GMD Reviews and Authors' Response concerning the paper "A parallel implementation of the confined-unconfined aquifer system model for subglacial hydrology: design, verification, and performance analysis (CUAS-MPI v0.1.0)"

Yannic Fischler[1], Thomas Kleiner[2], Christian Bischof[1], Jeremie Schmiedel[4], Roiy Sayag[4], Raban Emunds[1,2], Lennart Frederik Oestreich[1,2], and Angelika Humbert[2,3]

[1]Department of Computer Science, Technical University Darmstadt, Darmstadt, Hesse, Germany
[2]Alfred-Wegener-Institut, Helmholtz-Zentrum für Polar- und Meeresforschung, Bremerhaven, Bremen, Germany
[3]Faculty of Geosciences, University of Bremen, Bremen, Germany
[4]Department of Environmental Physics, BIDR, Ben-Gurion University of the Negev, Sde Boker, Israel

Correspondence: Yannic Fischler (yannic.fischler@tu-darmstadt.de)

1 General Comments

We thank the reviewer for his extensive work in the review process. We have tried to be more precise in our software section and adopted most of the minor issues raised.

2 Review 1 by an anonymous referee

2.1 General

The authors have made considerable efforts to implement the requested changes from the first round of reviews. I believe that the clarity of the message is better conveyed by the restructuring of the manuscript and the overall English improvements. The introduction is now complete with a well written state of the art, and the place and originality of the present study -in particular, with respect to previous publications by Beyer et al. 2018, has been more explicitly emphasized throughout. The additional figures are helpful, and the occasional rework of existing ones appropriate. Section 3 is also much easier to follow, with a clearer description of the chosen test cases backed up with proper references.

I believe the quality of this preprint is very close to being suitable for publication; however, I would still recommend a couple of "major" improvements beforehand.

First, I find that Section 2.2 still suffers from vagueness; and contains superfluous information that would be better merged with the Appendix "workflow": from l. 154 to line 168, not a lot is said that is going to help in the understanding of the rest of the paper. Instead, and in preparation for the discussion of Section 4.3, I think this would be a good place to provide more details about the model's implementation. For instance, the sentence starting l.150 "We designed the grids and kernels of CUAS-MPI. . . " should be elaborated upon. Can you provide more details about these "grids" that you design ? Is it an original storage feature or does it rely entirely on PETSc? In the same fashion, the paragraph starting l. 177 could be more descriptive. You could also be more technical about what "CUAS-MPI system kernels" versus "CUAS-MPI system matrix" contain, and what PETSc feature they rely upon/use.

We now explicitly mention the PETSc features we use: Distributed Array, Matrix & Vector. However, we do not consider it appropriate to go into more detail. This paper is intended to present the parallel implementation of the confined-unconfined aquifer system model. It is not intended to serve as PETSc tutorial. For readers familiar with PETSc and interested in these details, the source code, which is part of the submission package, provides the best explanation.

Indeed, being more descriptive in what you coded versus what PETSc is eventually in charge of would help make the analysis of Section 4.3 less generic: while I appreciate that CUAS-MPI is "an entirely different model simulating another compartment of the Earth system (than what is presented in Fischler et al., 2022)", I strongly disagree when the authors state that conclusions drawn from this previous paper do not apply here –unless they can provide a more detailed analysis of the culprit behind the observed throughput and scalability limits. An analysis culminating with the three paragraphs starting l. 347, incriminating the loss of performance "to the increasing communication overhead between MPI processes and decreasing computation performed on each MPI process" has to, at the very least, acknowledge similitudes with that previous work (Fischler et al., 2022) - which most likely also employed the same PETSc features and matrix/vector formatting. If anything else, to reinforce the statement made in Section 5 about how much improvements to the external libraries that you rely upon would improve your throughput in a multi-code environment (ie, ice sheet+hydrology)!

Indeed, the structure of performance studies is often similar. However, CUAS-MPI and ISSM are different codes with different properties, e.g.:

- Discretization: 2D structured grid versus 3D unstructured grid
- Physical model: Groundwater flow equation versus Stokes equation
- Numerics: Finite differences versus finite element

CUAS-MPI and ISSM both "use PETSc", but they use different PETSc data structures of this extensive software toolkit. As a result, they also have different bottlenecks: In ISSM we have seen the matrix assembly as the most significant scaling bottleneck, while in CUAS-MPI it is the grid boundary exchange. A second major difference is the fact that the linear solver is negligible in the performance of ISSM, but in CUAS-MPI it is of significant cost.

In this paper, we present a new simulation tool and the scaling study is intended to help scientists to determine which experiments they can afford with the CUAS-MPI implementation. While we do not compare scaling issues of CUAS-MPI

versus ISSM, we do want to set the throughput of CUAS-MPI in relation to that of ISSM as both simulations might be used to run joint ice shield and hydrology setups.

The comparison of the scaling of different PETSc implementations may be a useful study, but it is not the focus of this work. In particular, if one were to undertake such a study, the factors deciding the choice of software investigated would likely be based on their usage of PETSc features, not their application domain.

2.2 Other

There are also a few small typos/"questions" which require the authors' attention before publication:

- Overall, the equations, citations and sections referencing are messy (lots of misplaced parenthesis, sometimes it is "Eq." sometimes "equation", sometimes "section" sometimes "Sec." ...). Hopefully this is something that the editor will correct prior to publication? But it would be worth checking the Latex source file.
  Thank you. We have put effort into finding all odd usages.

- The first sentence in the abstract is still a bit difficult to follow. I suggest rewriting it ("due to sliding ..." → "through sliding and the location of lakes at the ice margin"? "as well as the ocean circulation" → "it also impacts the ocean circulation ..."?)
  Many thanks. There are four individual effects listed: (1) sliding, (2) location of lakes at the ice sheet margin, (3) freshwater input into the ocean from beneath the ice sheet at the grounding line and (4) freshwater input into the ocean arising from subglacial outflow over land. We have changed 'due to' in 'through' and have added (i)-(iv) to emphasise the different contributions better.

- l.4 "on the ice sheet scale" → "at the ice sheet scale"
  Done.

- l.7 "...the scaling behavior of..." → "the strong scaling"
  We prefer to keep our formulation as we think introducing the term "the strong scaling" without further context is confusing, and just "scaling" is general enough for an abstract.

- l.8 I would suggest completing the sentence with the type of supercomputers that you use
  We now mention the supercomputer we use in the abstract.

- In l. 31 the "inefficient system" is mentioned prior to introducing it. Perhaps you can add a descriptive term in parenthesis? We use the term "inefficient" already in l. 29 (first revised version) and now provide a reference to Fig. 1 additionally.

- l. 39 missing "the" before head?
  We adopt the referee's suggestion.

– l. 44 "..., advanced for seasonal evolution of the hydrological system" is awkward, perhaps the sentence could be rephrased

Indeed! We have rephrased to: "De Fleurian et al. (2014) use a double continuum approach with two different porous layers, one for the distributed system and one for the efficient system, with the second being important for the seasonal evolution of the hydrological system (de Fleurian et al., 2016)."

– l. 46 "... both systems: Sommers et al evolves . . . " → "... both systems. Sommers et al. solve for the water . . . " (splitting the sentence)

Done.

– l. 49 "lead" → "led"

Done.

– l. 50/51 "as in Smith-Johnsen et al. . . . " → "such as in Smith-Johnsen et al. . . . , for example."

Done.

– I would suggest referring the reader to Section 2 for a definition of the greek symbols in the legend of Fig. 1

Thank you. A reference is added as suggested.

– l. 69 "... will benefit if simulations of the subglacial system are available for the same time period." → "... will benefit from simulations of the subglacial system for the same time period."

Done.

– l. 71 perhaps provide a number for "fine resolution" (100-500m?)

– l. 72 "... software that uses parallel computers . . . " → "... softwares capable of using parallel computers . . . "

We adopt the referee's suggestion.

– The title 2.1 is still too short, please be more descriptive

We use "Description of the physical Model"

– l. 107 why not use the acronym CUAS that you have introduced?

We have dropped "used in the confined-unconfined aquifer" altogether.

– The numbers in Fig. 2, which will be very useful for Section 4, should be acknowledged in the caption of the figure

Thank you. We have changed the figure caption accordingly. "The numbers 1–5 in white circles are used again in Sec. 4, e.g. Fig. 8, for cross-referencing. "

– l. 148 sub-glacial → subglacial

Done.

– the sentence starting l. 156 could be dropped altogether
We see it worse to mention and moved it to appendix.

– l. 177 “We” → ”we” (no capital letter)
Done.

– l. 186 please provide a ref for MUMPS
We adopt the referee’s suggestion.

– Fig 5: the exact meaning of “difference” should be given. Is it an absolute error?
Thank you. We agree with the reviewer and change the caption of Fig. 5 to exactly describe the meaning of "difference".
Figure 5. ...(Bottom) The difference between the computed and predicted drawdown, shown in the top panel, indicates an overall small discrepancy between the numerical and theoretical values.

– The sentence starting l. 243 “To be able to compare performance . . . ” is awkward and should be rewritten
We now use: “We run the model for 24 time steps (one model day) using different grid resolutions to compare the model performance and scaling behavior."

– l. 245 “As linear solver, we use GMRES with a Jacobi preconditioner” → “We use the GMRES linear solver with a Jacobi preconditioner”
We adopt the referee’s suggestion.

– l. 257 you talk about GXX runs prior to introducing the notation Thank you. We now state the resolution (1200 m and 2400 m) explicitly here and keep the introduction of the “GXX” in section 4.1.

– The first sentence of Section 4.1 is awkward, consider rephrasing

We changed this to: “The model domain consists of a rectangular area containing the Greenland ice sheet. Grid points of the hydrological system for which Eq. 1 is solved are colored red in Figure 6.”

– l. 294 “sufficiently realistic” → “sufficiently realistically”
We adopt the referee’s suggestion.

– l. 331 “The CUAS-MPI system kernels category, which, like the code categories discussed in the following, is also identified in Fig. 2, includes all kernels running on CUAS-MPI grids.The CUAS-MPI system kernels category contains the characteristics of the EPM, such as the confined-unconfined scheme, transmissivity change and flux, as two dimensional fields.” → “The four remaining functional parts are identified in Fig. 2. In particular, the CUAS-MPI system kernels category includes all kernels running on CUAS-MPI grids (see the blue box in Fig. 2)”. This is just a rewrite suggestion
Thank you. We made the change accordingly.

– l. 337 “the preconditioned GMRES implementation”

We adopt the referee’s suggestion.

– l. 337 “Finally, NetCDF output includes . . . ” → “Finally, the NetCDF output category includes . . . ”

We adopt the referee’s suggestion.

– l. 370 “fix” → “fixed”

Done.

– Figure 10a should feature the same x-axis legend as Figure 9: 96, 192, 384 etc. It is very confusing otherwise. Consider enlarging the police for both axis also, it looks like there is the space to do so. Thank you. That is a good suggestion. We consider applying this change in the final version of the figures.

– The paragraph of l. 375 is redundant and could be dropped

We consider it relevant to make this point, even if it could be deduced from previous data.

– l. 383: “98” → “96”

Done.

– Can you elaborate more on the sentence starting l. 389? What is the cause of needing more iterations and how does that number grow with resolution? That is an interesting issue for further code improvement

This is of course of practical importance for simulation planning, and we elaborate on the implications in the text (l. 390 – 401, first revised version). We don’t think this would be worth another paragraph or figure. Especially as we expect this to be related to the model set-up and not the model implementation. On a higher resolution grid, more and more features of the model domain can be resolved, and thus, the gradients in the geometry might increase (bed elevation, ice thickness) as well as the complexity of the lateral boundary conditions, e.g. some outflow boundary conditions (small fjords) are only resolved on a higher resolution grid. We have not investigated those effects and prefer just to state our findings from the performance measurements we did. The set-ups we have chosen for the scaling analysis are on purpose based on the geometries that would use in science related to hydrology. The drawback is that we do not solve the same physical system using different resolutions as explained above.

– When discussing potential runtime improvements (l. 446-459), it would be interesting to say a word about what you, within your code implementation, could be able to improve: perhaps communication overhead could be avoided by recomputing ghost cells less often? Or could a different linear solver be potentially investigated - still within the PETSc framework? It would bring much to the discussion.

We already state out, that we think, that hybrid parallelism is a major potential, because the grid computation will be more efficient on shared memory, than the current distributed memory + communication. To use this a hybrid parallel solver is necessary as well. We have not elaborated on this yet, but it has potential.

We additionally state out, improvement potential of PETSc: using persistent MPI communication.

– l. 479 I would be careful in saying this

We agree. Given the difficulties in observing the subglacial hydrological system, this sentence probably sounds overconfident. We now write: "In contrast to the Greenland Ice Sheet, the hydrology below the Antarctic ice sheet is expected to be driven by basal water production without additional water from the ice sheet surface draining to the base during the summer season."