Review of manuscript egusphere-2026-1538

Summary

The authors present a new GPU-enabled aerosol module MAM4xx for the Earth system model EAMxx, derived from the four-mode module MAM4. They document their development steps: refactoring the Fortran code, porting the code to C++/Kokkos, and integrating the code into EAMxx. They evaluate the performance of EAMxx-MAM4xx with a global simulation at 12-kilometer resolution over 33 days. The paper is well written, well structured, and provides valuable insights into the development process and model performance. I have three general comments and a few specific comments.

General comments

- In the introduction, the authors should mention recent developments in kilometer-scale Earth system modeling. Segura et al. (2025) presented Earth system simulations with ICON and IFS-FESOM at 10-kilometer resolution in the atmosphere over 30 years. Klocke et al. (2025) presented simulations with ICON over short periods at 1-kilometer resolution. Weiss et al. (2025) presented a simulation with ICON coupled to the aerosol module HAM-lite at 5-kilometer resolution.
    - https://gmd.copernicus.org/articles/18/7735/2025
    - https://dl.acm.org/doi/10.1145/3712285.3771789
    - https://gmd.copernicus.org/articles/18/3877/2025
- In the methods, the authors could include a figure that shows the interface of EAMxx and MAM4xx and the time stepping. That would help the reader to understand the ordering and handling of processes and the computational costs.
- In the evaluation, the authors could include a global overview of the simulation. Quantities like global burdens, lifetimes, and optical depths would help the reader to get a basic understanding of how well the global aerosol cycle is represented.

Specific comments

- In the introduction, the authors could mention that GPU nodes are more energy efficient than CPU nodes.
- In line 117-118, are there plans to account for ice-nucleating particles as well?
- In line 220-221, how are the loops over columns parallelized in the new code?
- In line 283-289, how large are the conservation errors in the new code?
- In section 3.4, how do the computational costs scale with resolution? If not too much work, the authors could perform two additional tests at resolutions of 6 and 3 kilometers and evaluate how the total and relative costs change.

- In line 316-317, is the vertical mixing of aerosols still physically consistent with the vertical mixing of other species like water?

- In line 319-320, could you provide possible explanations for why the relative computational costs depend on the hardware?

- In line 339-341, how do the emissions of sea salt and dust change with resolution?

- In line 440-441, how flexible is the aerosol module? How much work would it be to include those species?