

Interactive comment on “CPMIP: Measurements of Real Computational Performance of Earth System Models” by V. Balaji et al.

C. Levy (Referee)

claire.levy@locean-ipsl.upmc.fr

Received and published: 5 October 2016

General comments I believe this paper is a major breakthrough for ESMs and climate studies: it introduces a set of metrics universally available for ESMs in the play for CMIP6. Using this suggestions for CMIP6 will be of great help for the future. All the required expertise is shown, allowing the authors to propose a set of very relevant metrics. The paper is nicely written and easy to read. It should be published and publicised in time (= as soon as possible now) to allow the whole community to set up the proposed metrics for the CMIP6 core experiments. It could be followed in the next years by some updates on those useful metrics (from adding something on temporal resolution (see below), to emerging ideas needing further discussions (single precision...))

Specific comments:

C1

Page 2 line 20: "This is because different computing architectures are less or more suitable to increasing problem size along any of these axes." My reading did not allow me to understand this sentence concluding the paragraph. Maybe a little more explanation would help (me, at least)

Page 3 line 26: "One issue was the limitations imposed by memory." I suggest replacing memory by memory bandwidth (e.g. not memory size)

Page 4 and 5: I'm not fully convinced by the Figure 1: it seems incomplete (and difficult to complete!). One could add a number of boxes (OcnDyn, OcnPhy, IceDyn, IceBio, etc...) so as a number of couplings, interactions and related processes. Would an "incomplete list" be better, or is there a better way to build a picture to make this complex ESM architecture visible?

Page 5 line 8: "but a model at fixed resolution is capped in terms of time to solution, absent advances in hardware or algorithm." I agree if you are still talking about one dycore in this paragraph, but this is not true for a whole ESM.

Page 6 Figure 2: Not clear to me what is the conclusion of paragraph ending line 14, nor precisely what is demonstrated on Figure 2.

Page 8: I believe the answers to the questions listed at top of page depend on how the computer platform is used: in dedicated mode, or not? But this seems to be taken in account in the ASYPD metric?

Page 11: taking the resolution in account Here, you have only taken the spatial resolution in account. I wonder if somehow (not simple though), temporal resolution should be taken in account in the G metric: using a large or a small time step does indeed widely change the number of floating point operation you need for one simulated year. Now, what time step to use (those are different in atmosphere, ocean, coupler...)? I do not have any simple answer, but to be able to compare metrics between ESMs, I suggest to add something in G to take in account number of time steps per year as an

C2

unavoidable constraint on the number of floating point operations. (Could also be in a next generation of those metrics)

Page 19: Reference for XIOS: is this paper from Jousaume et al. the appropriate reference for XIOS today?

Interactive comment on Geosci. Model Dev. Discuss., doi:10.5194/gmd-2016-197, 2016.