

Interactive comment on “Optimizing High-Resolution Community Earth System Model on a Heterogeneous Many-Core Supercomputing Platform (CESM-HR_sw1.0)” by Shaoqing Zhang et al.

Shaoqing Zhang et al.

lingan@tsinghua.edu.cn

Received and published: 7 June 2020

Interactive comment on “Optimizing High-Resolution Community Earth System Model on a Heterogeneous Many-Core Supercomputing Platform (CESM-HR_sw1.0)” by Shaoqing Zhang et al. Carlos Osuna (Referee) carlos.osuna@meteoswiss.ch Received and published: 11 May 2020

This paper provides a comprehensive description of the efforts to port large legacy code of the CESM model to the Sunway TaihuLight processors. The text gives a very

C1

detailed description of the improvements made to the model, the parallelization and optimization techniques employed as well as the programming models used (OpenACC and Athread). A pre-industrial simulation over 400 years is being performed and the main result is an optimization from 1 SYPD to 3.4 SYPD. The paper is well structured and provides comprehensive information about the model, and the experiments for reproducibility. The software is open source and referenced within the paper.

General comments: the text is sometimes too dense and hard to follow. (1) I would recommend to interleave the relevant tables and figures within the text, next to the references and the discussion. Otherwise it is hard to follow references to tables and figures that are placed at the end. (2) For a paper running on such a large system (65000CGs), a scalability plot is missing. (3) The paper emphasizes in several places the energy consumption point of view and the advantage of hybrid architectures like TaihuLight. However, there is no real data for this experiment presented, therefore there is no data that support some strong statements presented in the text.

RE: Thanks for the reviewer’s thorough examination of our manuscript (MS) and positive comments. We all agree that your comments are very constructive for us to improve presentation of the MS, and all your major comments and other points have been fully addressed in the revision. Specifically, in the revision, (1) the Tables and Figs. are inserted into the relevant places; (2) a scalability plot is added; (3) the energy consumption statement is appropriately re-written and more discussions about the uncertainties of the current work on power efficiency are added.

The point-by-point replies are followed.

Line105 (89 in revised version): the argument is valid only for applications that maximize the FLOPs provided by a computer, which does not hold for weather and climate applications. Same happens with Table 2 which presents general specs measured for different machines. If 3.4 SYPD are achieved with 65000, while the benchmarks with 11000 Intel processors runs at 1 SYPD I conclude that for the same SYPD, the Taihu-

C2

Light requires 2 times more processors than the Intel system. How does the energy efficiency of the TaihuLight chip compares to the Intel? I encourage to backup the energy arguments with data from the experiments or make a clearer link. The authors (very rightly) emphasize the large efforts required to port this large model to a new architecture. Here the reviewer is missing a more general discussion about the cost of these effort and performance portability. Was performance portability an important metric of this work? Considering that the TaihuLight is a unique system that did not go into the market of supercomputing, what is the cost of such refactoring? It would be interesting for the paper to provide a number for financial costs of the porting effort.

RE: We all agree that the power efficiency is currently an issue with rich uncertainties. The statement of power efficiency is modified, and more discussions on the uncertainties and future work direction are added in the revision. Please see L225-230. The discussions on the performance-portability issue are added in the revision. Please see L128-131; L370-374; L375; L387-388; L532-533; L540-543. Thanks a lot!

More specific comments:

The introduction is too long. It can be simplified and emphasize the contributions of this work. The related work part of the introduction is well written.

RE: The introduction is condensed, focusing on relevant project now. Please see the new introduction. Thanks.

line72: I don't think there is any major system in the list of supercomputers with FPGAs.

RE: The related statement is removed. Thanks.

line: 158-159 This is not very correct. The authors already mentioned in the literature NIM that was ported to GPU and XeonPhi. Other models like ICON are GPU ready.

RE: The sentence has been removed in the revision. Thanks.

line 265 (218): This is a very strong statement. It is true that the systems are very different. But from these differences is not obvious that a GPU system with multiple

C3

GPUs connected with a high-throughput low latency NVLink is not more suitable for scientific computations.

RE: The statement has been re-written in a more appropriate way. Please see L218-221.

Section 2.2.2: as already mentioned, this refer to general specifications of the system which are for sure different than the FLOP/energy consumption of weather applications. I would consider providing experiment data or removing this section.

RE: Agree that energy consumption is a complex issue and at the current stage, the Sunway machine has not shown any advantage on power efficiency. We may pursue really-greener utilization in the future when more experiences and optimization skills can make the computation with much higher efficiency. Please see L226-230.

Section 3.1: It is hard to follow the text and what are the major improvement. Considering supporting with data, figures and simplifying the list.

RE: The major improvement of CESM1.3-beta17_sehires38 has been reorganized. Please see the new Section 3.1. Thanks.

Interactive comment on Geosci. Model Dev. Discuss., <https://doi.org/10.5194/gmd-2020-18>, 2020.

C4