

Response to Reviewers

Title: High-performance coupled surface-subsurface flow simulation with SERGHEI-SWE-RE
Authors: Na Zheng, Zhi Li, Gregor Rickert, Mario Morales-Hernández, Ilhan Özgen-Xian, Daniel Caviedes-Voullième.

Response to Reviewer #2

The paper presents SERGHEI-SWE-RE, a coupled surface–subsurface model built on the Kokkos framework that targets performance portability across CPUs and GPUs. The work is clearly within GMD’s scope, and I found the asynchronous coupling strategy and the multi-platform scalability analysis to be the most interesting contributions. The benchmark suite covers a reasonable range of conditions, and the comparisons with ParFlow, CATHY, HGS and others are generally convincing. That said, I have a few concerns I would like the authors to address. Two of them are substantive and touch on aspects of the model that I think are currently under-described; the third is a straightforward request for numbers that the authors say they have but chose not to show. None of these require new simulations. My recommendation is minor revision, but I want to be clear that I expect concrete answers rather than a restatement of what is already in the paper.

Major comments

Comment 1

My main technical concern is with Section 2.2. The four-case boundary condition scheme is described clearly enough, but the manuscript is completely silent on what happens numerically when the system moves between these cases, for instance when a dry cell receives exfiltration and becomes wet, or when the top subsurface layer saturates and the pressure head approaches the surface elevation. These are precisely the transitions that are known to cause convergence problems in this class of models. Coon et al. (2020) document this explicitly for the ATS model, noting that convergence degrades near the dry-to-ponding transition, and Furman (2008) gives a broader review of the numerical difficulties associated with boundary condition switching. The present manuscript does not acknowledge any of this.

I should also point out that more elegant solutions to this problem exist in the literature and are not cited. Casulli (2017) derived a coupled surface–subsurface formulation in which the governing equations for both domains are solved simultaneously using the NCZ algorithm, the same algorithm on which SERGHEI-RE is based, entirely avoiding the switching problem. Earlier, Casulli (2009) introduced the idea of an additional computational node at the soil surface governed by the equation state $H(\psi)$, which allows rainfall to be prescribed as a permanent Neumann boundary condition regardless of ponding state. Tubini and Rigon (2022) implemented exactly this approach in WHETGEO-1D and showed that it handles both infiltration-excess and saturation-excess ponding without any case logic or numerical oscillations. I am not asking the authors to rewrite their code. I am asking them to engage with this literature and explain, in the paper, whether their four-case scheme has been tested under conditions of rapid state transitions and what its behaviour is. If fully saturated conditions with return flow are never actually reached in any of the five benchmark cases,

that should be stated explicitly, because right now the reader cannot tell.

Response:

We apologize for not clearly explaining the coupling strategy of SERGHEI-SWE-RE in the original manuscript. SERGHEI adopts a modular structure, with SWE and RE modules sequentially (or externally) coupled. This is fundamentally different from the fully coupled approaches of Casulli (2017); Coon et al. (2020). In a sequentially coupled framework, the subsurface and surface domains are solved separately, with the surface state serving as a boundary condition for the subsurface solver. As a result, the model does not encounter the same type of nonlinear convergence difficulties during dry-to-ponding transitions reported in Coon et al. (2020). Moreover, in SERGHEI, the surface module uses an explicit solver, while the subsurface module uses a predictor-corrector method. Neither of these methods relies on internal iterative schemes within a time step, which inherently avoids the convergence failure loops commonly associated with boundary switching in fully implicit systems, as discussed in Furman (2008). Our Case 3 and Case 5 both involve saturation transition and dry-to-ponding transition. The simulation results (Section 3.3 and 3.5) show that these processes are modeled reasonably well with our four-case boundary condition switching scheme.

In the revised manuscript, we have added a paragraph in Section 2.2 to discuss our coupling strategy in relation to the fully coupled approaches suggested by the reviewer. We highlight that the sequential strategy was chosen to preserve modularity and support asynchronous time stepping.

--- **Revisions in Main Manuscript (Section 2.2)** -----

The coupling strategy of SERGHEI-SWE-RE follows an externally coupled, sequential (non-iterative) approach. This means, both solvers are solved independently, and they exchange states and fluxes in sequence, without iterating during a coupling time step.

...

It should be noted that sequential coupling, which relies on explicit boundary condition switching, has historically been reported to cause numerical difficulties during rapid state transitions (Furman, 2008). A prominent alternative is to use a fully coupled scheme that eliminates the need for switching. While some fully coupled approaches still face convergence degradation near dry-to-ponding transitions (Coon et al., 2020), recent advances have proposed elegant formulations to bypass these issues completely (e.g., Casulli, 2017; Tubini and Rigon, 2022). Despite these advances, SERGHEI-SWE-RE deliberately uses a sequential coupling scheme to maximize flexibility, as the framework is designed to be highly modular and open-source. Crucially, this externally coupled structure is a prerequisite for supporting asynchronous time-stepping. Furthermore, SERGHEI adopts a non-iterative predictor-corrector scheme for subsurface flow, which inherently avoids the nonlinear convergence failure loops common to popular iterative methods. As demonstrated in Section 3, for benchmark tests involving boundary condition switching (Section 3.3 and 3.5), SERGHEI-SWE-RE successfully navigates these transitions without generating numerical oscillations, yielding solutions that agree well with established reference data.

Comment 2

The introduction lists modularisation and flexibility as key advantages of SERGHEI-SWE-RE. I agree the internal modularisation between the SWE and RE solvers is well demonstrated, but the paper says nothing about whether the model can be coupled with external codes. This matters because the community has largely converged on standard interfaces for this purpose, the Basic Model Interface (BMI; Hutton et al., 2020) being the most widely adopted, and models that do not address interoperability are increasingly difficult to integrate into larger workflows. WHETGEO (Tubini and Rigon, 2022) does this through OMS3. The Next-Generation Water Resources Modeling Framework uses BMI natively. The paper should state clearly whether SERGHEI-SWE-RE supports any such interface, and if not, acknowledge this as a current limitation rather than implying that flexibility is already there.

Response:

The reviewer makes a fair point. The modularisation strategy within SERGHEI is, so far, internal. That is, the framework is designed in such a way that the different modules (SWE, RE, LPT, ADE, ST) are interoperable, but are not required. This is the flexibility we implicitly refer to.

The reviewer is addressing what we would call "external" modularisation, i.e., using SERGHEI as a module in broader frameworks/workflows, or uptaking other modules from elsewhere. This is a long term objective in SERGHEI, but currently not being addressed. We appreciate highly the reference to community interfaces.

We have addressed this distinction in the manuscript, and acknowledged that the external modularisation the reviewer rightly points to is not yet available in SERGHEI.

--- Revisions in Main Manuscript (Section 2.2) ---

It must be acknowledged that the modularization of SERGHEI is currently limited to the internal framework. At present, SERGHEI does not natively support standard community interfaces for external coupling, such as the Basic Model Interface (BMI) or OMS3 (Hutton et al., 2020; Tubini and Rigon, 2022; Peckham et al., 2013). Establishing external interoperability through such standardized interfaces remains a long-term development objective for the SERGHEI framework.

Comment 3

A related point: the consistent domain decomposition (same MPI partition for SWE and RE) is already shown in Section 4 to be a bottleneck at high GPU counts. This is an inherent consequence of tying the two solvers' parallelisation together, and it would also complicate any future attempt to expose an external coupling interface. The authors discuss this briefly but do not connect it to the broader architectural question of whether the design supports or hinders future extension. A paragraph addressing this would strengthen the Discussion considerably.

Response:

We agree. The high level parallelisation and domain decomposition structure between the SWE and RE modules currently implemented in SERGHEI is not very flexible. From our own results, it becomes apparent that load imbalances may hinder the efficiency of the coupled model. Indeed, this clearly points to a more flexible solution.

The initial reasoning was to favour a trivial mapping from surface cells to the top layer of subsurface cells, so that no MPI communication would be required between the surface (SWE) and subsurface (RE) modules, and the pressure and mass exchange would happen inside a computational device (in particular, in the same GPU).

We do intend to move away into a more flexible approach allowing:

1. different horizontal resolution between the surface and subsurface grids, limiting it to nested grids. This impacts the load balance, of course.
2. different horizontal domain decomposition among the surface and subsurface domain (as the reviewer implies in this comment). This requires the definition of a halo layer to map the surface/subsurface states for the exchange, and is not in contradiction with the first point.

At this point we have not engaged in the implementation of these features for two reasons. First, the initial objective, as reported in this manuscript, is the implementation of the coupling strategy, in a first iteration. Second, as part of the broader development of SERGHEI, we intend in general to deal with domain decomposition in more flexible ways, even for the uncoupled modules. We have revised Section 4.1 to explain our choices and future directions of model development.

--- **Revisions in Main Manuscript (Section 4.1)** -----

Consequently, the optimal parallelization configuration for the RE solver becomes inefficient for the SWE solver due to significant disparities in the computational workloads. It indicates that the simple consistent surface-subsurface cell mapping and domain decomposition strategy (Fig. 3 and 4) might be sub-optimal for large-scale parallelization. ~~Future enhancements to SERGHEI-SWE-RE should incorporate either (i) different mesh resolutions for the SWE and RE solvers or (ii) workload-based domain decomposition schemes to improve scaling performance.~~

This scaling bottleneck highlights a broader architectural limitation in SERGHEI-SWE-RE. The consistent domain decomposition was initially chosen to provide a simple one-to-one mapping between surface and subsurface cells, allowing efficient local data exchange without additional inter-module MPI communication. It limits scalability when their optimal configurations differ and hinders future extensibility, including the potential integration of external coupling interfaces. To address these load imbalances and support broader interoperability, future versions will transition to a flexible exchange architecture that supports (i) different horizontal mesh resolutions and (ii) independent, workload-based domain decompositions.

Comment 4

Lines 173–176 state that "further quantitative examination (not shown) confirms" that the C-property is satisfied. This is not acceptable in a model description paper submitted to GMD. The C-property is a fundamental numerical requirement and its verification should be documented with actual numbers, the maximum absolute deviation from the initial condition is the standard metric and takes one line to report. The authors clearly have these numbers. They should include them, either in the text or in a supplementary table. Reviewer 1 made the same point and I fully agree.

Response:

We sincerely apologize for the vague statement. We have now added the detailed simulation results to the revised manuscript to illustrate that the C-property is satisfied.

--- Revisions in Main Manuscript (Section 3.1) -----

~~A quantitative examination (not shown) further confirms that both the free surface elevation and the subsurface pressure fields maintain their initial states during the entire simulation as expected, indicating that the C-property is satisfied.~~ Throughout the entire simulation, the maximum absolute deviation of the surface water level from its initial condition remains zero. Fig. 5 shows the distributions of surface water level and water table along the transect at $y = 90$ m at $T = 0, 500$, and 1000 s, demonstrating that both the free-surface elevation and subsurface pressure fields are preserved at their initial states. These results confirm that the C-property is satisfied to machine precision.




Figure 5. The simulated free surface and groundwater table elevations for Case 1 at (a) $T=0$ s, (b) $T=500$ s, (c) $T=1000$ s on the transect of $y = 90$ m.

Minor comments

Comment 5

The rainfall routing assumption (Section 2.2, lines 140–143) is validated only against Case 4, which was designed with no ponding. This is not a meaningful test of the assumption. The authors should acknowledge the scenarios where routing rain through the surface first introduces a numerical lag, the obvious case being a very dry soil with high infiltration capacity.

Response:

We thank the reviewer for this insightful comment. We realize that our original manuscript did not explain

the rainfall-routing algorithm clearly enough. In the SWE solver, rainfall is implemented as a source term in the continuity equation and does not directly affect the momentum equation within the same surface time step. As a result, under synchronous coupling, rainfall introduced during a given surface step is not converted into immediate runoff before the updated surface state is used to drive infiltration in the RE solver. In this sense, the numerical sequencing does not necessarily translate into a physical lag or a model discrepancy. A limitation arises when asynchronous coupling is activated, because the RE solver is not advanced at every surface time step. In this case, runoff may occur before infiltration is updated. This behavior is examined in the benchmark tests, and Case 5 in particular evaluates the asynchronous coupling strategy (with ponding). The results show good agreement with the reference solutions and only a small difference associated with asynchronous coupling for the tested conditions. We have therefore revised Section 2.2 to explain the coupling algorithm more clearly and to discuss this limitation explicitly.

--- Revisions in Main Manuscript (Section 2.2) -----

A special case is the rainfall condition. When the ground surface is dry, rainfall is typically a flux boundary condition to the subsurface domain. In SERGHEI-SWE-RE, for simplicity, it is assumed that rainfall is always applied to the surface domain, accumulates as surface ponding, then infiltration is computed following the above procedure. This avoids the need to handle rainfall in the subsurface solver. *It should be noted that numerically, rainfall acts as a source term in the SWE continuity equation. This source term is added after the depth and flux are solved, meaning that it does not affect the SWE momentum equation within the same time step. This ensures that rainfall cannot generate surface runoff before the updated water depth is passed to the RE solver to drive infiltration, indicating that this treatment does not translate into an unphysical lag between rainfall and infiltration.* ~~The validity of this approach will be demonstrated in Section 3.4.~~

...

An associated issue with asynchronous rainfall-runoff simulation is that because the RE solver is not executed at every surface time step under asynchronous coupling, and that rainfall is applied to the surface domain first, rainfall may accumulate and trigger surface runoff prior to the computation of infiltration. However, it will be shown in Section 3.5 that when $\Delta t_{re,max}$ is set to roughly an order of magnitude larger than the surface time step, this potential discrepancy remains negligible.

Comment 6

On the asynchronous coupling: Case 5 tests two values of $\Delta t_{sub,max}$ and both give satisfactory results for that specific benchmark. But the paper offers no guidance on how to choose this parameter in general. Users running scenarios with fast-moving wetting fronts or rapidly alternating wet/dry conditions will need some basis for this choice. Even a qualitative discussion would help.

Response:

We agree that providing guidance on choosing $\Delta t_{re,max}$ is highly valuable for users of SERGHEI-SWE-RE (Note: We have updated the notation from $\Delta t_{sub,max}$ to $\Delta t_{re,max}$ to better align with the RE solver

naming convention). Unfortunately, since the universally optimal value for $\Delta t_{re,max}$ depends heavily on problem-specific parameters, such as soil permeability, top layer thickness, and rainfall intensity, providing a strict quantitative formula is challenging. In the revised manuscript, we have expanded the analysis of Case 5 to vary $\Delta t_{re,max}$ from 1 s to 4 s, resulting in negligible impact on the model results. Drawing from this and our broader experience, for typical rainfall-runoff simulations, we believe $\Delta t_{re,max}$ can be safely set to roughly an order of magnitude larger than the surface flow Δt_{swe} . We have now explicitly included this heuristic guideline in Section 2.2 as a recommended starting point. However, we also found that for rapidly varying surface dynamics, such as very high frequency tidal oscillations in Li et al. (2023), variations in $\Delta t_{re,max}$ have significant effects on simulation results. We have revised the Conclusions section to discuss the physical factors (like high K_s and rapid wetting cycles) that necessitate tighter coupling.

--- **Revisions in Main Manuscript (Section 2.2)** -----

Asynchronous coupling significantly enhances computational efficiency by reducing the execution times of the RE module. However, it introduces a time lag of the subsurface module as illustrated in Fig. 3. This time lag inevitably results in coupling errors that depend on both the time scale difference between surface and subsurface flow and the lag duration, an issue extensively discussed in Li et al. (2023) and later in Section 3.5.

...

The selection of $\Delta t_{re,max}$ should be guided by the characteristic time scales of the surface and subsurface dynamics. For environments with highly permeable soils, thin top subsurface layers, or scenarios featuring rapidly varying wetting and drying cycles, a smaller $\Delta t_{re,max}$ is required to resolve fast-moving wetting fronts. Conversely, for more stable rainfall-runoff simulations, empirical testing (Section 3.5) indicates that $\Delta t_{re,max}$ can often be safely set to roughly an order of magnitude larger than the average surface time step without a noticeable loss of accuracy.

--- **Revisions in Main Manuscript (Section 3.5)** -----

In synchronous coupling approach, the surface and subsurface solvers are advanced with the same time step. The surface time step ($\Delta t_{sur}\Delta t_{swe}$) typically ranges from 0.4 to 0.55 s for this case. In asynchronous coupling simulations, the maximum time step of the RE solver ($\Delta t_{sub,max}\Delta t_{re,max}$) is set to 1 s, 2 s, 3 s and 4 s, with 4 s roughly an order of magnitude larger than the SWE solver time steps. Due to the high degree of overlap among the results, only the $\Delta t_{re,max} = 4s$ case is presented for clarity.

--- **Revisions in Main Manuscript (Section 5)** -----

Our tests demonstrate that the asynchronous coupling strategy significantly improves computational efficiency, though it inherently introduces a temporal lag in surface-subsurface exchange. The choice of the coupling time step must balance computational gains with model discrepancies. Future work will comprehensively explore this trade-off and provide more quantitative criteria for optimizing time-stepping for the coupled simulation.

Comment 7

The scalability breakdown in Section 4 isolates SWE, RE, and MPI communication, but never the coupling overhead itself (halo transfer, four-case logic). Given that the paper’s main claim includes efficiency, it would be worth confirming that this overhead is indeed negligible.

Response:

We would like to clarify that there is no halo transfer in the current implementation of the coupling strategy. This was intended, as a design choice, as discussed in response to the reviewer’s comment . We sincerely apologize for the misunderstanding caused by the schematic representation in Figure 1. We have redrawn Figure 1 and updated its corresponding description in Section 2.2 of the revised manuscript. The operations required for the coupling are simply fetching depth from surface arrays, imposing it as pressure boundary conditions in the subsurface arrays, and adding/subtracting a water column from the surface arrays. Since they all occur within a single computational unit and memory space, the overhead is negligible.

To prove this, we performed additional tests in Section 4.3, where the same test problem is modeled with the RE module alone, and with the coupled SWE-RE model. A comparison of the computational cost and parallel scalability indicate that the overhead caused by surface-subsurface data transfer is indeed negligible.

Revisions in Main Manuscript (Section 2.2)

Spatially, the surface and subsurface domains share the same horizontal discretization, both using structured Cartesian grids with uniform grid spacing (Fig. 1). After solving the SWE, the surface water depth (h_s) is mapped to the ghost cells of the subsurface grid (gray cells in Fig. 1) to establish the upper boundary condition for the RE solver. Upon completion of the RE solution, the SSE flux (q_{ss}) is evaluated from top-layer and ghost cell pressure heads (h_g), then passed back to the surface module. Finally, the surface water depth is updated by adding an equivalent depth calculated by integrating q_{ss} over the coupling time step.




Figure 1. The schematic of data exchange between subdomains.

Revisions in Main Manuscript (Section 4.3)

The coupling between the SWE and RE modules may introduce additional computational overhead that affects scalability. An additional scaling test is conducted to compare the performance of the coupled model and the standalone SERGHEI-RE module. The problem setup is identical to the lateral flow test (Section 3.4), in which rainfall infiltrates into the subsurface without causing surface runoff. Section 3.4 has demonstrated that SERGHEI-RE and SERGHEI-SWE-RE

produce indistinguishable water table evolution. However, in SERGHEI-RE, rainfall is applied directly to the subsurface domain as a flux boundary condition, whereas in SERGHEI-SWE-RE, rainfall is first applied to the surface domain, inducing ponding before infiltrating into the subsurface through surface–subsurface exchange computations. To understand the scaling behaviors of these two approaches, the original 3D model domain is used, which spans 4000 m \times 5000 m in the horizontal directions and 15 m in the vertical direction. The domain is discretized with uniform grid spacings of $\Delta x = \Delta y = 5m$ and $\Delta z = 1.5m$, resulting in a total of 8 million grid cells. Fig. 20 shows the simulation times and speedups of the two approaches as the number of CPU threads increases. The SERGHEI-SWE-RE and SERGHEI-RE models exhibit negligible difference in terms of both the computational time and the speedup. This behavior can be attributed to the relatively small number of surface grids cells (as discussed in Section 4.1) and the absence of surface runoff generation in this case. These results indicate that, for this particular test case, surface–subsurface coupling does not degrade model scalability.

It should be noted, however, that this conclusion applies only to this specific test scenario. In a broader context, surface-subsurface coupling is not expected to alter the scalability of the subsurface module because the SERGHEI-RE solver treats surface flow as a boundary condition (Section 2), which must be defined regardless whether SERGHEI-SWE is activated or not. Conversely, the surface module performs an additional parallel loop over all surface grid cells to incorporate the exchange flux as a source/sink term in the mass conservation equation. Thus, if SERGHEI-SWE scales poorly (as shown in Section 4.1 on JUWELS), surface-subsurface coupling is expected to further deteriorate the overall scaling performance.

From an HPC perspective, what is considered an overhead may vary. Arguably, inefficiencies due to load imbalance can be viewed as an overhead, e.g., attempting to use a resource set which provides high parallel efficiency for both solvers would likely result in a longer RE runtime. Whether this is an overhead or an inefficiency is a gray area. In any case, this can of course be alleviated by having different sets of resources for each solver (which is not explored in this paper). This in turn links to an undeniable overhead: the exchange of information between the solvers. When both solvers are mapped to the same hardware (i.e., with the same horizontal domain decomposition) as in these tests, this overhead is negligible. However, if states and fluxes must be exchanged through MPI across different resources (CPUs or GPUs), this overhead will grow.

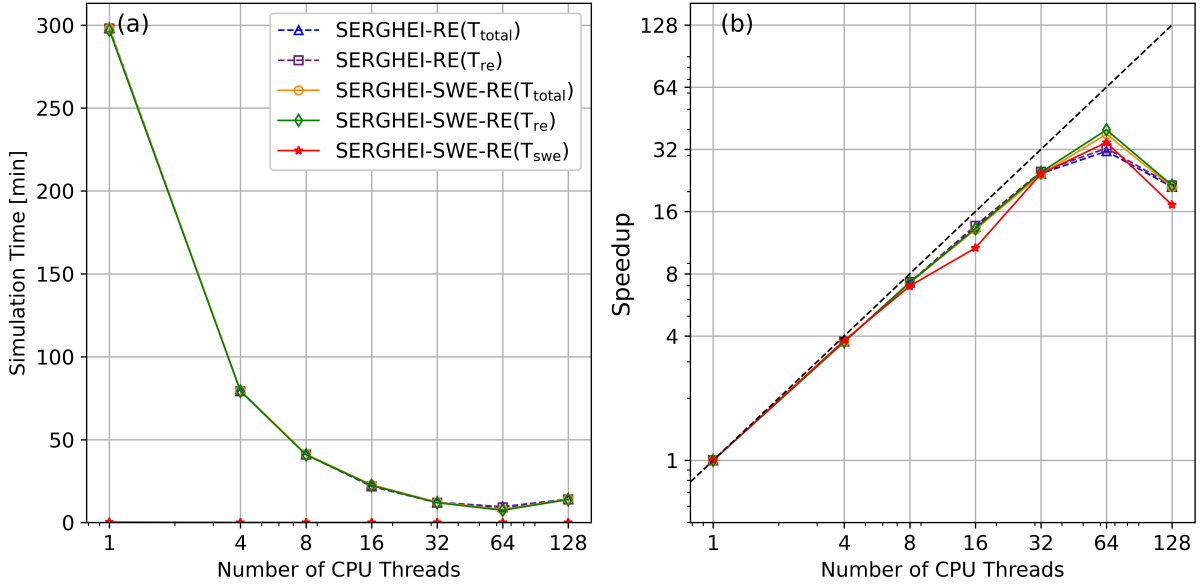


Figure 20. (a) simulation time and (b) speedup for the Case 4 (3D domain). The black dashed line represents the ideal scaling

Comment 8

Figure 3 needs a clearer caption. The meaning of the tick mark spacing and the shaded tsub interval is not self-evident from the figure or the surrounding text.

Response:

We apologize for the lack of clarity in Figure 3. In the revised manuscript, we have redrawn Figure 3 and added supplementary explanations in Section 2.2 to improve its interpretability.

Revisions in Main Manuscript (Section 2.2)

SERGHEI-SWE-RE supports both synchronous and asynchronous coupling between surface and subsurface modules. In synchronous coupling, the surface and subsurface modules are executed at every global time step, defined as the smaller of their individual time steps, ensuring both solvers advance on a common timeline. In asynchronous coupling, as shown in Fig. 3, the two modules operate on their own independent timelines, denoted as $T_{swe}^p = \sum_{i=0}^p \Delta t_{swe}^i$ and $T_{re}^q = \sum_{i=0}^q \Delta t_{re}^i$, where p and q are the time step indices and Δt_{swe}^p and Δt_{re}^q are the corresponding time step sizes. Both Δt_{swe}^p and Δt_{re}^q vary in time. The former is restricted by the CFL condition, and the latter is adjusted based on either the maximum change in water content (in the PC scheme) or the number of linearization iterations (in the MP scheme). A user-defined upper bound, $\Delta t_{re,max}$, is enforced to further constrain Δt_{re}^q (Li et al., 2025). The subsurface module is executed only when $T_{re} + \Delta t_{re}^q < T_{swe}$. With this approach, the SSE flux (q_{ss}) is computed using the current surface flow state and the previous subsurface flow state, shown as the red arrows in Fig. 3. Similarly, the latest surface water depth is sent backward to the subsurface module as boundary conditions, shown as the blue arrows in Fig. 3. Asynchronous

coupling significantly enhances computational efficiency by reducing the execution times of the RE module. However, it introduces a time lag of the subsurface module as illustrated in Fig. 3. This time lag inevitably results in coupling errors that depend on both the time scale difference between surface and subsurface flow and the lag duration, an issue extensively discussed in Li et al. (2023) and later in Section 3.5. Practically, users can tune $\Delta t_{re,max}$ to balance coupling accuracy and computational efficiency. It is, in principle, also possible to set the $\Delta t_{re}^q = N\Delta t_{swe}^p$ to prevent an excessive difference between the two time-step sizes. However, investigation of this approach, and a full investigation of the implications of synchronicity/asynchronicity is left for future work. Finally, it is worth noting that the synchronous/asynchronous coupling strategies apply to both PC and MP schemes of the RE solver. For both schemes, the surface water depth is used as a boundary condition for implicitly solving the subsurface pressure head. Consequently, the data exchange pathways between modules, shown as red and blue arrows in Fig. 3, are identical for both solution schemes.

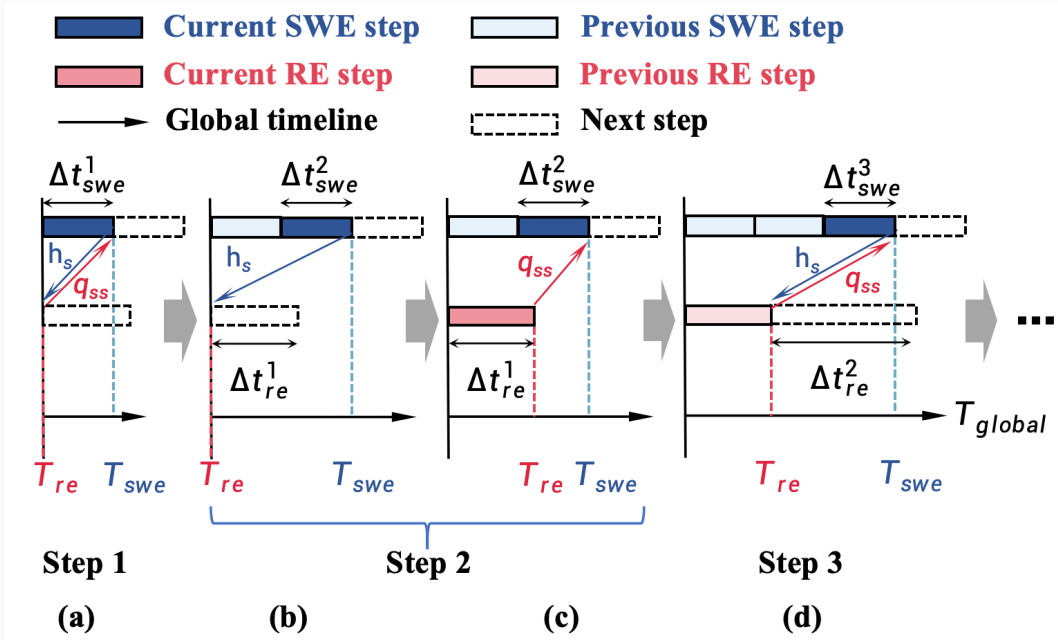


Figure 3. Schematic of the asynchronous coupling strategy between SWE and RE modules over time: (a) When $T_{re} + \Delta t_{re} > T_{swe}$, only the SWE module executes. The surface depth (h_s) is passed to the RE module, while the exchange flux (q_{ss}) is calculated using the lagged subsurface head from the previous step. This q_{ss} is passed back to the SWE module for finalizing the surface water depth; (b) As the SWE module advances, the execution condition for the RE module ($T_{re} + \Delta t_{re} < T_{swe}$) is met; (c) The RE module executes using the most recent h_s as a boundary condition. Once the subsurface state is updated, a q_{ss} is calculated and returned to the SWE module for updating surface water depth; (d) If the next RE step would again exceed the updated surface timeline, the RE module remains stationary while the SWE module continues to advance, re-entering the state described in (a).

Comment 9

Line 179: the y-range of the inlet boundary in Case 2 is not specified.

Response:

We have added the following clarification in Case 2:

--- **Revisions in Main Manuscript (Section 3.2)** -----

This test case maintains the same domain configuration as Case 1, except that a fluctuating surface water level is enforced at the inlet boundary ~~($x=260$ m)~~ located at $x=260$ m for y ranges from 0 to 260 m, following the hydrograph shown in Fig. 6.

Comment 10

Section 2.3 should state explicitly whether the model requires a structured Cartesian grid. This is relevant context for the domain decomposition limitations discussed in Section 4.

Response:

Yes. SERGHEI currently strictly requires a structured Cartesian grid. The domain decomposition strategy relates to this, but is become more flexible in developments happening in parallel. There is an outlook towards more flexible meshing strategies, but not in the immediate future, and thoroughly out of scope of the developments concerning this paper. We have revised Section 2.2 to explicitly explain the grid system of SERGHEI.

--- **Revisions in Main Manuscript (Section 2.2)** -----

The coupling strategy of SERGHEI-SWE-RE follows an externally coupled, sequential (non-iterative) approach. This means, both solvers are solved independently, and they exchange states and fluxes in sequence, without iterating during a coupling time step. Spatially, the surface and subsurface domains share the same horizontal discretization, both using structured Cartesian grids with uniform grid spacing (Fig. 1).

Comment 11

The finding in Case 5 that full SWE produces systematically higher ponding depths than kinematic or diffusive wave approximations (lines 247–252) is more interesting than the current discussion suggests. This has direct relevance for flood hazard applications and deserves more than a brief mention. Quantifying the discrepancy and commenting on the conditions under which it is largest would strengthen the paper.

Response:

We agree that this is an interesting and relevant finding. We would not label this as "ponding depths", rather just "depths". It is a subtlety, but meaningful. As we have pointed out in the manuscript, this finding is consistent with specific comparisons elsewhere in the literature, where steady flows computed with simplified SWE forms results in somewhat different depths (which has been quantified in other

papers). Under unsteady conditions this is difficult to gauge. We have revised Section 3.5 and the Conclusions to discuss this finding with quantitative evidences.

--- **Revisions in Main Manuscript (Section 3.5)** -----

This indicates that in this case, the simplified equations underestimate the water depth compared to the full SWE, thereby causing less ponding storage. Consistent with this, Caviedes-Voullième et al. (2020) and Li and Hodges (2021) reported that the diffusive wave approximation systematically underestimates water depths compared to the full SWE model in rainfall-runoff simulations. Similarly, de Almeida and Bates (2013) showed that simplified models (including diffusive wave equations and local inertial models) underestimate water depth gradients under unsteady conditions, leading to reduced ponding. Therefore, the model discrepancy in Fig. 15(b) does not undermine the validation of SERGHEI-SWE-RE for reliable surface–subsurface exchange simulations. In contrast, it underscores the need for the full SWE to resolve surface water dynamics, which is essential for accurate surface–subsurface exchange modeling. It also prompts the need to further conduct cross-model comparisons for coupled surface–subsurface models under more dynamic flow conditions, since kinematic and diffusive wave approaches are far more prevalent.

--- **Revisions in Main Manuscript (Section 5)** -----

The results also suggest potential differences between surface–subsurface models depending on how surface flow is formulated, i.e., typically kinematic and diffusive wave approaches in contrast to the fully dynamic approach adopted herein. This warrants further investigation.

Comment 12

Line 246: "Fig. 14(b)" should be "Fig. 15(b)". Figures 15(a) and 19 have overlapping lines/markers that are hard to distinguish; different line styles or marker sizes would help.

Response:

We have corrected the figure reference and optimized Figures 15 and 19 to improve the distinguishability of overlapping data series.

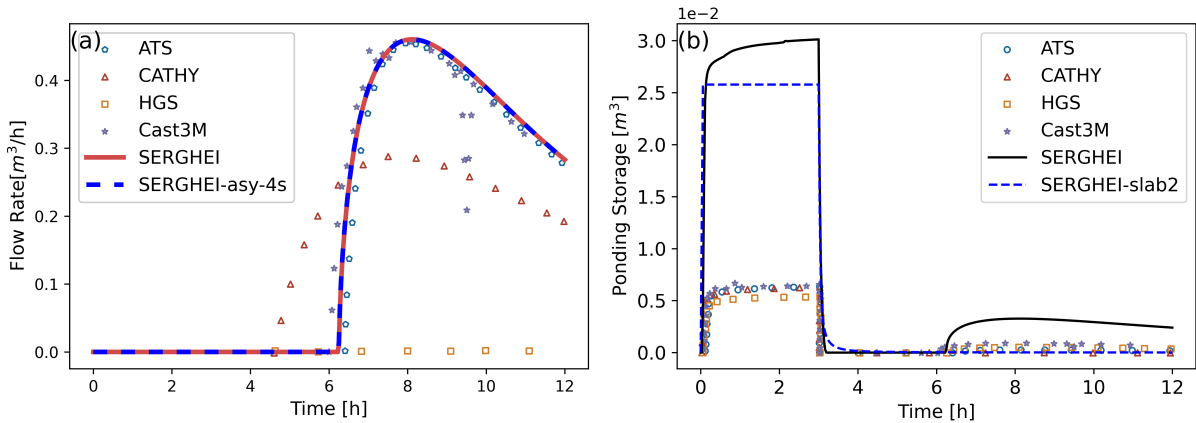


Figure 15. The temporal variations in (a) outlet discharge and (b) surface water volume. SERGHEI represents the synchronous calculation; SERGHEI-asy-4s represents the asynchronous calculation with the $\Delta t_{re,max} = 4s$. SERGHEI-slab2 is a customized SWE-only simulation on the slab 2, with the upslope and downslope sections truncated.




Figure 19. Parallel efficiency for the weak scaling tests performed with the extended lateral flow problem, on the JUWELS Booster with Nvidia A100 GPUs. Full, SWE and RE in the legends represent the efficiency of the entire simulation, the SWE solver and the RE solver, respectively. The dashed line represent the ideal efficiency.

Comment 13

Line 125: "that" should be "than".

Response:

Fixed.

References

- Casulli, V.: A coupled surface-subsurface model for hydrostatic flows under saturated and variably saturated conditions, *International Journal for Numerical Methods in Fluids*, 85, 449–464, <https://doi.org/https://doi.org/10.1002/flid.4389>, 2017.
- Caviedes-Voullième, D., Fernández-Pato, J., and Hinz, C.: Performance assessment of 2D Zero-Inertia and Shallow Water models for simulating rainfall-runoff processes, *J. Hydrol.*, 584, 124 663, <https://doi.org/10.1016/j.jhydrol.2020.124663>, 2020.
- Coon, E. T., Moulton, J. D., Kikinzon, E., Berndt, M., Manzini, G., Garimella, R., Lipnikov, K., and Painter, S. L.: Coupling surface flow and subsurface flow in complex soil structures using mimetic finite differences, *Adv. Water Resour.*, 144, 103 701, <https://doi.org/10.1016/j.advwatres.2020.103701>, 2020.
- de Almeida, G. A. M. and Bates, P.: Applicability of the local inertial approximation of the shallow water equations to flood modeling, *Water Resources Research*, 49, 4833–4844, <https://doi.org/https://doi.org/10.1002/wrcr.20366>, 2013.
- Furman, A.: Modeling coupled surface-subsurface flow processes: A review, *Vadose Zone Journal*, 7, 741–756, <https://doi.org/{10.2136/vzj2007.0065}>, 2008.
- Hutton, E. W. H., Piper, M. D., and Tucker, G. E.: The Basic Model Interface 2.0: A standard interface for coupling numerical models in the geosciences, *Journal of Open Source Software*, 5, 2317, <https://doi.org/10.21105/joss.02317>, 2020.
- Li, Z. and Hodges, B. R.: Revisiting surface-subsurface exchange at intertidal zone with a coupled 2D hydrodynamic and 3D variably-saturated groundwater model, *Water.*, 13, <https://doi.org/10.3390/w13070902>, 2021.
- Li, Z., Hodges, B. R., and Shen, X.: Modeling hypersalinity caused by evaporation and surface–subsurface exchange in a coastal marsh, *J. Hydrol.*, 618, 129 268, <https://doi.org/10.1016/j.jhydrol.2023.129268>, 2023.
- Li, Z., Rickert, G., Zheng, N., Zhang, Z., Özgen-Xian, I., and Caviedes-Voullième, D.: SERGHEI v2.0: introducing a performance-portable, high-performance three-dimensional variably-saturated subsurface flow solver (SERGHEI-RE), *Geosci. Model Dev.*, 18, 547–562, <https://doi.org/10.5194/gmd-18-547-2025>, 2025.
- Peckham, S. D., Hutton, E. W., and Norris, B.: A component-based approach to integrated modeling in the geosciences: The design of CSDMS, *Computers & Geosciences*, 53, 3–12, <https://doi.org/10.1016/j.cageo.2012.04.002>, 2013.
- Tubini, N. and Rigon, R.: Implementing the Water, HEat and Transport model in GEOframe (WHETGEO-1D v.1.0): algorithms, informatics, design patterns, open science features, and 1D deployment, *Geoscientific Model Development*, 15, 75–104, <https://doi.org/10.5194/gmd-15-75-2022>, 2022.