

Interactive comment on “The GPU version of LICOM3 under HIP framework and its large-scale application” by Pengfei Wang et al.

Anonymous Referee #2

Received and published: 15 February 2021

Numerical climate modeling is a key method for scientists and researchers to better understand our planet, and one of the most popular applications that greatly challenges the most state-of-the-art high performance computing (HPC) systems. In this work, LICOM3, a standard ocean model is selected and scaled onto the GPU-based heterogeneous supercomputing system. The authors have done lots of porting and optimizing work to put almost all of the time-consuming computation processes into the GPU side, and greatly reduce the communication overhead. As a result, both the dynamic core and the physics part are ported and parallelized on GPUs. A speedup of 42x is achieved when using 284 AMD GPUs VS 384 CPU cores. Excellent scalability is also achieved. A test of 1/20 degree LICOM3-HIP is reached using 6550 nodes and 26200 GPUs, 2.72 SYPD in time-to-solution.

Printer-friendly version

Discussion paper



As a computer scientist who also focuses on porting and tuning climate models onto different HPC platforms, dealing with a complete model with lots of code legacies using a new accelerator is obviously not an easy work. Sometimes rewriting and redesigning are necessary to obtain a satisfactory performance. In this work, the optimizing techniques provided are sound and solid, and can be used as a good guidance for corresponding work. AMD GPU and HIP, though not as popular as Nvidia GPU and CUDA for now, are still very promising GPU accelerators for current generation supercomputers. Moreover, it is likely that some of the forthcoming Exa-scale supercomputers, will also be adopting AMD GPUs. So this work is also a good trials ahead of time. More specifically to the strategies: only the most time-consuming parts (seven subroutines) are translated into HIP C, deeply re-coded, ported onto the GPUs, and fully optimized (such as the usage of temporary arrays to avoid data dependency, the change of data structure of original Fortran arrays, etc.). Halos that contain partial communications are handled by CPU part. Therefore, a hybrid computing model is performed, to further improve the overall performance. This is also a very popular strategy when dealing with numerical problems with inter-node or inter-process communications. Besides, The IO part is also considered and tuned by rewriting the data reading strategies and doing parallel scattering.

Overall, the paper is well-structured, with sufficient figures and tables to help better illustrate the ideas where necessary. But there are grammar errors and misleading descriptions here and there. So I suggest the authors ask help from native speakers for further proofreading.

Here are some other suggestions,

The authors mentioned the dynamic and physics parts, but lacks further explanations to what they are. I understand that most communications exist within the dynamic part, but could the authors be more specific in pointing out the optimizing strategies for the dynamic part and the physics part, respectively? Are there any differences?

[Printer-friendly version](#)[Discussion paper](#)

In the end, the authors claimed that the 1/20° LICOM3-HIP version can not only reproduce the observations, but also produce much smaller scale activities, such as sub-mesoscale eddies and frontal scales structures. Could the authors explain how they obtain the observation version?

Porting a complete model is not an easy work. In this work, approximately 12000 lines Fortran code were rewrote from fortran to C. Could the authors estimate the cost? For example, the number and time cost of persons in the whole project.

The following work is suggested to be cited and comment as well, to enrich the related work part. Optimizing high-resolution Community Earth System Model on a heterogeneous many-core supercomputing platform.

Line 45, I suggest to update the the TOP 500 list using the latest one (Nov. 2020);

Line 49, I don't think the energy result is provided in the work of Xu et al. (2015). Please double check.

Line 71, and the conclusions is in Section 6 → and the conclusions are in Section 6

Line 74: which started to develop → which has been developed

Line 83: That makes the coupler is suitable to apply to high resolution modelling. → It makes the coupler suitable to be applied to high resolution modeling.

Line 85: improve → improves

Line 103: remove totally

Part 2.3: add some citations or links with more detailed introductions to the supercomputer used in this work.

Line 140: place → replacing

Line 143: Some... the others... → Some... some...

Lots of professional words are used in this article, such as theses syntax and macros

[Printer-friendly version](#)[Discussion paper](#)

used in HIP or CUDA. Please use a different syntax (e.g. *italic*) for these professional word, to help better identify them. For example, I suspect that 'for tracer', 'baroclinic' and 'barotropic' may refer to professional processes or subroutines in LICOM3, and 'including', 'cuda_', 'hip_', may refer to professional designations of CUDA or HIP framework.

Line 244, an → a

Figure 5, I suspect the IO time is the result after the IO optimization (part 3.4) being applied. Is that right?

Line 263, times → time

Please provide more details about the hardware configurations, e.g., the version of CPU and GPUs, the version of compilers, OS, etc.

Please replace Flops/s with Flops.

Interactive comment on Geosci. Model Dev. Discuss., <https://doi.org/10.5194/gmd-2020-323>, 2020.

Printer-friendly version

Discussion paper

