Geoscientific
Model Development
Discussions

[GMDD]

Interactive
comment

Interactive comment on "Scalability and some optimization of the Finite-volumE Sea ice-Ocean Model, Version 2.0 (FESOM2)" *by* Nikolay V. Koldunov et al.

Leonidas Linardakis (Referee)

leonidas.linardakis@mpimet.mpg.de

Received and published: 28 April 2019

This paper falls under the category of "model description". The performance on 2D unstructured grids has been known to be on par with the structured ones in the small community of unstructured grid modeling, but has not been publiced in a comprehensive way for GC models. This paper serves the purpose of discussing and communicating this knowledge, demonstrating that an operational unstructured-grid ocean model is on par, performance-wise, with structured-grid ones. I have major concerns though about the protocol followed in the design of the scaling experiments, and the numbers obtained from them, so I recommend the paper to be published after a major revision.

Printer-friendly version

Discussion paper

1. The authors present two sets of scaling experiments: an "idealized" set and a "realistic" one, and they discuss their differences at page 17, lines 7-14. In the idealized setup the SSH solver is four times faster than the realistic one (STORM experiment on 13824 cores), which renders the related scaling numbers and analysis of the idealized setup at the very least questionable. On the other hand, the realistic setup is presented under the title "Testing throughput in an operational configuration", implying that this is an operational setup. The original operational STORM setup though is very different than the one described by the authors, with a large number of specific diagnostics and high frequency output, rendering the term "operational" questionable. I propose a single set of realistic experiments, that demonstrate the model's scaling characteristics in realistic conditions, but no "operational" ones. As the authors point out, the cost of I/O and diagnostics in operational setups varies widely with the requirements, and does not reveal the performance of the core model, which is the interest in the scaling analysis of the model. In the model scaling experiments I/O and additional diagnostics are not needed. The experiments should be restarted from a one year run; running for the duration of 15days – 1month should be sufficient.

2. The information for the model and experiment setups is scattered, and I found it difficult to piece it together. A table would be useful, that describes the parametrizations used, type of of vertical coordinates, solver tolerance, timestep, run duration, etc.

3. A very interesting and important point is raised on page 12, lines 25 – 30: how to tune the solver so it provides the required quality at a minimal cost. It would be interesting to know how the authors define "robust results" (line 25), what is the analysis and the metrics (max error, biases, conservation) that they used for this tuning.

4. A minor comment: In Fig 10, why the number of the solver iterations depend on the number of cores?

5. I recommend to cite: A. E. MacDonald et al. "A general method for modeling on irregular grids", 2010, doi:10.1177/1094342010385019

Interactive comment on Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2018-334, 2019.