

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

29 HADVPPM on the different GPU ~~accelerator accelerators~~ is ~~more obvious~~more apparent when the
30 computing scale is ~~larger~~more extensive, and the maximum speedup of GPU-HADVPPM on the
31 domestic GPU-like accelerator is 28.9 times. The hybrid parallelism with a message passing
32 interface (MPI) and HIP enables ~~achieve up to 17.2 times speedup when configure~~achieving up to
33 17.2 times speedup when configuring 32 CPU cores and GPU-like accelerators on the domestic
34 heterogeneous cluster. ~~And the OpenMP technology is introduced to further reduce the~~
35 ~~computation time of~~The OpenMP technology is introduced further to reduce the computation time
36 of the CAMx-HIP model by 1.9 times. More importantly, by comparing the simulation results of
37 GPU-HADVPPM on NVIDIA GPUs and domestic GPU-like accelerators, it is found that the
38 simulation results of GPU-HADVPPM on domestic GPU-like accelerators have less difference
39 than the NVIDIA GPUs. Furthermore, we also exhibit that the data transfer efficiency between
40 CPU and GPU has ~~an important~~a meaningful ~~essential~~ impact on heterogeneous computing, ~~and~~
41 point out that optimizing the data transfer efficiency between CPU and GPU is one of the
42 ~~important~~critical directions ~~directions~~ to improve the computing efficiency of geoscience
43 numerical models in heterogeneous clusters in the future.

44 1. Introduction

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

57 ~~In terms of the heterogeneous porting for air quality model, most scholars select the chemical~~
58 ~~module, one of the hotspots, to implement heterogeneous porting, and porting the computational~~
59 ~~process originally on the CPU processes to the GPU accelerator, in order to~~Regarding the
60 ~~heterogeneous porting for air quality model, most scholars select the chemical module, one of the~~
61 ~~hotspots, to implement heterogeneous porting, and porting the computational process initially on~~
62 ~~the CPU processes to the GPU accelerator, to~~ improve the computing efficiency. For example, Sun
63 et al. (2018) used CUDA technology to port the second-order Rosenbrock solver of ~~the~~ chemistry
64 module of CAM4-Chem to NVIDIA Tesla K20X GPU, ~~and achieved. They achieved up to~~ 11.7x
65 speedup compared to the AMD Opteron™ 6274 CPU (16 cores) using one CPU core. Alvanos and
66 Christoudias (2017) developed ~~a software that automatically generates CUDA kernels to solve~~
67 ~~chemical kinetics equation in the chemistry module for the global climate model ECHAM/MESSy~~
68 ~~Atmospheric Chemistry (EMAC)~~~~software that automatically generates CUDA kernels to solve~~
69 ~~chemical kinetics equations in the chemistry module for the global climate model~~
70 ~~ECHAM/MESSy Atmospheric Chemistry (EMAC)~~, and performance evaluation shows a 20.4x
71 speedup for the kernel execution. Linford et al. (2011) presented the Kinetic PreProcessor (KPP)
72 to generate the chemical mechanism code in CUDA language, which can be implemented on ~~the~~
73 NVIDIA Tesla C1060 GPU. The KPP-generated SAPRC'99 mechanism from ~~CMAQ model~~
74 ~~achieved a maximum speedup of 13.7x and KPP-generated RADM2 mechanism from the CMAQ~~
75 ~~model achieved a maximum speedup of 13.7x, and the KPP-generated RADM2 mechanism from~~
76 ~~the~~ WRF-chem model achieved an 8.5x speedup both compared to the Intel Quad-Core Xeon
77 5400 series CPU. Similarly, the advection module is also one of the hotspot modules in the air
78 quality model,; Cao et al. (2023) adopted the Fortran-C-CUDA C scheme and implemented a
79 series of optimizations, including ~~reducing reduction~~ the CPU-GPU communication frequency,
80 ~~optimize optimizing~~ the GPU memory access, and thread and block co-indexing, to increase the
81 computational efficiency of the HADVPPM advection solver. It can achieve up to the 18.8x
82 speedup on the NVIDIA Tesla V100 GPU compared to the Intel Xeon Platinum 8168 CPU.

83 The CUDA technology was implemented to carry out heterogeneous porting for the
84 atmospheric chemical models from the CPU processors to different NVIDIA GPU accelerators. In
85 this study, the Heterogeneous-computing Interface for Portability (HIP) interface was introduced

86 to implement the porting of GPU-HADVPPM from the NVIDIA GPU to ~~the China's~~China's
87 domestically GPU-like accelerators based on the research of Cao et al. (2023). The domestic
88 GPU-like accelerator plays the same role as the NVIDIA GPU, which is also used to accelerate the
89 advection module in the CAMx model, so we refer to it as a GPU-like accelerator. First, we
90 compared the simulation results of the Fortran version CAMx model with the CAMx-CUDA and
91 CAMx-HIP ~~model which were coupled with the CUDA and HIP versions of models, which were~~
92 ~~coupled with the CUDA and HIP versions of the~~ GPU-HADVPPM program, respectively. ~~And~~
93 ~~then, the computing performance of GPU-HADVPPM programs on different GPUs were~~Then, the
94 computing performance of GPU-HADVPPM programs on different GPUs were compared. Finally,
95 we tested ~~total coupling performance of CAMx-HIP model with multi-level hybrid parallelization~~
96 ~~on the China's~~ the total performance of the CAMx-HIP model with multi-level hybrid
97 parallelization on China's domestically heterogeneous cluster.

98 2. Model and experimental platform

99 2.1. The CAMx model description and configuration

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

112 In terms of the horizontal advection term on the right-hand side, this equation is solved using

113 either the Bott (1989) scheme or the Piecewise Parabolic Method (PPM) (Colella and Woodward,
114 1984; Odman and Ingram, 1996) scheme. The PPM horizontal advection scheme (HADVPPM)
115 was selected in this study because it provides higher accuracy with minimal numerical diffusion
116 (ENVIRON, 2014). The other numerical schemes selected during the CAMx model testing are
117 listed in Table S1. As described by Cao et al. (2023), the -fp-model precise compile flag which can
118 force the compiler to use the vectorization of some computation under value safety, is 41.4%
119 faster than the -mieee-fp compile flag, which comes from the Makefile of the official CAMx
120 version with the absolute errors of the simulation results are less than ± 0.05 ppbV. Therefore, the -
121 fp-model precise compile flag was selected when compiling the CAMx model in this research.

122 2.2. CUDA and ROCm introduction

123 Compute Unified Device Architecture (CUDA; NVIDIA, 2020) is a parallel programming
124 paradigm ~~which was~~ released in 2007 by NVIDIA. CUDA is a proprietary application
125 programming interface (API) and ~~as such~~ is only supported on NVIDIA's GPUs. ~~For the CUDA~~
126 ~~programming, it~~CUDA programming uses a programming language similar to standard C, which
127 achieves efficient parallel computing of programs on NVIDIA GPUs by adding some keywords.
128 ~~In the previous study, CUDA technology was implemented~~The previous study implemented
129 CUDA technology to port the HADVPPM program from CPU to NVIDIA GPU (Cao et al., 2023).

130 Radeon Open Compute platform (ROCm; AMD, 2023) is an open-source software platform
131 developed by AMD for HPC and hyperscale GPU computing. ~~In general, ROCm for the AMD~~
132 ~~GPU is~~The ROCm for the AMD GPU is generally equivalent to CUDA for NVIDIA GPU. ~~On the~~
133 ~~ROCm software platform, it uses the AMD's HIP interface which is a~~The ROCm software
134 platform uses the AMD's HIP interface. a C++ runtime API allowing developers to run programs
135 on AMD GPUs. In general, they are very similar ~~and their code can be converted directly by~~
136 ~~replacing the string "cuda" with "hip" in the~~, and their code can be converted directly by
137 replacing the string "cuda" with "hip" in most cases. More information about HIP API is available
138 on the AMD ROCm website (ROCm, 2023). Similar to AMD GPU, developers can also use
139 ~~ROCM-HIP programming interface to implement programs running on the China's~~the ROCm-
140 HIP programming interface to implement programs running on China's domestically GPU-like

141 accelerator. The CUDA code cannot run directly on domestic GPU-like accelerators, ~~and it needs~~
 142 ~~to~~ and must be transcoded into HIP code.

143 2.3. Hardware components and software environment of the testing system

144 Table 1 lists four GPU clusters where we conducted the experiments, two NVIDIA
 145 heterogeneous clusters ~~which have the same hardware configuration as Cao et al. (2023) and two~~
 146 ~~China's~~ that have the same hardware configuration as Cao et al. (2023), and two China's
 147 domestically heterogeneous clusters newly used in this research, namely “Songshan”
 148 supercomputer and “Taiyuan” computing platform. Two NVIDIA heterogeneous clusters are
 149 equipped with NVIDIA Tesla K40m and V100 GPU accelerators, ~~respectively~~. Both ~~two~~-domestic
 150 clusters include thousands of computing nodes, ~~and~~ each contains ~~ing~~ one China's ~~s'~~ domestically
 151 CPU processor, four China's ~~s'~~ domestically GPU-like accelerators, and 128 GB of DDR4 2666
 152 memory. The domestic CPU has four NUMA nodes, ~~each NUMA node has eight X86 based and~~
 153 each NUMA node has eight X86-based processors. The accelerator adopts a GPU-like architecture
 154 consisting of a 16 GB HBM2 device memory, ~~and~~ many compute units. The GPU-like
 155 accelerators ~~connected to CPU with PCI-E,~~ are connected to the CPU with PCI-E, and the peak
 156 bandwidth of the data transfer between main memory and device memory is 16 GB/s.

157 It is worth noting that the “Taiyuan” computing platform, ~~which~~ has been updated in three
 158 main aspects compared to the “Songshan” supercomputer. The CPU clock speed has been
 159 increased from 2.0 GHz to 2.5 GHz, the number of GPU-like computing units has been increased
 160 from 3,840 to 8,192, and the peak bandwidth between main memory and video memory has been
 161 increased from 16 GB/s to 32 GB/s. ~~In terms of~~ Regarding the software environment, the NVIDA
 162 GPU is programmed using the CUDA toolkit, and the domestic GPU-like is programmed using
 163 the ROCm-HIP toolkit developed by AMD (ROCm, 2023). More details about the hardware
 164 composition and software environment of the four heterogeneous clusters are presented in Table 1.
 165 **Table 1.** Configurations of the NVIDIA K40m cluster, NVIDIA V100 cluster, “Songshan” supercomputer, and
 166 “Taiyuan” computing platform.

	Hardware components					
	CPU			GPU		
NVIDIA K40m cluster	Intel	Xeon	E5-2682	v4	CPU	NVIDIA Tesla K40m GPU, 2880 CUDA

	@2.5 GHz, 16 cores	cores, 12 GB video memory
NVIDIA V100 cluster	Intel Xeon Platinum 8168 CPU	NVIDIA Tesla V100 GPU, 5120 CUDA cores, 16 GB video memory
	@2.7 GHz, 24 cores	
Songshan supercomputer	China' s 's domestically CPU processor A, 2.0GHz, 32 cores	China' <u>China's</u> 's domestically GPU-like accelerator A, 3840 computing units, 16 GB memory
Taiyuan computing platform	China' s 's domestically CPU processor B, 2.5GHz, 32 cores	China' <u>China's</u> '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

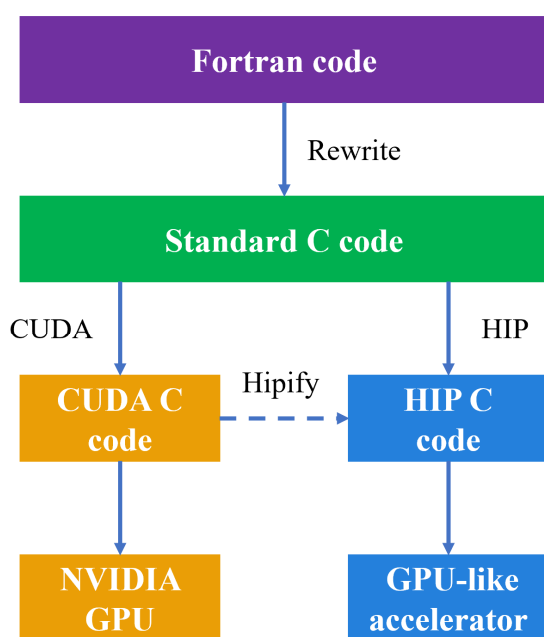
167 3. Implementation details

168 This section mainly introduced the strategy of porting [the](#) HADVPPM program from CPU to
169 NVIDIA GPU and domestic GPU-like accelerator, as well as the proposed multi-level hybrid
170 parallelism technology to make full use of computing resources.

171 3.1. Porting the HADVPPM program from CPU to NVIDIA GPU and domestic 172 GPU-like accelerator

173 Fig. 1 shows the heterogeneous porting process of HADVPPM from CPU to NVIDIA GPU
174 and domestic GPU-like accelerator. First, the original Fortran code was refactored using standard
175 C language. Then, ~~the~~ the CUDA and ROCm HIP technology were used to convert the standard C
176 code into CUDA C and HIP C code to make it computable on the NIVIDA GPU and domestic
177 GPU-like accelerator. Similar to CUDA technology, ~~the~~ HIP technology is implemented to convert
178 the standard C code to HIP C code by adding related built-in functions (such as hipMalloc,
179 hipMemcpy, hipFree, etc.). To facilitate the portability of applications across different GPU
180 platforms, ROCm provides hipify toolkits to help transcode. The ~~hipify~~Hipify toolkit is
181 essentially a simple script written in the Perl language, and its function is text replacement, which
182 replaces the function name in CUDA C code with the corresponding name in HIP C code
183 according to [certain-specific](#) rules. For example, ~~for the memory allocation function cudaMalloc in~~

184 ~~CUDA, the hipify toolkit can automatically recognize and replace it, the hipify toolkit can~~
 185 ~~automatically recognize and replace the memory allocation function cudaMalloc in CUDA~~ with
 186 hipMalloc. Therefore, the thread and block configuration of ~~the~~ GPU ~~remain~~-remains unchanged
 187 due to the simple text substitution during the transcoding. In this study, the ROCm HIP technology
 188 was used to implement the operation of GPU-HADVPPM on ~~the~~ domestic GPU-like accelerator
 189 based on the CUDA version of GPU-HADVPPM ~~which was~~ developed by Cao et al. (2023). The
 190 HIP code was compiled using the “hipcc” compiler driver with the library flag “-lamdhip64”.



191
 192 **Figure 1.** The heterogeneous porting process of HADVPPM Fortran code from CPU to NVIDIA GPU and
 193 domestic GPU-like accelerator.

194 **3.2. Multi-level hybrid parallelization of CAMx model on heterogeneous**
 195 **platform**

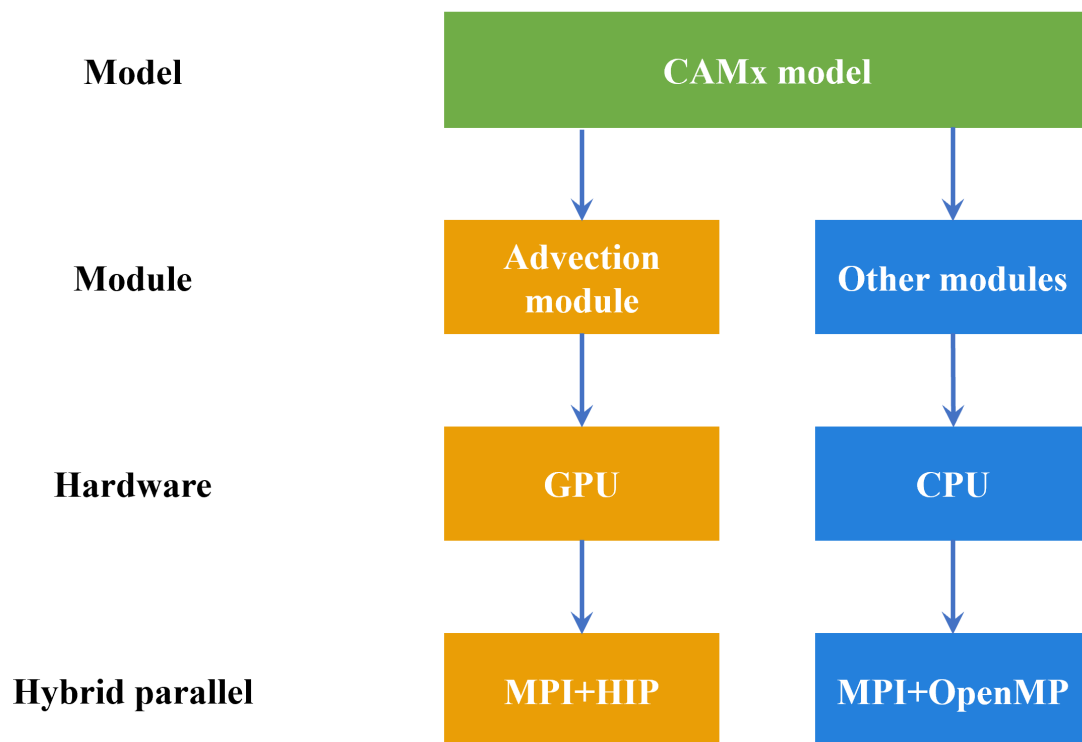
196 The original CAMx model running on the CPUs supports two types of parallelization
 197 (ENVIRON, 2014): (1) OpenMP (OMP), which supports multi-platform (e.g., multi-core) shared-
 198 memory programming in C/C++ and Fortran; (2) Message Passing Interface (MPI), which is a
 199 message passing interface standard for developing and running parallel applications on the

200 distributed-memory computer cluster. During the process of CAMx model simulation, MPI and
201 OMP hybrid parallelism can be used, several CPU processes can be launched, and each process
202 can spawn several threads. This hybrid parallelism can significantly improve the computational
203 efficiency of the CAMx model.

204 As mentioned, the original CAMx model supports message passing interface (MPI) parallel
205 technology running on the general-purpose CPU. The simulation domain is divided into several
206 sub-regions by MPI, and each CPU process is responsible for the computation of its sub-region.
207 To expand the heterogeneous parallel scale of the CAMx model on the Songshan supercomputer, a
208 hybrid parallel architecture with an MPI and HIP was adopted to make full use of GPU computing
209 resources. Firstly, we use the ROCm-HIP library function hipGetDeviceCount to obtain the
210 number of GPU accelerators configured for each compute node. Then, the total number of
211 accelerators to be launched and the ID number of accelerator cards in each node were determined
212 according to the MPI process ID number and the remainder function in standard C language.
213 Finally, the hipSetDevice library function in ROCm-HIP is used to configure an accelerator for
214 each CPU core.~~As mentioned above, the original CAMx model supports message passing~~
215 ~~interface (MPI) parallel technology running on the general-purpose CPU. The simulation domain~~
216 ~~is divided into several sub-regions by MPI, and each CPU process is responsible for computation~~
217 ~~of its sub-region, which includes the computation tasks of advection module and other modules~~
218 ~~such as photolysis module, deposition module, chemical module, etc. In the previous studying,~~
219 ~~Cao et al. (2023) adopt a parallel architecture with an MPI and CUDA (MPI+CUDA) hybrid~~
220 ~~paradigm to configure one GPU accelerator for each CPU process. For the advection module, the~~
221 ~~simulation originally implemented by the CPU is handed over to the GPU. Other module~~
222 ~~computing tasks continue to be completed on the CPU.~~

223 ~~In this study, GPU HADVPPM with an MPI and HIP heterogeneous hybrid programming~~
224 ~~technology can~~This study uses GPU-HADVPPM with an MPI and HIP heterogeneous hybrid
225 programming technology to run on multiple domestic GPU-like accelerators. However, the
226 number of GPU-like accelerators in a single compute node is usually much smaller than the
227 number of CPU cores in ~~the~~-heterogeneous HPC systems. Therefore, ~~in order to~~to make full use of
228 the remaining CPU computing resources, the OMP API of the CAMx model is further introduced

229 to realize the MPI+OMP hybrid parallelism of other modules on the CPU. A schematic of the
 230 multi-level hybrid parallel framework is shown in Fig.2. For example, ~~in a computing node, four~~
 231 ~~CPU processes and four GPU-like accelerators are launched, four CPU processes and four GPU-~~
 232 ~~like accelerators are launched in a computing node~~, and each CPU process spawns four threads.
 233 Then the advection module is simulated by 4 GPU-like accelerators, and ~~the other modules are~~
 234 ~~done by 4*4 threads spawned by CPU processes~~ 4*4 threads spawned by CPU processes do the
 235 other modules.



236
 237 **Figure 2.** A schematic of the multi-level hybrid parallel framework.

238 **4. Results and evaluation**

239 The computational performance experiments of CUDA and HIP version GPU-HADVPPM
 240 are reported in this section. First, we compared the simulation result of the Fortran version CAMx
 241 model with ~~CAMx-CUDA and CAMx-HIP model which were coupled with CUDA and HIP~~
 242 ~~version of the CAMx-CUDA and CAMx-HIP models, which were coupled with the CUDA and~~
 243 ~~HIP versions of the~~ GPU-HADVPPM program, respectively. Then, the computational
 244 performance of GPU-HADVPPM programs on the NVIDIA GPU and domestic GPU-like

245 accelerator are compared. Finally, we tested ~~total performance of CAMx-HIP model with multi-~~
246 ~~level hybrid parallelization on the the~~ the total performance of the CAMx-HIP model with multi-
247 level hybrid parallelization on the "Songshan" supercomputer. For ease of description, the CAMx
248 versions of the HADVPPM program written in Fortran, CUDA C ~~and HIP C code are named~~
249 Fortran, CUDA, and HIP C code are named Fortran, CUDA, and HIP, respectively.

250 4.1. Experimental setup

251 ~~There are three~~ Three test cases were used to evaluate the performance of CUDA and HIP
252 version GPU-HADVPPM. The experimental setup for the three test cases is shown in Table 2. In
253 the previous study of Cao et al. (2023), the BJ case was used to carry out the performance tests,
254 ~~HN case and ZY case are~~ and the HN case and ZY case was the newly constructed test cases in this
255 study. The Beijing case (BJ) covers Beijing, Tianjin, and part of the Hebei Province with $145 \times$
256 157 grid boxes, and the simulation of the BJ case starts on 1 November, 2020. The Henan case
257 (HN) mainly covers the Henan Province with 209×209 grid boxes. The starting date of
258 simulation in the HN case is 1 October, 2022. The Zhongyuan case (ZY) has the widest coverage
259 of the three cases, with Henan Province as the center, covering the Beijing-Tianjin-Hebei region,
260 Shanxi Province, Shaanxi Province, Hubei Province, Anhui Province, Jiangsu Province, and
261 Shandong Province, with 531×513 grid boxes. ZY case started simulation on 4 January, 2023.
262 ~~All of the three performance test cases are~~ three performance test cases have a 3km horizontal
263 resolution, 48 hours of simulation, and 14 vertical model layers. The number of three-dimensional
264 grid boxes in BJ, HN, and ZY cases ~~are totally 318,710, 611,534~~ total 318,710, 611,534, and
265 3,813,642, respectively. The meteorological fields inputting the different versions of the CAMx
266 model in the three cases were provided by the Weather Research and Forecasting Model (WRF).
267 In terms of emission inventories, the emission for ~~BJ case is consistent with the Cao et al. (2023);~~
268 ~~HN case uses the Multi-resolution Emission Inventory for China (MEIC) and the BJ case is~~
269 consistent with the Cao et al. (2023), the HN case uses the Multi-resolution Emission Inventory
270 for China (MEIC). The ZY case uses the emission constructed by the Sparse Matrix Operator
271 Kernel Emission (SMOKE) model in this study.

272 **Table 2.** The experimental setup for the BJ, HN, and ZY ~~cases~~ cases.

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

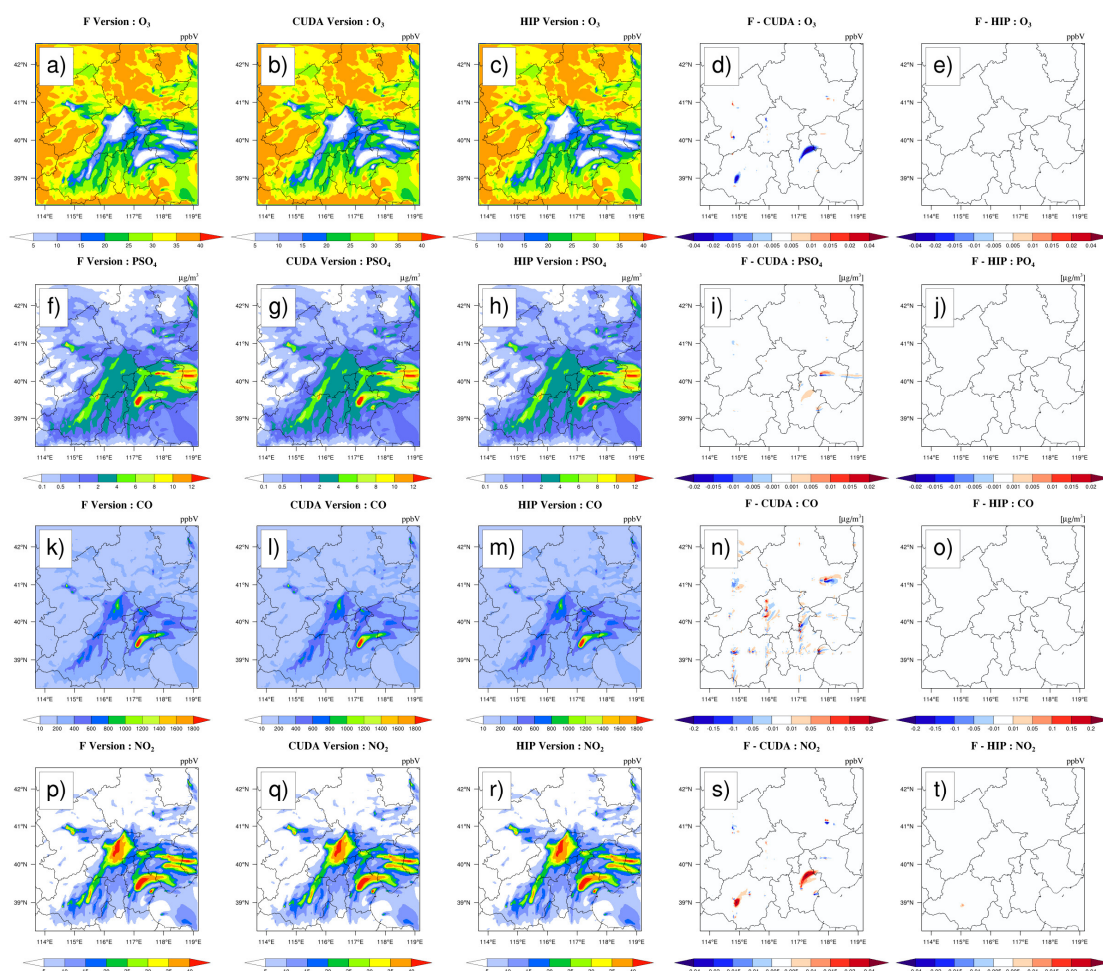
273 4.2. Error analysis

274 The hourly concentrations of four major species, i.e., O₃, PSO₄, CO, and NO₂, outputted by
275 the Fortran, CUDA, and HIP versions of CAMx for the BJ case are compared to verify the results
276 correctness before testing the computational performance. Fig. 3 shows the four major species
277 simulation results of the three CAMx ~~version versions~~, including the Fortran version on the Intel
278 E5-2682 v4 CPU, the CUDA version on the NVIDIA K40m cluster, and the HIP version on the
279 “Songshan” supercomputer, after 48 hours integration, as well as the absolute errors (AEs) of their
280 concentrations. As described by Cao et al. (2023), the parallel design of the CAMx model adopts
281 the primary/secondary mode, and the P0 process is responsible for inputting and outputting the
282 data and calling the MPI_Barrier function to synchronize the process, and the other processes are
283 ~~responsible-accountable~~ for simulation. When comparing the simulation results, we only launched
284 2 CPU processes on the CPU platform, ~~and launched 2 CPU processes~~ launched 2 CPU processes
285 and configured 2 GPU accelerators on the NVIDIA K40m cluster and “Songshan” supercomputer,
286 respectively.

287 The species’ spatial pattern of ~~three CAMx versions on different platform are visually very~~
288 ~~consistent, and the AEs between the HIP and Fortran version is much smaller than the CUDA and~~
289 ~~Fortran version~~ the three CAMx versions on different platforms are visually very consistent. The
290 AEs between the HIP and Fortran versions are much smaller than the CUDA and Fortran versions.
291 For example, the AEs between the CUDA and Fortran ~~version versions~~ for O₃, PSO₄, and NO₂ are
292 in the range of ±0.04 ppbV, ±0.02 μg · m⁻³, and ±0.04 ppbV. ~~And the~~ The AEs between the HIP
293 and Fortran versions ~~for above~~ the three species ~~are~~ fall into the range of ±0.01 ppbV, ±0.005 μg ·
294 m⁻³, and ±0.01 ppbV. For CO, AEs ~~is relatively large due to its~~ are relatively large due to their
295 high background concentration. However, the AEs between the HIP and Fortran versions ~~are~~ also
296 less than ~~those~~ at between the CUDA and Fortran versions, ~~which where~~ were in the range of ±0.4

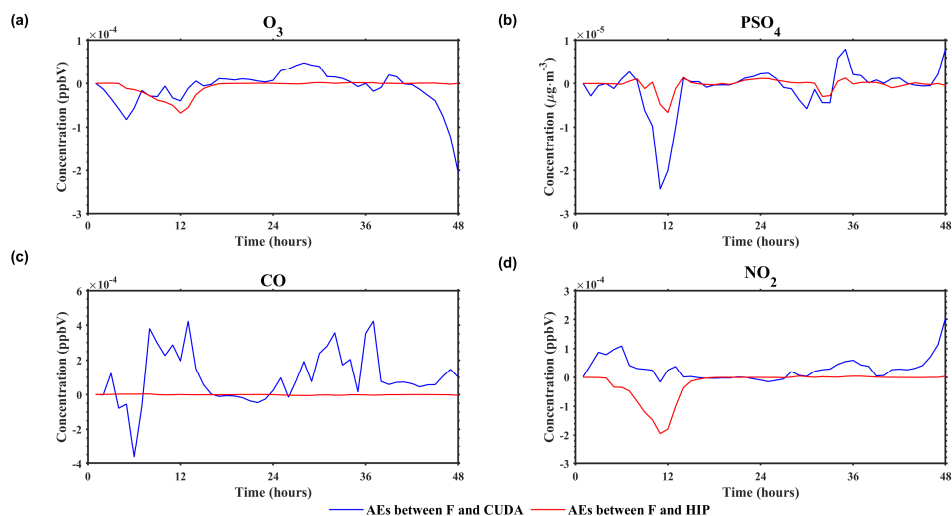
297 ppbV and ± 0.1 ppbV, respectively.

298 Considering the situation of ~~AEs accumulate and grow~~AEs accumulation and growth, Fig. 4
299 highlights the time series of AEs between Fortran and CUDA versions and between Fortran and
300 HIP versions after grid averaging. As is shown in Fig. 4, the AEs of O₃, PSO₄, CO, and NO₂
301 between the Fortran version and the CUDA version are -0.0002 to 0.0001 ppbV, -0.00003 to
302 0.00001 $\mu\text{g} \cdot \text{m}^{-3}$, -0.0004 to 0.0004 ppbV, and -0.0002 to 0.0002 ppbV, respectively, and
303 fluctuate. Although the AEs of the above four species between the Fortran and the HIP version
304 also ~~fluctuates~~fluctuate, the fluctuation range is much smaller than that of the CUDA version.
305 ~~Importantly~~Notably, the AEs between Fortran and CUDA versions and between Fortran and HIP
306 versions ~~both~~ do not accumulate and grow over prolonged simulation periods.



307
308 **Figure 3.** O₃, PSO₄, CO, and NO₂ concentrations outputted by the CAMx Fortran version on the Intel E5-2682 v4
309 CPU, CUDA version on the NVIDIA K40m cluster, and HIP version on the "Songshan" supercomputer under the
310 BJ case. Panels (a), (f), (k), and (p) are from the Fortran version of simulation results for four species. Panels (b),
311 (g), (l), and (q) are from the CUDA version of simulation results for four species. Panels (c), (h), (m), and (r) are

312 from the HIP version of simulation results for four species. Panels (d), (i), (n), and (s) are the AEs between the
 313 Fortran and CUDA versions. Panels (e), (j), (o), and (t) are the AEs between the Fortran and HIP versions.



314
 315 **Figure 4.** The time series of AEs between Fortran and CUDA versions (solid blue line) and between Fortran and
 316 HIP versions (solid red line) after grid averaging. After grid averaging, the time series of AEs between Fortran and
 317 CUDA versions (solid blue line) and between Fortran and HIP versions (solid red line). Panel (a)~(d) represent the
 318 AEs of O₃, PSO₄, CO, and NO₂, respectively.

319 To further detail the differences in the simulation results, we supplement the offline
 320 experimental results of the advection module on the NVIDIA K40m cluster and the Songshan
 321 supercomputer. First, we construct the Fortran programs to provide consistent input data for the
 322 advection module written in CUDA C code and HIP C code on NVIDIA Tesla K40m GPU and
 323 domestic GPU-like accelerator, respectively. The accuracy of the input data is kept at 12 decimal
 324 places. Then, the advection module outputs and prints the computing results after completing one
 325 integration operation on different accelerators. Finally, the results of the various accelerators were
 326 compared with those of the Fortran code on the Intel Xeon E5-2682 v4 CPU processor. The
 327 specific results are shown in the Fig.5. The difference in the computing results of the advection
 328 module written in HIP C code on the domestic GPU-like accelerator is smaller than that of the
 329 CUDA C code on the NVIDIA Tesla K40m GPU. The mean relative errors (REs) and AEs of the
 330 computing results on the NVIDIA Tesla K40m GPU are $1.3 \times 10^{-5} \%$ and 7.1×10^{-9} ,
 331 respectively, while on the domestic GPU-like accelerator, the mean REs and AEs of the results are
 332 $5.4 \times 10^{-6} \%$ and 2.6×10^{-9} , respectively.

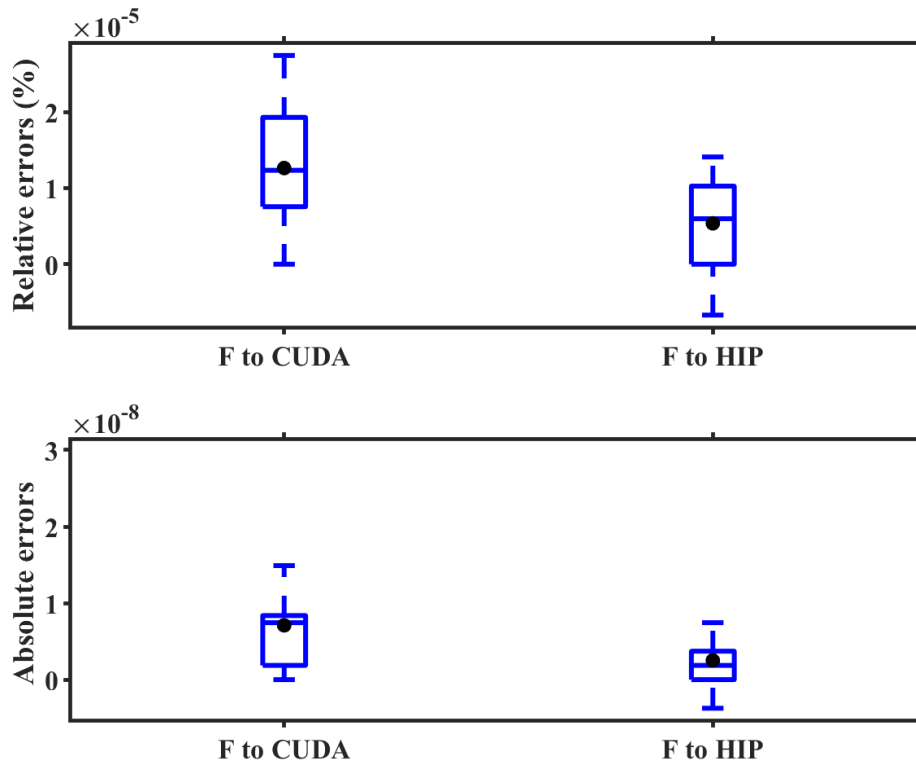



Figure 5. The boxplots of REs and AEs between the Fortran code on Intel Xeon E5-2682 v4 CPU and CUDA C code on NVIDIA Tesla K40m GPU, and between HIP C code on domestic GPU-like accelerator, respectively, in the case of offline testing.

Fig. 65 further presents the boxplot of the REs in all grid boxes for the PSO_4 , PNO_3 , PNH_4 , O_3 , CO , and NO_2 during the 48-hour simulation under the BJ case. Statistically, the REs between the CUDA version on the NVIDIA K40m cluster and Fortran version on the Intel E5-2682 v4 CPU for the above six species are in the range of $\pm 0.006\%$, $\pm 0.01\%$, $\pm 0.008\%$, $\pm 0.002\%$, $\pm 0.002\%$, and $\pm 0.002\%$. In terms of REs between the HIP version on the “Songsshan” supercomputer and the Fortran version on the Intel E5-2682 v4 CPU, the values are much smaller than REs between CUDA and Fortran versions which are fall into the range of $\pm 0.0005\%$, $\pm 0.004\%$, $\pm 0.004\%$, $\pm 0.00006\%$, $\pm 0.00004\%$, and $\pm 0.00008\%$, respectively. In the air quality model, the initial concentration of secondary fine particulate matter such as PSO_4 , PNO_3 , and PNH_4 is very low and is mainly generated by complex chemical reactions. The integration process of the advection module is ported from the CPU processor to the GPU accelerator, which will lead to minor differences in the results due to different hardware. The low initial concentration of

349 secondary fine particulate matter is sensitive to these minor differences, which may eventually
 350 lead to a higher difference in the simulation results of secondary particulate matter than other
 351 species.In the air quality model, the secondary particulate matter, such as PNH_4 , PNO_3 , and PSO_4 ,
 352 have a common characteristic: their initial concentration is very low and they are mainly generated
 353 through complex chemical reactions. Therefore, when calculating the relative error on different
 354 hardware platforms, because the value in the denominator is very small, it is very sensitive to a
 355 small difference in the numerator, resulting in a large relative error. But from the absolute error in
 356 Fig.3, the absolute error of PSO_4 on different hardware platforms is smaller than that of other
 357 species. For gaseous pollutants such as CO , O_3 , and NO_2 , the initial concentration is large due to
 358 emission, and the denominator value is large when calculating the relative error, which is
 359 insensitive to small differences in the numerator.



360
 361 **Figure 65.** The distribution of REsREs distribution in all grid boxes for the PSO_4 , PNO_3 , PNH_4 , O_3 , CO , and NO_2
 362 under the BJ case. The red boxplot represents the REs between the CUDA version on the NVIDIA K40m cluster
 363 and Fortran version on the Intel E5-2682 v4 CPU, and blue boxplot represents the REs between the HIP version on
 364 the “Songshan” supercomputer and the Fortran version on the Intel E5-2682 v4 CPU, and the blue boxplot
 365 represents the REs between the HIP version on the “Songshan” supercomputer and the Fortran version on the Intel
 366 E5-2682 v4 CPU.

367 Wang et al. (2021) verified the applicability of the numerical model in scientific research by
 368 computing the ratio of root mean square error (RMSE) between two different model versions to

369 system spatial variation (standard deviation, std). If the ratio is smaller, it is indicated that the
 370 difference in the simulation results of the model on the GPU is minimal compared with the spatial
 371 variation of the system, ~~that~~. That is to say, the simulation results of the model on the GPU are
 372 accepted for scientific research. Here, we calculate the standard deviation of O₃, PSO₄, CO ~~and~~
 373 ~~NO₂ on the Intel Xeon E5-2682 v4 CPU, and their RMSE between the NVIDIA V100 cluster,~~
 374 ~~NVIDIA K40m cluster, and NO₂ on the Intel Xeon E5-2682 v4 CPU and their RMSE between the~~
 375 ~~NVIDIA V100 cluster, NVIDIA K40m cluster,~~ and "Songshan" supercomputer and the Intel Xeon
 376 E5-2682 v4 CPU, which are presented in Table 3. The std for the above four species on the Intel
 377 Xeon E5-2682 v4 CPU are 9.6 ppbV, 1.7 $\mu\text{g} \cdot \text{m}^{-3}$, 141.9 ppbV, and 7.4 ppbV, respectively, and
 378 their ratios of RMSE and std on the "Songshan" supercomputer are $5.8 \times 10^{-5}\%$, $4.8 \times 10^{-6}\%$,
 379 $5.7 \times 10^{-8}\%$, and $2.1 \times 10^{-4}\%$, which are smaller than two NVIDIA clusters, ~~especially~~
 380 ~~significantly~~ much smaller than the NVIDIA V100 cluster. For example, the ratio on the NVIDIA
 381 K40m cluster for four species are $1.2 \times 10^{-4}\%$, $6.6 \times 10^{-5}\%$, $7.0 \times 10^{-5}\%$, and $4.1 \times 10^{-4}\%$,
 382 and ratio on the NVIDIA V100 cluster are $1.5 \times 10^{-2}\%$, $2.5 \times 10^{-3}\%$, $6.4 \times 10^{-3}\%$, and
 383 $1.3 \times 10^{-3}\%$, respectively.

384 **Table 3.** The standard deviation (std) of O₃, PSO₄, CO₂ and NO₂ on the Intel Xeon E5-2682 v4 CPU, root mean
 385 square error (RMSE), and its ratio on the NVIDIA V100 cluster, NVIDIA K40m cluster, and "Songshan"
 386 supercomputer

	std	NIVIDA V100 cluster		NIVIDA K40m cluster		"Songshan" supercomputer	
		RMSE	RMSE/std	RMSE	RMSE/std	RMSE	RMSE/std
O ₃ (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}
PSO ₄ ($\mu\text{g} \cdot \text{m}^{-3}$)	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}
NO ₂ (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}

387 From AEs, REs, and the ratio of RMSE and std between different CAMx versions, ~~it~~ there is
 388 less difference that the GPU-HADVPPM4HIP program runs on the "Songshan" supercomputer.
 389 Because the simulation accuracy of geoscience numerical model is closely related to the model
 390 efficiency, and many model optimization works improve the computational performance by
 391 reducing the precision of the data, such as Váňa et al. (2017) changed some variables precision in
 392 the atmospheric model from double precision to single precision, which increased the overall
 393 computational efficiency by 40%, and Wang et al. (2019) improved the computational efficiency

394 of the gas-phase chemistry module in the air quality mode by 25%~28% by modifying the
395 floating-point precision compile flag. Therefore, we speculate that this may be related to the
396 manufacturing process of NVIDIA GPUs and domestic GPU-like accelerators, which may use
397 unknown optimizations to improve GPU performance efficiency by losing part of the accuracy. In
398 this study, we mainly focus on numerical simulation. Of course, we also want to know the specific
399 reasons for this, ~~but we~~. Still, we are not professional GPU research and development designers
400 after all and do not know the underlying design logic of the hardware, so we can only present our
401 experimental results in the air pollution model to you, and discuss with each other to jointly
402 promote the application of GPU in the field of geoscience numerical models.

403 **4.3. Application performance**


404 **4.3.1. GPU-HADVPPM on a single GPU accelerator**

405 As described in Sect. 4.2, we validate the ~~48 hours~~48-hour simulation results outputted by the
406 Fortran, CUDA, and HIP versions of CAMx model. Next, computational performance was
407 compared for the Fortran version of HADVPPM on the Intel Xeon E5-2682 v4 CPU and domestic
408 CPU processor A, the CUDA version of GPU-HADVPPM on the NVIDIA Tesla K40m and V100
409 GPU, and the HIP version of GPU-HADVPPM on the domestic GPU-like accelerator A, under the
410 BJ, HN and ZY case. The simulation time in this section is 1 hour unless otherwise specified.

411 ~~Similarity, since the CAMx model adopts the primary/secondary mode, two CPU processes~~
412 ~~P0 and P1 are launched on the CPU, and the system_clock functions in the Fortran language are~~
413 ~~used to test the elapsed time of the advection module in~~ Similarly, since the CAMx model adopts
414 the primary/secondary mode, two CPU processes, P0 and P1, are launched on the CPU, and the
415 system_clock functions in the Fortran language are used to test the elapsed time of the advection
416 module in the P1 process. When testing the computation performance of the advection module on
417 the GPU-like accelerator, we ~~also~~ only launch 2 CPU processes and 2 GPU-like accelerators.
418 When a P1 process runs to the advection module, the original computation process is migrated
419 from the CPU to the GPU, and the hipEvent_t function in HIP programming is used to test the
420 running time of the advection module on the GPU-like accelerator. When comparing the speedup
421 on different GPU accelerators, the elapsed time of ~~advection module launched one CPU process~~


422 ~~(P1) on the domestic CPU processor A is taken as the benchmark, that~~the advection module
423 ~~launched one CPU process (P1) on the domestic CPU processor A is taken as the benchmark; that~~
424 is, the speedup is 1.0x. The runtime of the advection module on Intel CPU processor and different
425 GPU accelerators is compared with the baseline to obtain the speedup.

426 Fig. 76(a) and (b) ~~shows show~~ the elapsed time and speedup of the different versions of
427 HADVPPM on the CPU processors and GPU accelerators for BJ, HN, and ZY cases, respectively.
428 The results show that ~~using~~CUDA and HIP technology to port HADVPPM from CPU to GPU can
429 significantly improve its computational efficiency. For example, the elapsed time of the advection
430 module on the domestic processor A is 609.2 seconds under the ZY case. After it is ported to the
431 domestic GPU accelerator and NVIDIA Tesla V100 GPU, it only takes 21.1 seconds and 7.6
432 seconds to complete the computing, and the speedups are 28.9x and 80.2x, respectively. The ZY
433 case had the ~~largest-most significant~~ number of grids in the three cases ~~and exceeded. It exceeded~~
434 the memory of a single NVIDIA Tesla K40m GPU accelerator, so it was not possible to test its
435 elapsed time on it. Moreover, the optimization of thread and block co-indexing is used ~~to~~
436 ~~simultaneously compute the grid point in the horizontal direction~~to compute the grid point in the
437 ~~horizontal direction simultaneously~~ (Cao et al., 2023). Therefore, it can be seen from Fig. 6(b) that
438 the larger the computing scale, the ~~more-obvious~~more pronounced the acceleration, which
439 indicates that GPU is more suitable for super-large scale parallel computing, ~~and~~ provides
440 technical support for accurate and fast simulation of ultra-high-resolution air quality at the meter
441 level in the future.



442
 443 **Figure 76.** The elapsed time (a) and speedup (b) of the Fortran version of HADVPPM on the Intel Xeon E5-2682
 444 v4 CPU and the domestic CPU processor A, the CUDA version of GPU-HADVPPM on the NVIDIA Tesla K40m
 445 GPU, NVIDIA Tesla V100 GPU, and the HIP version of GPU-HADVPPM on the domestic GPU-like accelerator
 446 A for BJ, HN, and ZY case. The unit of elapsed time is in seconds (s).

447 The ~~timestep of BJ, HN and ZY case~~ BJ, HN, and ZY case timestep were 59, 47, and 61,
 448 respectively. Fig. 87 shows the GPU-HADVPPM4HIP acceleration in each time step on a single
 449 domestic GPU-like accelerator A. It can be seen from the figure that all three cases have the
 450 smallest speedup of 8.2x, 11.2x, and 27.8x at the first timestep, which is related to the time
 451 required for GPU-like accelerator startup. When the GPU-like is started and operating normally,
 452 the speedup of the three cases ~~tend to be stable in the following time steps, and stabilize around~~
 453 8.5x, 11.5x and 28.0x ~~extends to be stable in the following time steps and stabilize around~~
 454 and 28.0x, respectively.



455 **Figure 87.** The GPU-HADVPPM4HIP acceleration in each time step on a single GPU-like accelerator for BJ, HN,
 456 and ZY cases. The timestep of the above three cases are 59, 47, and 61, respectively.
 457

458 Table 4 further lists the total elapsed time of CAMx Fortran and HIP versions for BJ case on
 459 the "Songshan" supercomputer and "Taiyuan" computing platform, ~~and the computing time of~~
 460 ~~and the computing time of the~~ advection module with or without data transfer. By coupling the
 461 GPU-HADVPPM4HIP to the CAMx model and adopting a series of optimizations (Cao et al.,
 462 2023), such as communication optimization, memory access optimization, and 2D thread
 463 optimization, the overall computation time of the CAMx-HIP model on a single domestic GPU-
 464 like accelerator is faster than that of the original Fortran version on a single domestic CPU core.
 465 For example, on the "Songshan" supercomputer, one hour of simulation in the CAMx-HIP model
 466 takes 469 seconds, and the Fortran version takes 481 seconds. On the "Taiyuan" computing
 467 platform, the acceleration effect is more ~~obvious-evident~~ due to the upgrade of hardware and
 468 network bandwidth, ~~and the.~~ The integration time of the CAMx-HIP model is 433 seconds when
 469 maintaining the same software environment, and the integration time of the Fortran version is 453
 470 seconds.

471 The elapsed time of GPU-HADVPPM given in Table 4 on NVIDIA GPU and domestic GPU-
 472 like accelerator does not consider the data transfer time between CPU and GPU. However, the
 473 communication bandwidth of data transfer between the CPU and GPU is one of the most
 474 significant factors that restrict the performance of the numerical model on the heterogeneous
 475 cluster (Mielikainen et al., 2012; Mielikainen et al., 2013; Huang et al., 2013). To illustrate the
 476 significant impact of CPU-GPU data transfer efficiency, the computational performance of GPU-

477 HADVPPM with and without data transfer time for the BJ case is tested on the “Songshan”
478 supercomputer and “Taiyuan” computing platform with the same DTK version 23.04 software
479 environment, and the results are further presented in Table 6. For convenience of description, we
480 refer to the execution time of GPU-HADVPPM program on GPU kernel as kernel execution time,
481 and the time of GPU-HADVPPM running on GPU as total runtime, which contains two parts,
482 namely, kernel execution time and data transfer time between CPU and GPU. After testing, the
483 kernel execution time and total running time of the GPU-HADVPPM4HIP program on domestic
484 GPU-like accelerator A are 6.8 and 29.8 seconds, respectively. In other words, it only takes 6.8
485 seconds to complete the computation on the domestic accelerator, ~~but it.~~ Still, it takes 23.0
486 seconds to complete the data transfer between the CPU and the domestic GPU-like accelerator,
487 which is 3.4 times the computation time. The same problem exists in the more advanced the
488 “Taiyuan” computing platform, where the GPU-HADVPPM4HIP takes only 5.7 seconds to
489 complete the computation, while the data transmission takes 18.2 seconds, ~~which is~~ 3.2 times the
490 computation time.

491 By comparing the kernel execution time and total running time of GPU-HADVPPM4HIP on
492 the domestic accelerator, it can be seen that the data transfer efficiency between CPU and GPU is
493 ~~really~~ inefficient, which seriously restricts the computational performance of numerical models in
494 heterogeneous clusters. On the one hand, improving the data transfer bandwidth between CPU and
495 GPU can improve the computational efficiency of the model in heterogeneous clusters. On the
496 other hand, ~~the~~ optimization measures can be implemented to improve the data transfer efficiency
497 between CPU and GPU. For example, (1) Asynchronous data transfer ~~is used to reduce~~ reduces the
498 communication latency between CPU and GPU. Computation and data transfer are performed
499 simultaneously to hide communication overhead; (2) Currently, some advanced GPU architectures
500 support a unified memory architecture, so that the CPU and GPU can share the same memory
501 space and avoid frequent data transfers. This reduces the overhead of data transfer and improves
502 data transfer efficiency; (3) Cao et al. (2023) adopted communication optimization measures to
503 ~~reduce~~ minimize the communication frequency in one-~~time~~ integration step to one, but there is
504 still the problem of high communication frequency in the whole simulation. In the future, we will
505 consider porting other hotspots of ~~CAMx model, or even the whole integral module except I/O, to~~

~~GPU-like accelerators for increasing the proportion of code on the GPU and reduce the CAMx model or even the entire integral module except I/O, to GPU-like accelerators for increasing the proportion of code on the GPU and reducing~~ the frequency of CPU-GPU communication.

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

Table 4. The total elapsed time of CAMx Fortran and HIP versions for the BJ case on the "Songshan" supercomputer and "Taiyuan" computing platform, and the computing time of the advection module with or without data transfer. The 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

4.3.2. CAMx-HIP model on the heterogeneous cluster


Generally, heterogeneous HPC systems have thousands of compute nodes equipped with one or more GPUs on each compute node. To fully use multiple GPUs, the hybrid parallelism with an MPI and HIP paradigm was used to implement the HIP version of GPU-HADVPPM run on multiple domestic GPU-like accelerators. During the simulation of the CAMx model, the emission, advection, dry deposition, diffusion, wet deposition, photolysis process, and chemical process will be computed sequentially. In heterogeneous computing platforms, except for the advection process, the CPU processor completes the simulation of the rest of the processes, and the advection process is completed on the GPU accelerator. For example, using MPI and HIP hybrid parallel technology to launch four CPU processes and four GPU accelerators simultaneously, the advection process is

530 ~~completed on four GPUs, and the other processes are still completed on four CPU~~
531 ~~processes. Generally, the heterogeneous HPC systems have thousands of compute nodes which are~~
532 ~~equipped with one or more GPUs on each compute node. To make full use of multiple GPUs, a~~
533 ~~parallel architecture with an MPI and CUDA hybrid paradigm was implemented to improve the~~
534 ~~overall computational performance of CAMx-CUDA model (Cao et al., 2023). In this studying,~~
535 ~~the hybrid parallelism with an MPI and HIP paradigm was used to implement the HIP version of~~
536 ~~GPU-HADVPPM run on multiple domestic GPU-like accelerators.~~

537 Fig.98 shows the total elapsed time and speedup of ~~CAMx-HIP model which the~~ CAMx-HIP
538 model, which is coupled with the HIP version GPU-HADVPPM on the "Songshan"
539 supercomputer under the BJ, HN, and ZY cases. The simulation of ~~above three cases for one hour~~
540 ~~took 488 seconds, 1135 seconds and 5691 seconds respectively~~ the above three cases for one hour
541 took 488 seconds, 1135 seconds, and 5691 seconds, respectively, when launching two domestic
542 CPU processors and two GPU-like accelerators. For the BJ and HN case, the parallel scalability is
543 highest when configured with 24 CPU cores and 24 GPU-like accelerators, with speedup of 8.1x
544 and 11.6x, respectively. ~~In terms of~~ Regarding the ZY case, due to its large number of grids, the
545 parallel scalability is the highest when 32 CPU cores and 32 GPU-like accelerators are configured,
546 and the acceleration ratio is 17.2x.

547 As mentioned above, data transfer between CPU and GPU takes several times more time than
548 computation. Regardless of the CPU-GPU data transfer consumption, GPU-HADVPPM4HIP can
549 achieve up to 28.9x speedup on a single domestic GPU-like accelerator. However, in terms of the
550 total time consumption, the CAMx-HIP model is only 10~20 seconds faster than the original
551 Fortran version when one GPU-like accelerator is configured. ~~And as~~ As the number of CPU cores
552 and GPU-like accelerators increases, the overall computing performance of the CAMx-HIP model
553 is lower than that of the original Fortran version. The main reason is related to the amount of data
554 transferred to GPU. As the number of MPI processes increases, the number of grids responsible
555 for each process decreases, and the amount of data transmitted by the advection module from CPU
556 to GPU decreases. However, GPUs are suitable for large-scale matrix computing. When the data
557 scale is small, the performance of ~~GPU is low, and the communication efficiency between the~~
558 GPU is low, and the communication efficiency between the CPU and GPU is the biggest

559 bottleneck (Cao et al., 2023). Therefore, the computational performance of the CAMx-HIP model
 560 is not as good as the original Fortran version when MPI processes increase. According to the
 561 characteristics of GPUs suitable for large-scale matrix computing, the model domain can be
 562 expanded ~~and the model resolution can be increased in the future to ensure that the amount of data~~
 563 ~~transferred to each GPU reaches the maximum video memory occupation,~~ and the model
 564 resolution can be increased in the future to ensure that the amount of data transferred to each GPU
 565 reaches the maximum video memory occupation so as to make efficient use of GPU. In addition,
 566 the advection module only accounts for about 10% of the total time consumption in the CAMx
 567 model (Cao et al., 2023), ~~and in~~ In the future, ~~it is considered to port the entire integration module~~
 568 ~~except I/O to the GPU~~ porting the entire integration module except I/O to the GPU is supposed to
 569 minimize the communication frequency.




570
 571 **Figure 98.** The total elapsed time and speedup of the CAMx-HIP model on the "Songshan" supercomputer under
 572 the BJ, HN, and ZY cases. The unit is in seconds (s).

573 The number of GPU accelerators in a single compute node is usually much smaller than the
 574 number of CPU cores in ~~the~~ heterogeneous HPC systems. Using the hybrid parallel paradigm with
 575 MPI and HIP to configure one GPU accelerator for each CPU process results in idle computing
 576 resources for the remaining CPU cores. Therefore, the multi-level hybrid parallelism scheme was

577 introduced ~~to further improve the total computational performance of the CAMx-HIP~~
 578 ~~model~~ further to improve the total computational performance of the CAMx-HIP model. As
 579 described in the Sect. 3.2, ~~the horizontal advection module is accelerated by MPI and HIP~~
 580 ~~technology~~ MPI and HIP technology accelerates the horizontal advection module, and the other
 581 modules, such as the photolysis module, deposition module, chemical module, etc., which ~~runs~~ run
 582 on the CPU are accelerated by MPI and OMP under the framework of the multi-level hybrid
 583 parallelism.

584 The ZY case achieved the maximum speed-up when launching the 32 domestic CPU
 585 processors and GPU-like accelerators. ~~In the same configuration, Fig. 9 shows the total elapsed~~
 586 ~~time and speedup of CAMx-HIP model~~ Fig. 10 shows the total elapsed time and speedup of
 587 CAMx-HIP model in the same configuration when further implementing the multi-level hybrid
 588 parallelism on the "Songshan" supercomputer. The AEs of the simulation results between the
 589 CAMx-HIP model and CAMx-HIP model with the OMP technology ~~is~~ are within ± 0.04 ppbV, and
 590 the specified results are shown in Figure S1. As the number of threads increases, the elapsed time
 591 of the CAMx-HIP model is further reduced. When a CPU core ~~launching 8 threads, the one-hour~~
 592 ~~integration time in~~ launches 8 threads, the one-hour integration time in the CAMx-HIP model has
 593 been reduced from 338 seconds to 178 seconds, with a maximum acceleration of 1.9x.



594
 595 **Figure 109.** The total elapsed time and speedup of the CAMx-HIP model when implementing the multi-level
 596 hybrid parallelism in the ZY case. The unit is in seconds (s).

597 5. Conclusions and discussion

598 GPUs have become an essential part of providing processing power for high performance
599 computing applications, especially in ~~the field of~~ geoscience numerical models, ~~implementing~~.
600 Implementing super-large scale parallel computing of numerical models on GPUs has become one
601 of the significant directions of its future development. ~~In this study, the ROCm HIP technology~~
602 ~~was implemented~~ This study implemented the ROCm HIP technology to port the GPU-
603 HADVPPM from the NVIDIA GPUs to China's domestically GPU-like accelerators, ~~and further~~
604 ~~introduced~~. Further, it introduced the multi-level hybrid parallelism scheme to improve the total
605 computational performance of the CAMx-HIP model on the China's domestically heterogeneous
606 cluster.

607 The consistency of model simulation results is a significant prerequisite for heterogeneous
608 porting, ~~although~~. However, the experimental results show that the deviation between the CUDA
609 version and the Fortran version of the CAMx model, and the deviation between the HIP version
610 and the Fortran version of the CAMx model, are within the acceptable ~~range~~, the simulation
611 difference between the HIP version of CAMx model and Fortran version of CAMx model is
612 ~~smaller~~ more minor. Moreover, the BJ, HN, and ZY test cases can achieve 8.5x, 11.5x, and 28.9x
613 speedup, respectively, when the HADVPPM program is ported from the domestic CPU processor
614 A to the domestic GPU-like accelerator A. The experimental results of different cases show that
615 the larger the computing scale, the ~~more obvious the acceleration effect of the GPU-HADVPPM~~
616 ~~program, indicating that GPU is more suitable for super-large scale parallel computing~~, obvious
617 more pronounced the acceleration effect of the GPU-HADVPPM program, indicating that GPU is
618 more suitable for super-large scale parallel computing and provides technical support for accurate
619 and fast simulation of ultra-high-resolution air quality at the meter level in the future. The data
620 transfer bandwidth between CPU and GPU is one of the most important factors affecting the
621 computational efficiency of numerical model in heterogeneous clusters, as shown by the fact that
622 the elapsed time of GPU-HADVPPM program on GPU only accounts for 7.3% and 23.8% when
623 considering the data transfer time between CPU and GPU on the the "Songshan" supercomputer
624 and "Taiyuan" computing platform. Therefore, optimizing the data transfer efficiency between
625 CPU and GPU is one of the important directions for the porting and adaptation of geoscience

626 numerical models on heterogeneous clusters in the future.

627 There is still potential to further improve the computational efficiency of the CAMx-HIP
628 model in the ~~future~~^{further}. First, improve the data transfer efficiency of GPU-HADVPPM
629 between the CPU and the GPU and reduce the data transfer time. Secondly, increase the
630 proportion of HIP C code in CAMx-HIP model on the domestic GPU-like accelerator, and port
631 other modules of CAMx-HIP model to the domestic GPU-like accelerator for computing. Finally,
632 the data type of some variables could be changed from double precision to single precision, and
633 the mixing-precision method is used to further improve the CAMx-HIP computing performance.

634

635

636 *Code and data availability.* The source codes of CAMx version 6.10 are available at [https://camx-](https://camx-wp.azurewebsites.net/download/source/)
637 [wp.azurewebsites.net/download/source/](https://camx-wp.azurewebsites.net/download/source/) (ENVIRON, 2023). ~~The datasets, the CAMx-HIP codes,~~
638 ~~as well as the offline test code related to this paper are available online via ZENODO~~ ~~The datasets~~
639 ~~related to this paper and the CAMx-HIP codes are available online via ZENODO~~
640 (<https://zenodo.org/doi/10.5281/zenodo.10158214>~~https://doi.org/10.5281/zenodo.10158214~~), and
641 the CAMx-CUDA code is available online via ZENODO (<https://doi.org/10.5281/zenodo.7765218>,
642 Cao et al., 2023).

643

644 *Author contributions.* KC and QW conducted the simulation and prepared the materials. QW, LiW,₂
645 and LaW planned and organized the project. KC, QW, HG, HW, XT,₂ and LL refactored and
646 optimized the codes. LiW, NW, HC, DXL,₂ and DQL collected and prepared the data for the
647 simulation. KC, HW, QW, and HG validated and discussed the model results. KC, QW, LiW, NW,
648 XT, HG, and LaW took part in the discussion.

649

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

651

652 *Acknowledgements.* The National Key R&D Program of China (grant no. 2020YFA0607804), the
653 National Supercomputing Center in Zhengzhou Innovation Ecosystem Construction Technology
654 Special Program (grant no. 201400210700), GHfund A (grant no. 202302017828), and the Beijing

655 Advanced Innovation Program for Land Surface funded this work. The authors would like to
656 thank the High Performance Scientific Computing Center (HSCC) of Beijing Normal University
657 for providing some high-performance computing environment and technical support.

658

659 **Reference**

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

663 AMD: ROCm Documentation Release 5.7.1,
664 https://rocm.docs.amd.com/_/downloads/en/latest/pdf/ (last access: 20 October 2023), 2023.

665 Bott, A.: A Positive Definite Advection Scheme Obtained by Nonlinear Renormalization of the
666 Advective Fluxes, Monthly Weather Review - MON WEATHER REV, 117, 10.1175/1520-
667 0493(1989)117<1006:APDASO>2.0.CO;2, 1989.

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

670 Cao, K., Wu, Q., Wang, L., Wang, N., Cheng, H., Tang, X., Li, D., and Wang, L.: GPU-
671 HADVPPM V1.0: a high-efficiency parallel GPU design of the piecewise parabolic method
672 (PPM) for horizontal advection in an air quality model (CAMx V6.10), Geosci. Model Dev.,
673 16, 4367-4383, 10.5194/gmd-16-4367-2023, 2023.

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

678 Colella, P. and Woodward, P. R.: The Piecewise Parabolic Method (PPM) for gas-dynamical
679 simulations, Journal of Computational Physics, 54, 174-201, [https://doi.org/10.1016/0021-
680 9991\(84\)90143-8](https://doi.org/10.1016/0021-9991(84)90143-8), 1984.

681 ENVIRON: User Guide for Comprehensive Air Quality Model with Extensions Version 6.1,
682 https://camx-wp.azurewebsites.net/Files/CAMxUsersGuide_v6.10.pdf (last access: 20

683 October 2023), 2014.

684 ENVIRON: CAMx version 6.1, ENVIRON [code], available at: [https://camx-](https://camx-wp.azurewebsites.net/download/source/)

685 [wp.azurewebsites.net/download/source/](https://camx-wp.azurewebsites.net/download/source/), last access: 20 October 2023.

686 Huang, M., Huang, B., Mielikainen, J., Huang, H. L. A., Goldberg, M. D., and Mehta, A.: Further

687 Improvement on GPUBased Parallel Implementation of WRF 5-Layer Thermal Diffusion

688 Scheme, in: 2013 International Conference on Parallel and Distributed Systems, Seoul, South

689 Korea, 15–18 December 013, <https://doi.org/10.1109/icpads.2013.126>, 2013.

690 Linford, J. C., Michalakes, J., Vachharajani, M., and Sandu, A.: Automatic Generation of

691 Multicore Chemical Kernels, IEEE Transactions on Parallel and Distributed Systems, 22,

692 119-131, 10.1109/tpds.2010.106, 2011.

693 Mielikainen, J., Huang, B., Huang, H.-L. A., and Goldberg, M. D.: GPU Implementation of Stony

694 Brook University 5-Class Cloud Microphysics Scheme in the WRF, IEEE Journal of Selected

695 Topics in Applied Earth Observations and Remote Sensing, 5, 625-633,

696 10.1109/jstars.2011.2175707, 2012.

697 Mielikainen, J., Huang, B., Wang, J., Allen Huang, H. L., and Goldberg, M. D.: Compute unified

698 device architecture (CUDA)-based parallelization of WRF Kessler cloud microphysics

699 scheme, Computers & Geosciences, 52, 292-299, 10.1016/j.cageo.2012.10.006, 2013.

700 News, Frontier Remains as Sole Exaflop Machine and Retains Top Spot, Improving Upon Its

701 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/)

702 [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:

703 20 October 2023), 2023.

704 NVIDIA: CUDA C++ Programming Guide Version 10.2,

705 https://docs.nvidia.com/cuda/archive/10.2/pdf/CUDA_C_Programming_Guide.pdf (last

706 access: 20 October 2023), 2020.

707 Odman, M. and Ingram, C.: Multiscale Air Quality Simulation Platform (MAQSIP): Source Code

708 Documentation and Validation, 1996.

709 ROCm, AMD ROCm-HIP documentation, available at: <https://rocm.docs.amd.com/en/docs-5.0.0>

710 (last access: 20 October 2023), 2023.

711 Sun, J., Fu, J. S., Drake, J. B., Zhu, Q., Haidar, A., Gates, M., Tomov, S., and Dongarra, J.:

712 Computational Benefit of GPU Optimization for the Atmospheric Chemistry Modeling,
713 Journal of Advances in Modeling Earth Systems, 10, 1952-1969,
714 <https://doi.org/10.1029/2018MS001276>, 2018.

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

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

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

724 Wang, P., Jiang, J., Lin, P., Ding, M., Wei, J., Zhang, F., Zhao, L., Li, Y., Yu, Z., Zheng, W., Yu, Y.,
725 Chi, X., and Liu, H.: The GPU version of LASG/IAP Climate System Ocean Model version 3
726 (LICOM3) under the heterogeneous-compute interface for portability (HIP) framework and
727 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),
728 2021.