Response to reviewer 1

In the scalability discussion, I wonder if adding some cost analysis can enhance the understanding of the resulting curve seen in Figure 7? I provide what I meant as below, but it is up to authors to decide if such an extension is desirable. [...]

We thank the reviewer for this suggestion. We agree that adding a cost analysis improves the interpretation of the scalability curve and provides a useful way to estimate both computational cost and runtime.

We have therefore added this analysis in a new appendix, which holds as presented in the reviewer comments for the synchronous strategy. For the asynchronous strategy, the original formulation was incomplete. We have corrected and extended it accordingly. The revised manuscript now provides a consistent and accurate derivation that matches observed behaviour.

Related modifications on revised manuscript: Appendix B at lines 503-543. References to Appendix B at lines 311 and 334.

Following this addition, we also acknowledged the reviewer contribution in the acknowledgements section. We would like to ask whether the reviewer would be comfortable with removing anonymity so that its contribution can be explicitly named in the paper.

Related modifications on revised manuscript: Line 573

1. Line 80: “Another category”: I suggest “Second category” to enhance clarity.

We have modified the text accordingly.

Related modifications on revised manuscript: Line 80

2. Lines 221-222, Figures 5 and 6: I think “grids line” and “exchs line” are a bit awkward. Maybe adding line numbers to figures such that can directly refer as “code lines 6 and 7”.

We have revised the figure annotations to include explicit line numbers. Content of the revised manuscript now directly refers to the code lines as suggested by the reviewer, which improved readability.

Related modifications on revised manuscript: Figures 5 and 6. Lines 245 and 246.

3. Line 268-269: “the 2D cases are doubled”: needs clarification. What in 2D cases are doubled?

The intended meaning is that all 2D cases are executed twice: once using a pure CPU implementation and once using GPU-based inference. The text has been revised to explicitly state that the cases are “run twice” under these two execution modes.

Related modifications on revised manuscript: Lines 292.

4. For section 5.2 Future perspectives, I wonder if authors can provide some views on the recent effort to revive the usage of model hierarchy by Shaw et al (2025) which seems relevant?

We have extended Section 5.2 to reflect this perspective with the following paragraph:.

"In the long term, the need for bridging strategies (Section 2.1) may indeed be transitional. As discussed in Section 2.1, several components of Earth System Models (ESMs) are increasingly being implemented in high-level programming environments, sometimes assisted by AI-based coding tools. This trend suggests that entire ESMs may eventually be developed in environments such as JAX, Julia, or PyTorch.

In such a setting, cross-language bridging layers would no longer be necessary, as neural networks and numerical components could coexist within a unified runtime. We believe this direction aligns with the vision of hierarchical model integration discussed by Shaw et al. (2025), potentially enabling a broader paradigm shift where hierarchies of ESMs with varying levels of complexity can be more easily constructed and deployed."

Related modifications on revised manuscript: Paragraph at lines 393-398.

5. Figure 5: Is "ESM, = eophis.register_tunnels(…)" a typo? It looks like the comma on the left hand side is a bug. In python it should be "ESM, _ = …" or "ESM = …".

The reviewer points out a confusing notation. The function returns a list because it takes as input a list of Tunnel configurations. In the example, only one Tunnel is specified. The comma is used to unpack the first element of the returned list. The notation was indeed misleading. We now explicitly refer to the first element of this list for clarity in the code presentation.

Related modifications on revised manuscript: code line 10 on Figure 5.

6. Figure 1

- Since in the later example you are going to us variables u, v, maybe you can replace a, b with u, v to help the readers?
- I would suggest consistently use OASIS3-MCT as in the text throughout the manuscript, rather than "OASIS3" as if this is a different idea.
- I also think add text to differentiate two "OASIS API" boxes. Such as "OASIS3-MCT API (Python)" and "OASIS3-MCT API (C-Python adapter)". Current Figure 1 confuses me easily.

We thank the reviewer for these relevant suggestions aiming to improve both the consistency of the manuscript and clarity of Figure 1. We have modified Figure 1 to apply these recommendations directly. The caption has also been updated accordingly to better describe the workflow and the role of each component.

Related modifications on revised manuscript: figure 1 and caption of Figure 1.