

Interactive comment on “Parallel computing efficiency of SWAN” by Christo Rautenbach et al.

Anonymous Referee #1

Received and published: 4 January 2021

General comments

The manuscript highlights the parallel performance of the SWAN model, as a part of the operational weather forecasting in South Africa.

The manuscript is well-written and the authors set out their objectives and results very clearly. I find the results of the manuscript interesting which, in my opinion, it should be typified as a model evaluation paper rather than a development and technical paper. However, I am not sure whether the results are reproducible as they also depend on the used hardware. It might be good to highlight some additional info with respect to this aspect. (See also my comment no. 2 below.)

Specific comments

1) The notion "thread" is a bit confusing for the MPI adepts, which should be "core".

C1

May be a combi "thread/core" would be a better wording.

2) In general, within a single node (containing a number of threads/cores) OpenMP is more efficient than MPI. So, contrary to the study of Genseberger and Donner (2015, 2020), the results of your study contradict this general statement. Do you have an explanation for this? Perhaps, you may add some technical info concerning the used hardware with respect to this aspect (memory I/O, network, etc.)

3) One of the possible reason why a perfect speed-up cannot be obtained (see Fig. 1) is the domain partition of the computational grid and also the wet/dry (or active/inactive) points. The employed partition is the stripwise one which is because of the underlying parallel technique, namely the wavefront method. See Genseberger and Donner (2015) and also Zijlema (2005). The stripwise partition might not be the most optimal one with respect to the speed-up. In this specific case, it leads to a maximum of 6 threads/cores without too much sacrificing parallel efficiency.

It would be good to highlight this aspect.

Added reference: M. Zijlema. Parallelization of a nearshore wind wave model for distributed memory architectures. In *Parallel Computational Fluid Dynamics - Multidisciplinary applications*, pages 207-214. Elsevier Science, 2005.

Do you have active/non-active grid points in your model schematization? Can you comment on this?

Technical corrections, etc.

line 103: ration -> ratio

line 108: compliers -> compilers

line 156: (16 x 25 threads) -> (16 x 24 threads)

line 157: (16 x 24 threads) -> (64 x 24 threads)

C2

line 234: please change version number; also suggested to add the Technical Manual of SWAN besides the User Manual, as it contains the details of both physics and numerics

Interactive comment on Geosci. Model Dev. Discuss., <https://doi.org/10.5194/gmd-2020-314>, 2020.