

Response to Reviewers' comments

We thank the two reviewers for their thoughtful and constructive comments that helped us substantially improve the manuscript. We have revised the manuscript accordingly. Listed below is our blue point-to-point response to each comment the reviewers offered.

Response to Reviewer #1

General Comments

The authors present a hardware-accelerated implementation of the EPICC air quality model's horizontal advection scheme, demonstrating impressive computational performance improvements achieved through clever and careful source code translation and refactoring, although these improvements are offset by data transfer overhead. The authors couple the accelerated advection module back into the full model and demonstrate that this heterogeneous model implementation retains code integrity to produce reasonable simulation results. Moreover, the authors present a thorough performance analysis. Taken together, this is a comprehensive and impressive effort that may offer an impetus for the air quality modeling community to collectively modernize model codebases towards enabling the kind of ultra-high-resolution simulations that are necessary to characterize urban- and neighborhood-scale pollution, the scales at which communities experience harmful exposures. My comments are brief and minor.

Response: We sincerely thank the referee for the positive and encouraging assessment of our work. We also appreciate the referee's constructive comments, which have helped us improve the clarity and technical completeness of the manuscript. Our point-by-point responses are provided below.

Specific comments

Line 130: Please include, along with brief term definitions, the species continuity equation explicitly.

Response: Thank you for this helpful suggestion. We agree that explicitly presenting

the species continuity equation and defining its terms will improve the clarity and completeness of the model description. We will add the species continuity equation and definitions of the relevant terms in Sect. 2.1 in **lines 136-141**, as follows:

The model framework is fundamentally based on the following species continuity equation and is used to simulate the complex physical and chemical processes of pollutants in the atmosphere.

$$\begin{aligned} \frac{\partial C_i}{\partial t} + u \frac{\partial C_i}{\partial x} + v \frac{\partial C_i}{\partial y} + \omega \frac{\partial C_i}{\partial z} \\ = \frac{\partial}{\partial x} \left(K_x \frac{\partial C_i}{\partial x} \right) + \frac{\partial}{\partial y} \left(K_y \frac{\partial C_i}{\partial y} \right) + \frac{\partial}{\partial z} \left(K_z \frac{\partial C_i}{\partial z} \right) + S + P - R_d - W_{ash} \end{aligned}$$

where C_i is the concentration of pollutant species i , u , v , and ω are the wind-velocity components in the x , y , and z directions, respectively. K_x , K_y , and K_z are the turbulent diffusion coefficients in the longitudinal, latitudinal, and vertical directions, respectively. S represents the emission, P represents chemical transformation, R_d represents dry deposition, and W_{ash} represents wet scavenging.

Line 217, 599-600: Considering the substantial CPU-GPU data transfer overhead, could a cluster deploying EPICC with GPU-HADVPPM4HIP V1.0 on a cluster with PCIe 5.0 buses, offering double the bandwidth of PCIe 4.0, expect to see additional efficiency improvements? Please include a brief discussion of hardware outlook. Regardless, I commend the authors' focus on software optimization rather than relying purely on hardware capabilities, so that the model is maximally performant across diverse hardware and not only the latest generation.

Response: Thank you for this valuable and forward-looking comment. We fully agree with the referee that upgrading the interconnect from PCIe 4.0 to PCIe 5.0 would improve CPU-GPU communication efficiency. CPU-GPU communication is one of the key factors limiting model performance on heterogeneous clusters. In this study, we reconstructed the loop and array structures of the advection module, reducing the number of CPU-GPU data transfers from approximately 9.8 million to only 10 per time step. In future work, we will further investigate chunked transfers using pinned

memory, shared-memory and common-field caching, asynchronous communication, and unified memory to reduce data-transfer overhead and overlap CPU-GPU communication with GPU computation.

In addition to these software-level optimizations, advances in hardware interconnects provide a complementary means of improving CPU-GPU communication efficiency. PCI Express (PCIe) is a general-purpose, high-speed interconnect widely used for data exchange between host CPUs and accelerator devices. Its nominal per-lane raw data rate has increased from 16 GT/s with PCIe 4.0 to 32 GT/s with PCIe 5.0 and 64 GT/s with PCIe 6.0, while the PCIe 7.0 specification further increases it to 128 GT/s (Lender, 2019; Yanes, 2022; PCI-SIG, 2025). The EarthLab heterogeneous cluster used in this study is currently equipped with PCIe 4.0 buses; therefore, we cannot directly evaluate the performance improvement offered by PCIe 5.0. If access to a heterogeneous computing platform equipped with PCIe 5.0 or a later-generation interconnect becomes available, we will systematically evaluate the performance of the EPICC-Model coupled with GPU-HADVPPM4HIP V1.0 on that platform. By combining software optimization with advances in hardware interconnects, we aim to further reduce CPU-GPU communication overhead and improve model performance across different generations of heterogeneous computing platforms. We have added a discussion of PCIe development in **lines 849–864**, as follows:

First, future work will further improve CPU-GPU communication efficiency through chunked transfers using pinned memory, shared-memory and common-field caching, asynchronous communication, and unified memory. These strategies will reduce repeated data transfers and increase the overlap between CPU-GPU communication and GPU computation. In addition to these software-level optimizations, advances in hardware interconnects provide a complementary means of improving CPU-GPU communication efficiency. PCI Express (PCIe) is a general-purpose, high-speed interconnect widely used for data exchange between host CPUs and accelerator devices. Its nominal per-lane raw data rate has increased from 16 GT/s with PCIe 4.0 to 32 GT/s with PCIe 5.0 and 64 GT/s with PCIe 6.0, while the PCIe 7.0 specification

further increases it to 128 GT/s (Lender, 2019; Yanes, 2022; PCI-SIG, 2025). The EarthLab heterogeneous cluster used in this study is currently equipped with PCIe 4.0 buses; therefore, the performance improvement offered by PCIe 5.0 cannot be directly evaluated in the present study. If access to a heterogeneous computing platform equipped with PCIe 5.0 or a later-generation interconnect becomes available, we will evaluate GPU-HADVPPM4HIP V1.0 on that platform and further improve CPU-GPU communication efficiency through hardware-software co-optimization.

Technical corrections

Line 45: “Air pollution, a source of...” may be better worded as “Air pollution, including...”

Response: Thanks for this correction. We have revised the sentence accordingly at the original **line 45**, which are as follows:

Air pollution, including fine particulate matter in both urban and rural areas, is associated with an elevated risk of strokes, heart diseases, lung cancer, and acute and chronic respiratory diseases (Atkinson et al., 2010; Kim et al., 2015; Liu et al., 2016; Milton and White, 2020).

Line 83: Please define “HPL score” .

Response: Thank you for pointing this out. We will define HPL and briefly explained what the benchmark measures at the **lines 83-86**. The sentence has been revised as follows:

In December 2024, the 64th TOP500 list ranked the El Capitan system first, with a performance of 1.742 EFLOP/s in the High Performance Linpack (HPL) benchmark, which measures the sustained double-precision floating-point performance achieved when solving a dense system of linear equations (TOP500, 2024).

Line 92: Reword “enabling automated generates...” to “enabling automated generation of...” .

Response: Thank you for identifying this grammatical error. We have revised the

sentence at the original **line 95** as follows:

Alvanos and Christoudias (2017) developed a software package for the global atmospheric chemistry model ECHAM/MESSy Atmospheric Chemistry (EMAC), enabling automated generation of CUDA kernels for the numerical integration of atmospheric chemical kinetics using the Kinetic PreProcessor (KPP; Damian et al., 2002).

Line 108: “Air quality models-the advection module...” may be scan better with a comma instead of a hyphen.

Response: Thank you for this suggestion. We have replaced the hyphen with a comma at **lines 111-112**, which are as follows:

Regarding the advection module, another computational hotspot in air quality models, Cao et al. (2023, 2024) implemented GPU acceleration for the CAMx advection module using CUDA and HIP technologies, respectively.

Line 127: “Air pollution complex in China” maybe is meant to be “complex air pollution in China” ?

Response: Thanks for this correction. We have changed “air pollution complex in China” to “complex air pollution in China” at **line 131**, which are as follows:

The emission and atmospheric Processes Integrated and Coupled Community Model version 1.6.0 (EPICC-Model V1.6.0; EPICC-Model Working Group, 2025; Wang et al., 2025) is an air quality modeling system specifically designed for complex air pollution in China (Zhu et al., 2023).

Line 133: Change “developed using Fortran program language...” to include “the: between “using” and “Fortran” .

Response: Thanks for this correction. We have revised the sentence at **line 144**, which are as follows:

The EPICC-Model V1.6.0 adopts a modular architecture developed using the Fortran programming language, a high-performance computing language designed for

scientific applications.

Figure 1: Wet scavenging typo for RADM; aqueous-phase chemistry seems to reference the vertical diffusion schemes.

Response: Thank you for pointing out these errors, and we apologize for the plotting and spelling mistakes. We have corrected the corresponding entries in **Figure 1** as follows:

Wet scavenging: Regional Acid Deposition Model (RADM; Chang et al., 1987).

Aqueous-phase chemistry: Regional Acid Deposition Model (RADM; Chang et al., 1987).




Figure 1. The computational framework and workflow of the EPICC-Model V1.6.0. In the section of physical-chemical processes, yellow represents the physical module, and orange indicates the chemical module.

Line 216: Change “two China’ s domestic CPU…” to “two of China’ s domestic CPU…”

Response: Thank you for this correction. We have revised the sentence at the original lines 241-242 as follows:

As summarized in Table 1, each GPU-like computing node contains two of China’s domestic CPU processors and two GPU-like accelerators connected through PCIe

4.0 buses.

Response to Reviewer #2

The authors have addressed all of my major comments from the original submission. The revised manuscript shows significant improvements. I believe that GPU-accelerated computing will be important in the future if the data transfer bottleneck between the host and device can be resolved. Future generations of accelerated devices and CPUs are expected to be more tightly integrated.

Response: We sincerely thank the referee for the careful re-evaluation and positive assessment of the revised manuscript. We also appreciate the constructive comments, which have helped us improve the presentation of the model evaluation results. Our point-by-point responses are provided below.

Figures 6 and 7 have not been referenced or discussed in the main text.

Response: Thanks for pointing this out. We have now explicitly referenced and discussed Figures 6 and 7 in Sect. 4.2.1. Visually, the spatial distributions and concentration magnitudes calculated by the Fortran and HIP-Opt2 versions are nearly identical. Further comparison shows that the differences in HONO concentrations are within ± 0.0001 ppbV, while the differences for the other species are within ± 0.001 ppbV for gases or $\pm 0.001 \mu\text{g} \cdot \text{m}^{-3}$ for aerosols. Cao et al. (2024) reported that the differences between the HIP and Fortran versions of CAMx were within ± 0.01 ppbV for O_3 and NO_2 and within $\pm 0.005 \mu\text{g} \cdot \text{m}^{-3}$ for ASO_4 . By comparison, in the present study, the differences in O_3 and NO_2 are within ± 0.001 ppbV, while the difference in ASO_4 is within $\pm 0.001 \mu\text{g} \cdot \text{m}^{-3}$, all of which are smaller than the corresponding differences reported by Cao et al. (2024). For each species, the grid-cell differences were further averaged. The resulting mean differences are 5.1×10^{-9} , 3.6×10^{-7} , 1.7×10^{-6} , and 3.2×10^{-4} ppbV for HONO, SO_2 , NO_2 , and O_3 , respectively, and

1.1×10^{-6} , 1.8×10^{-5} , 4.0×10^{-6} , and $1.0 \times 10^{-6} \mu\text{g} \cdot \text{m}^{-3}$ for BC, PM_{2.5}, ASO₄, and ANH₄, respectively. Taken together, the smaller difference ranges and near zero domain mean differences confirm the close numerical agreement between the Fortran and HIP-Opt2 versions. We have added this part in **lines 490-503**, which are as follows:

Figures 6 and 7 compare the gas and aerosol concentrations calculated by the original Fortran and HIP-Opt2 versions after the 24 h integration. Visually, the two versions produce nearly identical spatial distributions and concentration magnitudes. The differences in HONO concentrations are within ± 0.0001 ppbV, while the differences for the other species are within ± 0.001 ppbV for gases or $\pm 0.001 \mu\text{g} \cdot \text{m}^{-3}$ for aerosols. Cao et al. (2024) reported that the differences between the HIP and Fortran versions of CAMx were within ± 0.01 ppbV for O₃ and NO₂ and within $\pm 0.005 \mu\text{g} \cdot \text{m}^{-3}$ for ASO₄. By comparison, in the present study, the differences in O₃ and NO₂ are within ± 0.001 ppbV, while the difference in ASO₄ is within $\pm 0.001 \mu\text{g} \cdot \text{m}^{-3}$, all of which are smaller than the corresponding differences reported by Cao et al. (2024). For each species, the grid-cell differences were further averaged. The resulting mean differences are 5.1×10^{-9} , 3.6×10^{-7} , 1.7×10^{-6} , and 3.2×10^{-4} ppbV for HONO, SO₂, NO₂, and O₃, respectively, and 1.1×10^{-6} , 1.8×10^{-5} , 4.0×10^{-6} , and $1.0 \times 10^{-6} \mu\text{g} \cdot \text{m}^{-3}$ for BC, PM_{2.5}, ASO₄, and ANH₄, respectively. Taken together, the smaller difference ranges and near zero domain mean differences confirm the close numerical agreement between the Fortran and HIP-Opt2 versions.

Figure 6, the caption should be expanded to describe the left, middle, and right panels. Specify the units of color bar (ppbV, ugm-3, %). Reduce the color bar scale to see the difference between the two models. Include O3 in your evaluation because it is an important air pollutant.

Response: Thanks for this helpful suggestion. We have redrawn Figures 6 and 7 to provide a clearer comparison between the two model versions. In both figures, the left, middle, and right columns show the Fortran results, HIP-Opt2 results, and differences between the two versions, respectively. The units are now marked in the upper-right

corner of every panel. To make the small differences more distinguishable, we reduced the color-bar ranges of the difference panels from the original ± 0.1 to ± 0.0001 for HONO and ± 0.001 for the other species, with units of ppbV for gases and $\mu g \cdot m^{-3}$ for aerosols. In addition, NH_3 has been replaced by O_3 in Figure 6, and the corresponding O_3 difference has been included. The captions of Figures 6 and 7 have also been expanded to describe the left, middle, and right panels and their units. The updated **Figures 6 and 7** are provided below.




Figure 6. Spatial distributions of HONO, SO₂, NO₂, and O₃ concentrations calculated by the original Fortran and HIP-Opt2 versions after the 24 h integration. Panels (a), (d), (g), and (j) show the Fortran results; panels (b), (e), (h), and (k) show the HIP-Opt2 results; and panels (c), (f), (i), and (l) show the differences between the two versions. The units of the concentrations and differences are ppbV.




Figure 7. Spatial distributions of BC, PM_{2.5}, ASO₄, and ANH₄ concentrations calculated by the original Fortran and HIP-Opt2 versions after the 24 h integration. Panels (a), (d), (g), and (j) show the Fortran results; panels (b), (e), (h), and (k) show the HIP-Opt2 results; and panels (c), (f), (i), and (l) show the differences between the two versions. The units of the concentrations and differences are $\mu\text{g} \cdot \text{m}^{-3}$.

Tables 5 and 6, include the normalized mean bias, as it is an important metric. Since each species has different concentration range. For example, the mean concentration of O₃ is about 30 ppbV, while HONO is around 1 ppbV.

Response: Thanks for this important suggestion. We agree that normalized mean bias (NMB) is useful for comparing species with substantially different concentration ranges. We have therefore added NMB to Tables 5 and 6 and incorporated its

interpretation into the corresponding discussions. For Table 5, NMB represents the overall signed difference of the HIP-Opt2 results relative to the Fortran results. The NMB values range from $-1.0 \times 10^{-4}\%$ for HONO to $2.9 \times 10^{-3}\%$ for ASO₄, indicating negligible systematic bias between the two model versions. For Table 6, NMB represents the normalized difference between the simulated and observed concentrations. Positive and negative NMB values indicate model overestimation and underestimation, respectively. We recalculated RMSE, r, and NMB using the same city-level simulated and observed time series and updated both Table 6 and the accompanying discussion. The contents of the above two parts have been revised in **lines 518-529** and **lines 593-608**, which are as follows:

Lines 518-529:

To further evaluate the suitability of HIP-Opt2 for scientific applications, we followed Wang et al. (2021) and Cao et al. (2024) and calculated the root mean square error (RMSE), standard deviation (std), RMSE/std ratio, and normalized mean bias (NMB). RMSE quantifies the differences between HIP-Opt2 and the original Fortran version, while std represents the spatial variability of each species in the Fortran results. The RMSE/std ratio therefore measures the magnitude of the computational differences relative to the spatial variability of the corresponding species. NMB measures the overall signed difference of the HIP-Opt2 results relative to the Fortran results. As shown in Table 5, the RMSE/std ratios range from $6.9 \times 10^{-4}\%$ for NO₂ to $9.2 \times 10^{-2}\%$ for ANH₄. The NMB values range from $-1.0 \times 10^{-4}\%$ for HONO to $2.9 \times 10^{-3}\%$ for ASO₄. The small RMSE/std ratios and near-zero NMB values demonstrate that the computational differences introduced by the HIP implementation are negligible relative to the spatial variability and concentration magnitudes of the evaluated species.

Table 5. Root mean square error (RMSE) between the Fortran and HIP-Opt2 versions, standard deviation (std) of the Fortran results, RMSE/std ratio, and normalized mean bias (NMB) of HIP-Opt2 relative to the Fortran version.

	RMSE	std	RMSE/std (%)	NMB (%)
--	------	-----	--------------	---------

<i>HONO (ppbV)</i>	1.1×10^{-7}	<i>0.01</i>	1.1×10^{-3}	-1.0×10^{-4}
<i>SO₂ (ppbV)</i>	1.2×10^{-5}	<i>0.5</i>	2.4×10^{-3}	1.9×10^{-4}
<i>NO₂ (ppbV)</i>	1.8×10^{-5}	<i>2.7</i>	6.9×10^{-4}	1.9×10^{-4}
<i>O₃ (ppbV)</i>	3.6×10^{-4}	<i>14.2</i>	2.5×10^{-3}	9.6×10^{-4}
<i>BC ($\mu\text{g} \cdot \text{m}^{-3}$)</i>	4.5×10^{-6}	<i>0.1</i>	3.6×10^{-3}	2.0×10^{-3}
<i>PM_{2.5} ($\mu\text{g} \cdot \text{m}^{-3}$)</i>	4.3×10^{-4}	<i>2.1</i>	2.2×10^{-2}	2.0×10^{-3}
<i>ASO₄ ($\mu\text{g} \cdot \text{m}^{-3}$)</i>	2.1×10^{-5}	<i>0.3</i>	7.9×10^{-3}	2.9×10^{-3}
<i>ANH₄ ($\mu\text{g} \cdot \text{m}^{-3}$)</i>	9.6×10^{-5}	<i>0.1</i>	9.2×10^{-2}	2.4×10^{-3}

Lines 593-608:

Figures 8 and 9, together with Table 6, present the time-series comparisons and statistical evaluation of the simulated and observed PM_{2.5} and O₃ concentrations. The formulas used to calculate the statistical metrics are provided in the Supplement. Positive and negative NMB values indicate model overestimation and underestimation, respectively. For daily PM_{2.5}, the EPICC-Model reproduced the temporal variations reasonably well in most cities. Fuzhou had an RMSE of $3.9 \mu\text{g} \cdot \text{m}^{-3}$, an *r* of 0.87, and an NMB of -14.6%. Across the eight cities, the mean RMSE, *r*, and city-level NMB were $9.4 \mu\text{g} \cdot \text{m}^{-3}$, 0.70, and 19.7%, respectively. The positive NMB values in Beijing and Hefei, 89.1% and 74.0%, indicate overestimation, whereas the negative NMB values in Guangzhou and Fuzhou, -28.1% and -14.6%, indicate underestimation. For MDA8 O₃, the mean RMSE, *r*, and city-level NMB were 12.1 ppbV, 0.76, and -17.7%, respectively, indicating an overall tendency toward underestimation. The largest negative NMB values occurred in Taiyuan and Hangzhou, at -41.5% and -34.3%, respectively, while Qingdao showed a small positive bias of 4.7%. Qingdao and Guangzhou had *r* values of 0.91 and 0.87 and RMSE values of 12.2 and 7.8 ppbV, respectively. These statistical results indicate that the EPICC-Model V1.6.0, integrated with GPU-HADVPPM4HIP V1.0, reasonably reproduces the spatiotemporal characteristics of PM_{2.5} and O₃ concentrations on China's domestic heterogeneous computing clusters.

Table 6. Statistical evaluation of the PM_{2.5} and O₃ simulations of the EPICC-Model in eight Chinese cities. NMB denotes normalized mean bias. The average row represents

the arithmetic mean of the city-level metrics.

	<i>PM_{2.5}</i>			<i>O₃</i>		
	<i>RMSE</i> <i>($\mu\text{g} \cdot \text{m}^{-3}$)</i>	<i>r</i>	<i>NMB</i> <i>(%)</i>	<i>RMSE</i> <i>(ppbV)</i>	<i>r</i>	<i>NMB</i> <i>(%)</i>
<i>Beijing</i>	<i>17.3</i>	<i>0.84</i>	<i>89.1</i>	<i>13.2</i>	<i>0.77</i>	<i>-7.0</i>
<i>Taiyuan</i>	<i>11.0</i>	<i>0.54</i>	<i>18.6</i>	<i>21.6</i>	<i>0.78</i>	<i>-41.5</i>
<i>Hangzhou</i>	<i>7.4</i>	<i>0.50</i>	<i>-3.7</i>	<i>14.8</i>	<i>0.80</i>	<i>-34.3</i>
<i>Hefei</i>	<i>12.4</i>	<i>0.55</i>	<i>74.0</i>	<i>14.2</i>	<i>0.70</i>	<i>-15.3</i>
<i>Fuzhou</i>	<i>3.9</i>	<i>0.87</i>	<i>-14.6</i>	<i>7.5</i>	<i>0.77</i>	<i>-17.0</i>
<i>Qingdao</i>	<i>11.0</i>	<i>0.90</i>	<i>20.6</i>	<i>12.2</i>	<i>0.91</i>	<i>4.7</i>
<i>Guangzhou</i>	<i>7.2</i>	<i>0.65</i>	<i>-28.1</i>	<i>7.8</i>	<i>0.87</i>	<i>-14.9</i>
<i>Kunming</i>	<i>5.3</i>	<i>0.76</i>	<i>2.0</i>	<i>6.0</i>	<i>0.46</i>	<i>-16.4</i>
<i>Average</i>	<i>9.4</i>	<i>0.70</i>	<i>19.7</i>	<i>12.1</i>	<i>0.76</i>	<i>-17.7</i>

Provide the definition of GPU-like accelerator.

Response: Thanks for requesting clarification of this terminology. In this study, a GPU-like accelerator refers to a domestically developed general-purpose accelerator in China that performs the same massively parallel acceleration role as a conventional GPU and supports GPU-style programming through the HIP framework. Its architecture and programming model are similar to those of AMD GPUs, including the single-instruction multiple-thread (SIMT) execution model and the grid-block-thread parallel hierarchy. The GPU-like accelerator used in this study has 3840 computing units and 16 GB of device memory, as summarized in Table 1. The term “GPU-like accelerator” follows Cao et al. (2024), who used HIP to port and accelerate the PPM advection solver of the CAMx air quality model on the same type of domestically developed accelerator in China. The broader applicability of HIP to geoscientific models was also demonstrated by Wang et al. (2021), who developed and applied the HIP version of the LICOM3 ocean model on large-scale AMD GPU systems. Accordingly, we have clarified the definition and application context of GPU-like accelerators in Sect. 2.2. We have added this part in **lines 232-248**, which are as follows:

All performance testing of the EPICC-Model V1.6.0 and the heterogeneous adaptation and optimization of its advection module were conducted at the Earth System Numerical Simulation Facility (EarthLab; Chai et al., 2021). Jointly

developed by the Institute of Atmospheric Physics, Chinese Academy of Sciences, and collaborating institutions, EarthLab is designed for Earth system modeling and high-resolution regional environmental simulations and employs a CPU-GPU heterogeneous architecture. Within this platform, the term “GPU-like accelerator” refers to a domestically developed general-purpose accelerator in China that performs the same massively parallel acceleration role as a conventional GPU and supports GPU-style programming through the HIP framework. The GPU-like accelerator used in this study has 3,840 computing units and 16 GB of device memory. As summarized in Table 1, each GPU-like computing node contains two of China’s domestic CPU processors and two GPU-like accelerators connected through PCIe 4.0 buses, and both the CPUs and GPU-like accelerators are first-generation devices. Cao et al. (2024) previously used HIP to accelerate the PPM advection solver of the CAMx air quality model on the same type of domestic GPU-like accelerator, while Wang et al. (2021) demonstrated the large-scale application of the HIP-based LICOM3 ocean model. The software environment includes the Intel OneAPI 2021.3.0 toolkit for compiling the CPU code and the dtk-23.04.1 toolkit for compiling the GPU-like code.

References:

- Cao, K., Wu, Q., Wang, L., Guo, H., Wang, N., Cheng, H., Tang, X., Li, D., Liu, L., Li, D., Wu, H., and Wang, L.: GPU-HADVPPM4HIP V1.0: using the heterogeneous-compute interface for portability (HIP) to speed up the piecewise parabolic method in the CAMx (v6.10) air quality model on China’s domestic GPU-like accelerator, *Geosci. Model Dev.*, 17, 6887–6901, <https://doi.org/10.5194/gmd-17-6887-2024>, 2024.
- Chang, J. S., Brost, R. A., Isaksen, I. S. A., Madronich, S., Middleton, P., Stockwell, W. R., and Walcek, C. J.: A three-dimensional Eulerian acid deposition model: Physical concepts and formulation, *J. Geophys. Res.-Atmos.*, 92, 14681–14700,

<https://doi.org/10.1029/JD092iD12p14681>, 1987.

Lender, K.: *PCI-SIG in 2019: A year in review*, PCI-SIG, <https://pcisig.com/blog/pci-sig%C2%AE-2019-year-review> (last access: 6 August 2026), 2019.

PCI-SIG: *The PCIe 7.0 specification, Version 1.0 is now available to members*, PCI-SIG, <https://pcisig.com/specifications/pcie-70-specification-version-03-now-available-members> (last access: 6 August 2026), 2025.

Wang, P., Jiang, J., Lin, P., Ding, M., Wei, J., Zhang, F., Zhao, L., Li, Y., Yu, Z., Zheng, W., Yu, Y., Chi, X., and Liu, H.: *The GPU version of LASG/IAP Climate System Ocean Model version 3 (LICOM3) under the heterogeneous-compute interface for portability (HIP) framework and its large-scale application*, *Geosci. Model Dev.*, 14, 2781–2799, <https://doi.org/10.5194/gmd-14-2781-2021>, 2021.

Yanes, A.: *PCIe 6.0 specification released to members: Double the bandwidth for next-generation applications*, PCI-SIG, <https://pcisig.com/blog/pcie%C2%AE-60-specification-released-members-double-bandwidth-next-generation-applications> (last access: 6 August 2026), 2022.