Review of the Manuscript: New generic coupling adapters for ice sheet and subglacial hydrology models (ISSM-preCICE Adapter 0.4, CUAS-MPI 0.1)

March 18, 2026

Summary

The revised manuscript presents an updated and improved version of the original work and the manuscript now demonstrates a clearer scientific contribution in addition to its technical achievements.

The authors have addressed the major concerns raised in the initial review. The revised manuscript now provides both a solid technical contribution and a clear scientific rationale for the coupling approach. I recommend the manuscript for publication in GMD after minor editorial corrections.

The main remaining concern with the current revision of the paper is the comparison of preCICE with coupling libraries specialized for Earth System Modeling. The manuscript lacks a sufficiently critical and balanced analysis of the differing characteristics of these couplers. At times, the discussion appears notably biased, possibly due to limited background information on specialized ESM couplers. Further details and specific points are provided in the "Special comments" section.

Comparative measurements

To provide context for my performance-related comments, I created a YAC-based toy setup with two components representing CUAS and ISSM [1]. The CUAS component uses the G600 grid (2,492 × 4,236 cells), while the ISSM component uses a G1000 grid (4.1 million cells), generated according to the information in the manuscript. As in the preCICE measurements, the number of ISSM processes was set to be 2.2 times higher than for CUAS. Interpolation was configured as 1-nearest-neighbor for two-way exchange (the actual number of fields does not affect initialization time, since weights are computed once and reused). All measurements were performed on DKRZ's Levante supercomputer and reflect the time YAC requires to compute interpolation weights and establish communication patterns for subsequent data exchanges, based on the local grid information provided by each process. The results are presented in Table 1 and shown together with preCICE results in Figure 1.

General Comments

The paper compares preCICE primarily with the ESM couplers OASIS and YAC, which are, for good reasons, quite similar. However, several other relevant couplers are not considered in

[1] `https://gitlab.dkrz.de/YAC/YAC-dev/-/tree/example_issm_cuas?ref_type=heads`

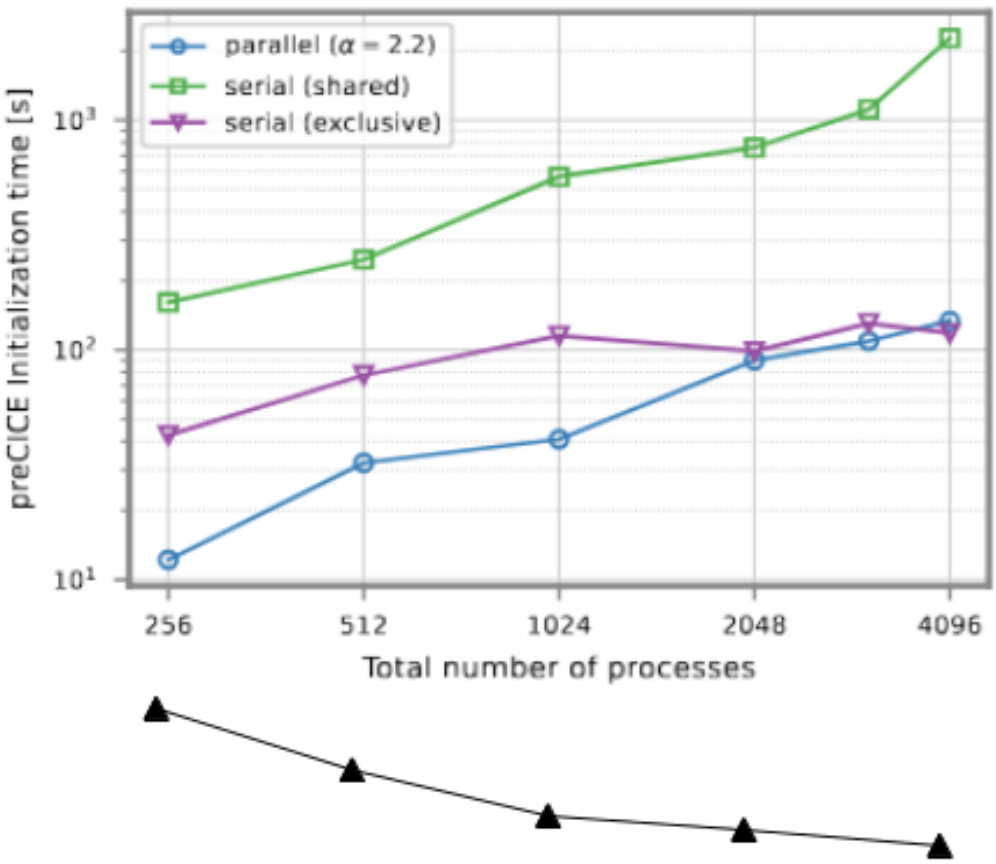

Figure 1: Initialisation time in seconds for preCICE and YAC (in black). Figure based on Figure 9 of the manuscript.

the discussion. As a result, any general conclusions about ESM couplers should be made with caution.

I recommend avoiding sentences that begin with "preCICE," as this may conflict with language rules and, depending on GMD's editorial standards, could result in the use of "PreCICE."

The manuscript does not provide a clear rationale for the choice of coupler in this setup. It appears that an ESM-specific coupler could have been used as well. A statement such as "Any coupler would have sufficed and we selected preCICE," or "We chose preCICE specifically to demonstrate its applicability in an Earth System Modeling context," would be a valid and transparent justification.

Additionally, "Earth System Modeling" is typically capitalized.

Specific Comments

- Line 51-54: "Hocks and Uekermann (2026) compares the general data mapping of preCICE to those of specialized couplers using an ESM mapping benchmarks and concludes that preCICE is competitive. While specifically the radial-basis-function interpolation performs well for the considered smooth test functions, preCICE still lacks specialized conservative data mappings – a limitation that can, however, be overcome by bespoke pre- and postprocessing."

  Hocks and Uekermann (2026) is currently under review, and I am one of the referees. The con-

Total number of processes	Initialisation time (s)
256	2.933
512	1.602
1024	1.021
2048	0.890
4096	0.762

Table 1: Measured initialisation times for different total process counts in the YAC-based toy setup.

clusions of that manuscript have received criticism from multiple reviewers, including myself. Therefore, I recommend refraining from citing these conclusions at this stage.

- Line 55-56: "The missing ESM specialization of preCICE is less important in non-global ESM scenarios, such as the setup of this paper, or scenarios where more flexibility is needed such as adding further models."

  This statement could be interpreted as: "The missing ESM specialization of preCICE is less important in non-global ESM scenarios (such as the setup in this paper), but also in global ESMs if greater flexibility is required, for example, to add additional models."

  If this is not the intended meaning, please clarify. As currently phrased, it is unclear why the need for more flexibility would make ESM specialization less important, especially in the context of global ESMs.

- Line 57: "The functionality of preCICE goes beyond communication and data mapping, as specialized ESM couplers typically offer."

  The wording here is somewhat misleading. ESM couplers can also provide additional functionalities. Consider rephrasing, for example: "Functionality not typically offered by ESM couplers, but provided by preCICE, includes, among others, ..."

- Line 58-59: "Implicit coupling schemes including quasi-Newton acceleration (Mehl et al., 2016), time interpolation (Rodenberg and Uekermann, 2025), or multi-scale coupling (Desai et al., 2023) are all features that are relevant for ESM."

  Implicit coupling schemes are a niche area in the ESM community; explicit coupling remains the standard for ESM component interaction. One reason for this is the additional resource consumption required by implicit schemes, which is a significant consideration. YAC, for example, does not implement implicit coupling due to lack of demand, but this could be added if needed.

  Similarly, time interpolation could be implemented in YAC, but until recently there was no demand for it.

  The meaning of "multi-scale coupling" is unclear to me. ICON, for instance, supports nested grids with different resolutions and timesteps, which could be seen as multi-scale. However, exchanges between nests currently occur only at the largest timestep, not due to coupler limitations. YAC supports different timesteps and multiple grids with arbitrary resolution. Would this qualify as multi-scale coupling?

  Overall, I do not agree with the argument that preCICE is superior to ESM couplers due to a lack of certain features, as these features may not be as important or may already be available in ESM couplers.

- Line 57-62: Whole paragraph

  This paragraph appears to argue for the superiority of preCICE over ESM couplers, but omits counterarguments.

  Some of the points made are debatable: missing functionality may already be available or may not be relevant; the claim of limited documentation for ESM couplers is questionable; and the assertion of unsustainable software development is not convincing (OASIS, for example, has been maintained for over 30 years). The integration of OpenFOAM and FEniCS is certainly valuable, but its relevance to the ESM community should be clarified.

  As noted in the abstract, using a generic coupler versus an ESM-specific coupler has both advantages and disadvantages for the ESM community.

  A balanced discussion of the pros and cons of both coupler types would be more appropriate in the discussion section. In the introduction, it would be better to focus on preCICE's strengths and unique features that make it suitable for coupling ISSM and CUAS. Additional features could be mentioned in the "2.1 preCICE" section, but unless you are sure that certain features are unique to preCICE, please choose your wording carefully.

- Line 78-80: "All preCICE configuration options can be set at runtime in a configuration file that is shared by the participating solvers, via which the respective algorithms for communication, data mapping, and time stepping are selected."

  This wording could be misunderstood as: "During runtime, participants generate a configuration file, which is then read by preCICE and shared with other participants."

  Is this what you intended? Or did you mean: "Configuration options such as the communication algorithm and data mapping are specified in a configuration file, which is read by preCICE during initialization and then shared with all participants."

- Line 111-114: Whole paragraph

  Just a remark: The CalculiX-preCICE adapter uses what is often called a "library approach" in the ESM context, while the CAMRAD II-preCICE adapter is an example of a "framework approach."

- Figure 3

  Remark: The figure clearly illustrates the design and dependencies of the ISSM preCICE adapter. Since the "Adapter" is the entry point, I would have placed it at the top-left or top-middle, but this is just a personal preference.

- Line 185-186: "Mesh connectivity is added to support mapping schemes like linear cell interpolation (Chourdakis et al., 2022)."

  I could not find a reference to linear cell interpolation in Chourdakis et al. (2022). You may want to cite https://mediatum.ub.tum.de/1685618 instead.

  Remark: In 2D, this method is likely similar to YAC's "average with barycentric coordinates" [2] or OASIS' linear interpolation.

- Line 189-190: "For best results, this would require multiple coupling interfaces for different finite element types, and mesh connectivity would not be available."

[2] `https://dkrz-sw.gitlab-pages.dkrz.de/yac/d0/d3a/interp_method_avg.html`

Remark: Alternatively, you could provide a callback mechanism for a weight computation function. This function, set by the user, would compute the interpolation weights for a given target point. Such an approach has already been used successfully with YAC.

- Line 195-196: "But it may be beneficial for, e.g., ice-ocean coupling to only couple over floating ice."

  Or, for example, when coupling with a globally defined atmosphere model.

- Line 420: "No energy is wasted"

  Remark: Even in an idle state, the power consumption of HPC nodes is non-negligible.

- Line 436-437: "In earth system models, initialization of the coupled setup is often a significant part of the runtime, see for example the runtime analysis of the OASIS3-MCT coupling library in Craig et al. (2017)."

  The phrase "initialization of the coupled setup" is ambiguous. It could refer to the setup of the coupling library (including weight computation, reading weights from file, and establishing communication patterns) or to the entire setup, including model initialization. Please be more specific.

  Using this paper as a general reference for initialization in ESMs is not ideal. Other couplers have demonstrated much faster initialization [3] [4]. The referenced paper is also rather old and may not reflect the current state of the art.

  While this was a significant issue in the past, such measurements remain relevant.

- Figure 9:

  This figure suggests a potentially serious issue in the initialization of preCICE. For this problem size and these process counts, I would have expected a nearly linear decrease in runtime as the number of processes increases.

- Line 440: "For comparison, Fig. 10 and 9 show the time required to initialize the solver."

  You may have referenced the wrong figure for the CUAS data here.

- Figure 12 and 15

  Consider presenting both diagrams side by side in the same figure.

- Figure 18:

  Out of curiosity: can you explain the long runtime of LCI? Based on the method, I would have expected only a minor increase compared to NN. Perhaps a detailed performance analysis is warranted. In YAC, a similar method is among the fastest in terms of runtime.

- Line 574: "We believe preCICE is at least on par regarding the basic functionality."

  This statement directly contradicts the following from the manuscript: Line 581-582: "However, we have identified several missing features (spherical coordinate systems, masking, specialized mapping methods) that may be highly relevant for other ESM applications."

[3] https://doi.org/10.5194/gmd-16-2833-2023
[4] https://doi.org/10.5676/DWD_pub/nwv/icon_003

I would consider these to be basic functionality. Furthermore, as discussed later, performance also does not appear to be on par.

- Line 577-578: "The regridding benchmark by Hocks and Uekermann (2026) already demonstrated that numerical performance is equivalent."

  As mentioned above, I would not cite this paper in its current form as a reference here.

- Line 578-579: "The results in Sect. 3.2 are in line with those reported for other libraries (Craig et al., 2017; Hanke et al., 2016)"

  Using Craig et al., 2017 as a performance reference is misleading. The measurements were performed on a slower machine, and the version used in that study contained a bug that negatively affected initialization performance. After fixing this bug, performance improved significantly [5].

  Note that in Hanke et al. (2016), the x-axis is "number of nodes per component." To get the total number of processes, multiply by 48 (two components times 24 processes per node). This was also on an older machine. Nevertheless, Figure 3b should represent values for a problem size similar to that in Figure 9 of the preCICE paper. At high scales (4096 processes), YAC initialization time is about 1s, while preCICE is roughly 100s. The scaling behavior is also quite different. (As a side note, the YAC paper used version 1.2; the latest version is much faster.)

  Therefore, I do not support the statement that preCICE's performance is in line with the couplers mentioned here. When taking my comparative measurements (see above) into account, it seems to be quite the opposite.

- Line 583-585: "The main advantage of preCICE due to its generic nature is the potentially much larger community, collaborating in the development of adapters and the coupling library itself."

  According to a survey on the OASIS3-MCT website, at least 67 modeling groups are using it. Comparing actual community sizes is challenging.

  Additionally, ESM couplers often benefit from institutional funding. For example, ESMF receives substantial support, and there are dedicated national and international projects (such as natESM) that provide professional support for coupling infrastructure development. DKRZ, for instance, exists solely to provide compute power and services to the ESM community.

  The advantage of being generic may be offset by these other factors. Thus, the argument as stated is one-sided and should be revised.

- Line 596-597: "We found preCICE to be competitive in all aspects."

  Considering the information presented in the manuscript and my comments above, I reach a different conclusion.

- Line 598-599: "We also provided arguments that either approach is superior in the long run to integrating different models into the same monolithic code."

  This sentence seems out of place in this context.

[5] https://oasis.cerfacs.fr/wp-content/uploads/sites/114/2021/09/2020_OASIS_CW2020_2.pdf