Geosci. Model Dev. Discuss.,
https://doi.org/10.5194/gmd-2020-18-AC1, 2020

Geoscientific Model Development Discussions

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. Mark Govett (Referee) mark.w.govett@noaa.gov Received and published: 25 March 2020

General Comments: This was a very well organized and written paper. The paper described efforts to port a large legacy, climate code to the Sunway TaihuLight sys-

tem. The unique architecture of the Sunway processor was described which helps the reader understand some of the changes that were needed for the climate application to run efficiently. The original code was designed and run on Intel-based processors. Performance and baseline scientific results were made prior to porting the code. The work described detailed efforts to optimize the code, while maintaining sufficient accuracy in the solutions. As the authors admit, determining an acceptable level of scientific accuracy is an ongoing process determined by many factors. With a complex scientific application, thorough testing and evaluation using multiple criteria is needed to build confidence in the solution.

RE: Thank you for the thorough examination on our manuscript (MS) and constructive comments. We agree that all of the comments are very useful for us to improve the presentation of the MS, and we have fully addressed them in the revision.

What follows is a point-by-point reply to each comment.

Specific Comments:

(1) The authors describe extensive efforts to optimize the code, which included many common techniques. You spent a lot of time optimizing for the Sunway processor, but did not apply that level of effort to the original code. Some of the changes could have been applied to original code, making the comparison of performance more fairly represented.

RE: This is a good point that needs to be addressed in the MS. The work presented in this paper represents the first step of running CESM on the new architecture machine. To minimize the coding uncertainties, we keep the original CPU code unchanged for both accuracy verification and performance evaluation. It is true that many of the applied techniques can be generalized to CPU or even GPU architectures. That would be the focus of a later-stage study, which is to extract both general practices and tools that would help the transition from current multi-core CPUs to many-core accelerators. Redesigning some of the original algorithms is expected to further improve the efficiency.

Some clarifications and discussions are added in the revision. Please see L43-44; L125-126; L432-433; L532-533. Thanks a lot!

(2) The work appears specifically designed to target a single system with a unique processor. Was performance portability considered as a factor in this effort? Could the modified code run on an Intel-based system and how did the results compare to the original code. Given the fine-grained nature of the parallelization, would GPUs or high core count CPU processors be a target for this work? Addressing performance-portability would make the impact of this work much greater than the results you achieved targeting a single system.

RE: We agree that addressing the issue of general performance-portability is an important and interesting topic in using heterogeneous many-core HPC systems and should be further clarified and discussed more in the MS. In the revision we expanded the discussion of this topic. Please see L128-131; L370-373; L375; L387-388; L532-533; L540-543. Thanks a lot!

(3) The performance impact of the different types of code optimizations you made were not described. This would be a useful way to determine the tradeoff between portability and performance. For example, specific optimizations described in the Stage 1-3 optimizations were closely aligned to the Sunway processor. How much performance benefit were there for each of the stages and was it applied to a large portion of the code?

RE: Combined with comment (2), this is a very good comment addressing the general performance-portability issue. The description of performance impact of the different types of code optimizations is now added. Please see L370-373; L375; L387-388. The three-stage optimizations described in section are proposed to solve a specific prefix summation calculation using register communication. Our method is 27 times faster than the original code that has only a single loop. More discussions about its performance portability are added in the revision. Please see L370-373.

(4) Regarding porting the code to the TaihuLight system, your approach seemed to be first port the code to two Intel based systems (TAMU and QNLM). It was unclear why you felt the need to port to both systems. Further, your comparisons were made after only 9 timesteps seemed arbitrary and perhaps no sufficient. Please include some justification. It was also unclear what field were compared in the UF-CAM-ECT test. A summary of the relevant details from the paper would be useful here.

RE: Given the totally new architecture of the Sunway machine, to minimize the uncertainties of code porting, before we port the CESM-HR to the Sunway TaihuLight, we first ran and tested its correctness on the Intel multi-core supercomputing platforms available to iHESP, serving as the first benchmark. The justification and adjusted statement are made in the revision. Please see L284-287.

More descriptions and discussions about UF-CAM-ECT tools are added in the revision. Please see L293-295; L461-464. Thanks.

(5) Are there references for the tools given in Table 1? Most of the tools listed were not referenced in the manuscript. They should be either introduced to the reader in the paper if they add value to the manuscript. For example, you could state how you used them and how it helped identify

RE: The point is well taken! One more column describing the role of each tool in this project is now added in the Table 1. More reference information about Table 1 is added in the revision. Please see the new Table 1 and L193-194. Thanks.

Technical Corrections: Line 103-105: Awkward sentence. Perhaps break into two sentences?

RE: The entire sentence is removed in the revision. Thanks.

Line140 (104 in revised version): Old reference (Govett, 2010) should be replaced with a more comprehensive paper (Govett, et al. 2015) in the Bulletin of the AMS: https://doi.org/10.1175/BAMSD-15-00278.1

RE: Done. Please see L104; L615-617. Thanks.

Lines 212-218 (170-175): The authors don't describe what the speed and bandwidth of connections between super-nodes, within a cabinet, and between cabinets. This is essential in understanding the limitations of the Sunway TaihuLight system at scale.

RE: Good comment! The TaihuLight compute nodes are connected via a 2-level InfiniBand network. A single-switch with full bisection bandwidth connects all 256 nodes within a super-node, while a fat-tree with 1/4 of the full bisection bandwidth connects all super-nodes (as shown the attached Figure 1). Table II (as shown the attached Figure 2) shows measurements of bisection communication bandwidth at different levels of the system.

Such essential information is added in the revision. Please see L175-178. Thanks.

Line 243 (198): change "details Section 3.3" to: "detail in Section 3.3" Line 245-250 (200-205): unclear if CPE based parallelism is with MPI or something else??? Line 251 (207): change "details" to "detail"

RE: Line 198: Done; Line 207: Done. Line 245-250: The statements are revised for clarification. Please see L204-207.

Lines 260-266 (215-221): It appears that an MPI-based intelligent programming model would work here. The intelligence would be knowing when comes within a CPU task group or to an MPE are needed.

RE: Thanks for reminding the clarification. The statement is further clarified based on the comment. Please see L218-221.

Lines 268-272 (223-230): Regarding the power efĩňĄciency comparison in Table 2, is this a fair comparison? It seems like there would be size benefits favoring the larger systems especially in terms of infrastructure required no matter the size of the system.

RE: This is a good point! We all agree that the power efficiency is currently an issue

with rich uncertainties. More discussions on the uncertainties and future work direction are added in the revision. Please see L225-230. Thanks

Line 307 (257): swlu was italicized but not introduced in section 2.1.2. It appears in capitals as SWLU in Table 1.

RE: The tool name in Table 1 is modified to italicized, consistent with the context now. Thanks.

Line 332 (279): TAMU and QNLM are undefined (QNLM is defined on line 602, TASMU on line 603)

RE: In the revision, TAMU and QNLM are first defined at L116-117. Thanks.

Line 341 (288): It seems that such short runs are not sufficient. Did you make a similar test with more than nine time steps. There remains a high potential for variations to show up later in the simulation experiments Optimizations were made to achieve 1 SYPD on the Sunway system. Did you attempt to incorporate these changes and optimization techniques back into the original model?

RE: We also compare the results of 1-yr long runs in different perturbation scenarios (as shown in Fig. 10) to comprehend the integrations on Sunway machine. More discussions on the 9 timesteps are added (please see L460-463). To minimize the coding-caused uncertainties in the porting and optimizing process, at this stage, the current work doesn't apply any change to the original code. More statement and discussions are added in the revision. Please see L125-126. Thanks.

Line 372 (316): Can you provide more details or analysis regarding why -O2 fails but -O3 passes? How long were the runs made before comparisons were done?

RE: While more discussions on the metrics of the ECT tool are added in L460-463, the possible reason for which -O2 fails but -O3 passes is given in the revision. Please see L318-319; L460-463.

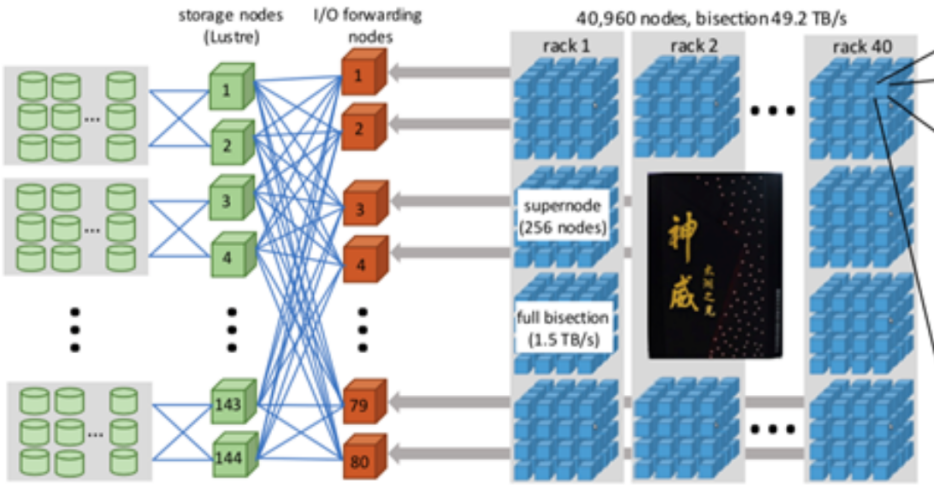

Fig. 1. TaihuLight Network System

TABLE II

TAIHULIGHT SPECIFICATIONS AND BANDWIDTH MEASURED

Component	Configurations	BW measured (% peak)
MPE	1.45 GHz, 32/256KB L1/L2	DRAM 8.0 GB/s (80%)
CPE	1.45 GHz, 64KB SPM	Reg. 630 GB/s (85%)
CG	1 MPE + 64 CPEs	DRAM 28.98 GB/s (85%)
Node	1 CPU (4 CGs), 4×8GB RAM	Net: 6.04 GB/s (89%)
Super Node	256 nodes, FDR 56 Gbps IB	Bisection 1.5 TB/s (72%)
TaihuLight	160 supernodes	Bisection 49.2 TB/s (68.6%)
Agg. memory	1.3 PB	4.6 PB/s (85%)
Storage	5.2 PB (Online2)	70 GB/s (54.35%)

Fig. 2. TaihuLight Specifications and Bandwidth Measured