LR	3.12 E+06		55.6	22.7	388.7	878	2.56 E+07					18860
MPI-ESM1-LR-ATM	8.67 E+05		45.9	25.2	163.2	312	1.11 E+07					991
MPI-ESM1-LR-LAND	8.67 E+05		282.8	265.4	3.1	36	1.39 E+05					2460
UKESM1-0-LL-NERC	1.14 E+07	252	2.02	1.1	8568	720	3.18 E+08	0.078	28	1.19	39	195
UKESM1-AMIP-NERC	2.35E+06	202	1.64	1.41	7358	504	1.04 E+08	0	52.5	1.31	25.7	45
HadGEM3-GC3.1-SS_NERC	1.14E+07	150	4.25	1.06	12268	2160	4.33 E+08	0.047	56.8	1.41	194	70
HadGEM3-GC3.1-MM_NERC	1.97E+08	54	0.58	0.46	192412	4656	7.70 E+09	0.21	154		107	65
HadGEM3-GC3.1-LL_NERC	1.24 E+09	54	0.49	0.34	585540	12024	2.30 E+10		183	1.41	207	65
IITM-ESM	4.8 E+06	168	8	7	996	332	3.81 E+07		36.7			845
IMPE- BESM	6.9 E+06	132	3.6	3.4	1857	278					37.1	360

Table 5. CPMIP complete results for each CMIP6 configuration.

As the reader can see, specific institutions/models are missing some metrics. Usually, the metrics missing are Coup. C., Mem. B., and DO. This is due to the difficulties that can arise when trying to collect these metrics, a process much more complicated than collecting, for example, SYPD. These metrics were provided by the institutions at the end of the first and the second stage. While some institutions required support to produce the data, three of them were not able to produce those figures at all, since they did not have the time or the resources during or after the CMIP6 runs and/or the machine was changed.

5 Methodology

We gathered all the metrics from Table 5 to statistically prove the correlations between them. However, some outliers were found (Figure 1) due to the diversity among configurations and models (resolutions, number of components of the model used, etc.).

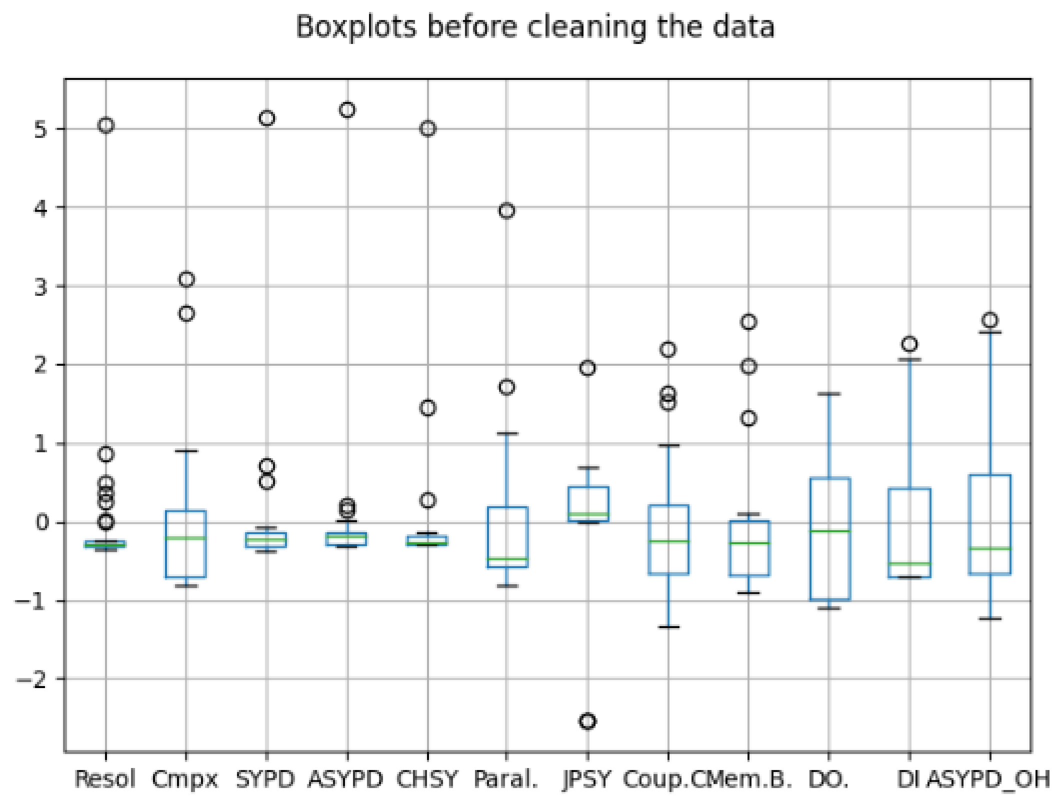

Figure 1. Metric results grouped. ASYPD_OH (ASYPD overhead) is the extra percentual time comparing the ASYPD (simulation + queues, interruptions, failures...) and useful SYPD, introduced to quantify ASYPD properly for the analysis.

Keep in mind that the goal of this section is to prove the usefulness of the CPMIP metrics and not to carry out their complete analysis. The results presented will be analyzed in full details in the aforementioned paper, still pending for publication. We started excluding configurations which were more than two standard deviations apart from the mean. After that, we used Pearson's correlation coefficient (higher than 0.7) to get some insights in the relations between the metrics and to provide some examples for this deliverable (Figure 2). Pearson's correlation coefficient is a measure of linear correlation between two sets of data, being the ratio between the covariance of two variables and the product of their standard deviations Although the data-set is limited (around 30 configurations) for a Pearson’s correlation, the results helps us to support our expectations and assist us at understanding the data we have.

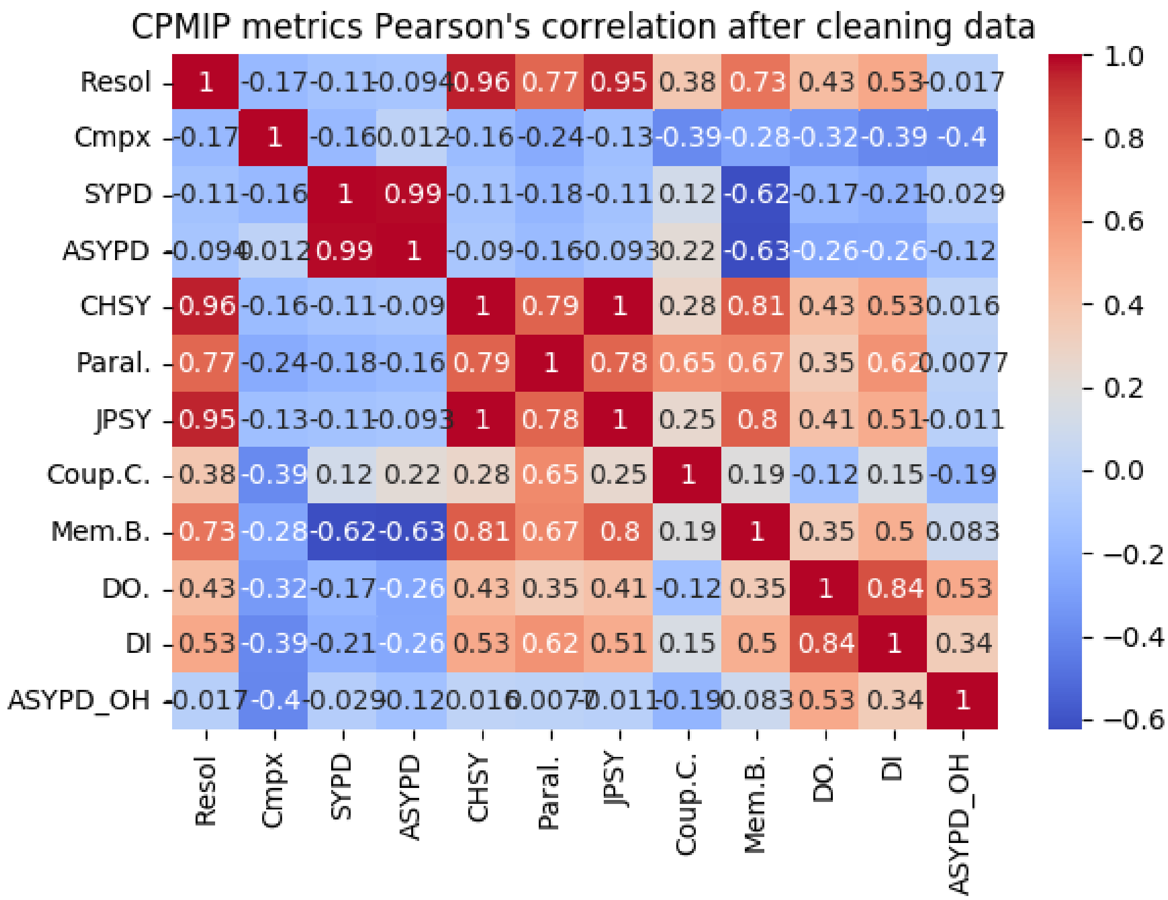

Figure 2. Pearson’s correlation among all metrics after cleaning the data.

6 Analysis examples

Pearson's correlation helps to support some theoretical expectations, for example, to prove for all the configurations studied the clear relation between DI and DO. or between Resolution and CHSY.

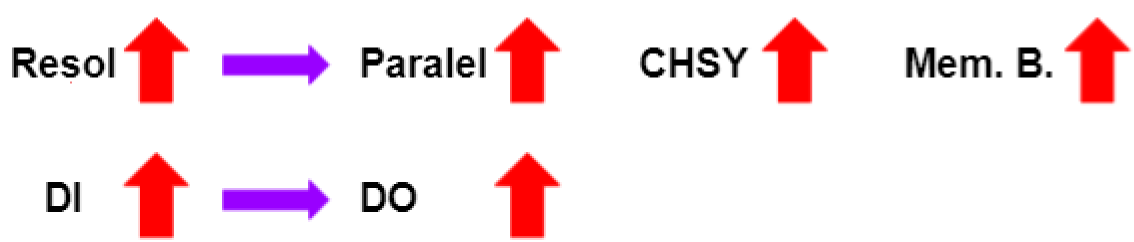

There are two direct relations:

1 The higher Resolution, the higher the values of: Paralel, CHSY, and Mem. B.
2 The higher the DI, the higher the DO..

This result is quite logical. However, it is particularly interesting to find this clear relation between Resol and Mem.B. from the beginning, which supports the hypothesis that all ESMs are memory bound and lose efficiency in the memory access when using a high number of parallel resources, at least for modern supercomputers.

However, for more specific conclusions, not only dropping configurations clearly outside the spread was needed. The analysis has been done by grouping manually experiments with similarities (as resolution or complexity) and adding thresholds (DO. higher than 20% for example) to study the results properly. It is important to highlight that the correlation study has been useful to find connections. However, the next examples do not imply only correlation but also causality between metrics. This means that some of them as complexity are independent but the increment of those independent metrics imply causality in other metrics looking for a specific goal (such as maintain a similar SYPD when we introduce more complexity in a model, increasing the number of parallel resources).

6.1 Resolution impact

Once we group some models/configurations with similar complexity, we find a new connection.

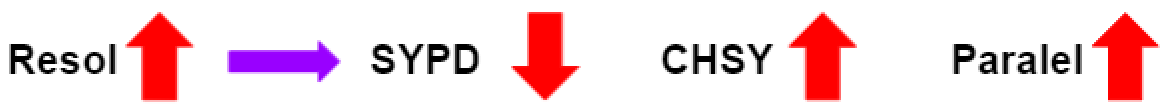

Although with the highest resolutions we are talking about the most demanding configurations from CMIP6, this supports the hypothesis that there is no model using those highest resolutions which is scaling ideally for the modern supercomputers. There are some possible reasons explaining this and not necessarily related to an inefficiency in the model. This could be due to hardware limitations (high latency or low bandwidth), the high cost of other phases (Coup C, DO., etc.), the inherent overhead introduced for a higher number of parallel resources and/or the memory footprint… or a combination of those.

6.2 Complexity impact

From the results, we can also infer the complexity impact. A higher complexity for a specific model (i.e., adding a new component to a model to obtain a coupled version) requires to increase CHSY to maintain a similar SYPD (considering that other phases have not a significant impact, such as DO. or CO)

However, we notice for experiments such as EC-Earth and CNRM where

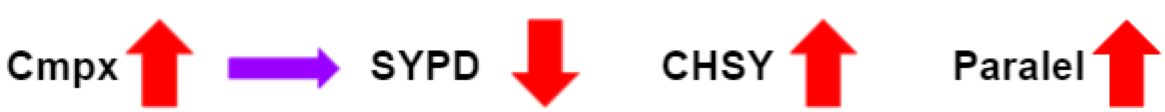

Increasing Cmpx reduces SYPD and increases CHSY (comparing coupled and amip versions between LR and HR configurations when we increase Paralel looking for maintain a similar performance). In these two cases it is likely that the component added (the ocean) is less efficient or more computationally expensive. On the other hand, we also have cases such as GFDL experiments, where it is seen that a constant SYPD is achieved if the complexity and CHSY increase.

6.3 Data Output impact

We can also use the model intercomparison to identify when a value for a certain metric indicates a bottleneck. Take for example DO.: as explained in the beginning, there is a direct relation between DI and DO. Additionally, notice that a DI higher than 1 GB per Hour implies a DO. higher than 20%.

If we classify the results according to this new threshold (DO. higher than 20%) we find a new relation connecting all the experiments remaining in this group, i.e., DO affects SYPD negatively.

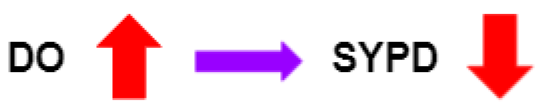

This could prove that configurations with a DO. higher than 20% should be studied and the performance of the IO evaluated separately. Moreover, this approach can be used to identify those values starting from which a particular metric represents a significant bottleneck for the execution.

6.4 ASYPD analysis

ASYPD_OH measures the waste of time in percentage comparing ASYPD and SYPD. This number is around 10-50% for most of the cases studied, and it can be classified in two groups

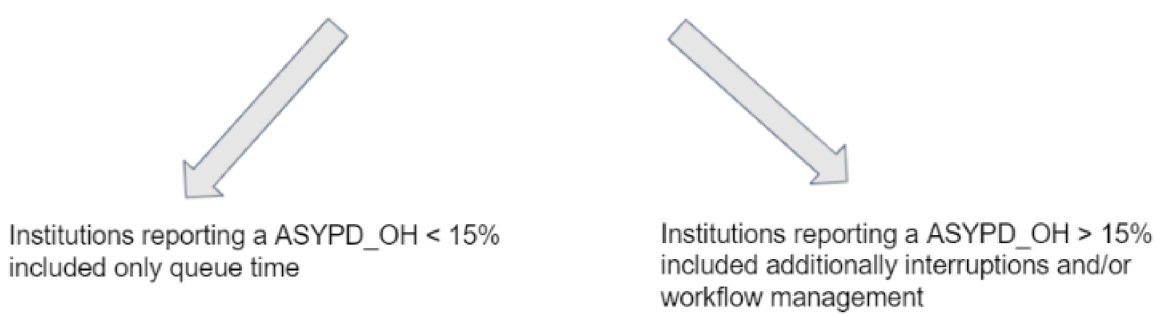

The results could support the idea that queuing time represents an increment around 10-20% compared to SYPD. On the other hand, adding interruptions and workflow management could increase up to 40-50% the final execution time. However, ASYPD is a metric where a finer granularity and the use of sub-metrics can be useful. BSC CMIP6 results using the same configuration on two different platforms (Marenostrum and CCA) proved that the percentage of each part (queue time, interruptions or post-processing) could change among platforms even though the CMIP6 experiment is the same (see first GA presentation for more details).

6.5 Coupling cost impact

In most of the cases Coup. C. is around 3-15% of the execution time of the model simulation, we can find a common connection when Coup. C. is studied:

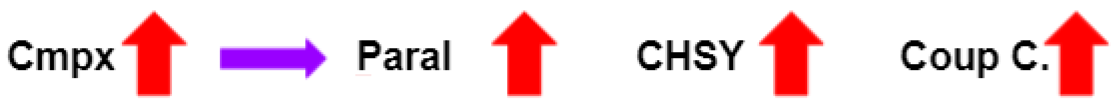

The increase in Coup C. is not necessarily related to a decrease in the performance. However, a specific case has been identified, for example for NERC or CNRM experiments, which should be studied:

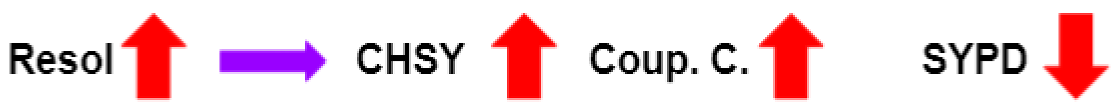

This case could reflect a problem in the coupling phase. It can occur that the coupling algorithm is not scaling correctly or that the higher resolution configuration is not well-balanced among coupled components (the waiting time between components increases when one or more components are faster than others). This also means that for high resolutions and not well-balanced cases, institutions will spend more resources to balance the coupling, since the higher is the resolution, the higher will the CHSY needed to do the tests. Since there are no specific tools to balance a coupled model, these institutions are forced to use a trial and error approach, which is not trivial for complex configurations with several components or differences in the time step among different components.

For these cases, a finer granularity in the Coup.C. metric and new ways to achieve a well-balanced configuration could be needed, splitting interpolation algorithm and waiting time in different sub-metrics or providing other sub-metrics (SYPD, CHSY…) not only for the coupled version but also per component.

7 General recommendations

Below are the four main drawbacks that were detected during the coordination of the data collection. Additionally, recommendations are included on how to successfully solve these problems in future collections.

Drawback N. 1	Although the CPMIP collection is important, it is secondary during the CMIP execution. During the execution itself all the resources (humans and computational) are spent to finish the simulations successfully. This makes clear that the main goal is to finish the experiments and produce the scientific outputs. The institutions rely on collecting the metrics after the CMIP6 experiments, re-running some chunks. However, the time and resources after the CMIP6 execution could not be enough. Some institutions have reported that re-running experiments (even partially) is too expensive. In some extreme cases, the machine was changed before completing the collection, making it impossible to finalize it.
Actions recommended	• Collect CPMIP metrics before or during the CMIP experiment: Spend some resources before the CMIP experiments or at least during the spin-up/tuning process. This will avoid wasting resources only for the performance collection. Moreover, the preliminary analysis during the tuning could be useful to improve the throughput of the final executions. • Only ISENES3 partners provided the complete set of metrics, proving that a coordinated collection (as the collection done from WP4 in ISENES3) is useful and critical to get the performance results. This coordination should include reporting, support and tracking from each partner. • To facilitate this process, the development of portable and automatic processes such as the integration with workflow managers could be a solution. • Even following the previous recommendations, some metrics as Coup C could be difficult to collect without the proper expertise or without wasting resources. For this reason, the HPC community and tools developers should provide new tools to facilitate the collection of these metrics. One example is the new load balance analysis tool for OASIS which will provide the coupling cost without additional runs.
Drawback N. 2	There was a strong commitment from the institutions involved to collect the metrics. However, data was missing from some institutions eventually. In some cases this was justified for technical reasons. However, the technical reasons were not solved since the extra effort and resources required were not justified for a dissemination activity (provide the numbers as the final goal). This proves that the institutions do not yet consider these metrics as a

	powerful tool to analyze the computational efficiency of a model across platforms, configurations or for multi-model comparisons. Additionally, there are other reasons why institutions could not obtain conclusions directly from the collection done for their model. First of all, because some HPC expertise was needed to interpret the results. Secondly, the limitations of the hardware can be a bottleneck difficult to distinguish from the inefficiency of a model.
Actions recommended	• CMIP6 was the first coordinated collection and the results are still limited. However, there will be more collections and the bigger the data-base will be, the easier it will be to follow recommendations coming from the analysis done previously. • Create new surveys and tools to provide details about the components of a coupled model and the configurations. For example, ES-DOC platform will be updated to provide metrics per component. • Running the same model across different platforms will facilitate the differentiation of the limitations of the hardware from the computational problems of the model. Additionally, the integration of the CPMIP metrics in complex and representative benchmarkings as dwarfs (i.e. ESCAPE2) could identify limitations of the platform before the real executions are used.
Drawback N. 3	We identify inconsistencies in the way some metrics are collected, which makes inter-model comparison difficult. Some differences are coming from technical limitations and are not controlled from the person in charge of the collection. The identified metrics are: • ASYPD: Many institutions are including queuing time and interruptions, while others are including only queuing time or adding data movement, data leaning, etc. Some results from BSC [4] proved that most of the penalization from SYPD to ASYPD could come from different sources, such as the inclusion of queuing time, executions interruptions or post-processing time. The aggregate value alone cannot explain where the problem is coming from. • JPSY: The value is taken from different sources: usually it is provided by the department in charge of operating the machine and not decided by the person in charge of the collection. The source is not known for some platforms, making it difficult to know if the values are comparable or not across different machines.
Actions recommended	If there is a coordinated collection before the CMIP runs, it will be recommended to normalize the way to collect some specific metrics. Additionally, new tools and approaches could be developed which could collect the metrics across different platforms and facilitate the

	normalization. The integration of the collection through workflow managers for long periods of simulations could facilitate the process too.
Drawback N. 4	Even though some metrics point out a clear bottleneck, the lack of expertise or additional information (per component for instance) could make it the task of identifying and improving the computational bottlenecks even harder. For example, when we have a very high coupling cost: Is the Coup C coming from the coupling algorithm itself or simply a load imbalance issue? Should the institution waste resources using a trial and error approach to minimize Coup C? Yet a similar example: if we have an important reduction from SYPD to ASYPD, is the reduction coming from the queue time, the interruptions or the postprocessing? Could it affect other parts (such as cmorization) of the workflow to ASYPD? Could the real problem change among configurations or platforms and should it be studied independently?
Actions recommended	• If some metrics are sensitive, we recommend to create a finer granularity, splitting them in different metrics or sub-metrics. • ASYPD should have three sub-metrics to measure queue time, interruptions and post-processing separately. • Coup C should identify coupling algorithm cost and waiting time separately, as the new load balance analysis tool being developed for OASIS. • Longer runs are recommended to evaluate interruptions.

8 Future Steps

Although the main goal for ISENES3 has been accomplished with the metrics' collection and this deliverable, some questions remain. The break-out session in the second GA was useful to determine if the partners recognize the usefulness of the CPMIP metrics. However, the collection of some of them is still difficult in a critical moment where all the resources are dedicated to the CMIP execution itself.

According to this fact, a final coordination will be done to collect feedback from all the institutions to ensure that the difficulties found are solved. A new survey could be the solution to ensure active participation from all the institutions and not only from those more interested in this topic.

Additionally, the new data-base from CMIP6 will be used to do a complete analysis and disseminate to the community the new CPMIP results. A scientific paper will be published in the GMD journal in collaboration with V. Balaji and all the people involved in the collection. The paper should be finished before the end of February 2022.

The coordination done for ISENES3 has proved to be very useful to ensure the collection in a multi-model context, providing support from HPC experts and tracking the collection during a critical moment as it is the CMIP execution. For these reasons, a new coordination is recommended for future CMIP iterations, starting before or during the set-up process. The goal should be to confirm that each institution is ready to provide each metric before the current CMIP iteration starts. Moreover, additional collections will be useful to increase the capacity of analysis of the novel data-base, adding new results and confirming with new data any conclusion achieved for the actual set of metrics. This data-base could be used for the communication with vendors and improve the decision-making when a novel hardware has to be designed or bought.

Finally, the collection itself has been useful to determine which metrics are not trivial to collect without domain specific tools. Another goal will be to collaborate with developers and HPC experts to provide new tools which will automatize and facilitate the collection of some metrics. For example, LUCIA has been improved during ISENES3 to improve the collection and comprehension of Coup C for models using the OASIS coupler. Additionally, a new collaboration with ES-DOC developers will be valuable to include some CPMIP metrics per component (of a coupled model), facilitate the analysis and focus the performance evaluation only on the components which should be considered as the main bottlenecks.

9 Conclusions

This deliverable shows the main results obtained from the CPMIP collection done for CMIP6, including ISENES3 and some external partners. The document provides the list of partners involved and the CPMIP metrics per institution/model, as well as the approach used for the collection and the coordination behind this process. Additionally, a section has been included to explain the collaborations done with other groups (HPC-TF and carbon footprint) and the results produced, proving that the possibilities using CPMIP metrics go beyond a simple performance evaluation.

This document also includes a general analysis and some examples to prove opportunities for the performance evaluation possible with the new CPMIP results. Although the analysis will be improved for a final peer-reviewed publication, the examples can give the users an idea about how useful the metrics can be beyond the idea about fill in a table including results of several models.

The experience from the collection and coordination has also been resumed in four main issues. These points should be of special interests for the partners in order to improve and facilitate future collections and be ready for CMIP7. Some recommendations have been given to solve these problems, which could be adopted from the partners or used as the base for the development of novel tools intended to facilitate this work.

Finally, some ideas have been discussed to continue this work: from new collections of metric to tool developments intended to facilitate metrics collection, improve metrics granularity or even the creation of a complete database to be shared with the community or vendors.

10 References

[1] Balaji, V., Maisonnave, E., Zadeh, N., Lawrence, B. N., Biercamp, J., Fladrich, U., Aloisio, G., Benson, R., Caubel, A., Durachta, J., Foujols, M.-A., Lister, G., Mocavero, S., Underwood, S., and Wright, G.: CPMIP: measurements of real computational performance of Earth system models in CMIP6, Geosci. Model Dev., 10, 19–34, https://doi.org/10.5194/gmd-10-19-2017, 2017.

[2] Machine and Performance version 1 for ISENES3 (final results 2020): https://docs.google.com/spreadsheets/d/1rNSGb4uZHMlx4jnrlBMtIQ6mGWXWDUTeE-9ntjn4FxE/edit#gid=1852886891

[3] Pascoe, C., Lawrence, B. N., Guilyardi, E., Juckes, M., and Taylor, K. E.: Documenting numerical experiments in support of the Coupled Model Intercomparison Project Phase 6 (CMIP6), Geosci. Model Dev., 13, 2149–2167, https://doi.org/10.5194/gmd-13-2149-2020, 2020.

[4] Carbon footprint actions: https://drive.google.com/drive/folders/1DJ-lNBjKPeml9Zj6o96LGCt7qTFSxNhy