

Interactive comment on “Coupling framework (1.0) for the ice sheet model PISM (1.1.1) and the ocean model MOM5 (5.1.0) via the ice-shelf cavity module PICO” by Moritz Kreuzer et al.

Rupert Gladstone (Referee)

rupertgladstone1972@gmail.com

Received and published: 7 January 2021

The paper describes an implemented approach to interactive coupling between an ice sheet model and an ocean model through a reduced complexity ocean cavity model. The project provides a compromise between complexity and efficiency, allowing large scale coupled simulations in which at least some aspects of cavity circulation are represented. This is a useful contribution to global modelling efforts. For the most part, the work is very clearly described and presented. The figures are well prepared and appropriate. I would recommend publication of this paper with some fairly minor modifications.

C1

General comments

I'm not entirely sure what other models might be competing with the current study in terms of global coupled models of intermediate complexity, but I think maybe the UKESM falls into this category? Are there publications about UKESM and how does it compare to your approach? Perhaps it is significantly more computationally expensive than your setup? I think it uses BISICLES and NEMO for ice and ocean models.

The first reviewer recommended including model validation in this paper. Robust model validation is a very large challenge and is not typically included in model description papers. I would prefer to see separate studies present some level of validation rather than try to include it here. Note that this is a highly tunable model, and so a good match to observations over a short time period should be straight forward to achieve. The real challenge will be in quantifying the uncertainty in the model as conditions evolve significantly over long periods of time, and this challenge is well beyond the scope of the current paper. The paper does present some level of model verification, especially regarding conservation, and this is useful.

I am aware that some aspects of my review read like an advertisement for the Earth System Modelling Framework (ESMF)! I should clarify I am fully independent of ESMF; I'm just an end user. In my experience ESMF is a very robust, user-friendly, well-documented code. So I guess I'm just a fan! It is a valid point though. Existing coupling frameworks (and I talk about ESMF rather than OASIS or others simply because I have the most experience with it) do provide more features and flexibility than the authors convey knowledge of in some places (see also my specific points below).

It is very interesting to learn that the PISM initialisation time becomes a significant factor as coupling timestep is shortened. It makes the question "what coupling timestep do you need?" rather important. I couldn't find a comparison of important result metrics between the 1 year and 10 year coupling timesteps. Information about computational cost is given, but is the actual behaviour different? For example, do they give the same

C2

total melt rate over the 200 years? Does the ice shelf thickness evolution look the same in both simulations? And grounding line and ice front evolution? You can't really address whether your framework meets requirement 2 until you've established whether or not a 10 year coupling timestep is sufficient.

I didn't find information about parallelism of the coupling. I infer from this lack of information that the coupling (bash script and file manipulation) all occurs on one processor. If this is the case, can you comment on how many processors you intend to use for production runs, and whether this might become a bottle neck for larger parallel simulations? For example, the ocean postprocessing and intermodel processing took 15

The Earth System Modelling Framework (ESMF) community adopts some terminology that I find quite beneficial in that it offers clarity in certain areas that can otherwise become slightly confusing. Individual models in a coupled system are referred to as components. This avoids confusion between the coupled model and the individual ice and ocean models. I like this clarity and I think it would be nice the the authors adopt it, but I do not wish to make this a requirement, just a suggestion.

Specific comments

Line 107: I don't quite understand the use of two horizontal dimensions. If I understood right, PICO just represents the overturning circulation. How do two horizontal dimensions come into play here? Or perhaps I should go and read Ronja's 2018 paper on PICO (lazy reviewer!)

Lines 113-114: This is also known as "sequential coupling".

Line 115: In this context, does an "integration step" mean running the model for a coupling time step? Can you clarify this in the text?

Line 118: "the last of these averaged fields" is a slightly confusing expression. Presumably PISM receives the average of the fields over the full (10 years in this example)

C3

coupling timestep? But this isn't clear to me from the chosen wording.

Figure 4: There appears to be a temporal offset between t_{ice} and t_{ocean} in the way that the Figure is presented. But the text suggests that the time period over which the two components are integrated should not be offset (line 113 refers to the "same model time"). Can you clarify this? It is not clear to me what is meant by "Sharing the same time axis", perhaps this is related to my question?

Line 134: I don't know what "entanglement" means in this context.

Line 139: Is "surrounding ice sheets" a typo in this context? I mean, the Zwally basins divide up the Antarctic Ice Sheet... perhaps you mean shelves not sheets here?

Lines 148-149: If you use only adjacent cells to populate missing values, presumably you iterate until all cells have a value? I mean there must, to start with, be plenty of cells that are not adjacent to a cell with a value.

Lines 159-160: I'm fairly confident that ESMF's "common" regridding algorithms include masked nearest neighbour remapping options that are very similar to what is described here. I intend to use these for remapping subglacial outflow from an ice model to an ocean model, though I haven't actually implemented this yet, and it looks like the required functionality to do this in a mass conserving way is already in place.

Section 5.1. I can't quite make the numbers add up here. The run with a 1-year coupling timestep takes 22700s. The ocean post-processing, interprocessing and PISM percentages given add up to 35

Line 193: I presume "total runtime" is the elapsed time not the cpu hours? i.e. the total computational time would be this number times 32?

Equation 1: Surely mass has dimensions of mass and smb has dimensions of mass/time. So how can you simply add these? I don't understand why there would be rates in this equation. Surely the total mass is the sum of the mass from each component at any given moment in time? The same comment for lines 214-215. How can

C4

you subtract a flux from a mass? I can't make sense of it!

Line 222: It is not clear from this description whether the dimensions of d_{osi} should be mass or mass/time.

Figure 9: To put this apparently small error into context, it would be useful to give some indication of how much mass is transferred between ice and ocean. I think the total mass transfer is probably a more relevant figure here than the total ice or ocean mass. We need to know that the error measures are small not only compared to the total mass of the coupled system, but also small compared to the amount of mass being transferred between ice and ocean over the intergration period.

Line 239: I don't think you can make this unqualified statement that the framework fulfills all three requirements. 1 and 3, yes, sure, but requirement 2 is really only partially fulfilled, as you continue on to discuss in the following paragraph.

Paragraph 248-267. This paragraph completely omits to discuss alternative forms of online coupling in which a single executable links to component runtime libraries and fulfills a role equivalent to your bash script as a master (or parent) program. The implication in this paragraph is that one must choose between an offline coupling in which components are called independently and an online coupling in which one component must be the master and the other the slave. But this is not the choice that a coupled model developer faces, there are many more options available. For example, the Earth System Modelling Framework (ESMF) offers full flexibility in this sense. The developer can choose whether to create a new parent routine or to establish one component as master. You might want to cite Gladstone et al GMDD FISOC paper (published in June 2020 as a discussions paper and now accepted pending minor revisions for GMD), which is essentially an online equivalent of your coupling structure, with a new (Fortran in this case) master program that calls the child components (which have been made ESMF compatible and compiled as libraries). This approach also allows flexibility in terms of switching between components (indeed we currently have a choice of two

C5

ocean models coupled through FISOC, and two furtehr ice sheet models are in the process of being incorporated). Of course there is still some overhead in terms of ensuring each component is compatible with the coupling framework (ESMF in the case of FISOC), so it is a longer development path than your bash script plus file manipulations, but perhaps not as onerous or restrictive as implied in the current version of the text.

Line 256. ESMF (for example) handles all parallel regridding in an efficient manner. So long as component mesh and field information can be made available in ESMF runtime data structures (which I acknowledge requires some coding and may not be trivial), the regridding between different partitioning is all handled automatically.

Line 260. The C/Fortran issue is a fairly small technical issue. There are plenty of codes around that use both at runtime.

Lines 269 – 274. The physical implications of this issue could be fairly interesting. Could the input ever come in beneath the turbulent mixed layer? Could the input of fresh water at the surface have a stabilising effect that would not occur if it was mixing up from lower down? Do you have any plans to investigate this further? This could all raise interesting questions that have been brushed over very lightly here.

Lines 285 – 288. This is actually a very important issue, especially because you intend this framework to be applied to long timescale simulations. You mention that work is in progress, but can you say a few lines about how this will be implemented? Can you also ive an example or two of the how you intend to use the model in its current form so that the reader can start to envisage how much of an issue (or not) the lack of evolving active ocean domain is?

Lines 293-295. These lines seem to imply that tuning PICO will somehow make up for the lack of a full representation of the complex 3D ocean circulation over the continental shelf and under ice shelves. This will not, in general, be the case. It is clear that this model is a compromise approach, a model of intermediate complexity with the benefit

C6

of efficiency. This has value; you don't need to try too hard to defend or justify limitations of PICO. I would prefer to see the limitations presented directly without implying that they can be overcome (through tuning for example).

311. "appropriate" is not well defined in this context, and perhaps not yet fully justified. How about just "intermediate" instead?

Interactive comment on Geosci. Model Dev. Discuss., <https://doi.org/10.5194/gmd-2020-230>, 2020.