Response to RC1:

In this work, the authors compare the performances of several different direct and iterative solvers, provided by MUMPS and PETSc libraries respectively, for solving the transient ice flow model using ISSM. Specifically, the authors target a well known transient benchmark problem (ISMIP-HOM, test F), in the case of a frozen bed or sliding bed. The flow model is constituted by the ice velocity part (Blatter-Pattyn model with constant viscosity), a part for reconstructing vertical velocities, and the mass transport part. The authors highlight some of the solvers that perform better on different mesh resolutions, for both frozen or sliding bed.
The detailed comparison of the solvers available in ISSM is certainly useful for the several ISSM users. However, I have a few reservations about the impact that this work can have on a broader community.

- The benchmark problem addressed in this work has several simplifications that makes it not very representative of real problems, most notably: 1. Geometry is very simple (in constrast with complex margins or bed roughness encountered in real ice sheets). 2. Viscosity is constant, making the model linear. In real problems viscosity strongly depends on velocity and temperature, which makes the problem much harder to be solved numerically. 3. A relatively high basal friction coefficient is considered, which is not representative of what can be found in ice streams and ice shelves.

Thank you for your comments and review. We chose to test a suite of solvers using a commonly used transient ice flow benchmark test (ISMIP-HOM experiment F), which makes the simplifications that you list, so that other researchers could reproduce the results as well as conduct their own tests using different, potentially customized, solvers or other ice sheet codes with a common and well-known model setup. In addition to updating the manuscript to highlight the limitations of our initial benchmark tests relative to real-world problems, we are including results from applying solvers to another benchmark test (ISMIP-HOM experiment A) in order to explore the impact of using a more realistic nonlinear viscosity model for solving the stress balance equations.

- The authors consider only off-the-shelf solvers that "naturally fit the ISSM framework", whereas several efforts (not mentioned by the authors) have been done in recent years in order to build efficient solvers/preconditioners tailored on the ice sheet problems. Some of these solvers have been demonstrated on large-scale simulations of Greenland or Antarctic ice sheets. See, for example, T. Isaac et al., SIAM J. Sci. Comput., 37(6), B804–B833; Tezaur et al., Procedia Computer Science, 51:2026-2035, ICCS, 2015, S. Cornford at al. J. Comput. Phys, 232(1):529-549, 2013; plus the one by Brown et al. already cited, but not discussed, by the authors. I recommend that the authors make it clear in the abstract that they are only considering the off-the-shelf solvers readily available in ISSM. I also recommend to consider more realistic problems and to mention relevant work in the literature.

We updated the text to include the recommended citations relevant to this work and specified that the focus of this work was to test readily available solvers in the abstract. As you mentioned, our results using PETSc solvers within ISSM are relevant to ISSM users. However, since many numerical models use PETSc, including the Parallel Ice Sheet Model (PISM) and the Community Ice Sheet Model (CISM), which has the ability to leverage PETSc solvers through the Trilinos package, we anticipate that our results should extend to other ice sheet models and benefit modelers beyond the ISSM community.

Minor comments:
- At line 144 the authors mention that they apply single-point constraints on velocity and thickness equations. I'd like the author to expand on this, mentioning how/with what values they constrain in a single point the velocity and thickness. Typically, single-point constraining is used in presence of a singular problems (which should not happen here), and it is known to artificially modify the spectrum of the matrix, which in turn can deteriorate the convergence of iterative solvers.

The text referring to using single point constraints was misstated and is corrected in the updated manuscript to note that we impose Dirichlet boundary conditions. Removing these entries from the matrix does not adversely impact the condition number of the stiffness matrix.

- Line 150, how the vertical velocity is reconstructed? With an L2 projection?

We solve the incompressibility equation to recover the vertical velocity, which is constant per element. Subsequently, we use an L2 projection to evaluate the nodal velocity. We updated the manuscript to clarify this point.

- The time reported in the tables is solver time, or total time (including assembly, linear solvers and I/O)? The weak scaling results for the iterative solvers are not very good and it would be useful to understand what is causing this.

We updated the text in the manuscript to note that the timing results include solving the system of equations, assembling the stiffness matrix, load vector, and updating the input from the solution.