

This manuscript is well written and presents a solid summary of work to incorporate reservoirs as simple idealized model objects. I have a few comments that may be worth addressing, but they are all suggestions and should not hold up publication of this work.

I've attached a copy of the pdf with inline comments.

*Thank you for taking the time to provide thoughtful feedback! We appreciate your positive evaluation and have included responses to each of your specific suggestions below in italics.*

Line 81: Not sure this sweeping statement is justified

*We agree and will modify this line to be less general and more specific to the particular advantages of Integrated hydrologic models*

Lines 178-181: I feel like the design philosophy of starting with simplicity and extensibility in mind could be stated more succinctly at the start of this section

*We agree and will revise this per your suggestion.*

Lines 189-190: scraped?

*We are using "scraped" here to denote the process of removing water from the domain and placing it into storage. We agree this is not the most intuitive word choice and will instead use the "diverted into reservoir storage" in the revised version.*

Lines 241-242: Nice touch – might want to mention this up in the intro? Accessibility is king and it would be good to make clear that this approach aligns with existing data.

*Thank you! We worked really hard to make this accessible and agree that this could be highlighted a bit more. In response to this comment and two similar comments, we will add a short paragraph describing our design philosophy, along with a brief description of our strategy of allowing our reservoirs to have their actual released amount calculated by a separate process in order to support both existing and future ways of representing operations.*

Line 425, Figure 9: Those are some really little reservoirs if these are 10m grid cells! Is this maybe too much of a toy example?

*We agree that 10m is small for a reservoir, however we would like to note that the purpose of this test case is only for testing scaling. Our goal here is to measure the efficiency of the underlying additions to Parflow's code both in terms of extra calculations being done, and extra memory being stored and accessed. Neither of these should be tied to the dimensions of individual grid cells. In this case the size of each grid cell does not impact the difficulty of the problem, it is really the number of grid cells being solved and the number of reservoirs we are applying, therefore, changing the size of our grid cells does not actually change the size of the problem. That said, we do agree that it may appear odd to some readers to have such small grid cells. We will add a few sentences to the description of this test case noting that we are scaling the number of unknowns in the problem and that similar results would be expected regardless of the grid cell sizes.*