

Interactive comment on “A distributed computing approach to improve the performance of the Parallel Ocean Program (v2.1)” by B. van Werkhoven et al.

Anonymous Referee #1

Received and published: 9 December 2013

general comments:

In this paper two methods for improving the performance of the Parallel Ocean Program are presented and evaluated. The first method consists of the application of a hierarchical partitioning scheme to improve the domain decomposition of the model. The second method uses GPUs to improve the performance of some parts of the model.

In addition to tests on a single cluster, configuration are run, for which the model processes of a run are distributed across computing nodes of two different clusters, which are connected via a relatively slow network connection. Due to the hierarchical nature of the partitioning scheme, it allows the reduction of communication between different

C1632

sets of processes. This is especially beneficial in this kind of configuration.

The paper describes well structured and in high detail the work that has been done. Improving model performance is always an important topic. Doing this by changing the load balancing scheme and by using accelerator hardware are common approaches[1]. Therefore, the results of the paper present no novelty to the general modelling community. However, the paper shows that these approaches can be successfully applied to the Parallel Ocean Program.

The description of how the accelerator hardware was utilised is well written and can be a good guideline for similar attempts.

special comments:

1. Some reference to or comparison with similar work done in the area of load balancing and GPGPU in climate modelling would be nice.
2. Regarding the hierarchical partitioning scheme see comment by Ilja Honkonen
3. Regarding Fig. 8 a. Could you add the measurements for the CPU? b. Do these measurements include the time spend in functions called by these functions?
4. Regarding Fig. 9 a. Could you add measurements for the original POP version?
5. Could you document how many sample measurements you did for your performance results.
6. A figure showing the differences for your three versions (Explicit, Implicit and Stream) might help understanding the methods.
7. For someone interested in porting other models to CUDA it might be interesting to hear about the effort required for the different steps.
8. Comparing the performance of the model for the “CPU only”-run and the “CPU+GPU”-run is not really fair. One could for example take the hardware costs

C1633

and/or power consumption into account when interpreting the performance results.

[1]: J. Michalakes and M. Vachharajani, "GPU acceleration of numerical weather prediction." in IPDPS. IEEE, 2008, pp. 1-7.

Interactive comment on Geosci. Model Dev. Discuss., 6, 4705, 2013.