

Author Response to Reviewer 3

We thank the editor and reviewers for their comments, which helped improve this manuscript. Please note: responses to reviewer comments are italicized and bolded text below each reviewer's comments.

General Comments:

This paper provides a comprehensive overview of the new updates and major revisions made to the Variable Infiltration Capacity (VIC) macroscale hydrology. This new version 5 of the model allows for different instances of the model driver, preserving legacy aspects of VIC and paving ways for coupling with other models, e.g., atmospheric models. A major change involves allowing the model to run in space first then in time, which was done in reverse in its legacy versions (prior to version 5). Also, the authors describe and provide some examples of several new test cases, which can be used for either unit and more scientific testing.

Overall it is very well written paper and provides many details on the way new model drivers and components can be run and interact in VIC. Not many papers prior to this have described the main software components of VIC, so in that of itself makes the paper more novel for the VIC and land surface model communities to reference in the future.

Thank you for taking the time to review this paper. We have appreciated your comments.

Specific Comments:

One area that is not addressed well in this paper is how the VIC distributed parameters (e.g., LAI or soils information) are handled and provided for the different drivers. The legacy “Classic” driver’s original I/O is mentioned to still be supported, but how is that information changed (or not) for the other formats (e.g., NetCDF) to support the other space-before-time or “image-based” drivers. The VIC parameters for previous versions have been another problem for the community as they are hard to generate or customize for each user’s needs. Please address how parameters can be generated (e.g., any tools) to support the “Image” type drivers and how some of the parameter tuning processes and practices (e.g., for soil parameters) can be used to optimize streamflow estimates. The handling of the parameters and preprocessing to other resolutions and grids is an important part of any model structure and release.

We agree that a new set of tools is needed in the “VIC ecosystem” to enable community wide adoption of new features in VIC, specifically the Image Driver. For example, the Image Driver now uses a netCDF file format for distributed input model parameters (e.g. soil and vegetation). We are aware of a number of efforts across the VIC community to develop tools supporting the new VIC.5 file formats. We expect those to be published separate from this paper.

Minor Comments:

Page 2, lines 13-15: Part c): Previous VIC versions that have been run, prior to 5, were very slow, especially when run for large-scale simulations (e.g., NLDAS), since each gridcell was run separately through time (as referred to as “time-before-space” in the paper). Many of the other LSMs that you mention are not always run in “coupled” mode and tend to run faster than VIC in uncoupled mode. You may want to consider removing this statement or addressing the deficiency of it in relation to computation speed and how truly it was able to run faster.

The statement in question is already conditioned on the comparison to larger coupled models. Moreover, our experience with previous versions of VIC taught us that the model was typically I/O bound – a fact that was exasperated by running each grid cell independent and writing separate files for each grid cell. Nevertheless, the computational resource required to run the model were in fact reasonably achieved without the requirement of large HPC systems.

Page 4, line 15: Please try to indicate what the acronym, “MT-CLIM”, stands for.

We have indicated in the manuscript that MT-CLIM name is derived from the “Mountain Climate Simulator”.

Page 5, line 6: Does “RVIC” simply stand for “routing VIC”? Please include any additional information to specify this acronym.

We have indicated in the manuscript that RVIC is short for “Routing VIC”.

Page 7, lines 30-32: The terminology convention of using “restarts” can be specific to certain communities, like the hydrometeorology community. It may be useful here to provide some short background on, first, what the “restart” represents and then to further specify what “exact” and “near-exact” indicate. Also, do all VIC versions, including 5, have an option to write out “instant” states and fluxes versus time-averaged?

Thank you for this comment. We have clarified what we mean by the terms restarts/exact/near-exact/instant.

Page 8, lines 9-13: The authors state here that the “input forcings must now have the same time step length as the model simulation” due to the removal of the MT-CLIM temporal disaggregation code, which has now become a separate preprocessing step. The reviewer wonders why temporal interpolation options have not been considered at the driver level for the forcing inputs, or why the MT-CLIM could not be rewritten to be incorporated more at the driver level. Please address in a statement or two.

The specific reasons MT-CLIM was omitted VIC at the driver level are described in the numbered list in section 2.6.

Page 8, lines 25-27: Which libraries and compilers has the new VIC version been tested with? This information would be useful to the user-audience.

VIC is routinely tested with GNU (C and Fortran) and Clang (C) compilers, and the Open-MPI, and netCDF libraries. Though not part of the continuous integration system for licensing reasons, VIC has also been tested with Intel, PGI, and Cray compilers and the MPICH and Intel-MPI libraries. We have indicated these details to the manuscript.

Page 8, lines 28-29: Are unit tests also available to be tested with the other drivers? What if changes to the drivers affect the unit tests? Please address this concern.

At the time of writing, VIC's unit tests only test the physical core of the model, and the shared driver code. The system tests are designed to test the driver functions and they are tested using a similar python testing framework. This point has been clarified in the manuscript.

Page 9, lines 1-5: Why are the science-level tests only available with one driver type? Do you have science tests also set up for the other driver types?

We have scoped the science test category for those simulations that can be compared directly to observations (i.e. from flux towers). Within the system tests are configurations that test for equivalency between drivers (image vs. classic). By the transitive relation, we can say that the image driver will behave just like the classic driver. Therefore, running the science tests again for other drivers is not necessary. However, the release tests do mainly target the image driver and provide the opportunity for additional verification through model evaluation. We have tried to clarify these points in the manuscript.

Page 9, lines 1-15: Authors may want to mention that further descriptions of the different test figures and configuration setup are found also in Section 3.

Thank you for the suggestion. We have pointed to section 3 in this case.

Page 9, line 16: The “Travis CI” system is outlined for how it is used with VIC workflow and testing. Can the authors provide some additional information on the origins and source of Travis?

Travis CI is a hosted service for testing open-source software. The TravisCI Wikipedia page (https://en.wikipedia.org/wiki/Travis_CI) and the “Travis CI About Page” (<https://about.travis-ci.com/>) provide good overviews of the service, its typical uses, and its adoption across the software development community.

Page 11, lines 16-19, and Figure 3: Please either indicate here that the VIC-5 version performs at or above 0.99 r^2 values relative to previous VIC-4.x version(s), as you did C3 with the observed SNOTEL peak SWE comparison in Figure 2, or include a third time series of the VIC-4 version to show overlap with VIC-5.

We have added the VIC.4 results to this figure. The reviewers may note that the r^2 values comparing VIC-4 and VIC-5 are 1.00 for all sites.

Page 11, lines 24-29: Please add to this discussion how the information in Table 3 may relate to the tested parallelization performance metrics shown in Figure 5.

We have elaborated on our previous point that as VIC is scaled out, it becomes more and more I/O bound. Therefore, a more performant disk system (e.g. fast SSDs) along with faster interconnects (e.g. infiniband) would likely yield improved scalability.

Figure 1: Please provide more information in the caption as to what is represented in the “Shared Driver Utilities”. This part of the figure was not well addressed in the figure or accompanying documentation.

We have added a brief paragraph in section 2.3 describing what we mean by the Shared Driver Utilities.

Figure 3: Mention in figure caption what the red and blue shading represent.

We have added a sentence describing the meaning of this shading (interannual variability represented as +/- 0.5 standard deviation).

Table 4: Replace “they” with “that” and don’t need the parentheses around “Liang et al. (1994)”, just 1994.

We have correct these two typos.