Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Reducing Carbon Footprint in Climate Modeling: Unveiling the Efficiency and Energy Costs of Our Models

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

ENES MHPCTF

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

Mario Acosta and Sergi Palomas

September 2024

Why?

- IPCC 1.5°C report (2018): reduce CO_2 emissions by 7% every year to meet 0 emissions objective by 2050
- Increasing expectations from the society to meet these goals
- Growing initiatives within the scientific community
- First step: carbon footprint

Measure greenhouse gases emitted by activities at BSC-Earth in CO_2 equivalents (CO_2e)

Measure energy cost by activities such as coordinated experiments as CMIP6 using a common set of metrics (CPMIP)

Why?

- IPCC 1.5°C report (2018): reduce CO_2 emissions by 7% every year to meet 0 emissions objective by 2050
- Increasing expectations from the society to meet these goals
- Growing initiatives within the scientific community
- First step: carbon footprint

Measure greenhouse gases emitted by activities at BSC-Earth in CO_2 equivalents (CO_2e)

Measure energy cost by activities such as coordinated experiments as CMIP6 using a common set of metrics (CPMIP)

Carbon footprint of BSC-Earth 2018

Taken into account:

- Commuting
- Computing infrastructure
- Building & Infrastructures
- Travels

Total budget of BSC-Earth: **?**

Equivalent per person: **?**

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Commuting

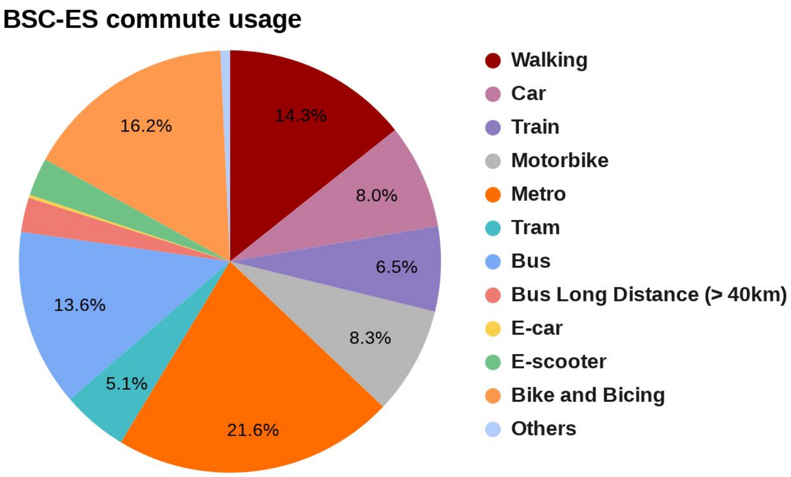

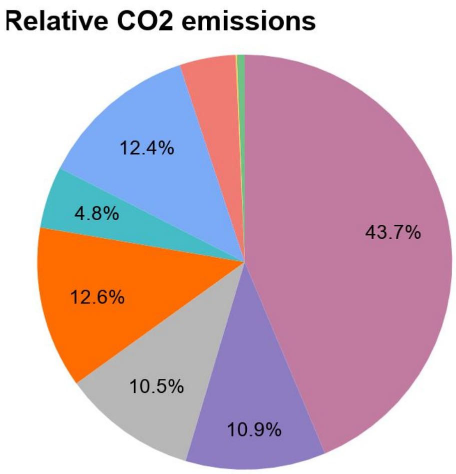

Total budget of BSC-ES :

29 tCO_2e / yr

Equivalent per person:

0.4 tCO_2e / person / yr

- **Walking and biking** produce **0** emissions (31% of BSC-ES people).
- **44%** of commuting emissions due to **cars** (8% of BSC-ES people).

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Building & Infrastructures

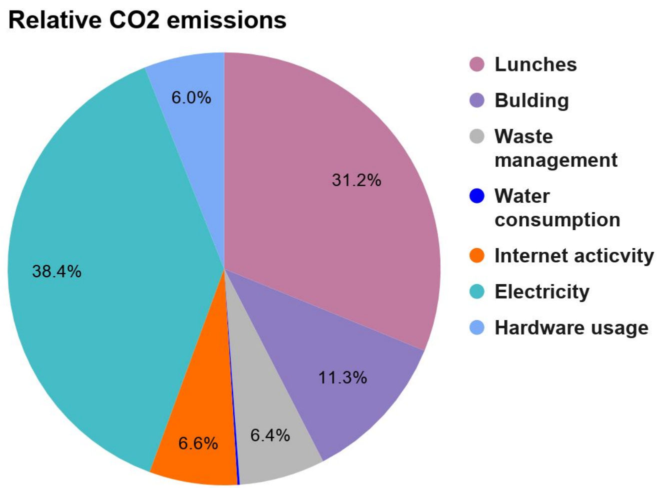

Total budget of BSC-ES :

117 tCO_2e / yr

Equivalent per person:

1.5 tCO_2e / person / yr

- **Electricity and meals** account for 2/3 of emissions.
- High uncertainty!

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Travels

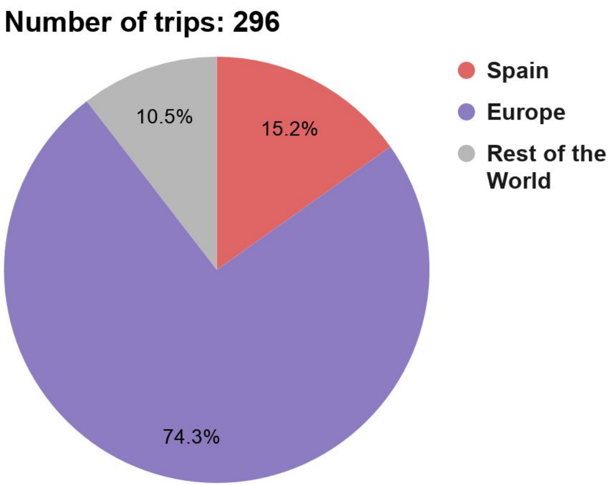

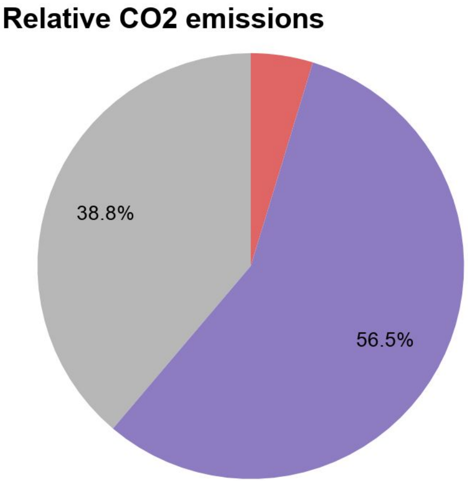

Total budget of BSC-ES :

255 tCO_2e / yr

Equivalent per person:

3.2 tCO_2e / person / yr

- **Plane >> train**: e.g. Bcn-Madrid, 300 >> 40 kg CO_2e / trip
- **Overseas trips** are high contributors

Computing infrastructure

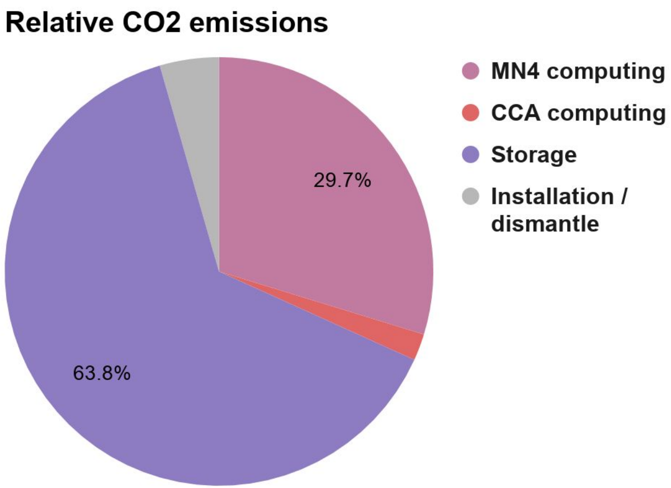

Total budget of BSC-ES :

397 tCO_2e / yr

Equivalent per person:

5 tCO_2e / person / yr

- Storage main contributor (tape system will ↓ emissions)
- Highly dependent on electricity sources

Carbon footprint of BSC-Earth 2018

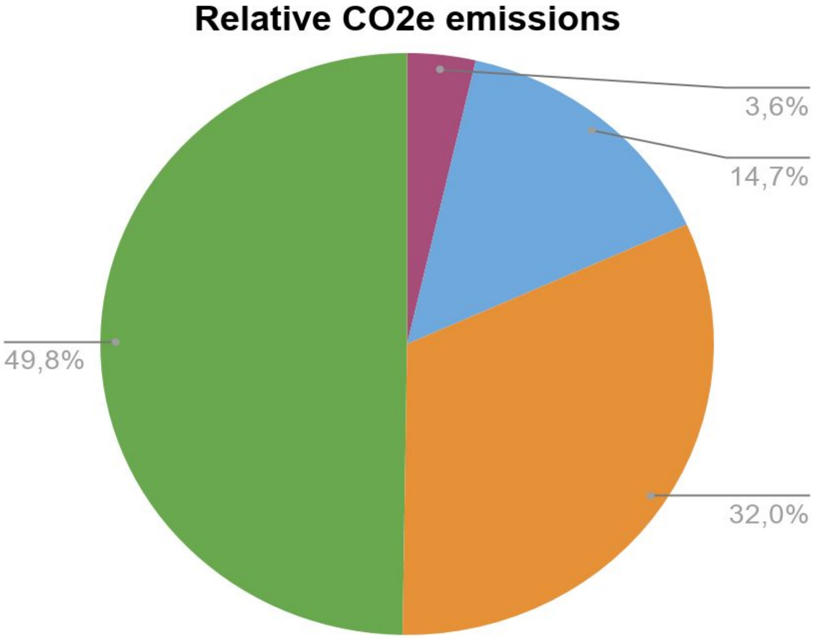

Taken into account:

- Commuting **29** tCO_2e / yr
- Computing infrastructure **397** tCO_2e / yr
- Building & Infrastructures **117** tCO_2e / yr
- Travels **255** tCO_2e / yr

Total budget of BSC-Earth **798** tCO_2e / yr

Equivalent per person **10** tCO_2e / yr

Bcn: 2; Spain: 5; World: 5 (tCO_2e / yr)

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

From awareness to action?

- Carbon footprint: what’s next?
  - Updating annually, compare pandemic and non-pandemic years
  - Incorporating CO_2 accounting to management tools, ↓ uncertainty
- Taking action, individual and collective. What if we…
  - Use 100% renewable energy: ↓60% CO_2e
  - Fly less:
    - ➢ Use video conference: cost-effective (win-win)
    - ➢ Take the train

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Why?

- IPCC 1.5°C report (2018): reduce CO_2 emissions by 7% every year to meet 0 emissions objective by 2050
- Increasing expectations from the society to meet these goals
- Growing initiatives within the scientific community
- First step: carbon footprint

Measure greenhouse gases emitted by activities at BSC-Earth in CO_2 equivalents (CO_2e)

Measure energy cost by activities such as coordinated experiments as CMIP6 using a common set of metrics (CPMIP)

Computational Collection - Why?

- 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 and complete workflow 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 metrics

Model and platform

- **Resolution**

  Number of grid-points

- **Complexity**

  Number of prognostic variables (i.e. number of variables to contain the complete state of the ESM)

- **Platform**
  - Core count
  - Clock frequency
  - Double-precision operations per cycle
  - Memory per core

The metrics

Model and platform

Resolution

ESMs in CMIP6 often use more grid-points for the OCN than for the ATM

Useful to:
- Compare models that have a common target
- Evaluate the costs of increasing the resolution

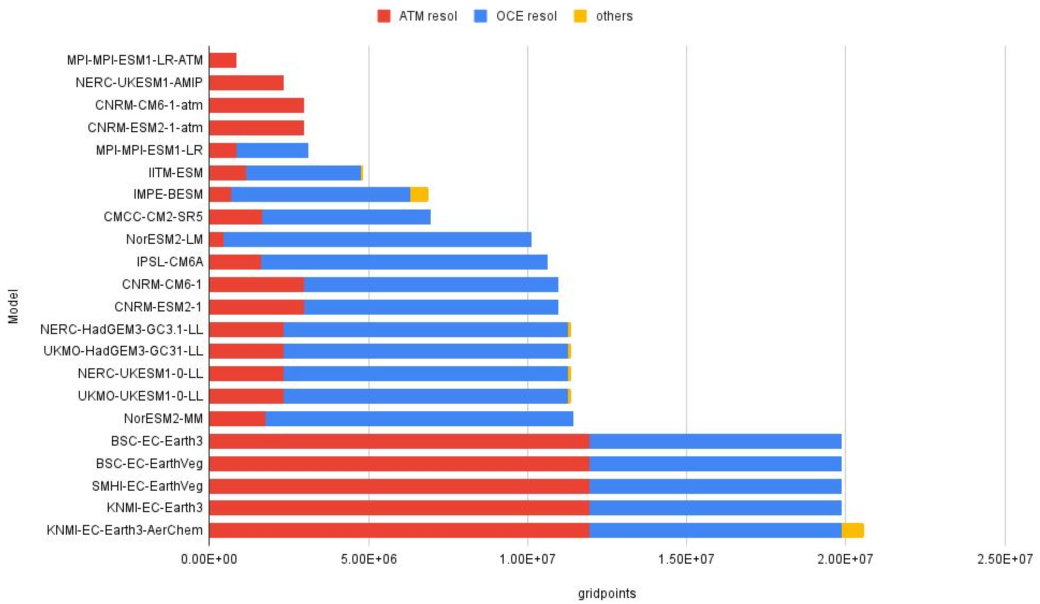

The metrics

Model and platform

Complexity

Increasing the Complexity of a model reduces the parallel efficiency

Hard to compare across different ESMs, useful when we compare the same model and different configurations.

Increasing the resolution does not change the Complexity

Institution	Model	Resol	SYPD	CHSY	Paral	Cmplx	Cpl C
GFDL	OM4-p25	1.26E+08	11.5	9748	4671	11	0.13
	OM4-p5	3.32E+07	15.9	1962	1300	13	0.14
	CM4	1.28E+08	9.98	15388	6399	31	0.26
	ESM4	3.76E+07	8.65	13576	4893	140	0.27
CNRM	CM6-1-atm	2.98E+06	7.3	1292	393	128	
	CM6-1	1.10E+07	8.1	1541	520	181	
	CM6-1-HR-atm	2.36E+07	2.2	8520	781	128	
	CM6-1-HR	1.37E+08	1.5	21552	1347	181	
	ESM2-1-atm	2.98E+06	7.1	1352	400	335	
	ESM2-1	1.10E+07	4.7	4289	840	393	

Cmpx ↑ → SYPD ↓ CHSY ↑ Paralel ↑ Cpl. C ↑

The metrics

Computational cost

- **SYPD**

  Simulated years per day of execution (24h). Measure of the speed of the ESM

- **CHSY**

  Core hours per simulated year. Measure of the execution cost to simulate 1 year

- **Parallelization**

  Total number of cores

nproc	SYPD
48	4.1
128	10.4
208	16.2
288	22.8
368	27
432	32.2
528	35.8
608	39
688	43.3
768	46.8
848	46.6

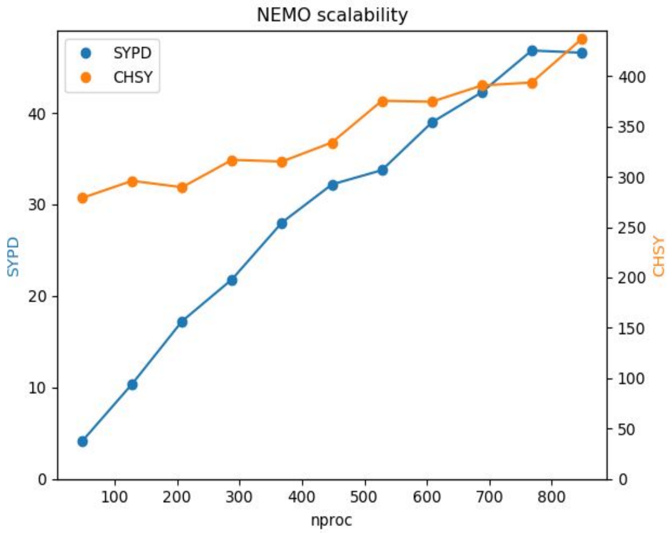

The metrics

Computational cost

- **ASYPD**

  Actual SYPD. ESM speed including system interruptions, queue time, post-processing, workflow management

- **JPSY**

  Energy cost of the simulation. Joules consumed to simulate 1 year.

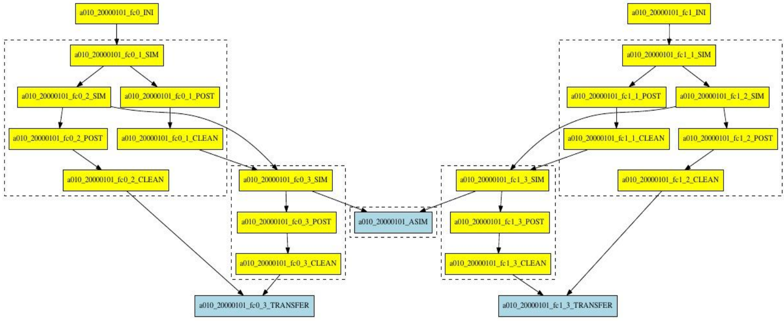

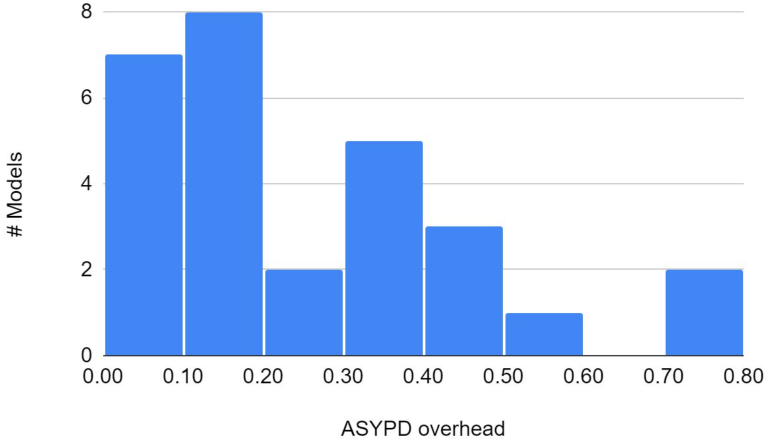

The metrics

Coupling, memory and I/O

- **Coupling Cost**

  Resources wasted due to the coupling (algorithm and load-balance)

- **Memory Bloat**

  Ratio between actual and ideal memory size

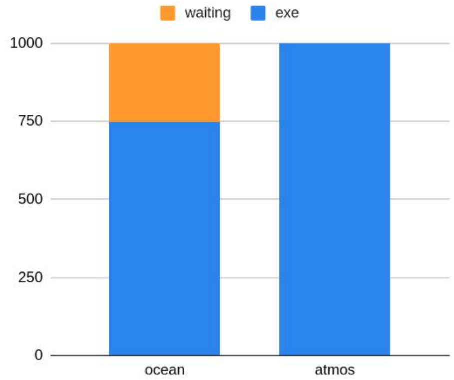

The metrics

Coupling, memory and I/O

- **Coupling Cost**

  Resources wasted due to the coupling (algorithm and load-balance)
- **Memory Bloat**

  Ratio between actual and ideal memory size
- **Data output cost**

  Cost associated to the I/O operations
- **Data intensity**

  Ratio between the amount of data produced (in GB) and the CHSY

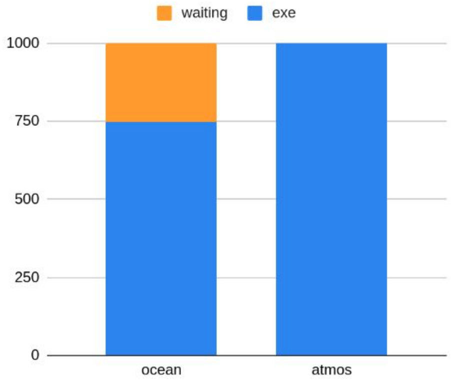

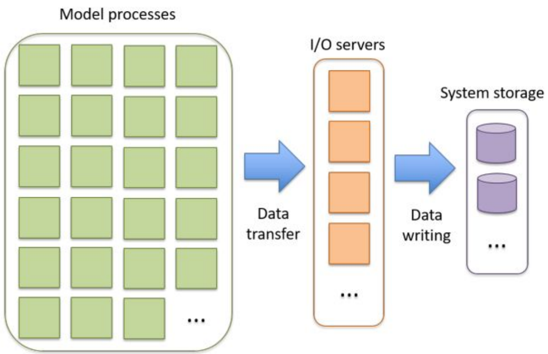

How should we collect the metrics…

…using workflow managers as Autosubmit (AS)

Directly obtained from AS

- SYPD
- ASYPD
- CHSY
- Paral

AS Databases

AS & Slurm & filesystem

- JPSY
- Memory bloat
- Data Intensity

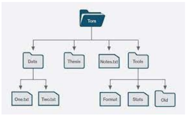

Manually collected

- Resolution
- Complexity
- Platform
- Coupling Cost
- Data Output cost

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

CMIP6 collection: Community List

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

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

CMIP6 General numbers from ENES-HPCTF

CMIP6 Experiments: Institutions/Models	Useful SY	Total SY	Useful Data Produced (PB)	Total Data Produced (PB)	Useful CH (Mh)	Total CH (Mh)	Total Energy Cost (Joules)
EC-Earth	17,598	27,568	0.73	1.34	27.3	41.8	1.27×10^{12}
CNRM-CERFACS	23,620	72,000	1.2	1.98	106.4	325	3.13E+12
IPSL	75,000	165,000	1.8	7.6	150	320	6.16E+12
CMCC	965	2049	0.965	0.281	1.99	4.86	1.61E+12
UKMO	23,431	49747	7.3	21.35	473	1155	1.76E+13
DKRZ	1,276	1,321	0.606	1.78	5.52	5.90	4.09E+11
NCC-NORESM2	23,096	49,036	0.596	1.74	27.23	80	4.75E+11
NERC	640	1359	0.460	1.35	55.497	135.5	2.17E+12
MPI	24,175	35,000	1.9	5.63	968.116	2365	6.20E+11

* Red Values are proportional estimations

Complete results here: https://gmd.copernicus.org/articles/17/3081/2024/gmd-17-3081-2024.html

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

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

Further analysis

Carbon Footprint

Carbon Footprint = Total Energy Cost × PUE × CF

Total Energy Cost from JPSY
PUE to account for the platform efficiency
Conversion Factor (CF) from MWh to CO2

The total Carbon Footprint is **1692 tCO2** (only from 8 out of 45 institutions)
Equivalent to driving 377 gasoline cars 24h a day during 1 year

CMIP6 Experiment	Total Energy Cost (Joules)	PUE	CF (g CO2/KWh)	Total Carbon Footprint (CO2)
EC-Earth	1.24E+12	1.35	357	165t
CNRM-CERFACS	6.18E+12	1.43	40	97t
IPSL	8.72E+12	1.43	50	172t
CMCC	1.61E+12	1.84	408	329t
UKMO	2.67E+13	1.35	87	868t
DKRZ	4.09E+11	1.19	184	24t
NERC	2.17E+12	1.10	0*	0
MPI	7.10E+11	1.19	184	42t

*Green tariff according to NERC

Further analysis

Institutions Useful core-hours vs Total Energy consumption. The size of the bubbles is proportional to the reported Carbon Footprint (Tons of CO2). Note that the x-axis is in logarithmic scale.

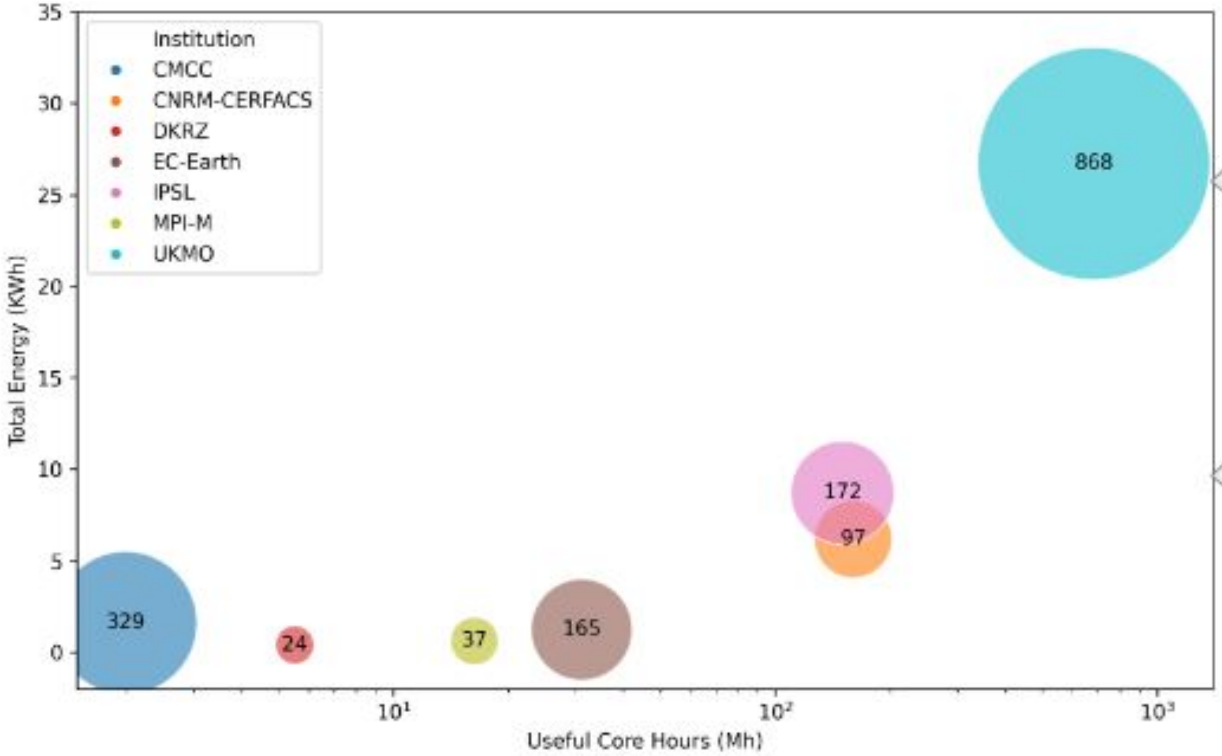

Increased core-hour usage leads to higher energy consumption, which in turn increases carbon footprint estimates

Machine energy efficiency (PUE) and the conversion factor (MWh - kg CO2) significantly impact the Carbon footprint, as noted by CMCC

From awareness to action?

- Carbon footprint: what's next?
    - Updating annually, compare pandemic and non-pandemic years
    - Incorporating CO_2 accounting to management tools, ↓ uncertainty
- Taking action, individual and collective. What if we…
    - Use 100% renewable energy: ↓60% CO_2e
    - Fly less:
        - ➢ Use video conference: cost-effective (win-win)
        - ➢ Take the train
    - Use the analysis to understand and introduce optimizations in the computing part

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

CPMIP Results

- A higher complexity for a specific model (add a new component as a coupled version) require increasing CHSY to maintain a similar SYPD (always that other phases maintain a low profile (DI, CO…))

Cmpx ↑ → **SYPD** ↓ **CHSY** ↑ **Paralel** ↑

- If you find the previous relation (such as NERC and CNRM experiments, comparing coupled and amip version between LR and HR), it is probably that the component added (the ocean) is less efficient or more computational expensive.

- We have positive cases as GFDL experiments, where they achieve a similar SYPD.

Cmpx ↑ → **SYPD** = **CHSY** ↑

CPMIP Results

- We can use the inter-model comparison to identify when the value of a metric could represent a problem

DI ↑ → DO ↑

DI > 1GB/H → DO > 20%

- We classify the results according to a new threshold → DO >= 20% and we find that there is a clear relation for all the experiments in this group

DO ↑ → SYPD ↓

- This could prove that configurations with a DO higher than 20% should be studied and evaluate the performance of the IO approach

CPMIP Results

- ASYPD_OH is around 10-50% for the most of the cases.
- ASYPD can be classified in two clear groups

Institutions reporting a ASYPD_OH < 15% included only queue time

Institutions reporting a ASYPD_OH > 15% included additionally interruptions and/or postprocessing/workflow management

- The results could support that queue time represents around 10-20% and adding interruptions and post-processing/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.

CPMIP Results

- In most of the cases Coup C. is around 3-15% where:

Cmpx ↑ ➡ **Paral** ↑ **CHSY** ↑ **Coup C.** ↑

- The increase in Coup C. is not necessarily related to decrease the performance. However, a specific case has been identified and it should be studied:

Resol ↑ ➡ **CHSY** ↑ **Coup. C.** ↑ **SYPD** ↓

- It could be a problem in the Coupling phase, maybe because the coupling algorithm is not scaling correctly or maybe simply because the higher resolution configuration is not well-balanced.
- This means also that for high resolutions and not well-balanced cases, institutions should spend more resources to find the balanced version using the trial and error approach… and not always it is trivial.
- A finer granularity and new ways to achieve a well-balanced configuration could be needed

CMIP6 analysis: Data Transfer

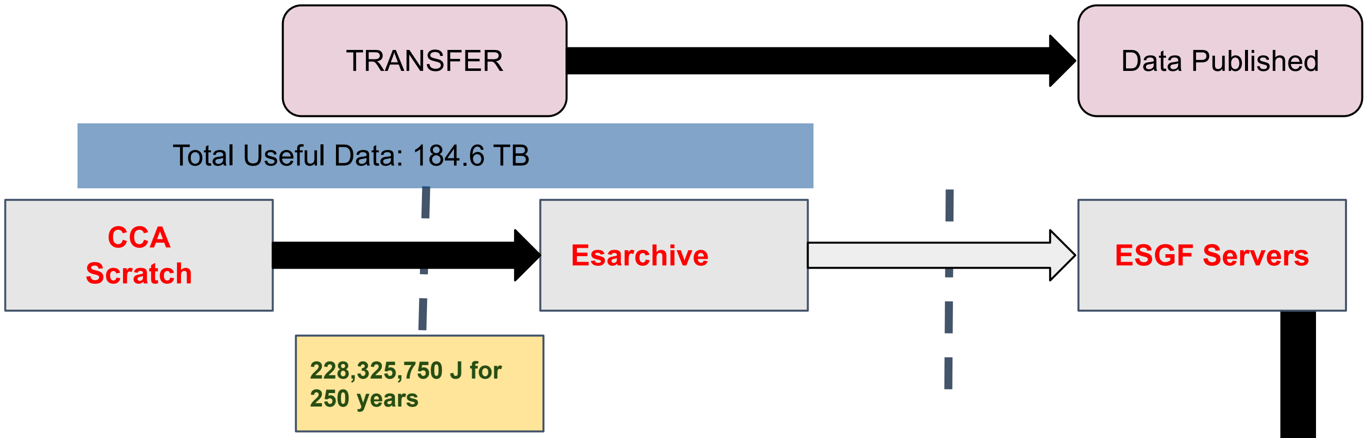

Transfer process and data availability (10 years) represent a **5.13%** of additional consumption compared to the production of the data.

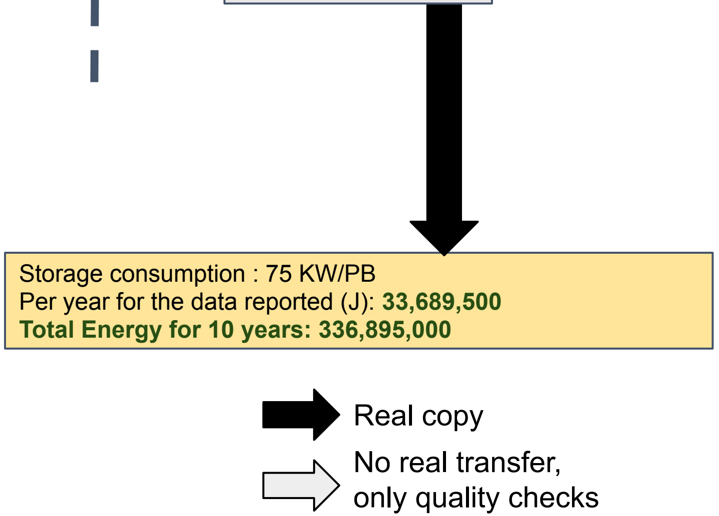

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

CMIP6 analysis: Data Transfer

TRANSFER using DT machine → Data Published

Total Useful Data: 184.6 TB

Marenostrum4 Scratch → **Esarchive** → **ESGF Servers**

Consumption transfer per hour (W) : 13,75
Total number of hours transfer: 3094
Total Energy: 42,544 J

Storage consumption : 75 KW/PB
Per year for the data reported (J): **33,689,500**
Total Energy for 10 years: 336,895,000

Transfer process and data availability (10 years) represent a **3.05%** of additional consumption compared to production on the data.
Transfer process is reduced a 99%.

Real copy using DT machine

No real transfer, only quality checks

From awareness to action?

- Carbon footprint: what's next?
  - Updating annually, compare pandemic and non-pandemic years
  - Incorporating CO_2 accounting to management tools, ↓ uncertainty
- Taking action, individual and collective. What if we…
  - Use the analysis to introduce optimizations in the computing part
    - ➢ A standard as CPMIP metrics and a coordinated collection allow us to work in the reduction of the most expensive parts of the workflow.

Lessons learned and future

Problem

Just requesting the metrics is insufficient, not all institutions will comply

Reported metrics may be inaccurate or even inconsistent

Some metrics (as energy) sometimes difficult to obtain and not representative

Actions

Improve the coordination:
- Follow Up meetings
- Normalise metric collection
- Supervise reported values
- Coordinated collection for CMIP7 (WIP-TF)

Investigating more reliable options to obtain the energy cost

Outcome

Compare the computational performance in a **multi-model multi-platform in the CMIP** context

Measure the **real** computational and energy cost

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Lessons learned and future

Problem

Some metrics are machine or configuration dependent

High variability between different models and platforms used in CMIP6

Actions

- Create a common and automatic approach using workflow managers and a common benchmark
- Multi-model comparisons will be better-grounded once more data is available

Outcome

Compare the computational performance of the common HPCW benchmark in different platforms

Comparison of representative ESMs benchmarks including models around the world

Workflow manager solutions

Automatic scripts for queue systems (IFS, NEMO, ICON, FESOM)

Solutions for common community tools as MultIO, XIOS, YAC or OASIS

Lessons learned and future

Problem

Data management is not covered

Some metrics are not representative enough and mixing different problems

Actions

- ASYPD should differ between queue and initialization/post-processing time/data movement time
- Cpl. C and Output C. due to explain inherent cost and load imbalance.
- Data evaluation: transfer, storage, availability

Outcome

Finer representation to understand better the efficiency of the models and the workflow

New metrics to analyze the efficiency of the complete workflow as a whole

Proposed collection

Tier	Level	Description	Metrics
01	Institution	Overall consumption and platform information	• Number of simulated years • Core-hours consumed • Data output • Energy cost (Carbon footprint) • HPC platform
02	Experiment	Metrics by experiment that require close to no effort to obtain	• Simulated years per day (SYPD) • Actual SYPD • Core-hours per simulated year (CHSY) • Joules per simulated year • Parallelisation • Data Intensity • Resolution
03	Experiment	Detailed metrics, not required but collected from institutions willing to collaborate (ENES)	• Data output cost • Coupling Cost • Memory Bloat • Complexity

A common approach using automatic tools

Autosubmit (AS) as a common workflow manager

Directly obtained from AS

- SYPD
- ASYPD
- CHSY
- Paral

AS Databases

AS & Slurm & filesystem

- JPSY
- Memory bloat
- Data Intensity

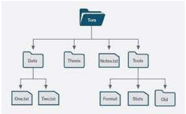

Manually collected

- Resolution
- Complexity
- Platform
- Coupling Cost
- Data Output cost

Actions for CMIP7

Supervision during the collection

- Ensure that the all institutions report the same for each metric
- Provide support, a roadmap, a collection approach and deadlines
- Encourage institutions to conduct the collection during the initial phase (e.g., spin-up), as this may lead to increased throughput in subsequent runs

Publish the collection and **analysis** openly and accessible to the entire community:

- ES-DOC
- Common **spreadsheet** template shared among all institutions for recording results (useful for milestones and monitoring)
- Visualization on ESGF Data Statistics

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Reducing Carbon Footprint in Climate Modeling: Unveiling the Efficiency and Energy Costs of Our Models

Mario Acosta and Sergi Palomas

mario.acosta@bsc.es
sergi.palomas@bsc.es