Dear Editor and Referee #1,

We thank you for the time spent to evaluate and review our manuscript and for your positive recommendations.

Referee #1 raises a number of minor comments, which were all addressed in the revised manuscript. Point-by-point answers to the Referee #1 are provided below, whereby we explain how and where each point raised was incorporated into the revised manuscript.

Answers are in blue text and appear after each reviewer comment, followed by

a box showing the associated changes applied to the revised manuscript.

Sincerely,

Georges Kesserwani and James Shaw

"Congratulations to this very clear and concise piece of work. The overall goal, methods, and results of the model development and model application are described with sufficient detail, and the overall structure of the paper supports reading and understanding of the content. Also, the presentation of the results is in most cases clear and informative."
We thank the reviewer for their interest in our paper, and we are glad the manuscript was clear and concise. The manuscript has been revised to address all the referee's comments, as we clarify in the following.

"The fact that model code and study results are provided on Zenodo is very much appreciated, great! A short readme or similar would help, however, to guide the user through the folder structure of the repository."
The Zenodo record description at `https://zenodo.org/record/4066824` has been substantially expanded to document the folder structure, file format and filename conventions. Also, the website, `https://www.seamlesswave.com/LISFLOOD8.0`, was cited in the abstract and introduction. This website allows readers to access step-by-step documentation of how to run LISFOOD-FP8.0 over a series of test cases.

"1. Page 2 / lines 52-43: maybe good to mention the license under which the model is distributed? This gives directly an idea of what is permitted and not. Additionally, I strongly encourage you to add a DOI in the manuscript linking to the Zenodo-version used in this manuscript."
The abstract and the introduction section has been revised to mention the license under which the model is distributed.

Abstract

LISFLOOD-FP 8.0 therefore marks a new step towards operational DG2 flood inundation modelling at the catchment scale. LISFLOOD-FP 8.0 is freely available under the GPL v3 license, with additional documentation and case studies at `https://www.seamlesswave.com/LISFLOOD8.0`.

1 Introduction

This paper presents a new LISFLOOD-DG2 solver of the full shallow water equations, which is integrated into LISFLOOD-FP 8.0 and freely available ~~for non-commercial use~~ under the GNU GPL v3 license (LISFLOOD-FP developers, 2020).

"2. Page 11 / line 213: why are these recent optimizations lacking? Please provide a brief explanation."
The optimised LISFLOOD-ACC solver specified by Neal et al. (2018) implements a sub-grid channel model (Neal et al., 2012a) and CPU-specific optimisations that do not translate naturally to GPU architectures, and this optimised ACC solver does not yet support the rain-on-grid features used later in Sect. 3.3. To

facilitate an intercomparison between solvers, the LISFLOOD-ACC solver used here is the version specified by Neal et al., (2012b), which supports the necessary rain-on-grid features and shares the same algorithmic approach as the FV1 and DG2 solvers.

3 Numerical results

The optimised LISFLOOD-ACC solver ~~used here is the version specified by Neal et al., 2012b, which supports~~ specified by Neal et al., 2018 implements a sub-grid channel model (Neal et al., 2012a) and CPU-specific optimisations that do not translate naturally to GPU architectures. Additionally, at the time that model runs were performed, the optimised ACC solver did not yet support the rain-on-grid features used later in Sect. 3.3~~, but lacks the recent optimisations for multi-core CPUs, as documented by Neal et al., 2018. Rain-on-grid support will be added to the optimised ACC solver in a future LISFLOOD-FP release~~[a]. To facilitate a like-for-like intercomparison between solvers, the LISFLOOD-ACC solver used here is the version specified by Neal et al., 2012b, which already supports the necessary rain-on-grid features and shares the same algorithmic approach as the FV1 and DG2 solvers.

---

[a] Rain-on-grid features have since been added to the optimised ACC solver, and will be available in a future LISFLOOD-FP release.

3.1.3 Solver runtimes for a varying number of elements

To facilitate a like-for-like comparison with FV1 and DG2, ACC solver runtimes were obtained for the ACC implementation of Neal et al., 2012a.

...

It is expected that additional performance can be gained by using the alternative, CPU-optimised ACC implementation (Neal et al., 2018), and these CPU-specific optimisations are also under consideration for future enhancement of the DG2 and FV1 solvers.

“3. Section 3.1 and 3.2: for both test cases, it is stated that they are widely used to assess model predictions. What I thus would like to see is an extended discussion of how model results here generally compare to other modeling studies. For instance, are the trend presented here and their magnitude comparable? A bit more depth is needed to move the manuscript more towards a scientific publication rather than an extended model documentation logbook.”

In section 3.1.1, the previous version of the manuscript already mentioned that “DG2, FV1 and ACC predictions of water depth and velocity agree closely with existing industrial model results (Fig. 4.10 and 4.11 in Néelz and Pender (2013)).”. A sentence was added to further discuss the modelling studies in (Huxley et al. 2017, Table 11)

3.1.1 Water depth and velocity hydrographs

Differences in velocity predictions are more pronounced (Fig. 6e–h). The biggest differences are seen at point 1 (Fig. 6e), located only 50 m from the breach, since the flow at this point is dominated by strong inflow discharge with negligible retardation by frictional forces. At point 1, ACC and DG2 velocity predictions agree closely with the majority of industrial models (Fig. 4.11 in Néelz and Pender (2013)). LISFLOOD-FV1 predicts faster velocities up to 0.5 ms^{-1}, which is close to the prediction of TUFLOW FV1 (Huxley et al., 2017, Table 11).

For section 3.2.1, two sentences have been added to further discuss the modelling studies in (Huxley et al. 2017) for TUFLOW solvers and the (Ayog et al. 2021) including results from a MUSCL-FV2 solver at 10 m resolution.

3.2.1 Water level and velocity hydrographs

Predicted water level and velocity hydrographs are shown in Fig. 11. The water level hydrographs show that water ponds in small topographic depressions at point 1 (Fig. 11a), point 3 (Fig. 11b) and point 5 (Fig. 11c). Point 7 is positioned near the steep valley slope, and is only inundated between $t = 1$ hour and $t = 8$ hours (Fig. 11d). At both resolutions, water levels predicted by all solvers agree closely with existing industrial model results at points 1, 3 and 7 (Fig. 4.16 in Néelz and Pender, 2013). Small water level differences accumulate as water flows downstream and at point 5, positioned farthest downstream of the dam break, differences of about 0.5 m are found depending on the choice of resolution and solver (Fig. 11c). Similar water level differences have been found amongst the suite of TUFLOW solvers (Huxley et al., 2017) and amongst other industrial models (Néelz and Pender, 2013).

Bigger differences are found in velocity predictions, particularly at locations farther downstream at point 3 (Fig. 11f), point 4 (Fig. 11g) and point 7 (Fig. 11h). At point 3, DG2 predicts small, transient velocity variations at $\Delta x = 50$ m starting at $t = 1$ hour; these variations are not captured by the FV1 or ACC solvers, but have been captured by a FV2-MUSCL solver at the finest resolution of $\Delta x = 10$ m, as reported by Ayog et al., 2021.

"4. Chapter 3.1.2: Here, the cases are benchmarked for two different resolutions. The initial 5 m resolution and a finer 1 m resolution. Question 1 is, how did you derive the new geographical data, how did you perform the resampling? And question 2, why did you not look into grid coarsening - wouldn't the chance of having abrupt water level changes between cells become larger when applying grids with, admittedly probably very much, coarser spatial resolution? This would be less relevant to assess speed of the solvers, but accuracy."
Reply to Question 1: Spatial grid convergence is studied by modelling at grid resolutions between $\Delta x = 5$ m and $\Delta x = 1$ m. Since the floodplain is flat, no topographic resampling is required.

3.1.2 Spatial grid convergence

~~Next~~Spatial grid convergence is studied by modelling at grid resolutions of $\Delta x = 5$ m, 1 m, and 0.5 m. Since the floodplain is flat, no topographic resampling is required. On each grid, the water depth cross-section is measured along the centre of the domain (Fig 7). DG2, FV1 and ACC cross-sectional profiles at the standard grid spacing of $\Delta x = 5$ m agree well with industrial model results (Fig 4.13 in Néelz and Pender (2013)).

Reply to Question 2: Coarsening beyond $\Delta x = 5$ m yielded a similar conclusion to that reached by the results at $\Delta x = 5$ m, so additional coarse results were excluded to ensure the clarity of Fig. 7.

"5. Chapter 3.3: While for the two Environment Agency test cases a motivation was stated, it does not become clear where why exactly this case study was selected. Please elaborate briefly why you decided to use this test case and not another one from the rich literature of LISFLOOD-FP studies, for example."
The Storm Desmond case study was selected for its spatially- and temporally-varying rainfall across a large catchment. This dynamic rain-on-grid capability was unavailable in earlier versions of LISFLOOD-FP and therefore unavailable in earlier LISFLOOD-FP studies.

3.3 Catchment-scale rain-on-grid simulation

In December 2015, Storm Desmond caused extensive fluvial flooding across the Eden catchment in North West England (Szönyi et al., 2016). This storm event has previously been simulated using a first-order finite volume hydrodynamic model (Xia et al., 2019), with ~~flow~~ overland flow and fluvial flooding driven entirely by ~~rainfall data. The simulation is revisited~~ spatially- and temporally-varying rainfall data over the 2500 km^2 catchment. As such, this simulation is ideally-suited to assess the ~~capability of DG2, FV1 and ACC in reproducing~~ new rain-on-grid capabilities in LISFLOOD-FP 8.0, and represents one of the first DG2 hydrodynamic modelling studies of rainfall-induced overland flow ~~and fluvial flooding over the 2500~~ across a large catchment.

"6. Page 27 / Line 423: this is a very important aspect and should be highlighted more prominently in the manuscript (e.g. abstract and/or summary). One of the key reasons many scholars/practitioners use LISFLOOD-FP is its sub-grid scheme. This also holds for the comment made in point 2."

The revised introduction clarifies that, "Since the new DG2 and FV1 solvers are purely two-dimensional and parallelised for multi-core CPU and GPU architectures, the new solvers do not currently integrate with the LISFLOOD-FP sub-grid channel model (Neal et al., 2012a) or incorporate the CPU-specific optimisations available to the ACC solver(Neal et al., 2018)."

1 Introduction

~~All solvers can load spatially- and temporally-varying rainfall data in TUFLOW NetCDF format (?), enabling real-world rain-on-grid simulations in~~ Since the new DG2 and FV1 solvers are purely two-dimensional and parallelised for multi-core CPU and GPU architectures, the new solvers do not currently integrate with the LISFLOOD-FP ~~8.0~~sub-grid channel model (Neal et al., 2012a) or incorporate the CPU-specific optimisations available to the ACC solver (Neal et al., 2018).

The revised opening of Sect. 3 also clarifies that, "The optimised LISFLOOD-ACC solver specified by Neal et al., 2018 implements a sub-grid channel model (Neal et al., 2012a) and CPU-specific optimisations that do not translate naturally to GPU architectures." (recall the changes made in response to point 2 above).

Abstract

The solvers are parallelised on multi-core CPU and Nvidia GPU architectures and run existing LISFLOOD-FP modelling scenarios without modification. These new, fully two-dimensional solvers are available alongside the existing local inertia solver (called ACC), which is optimised for multi-core CPUs and integrates with the LISFLOOD-FP sub-grid channel model.

"7. Chapter 3.3.2: For the analyses of flood extent, would it not be useful to include metrics like the hit rate, false alarm ration and critical success index to quantify the actual (dis)agreement between simulations and results?"

We followed the reviewer recommendations and investigated the hit rate, false alarm and critical success index metrics, given their relevance to quantify the analysis of floodplain comparisons.

Table 3

Hit rate (H), false alarm ratio (F) and critical success index (C) for the DG2 and FV1 predictions of maximum flood extent calculated against the reference solution of the ACC solver at $\Delta x = 10$ m.

	$\Delta x = 40$ m			$\Delta x = 20$ m			$\Delta x = 10$ m		
	H	F	C	H	F	C	H	F	C
DG2	0.83	0.32	0.59	0.86	0.55	0.40	0.85	0.56	0.40
FV1	0.77	0.27	0.59	0.80	0.19	0.67	0.93	0.03	0.90

3.3.2 Maximum flood extent over Carlisle

The DG2 and FV1 predictions of maximum flood extent can be quantified against the ACC prediction at $\Delta x = 10$ m, which is treated as the reference solution. The hit rate measures flood extent underprediction as the proportion of wet elements in the reference solution that were also predicted as wet. The false alarm ratio measures flood extent overprediction as the proportion of predicted wet elements that were dry in the reference solution. The critical success index measures both over- and underprediction. All three metrics range between 0 and 1, and further details are provided by Wing et al., 2017.

The hit rate (H), false alarm ratio (F) and critical success index (C) are given in Table 3. At $\Delta x = 40$ m, the critical success index is 0.59 for both DG2 and FV1, but DG2 has a higher hit rate and false alarm ratio, suggesting that DG2 predicts a wider flood extent than ACC or FV1. At $\Delta x = 20$ m and $\Delta x = 10$ m, the false alarm ratio and critical success index for DG2 deteriorate, but a hit rate of 0.83–0.86 is maintained, which is acceptable given that high-resolution predictions are downscaled from the DG2 piecewise-planar solution at $\Delta x = 40$ m. FV1 predictions at $\Delta x = 20$ m and $\Delta x = 10$ m are obtained directly without downscaling, and FV1 predictions converge towards ACC predictions with successive grid refinement. This convergence is evidenced in all three metrics, with FV1 at $\Delta x = 10$ m achieving a high hit rate (0.93), low false alarm ratio (0.03), and high critical success index (0.90).

“8. Figure 16 and Figure 17: This figure is hard to read. While adding the OSM background map is appreciated for geographical reference, (as done in a figure above), it’s diluting the actual information about the flood maps in the current form. Please consider revising this figure.” The OpenStreetMap background was removed from Figures 16 and 17.

“9. Summary/Conclusions: this section nicely wraps up the manuscript. However, recommendations for further studies and improvements, and the shortcomings of the current version (both feature-wise and technologically), are missing. Please set your work in context of what was done so far, how your work adds to that and opens up new possibilities, and what challenges are still lying ahead to fast hydrodynamic simulation over coarse and large domains”
The challenges that lie ahead to improve DG2 flood modelling over coarse and large domains as well as the possibilities for future improvements have been discussed in the revised ”Summary and conclusions” section.

4 Summary and conclusions

However, FV1 and DG2 are the first solvers in LISFLOOD-FP to gain a dynamic rain-on-grid capability, with this capability being added to the optimised ACC solver in a future release. To further improve efficiency and accuracy at coarse resolutions over large catchments, one future direction would be to port the sub-grid channel model—currently integrated with the CPU-optimised ACC solver—to GPU architectures. Another useful direction would be to enable a multi-resolution solver based on Kesserwani and Sharifian, 2020, and introduce a hybrid DG2/FV1 solver that downgraded the DG2 formulation to FV1 in regions of very thin water layer, or in regions of finest grid resolution, to further reduce the computational cost. Both directions are being investigated for inclusion in future LISFLOOD-FP releases.