

Geosci. Model Dev. Discuss., author comment AC1
<https://doi.org/10.5194/gmd-2022-208-AC1>, 2022
© Author(s) 2022. This work is distributed under
the Creative Commons Attribution 4.0 License.

Reply on RC1

Daniel Caviedes-Voullième et al.

Author comment on "SERGHEI (SERGHEI-SWE) v1.0: a performance-portable high-performance parallel-computing shallow-water solver for hydrology and environmental hydraulics" by Daniel Caviedes-Voullième et al., Geosci. Model Dev. Discuss., <https://doi.org/10.5194/gmd-2022-208-AC1>, 2022

- Title: environmental hydraulics is mentioned; the reader may expect more than 2D hydraulics and rainfall runoff -> rethink

We understand the reviewer's point. However, we'd like to keep this as the broader idea of environmental hydraulics is where we envision the applicability of SERGHEI-SWE, although SERGHEI's capabilities are not yet as broad as imagined.

- Abstract: it describes to a large extent the overall aim of the whole project where of course only parts are included in this paper -> reduce this a bit and therefore comment more on the results you obtained here

Done. The abstract has been reformulated somewhat.

- There are many abbreviations in the text which are not all well known to the readers -> introduce a list of abbreviations

Indeed, this is particularly relevant considering the interdisciplinary target audience. We have included a list of abbreviations in the form of a glossary.

- P. 4: as you are also addressing urban environments, I suggest to also comment on the linkage to sewer systems – although this might not be planned here

Indeed. It is currently not a high priority, but it is something that we aim for. We have included a comment in the text "*...both natural and urban environments (for which coupling to sewer system models is a longer term objective)..."*.

- 5: your framework is quite fundamental; why are viscous / turbulent terms not included, e.g. to also in the perspective include eg turbulence models (eg algebraic, k - ϵ); this also might be relevant in the view of environmental hydraulics (title); this might be extended later (include in chap. 7) ?

The reviewer brings up an important point. From a theoretical point of view, the addition of viscous / turbulent terms into the shallow water equations turns the hyperbolic PDE into a degenerate parabolic-hyperbolic PDE, which has some implications for the solution procedure. In general, the addition of such terms in SERGHEI is straight-forward. We did

not consider these terms, because we targeted rainfall-runoff simulations in this work. The term “environmental hydraulics” might have been used too broadly here. We used this term to refer to the use of computational hydraulics for environmental flow problems, in our case surface hydrology at the catchment scale. However, we are aware that the term often is used for river hydraulics—ecology relation investigations. We have written: *The observant reader will note that in Equation 1, viscous and turbulent fluxes have been neglected. The focus here is on applications (rainfall-runoff, dam-breaks) where the influence of these can be safely neglected. Turbulent viscosity may become significant for ecohydraulic simulations of river flow, and turbulent fluxes play of course an important role in mixing in transport simulations. We will address these issues in future implementations of the transport solvers in SERGHEI.*

- p. 5: there are also friction laws for small water depths; as rainfall runoff has an important role, should 1 or 2 be mentioned ? Ilhan knows

This is indeed a very relevant point. We have included some comments in the text for this. We do not wish to explore this in this manuscript, but it is certainly very relevant. We now state:

In addition, specialised formulations of the friction slope exist to consider the effect of microtopography and vegetation for small water depths, e.g., variable Manning’s coefficients (Jain and Kothyari, 2004; Mügler et al, 2011} or generalised friction laws (Özgen et al., 2011). A recent systematic comparison and in-depth discussion of several friction models with a focus on rainfall-runoff simulations is given in Crompton et al (2020). Implementing additional friction models is of course possible, and relevant, especially to address the multiscale nature of runoff in catchments, but not essential to the points in this paper.

- p. 5: a comment that there are also many 2nd order schemes and unstructured grids

This is a good point. We have included the following:

It is relevant to acknowledge that second (and higher) order schemes for SWE are available (e.g. Buttinger-Kreuzhuber et al., 2019; Caviedes-Voullième et al., 2020b; Hou et al., 2015; Navas-Montilla and Murillo, 2018). However, first order schemes are still a pragmatic choice (Ayog et al., 2021), especially when dealing with very high resolution (as targeted with SERGHEI) which offsets their higher discretisation error and numerical diffusivity in comparison to higher order schemes. Similarly, robust schemes for unstructured triangular meshes are well established, together with their well-known advantages in reducing cell counts and numerical diffusion (Bomers et al., 2019; Caviedes-Voullième et al., 2012, 2020a). As these advantages are less relevant at very high resolution, we opt for Cartesian grids to avoid issues with memory mapping, coalescence and cache misses in GPUs (Lacasta et al., 2014) and additional memory footprint, while also making domain-decomposition simpler. Both higher order schemes and unstructured (and adaptive) meshes may be also implemented within SERGHEI.

- p. 12, Fig 5: add axis description specific discharge, why is discharge 0 ? or do you show the error or a relative error, unit ? in case it is a (relative) error, add description how it is computed in the text; supercritical flow: is this elevation or water depth ? the water depth goes down to 0 ? add a comment if it becomes very small and is no more visible

Done. The axis labels have been corrected.

- p. 13, Tab 2: I suggest to explain the computation of the norms in the text add a, b, c, d also in Fig 5; I do not understand the numbers in Tab 2, they are far away from machine accuracy?

The norms in Tab. 2 are not normalised by the number of elements. Thus, in order to obtain an average error, values in the table have to be divided by the number of elements, in this case $N = 1000$. We added the equations to calculate the L-norms to clarify. There was also an error in the norms concerning the C-property. It is now zero, as expected.

- chap. 4.3.1: a bit more description here

We have moved this case to the appendix and extended the description a bit.

- p. 17, Fig. 9: the legends must be larger, add [m] for the axis

Done

- p. 18, chap. 4.4: please comment a bit more on this rainfall; otherwise the results cannot be understood

We have now included a brief description of the rainfall: "*a piecewise constant rainfall with two periods of alternating low and high intensities (50.8 and 101.6 mm/h) up until 2400 s.*"

- p. 22, Fig 15: legend text a bit larger here

Done

- p. 25, chap. 5.7: in the previous case 1cm, what resolution here ?

Same resolution. It is now explicitly stated and we have slightly reformulated the text. We meant to say that in principle, this case really demands this resolution, whereas in the previous case it is not really a requirement of the topology. We have also moved this to the appendix.

- p. 26: add 1 or 2 sentences to the agreement of model and field results; Fig 20: explain these 3 terms in the text, otherwise the reader cannot understand

Thanks for pointing this out. We have now included a short description of the results, their visualisation and interpretation.

- chap 4+5: there is a large number of test cases which of course is very impressive; I would include all these test cases in a validation document, however I suggest to reduce this number in a journal paper and possibly shift some cases in the appendix, for example: 4.1: 2 instead of 4 hydraulic cases; 4.2 show only 4.2.2; 4.3 show only 4.3.2; 5.3 show only 1 case; 5.6+5.7: show only 1; please rethink, this does not mean that all should be shifted to the appendix or removed from the paper

Indeed, it is our goal to be very thorough in the verification strategy. We do agree with the comments that the volume of tests distracts from the core message of the paper. Consequently, and since there is a nice consensus from the comments, we have moved many of the tests into the appendix. We still keep a strong verification content, as it is essential to demonstrate the validity and robustness of the implementation. It is also relevant to highlight that most of the case setups (except those which require data we are not at liberty to distribute) are available for the community to also test. We leave one test to show how SERGHEI-SWE conserves the C-property with a wet/dry front, two steady transcritical flumes, one analytical dam-break and one transient flow in a parabolic bowl. We leave one analytical rainfall-runoff case, one experimental dam-break flow in a flume, and one experimental steady flow in a flume. We also keep one

experimental smooth transient flow over an island and one experimental rainfall-runoff case in a flume. Finally, we also keep a field runoff experiment and the Malpasset dam break. All others have been moved to the appendix.

- chap. 6: from small scale to large scale: show 6.2 as 6.1 ?

Yes, this is a good suggestion. We have changed the order.

- 4, 5 + 6: I understand and it is good that the cases descriptions are very short as most cases have been widely used in the literature; however not every reader know every cases; to avoid that the reader has to search for several other papers to understand the systems and results in the test cases I suggest to add core information to each test case in an appendix, eg: figure of system, initial and boundary condition, core parameters; it is not a must, but would be helpful for the reader

We have extended some of the descriptions, to ensure that all the critical information is available. However, given that there are plenty of reproductions of the setups in the literature, we believe it is not particularly critical to contribute with yet-one-more detailed description or sketch of the cases. We also highlight that most of the cases are openly available to download and run with SERGHEI, where all the details of the setups are necessarily included.

- chap. 7.2+3: you had that many cases before, therefore I am wondering why not one of the previous cases was chosen or the cases of 7.2+3 should be included earlier -> rethink

This test is particularly interesting for scalability analysis. We have now written out this explicitly and explained why:

"This is a simple analytical verification test in the shallow water literature, which generalises the 1D dam-break solution. We purposely select this case (instead of one of the many verification problems) for its convenience for scaling studies. Firstly, resolution can be increased as will. Additionally, the square domain allows for trivial domain decomposition, which together with the fully wet domain and the radially-symmetric flow field minimises load balancing issues. Essentially, it allows for a very clean scalability test with minimal interference from the problem topology, which facilitates scalability and performance analysis (in contrast to the limitations of the Malpasset domain discussed in section 7.1). "

For verification, this test overlaps with the dam break over a wet bed, which is why we choose not to include it.

- p. 32, close to end: add references to Ecosys, Ech20

Done

- chap. 8: is further (substantial) performance improvement possible using domain decomposition with MPI ?

Our multi-GPU and multi-node simulations already rely on domain decomposition and MPI. We have made this a bit more clear in the HPC implementation section for which we have written a new short introductory paragraph explaining the different layers of parallelism:

"In this section we describe the key ingredients of the HPC implementation of SERGHEI. Conceptually, this requires, firstly, handling parallelism inside a computational device (multicore CPU or GPU) with shared memory, and the related portability and

corresponding backends (i.e., OpenMP, CUDA, HIP, etc.). On a higher level of parallelism, distributing computations across many devices requires domain decomposition and a distributed memory problem, implemented via MPI. The complete implementation of SERGHEI encompasses both, distributing parallel computations into many subdomains, each of which is mapped onto a computational device. Here we start the discussion from the higher level of domain decomposition.”

■

Further minor comments are in the commented pdf.
We have addressed the typos and suggestions.

■

Experimental idealised urban dam-break: comment on model results and measurements, does a refined grid improve the agreement. similar deviations found in the literature ?

We included some qualitative comments on the quality of the solution. The results are comparable to those in the literature. We do not wish to go deeper the analysis in terms of grid convergence. Proper quantitative analysis of the results, especially for different grids is rather complex for this case, as there are many dynamics and many different factors playing a role. This would require a dedicated effort which is not the point of this paper.