

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 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

Printer-friendly version

Discussion paper



reproducibility. The software is open source and referenced within the paper.

General comments: the text is sometimes too dense and hard to follow. 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.

For a paper running on such a large system (65000CGs) , a scalability plot is missing.

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.

Line 105: 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 TaihuLight 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 more clear 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 didnt 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.

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.

line 72: I dont think there is any major system in the list of supercomputers with FPGAs.

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 .

line 265: 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 GPUs connected with a high-throughput low latency NVLink is not more suitable for scientific computations.

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.

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.

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

Printer-friendly version

Discussion paper

