Dear editor and reviewer,

First of all, we would like to express our sincere appreciation for your valuable feedback. Your review is not only highly insightful but also extremely meticulous. You have provided us with many important suggestions, and you have also pointed out numerous formatting errors in detail. Your comments will help us to substantially improve the quality of our manuscript. Below are our point-by-point responses to all the comments and our plans to revise the manuscript.

Major comments:

1. Please clarify the reinforcement learning methodology. Provide detailed information on the RL agent's architecture (e.g., algorithm type, reward function specifics, training duration, and hyperparameters). Specify how performance data (network speed, PnetCDF rates) is collected and preprocessed for the virtual environment. It would be helpful to add pseudocode or a flowchart for the RL training process.

[Response]:

In this work, we employ the Proximal Policy Optimization (PPO) model for reinforcement learning. The reward function is defined as the total time taken to complete all I/O operations. The training was conducted on an NVIDIA 3090 GPU and lasted approximately 30 minutes. Specific hyper-parameters and technical details will be explicitly outlined in the revised version of the paper. Additionally, we will include pseudo-code and a flowchart to illustrate the RL training process for better clarity.

2. The evaluation system can be expanded by including metrics beyond speedup, such as CPU/memory utilization during I/O phases, network overhead, or buffer management efficiency.

[Response]:

Thank you for the suggestion. We will conduct additional experiments and include these metrics (e.g., CPU/memory utilization, network overhead, and buffer management efficiency) in the revised version to provide a more comprehensive evaluation.

3. Some discussions about trade-offs between resource consumption and performance gains can be added in the revised version.

[Response]:

We appreciate your suggestion. In the revised version, we will add a detailed discussion on the trade-offs between resource consumption and performance gains to provide deeper insights into the efficiency of our approach.

4. Please explicitly address limitations of CFIO2.0, such as dependency on pre-collected data for RL, scalability across heterogeneous file systems (e.g., Lustre vs. ParaStor), or adaptability to non-climate modeling workloads.

[Response]:

Thank you for pointing this out. We will include a discussion on the limitations of CFIO2.0 in the revised version. However, we would like to clarify that the dependency on pre-collected data for RL training is optional. Even without RL pre-training, the default parameters can still achieve satisfactory performance. Regarding file system support, since our work is built on PnetCDF for

file I/O, the underlying logic and strategies are algorithm-based and not limited by the file system. Therefore, our method theoretically supports any file system that PnetCDF supports, ensuring compatibility with most file systems. Additionally, our work is specifically designed for models that using NetCDF format file as output. Any MPI parallel program using this file format can be supported, while other formats or non-MPI programs are not within the scope of this study. We will explicitly state these points in the revised manuscript.

5. Ensure that all parameters (e.g., stride values in PIO experiments) are explicitly defined in tables or appendices.
[Response]:
Thank you for the suggestion. We will carefully review the manuscript and ensure that all parameters, including stride values in PIO experiments, are explicitly defined in tables or appendices.

Minor suggestions:
1. P1. Line 14-15: it is unclear whether the increases in ouput efficiency are across the two models or two I/O strategies.
[Response]:
The 1.54x speedup refers to the overall runtime improvement of the GOMO model when using our method compared to not using it. The 13.1x speedup refers to the overall runtime improvement of the LICOM model when using our method compared to not using it. We will clarify this in the revised version.

2. Page 3, Line 70: "Figure 1: Concurrent I/O and Computation." → Align with figure numbering (e.g., "Figure 1: (a) Alternating... (b) Parallel...").
[Response]:
We will correct this to align with the figure numbering.

3. Page 18, Line 385: "LICOM case study" → Add a colon for consistency ("LICOM Case Study:").
[Response]:
We will add a colon for consistency.

4. Figure 3 Caption (Page 4, Lines 100–110): Move the lengthy step-by-step description from the caption to the main text or a supplementary note to improve readability.
[Response]:
Thank you for the suggestion. We will move the detailed step-by-step description from the caption to the main text or a supplementary note to enhance readability.

5. Section 5.3 (Page 19, Lines 410–415): Label subfigures in Figure 13 as "Fig. 13a" and "Fig. 13b" instead of "(a)" and "(b)" to avoid confusion.
[Response]:
We will update the labels to "Fig. 13a" and "Fig. 13b" for clarity.

6. References (Page 23–24): Correct "Corbetty et al. (1996)" to "Corbett et al. (1996)" in the text (Page 2, Line 45).
[Response]:
We will correct this typo in the revised version.

7. Ensure all citations (e.g., "Kang et al. (2019)") have corresponding entries in the References section.
[Response]:
We will verify that all citations have corresponding entries in the References section.

8. Code Availability (Page 22, Lines 525–530): Verify that all Zenodo links are functional and datasets are publicly accessible.
[Response]:
We have verified that all Zenodo links are functional, and the datasets are publicly accessible.

We really appreciate your highly constructive comments.

Best wishes,
Dong Wang, Xiaomeng Huang