

1 **GPU-HADVPPM4HIP V1.0: using the heterogeneous interface for**
2 **portability (HIP) to speed up the piecewise parabolic method in the**
3 **CAMx (v6.10) air quality model on China's domestic GPU-like**
4 **accelerator**

5 **Kai Cao¹, Qizhong Wu^{1,5}, Lingling Wang², Hengliang Guo³, Nan Wang², Huaqiong Cheng^{1,5},**
6 **Xiao Tang⁴, Dongxing Li^{1,5}, Lina Liu³, Dongqing Li¹, Hao Wu³, and Lanning Wang^{1,5}**

7 ¹College of Global Change and Earth System Science, Faculty of Geographical Science, Beijing
8 Normal University, Beijing 100875, China

9 ²Henan Ecological Environmental Monitoring Centre and Safety Center, Henan Key Laboratory
10 of Environmental Monitoring Technology, Zhengzhou 450008, China

11 ³National Supercomputing Center in Zhengzhou, Zhengzhou, 450001, China

12 ⁴State Key Laboratory of Atmospheric Boundary Layer Physics and Atmospheric Chemistry,
13 Institute of Atmospheric Physics, Chinese Academy of Science, Beijing 100029, China

14 ⁵Joint Center for Earth System Modeling and High Performance Computing, Beijing Normal
15 University, Beijing, 100875, China

16

17 **Correspondence to:** Qizhong Wu (wqizhong@bnu.edu.cn); Lingling Wang(928216422@qq.com);
18 Lanning Wang (wangln@bnu.edu.cn)

19

20 **Abstract.** The graphics processing units (GPUs) are becoming a compelling acceleration strategy
21 for geoscience numerical model due to their powerful computing performance. In this study,
22 AMD's heterogeneous compute interface for portability (HIP) was implemented to port the GPU
23 acceleration version of the Piecewise Parabolic Method (PPM) solver (GPU-HADVPPM) from
24 the NVIDIA GPUs to China's domestically GPU-like accelerators as GPU-HADVPPM4HIP, and
25 further introduced the multi-level hybrid parallelism scheme to improve the total computational
26 performance of the HIP version of CAMx (CAMx-HIP) model on the China's domestically
27 heterogeneous cluster. The experimental results show that the acceleration effect of GPU-
28 HADVPPM on the different GPU accelerator is more obvious when the computing scale is larger,

29 and the maximum speedup of GPU-HADVPPM on the domestic GPU-like accelerator is 28.9
30 times. The hybrid parallelism with a message passing interface (MPI) and HIP enables achieve up
31 to 17.2 times speedup when configure 32 CPU cores and GPU-like accelerators on the domestic
32 heterogeneous cluster. And the OpenMP technology is introduced to further reduce the
33 computation time of CAMx-HIP model by 1.9 times. More importantly, by comparing the
34 simulation results of GPU-HADVPPM on NVIDIA GPUs and domestic GPU-like accelerators, it
35 is found that the simulation results of GPU-HADVPPM on domestic GPU-like accelerators have
36 less difference than the NVIDIA GPUs, and the reason for this difference may be related to the
37 fact that the NVIDIA GPU loss part of the accuracy for improved computing performance.
38 Furthermore, we also exhibit that the data transfer efficiency between CPU and GPU has an
39 important impact on heterogeneous computing, and point out that optimizing the data transfer
40 efficiency between CPU and GPU is one of the important directions to improve the computing
41 efficiency of geoscience numerical models in heterogeneous clusters in the future.

42 **1. Introduction**

43 Over the recent years, GPUs have become an essential part of providing processing power for
44 high performance computing (HPC) application, and heterogeneous supercomputing based on
45 CPU processors and GPU accelerators has become the trend of global advanced supercomputing
46 development. The 61st edition of the top 10 list, released in June 2023, reveals that 80% of
47 advanced supercomputers adopt the heterogeneous architectures (Top500, 2023), and the Frontier
48 system equipped with AMD Instinct MI250X GPU at the Oak Ridge National Laboratory remains
49 the only true exascale machine with the High-Performance Linpack benchmark (HPL) score of
50 1.194 Exaflop/s (News, 2023). How to realize the large-scale parallel computing and improve the
51 computational performance of geoscience numerical models on the GPU has become one of the
52 significant directions for the future development of numerical models.

53 In terms of the heterogeneous porting for air quality model, most scholars select the chemical
54 module, one of the hotspots, to implement heterogeneous porting, and porting the computational
55 process originally on the CPU processes to the GPU accelerator, in order to improve the
56 computing efficiency. For example, Sun et al. (2018) used CUDA technology to port the second-

57 order Rosenbrock solver of chemistry module of CAM4-Chem to NVIDIA Tesla K20X GPU, and
58 achieved up 11.7x speedup compared to the AMD Opteron™ 6274 (Interlagos) CPU (16 cores)
59 using one CPU core. Alvanos and Christoudias (2017) developed a software that automatically
60 generates CUDA kernels to solve chemical kinetics equation in the chemistry module for the
61 global climate model ECHAM/MESSy Atmospheric Chemistry (EMAC) and performance
62 evaluation shows a 20.4x speedup for the kernel execution. Linford et al. (2011) presented the
63 Kinetic PreProcessor (KPP) to generate the chemical mechanism code in CUDA language which
64 can be implemented on NVIDIA Tesla C1060 GPU. The KPP-generated SAPRC'99 mechanism
65 from CMAQ model achieved a maximum speedup of 13.7x and KPP-generated RADM2
66 mechanism from WRF-chem model achieved an 8.5x speedup both compared to the Intel Quad-
67 Core Xeon 5400 series CPU. Similarly, the advection module is also one of the hotspot modules in
68 the air quality model, Cao et al. (2023) adopted the Fortran-C-CUDA C scheme and implemented
69 a series of optimizations, including reduction the CPU-GPU communication frequency, optimize
70 the GPU memory access, and thread and block co-indexing, to increase the computational
71 efficiency of the HADVPPM advection solver. It can achieve up to the 18.8x speedup on the
72 NVIDIA Tesla V100 GPU compared to the Intel Xeon Platinum 8168 CPU.

73 The CUDA technology was implemented to carry out heterogeneous porting for the
74 atmospheric chemical models from the CPU processors to different NVIDIA GPU accelerators. In
75 this study, the Heterogeneous-computing Interface for Portability (HIP) interface was introduced
76 to implement the porting of GPU-HADVPPM from the NVIDIA GPU to the China's domestically
77 GPU-like accelerators based on the research of Cao et al. (2023). The domestic GPU-like
78 accelerator plays the same role as the NVIDIA GPU, which is also used to accelerate the
79 advection module in the CAMx model, so we refer to it as a GPU-like accelerator. First, we
80 compared the simulation results of the Fortran version CAMx model with the CAMx-CUDA and
81 CAMx-HIP model which were coupled with the CUDA and HIP versions of GPU-HADVPPM
82 program, respectively. And then, the computing performance of GPU-HADVPPM programs on
83 different GPUs were compared. Finally, we tested total coupling performance of CAMx-HIP
84 model with multi-level hybrid parallelization on the China's domestically heterogeneous cluster.

85 **2. Model and experimental platform**

86 **2.1. The CAMx model description and configuration**

87 The Comprehensive Air Quality Model with Extensions version 6.10 (CAMx v6.10;
88 ENVIRON, 2014) is a state-of-the-art air quality model which simulates the emission, dispersion,
89 chemical reaction, and removal of the air pollutants on a system of nested three-dimensional grid
90 boxes (CAMx, 2023). The Eulerian continuity equation is expressed as shown Cao et al. (2023),
91 the first term on the right-hand side represents horizontal advection, the second term represents net
92 resolved vertical transport across an arbitrary space and time varying height grid, and the third
93 term represents turbulent diffusion on the sub-grid scale. Pollutant emission represents both point
94 source emissions and grided source emissions. Chemistry is treated by solving a set of reaction
95 equations defined by specific chemical mechanisms. Pollutant removal includes both dry
96 deposition and wet scavenging by precipitation.

97 In terms of the horizontal advection term on the right-hand side, this equation is solved using
98 either the Bott (1989) scheme or the Piecewise Parabolic Method (PPM) (Colella and Woodward,
99 1984; Odman and Ingram, 1996) scheme. The PPM horizontal advection scheme (HADVPPM)
100 was selected in this study because it provides higher accuracy with minimal numerical diffusion
101 (ENVIRON, 2014). The other numerical schemes selected during the CAMx model testing are
102 listed in Table S1. As described by Cao et al. (2023), the -fp-model precise compile flag which can
103 force the compiler to use the vectorization of some computation under value safety is 41.4% faster
104 than -mieee-fp compile flag which comes from the Makefile of the official CAMx version with
105 the absolute errors of the simulation results are less than ± 0.05 ppbV. Therefore, the -fp-model
106 precise compile flag was selected when compiling the CAMx model in this research.

107 **2.2. CUDA and ROCm introduction**

108 Compute Unified Device Architecture (CUDA; NVIDIA, 2020) is a parallel programming
109 paradigm which was released in 2007 by NVIDIA. CUDA is a proprietary application
110 programming interface (API) and as such is only supported on NVIDIA's GPUs. For the CUDA
111 programming, it uses a programming language similar to standard C, which achieves efficient

112 parallel computing of programs on NVIDIA GPUs by adding some keywords. In the previous
113 study, CUDA technology was implemented to port the HADVPPM program from CPU to
114 NVIDIA GPU (Cao et al., 2023).

115 Radeon Open Compute platform (ROCm; AMD, 2023) is an open-source software platform
116 developed by AMD for HPC and hyperscale GPU computing. In general, ROCm for the AMD
117 GPU is equivalent to CUDA for NVIDIA GPU. On the ROCm software platform, it uses the
118 AMD’s HIP interface which is a C++ runtime API allowing developers to run programs on AMD
119 GPUs. In general, they are very similar and their code can be converted directly by replacing the
120 string “cuda” with “hip” in the most cases. More information about HIP API is available on the
121 AMD ROCm website (ROCm, 2023). Similar to AMD GPU, developers can also use ROCM-HIP
122 programming interface to implement programs running on the China’ s domestically GPU-like
123 accelerator. The CUDA code cannot run directly on domestic GPU-like accelerators, and it needs
124 to be transcoded into HIP code.

125 **2.3. Hardware components and software environment of the testing system**

126 Table 1 lists four GPU clusters where we conducted the experiments, two NVIDIA
127 heterogeneous clusters which have the same hardware configuration as Cao et al. (2023) and two
128 China’ s domestically heterogeneous clusters newly used in this research, namely “Songshan”
129 supercomputer and “Taiyuan” computing platform. Two NVIDIA heterogeneous clusters are
130 equipped with NVIDIA Tesla K40m and V100 GPU accelerators, respectively. Both two domestic
131 clusters include thousands of computing nodes and each containing one China’ s domestically
132 CPU processor, four China’ s domestically GPU-like accelerators, and 128 GB of DDR4 2666
133 memory. The domestic CPU has four NUMA nodes, each NUMA node has eight X86 based
134 processors. The accelerator adopts a GPU-like architecture consisting of a 16 GB HBM2 device
135 memory and many compute units. The GPU-like accelerators connected to CPU with PCI-E, the
136 peak bandwidth of the data transfer between main memory and device memory is 16 GB/s.

137 It is worth noting that the “Taiyuan” computing platform, which has been updated in three
138 main aspects compared to the “Songshan” supercomputer. The CPU clock speed has been
139 increased from 2.0 GHz to 2.5 GHz, the number of GPU-like computing units has been increased

140 from 3,840 to 8,192, and the peak bandwidth between main memory and video memory has been
 141 increased from 16 GB/s to 32 GB/s. In terms of the software environment, the NVIDIA GPU is
 142 programmed using the CUDA toolkit, and the domestic GPU-like is programmed using the
 143 ROCm-HIP toolkit developed by AMD (ROCm, 2023). More details about the hardware
 144 composition and software environment of the four heterogeneous clusters are presented in Table 1.
 145 **Table 1.** Configurations of the NVIDIA K40m cluster, NVIDIA V100 cluster, “Songshan” supercomputer, and
 146 “Taiyuan” computing platform.

	Hardware components	
	CPU	GPU
NVIDIA K40m cluster	Intel Xeon E5-2682 v4 CPU @2.5 GHz, 16 cores	NVIDIA Tesla K40m GPU, 2880 CUDA cores, 12 GB video memory
NVIDIA V100 cluster	Intel Xeon Platinum 8168 CPU @2.7 GHz, 24 cores	NVIDIA Tesla V100 GPU, 5120 CUDA cores, 16 GB video memory
Songshan supercomputer	China’ s domestically processor A, 2.0GHz, 32 cores	China’ s domestically GPU-like accelerator A, 3840 computing units, 16 GB memory
Taiyuan computing platform	China’ s domestically processor B, 2.5GHz, 32 cores	China’ s domestically GPU-like accelerator B, 8192 computing units, 16 GB memory
	Software environment	
	Compiler and MPI	Programming model
NVIDIA K40m cluster	Intel Toolkit 2021.4.0	CUDA-10.2
NVIDIA V100 cluster	Intel Toolkit 2019.1.144	CUDA-10.0
Songshan supercomputer	Intel Toolkit 2021.3.0	ROCm-4.0.1/ DTK-23.04
Taiyuan computing platform	Intel Toolkit 2021.3.0	DTK-23.04

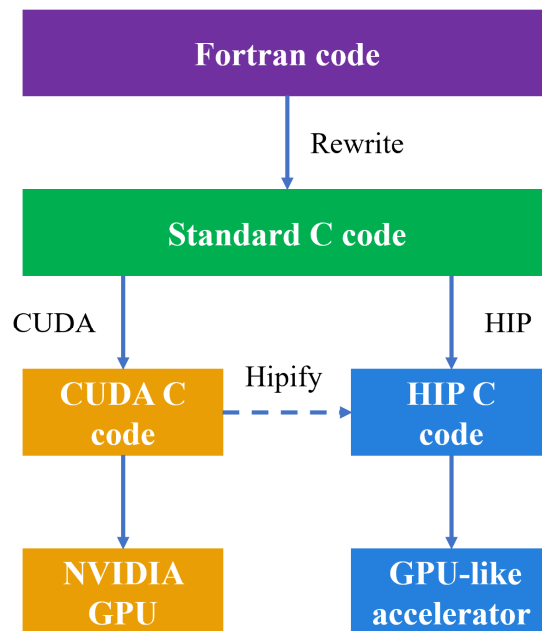
147 3. Implementation details

148 This section mainly introduced the strategy of porting HADVPPM program from CPU to
 149 NVIDIA GPU and domestic GPU-like accelerator, as well as the proposed multi-level hybrid
 150 parallelism technology to make full use of computing resources.

151 3.1. Porting the HADVPPM program from CPU to NVIDIA GPU and domestic 152 GPU-like accelerator

153 Fig. 1 shows the heterogeneous porting process of HADVPPM from CPU to NVIDIA GPU
 154 and domestic GPU-like accelerator. First, the original Fortran code was refactored using standard
 155 C language. Then the CUDA and ROCm HIP technology were used to convert the standard C

156 code into CUDA C and HIP C code to make it computable on the NVIDIA GPU and domestic
157 GPU-like accelerator. Similar to CUDA technology, the HIP technology is implemented to convert
158 the standard C code to HIP C code by adding related built-in functions (such as hipMalloc,
159 hipMemcpy, hipFree, etc.). To facilitate the portability of applications across different GPU
160 platforms, ROCm provides hipify toolkits to help transcode. The hipify toolkit is essentially a
161 simple script written in the Perl language, and its function is text replacement, which replaces the
162 function name in CUDA C code with the corresponding name in HIP C code according to certain
163 rules. For example, for the memory allocation function cudaMalloc in CUDA, the hipify toolkit
164 can automatically recognize and replace it with hipMalloc. Therefore, the thread and block
165 configuration of GPU remain unchanged due to the simple text substitution during the transcoding.
166 In this study, the ROCm HIP technology was used to implement the operation of GPU-
167 HADVPPM on domestic GPU-like accelerator based on the CUDA version of GPU-HADVPPM
168 which was developed by Cao et al. (2023). The HIP code was compiled using the “hipcc”
169 compiler driver with the library flag “-lamdhip64”.



170

171 **Figure 1.** The heterogeneous porting process of HADVPPM Fortran code from CPU to NVIDIA GPU and

172 domestic GPU-like accelerator.

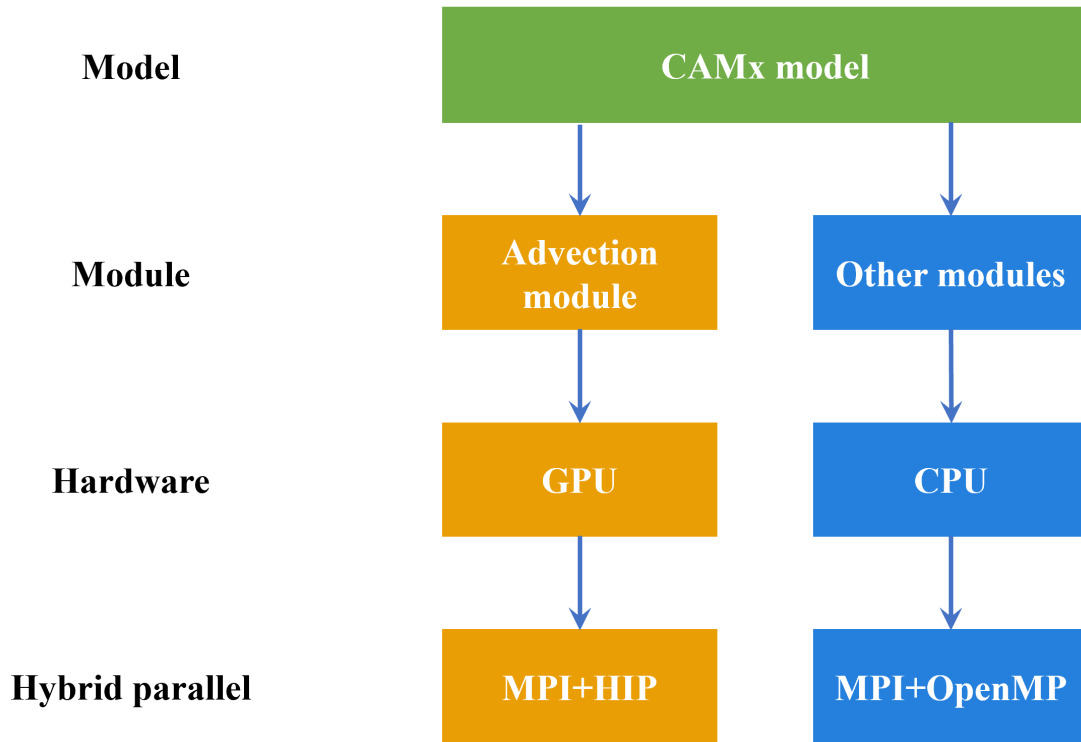
173 **3.2. Multi-level hybrid parallelization of CAMx model on heterogeneous** 174 **platform**

175 The original CAMx model running on the CPUs supports two types of parallelization
176 (ENVIRON, 2014): (1) OpenMP (OMP), which supports multi-platform (e.g., multi-core) shared-
177 memory programming in C/C++ and Fortran; (2) Message Passing Interface (MPI), which is a
178 message passing interface standard for developing and running parallel applications on the
179 distributed-memory computer cluster. During the process of CAMx model simulation, MPI and
180 OMP hybrid parallelism can be used, several CPU processes can be launched, and each process
181 can spawn several threads. This hybrid parallelism can significantly improve the computational
182 efficiency of CAMx model.

183 As mentioned above, the original CAMx model supports message passing interface (MPI)
184 parallel technology running on the general-purpose CPU. The simulation domain is divided into
185 several sub-regions by MPI, and each CPU process is responsible for computation of its sub-
186 region, which includes the computation tasks of advection module and other modules such as
187 photolysis module, deposition module, chemical module, etc. In the previous studying, Cao et al.
188 (2023) adopt a parallel architecture with an MPI and CUDA (MPI+CUDA) hybrid paradigm to
189 configure one GPU accelerator for each CPU process. For the advection module, the simulation
190 originally implemented by the CPU is handed over to the GPU. Other module computing tasks
191 continue to be completed on the CPU.

192 In this study, when the CUDA C code of GPU-HADVPPM is converted to HIP C code, GPU-
193 HADVPPM with an MPI and HIP (MPI+HIP) heterogeneous hybrid programming technology can
194 also run on multiple domestic GPU-like accelerators. However, the number of GPU-like
195 accelerators in a single compute node is usually much smaller than the number of CPU cores in
196 the heterogeneous HPC systems. Therefore, in order to make full use of the remaining CPU
197 computing resources, the OMP API of CAMx model is further introduced to realize the
198 MPI+OMP hybrid parallelism of other modules on CPU. A schematic of the multi-level hybrid
199 parallel framework is shown in Figure 2. For example, in a computing node, four CPU processes
200 and four GPU-like accelerators are launched, and each CPU process spawns four threads. Then the
201 advection module is simulated by 4 GPU-like accelerators, and the other modules are done by 4*4

202 threads spawned by CPU processes.



203

204 **Figure 2.** A schematic of the multi-level hybrid parallel framework.

205 **4. Results and evaluation**

206 The computational performance experiments of CUDA and HIP version GPU-HADVPPM
207 are reported in this section. First, we compared the simulation result of the Fortran version CAMx
208 model with CAMx-CUDA and CAMx-HIP model which were coupled with CUDA and HIP
209 version of GPU-HADVPPM program, respectively. Then, the computational performance of
210 GPU-HADVPPM programs on the NVIDIA GPU and domestic GPU-like accelerator are
211 compared. Finally, we tested total performance of CAMx-HIP model with multi-level hybrid
212 parallelization on the the "Songsshan" supercomputer. For ease of description, the CAMx versions
213 of the HADVPPM program written in Fortran, CUDA C and HIP C code are named Fortran,
214 CUDA and HIP, respectively.

215 **4.1. Experimental setup**

216 There are three test cases were used to evaluate the performance of CUDA and HIP version

217 GPU-HADVPPM. The experimental setup for the three test cases is shown in Table 2. In the
 218 previous study of Cao et al. (2023), the BJ case was used to carry out the performance tests, HN
 219 case and ZY case are the newly constructed test cases in this study. The Beijing case (BJ) covers
 220 Beijing, Tianjin, and part of the Hebei Province with 145×157 grid boxes, and simulation of BJ
 221 case starts on 1 November, 2020. The Henan case (HN) mainly covers the Henan Province with
 222 209×209 grid boxes. The starting date of simulation in HN case is 1 October, 2022. The
 223 Zhongyuan case (ZY) has the widest coverage of the three cases, with Henan Province as the
 224 center, covering the Beijing-Tianjin-Hebei region, Shanxi Province, Shaanxi Province, Hubei
 225 Province, Anhui Province, Jiangsu Province, and Shandong Province, with 531×513 grid boxes.
 226 ZY case started simulation on 4 January, 2023. All of the three performance test cases are 3km
 227 horizontal resolution, 48 hours of simulation, and 14 vertical model layers. The number of three-
 228 dimensional grid boxes in BJ, HN, and ZY cases are totally 318,710, 611,534 and 3,813,642,
 229 respectively. The meteorological fields inputting the different versions of the CAMx model in the
 230 three cases were provided by the Weather Research and Forecasting Model (WRF). In terms of
 231 emission inventories, the emission for BJ case is consistent with the Cao et al. (2023), HN case
 232 uses the Multi-resolution Emission Inventory for China (MEIC) and ZY case uses the emission
 233 constructed by Sparse Matrix Operator Kernel Emission (SMOKE) model in this study.

234 **Table 2.** The experimental setup for the BJ, HN, and ZY case.

	BJ	HN	ZY
Start date	November 1, 2020	October 1, 2022	1 January, 2023
Horizontal resolution	3km	3km	3km
Grid boxes	$145 \times 157 \times 14$	$209 \times 209 \times 14$	$531 \times 513 \times 14$
Meteorological fields	WRF	WRF	WRF
Emission	Cao et al. (2023)	MEIC	SMOKE

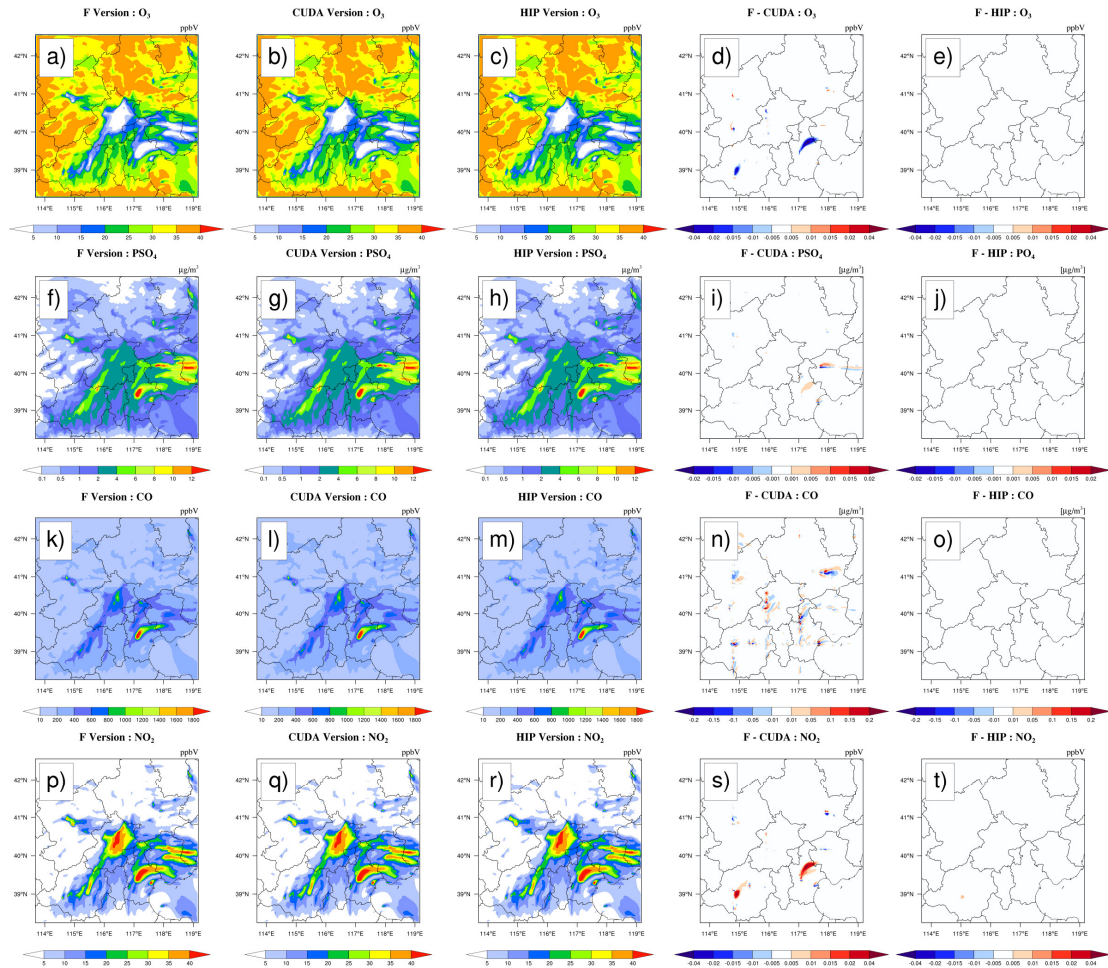
235 **4.2. Error analysis**

236 The hourly concentrations of four major species, i.e. O₃, PSO₄, CO, and NO₂, outputted by
 237 the Fortran, CUDA, and HIP versions of CAMx for the BJ case are compared to verify the results
 238 correctness before testing the computational performance. Fig. 3 shows the four major species
 239 simulation results of the three CAMx version, including Fortran version on the Intel E5-2682 v4
 240 CPU, CUDA version on the NVIDIA K40m cluster and HIP version on the “Songshan”

241 supercomputer, after 48 hours integration, as well as the absolute errors (AEs) of their
242 concentrations. As described by Cao et al. (2023), the parallel design of the CAMx model adopts
243 the primary/secondary mode, and P0 process is responsible for inputting and outputting the data
244 and calling the MPI_Barrier function to synchronize the process, and the other processes are
245 responsible for simulation. When comparing the simulation results, we only launched 2 CPU
246 processes on the CPU platform, and launched 2 CPU processes and configure 2 GPU accelerators
247 on the NVIDIA K40m cluster and “Songshan” supercomputer, respectively.

248 The species’ spatial pattern of three CAMx versions on different platform are visually very
249 consistent, and the AEs between the HIP and Fortran version is much smaller than the CUDA and
250 Fortran version. For example, the AEs between the CUDA and Fortran version for O₃, PSO₄, and
251 NO₂ are in the range of ± 0.04 ppbV, $\pm 0.02 \mu\text{g} \cdot \text{m}^{-3}$, and ± 0.04 ppbV. And the AEs between the
252 HIP and Fortran version for above the three species are fall into the range of ± 0.01 ppbV, ± 0.005
253 $\mu\text{g} \cdot \text{m}^{-3}$, and ± 0.01 ppbV. For CO, AEs is relatively large due to its high background
254 concentration. However, the AEs between the HIP and Fortran versions is also less than that
255 between the CUDA and Fortran versions where were in the range of ± 0.4 ppbV and ± 0.1 ppbV,
256 respectively.

257 Considering the situation of AEs accumulate and grow, Fig. 4 highlights the time series of
258 AEs between Fortran and CUDA versions and between Fortran and HIP versions after grid
259 averaging. As is shown in Fig. 4, the AEs of O₃, PSO₄, CO, and NO₂ between the Fortran version
260 and the CUDA version are -0.0002 to 0.0001 ppbV, -0.00003 to 0.00001 $\mu\text{g} \cdot \text{m}^{-3}$, -0.0004 to
261 0.0004 ppbV, and -0.0002 to 0.0002 ppbV, respectively, and fluctuate. Although the AEs of the
262 above four species between the Fortran and the HIP version also fluctuates, the fluctuation range
263 is much smaller than that of the CUDA version. Importantly, the AEs between Fortran and CUDA
264 versions and between Fortran and HIP versions both do not accumulate and grow over prolonged
265 simulation periods.



266

267

Figure 3. O₃, PSO₄, CO, and NO₂ concentrations outputted by the CAMx Fortran version on the Intel E5-2682 v4

268

CPU, CUDA version on the NVIDIA K40m cluster and HIP version on the "Songshan" supercomputer under the

269

BJ case. Panels (a), (f), (k), and (p) are from the Fortran version of simulation results for four species. Panels (b),

270

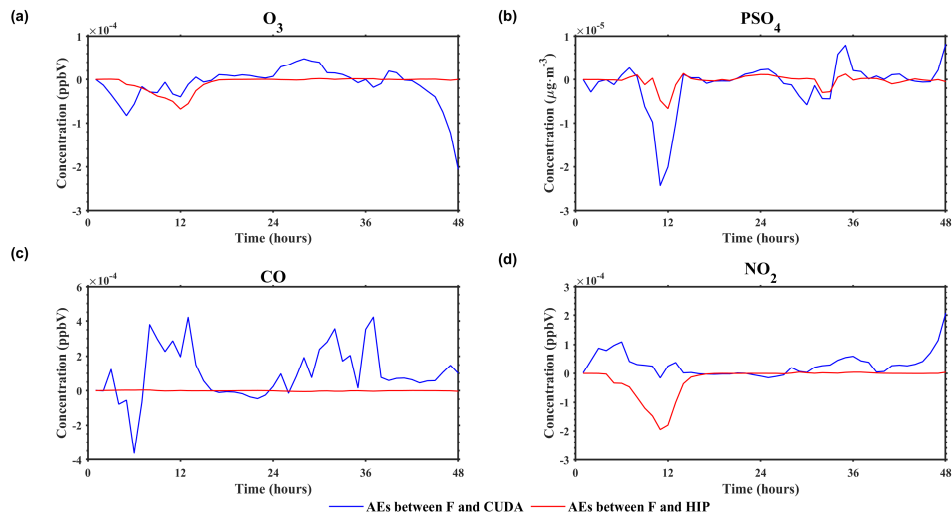
(g), (l), and (q) are from the CUDA version of simulation results for four species. Panels (c), (h), (m), and (r) are

271

from the HIP version of simulation results for four species. Panels (d), (i), (n), and (s) are the AEs between the

272


Fortran and CUDA versions. Panels (e), (j), (o), and (t) are the AEs between the Fortran and HIP versions.



273

274 **Figure 4.** The time series of AEs between Fortran and CUDA versions (solid blue line) and between Fortran and
 275 HIP versions (solid red line) after grid averaging. Panel (a)–(d) represent the AEs of O₃, PSO₄, CO, and NO₂,
 276 respectively.

277 Fig. 5 presents the boxplot of the relative errors (REs) in all grid boxes for the PSO₄, PNO₃,
 278 PNH₄, O₃, CO, and NO₂ during the 48 hours simulation under the BJ case. Statistically, the REs
 279 between the CUDA version on the NVIDIA K40m cluster and Fortran version on the Intel E5-
 280 2682 v4 CPU for the above six species are in the range of $\pm 0.006\%$, $\pm 0.01\%$, $\pm 0.008\%$,
 281 $\pm 0.002\%$, $\pm 0.002\%$, and $\pm 0.002\%$. In terms of REs between the HIP version on the “Songshan”
 282 supercomputer and Fortran version on the Intel E5-2682 v4 CPU, the values are much smaller
 283 than REs between CUDA and Fortran versions which are fall into the range of $\pm 0.0005\%$,
 284 $\pm 0.004\%$, $\pm 0.004\%$, $\pm 0.00006\%$, $\pm 0.00004\%$, and $\pm 0.00008\%$, respectively. In the air quality
 285 model, the secondary particulate matter, such as PNH₄, PNO₃, and PSO₄, have a common
 286 characteristic: their initial concentration is very low and they are mainly generated through
 287 complex chemical reactions. Therefore, when calculating the relative error on different hardware
 288 platforms, because the value in the denominator is very small, it is very sensitive to a small
 289 difference in the numerator, resulting in a large relative error. But from the absolute error in Fig.3,
 290 the absolute error of PSO₄ on different hardware platforms is smaller than that of other species.
 291 For gaseous pollutants such as CO, O₃, and NO₂, the initial concentration is large due to emission,
 292 and the denominator value is large when calculating the relative error, which is insensitive to small
 293 differences in the numerator.



294
 295 **Figure 5.** The distribution of REs in all grid boxes for the PSO_4 , PNO_3 , PNH_4 , O_3 , CO , and NO_2 under the BJ case.
 296 The red boxplot represents the REs between the CUDA version on the NVIDIA K40m cluster and Fortran version
 297 on the Intel E5-2682 v4 CPU, and blue boxplot represents the REs between the HIP version on the "Songshan"
 298 supercomputer and Fortran version on the Intel E5-2682 v4 CPU.

299 Wang et al. (2021) verified the applicability of the numerical model in scientific research by
 300 computing the ratio of root mean square error (RMSE) between two different model versions to
 301 system spatial variation (standard deviation, std). If the ratio is smaller, it is indicated that the
 302 difference in the simulation results of the model on the GPU is minimal compared with the spatial
 303 variation of the system, that is to say, the simulation results of the model on the GPU are accepted
 304 for scientific research. Here, we calculate the standard deviation of O_3 , PSO_4 , CO and NO_2 on the
 305 Intel Xeon E5-2682 v4 CPU, and their RMSE between the NVIDIA V100 cluster, NVIDIA K40m
 306 cluster and "Songshan" supercomputer and the Intel Xeon E5-2682 v4 CPU, which are presented
 307 in Table 3. The std for the above four species on the Intel Xeon E5-2682 v4 CPU are 9.6 ppbV, 1.7
 308 $\mu\text{g} \cdot \text{m}^{-3}$, 141.9 ppbV, and 7.4 ppbV, respectively, and their ratios of RMSE and std on the
 309 "Songshan" supercomputer are $5.8 \times 10^{-5} \%$, $4.8 \times 10^{-6} \%$, $5.7 \times 10^{-8} \%$, and $2.1 \times 10^{-4} \%$,
 310 which are smaller than two NVIDIA clusters, especially much smaller than the NVIDIA V100
 311 cluster. For example, the ratio on the NVIDIA K40m cluster for four species are $1.2 \times 10^{-4} \%$,
 312 $6.6 \times 10^{-5} \%$, $7.0 \times 10^{-5} \%$, and $4.1 \times 10^{-4} \%$, and ratio on the NVIDIA V100 cluster are
 313 $1.5 \times 10^{-2} \%$, $2.5 \times 10^{-3} \%$, $6.4 \times 10^{-3} \%$, and $1.3 \times 10^{-3} \%$, respectively.

314 **Table 3.** The standard deviation (std) of O_3 , PSO_4 , CO and NO_2 on the Intel Xeon E5-2682 v4 CPU, root mean

315 square error (RMSE) and its ratio on the NVIDIA V100 cluster, NVIDIA K40m cluster and "Songshan"
 316 supercomputer

	std	NVIDIA V100 cluster		NVIDIA K40m cluster		"Songshan" supercomputer	
		RMSE	RMSE/std	RMSE	RMSE/std	RMSE	RMSE/std
O₃ (ppbV)	9.6	1.5×10^{-3}	1.5×10^{-2}	1.1×10^{-5}	1.2×10^{-4}	7.4×10^{-6}	7.7×10^{-5}
PSO₄ ($\mu\text{g} \cdot \text{m}^{-3}$)	1.7	4.3×10^{-5}	2.5×10^{-3}	1.1×10^{-6}	6.6×10^{-5}	2.5×10^{-7}	1.5×10^{-5}
CO (ppbV)	141.9	9.0×10^{-3}	6.4×10^{-3}	1.0×10^{-4}	7.0×10^{-5}	4.4×10^{-7}	3.1×10^{-7}
NO₂ (ppbV)	7.4	9.3×10^{-5}	1.3×10^{-3}	3.0×10^{-5}	4.1×10^{-4}	2.0×10^{-5}	2.7×10^{-4}

317 From AEs, REs, and ratio of RMSE and std between different CAMx versions, it is less
 318 difference that the GPU-HADVPPM4HIP program runs on the "Songshan" supercomputer.
 319 Because the simulation accuracy of geoscience numerical model is closely related to the model
 320 efficiency, and many model optimization works improve the computational performance by
 321 reducing the precision of the data, such as Váña et al. (2017) changed some variables precision in
 322 the atmospheric model from double precision to single precision, which increased the overall
 323 computational efficiency by 40%, and Wang et al. (2019) improved the computational efficiency
 324 of the gas-phase chemistry module in the air quality mode by 25%~28% by modifying the
 325 floating-point precision compile flag. Therefore, we speculate that this may be related to the
 326 manufacturing process of NVIDIA GPUs and domestic GPU-like accelerators, especially NVIDIA
 327 Tesla V100 series GPUs, which may use unknown optimizations to improve GPU performance
 328 efficiency by losing part of the accuracy. In this study, we mainly focus on numerical simulation.
 329 Of course, we also want to know the specific reasons for this, but we are not professional GPU
 330 research and development designers after all and do not know the underlying design logic of the
 331 hardware, so we can only present our experimental results in the air pollution model to you, and
 332 discuss with each other to jointly promote the application of GPU in the field of geoscience
 333 numerical models.

334 4.3. Application performance


335 4.3.1. GPU-HADVPPM on a single GPU accelerator

336 As described in Sect. 4.2, we validate the 48 hours simulation results outputted by the Fortran,
 337 CUDA, and HIP versions of CAMx. Next, computational performance was compared for the

338 Fortran version of HADVPPM on the Intel Xeon E5-2682 v4 CPU and domestic CPU processor A,
339 the CUDA version of GPU-HADVPPM on the NVIDIA Tesla K40m and V100 GPU, and the HIP
340 version of GPU-HADVPPM on the domestic GPU-like accelerator A, under the BJ, HN and ZY
341 case. The simulation time in this section is 1 hour unless otherwise specified.


342 Similarity, since the CAMx model adopts the primary/secondary mode, two CPU processes
343 P0 and P1 are launched on the CPU, and the `system_clock` functions in the Fortran language are
344 used to test the elapsed time of the advection module in P1 process. When testing the computation
345 performance of the advection module on the GPU-like accelerator, we also only launch 2 CPU
346 processes and 2 GPU-like accelerators. When a P1 process runs to the advection module, the
347 original computation process is migrated from the CPU to the GPU, and the `hipEvent_t` function
348 in HIP programming is used to test the running time of the advection module on the GPU-like
349 accelerator. When comparing the speedup on different GPU accelerators, the elapsed time of
350 advection module launched one CPU process (P1) on the domestic CPU processor A is taken as
351 the benchmark, that is, the speedup is 1.0x. The runtime of the advection module on Intel CPU
352 processor and different GPU accelerators is compared with the baseline to obtain the speedup.

353 Fig. 6(a) and (b) shows the elapsed time and speedup of the different versions of HADVPPM
354 on the CPU processors and GPU accelerators for BJ, HN, and ZY cases, respectively. The results
355 show that using CUDA and HIP technology to port HADVPPM from CPU to GPU can
356 significantly improve its computational efficiency. For example, the elapsed time of the advection
357 module on the domestic processor A is 609.2 seconds under the ZY case. After it is ported to the
358 domestic GPU accelerator and NVIDIA V100 GPU, it only takes 21.1 seconds and 7.6 seconds to
359 complete the computing, and the speedups are 28.9x and 80.2x, respectively. The ZY case had the
360 largest number of grids in the three cases and exceeded the memory of a single NVIDIA Tesla
361 K40m GPU accelerator, so it was not possible to test its elapsed time on it. Moreover, the
362 optimization of thread and block co-indexing is used to simultaneously compute the grid point in
363 the horizontal direction (Cao et al., 2023). Therefore, it can be seen from Fig. 6(b) that the larger
364 the computing scale, the more obvious the acceleration, which indicates that GPU is more suitable
365 for super-large scale parallel computing, and provides technical support for accurate and fast
366 simulation of ultra-high-resolution air quality at the meter level in the future.



367
 368 **Figure 6.** The elapsed time (a) and speedup (b) of the Fortran version of HADVPPM on the Intel Xeon E5-2682
 369 v4 CPU and the domestic CPU processor A, the CUDA version of GPU-HADVPPM on the NVIDIA Tesla K40m
 370 GPU, NVIDIA Tesla V100 GPU, and the HIP version of GPU-HADVPPM on the domestic GPU-like accelerator
 371 A for BJ, HN, and ZY case. The unit of elapsed time is in seconds (s).

372 The timestep of BJ, HN and ZY case were 59, 47, and 61, respectively. Fig. 7 shows the
 373 GPU-HADVPPM4HIP acceleration in each time step on a single domestic GPU-like accelerator A.
 374 It can be seen from the figure that all three cases have the smallest speedup of 8.2x, 11.2x, and
 375 27.8x at the first timestep, which is related to the time required for GPU-like accelerator startup.
 376 When the GPU-like is started and operating normally, the speedup of the three cases tend to be
 377 stable in the following time steps, and stabilize around 8.5x, 11.5x and 28.0x respectively.



378

379 **Figure 7.** The GPU-HADVPPM4HIP acceleration in each time step on a single GPU-like accelerator for BJ, HN,
 380 and ZY case. The timestep of above three cases are 59, 47, and 61, respectively.

381 Table 4 further lists the total elapsed time of CAMx Fortran and HIP versions for BJ case on
 382 the "Songshan" supercomputer and "Taiyuan" computing platform, and the computing time of
 383 advection module with or without data transfer. By coupling the GPU-HADVPPM4HIP to CAMx
 384 model and adopting a series of optimizations (Cao et al., 2023) such as communication
 385 optimization, memory access optimization, and 2D thread optimization, the overall computation
 386 time of CAMx-HIP model on a single domestic GPU-like accelerator is faster than that of the
 387 original Fortran version on a single domestic CPU core. For example, on the "Songshan"
 388 supercomputer, one hour of simulation in CAMx-HIP model takes 469 seconds, and the Fortran
 389 version takes 481 seconds. On the "Taiyuan" computing platform, the acceleration effect is more
 390 obvious due to the upgrade of hardware and network bandwidth, and the integration time of
 391 CAMx-HIP model is 433 seconds when maintaining the same software environment, and the
 392 integration time of the Fortran version is 453 seconds.

393 The elapsed time of GPU-HADVPPM given in Table 4 on NVIDIA GPU and domestic GPU-
 394 like accelerator does not consider the data transfer time between CPU and GPU. However, the
 395 communication bandwidth of data transfer between the CPU and GPU is one of the most
 396 significant factors that restrict the performance of numerical model on the heterogeneous cluster
 397 (Mielikainen et al., 2012; Mielikainen et al., 2013; Huang et al., 2013). To illustrate the significant
 398 impact of CPU-GPU data transfer efficiency, the computational performance of GPU-HADVPPM
 399 with and without data transfer time for the BJ case is tested on the "Songshan" supercomputer and

400 “Taiyuan” computing platform with the same DTK version 23.04 software environment and the
401 results are further presented in Table 6. For convenience of description, we refer to the execution
402 time of GPU-HADVPPM program on GPU kernel as kernel execution time, and the time of GPU-
403 HADVPPM running on GPU as total runtime, which contains two parts, namely, kernel execution
404 time and data transfer time between CPU and GPU. After testing, the kernel execution time and
405 total running time of GPU-HADVPPM4HIP program on domestic GPU-like accelerator A are 6.8
406 and 29.8 seconds, respectively. In other words, it only takes 6.8 seconds to complete the
407 computation on the domestic accelerator, but it takes 23.0 seconds to complete the data transfer
408 between the CPU and the domestic GPU-like accelerator, which is 3.4 times the computation time.
409 The same problem exists in the more advanced the "Taiyuan" computing platform, where the
410 GPU-HADVPPM4HIP takes only 5.7 seconds to complete the computation, while the data
411 transmission takes 18.2 seconds, which is 3.2 times the computation time.

412 By comparing the kernel execution time and total running time of GPU-HADVPPM4HIP on
413 the domestic accelerator, it can be seen that the data transfer efficiency between CPU and GPU is
414 really inefficient, which seriously restricts the computational performance of numerical models in
415 heterogeneous clusters. On the one hand, improving the data transfer bandwidth between CPU and
416 GPU can improve the computational efficiency of the model in heterogeneous clusters. On the
417 other hand, the optimization measures can be implemented to improve the data transfer efficiency
418 between CPU and GPU. For example, (1) Asynchronous data transfer is used to reduce the
419 communication latency between CPU and GPU. Computation and data transfer are performed
420 simultaneously to hide communication overhead; (2) Currently, some advanced GPU architectures
421 support a unified memory architecture, so that the CPU and GPU can share the same memory
422 space and avoid frequent data transfers. This reduces the overhead of data transfer and improves
423 data transfer efficiency; (3) Cao et al. (2023) adopted communication optimization measures to
424 reduce the communication frequency in one time integration step to one, but there is still the
425 problem of high communication frequency in the whole simulation. In the future, we will consider
426 porting other hotspots of CAMx model, or even the whole integral module except I/O, to GPU-
427 like accelerators for increasing the proportion of code on the GPU and reduce the frequency of
428 CPU-GPU communication.

429 Video memory and bandwidth are the two most significant factors affecting GPU
 430 performance, and high video memory and high bandwidth can better play the powerful computing
 431 performance of GPUs. Usually, the memory and bandwidth of the GPU are already given at the
 432 factory. In this case, the amount of data transferred to the GPU can be roughly estimated before
 433 the data is transferred to the GPU, and the amount of data transferred to the GPU can be adjusted
 434 according to the size of the GPU memory to ensure that the amount of data transferred to the GPU
 435 each time reaches the maximum GPU video memory, so as to give full play to the GPU
 436 performance more efficiently.

437 **Table 4.** The total elapsed time of CAMx Fortran and HIP versions for BJ case on the "Songshan" supercomputer
 438 and "Taiyuan" computing platform, and the computing time of advection module with or without data transfer. The
 439 unit of elapsed time is in seconds (s).


	"Songshan" supercomputer		"Taiyuan" computing platform	
	Fortran version	HIP version	Fortran version	HIP version
Total elapsed time	481.0	469.0	453.0	433.0
Computing time of advection module without data transfer	57.8	6.8	47.8	5.7
Computing time of advection module with data transfer	57.8	29.8	47.8	23.9

440 **4.3.2. CAMx-HIP model on the heterogeneous cluster**

441 Generally, the heterogeneous HPC systems have thousands of compute nodes which are
 442 equipped with one or more GPUs on each compute node. To make full use of multiple GPUs, a
 443 parallel architecture with an MPI and CUDA hybrid paradigm was implemented to improve the
 444 overall computational performance of CAMx-CUDA model (Cao et al., 2023). In this studying,
 445 the hybrid parallelism with an MPI and HIP paradigm was used to implement the HIP version of
 446 GPU-HADVPPM run on multiple domestic GPU-like accelerators. Fig.8 shows the total elapsed
 447 time and speedup of CAMx-HIP model which coupled with the HIP version GPU-HADVPPM on
 448 the "Songshan" supercomputer under the BJ, HN, and ZY cases. The simulation of above three
 449 cases for one hour took 488 seconds, 1135 seconds and 5691 seconds respectively when launching
 450 two domestic CPU processors and two GPU-like accelerators. For the BJ and HN case, the
 451 parallel scalability is highest when configured with 24 CPU cores and 24 GPU-like accelerators,
 452 with speedup of 8.1x and 11.6x, respectively. In terms of the ZY case, due to its large number of

453 grids, the parallel scalability is the highest when 32 CPU cores and 32 GPU-like accelerators are
454 configured, and the acceleration ratio is 17.2x.

455 As mentioned above, data transfer between CPU and GPU takes several times more time than
456 computation. Regardless of the CPU-GPU data transfer consumption, GPU-HADVPPM4HIP can
457 achieve up to 28.9x speedup on a single domestic GPU-like accelerator. However, in terms of the
458 total time consumption, the CAMx-HIP model is only 10~20 seconds faster than the original
459 Fortran version when one GPU-like accelerator is configured. And as the number of CPU cores
460 and GPU-like accelerators increases, the overall computing performance of CAMx-HIP model is
461 lower than that of the original Fortran version. The main reason is related to the amount of data
462 transferred to GPU. As the number of MPI processes increases, the number of grids responsible
463 for each process decreases, and the amount of data transmitted by the advection module from CPU
464 to GPU decreases. However, GPUs are suitable for large-scale matrix computing. When the data
465 scale is small, the performance of GPU is low, and the communication efficiency between CPU
466 and GPU is the biggest bottleneck (Cao et al., 2023). Therefore, the computational performance of
467 CAMx-HIP model is not as good as the original Fortran version when MPI processes increase.
468 According to the characteristics of GPUs suitable for large-scale matrix computing, the model
469 domain can be expanded and the model resolution can be increased in the future to ensure that the
470 amount of data transferred to each GPU reaches the maximum video memory occupation, so as to
471 make efficient use of GPU. In addition, the advection module only accounts for about 10% of the
472 total time consumption in CAMx model (Cao et al., 2023), and in the future, it is considered to
473 port the entire integration module except I/O to the GPU to minimize the communication
474 frequency.




475

476 **Figure 8.** The total elapsed time and speedup of CAMx-HIP model on the "Songshan" supercomputer under the BJ,
 477 HN, and ZY cases. The unit is in seconds (s).

478 The number of GPU accelerators in a single compute node is usually much smaller than the
 479 number of CPU cores in the heterogeneous HPC systems. Using the hybrid parallel paradigm with
 480 MPI and HIP to configure one GPU accelerator for each CPU process results in idle computing
 481 resources for the remaining CPU cores. Therefore, the multi-level hybrid parallelism scheme was
 482 introduced to further improve the total computational performance of the CAMx-HIP model. As
 483 described in the Sect. 3.2, the horizontal advection module is accelerated by MPI and HIP
 484 technology, and the other modules, such as photolysis module, deposition module, chemical
 485 module, etc., which runs on the CPU are accelerated by MPI and OMP under the framework of the
 486 multi-level hybrid parallelism.

487 The ZY case achieved the maximum speed-up when launching the 32 domestic CPU
 488 processors and GPU-like accelerators. In the same configuration, Fig. 9 shows the total elapsed
 489 time and speedup of CAMx-HIP model when further implementing the multi-level hybrid
 490 parallelism on the "Songshan" supercomputer. The AEs of the simulation results between the
 491 CAMx-HIP model and CAMx-HIP model with the OMP technology is within ± 0.04 ppbV, and the
 492 specified results are shown in Figure S1. As the number of threads increases, the elapsed time of

493 CAMx-HIP model is further reduced. When a CPU core launching 8 threads, the one-hour
 494 integration time in CAMx-HIP model has been reduced from 338 seconds to 178 seconds, with a
 495 maximum acceleration of 1.9x.



496
 497 **Figure 9.** The total elapsed time and speedup of CAMx-HIP model when implementing the multi-level hybrid
 498 parallelism in the ZY case. The unit is in seconds (s).

499 5. Conclusions and discussion

500 GPUs have become an essential part of providing processing power for high performance
 501 computing applications, especially in the field of geoscience numerical models, implementing
 502 super-large scale parallel computing of numerical models on GPUs has become one of the
 503 significant directions of its future development. In this study, the ROCm HIP technology was
 504 implemented to port the GPU-HADVPPM from the NVIDIA GPUs to China's domestically
 505 GPU-like accelerators, and further introduced the multi-level hybrid parallelism scheme to
 506 improve the total computational performance of the CAMx-HIP model on the China's
 507 domestically heterogeneous cluster.

508 The consistency of model simulation results is a significant prerequisite for heterogeneous
 509 porting, although the experimental results show that the deviation between the CUDA version and
 510 the Fortran version of CAMx model, and the deviation between the HIP version and the Fortran

511 version of CAMx model, are within the acceptable rang, the simulation difference between the
512 HIP version of CAMx model and Fortran version of CAMx model is smaller. Moreover, the BJ,
513 HN, and ZY test cases can achieve 8.5x, 11.5x, and 28.9x speedup, respectively, when the
514 HADVPPM program is ported from the domestic CPU processor A to the domestic GPU-like
515 accelerator A. The experimental results of different cases show that the larger the computing scale,
516 the more obvious the acceleration effect of the GPU-HADVPPM program, indicating that GPU is
517 more suitable for super-large scale parallel computing, and provides technical support for accurate
518 and fast simulation of ultra-high-resolution air quality at the meter level in the future. The data
519 transfer bandwidth between CPU and GPU is one of the most important factors affecting the
520 computational efficiency of numerical model in heterogeneous clusters, as shown by the fact that
521 the elapsed time of GPU-HADVPPM program on GPU only accounts for 7.3% and 23.8% when
522 considering the data transfer time between CPU and GPU on the the “Songshan” supercomputer
523 and “Taiyuan” computing platform. Therefore, optimizing the data transfer efficiency between
524 CPU and GPU is one of the important directions for the porting and adaptation of geoscience
525 numerical models on heterogeneous clusters in the future.

526 There is still potential to further improve the computational efficiency of the CAMx-HIP
527 model in the further. First, improve the data transfer efficiency of GPU-HADVPPM between the
528 CPU and the GPU and reduce the data transfer time. Secondly, increase the proportion of HIP C
529 code in CAMx-HIP model on the domestic GPU-like accelerator, and port other modules of
530 CAMx-HIP model to the domestic GPU-like accelerator for computing. Finally, the data type of
531 some variables could be changed from double precision to single precision, and the mixing-
532 precision method is used to further improve the CAMx-HIP computing performance.

533

534

535 *Code and data availability.* The source codes of CAMx version 6.10 are available at [https://camx-](https://camx-wp.azurewebsites.net/download/source/)
536 [wp.azurewebsites.net/download/source/](https://camx-wp.azurewebsites.net/download/source/) (ENVIRON, 2023). The datasets related to this paper and
537 the CAMx-HIP codes are available online via ZENODO
538 (<https://doi.org/10.5281/zenodo.10158214>), and the CAMx-CUDA code is available online via
539 ZENODO (<https://doi.org/10.5281/zenodo.7765218>, Cao et al., 2023).

540

541 *Author contributions.* KC and QW conducted the simulation and prepared the materials. QW, LiW
542 and LaW planned and organized the project. KC, QW, HG, HW, XT and LL refactored and
543 optimized the codes. LiW, NW, HC, DXL and DQL collected and prepared the data for the
544 simulation. KC, HW, QW, and HG validated and discussed the model results. KC, QW, LiW, NW,
545 XT, HG, and LaW took part in the discussion.

546

547 *Competing interests.* The authors declare that they have no conflict of interest.

548

549 *Acknowledgements.* The National Key R&D Program of China (grant no. 2020YFA0607804), the
550 National Supercomputing Center in Zhengzhou Innovation Ecosystem Construction Technology
551 Special Program (grant no. 201400210700), GHfund A (grant no. 202302017828), and the Beijing
552 Advanced Innovation Program for Land Surface funded this work. The authors would like to
553 thank the High Performance Scientific Computing Center (HSCC) of Beijing Normal University
554 for providing some high-performance computing environment and technical support.

555

556 **Reference**

557 Alvanos, M. and Christoudias, T.: GPU-accelerated atmospheric chemical kinetics in the
558 ECHAM/MESSy (EMAC) Earth system model (version 2.52), Geoscientific Model
559 Development, 10, 3679-3693, 10.5194/gmd-10-3679-2017, 2017.

560 AMD: ROCm Documentation Release 5.7.1,
561 https://rocm.docs.amd.com/_downloads/en/latest/pdf/ (last access: 20 October 2023), 2023.

562 Bott, A.: A Positive Definite Advection Scheme Obtained by Nonlinear Renormalization of the
563 Advective Fluxes, Monthly Weather Review - MON WEATHER REV, 117, 10.1175/1520-
564 0493(1989)117<1006:APDASO>2.0.CO;2, 1989.

565 CAMx, A multi-scale photochemical modeling system for gas and particulate air pollution,
566 available at: <https://www.camx.com/> (last access: 20 October 2023), 2023.

567 Cao, K., Wu, Q., Wang, L., Wang, N., Cheng, H., Tang, X., Li, D., and Wang, L.: GPU-

568 HADVPPM V1.0: a high-efficiency parallel GPU design of the piecewise parabolic method
569 (PPM) for horizontal advection in an air quality model (CAMx V6.10), *Geosci. Model Dev.*,
570 16, 4367-4383, 10.5194/gmd-16-4367-2023, 2023.

571 Cao, K., Wu, Q., Wang, L., Wang, N., Cheng, H., Tang, X., Li, D., and Wang, L.: The dataset of the
572 manuscript “GPUHADVPPM V1.0: high-efficient parallel GPU design of the Piecewise
573 Parabolic Method (PPM) for horizontal advection in air quality model (CAMx V6.10)”,
574 Zenodo [data set], <https://doi.org/10.5281/zenodo.7765218>, 2023.

575 Colella, P. and Woodward, P. R.: The Piecewise Parabolic Method (PPM) for gas-dynamical
576 simulations, *Journal of Computational Physics*, 54, 174-201, [https://doi.org/10.1016/0021-](https://doi.org/10.1016/0021-9991(84)90143-8)
577 9991(84)90143-8, 1984.

578 ENVIRON: User Guide for Comprehensive Air Quality Model with Extensions Version 6.1,
579 https://camx-wp.azurewebsites.net/Files/CAMxUsersGuide_v6.10.pdf (last access: 20
580 October 2023), 2014.

581 ENVIRON: CAMx version 6.1, ENVIRON [code], available at: [https://camx-](https://camx-wp.azurewebsites.net/download/source/)
582 [wp.azurewebsites.net/download/source/](https://camx-wp.azurewebsites.net/download/source/), last access: 20 October 2023.

583 Huang, M., Huang, B., Mielikainen, J., Huang, H. L. A., Goldberg, M. D., and Mehta, A.: Further
584 Improvement on GPUBased Parallel Implementation of WRF 5-Layer Thermal Diffusion
585 Scheme, in: 2013 International Conference on Parallel and Distributed Systems, Seoul, South
586 Korea, 15–18 December 013, <https://doi.org/10.1109/icpads.2013.126>, 2013.

587 Linford, J. C., Michalakes, J., Vachharajani, M., and Sandu, A.: Automatic Generation of
588 Multicore Chemical Kernels, *IEEE Transactions on Parallel and Distributed Systems*, 22,
589 119-131, 10.1109/tpds.2010.106, 2011.

590 Mielikainen, J., Huang, B., Huang, H.-L. A., and Goldberg, M. D.: GPU Implementation of Stony
591 Brook University 5-Class Cloud Microphysics Scheme in the WRF, *IEEE Journal of Selected*
592 *Topics in Applied Earth Observations and Remote Sensing*, 5, 625-633,
593 10.1109/jstars.2011.2175707, 2012.

594 Mielikainen, J., Huang, B., Wang, J., Allen Huang, H. L., and Goldberg, M. D.: Compute unified
595 device architecture (CUDA)-based parallelization of WRF Kessler cloud microphysics
596 scheme, *Computers & Geosciences*, 52, 292-299, 10.1016/j.cageo.2012.10.006, 2013.

597 News, Frontier Remains as Sole Exaflop Machine and Retains Top Spot, Improving Upon Its
598 Previous HPL Score, available at: [https://www.top500.org/news/frontier-remains-sole-](https://www.top500.org/news/frontier-remains-sole-exaflop-machine-and-retains-top-spot-improving-upon-its-previous-hpl-score/)
599 [exaflop-machine-and-retains-top-spot-improving-upon-its-previous-hpl-score/](https://www.top500.org/news/frontier-remains-sole-exaflop-machine-and-retains-top-spot-improving-upon-its-previous-hpl-score/) (last access:
600 20 October 2023), 2023.

601 NVIDIA: CUDA C++ Programming Guide Version 10.2,
602 https://docs.nvidia.com/cuda/archive/10.2/pdf/CUDA_C_Programming_Guide.pdf (last
603 access: 20 October 2023), 2020.

604 Odman, M. and Ingram, C.: Multiscale Air Quality Simulation Platform (MAQSIP): Source Code
605 Documentation and Validation, 1996.

606 ROCm, AMD ROCm-HIP documentation, available at: <https://rocm.docs.amd.com/en/docs-5.0.0>
607 (last access: 20 October 2023), 2023.

608 Sun, J., Fu, J. S., Drake, J. B., Zhu, Q., Haidar, A., Gates, M., Tomov, S., and Dongarra, J.:
609 Computational Benefit of GPU Optimization for the Atmospheric Chemistry Modeling,
610 *Journal of Advances in Modeling Earth Systems*, 10, 1952-1969,
611 <https://doi.org/10.1029/2018MS001276>, 2018.

612 Top500, Supercomputing Top500 list, available at: <https://www.top500.org/lists/top500/2023/06/>
613 (last access: 20 October 2023), 2023.

614 Váňa, F., Düben, P., Lang, S., Palmer, T., Leutbecher, M., Salmond, D., and Carver, G.: Single
615 Precision in Weather Forecasting Models: An Evaluation with the IFS, *Mon. Weather*
616 *Rev.*, 145, 495–502, <https://doi.org/10.1175/mwr-d-16-0228.1>, 2017.

617 Wang, H., Lin, J., Wu, Q., Chen, H., Tang, X., Wang, Z., Chen, X., Cheng, H., and Wang, L.: MP
618 CBM-Z V1.0: design for a new Carbon Bond Mechanism Z (CBM-Z) gas-phase chemical
619 mechanism architecture for next-generation processors, *Geosci. Model Dev.*, 12, 749–764,
620 <https://doi.org/10.5194/gmd-12-749-2019>, 2019.

621 Wang, P., Jiang, J., Lin, P., Ding, M., Wei, J., Zhang, F., Zhao, L., Li, Y., Yu, Z., Zheng, W., Yu, Y.,
622 Chi, X., and Liu, H.: The GPU version of LASG/IAP Climate System Ocean Model version 3
623 (LICOM3) under the heterogeneous-compute interface for portability (HIP) framework and
624 its large-scale application, *Geosci. Model Dev.*, 14, 2781-2799, [10.5194/gmd-14-2781-2021](https://doi.org/10.5194/gmd-14-2781-2021),
625 2021.