

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 Fischler¹, Thomas Kleiner², Christian Bischof¹, Jeremie Schmiedel⁴, Roiy Sayag⁴, Raban Emunds^{1,2}, Lennart Frederik Oestreich^{1,2}, and Angelika Humbert^{2,3}

¹Department of Computer Science, Technical University Darmstadt, Darmstadt, Hesse, Germany

²Alfred-Wegener-Institut, Helmholtz-Zentrum für Polar- und Meeresforschung, Bremerhaven, Bremen, Germany

³Faculty of Geosciences, University of Bremen, Bremen, Germany

⁴Department of Environmental Physics, BIDR, Ben-Gurion University of the Negev, Sde Boker, Israel

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

Copyright statement. ©2023 all rights reserved

1 General Comments

We thank the anonymous reviewer for the helpful comments on the manuscript. We will extend the entire description of the state of the art of hydrology, the description of the hydrology of this model and the description of the modeling of the pumping tests. In return we moved Section 2.3 into appendix. We will incorporate the vast majority of the minor suggestions.

2 Review 1 by an anonymous referee

The authors present a novel HPC modelling tool for subglacial water fluxes, for which they produce a validation and a computational performance assesment on a case set up at the whole Greenland scale. Finally they discuss its perspectives of applications.

10 The paper is clearly organised and it includes a careful study of the parallel behaviour of the developed modelling tool, with a module-wise quantification of the scalability. This point of especially important due to the limitations of applicability of modelling strategies that are related to too long computation times. The potential fallouts of this work are important, and adequately discussed. Nevertheless the manuscript suffers from formal flaws and is sometimes too elliptical – see the general and specific comments below.

15 So I suggest that this manuscript should be accepted for publication in GMD after minor revisions for applying the recommended improvements.

2.1 General comments :

The section titles should be more explicit and informative. For instance CUAS MPI / model / workflow are a bit short as section titles. The english language needs also improvements.

20 We changed the titles to provide more information about the relevant sections and reviewed the text.

2.2 Specific comments:

2.2.1 Introduction

– 1 23 – 24 – 25 : " For simulating large areas, like entire ice sheets in adequate spatial resolution and in a temporal resolution to allow to represent changes on short temporal change, such as seasonal melt water input, an efficient numerical code is indispensable. " Clumsy and unprecises ('adequat' suffers of vagueness — adequat for what ?). To be rephrased.
25 The introduction will be rephrased and extended to also address a number of issues raised by RC2, e.g. "state of the art" for hydrological modelling.

– 1 26 – 30 : OK but it is not said properly – the number of time steps depends on the length of the time step and of the time interval to be simulated. Maybe better to talk of 'time discretization' rather than 'time step'? + english problems :
30 an efficient code.

We adopted the referee's suggestion.

– 1 33 : compute clusters → supercomputers.

We adopt the referee's suggestion.

Add parenthesis around the reference to Balay et al.

35 We adopt the referee's suggestion.

– 1 37 : repetitions (then . . . then)

We rewrote parts of this block.

2.2.2 Section 2

– 1 41 : The title of section 2 should be changed : a simple acronym cannot be a proper section title. At least the full name
40 of the CUAS-MPI tool should be detailed alongside with its nature (e.g. : modeling tool for sub-glacial hydrology).

We adopted the referee's suggestion.

– 1 51 – 52 " Areas with high permeability represent very efficient water transport, while low permeability represents an ineffective water system. " A bit tautological. Please be more precise (e.g., areas with high permeability are associated with higher density of channels or something like that, more descriptive.)

45 We will reformulate this sentence to be more descriptive.

– 1 55 – 56 : " To this end, an unconfined layer is incorporated, capturing the dynamics if the head is falling below the layer thickness allowing for further water drainage. " A little explaining figure (a sketch with both configuration, confined and unconfined layer) would be helpful for the reader and would enhanced the self-consistency of the paper. See also comments of the Appendix A.

50 The focus of this manuscript is about the performance of the new parallel MPI implementation of the often-cited CUAS physics described in (Beyer et al., 2018). We have extended Appendix A but otherwise we think it most appropriate for readers unfamiliar with the model specifics to read Beyer's paper rather than us trying to summarize relevant parts, as the issue of relevance is highly dependent on the reader's background.

– 1 58 : " an evolution equation for the hydraulic head for Darcy flow " : a governing PDE for hydraulic head spatio-temporal variations, analogous to the diffusivity equation for groundwater.

We would like to keep our formulation. The term "evolution equation" is very well established and we would like to preserve the connection to Darcy flow.

– 1 59 – 60 : " A second major equation is describing the change in transmissivity with time based on melting, creep and cavity formation. " This is a constitutive law for the hydraulic transmissivity of the equivalent porous medium, right ?

60 We have only two state variables in the model (head h , and transmissivity T). All other quantities can be derived from them using equations outlined in Beyer et al. (2018). We solve the evolution equation for h using a fully implicit scheme and rely on the linear equation solvers from PETSc for this task. In contrast to this, we use an explicit time step to update T . The evolution of T is thus not an constitutive law — it's our very simple model to account for subglacial channel evolution.

65 – 1 62 : reference to Appendix B would be better placed in the " Software-Design " section (once again the title should be improved). Besides, given its shortness, I think it could be introduced directly in the body of the text.

Appendix B has been included in Section 2.

– 1 62 : Since equations of the Appendix A are refered to explicitly in figure 1, I think that they should be introduced directly in the body of the text.

70 Appendix B has been included in Section 2.

– Figure 1 : How are derived the flux and effective pressure ? Operational links (e.g. : arrows) between the boxes would improve the schematization of the modeling tool I think.

75 This figure is intended to illustrate the physics solved in the model (left side) and how this is solved in principle using high-level pseudo-code (right side). The information about how the flux and the effective pressure is computed (based on hydraulic head, transmissivity, and geometry) is not explicitly stated in Beyer et al. (2018), unfortunately. Nevertheless, we think this information is too detailed to fit the purpose of the figure. We would prefer to keep the figure as is, with slight modifications in the terminology. We will also include the missing equations in the model description.

– 166 : " such as Greenland " : Not specific enough. You mean probably " such as a set up for modelling the sub-glacial flows under the whole Greenland Ice Sheet " ?

80 The sentence has been rephrased accordingly.

– 170 – 71 : " Time stepping parameters are optionally described by command line parameters or a time step file. " What time stepping parameters ? Is this only time step length, or is there any adaptive time step strategy ? In the first case, I don't feel that this sentence is really necessary.

We clarified the description.

85 – 172 : The MPI version uses " the same command line parameters " than the serial one. If you specify it, then you should add here a reference for finding the serial command line parameters.

We specified how the list of parameters can be obtained.

– 174 : " We use the well-known PETSc parallel math library " Here a reference for that tool should be given.

We will add an additional reference to PETSc.

90 – 182 : "if the problem size allows for comparison." Why do you say that? Please explicit what you have in mind.

Reprased the sentence to clarify.

– 186 – 87 : "which provides an uniform interface of the features we require in CUAS-MPI." Unclear, to be rephrased.

We adopt the referee's suggestion.

– 192 – section 2.3 Workflow : this section could be put in Appendix I think.

95 We adopt the referee's suggestion.

2.2.3 Section 3

– 1136 – 137 : I don't understand why the symbol infinity is used here. The boundary conditions are defined at the boundaries, and if using the method of the image wells these bopundaries should be at finite distances of the real well ?

100 Unfortunately, this notation raised confusion. It was meant to say that we have Dirichlet conditions at the southern and western boundaries, while Neumann conditions are applied at the northern and eastern boundaries. We will state this explicitly in the text and provide a figure for the setup.

– 1141 : equation 1 : Always give the units of all used variables and parameters.

We will change the text accordingly

105 – 1121 section 3 : this section should be rewritten by splitting it in two parts, one for the confined case, and one for the unconfined ones. A figure presenting the considered geometry (position of boundaries and of image wells) should be added.

We provide an additional figure showing the setup. Nevertheless we keep the section about constructing the solutions —

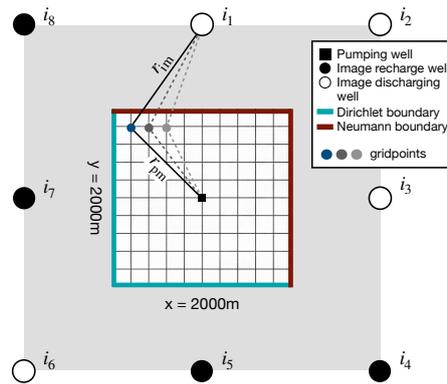


Figure 1. Pumping test configuration, with Dirichlet conditions at the southern and western boundaries, and Neumann conditions at the northern and eastern boundaries. The image well configuration outside of the domain is used for the analytical solutions, which can be extrapolated to achieve the F1962 (Ferris et al., 1962) bounded-low, -med, -high solutions, respectively.

using image wells — rather short, because this is beyond the topic of this manuscript. We already cite Ferris et al. (1962) for the implementation details.

110 2.2.4 Section 4

– 1 164 – 165 : " In case of a land terminating margin, the boundary condition is no flow ". Then rivers coming out of subglacial springs are neglected, right ? Adding a sentence about the consequences of this assumption would be interesting.

115 This is right. Outflow boundary conditions (non-homogeneous Neumann boundary conditions) for rivers coming out of the subglacial environment are neither considered in the setup nor in the model. Discharge through those rivers would lower the hydraulic head (water pressure) in the subglacial system. Outflow could be indirectly simulated by applying Dirichlet conditions for the hydraulic head at those river locations and some other locations, e.g. ice-marginal lakes. The model would be capable of doing so, but to our knowledge, no data set of the river and ice-marginal lake locations for the Greenland Ice Sheet exists, unfortunately. Even though outflow along the land terminating margin is not considered,
120 and this a very strong limitation to the hydrological modelling, we think the setup is realistic enough to draw conclusions about the model performance.

We will include this information in the manuscript.

2.2.5 Section 5

– 1 193 – 194 : " high-throughput simulations ". I don't know the throughput concept. It might be usefull to remind it.

125 In the current manuscript version it is explained in the throughput section. We will adapt to explain it earlier.

- l 198 – 199 : " We employ GCC 11.2, Open MPI 4.1.4, PETSc 3.17.4 and NetCDF 4.7.4 with HDF5 1.8.22 for our performance experiments. " Please give the nature all these software (e.g. : the compiler GCC, the message-passing library OpenMPI 4.1.4, etc)
We added information about all packages we used.
- 130 – l 202 : please give the topology of the network (e.g. : hypercube, fat tree, etc).
We added information about the fat tree topology.
- Figure 3 : please add a scale for the shaded bed topography.
This figure was meant to mainly show the mask. The shading was applied only to give the figure a more three-dimensional appearance. As this information is not relevant for the manuscript, we provide an updated version of the figure without shading.
- 135
- l 213 : " full, half, quarter and one-eighth occupied compute nodes " I think that using populated instead of occupied would be more idiomatic → full, half, quarter and one-eighth populated computing nodes.
We adopt the referee's suggestion.
- l 216-218 : " The size of the circle indicates the hardware investment. The smallest circle indicates that only one node was used, the next size up indicates two nodes, then four nodes, and lastly 16 nodes. " This should also be specified directly on Figure 4.
This will help understanding the figure and we will add the explanation to the capture.
- 140
- l 219 – 224 : I don't understand the runtimes mentionned in the text. For instance 4900 s is stated for 12 threads on 1 node, while on Figure 4 the vertical coordinate of the corresponding point is 140 s. 4900 is not a CPU time, since $140 \text{ s} * 12 \text{ cores} = 1680 \text{ cpu.s}$ (or $140 * 96 = 13440$, if we consider that the whole node is requested for the computation even if only 12 cores are really used), which is different from 4900.
Currently the numbers does not match. We will fix this and apologize for this unawareness.
- 145
- l 225 – 232 : Very interesting discussion, which concludes to the interest of half-populated runs with CUAS-MPI on the considered cluster. I think that the CPU times associated to the slurm requests (→ twice bigger for a half-populated run compared to the fully populated one for a given number of MPI processes - if the run times are equal !) should be also discussed.
We will extend our discussion on this topic.
- 150
- l 234 : I think that 'parts' would be a better word choice than 'categories'.
We adopt the referee's suggestion.
- 155 – l 249 : " which writes a single output of configuration "large" " : please say here what is the total size (in bytes) of the outputs written in each case by NetCDF.
We will at this, as it will help to get the context correctly.

- 1 254 : " scales up to 2304 MPI processes " : according to Figure 5 I would rather say that it scales up quasi-linearly ut to 1536 MPI processes.
- 160 We adapted the text to this comment.
- 1 259 – 262 : " In particular, we see approximately linear scaling also for the "CUAS-MPI system kernels" and the system matrix routine up to at least 768 MPI processes. Then the "CUAS-MPI system kernels" and thereafter the system matrix creation reach their scaling limit and their respective runtime increases." This is correct for CUAS-MPI system kernels, but not for the system matrix creation, which well scales up to 3072 MPI processes.
- 165 We adapted to a more precise description.
- 1 264 : " grid data exchange communication caused by PETSc " : I think that this term is misleading, and so is the associated legend in Figure 6. The problem as far as I understood is not the kernel computations, but the communication to PETSc of the results obtained by these kernel computations, so that the PETSc linear solver can use them for solving the diffusivity equation governing the sub-glacial flow. So I would keep 'kernel computations' for the green curve in Figure 6, but use 'kernel communications to PETSc' instead of 'PETSc communication' for the pink one.
- 170 The CUAS kernel operates on data structures from PETSc. So in performing the kernel computation, functions from the PETSc library are called, for example, to handle border exchanges. We follow the reviewer's suggestion and use the term "kernel communication" in the legends, and will clarify this issue in the text.
- 1 267 – 268 : " In our studies of throughput, i.e. how many simulated system years we can run in a day of compute time (simulated years per day, SYPD). " I didn't know this concept of throughput as a measure of result production rate – which does not mean that it does not already exist in the literature! But still I think that because it has been already used earlier in the manuscript (even in the abstract), its definition should have been explicated (or reminded) earlier as well.
- 175 We will place this earlier.
- 1 278 : The observation of a significant impact of I/O rate on computation time is interesting. This question has for instance also been investigated in another porous media related context in Orgogozo et al., 2014 (<http://dx.doi.org/10.1016/j.cpc.2014.08.004> see section 4.4).
- 180 The observation of Orgogozo et al. is similiar to our findings. We will consider it.
- Figure 7 : the letters a, b and c used in the caption are missing on the figure.
- We will add letters to other Figures to and emphasize it in this figure.
- 185 – 1 284 – 288 : This paragraph is hard to follow. I don't think that one can say that larger grids need more cells per MPI process than smaller grids for efficient parallelism. While using larger grid, one increases the computational load and thus the computation over communication ratio, which leads to better parallel efficiency for a given number of MPI processes.
- Our statement is related to the fact, that the graph uses "cells per MPI process" on the x-axis. We will clarify this.

190 And also I don't understand what is the sweet spot.

In this context, we assume that the sweet spot is the minimum runtime and will clarify this in the text.

– 1 291 – 292 : "The minimum of the total runtime per iteration is increasing with spatial resolution". Please discuss why?
We will explain it. It is related to the metric "cells per MPI process".

195 – 1 292-293 : " the spread is less than the total runtime per year, showing that the total runtime is driven by the increase in number of linear solver iterations with resolution. " : Hard to follow. Does panel b of Figure 7 present the total runtime per year ? If yes say it explicitly in the caption of the figure and refer explicitly to this figure in that discussion.
We rephrased the description in order to clarify the issue.

2.2.6 Section 6

– 1 295 : "but potentially not highest resolution." Unclear.

200 Currently the highest resolution is 150m as this is the current bedrock resolution. We will add this information in the text.

– 1 297 : " the amount of core-hours usually available for such runs. " Be more precise. What is the target of CPU hours for this kind of run, and why ?

205 This sentence is unfortunately not precise, and we would drop the second part of the sentence: "Although ...limits the number of years that can be simulated." The number of core-hours strongly depends on the scientific question. If someone wants to run the hydrology model as part of an Earth-System-Model over decadal to millennial time scales, the computational resources spent just for the subglacial hydrology would probably be less, than for a detailed simulation with very high resolution (time and space) for e.g., hydropower plants.

– 1 297-298 : " A simulation covering the 90 years from 2010 to 2100 in 600 m spatial and 1 hour temporal resolution requires a wall-clock time of 1350 hours (56 days) on 384 MPI processes. " Please precise again here that this quantification is obtained for the Lichtenberg HPC system.

We clarified it at this point.

One could discuss how it could be extrapolated to larger supercomputers, for instance European level ('tier-0') ones.

215 The data show that already on the Lichtenberg system, CUAS-MPI reaches its scaling limits, so using even bigger computers is not a productive use of such resources.

– Section 6 : english language problems ('might of interest', 'now have now', ...)

We checked the entire section again.

– 1 322 : G250 is too cryptic ; either add a more informative name or no name for this resolution.

We will add a brief description to G250 to give the reader context.

220 – 1 326-333 : This discussion is a bit confusing, to be rewritten more clearly.

We adopt the referee's suggestion.

Add also physics-based arguments for the statement that there is no necessity of high frequent data exchange.

225 The question on how often to couple both models depends on the phenomena that are investigated in both models. If someone for instance wants to study the effect of ocean tides on the subglacial environment, time steps in the hydrology model needs to resolve the tides. If the ice sheet model can afford the same time steps in terms of computational costs, why not? But in most cases, the ice sheet model will be computationally more expensive and a compromise must be found (e.g., daily coupling only). We will rephrase this sentence.

– 1 345 : " computational granularity " : please remind briefly the definition of this concept here

We will add an brief explanation: granularity is a measure of the amount of work which is performed by a task.

230 2.3 Appendix A :

We will incorporate Appendix A into the main text as suggested by both reviewers. This part will also be extended slightly to address a number of issues listed below. We expect that this will also improve the clarity in the validation section.

– Psi and b should be introduced in a little explaining figure as proposed above (see also the comment of lines 55–56).

235 Both reviewers ask for more information about the meaning of quantities in the hydrological system. We will briefly state, what every quantity is referring to.

– The units of all the used variables and parameters should be given.

Yes, we fully agree.

– 1 376 : " To allow a smooth transition between the confined and unconfined system, a range d is introduced. " This parameter should be a bit more discussed. What is its physical signification? What is its range of possible values?

240 The parameter d , with $0 \leq d \leq \Psi$ (unit: m), allows for a gradual transition of the effective storativity, S_e in the confined-unconfined-transition scheme (Eq. A2 in the manuscript and Ehlig and Halepaska (1976, Eqns. 6, 7 and Fig. 2). We will move this sentence and the additional information below Eq. A2. In all the simulations used here we use $d = 0$ as in (Beyer et al., 2018, Tab. 3).

– 1 382 : The difference between S_s and S_y should be more explicited.

245 We will give additional information about all hydrological quantities including S_s and S_y .

– 1 387 : There is surely (at least) one publication associated with this complex equation ; please refer to it here.

We cite Ehlig and Halepaska (1976) as this is the foundation for the confined-unconfined scheme used in Beyer et al. (2018) and implemented in CUAS-MPI.

250 – 1 392 : there should be a link between h and p_w ? How π is computed, or from where does it come from if it is a forcing ? Please explicit these point here.

The hydraulic head is linked to the water pressure via $h = p_w/(\rho_w g) + b$, with the density of water $\rho_w = 1000 \text{ kg m}^{-3}$, the acceleration due to gravity $g = 9.81 \text{ m s}^{-2}$, and bed elevation b (unit: m). The ice overburden pressure $p_i = \rho_w g H$, with the bulk density of ice $\rho_i = 910 \text{ kg m}^{-3}$ and the ice thickness H (unit: m). The ice thickness is taken from the BedMachine Greenland dataset (Morlighem, 2021; Morlighem et al., 2017, Version 4) as part of the model input.

255 2.4 Appendix B :

- 1 399 – 400 : ‘equidistant’ ; may be regular would be more appropriate ? And why using rectangular grid cells instead of square ones ?

We adopt the referee’s suggestion.

- 1 405 – 406 : " If an iterative solver is used, convergence is decided by the decrease of the residual norm relative to the norm of the right hand side (rtol) and the absolute size of the residual norm (atol). " Here an equation allowing to precisely identify the quantity rtol and atol should be provided. Typical ranges to be used should also be given.

260 We don’t think, that typical values for convergence values exist. They usually dependent on the physical problem to be solved, the boundary conditions and also the initial conditions in a specific setup. We state the values that we have used already in the main text and we will provide brief information about the tolerances.

- Beyer, S., Kleiner, T., Aizinger, V., Rückamp, M., and Humbert, A.: A confined–unconfined aquifer model for subglacial hydrology and its application to the Northeast Greenland Ice Stream, *The Cryosphere*, 12, 3931–3947, <https://doi.org/10.5194/tc-12-3931-2018>, 2018.
- Ehlig, C. and Halepaska, J. C.: A numerical study of confined-unconfined aquifers including effects of delayed yield and leakage, *Water Resources Research*, 12, 1175–1183, <https://doi.org/10.1029/WR012i006p01175>, 1976.
- 270 Ferris, J. G., Knowles, D. B., Brown, R. H., and Stallman, R. W.: *Theory of Aquifer Tests*, Tech. rep., U.S. Government Print. Office, <https://doi.org/10.3133/wsp1536E>, 1962.
- Morlighem, M., Williams, C. N., Rignot, E., An, L., Arndt, J. E., Bamber, J. L., Catania, G., Chauché, N., Dowdeswell, J. A., Dorschel, B., Fenty, I., Hogan, K., Howat, I., Hubbard, A., Jakobsson, M., Jordan, T. M., Kjeldsen, K. K., Millan, R., Mayer, L., Mouginot, J., Noël, B. P. Y., O’Cofaigh, C., Palmer, S., Rysgaard, S., Seroussi, H., Siegert, M. J., Slabon, P., Straneo, F., van den Broeke, M. R., Weinrebe, 275 W., Wood, M., and Zinglensen, K. B.: *BedMachine v3: Complete Bed Topography and Ocean Bathymetry Mapping of Greenland From Multibeam Echo Sounding Combined With Mass Conservation*, *Geophysical Research Letters*, 44, <https://doi.org/10.1002/2017gl074954>, 2017.
- Morlighem, M. e. a.: *IceBridge BedMachine Greenland, Version 4*, <https://doi.org/10.5067/VLJ5YXKCGXO>, 2021.