Geoscientific Model Development Discussions

Geosci. Model Dev. Discuss., author comment AC2
https://doi.org/10.5194/gmd-2022-208-AC2, 2022

Reply on RC2

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-AC2, 2022

- The paper deals with the development of, what the authors call, Simulation Environment for Geomorphology, Hydrodynamics and Ecohydrology (SERGHEI) - Shallow Water Equation (SWE) for hydrodynamics, ecohydrology, morphodynamics simulations. Although there are many shallow water equation models in the world, I think the SERGHEI-SWE is new in that it can handle even exa-scale problem such as with 122000000 computational cells (dx=0.5m) for the rainfall-runoff processes at lower Triangle region in the East River Watershed with fast parallel computing.

This is a fair summary. We have decomposed the reviewer's comments to address them in order.

- My impression of the overall contents of the paper was that actually the contents is basically for the evaluation of shallow water model, however, the authors are emphasizing more on the broad view of SERGHEI that it can handle hydrology and environmental hydraulics problems.

We intend to convey both aspects. We specifically describe the SERGHEI-SWE module, and evaluate in terms of accuracy, applicability and performance. Nonetheless, since SERGHEI-SWE is the core hydrodynamic module of the broader SERGHEI, and because the modularity and broader project shape both the motivation and the design of SERGHEI-SWE, it is relevant and interesting to discuss in this paper as well. We intend to convey that there is a clear outlook in proposing SERGHEI-SWE.

- I had a feeling that you could simply describe the importance of SERGHEI for the future of pure shallow water equation modelling in the paper, but you didn't.

This is a fair point. As the reviewer points out several times, there is some opportunity in the manuscript to extend on the outlook of HPC (and exascale) enabled shallow water modelling. We have written some remarks on this in the introduction and conclusions.

- Likewise, you do not have to say fluvial or urban flood modelling as classical engineering problem. If you model fluvial or urban flood modelling with e.g. 100 m resolution, you may be justified to say that it is established engineering problem, but I think it is still new if you model them with one-digit resolution and large area with scientific scope such as considering sub, super-critical flow distinction.

From a historical point of view, shallow water solvers received wide attention for fluvial and urban flooding applications, and a large part of the numerical developments, as well as software implementations were tailored to these applications. However, this trend has been systematically changing, opening up shallow water models for broader hydrological applications. This difference is what we mean to highlight here. By "classical", we meant that it is the more established application, but we did not intend to convey that it is one for which everything is solved and that there are no scientific nor technical questions to address.
We do agree that flood modelling applications with large domains and very high resolution are NOT an established thing, especially with complex transcritical flow conditions for which simplified models are not well suited. Such applications can of course benefit from this technology. We have included some remarks to clarify this, and better frame the applicability and interest of SERGHEI across different use cases.

-

  In this sense, turbulent modelling may be interesting to add more scientific essence in the analysis.
  Yes, this is of course an interesting point, and it will come up as more relevant as we move towards additional applications, especially those concerning transport. We intend to address this in the future, when we have the transport modules further developed. This is now also mentioned in the text, coinciding with a point from another reviewer.

-

  Anyway, I would appreciate it very much if you could explain more why the exa-scale handling is necessary for the shallow water modelling. Lower Triangle region is 14.82km2 which is small. Why do you need to model such catchment with 0.5m resolution?
  Concerning the first part of this question, please see the reply to point 7 below.
  There are two parts to the second question, of why do we need to model the Lower Triangle at 0.5m resolution:

a. for the purposes of this paper, there is no need to do this. We chose the Lower Triangle because a high-resolution DEM was available, and because it has a complex topography which should be general enough to test SERGHEI's performance under real conditions. It creates a large computational problem in space, and in time due to high velocities thanks to the high slopes. In this sense, we stress that this is simply a computational performance test. We set out to test the scalability of SERGHEI on something that was close to reality, or a plausible application. We note that the value of such high-resolution data has been recognised by the shallow water community and the Lower Triangle case has been used in a similar manner as a benchmark in Shaw et al. (2021) (https://doi.org/10.5194/gmd-14-3577-2021) to assess the parallel performance of LISFLOOD-FP 8.0.

b. From a broader perspective, and to frame the context of why the Lower Triangle is of interest. The Lower Triangle catchment is part of the U.S. Department of Energy's Watershed Scientific Focus Area (Watershed SFA) project and is intensively monitored, see Hubbard et al. (2018) (doi: 10.2136/vzj2018.03.0061). The catchment is strategically important because it collects snowmelt-based headwater from upstream catchments throughout the year, and thus, the discharge at the outlet of this catchment is an integrator of many hydrological signatures. From a water resources management point, the catchment is important because it contributes significantly to the water supply of the midwest of the United States. Locally, a pump house at the outlet of the catchment provides water for the municipality of Mt. Crested Butte, Colorado, see Maavara et al. (2021) (doi: 10.1371/journal.pone.0247907). One of the aims of Watershed SFA is to link hydrological signatures to geochemical processes. It is well-known that hydrological fluxes create hot spots and hot moments in watershed geochemistry, see, for example Arora et al. (2022) (https://doi.org/10.1007/978-3-030-95921-0_2), Zhao et al. (2021) (https://doi.org/10.1016/j.scitotenv.2020.144168), and Krause et al. (2017) (https://doi.org/10.1002/2016WR019516). These hydrological fluxes are dependent on small-scale watershed features that can only be captured at these high resolutions, for example, infiltration and exfiltration patterns are strongly controlled by microtopography. Thus, in order to properly detect potential geochemical hot spots and study how they scale up to the watershed scale, this high resolution information is required. In the future, we will continue to use this area to test additional modules oriented towards these purposes.

Overall, in the research outlook of the SERGHEI team, sites such as the Lower Triangle are interesting to simulate at very high resolution to study how small scale processes produce large scale hydrological signatures in general. Please see next point for further discussion.

- What is your vision for the exa-scale modelling using a shallow water equation?

We interpret this point as a request to better explain why there is a need for exascale-ready shallow water modelling. We appreciate this very much, as this is indeed at the core of our proposition with this model.

There is the opportunity to exploit increasingly better and highly resolved geospatial information (DEMs, land use, vector data of structures) which prompts the need for high resolution solvers. At the same time, the push towards the study of multiscale systems and integrated management warrants for increasingly larger domains. Together, these trends result in larger computational problems, and consequently HPC shallow water solvers.

Although it is likely that for purposes such as generating hydrographs (which would be the interest for flood modelling) such high resolution is unnecessary, to capture the spatial dynamics, pathways and transit times (all of which are of ecohydrological interest), it is most likely beneficial. High resolution is necessary for such purposes to avoid some type of a-priori averaging of properties/parameters/processes. We are interested in ecohydrological processes, which will arguably benefit from this larger resolution.

We have included these arguments in the "Conclusions and outlook" section.

- In addition, it is little hard to follow all the benchmark cases one by one. The total number of page is 44. You can remain for example only the essential benchmark for readers' sake. If you say, that every benchmark is necessary for the true evaluation of the shallow water model by SERGHEI, then you can remain all of them, but you can as well keep them in the appendix for example. Please consider them. Please sharpen the contents for what you really want to convey.

There is a strong consensus from the review comments on this point, and we actually agree that the verification section is far too long. We do keep all of them, but around half of them have been moved to the appendix. The ones we keep intend to show the robustness and correctness of the SERGHEI-SWE implementation for a variety of problems which can be tackled with SWE solvers.

- Overall, I think there are shallow water models which can simulate benchmarks with the same level accuracy with SERGHEI-SWE except that those are slower than SERGHEI-SWE.

This indeed is likely the case. We do not make any comparative arguments in terms of the accuracy of SERGHEI-SWE. The main reason is that we implement well-established numerical schemes, and we only wish to verify that the implementation is robust and accurate with our verification exercise.
Whether other solvers are "faster/slower" is a matter of debate. Our argument is not really one of speed, but of scalability and portability. That is, a very lightweight, well-implemented solver, running on a GPU and using potentially a cheaper numerical scheme may do the job very fast for any given application. However, it would be limited to off-the-shelf consumer hardware. With SERGHEI, what we are proposing is that the community knowledge, developed and gathered in the past decade, can be harnessed into state-of-the-art computational technology, and it can be scaled up, in the same way that many communities in the geosciences and engineering do. Since HPC implementations require additional attention to implementation details, and because we explicitly set out for it, SERGHEI is likely to be a comparatively efficient implementation — although we would not outright claim this, since we do not wish to engage in an intermodel performance comparison. However, how "fast" it is, is a function of the hardware, reason for which portability is relevant. The point is really, that with SERGHEI, we offer the community a scalable solver, which, sure, allows to tackle "normal" problems faster, but most importantly, enables larger problems and sets a tool that can harness upcoming computing technology.

- But other shallow water models could be more user-friendly. Likewise, the users of other shallow water models may consider that they do not need exa-scale computation for their purpose. The authors need to explain the vision how the SERGHEI-SWE needs to be used in the near future.

In general terms we agree, although this depends on the definition of user-friendliness. In the context of engineering-oriented shallow water solvers "user friendly" has often meant a sophisticated Graphical User Interface (usually topped with Windows support). This is something mostly out-of-scope for our purposes. We envision SERGHEI mostly as a scientific tool, and assume that users are capable of interacting with the model via a small set of input files. We also provide some R-based interfaces for input/output, which aside from being our tools to pre/post process SERGHEI input/output, also facilitate its usage. Our contribution in terms of user-friendliness is that SERGHEI is easily deployed on a local GPU-enabled workstation. Although we target large HPC facilities, the required dependencies and underlying software stack is quite lightweight. Therefore, setting up and building SERGHEI on a workstation is rather straightforward. Most importantly, our vision of SERGHEI as a community model has required for the framework to be simultaneously performance-portable and to facilitate development by the community (i.e., the code is quite readable, and not obfuscated by hardware specific code or complicated software abstractions). In this sense, it is a different type of user friendliness.
We of course don't argue that everyone needs exascale. But most likely every user of shallow water models will benefit from performance, and we hope, from a transparent, open-source community model.

- The authors do not necessary needs to write the answers to the above questions in the paper itself but I would like to know the vision of exa-scale computing.

As the reviewer can tell, we have addressed these points in an extended discussion. Additionally, some of these points are also succinctly treated in the manuscript, as these issues are relevant to the overall presentation of SERGHEI and its expected use and place in the community.