

Answer from authors to Reviewer 1 comments

Reducing Time and Computing Costs in EC-Earth: An Automatic Load-Balancing Approach for Coupled ESMs

by

Sergi Palomas, Mario C. Acosta, Gladys Utrera, and Etienne Tourigny

Reviewer (R)

[Authors \(A\)](#)

R:

The paper addresses a very significant issue of load imbalances on (large) parallel runs of coupled ESM models. The proposed solution and the developed tool represent a meaningful contribution to alleviating the waste of computational resources and giving the users of the ESM models a better control over the coupled simulation runs and their overhead.

The current version of the paper is quite polished and, with an exception of a handful of mostly technical errors, nearly ready for publication.

A:

[We sincerely appreciate the reviewer's insightful comments, which have helped improve the quality of our manuscript. Below, we provide detailed responses to each point raised.](#)

General comments

R:

A suggestion that the authors might want to briefly **discuss in the outlook:**** How difficult would be an extension of the proposed methodology to heterogeneous/hybrid architectures (e.g. CPU-GPU systems)?**

A:

[We agree that this should be discussed in the manuscript. We have added it in the outlook \(lines](#)

422-437). Now it reads:

““““

At the same time, the increasing adoption of GPU acceleration in Earth System modelling software reflects a broader shift towards hybrid computing infrastructures. A good example of this trend can be found in the new EuroHPC systems, where 7 out of the 8 integrate both CPU and GPU resources. Consequently, methodologies for load-balancing must evolve to account for these new hybrid architectures.

While the principles described for the auto-lb approach remain relevant, heterogeneous CPU-GPU codes introduce additional complexities.

The primary challenge lies in controlling the speed at which each component has to run to keep the load-balance. In a pure MPI setup, resource redistribution is straightforward, as coupled components share a common pool of processing elements (PEs, physical cores) and can reallocate them while keeping the total amount of parallel resources used constant.

In contrast, for components running on different hardware (e.g., CPUs and GPUs), the term "processing element" has different meanings, and resources are not directly interchangeable - a CPU core and a GPU core do not have a one-to-one equivalence.

The authors believe that the overall methodology described for auto-lb could be extended to hybrid CPU-GPU ESMs, provided that a standardised definition of the computational resources is established. This would allow the optimisation process to account for the equivalences and differences between CPUs and GPUs, potentially through an application-specific equivalent compute unit metric. Such a metric would involve profiling the performance characteristics of each component on both types of hardware to guide resource allocation decisions.

””””

Specific comments:

R:

I suggest to include a column with parallel efficiency in the Table 1.

A:

We have updated Table 1 to include the requested column for parallel efficiency. Additionally, we have added some missing information to the Table's caption.

Technical comments:

R:

'can not' should be 'cannot' at several places in text

A:

This has been corrected in all instances within the manuscript. (lines 152 and 176)

R:

lines 178-181: since a single node is taken as the baseline, processors and processes should be replaced with nodes

A:

We believe that comparing against a single node or the processors available within that node is conceptually equivalent. Since the granularity used in our tables, figures, and results is based on processors, we prefer to maintain the original terminology for consistency.

R:

line 201: remove the multiplication dot in the denominator

A:

This has been corrected in the manuscript.

R:

lines 313-314: it seems that Figures 5a and 5b should be referenced there instead of Figures 10a and 10b

A:

We confirm that there was an error in the references. This has been corrected in the appropriate section 5.1 (lines 313-314).

R:

line 332: 'worse the original...' should read 'worse than the original...'

A:

This has been fixed in the manuscript

R:

Section 5.3: Figure 10 should probably be referenced somewhere within this section

A:

As noted in a previous comment, there were inconsistencies in figure references in section 5.3. We have reviewed and corrected these references accordingly, as well as some minor improvements to the whole paragraph.