

We want to thank Moritz Hanke and Basile de Fleurian for their detailed review and suggestions. We first summarize the major revisions to the paper. This is in answer to both reviewers since we feel that they agree on the main issues.

We have

1. extended the performance analysis to include initialization and different mapping methods to make them more relevant for comparison to other solutions
2. replaced the synthetic verification setup with a comparison of coupled and uncoupled simulations and discussed its evolution over time. This will allow readers to better judge the correctness of the coupling solution.
3. added a more detailed discussion of the pros and cons of our solution, in particular compared to other coupling libraries
4. made the technical focus of the paper clearer in title, abstract, introduction and discussion.

In the following, we answer the specific comments that have been made. In bold font, we repeat verbatim the comment from the reviewer. Continuing in normal font, we give our reply and revisions. In some cases, the order of reviewer comments has been changed to reply to multiple related comments at the same time.

As a result of the major revisions, there have been significant changes in the text, in particular in sections 3 (Experiments) and 4 (Discussion). Some specific reviewer comments have become obsolete and we noted so below.

## Specific comments by M. Hanke:

### **Abstract:**

**Please formulate the goals and achievements of the work presented by this paper more clearly.**

Rephrased the abstract, see point 4 In the list above.

### **Line 1: "Accurate earth system models"**

**Is there such a thing as an accurate Earth system model?**

You are right that this wording is too strong. We have rephrased this as “adequate”.

### **Line 30-31: "on different spatial and temporal discretizations, which complicates the coupling"**

**Isn't handling these issues one of the main functions of a coupler?**

Your assertion is correct. The cited statement is to justify the use of a coupling library, in that without different discretizations it would not be as needed. We discuss the comparison to hydrology models that are directly integrated into ISSM. This section has been clarified in the revised manuscript.

### **Line 34: "time points where"**

**Maybe: "time intervals at which"**

You are correct that the interval is configured, while the time points are implicitly derived. But the statement has been removed anyway as part of larger revisions in the section.

**Line 35: Repeated start of a sentence**

Changed.

**Line 35-36: "provides sophisticated numerical methods"**

**Which "numerical methods" do you mean (e.g. implicit vs explicit time stepping, remapping methods, ...)? Please be more specific.**

We refer here to the methods that you mentioned, as well as additional ones like time interpolation. We have added details in the text.

**Line 37-38: "easy to either add component like an existing ocean circulation model" I do not think, that this is the case. Reasons will be discussed further below.**

You are correct that easy is not the right term. Coupling of a global circulation model would of course run into problems with incompatible coordinate systems as you mention elsewhere. Local ocean models would not have that problem. Other complexities remain that cannot be handled generically. For example, input variables for one model must be derived from output variables of the other model if they don't fit exactly. We describe how we do this in the section on the CUAS adapter. We argue that the adapters presented in our submission reduce the developmental effort for many use cases, but we do not wish to imply that it's trivial.

We have rephrased this statement here and added above points in the discussion section.

**Line 41 and 42: "Sect. 3"/"Sect. 4"**

**No abbreviation for "Section 2" was used.**

Changed.

**Line 46: "overview of the coupling"**

**Maybe: "overview of the coupling setup"**

Changed.

**Line 46: "existing codes"**

**Maybe: "existing three codes"**

Changed.

**Line 50: Add long form of preCICE abbreviation**

Changed. (adding "Precise Code Interaction Coupling Environment", <https://precice.org/docs.html>)

**Line 51-52: "handles communication, data mapping, and coordination of the solvers"**

**Maybe: "handles communication, data mapping, and coordination between the solvers" to explicitly exclude communication like halo/ghost cell exchanges within a solver.**

By coordination, we mean here the temporal order of computation (coupling intervals, implicit vs. explicit, serial vs. parallel schemes). Your suggestion is still good, but we think "between solvers" fits better earlier in the list after "communication". We have rephrased this statement and added a short clarification what communication is meant.

**Line 53: "calls the preCICE library"**

**Maybe something like: "call initialisation routine of ..." would be more specific.**

We refer here to all calls made by the solver to preCICE, not just initialization. It is always the solver that calls preCICE. The confusion probably comes from the start of the quoted sentence, which previously read "To start a coupled simulation...". We have now included specifics about which routines are meant.

**Line 53: "all options"**

**Maybe: "all preCICE configuration options"**

Changed.

**Line 53-54: "This approach requires..."**

**In my opinion this is something for the conclusion.**

This statement has been removed, the topic is covered in the discussion section.

**Line 54-55: "the respective algorithms are selected"**

**Be more specific on the algorithms. Which task will they perform.**

Changed to add coupling scheme and data mapping.

**Paragraph 2.1:**

**Could you add more information about the adapters? Are they part of the solver, preCICE or independent codes. Who usually implements them (preCICE or model developers)?**

We have a short discussion of this later in the paper when describing the ISSM adapters. We have moved this to a short section here that also includes the additional details you ask for.

**Line: 62-63: Sentence "To establish the communication ..."**

**To me these are implementation details not relevant for this paper. How about:**

**"Communication channels between the processes of both solvers are established using a highly scalable algorithm (Totounferoush et al., 2021)."**

We agree that the focus here should be on the advantage for the user instead of implementation details and have rephrased this as you suggested.

**Line 70-71: Sentence : "However, if we wanted..."**

**This might be more suited for the discussion.**

You are correct that statement does not fit here in its current form. In the revised manuscript, we added a reference to Hocks 2025 to give the current state of spherical grids in preCICE, and we expand on the topic in the discussion.

**Paragraph 2.1.4**

**Could be part of the introduction as part of "state-of-the-art".**

**Line 89-90: Sentence "One of the goals..."**

**Should be mentioned in the abstract and in my opinion would fit better in the introduction section.**

You are right. This whole section has been replaced with the aforementioned section on adapter development. We have added the points to the introduction and abstract.

**Line 133-134: Sentence "This could be resolved..."**  
**Is this relevant for the paper?**

This sentence is intended as further justification for the development of the adapters we present here. An alternative to the new development would be integrating CUAS into ISSM. Flexible decomposition is easier (although not impossible) with independent solvers coupled with preCICE (or another coupling library). We use this in some of the performance experiments in this paper, where CUAS is allotted fewer CPUs than ISSM in a parallel coupling scheme.

You are right that this is not clear from the text, and we have added clarification.

**L149-152: Sentence "There are different..."**  
**This general description of adapters could be included in section 2.1**

As mentioned above, this has been moved to the newly added adapter development section (2.4)

**L164: "linear cell interpolation"**  
**Add reference to clarify its meaning.**

Changed by adding a reference.

**L179: "Partial interfaces"**  
**In Earth system modelling, I would consider the term "masking" more appropriate.**

You are right, we have rephrased this using the more established term.

**L284: Section 3.1**  
**Add minimal text to introduce this section.**

Changed.

**Figure 6:**  
**What additional information is provided by this figure when compared to Figure 1?**

It is a presentation that is more familiar to users of preCICE. It is auto-generated by preCICE tools and adds some details about the precise data mapping configuration, in particular whether mapping is performed on read or on write. However, since these are not very important and can be seen in the supplementary material when necessary, we have removed the figure.

**Section 3.1.2 Results:**  
**If possible, this section could be improved by evaluating the results in terms of quality of simulation results: Does the coupled setup produce more accurate results than the uncoupled ones?**

There are no known or analytic results for the synthetic setup, so it's not possible to quantify absolute accuracy. A real-world setup could allow comparison to measured data, but the setup (beyond what is necessary for the performance experiments) is significant modelling effort and outside the scope of this technical paper.

We have added a comparison between a coupled and uncoupled synthetic setup in the revised manuscript. This will show the effect of the missing feedback in the uncoupled setup.

**Line 382: Sentence "This also includes..."**  
**The imbalance between solver initialisation times could be explicitly measured by introducing an additional synchronisation point.**

We have added a section explicitly concerned with measurements for solver (and coupler) initialization.

**Line 383-386: whole paragraph**

**Maybe simplify by: "3. Running the solver for one coupling window."**

You are right that the paragraph should be simplified. We still think it is important to mention that subcycling in time does not have significant computational overhead. However, this whole section has been reworked entirely and the specific comment is obsolete.

**Line 390: "i.e., everything that happens before the participants or the first participant in a serial coupling"**

**This part is unclear to me and looks incomplete. "everything that happens before the participants [...] in a serial coupling" do what?**

You are right, the sentence is badly worded and also missing specificity. The specific comment is obsolete because the section has been reworked entirely. We have taken care to be specific about what is included in initialization.

**Line 389-393: whole paragraph**

**Initially it was not clear to me, that in both cases (serial and parallel coupling) both models have their own dedicated processes, but that in the serial case the two models are allocated to the same resources, while in the parallel case have each have their own dedicated computing resources. This could potentially be phrased more clearly.**

We agree. The resource allocation has been more clearly stated in the revised manuscript.

**[Line 389-393: whole paragraph, cont'd]**

**In MPI receive operations are often implemented as busy-loops that wait until the requested data is available. If the advance step contains an MPI receive operation, while waiting for the result from ISSM, CUAS might generate a significant load on the shared resources in the serial case. Therefore, a potentially interesting experiment could be the comparison of a setup with serial coupling using shared or dedicated resources.**

We ran experiments and have not found this to be a problem during simulation, at least on our system. But initialization is strongly affected, since both solvers initialize at the same time on contested hardware resources. We have included a experiments with exclusive resources and a dedicated section on performance of initialization.

**Line 394-397: whole paragraph**

**Does this represent a processes which in climate science is often referred to as "spin-up"? Could this be avoided by initialising the models using restart files from a previous run?**

In our current process, both models first perform spin-up uncoupled. This produces setups that are not entirely compatible. A period of relaxation at the beginning of coupling is probably unavoidable. This could be called coupled spin-up.

Due to limited investigation of the model results of the Greenland setup in this paper, we cannot confidently discuss this effect. Depending on which model parameters change between runs, it would certainly be possible to start from the relaxed state that was precomputed. For these reasons we mostly limit our discussion to later coupling windows when execution times have stabilized.

We have added this point to the revised manuscript.

**Figure 9: Sentence "With few processes, ..."**

**I would assume that the idle time for the 768-process case is higher due to CUAS having better scaling properties than ISSM. With 192 processes you seem to have hit a sweet spot, where both models roughly require the same time for the simulation of one coupling window.**

This is true, but this is not accidental. We are using a the distribution of 2:1 for CPUs assigned to ISSM and CUAS precisely because first experiment results with low CPU counts (one or two cluster nodes) suggested this as a sweet spot. We scaled up from there to see if the factor should be changed for larger numbers of CPUs. The results show increasing wait times, but adjusting the CPU distribution to minimize wait times does not significantly improve the runtime anyway.

In the revised manuscript we have described more clearly how we arrived at the resource distributions that we tested.

**Line 401-403: whole paragraph**

**Do solver iteration counts differ between serial and parallel coupling? Could you quantify the differences between the execution time of a coupling window of both cases? Can these differences exclusively be explained by the imbalance between the two solvers, which can be measured by the "advance" timer?**

Iteration counts of the solvers (nonlinear solvers in stress balance/thermal core in ISSM, linear solver in CUAS) fluctuate but do not differ in the aggregate. We have added a statement to the revised manuscript.

It is difficult to compare parallel and serial in fine detail since the change in coupling scheme implies changes in the resource distribution and mesh partitions. We have figures on both coupling window and solver times that allow some inferences.

**Figure 12:**

**For 192 processes ISSM has an execution time for a single coupling window of around 450 s and CUAS 200 s. However in Figure 9 both models seem to have for the same number of processes and each an execution time of around 650 s. Did I misinterpret the figures or how do you explain the inconsistency between the two figures?**

Our figures are not clearly labelled here. In Fig. 9, 10 and 11, the CPU counts refer to the combined count for both ISSM and CUAS. In Fig. 12 and 13, the CPU counts refer to the count for a single model. In Fig. 9, ISSM is assigned 128 CPUs, CUAS is assigned 64 CPUs. Fig. 12 does not include a measurement for CUAS with 64 CPUs, but the measurement for ISSM with 128 CPUs is consistent with the execution times in Fig. 9. All figures are based on the same measurements.

However, the specific figures mentioned have been replaced entirely in the revised manuscript.

**[Figure 12, cont'd]**

**Based on the measurements, I would assume that a parallel coupling setup with 512 processes for ISSM and 192 processes for CUAS should perform quite well. Do you agree? Why did you not use similar setups for the measurements of Figure 9?**

**I personally would start the performance analysis with this data and base the other measurements on its results.**

We did basically what you suggested. First experiments with 192 total CPUs suggested a distribution of 2:1 to be optimal. For higher numbers of CPUs, the wait times added overhead, so we looked at the execution times of each solver to calculate a distribution that would minimize wait times. You will find that your suggested distribution of 512:192 is very close to the value of 2.7:1 that we used. While this almost eliminated wait times (see Fig. 13), total runtime was only marginally improved (Fig. 11).

In the revised manuscript we state more clearly how we arrived at the distribution we used.

**Figure 13:**

**Does this again contradict the measurements from Figure 9?**

The figure has been replaced entirely, obsolete.

**Line 407-412: whole paragraph**

**Instead of defining a process count ratio, wouldn't it be easier to determine the number of processes for each component for a fixed coupling window execution time based on the measurements from Figure 12?**

As explained above, this is basically what we did.

**Section 3.2.2 Results**

**The remapping and communication time between two coupling windows usually is negligible in climate setups. However, initialisation time can have a significant impact on the overall runtime and this is a factor, which is often of interest to users. Therefore, this should be analysed in more detail. (See [1])**

**Line 420-425: whole paragraph**

**Please include an analysis of preCICE's impact on initialisation time.**

A large part of initialization is IO, which exhibits far less consistent measured execution times for different clusters. Measurements are very noisy. But we agree that it should not be ignored and we have added measurements of initialization times to the manuscript, in particular the time required to initialize preCICE as you requested.

**Line 416-419. "CUAS and ISSM use [...] are easy to adapt."**

**This are properties which in my opinion are not unique to preCICE but shared with other coupling solutions. This could be clarified in the discussion.**

You are correct, and we have stated so more clearly in the revised discussion.

**Line 425: Sentence "This imbalance..."**

**Can you explain why it is not possible in this case?**

We tried to do fix the imbalance by adjusting the distribution of CPUs but this did not have significant positive effects on the overall runtime. Moving resources from CUAS (scales well) to ISSM (not scaling as well) mostly makes the CUAS slower and ISSM only marginally faster, while still balancing the wait times. We have rephrased this to make clear what we mean and that it is based on our experiments.

**Line 440-441: Sentence "For example, the ice sheet code..."**

**I was told that this example does not make much sense, since hydrology is already included in Elmer/Ice [3].**

Most ice sheet models include at least some hydrology model. We mention in the introduction that both ISSM and Elmer/Ice include multiple different hydrology models already. Models have advantages and disadvantages and it is worthwhile to compare them in the same context with the same ice sheet model. Coupling adapters makes this easier, since not every hydrology model has to be integrated into every ice sheet model. All the reasons for coupling CUAS with ISSM (allow different grid resolutions, independent development, ...) apply to Elmer/Ice as well. We did not intend to imply that hydrology would be new to Elmer/Ice. In the revisions, we have added an explanation and a statement regarding the value of diverse hydrology models.

#### **Section Discussion:**

**If I am not mistaken, your adapter design should allow for a single executable setup using serial coupling, in which ISSM is called by the CUAS-preCICE adapter within the CUASSolver. This could significantly improve performance of the serial setup. Did you consider this option? Maybe you can a discussion about this.**

This is worth thinking about. However, preCICE is not intended to be used in this way, nor does it provide the necessary API. One adapter writes data, the other adapter reads it. Since the computation performed by each model is still the same and overhead for data mapping is small, we would not expect performance improvements that justify the increased development effort. It would also be a less flexible/extensible solution, e.g., we cannot easily add more participants. Since there are so many barriers and drawbacks to actually implementing it, discussion of the specific suggestion would be distracting.

#### **[Section Discussion, cont'd]**

**I think in order to improve the scientific value of the paper you should extend the discussion on pros and cons of using a generic coupling library compared to one specialised for Earth system modelling. In case you are interested here are a few points:**

**\* Usually specialised couplers like OASIS3-MCT or YAC basically provide the coupling of 2D fields on a spherical surface. For global circulation models that simulate for example an atmosphere or an ocean, this constraint is adequate. For the coupling of these models specialised couplers are probably the best choice since they also provide user interfaces and remapping options that are optimised for these cases. In addition, these models have to my knowledge not the required software infrastructure to efficiently support implicit coupling provided by preCICE.**

**\* For the remapping of properties like surface freshwater flux conservative remapping as described in [2] is used. This remapping scheme requires the usage of spherical geometry during the interpolation weight computation. This is currently not supported by preCICE and its support could become difficult. This could be overcome by supporting a remapping scheme based on user-provided weight files. These weight files could be produced in a preprocessing step by specialised tool (e.g. CDO, ESMF, or YAC).**

**\* I think the best use-case for a general coupling library in earth system modelling are cases where the constrain to the coupling of 2D spherical fields is insufficient (e.g. glacier ground interaction as described in this paper). It could even be feasible and interesting to set up configuration where coupling between ISSM and CUAS is implemented using a general coupler like preCICE and in the same setup additional coupling with an atmosphere and or ocean model is implemented with a specialised coupler.**

As stated above, we have expanded the discussion to include a comparison of (types of) coupling libraries. Your points are highly appreciated and merits more discussion.

**Line 455-456 Sentence "The use of preCICE..."**

**See comment above.**

We will reference some initial results towards using preCICE

**Line 465: "<https://git.rwth-aachen.de/yannic.fischler/cuas-mpi>"**

**The associated repository is not publicly available. The link to the version used for the paper is available, so removing the link to the current version would not be a problem.**

We mixed up the URLs of our private and public repositories. We have replaced it with a link to the public repository, <https://github.com/tudasc/CUAS-MPI>.

**References [by Moritz Hanke, included for completeness] :**

**[1]: [https://raw.githubusercontent.com/IS-ENES3/IS-ENES-Website/main/pdf\\_documents/IS-ENES2\\_D10.3\\_FV.pdf](https://raw.githubusercontent.com/IS-ENES3/IS-ENES-Website/main/pdf_documents/IS-ENES2_D10.3_FV.pdf)**

**[2]: <https://data-ww3.ifremer.fr/TRAINING/DOC/SCRIPusers.pdf>**

**[3]: <https://kannu.csc.fi/s/6CRGEdSZPEajnL6>**

## Specific comments by B. de Fleurian:

### Title:

Depending on the focus of the final paper I think that the title should be modified to be a better reflection of the content of the study. In the current version I would expect a study about the results of the coupling itself rather than a presentation of the coupler. I would also exclude version numbers from the title and have them later in the text to avoid cluttering the title.

We changed the title to more accurately reflect the content as you suggested. The version numbers are a requirement of the journal. But we have moved them to the end. The new title is

“New generic coupling adapters for ice sheet and subglacial hydrology models (ISSM-preCICE Adapter 0.4, CUAS-MPI 0.1)”

The version of ISSM and preCICE are implied by the versions of the adapter.

### Introduction:

- **Line 19: The sentence on starting : "While the hydrological..." is not extremely clear to me, I would suggest to rephrase with something like : "While the hydrological system evolves on long time scales in the centre of ice sheets, at the margins and particularly in Greenland its evolution is faster especially during the melt season"**
- **Line 20 : Emphasis should be made here on the fact that the interaction between subglacial hydrology and ice dynamics are complex, there are a few studies that could be referenced here from observations or model e.g.(de Fleurian 2022, Ing 2024)**

We have rephrased this sentence. The references are appreciated, as they strengthen the need for a robust coupling solution.

- **Paragraph from Line 22 to 28: Here and later in the manuscript I do not understand the special focus on the SHAKTI model, looking through ISSM code I can count 7 implemented subglacial hydrology models of different complexity, I might have a bias in this matter as I am the author of one of those models but I think that this paragraph should be reworked. Also, in this paragraph GlaDS is missing a reference.**

SHAKTI is only used as one example of a hydrology model that is integrated in ISSM, but you are right that it does not have a special significance beyond that or over other models integrated into ISSM. We have clarified this revised manuscript and added more diverse examples and references.

- **Line 29 : missing a reference to Fischler.**

Changed.

- **Line 34 : I suggest to rephrase with : "multiple properties, such as the time point for example...."**

This sentence was also mentioned by reviewer Moritz Hanke. We have rephrased the entire sentence.

- **Line 41 : "next" should be removed**

Changed.

**Software :**

- **Figure 1: On the figure it is not completely clear what fields are actually used by the given models. For example I expect that the grounded ice melting rate is used as input for CUAS-MPI. That might be something that have been forgotten or just that I missed something in the design of the diagram.**

All data sent by one model to the other is used as input by the receiving model, either directly or indirectly after some transformation that is described in the text. We have updated the caption to make this clear.

- **Section 2.1.3: It is stated here that the coupler has a fixed time window. Is that something that could evolve. I expect that for an efficient coupling and in the case of subglacial hydrology you would lie to adapt the time window's length depending on the season to catch the different variability of the subglacial drainage system.**

The window was chosen conservatively and will be adapted when necessary for running real simulations. preCICE does not currently support adaptive coupling windows. Performance experiments (also mentioned by reviewer M. Hanke) show that data mapping is not a significant overhead, so we would expect coupling windows can be made significantly shorter in summer without much unnecessary overhead in winter. Also note that the solvers are allowed to do smaller time steps than the coupling window.

In our new experiments, we already included this, the coupling window is now 1 day, same as the ISSM time step. We have also discussed possibilities to get even shorter coupling windows if necessary.

- **Line 105: The simulation codes (G4000, G250) are used here without references, that should be fixed.**

The model names come from the article referenced at the end of the sentence. We have rephrased this statement to better connect the model and reference and explain the model names.

**line 110: The ordering of the figures is not respected with Figure 3 appearing before figure 2.**

We don't see this in our manuscript. We are not sure what is meant here, as none of the mentioned figures are around line 110. However, we have the order of the figure made this a point of emphasis before submission of the revised manuscript.

- **Line 120: I think that "set" here is a better term than "aligned"**

Changed.

- **Line 129: I am not sure that SHAKTI actually does subdivides time-steps (and I can not find where it would do it in the code), The Double Continuum approach in the other hand does it (in the src/c/core/hydrology\_core.cpp file)**

You are correct, this seems to be a mix up. The point remains that ISSM cores can in principle subdivide time steps. We have corrected the statement in the revised manuscript.

- **Section 2.3: While I approve the authors for stressing the limitations of the coupler I feel that the presentation they chose is disrupting the flow of the paper and sometimes makes for quite awkward sections (2.3.6 for example) and that including the limitations as standard text would be easier to read.**

We see your point, and also that supported features and limitations are not easily described separately in the way that we attempted. We have rephrased every subsection to include the limitations in normal text.

- **Line 180: Regarding the partial interfaces, I expect that this is also true for CUAS-MPI and that the hydrology model is also running on ocean and ice free nodes?**

This is correct. Both ISSM and CUAS have the ability to ignore inactive regions, so the impact on performance is small. Other models may have bigger problems with it, we added this to the discussion.

- **Line 185: I am missing the point of this remark on 3D variables.**

These extrusion and depth-averaging features are mentioned here because they are supported by the coupling adapter. We have rephrased this statement to make clear that the features were added to the adapter and the statement does not refer to the existing features in ISSM.

- **Line 202: I might be missing the point here but I don't see how setting the initial condition is really a limitation. To me the initialisation is part of the modelling workflow and the user is expected to set those values.**

By default, preCICE assumes that all exchanged data is initially zero everywhere. This is rarely a valid assumption for earth system models, and we wanted to highlight this for readers familiar with preCICE but not earth system models. In addition, some ISSM variables are not required to have real initial values since they are normally computed by one ISSM core before they are used by other cores. The adapter requires all coupled variables to be properly initialized because the first set of variables is sent before any computation is performed.

Limitation may be the wrong word here. The entire paragraph of limitations has been rewritten as mentioned above with a better explanation of this particular issue.

**Section 2.3.6: As stated ISSM performs multiple iteration in a single coupling window. However I wonder what strategy is used to pass the data to the coupler and hence the coupled model. Is the data passed only a timestamp at the end of the coupling window, a mean over the coupling window or the full history. For ice dynamics with a quite slow response time that should not be a big issue but it can be problematic when dealing with the faster evolution of the subglacial drainage system.**

Models normally write and read a single snapshot of the coupled variables at the end of the coupling window. Using subcycling in combination with the time interpolation feature of preCICE, the history of the variables over the coupling window can be used, at least in the serial coupling scheme. We are not currently using this, but you are correct that it may be necessary for real simulations and the adapters can be extended to support this. However, before that the size of coupling windows can also be decreased as we have already done in our new simulations. We have added these points to the manuscript.

- **Line 220: Perhaps the advance function behaviour should be described more in depth here.**

You are right. The “advance” function includes exchanging and mapping of data and advances the time. In addition, it may block to wait if the other solver has not finished its computation yet. However, in the revised manuscript we have avoided referring to specific routines of preCICE and instead phrased the statements more based on concepts and actual functionality than function names.

- **Line 230: A reference to the original CUAS model is missing here.**

Changed.

- **Line 263: Is there a reason to use the ice Thickness rather than the levelsets provided by ISSM to define the bndmask?**

We want the grounding line (GL) in CUAS to be consistent with the bed topography. The mesh resolution in ISSM at a given location could be much coarser than the grid resolution in CUAS. Via the preCICE coupler, we only obtain the interpolated level-set field representing the GL location on the coarser ISSM mesh, but this might not correspond to the GL position on the higher-resolution bed topography in CUAS.

We have mentioned the reasoning in the revised manuscript.

- **Line 274 : "provide" in place of "provided"**

Changed.

- **Paragraph line 276 to 279: You mean here that the forcing is splitted in surface input, friction generated water, GHF production of water and potentially other sources and that the coupling only provides friction? This sentence is not very clear and should be rephrased.**

CUAS until now only had one water source from input files that combined all sources, i.e., runoff and basal melt. In the coupled setup, we need to be able to change one source (basal melt) independent of the other because it is computed by ISSM. CUAS now has the ability to aggregate multiple water sources, e.g., basal melt from coupling, runoff from input files. We have described this more clearly in the revised manuscript.

### Experiments:

- **Line 281 : "a model of subglacial hydrology" should be replaced by "CUAS"**

Changed.

- **Line 311: It should be added here that the mapping scheme is used to convert between grids.**

Changed.

### Results:

- **I think that the choice of results here is not the more informative. The difference map do not really give a good idea of what the fields really look like and that is an issue for the comprehension of the impact of the coupled model. I get that the scope is more to assess is the coupler is actually behaving properly but we would need to have a better understanding of the results for that. I would like to see a steady state field for the different variables, at the end of the spin-up or at the end of the coupled simulation to confirm that the coupled model is actually behaving in a**

**proper way. It would also be interesting to see a time evolution of bulk variables to see if there are any transient changes. Connected to that last remark I wonder if the dynamic of the coupled system would be the same if the coupling window is changed?**

The main focus of the manuscript is on the technical details of the implementation and its performance. However, we agree that showing only the differences with respect to a reference field (Fig.7) is not very convincing. As both reviewers requested more details, we have updated our synthetic setup, adding seasonal forcing.

The new experiment is a comparison between coupled and uncoupled simulations. We have added figures at the end of the spinup and the peak of melt season as well as transitive bulk variables as requested. This should allow a much greater insight into the behavior of the system.

In-depth analysis will require a full real-world scenario that is out of scope of this work.

- **Line 322. Effective pressure is on panel 7b not 7a**

Figure replaced, obsolete.

- **Line 325. I do not agree with the analysis of the results in the cold based region. To my eyes, the interior part of the cold based region shows a decrease in effective pressure while the outside parts shows an increase. without the original field it is hard to figure out what the effective pressure field of the coupled simulation looks like and if that result is reasonable. I also wonder why there is a quite large decrease in effective pressure and ice thickness at the grounding line of the domain. Is that due to a retreat of the grounding line or an error associated with the mapping of the grids?**

Indeed, our statement was very incomplete. The interior part of the cold-based region shows a decrease in effective pressure, while the outer parts show an increase. In the revised version, we have shown differences only on the part of the mask where both simulations have grounded ice to address the differences in grounding line position in the figure and rephrased the text accordingly. As mentioned above, we also show the fields and not just the differences.

- **Line 328: I don't see a large interest in the anom. simulation. In my opinion it would be more interesting to give more details on the reference simulation and it's comparison with the uncoupled run. if the author decide to keep this simulation in the end I think that presenting it without a plot of the transient evolution is not very useful.**

We agree. We replaced this with a comparison of coupled and uncoupled runs with seasonal runoff forcing.

- **Line 330: The Budd friction law given here is the good one the one on line 291 should be changed.**

Changed.

- **3.2 Performance : I expect that the new set-up is needed because the synthetic one is running to fast to be relevant to spot performance issue? If that is the case that should probably be stated here.**

You are correct, the synthetic setup does not exhibit realistic performance characteristics (e.g., lower iteration counts) due to its low complexity. We have added this statement to the revised manuscript.

- **Line 346: I am not sure that not restricting CUAS to warm-based ice is the good way to go. If I am not mistaking the effective pressure can be computed everywhere even if there is no water and then it just becomes the ice overburden pressure. In this case wouldn't it be better to run CUAS only on warm nodes and then update the effective pressure to be equal to overburden ice pressure on the frozen nodes. My main concern with that is that frozen region could actually have an impact on the simulated subglacial drainage system as they are acting as dams preventing the flow of water which could change the water pressure distribution at the base of the ice.**

This could be a modeller's choice. The current design of the CUAS preCICE adapter would allow for later improvements to deal with frozen areas. For the simulation presented here, we decided not to introduce discontinuities in the effective pressure field at the cold/warm transitions. In the cold ice areas, the basal melt rate is zero, the effective pressure is relatively high, and the sliding velocity is low.

However, in our new simulations we have followed your advice and allowed the mask to be activated and deactivated based on basal temperature received by ISSM. Details are described in the revised manuscript.

- **Paragraph Line 361 to 363: This paragraph is not very clear and should be rephrased.**

Changed.

- **Line 368. I don't see why serial and parallel coupling would reasonably have different results if everything else is equal. The only difference I can see is that the coupled scheme as some idling time for one of the solver but don't they get the same result from the other participant anyway? I feel that I am missing something here that should probably be explained better.**

In serial coupling, the second participant immediately uses the results that were just computed by the first participant. Thus, feedback may be propagated faster. This is not an issue if the coupling window is sufficiently short. As explained above, we have used significantly shorter coupling windows in the updated simulations. We have removed the statement here since it is not relevant.

- **Figure 9-10: It would be nice to have the simulation with 384 or 792 CPUs presented in both those figure to have a better idea of the difference between the two coupling schemes.**

The number of CPUs in the figures is misleading. In Fig. 9, 10, and 11, the number refers to the combined CPUs used for both participants. In Fig. 12 and 13, the number refers to the CPUs used by each participant individually. The offending figures have been removed since they are no longer relevant to the revised text.

Comparison between parallel and serial coupling are difficult in general, since either the total number of CPUs or the CPUs per solver will be different. We hope the new figures are at least consistent in this regard.

## Discussion and Conclusion.

- **I think that a better point could be made to stress the pros and cons of this coupler against earth system specific couplers. The conclusion should also make a better job at synthesising the papers results rather than presenting future work. I also think that the reference to Keyes should probably appear in the introduction of the paper rather than here.**

We agree with all of these points. As stated in our opening remarks, we have expanded the comparison of different types of couplers. We have rephrased the conclusion to summarize the results of the paper in relation to the goals outlined in the introduction. We will have moved the reference to Keyes to the introduction.

- **Line 415: I am probably biased here but I think that citing DOCO (deFleurian 2022) here would be more relevant as this model has actually been coupled to an ice dynamics model but I am not aware of an application where SHAKTI was coupled.**

SHAKTI (as well as other hydrology models) is integrated into ISSM, and therefore coupled as well, albeit not using a coupling library. This makes it a relevant comparison. But you are right that it's not the only (or even most) relevant comparison, and we have added more diverse references and a clearer statement of the intent of the comparison.