Review of "A Scalability Study of the Ice-sheet and Sea-level System Model (ISSM, Version 4.18) by Y. Fischler, M. Ruckamp, C. Bischof, V. Aizinger, M. Morlinghem, A. Humbert

In the manuscript in question, the authors perform a detailed study of the overall scaling of the Ice-sheet and Sea-level System Model (ISSM) in the context of a Greenland ice sheet simulation and the higher-order (HO) Blatter-Pattyn model for the ice sheet velocities/momentum balance. The authors describe a low-overhead performance instrumentation using Score-P developed within this code base to enable continuous performance monitoring. The scalability study reveals that the matrix assembly part of the computation is the main bottleneck when it comes to scalability/performance, and should be examined further.

The manuscript in question is well-written, interesting and a good fit for GMD. My recommendation is publication following a minor revision. I ask that the authors please address the following questions/comments in their revision.

- The authors mention exascale readiness in the introduction, but there is no discussion in the paper of whether ISSM is portable to up-and-comping heterogenous architectures (GPUs). Is it? Can the present study be repeated on a set of GPUs? Some discussion of this is warranted.
- On line 77 of the introduction, the authors mention that they are using a GMRES linear solver preconditioned with a simple block Jacobi preconditioner. The HO Stokes equations are symmetric. Have the authors tried using Conjugate Gradient? I additionally worry that the Jacobi preconditioner is inadequate for problems with floating ice, e.g., Antarctica, as shown in the references by Tezaur et al. and additionally: (1) T. Isaac, G. Stadler, and O. Ghattas, Solution of nonlinear Stokes equations discretized by high-order finite elements on nonconforming and anisotropic meshes, with application to ice sheet dynamics, SIAM J. Sci. Comput., 37 (2015), pp. B804–B833, doi:10.1137/140974407 and (2) R. Tuminaro, M. Perego, I. Tezaur, A. Salinger, and S. Price. A Matrix Dependent/Algebraic Multigrid Approach for Extruded Meshes with Applications to Ice Sheet Modeling. SIAM Journal on Scientific Computing 2016 38:5, C504-C532. Are more sophisticated preconditioners required for Antarctica? Is scaling of those comparable to what you show for Greenland?
- Line 131-132: the linear solver convergence tolerances stated here seem loose to me. Have the authors verified that the solutions they have obtained at all their mesh resolutions are sufficiently converged/accurate? Accuracy/verification is an important thing to establish prior to studying scalability/performance.
- It is really great that you have set up a workflow and are doing performance monitoring! I agree that without this, it is inevitable that performance will be compromised in a big code with a lot of moving parts. It isn't entirely clear to me when the performance monitoring tests are run. Does it happen every time there is a PR merged into ISSM? Does it happen automatically or it must be run manually? Some further discussion of this is warranted. All I found was the following phrase: "it is quite feasible to periodically run an instrumented version of the code as part of the regular work of domain scientists", which suggests the performance testing is not run regularly or automatically, but perhaps I am misunderstanding.
- I did not really find discussion of how the load-balancing is done of the mesh on which the simulation proceeds, and how the mesh gets updated in a time-dependent simulation modeling

ice sheet evolution. How do you partition the mesh? Do you repartition every time the geometry (active mesh) changes? Or you partition a mesh including active and inactive cells once in the beginning? The latter approach has the potential of giving a lot of procs with no elements or poorly load-balanced meshes. Some discussion of this is warranted. I think the load imbalances you talk about in the paper have a different cause, unless I am misunderstanding.

- Is there any hope to improve the scalability of the matrix assembly?
- It would be interesting to compare ISSM performance to that of other open-source ice-sheet models based on the HO Stokes equations. I am not suggesting to do this in the paper, just commenting.

Minor comments:

- Change “to solve” in the first line of the abstract to “solving”.
- Line 12: remove commas around “thus”.
- Line 25: I don’t really understand the phrase, “standalone ice sheet projections are suffering from a large spread in climate forcing fields”. I don’t think “suffering” is the right word here. Please rephrase.
- After a colon, one does not use a capital letter. In many such instances, the colon should be replaced with a period.
- Lines 131-132: I suggest stating what is \epsilon_i here so the reader does not have to refer to the appendix.
- Line 140: there is a space missing before “For profiling and tracing”.
- Line 147: replace colon with period.
- Line 171: “data are stored” instead of “data is stored”. Change “scalars, both” to “scalars. Both”.
- Line 182: replace colon with period.
- Line 190: change “The instrumented” to “the instrumented”.
- Line 193: change “Sampling” to “sampling”.
- All strong scaling figures: please move the linear scaling line to be either below or above all the other curves. It’s very hard to see it with all the lines on top of it (e.g., in Fig. 3).
- Line 288: change “It” to “it”.
- Line 293: change “12.000” to “12 000”.
- Some of the figure captions say “(draft)”. Was that intentional? I suspect it was not.
- Line 420: change “to very a modest” to “to a very modest”.