

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.

Marshall L. Ward

marshall.ward@gmail.com

Received and published: 29 April 2019

Hi, I am the author of the Ward (2016) presentation cited in the Table 3. Overall the results in this paper are very nice, and am looking forward to learning more about the model. My comments are only restricted to the timings used in Table 3.

My main comment is that the numbers which I presented at ECMWF were attempting to assess the scalability of existing configurations of NEMO, MOM5, and MOM6 in use on our system (Raijin at NCI Australia), and were not necessarily representative measures of performance of the models.

Printer-friendly version

Discussion paper



For example, the NEMO configuration was quite old even at the time (v3.4), and two attendees pointed out potential implementation issues of the tripole that were resolved in v3.6.

Additionally, I have my own internal numbers showing that MOM5 0.25 can scale much higher than 25 yr/day. For example, I have shown results showing performance of sea-ice (SIS) bound timings of 35 yr/day, and can increase this to nearly 60 yr/day, when sea ice is running concurrently to MOM5 and we allow high inefficiency.

(This is my usual "reference", a presentation at AMOS in 2016 where the ocean model was individually profiled:

<https://marshallward.org/talks/amos2016.html#/11/>

though again it's just a presentation and hasn't necessarily been stress tested.)

I know how difficult it can be to find cited performance numbers for ocean models, and I have probably not published enough of my own numbers, but feel a bit uncomfortable being associated with these particular numbers and wonder if they could be replaced or revised in some way.

I do have one published paper of MOM performance numbers which may be suitable:

* https://link.springer.com/chapter/10.1007/978-3-319-15994-2_55

And Balaji et al's CPMIP paper may provide some useful comparison numbers:

* <https://doi.org/10.5194/gmd-10-19-2017>

Or, if you do not want to change the reference, maybe just disclose somewhere that the numbers are not necessarily reflective of peak performance but "typical" performance, in which case I'd request that you use the lower ~ 1000 CPU timings.

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

Printer-friendly version

Discussion paper

