



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

Christo Rautenbach et al.

rautenbachchristo@gmail.com

Received and published: 7 May 2021

Response: Thank you for your interest in our paper and for the discussion, which has helped to improve the manuscript. Herewith inline responses to the questions, comments, and corrections:

General Comments: The manuscript is nicely written with extensive reference to research articles and the authors clearly identifies 4 research questions which will be covered in the paper.

Scientific Comments: 1. However reading through the manuscript, missing the discussion over the scalability as shown in Figure 1c which would help about answering to the research question 4) What is the scalability of a rectangular grid, SWAN set-up?

Response: The scalability can be assessed in several ways. We use a variety of

metrics to help elucidate the concept. In general, the most dramatic increase in the time saving ratio occurs when the number of threads/ cores is below 6. Thereafter, increasing the number of cores doesn't substantially improve the time saving ratio. In general, the more thread/ cores one uses (for this SWAN configuration) the faster the total computation will be completed. This is evident from Fig. 2c (in the revised text). The point where the internal, domain-decomposed, communication time starts to dominate (and thus decrease run time with increase thread counts) is not reached within our thread counts (as explained on line 168). To be sure this is clear to the reader, a more detailed explanation has been added at the end of the discussion at line 176.

2. Further as mentioned on Line 169 quote "The scalability is presented via three performance metrics: the efficiency, speed-up ratio and the timesaving ratio" would like that authors touch upon all the scalability for all these metrics and not only speed-up ratio?

Response: Very good suggestion! This comment has been incorporated together with the previous comment at the end of the discussion section. Extra text has been added to discuss the performance in terms of these metrics.

3. Further, in line 66, it is mentioned quote "Here we build on the case study of Genseberger & Donners using results produced in the present study for southern Africa.." but to me it seems that the present study discussion is more and more comparing the results of current study to study of Genseberger & Donners - to me as mentioned earlier are of different domain. If this is not correct, please explain in Methodology and Background accordingly.

Response: We extend on the results of Genseberger & Donners rather than simply compare them. Specifically, we use different scalability/ benchmarking metrics than they did. We summarise their results because our studies build on their work (lines 66 -71). We make detailed comparisons between our study and theirs, to highlight the

new insights provided by our study.

4. To make this research article self standing – please include the case study domain of southern Africa figure here (instead of referring to the - model configuration can be found in (Rautenbach, et al., 2020 (a)) and (Rautenbach, et al., 2020 (b))). Also figure of case study of Genseberger & Donners can be included here to make understanding of the results and discussion clear to readers. Later , unless as pointed out in point 3 above.

Response: The South African SWAN model domain has been added as the new Figure 1. We could readily do this as this was our model. Due to copyright constraints I do not think its possible to add the model domain and extent of Genseberger and Donners. Due to the clear referencing, we recommend the readers to have a look at their publication, which is easily accessible.

5. Please include a table/figure in the Conclusion part to make conclusion more obvious and readable to the users. Refer “A hybrid SWAN version for fast and efficient practical wave modelling, Genseberger & Donners, (2020) paper section 4.2 to see what I mean by including a table to compare between OpenMP / MPI different metrics and/or with current study with the study of Genseberger & Donners.

Response: It was difficult to make a table that compares our results directly with Genseberger & Donners as they used multiple nodes and comparisons with their hybrid model. Our results were dependent on model resolution as well as number of threads/ cores (single node). We feel that a table could potentially confuse or obscure the key results. We note that Figure 1 (now Figure 2) now summarizes all the results, which, combined with the additional changes (added w.r.t the other comments above), hopefully improves the readability of the conclusions.

6. Can authors make the connection between Zafari, Larsson, & Tillenius, (2019) study of shallow water with the current study of SWAN Model clear. There is reference made to "gcc" - but current study "Methodoloy and Background" does not include details of

this current study being run on gcc except what is mentioned in lines 124-128. The reason for this comment is that seems that authors are hinting to gcc but no further references or discussion on this in later sections. Maybe I am missing something here?

Response: This paper reference was added for completeness to inform the reader of other research that might be related to this topic. This paragraph aimed to introduce the next paragraph discussing the model resolutions used in the present study. A sentence was added to the conclusion suggesting further investigations with regards to using a gcc compiler.

Technical Corrections: Line 55: SLOSH : Sea, Lake, and Overland Surges from Hurricanes. (though SLOSH can be NOAA official storm surge forecasting model - but this is not the official full name) Corrected Line 58 : Mexican golf : I think here the Gulf of Mexico is being referred. Corrected Line 103: ration should be changed to ratio. Correction Line 126: ggc should be changed to gcc. Corrected Line 156 : 16 nodes (16 × 25 threads) should be changed to 16 nodes (16 × 24 threads) Line 157 : 64 nodes (16 × 24 threads) should be changed to 64 nodes (64 × 24 threads) Corrected

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

Printer-friendly version

Discussion paper