

1 **GPU-HADVPPM V1.0: A high-efficiency parallel GPU design of**
2 **the piecewise parabolic method (PPM) for horizontal advection**
3 **in an air quality model (CAMx V6.10)**

4 **Kai Cao¹, Qizhong Wu¹, Lingling Wang², Nan Wang², Huaqiong Cheng¹, Xiao**
5 **Tang³, Dongqing Li¹, and Lanning Wang¹**

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

8 ²Henan Ecological Environment Monitoring and Safety Center, Henan Key
9 Laboratory of Environmental Monitoring Technology, Zhengzhou 450000, China

10 ³State Key Laboratory of Atmospheric Boundary Layer Physics and Atmospheric
11 Chemistry, Institute of Atmospheric Physics, Chinese Academy of Science, Beijing
12 100029, China

13
14 **Correspondence to:** Qizhong Wu (wqizhong@bnu.edu.cn); Lingling Wang
15 (928216422@qq.com); Lanning Wang (wangln@bnu.edu.cn)

16
17 **Abstract.** With semiconductor technology gradually approaching its physical and
18 thermal limits, graphics processing units (GPUs) are becoming an attractive solution
19 for many scientific applications due to their high performance. This paper presents an
20 application of GPU accelerators in an air quality model. We demonstrate an approach
21 that runs a PPM solver of horizontal advection (HADVPPM) for the air quality model
22 CAMx on GPU clusters. Specifically, we first convert the HADVPPM to a new
23 Compute Unified Device Architecture C (CUDA C) code to make it computable on the
24 GPU (GPU-HADVPPM). Then, a series of optimization measures are taken, including
25 reducing the CPU-GPU communication frequency, increasing the data size
26 computation on the GPU, optimizing the GPU memory access and using thread and
27 block indices to improve the overall computing performance of the CAMx model
28 coupled with GPU-HADVPPM (named the CAMx-CUDA model). Finally, a
29 heterogeneous, hybrid programming paradigm is presented and utilized with the GPU-
30 HADVPPM on the GPU clusters with a message passing interface (MPI) and CUDA.
31 The offline experimental results show that running GPU-HADVPPM on one NVIDIA
32 Tesla K40m and an NVIDIA Tesla V100 GPU can achieve up to a 845.4x and 1113.6x

33 acceleration. By implementing a series of optimization schemes, the CAMx-CUDA
34 model results in a 29.0x and 128.4x improvement in computational efficiency by using
35 a GPU accelerator card on a K40m and V100 cluster, respectively. In terms of the
36 single-module computational efficiency of GPU-HADVPPM, it can achieve 1.3x and
37 18.8x speedup on an NVIDIA Tesla K40m GPU and NVIDIA Tesla V100 GPU,
38 respectively. The multi-GPU acceleration algorithm enables a 4.5x speedup with 8 CPU
39 cores and 8 GPU accelerators on a V100 cluster.

40 **1. Introduction**

41 Since the introduction of personal computers in the late 1980s, the computer and
42 mobile device industry has created a flourishing worldwide market (Bleichrodt et al.,
43 2012). In recent years, improvements of the central processing unit (CPU) performance
44 have been limited by its heat dissipation, and the applicability of Moore's Law has
45 flattened. A common trend in high-performance computing today is the utilization of
46 hardware accelerators, which execute codes rich in data parallelism, to form high-
47 performance heterogeneous systems. GPUs are widely used as accelerators due to their
48 high peak performances. In the top ten supercomputing list released in December 2022
49 (<https://www.top500.org/lists/top500/list/2022/11/>, last access: 19 December 2022),
50 there were seven heterogeneous supercomputing platforms built with CPU processors
51 and GPU accelerators, of which the top one, Frontier at the Oak Ridge National
52 Laboratory, uses AMD's third-generation EPYC CPU and AMD's Instinct MI250X
53 GPU, and its computing performance reaches exascale levels (10^{18} calculations per
54 second) for the first time ([https://www.amd.com/en/press-releases/2022-05-30-world-
55 s-first-exascale-supercomputer-powered-amd-epyc-processors-and-amd](https://www.amd.com/en/press-releases/2022-05-30-world-s-first-exascale-supercomputer-powered-amd-epyc-processors-and-amd), last access:
56 19 December 2022). Such a powerful computing performance of the heterogeneous
57 system not only injects new vitality into high-performance computing but also
58 generates new solutions for improving the performance of geoscience numerical
59 models.

60 The GPU has proven successful in weather models such as the nonhydrostatic
61 icosahedral model (NIM; Govett et al., 2017), global/regional assimilation and
62 prediction system (GRAPES; Xiao et al., 2022), weather research and forecasting
63 model (WRF; Huang et al., 2011; Huang et al., 2012; Mielikainen et al., 2012a;
64 Mielikainen et al., 2012b; Mielikainen et al., 2013a; Mielikainen et al., 2013b; Price et
65 al., 2014; Huang et al., 2015), ocean models such as the LASG/IAP climate system
66 ocean model (LICOM; Jiang et al., 2019; Wang et al., 2021a) and Princeton ocean
67 model (POM; Xu et al., 2015) and earth system model of the Chinese Academy of
68 Sciences (CAS-ESM; Wang et al., 2016; Wang et al., 2021b).

69 Govett et al. (2017) used open accelerator (OpenACC) directives to port the
70 dynamics of NIM to the GPU and achieved a 2.5x acceleration. Additionally, using
71 OpenACC directives, Xiao et al. (2022) ported the PRM (piecewise rational method)
72 scalar advection scheme in GRAPES to the GPU, achieving up to 3.51x faster results
73 than 32 CPU cores. In terms of the most widely used WRF, several parameterization
74 schemes, such as the RRTMG_LW scheme (Price et al., 2014), 5-layer thermal
75 diffusion scheme (Huang et al., 2015), Eta Ferrier cloud microphysics scheme (Huang
76 et al., 2012), Goddard shortwave scheme (Mielikainen et al., 2012a), Kessler cloud
77 microphysics scheme (Mielikainen et al., 2013b), SBU-YLIN scheme (Mielikainen et
78 al., 2012b), WMS5 scheme (Huang et al., 2011) and WMS6 scheme (Mielikainen et al.,
79 2013a), have been ported heterogeneously using CUDA C and achieved 37x~896x
80 acceleration results. LICOM has conducted heterogeneous porting using OpenACC
81 (Jiang et al., 2019) and used heterogeneous-compute interface for portability C (HIP C)
82 technologies and achieved up to a 6.6x and 42x acceleration, respectively (Wang et al.,
83 2021a). For the Princeton ocean model, Xu et al. (2015) use CUDA C to conduct
84 heterogeneous porting and optimization, and the performance of gpu-POM v1.0 on four
85 GPUs is comparable to that on the 408 standard Intel Xeon X5670 CPU cores. In terms
86 of climate system models, Wang et al. (2016) and Wang et al. (2021b) used CUDA
87 Fortran and CUDA C to conduct heterogeneous porting of the RRTMG_SW and
88 RRTMG_LW schemes of the atmospheric component model of the CAS-ESM earth

89 system model and achieved a 38.88x and 77.78x acceleration, respectively.

90 Programming a GPU accelerator can be a difficult and error-prone process that
91 requires specially designed programming methods. There are three widely used
92 methods for porting programs to GPUs, as described above. The first method uses the
93 OpenACC directive (<https://www.openacc.org/>, last access: 19 December 2022), which
94 provides a set of high-level directives that enable C/C++ and Fortran programmers to
95 utilize accelerators. The second method uses CUDA Fortran. CUDA Fortran is a
96 software compiler that was codeveloped by the Portland Group (PGI) and NVIDIA,
97 and is a tool chain for building performance-optimized GPU-accelerated Fortran
98 applications targeting the NVIDIA GPU platform ([https://developer.nvidia.com/cuda-
99 fortran](https://developer.nvidia.com/cuda-fortran), last access: 19 December 2022). Using CUDA C involves rewriting the entire
100 program using the standard C programming language and low-level CUDA subroutines
101 (<https://developer.nvidia.com/cuda-toolkit>, last access: 19 December 2022) to support
102 the NVIDIA GPU accelerator. Compared to the other two technologies, the CUDA C
103 porting scheme is the most complex but it has the highest computational performance
104 (Mielikainen et al., 2012b; Wahib and Maruyama, 2013; Xu et al., 2015).

105 Air quality models are critical for understanding how the chemistry and
106 composition of the atmosphere may change throughout the 21st century, as well as for
107 preparing adaptive responses or developing mitigation strategies. Because air quality
108 models need to take into account the complex physicochemical processes that occur in
109 the atmosphere of anthropogenic and natural emissions, simulations are
110 computationally expensive. Compared to other geoscientific numerical models, few
111 studies have conducted a heterogeneous porting of air quality models. In this study, the
112 CUDA C scheme, implemented in this paper, conducted a hotspot module porting of
113 CAMx to improve the computation efficiency.

114 2. The CAMx model and experiments

115 2.1. Model description

116 The CAMx model is a state-of-the air quality model developed by Ramboll
117 Environ (<https://www.camx.com/>, last access: 19 December 2022). CAMx version 6.10
118 (CAMx V6.10; ENVIRON, 2014) is chosen in this study; it simulates the emission,
119 dispersion, chemical reaction and removal of pollutants by marching the Eulerian
120 continuity equation forward in time for each chemical species on a system of nested
121 three-dimensional grids. The Eulerian continuity equation is expressed mathematically
122 in terrain-following height coordinates as Formula (1):

$$123 \quad \frac{\partial c_i}{\partial t} = -\nabla_H \cdot V_H c_i + \left[\frac{\partial(c_i \eta)}{\partial z} - c_i \frac{\partial^2 h}{\partial z \partial t} \right] + \nabla \cdot \rho K \nabla (c_i / \rho)$$
$$124 \quad + \left. \frac{\partial c_i}{\partial t} \right|_{Emission} + \left. \frac{\partial c_i}{\partial t} \right|_{Chemistry} + \left. \frac{\partial c_i}{\partial t} \right|_{Removal} \quad (1)$$

$$125 \quad \nabla_H \cdot \rho V_H = \frac{m^2}{A_{yz}} \frac{\partial}{\partial x} \left(\frac{u A_{yz} \rho}{m} \right) + \frac{m^2}{A_{xz}} \frac{\partial}{\partial y} \left(\frac{v A_{xz} \rho}{m} \right) \quad (2)$$

126 The first term on the right-hand side represents horizontal advection. In numerical
127 methods, the horizontal advection equation (described in Formula (2)) is performed
128 using the area preserving flux-form advection solver of the piecewise parabolic method
129 (PPM) of Colella and Woodward (1984) as implemented by Odman and Ingram (1996).
130 The PPM horizontal advection solution (HADVPPM) was incorporated into the CAMx
131 model because it provides higher order accuracy with minimal numerical diffusion.

132 In the Fortran code implementation of the HADVPPM scheme, the CAMx main
133 program calls the emistrns program, which mainly performs physical processes such as
134 emission, diffusion, advection and dry/wet deposition of pollutants. Then, the
135 horizontal advection program is invoked by the emistrns program to solve the
136 horizontal advection equation by using the HADVPPM scheme.

137 2.2. Benchmark performance experiments


138 The first porting step is to test the performance of the CAMx benchmark version
139 and identify the model's hotspots. On the Intel x86 CPU platform, we launch two
140 processes concurrently to run the CAMx, and take advantage of the Intel trace analyser
141 collector (ITAC; [https://www.intel.com/content/www/us/en/docs/trace-analyzer-
142 collector/get-started-guide/2021-4/overview.html](https://www.intel.com/content/www/us/en/docs/trace-analyzer-collector/get-started-guide/2021-4/overview.html), last access: 19 December 2022) and
143 the Intel VTune profiler
144 (VTune;[https://www.intel.com/content/www/us/en/develop/documentation/vtune-
145 help/top.html](https://www.intel.com/content/www/us/en/develop/documentation/vtune-help/top.html), last access: 19 December 2022) performance analysis tools to collect
146 performance information during the CAMx operation.

147 The general MPI performance can be reported by the ITAC tool, and MPI load
148 balance information, computation and communication profiling of each process is
149 shown in Fig. 1a. During the running process of the CAMx model, Process 0 (P0)
150 spends 99.6% of the time on the MPI_Barrier function and only 0.4% of the time on
151 computation, while Process 1 (P1) spends 99.8% of its time computation and only 0.2%
152 of its time receiving messages from P0. It is apparent that the parallel design of the
153 CAMx model adopts the Master-Slave mode, and P0 is responsible for inputting and
154 outputting the data and calling the MPI_Barrier function to synchronize the process, so
155 there is a lot of MPI waiting time. The other processes are responsible for computation.

156 The VTune tool detects each module's runtime and the most time-consuming
157 functions on P1. As shown in Figure 1b, the top four time-consuming modules are
158 chemistry, diffusion, horizontal advection and vertical advection in the CAMx model.
159 In the above four modules, the top five most time-consuming programs are the ebrate,
160 hadvppm, tridiag, diffus and ebisolv programs, and the total runtime of P1 is 325.1
161 seconds. Top1 and Top2's most time-consuming programs take 49.4 and 35.6 seconds,
162 respectively.

163 By consideration, the hadvppm program was selected to conduct heterogeneous
164 porting for several reasons. First, the advection module is one of the air quality model's
165 compulsory modules, and is mainly used to simulate the transport process of air

166 pollutants, additionally it is also a hotspot module detected by the Intel VTune tool. The
 167 typical air quality models, CAMx, CMAQ and NAQPMS, include advection modules
 168 and use the exact PPM advection solver. The heterogeneous version developed in this
 169 study can be directly applied to the above models. Furthermore, the weather model (e.g.,
 170 WRF) also contains an advection module, so this study's heterogeneous porting method
 171 and experience can be used for reference. Therefore, a GPU acceleration version of the
 172 HADVPPM scheme, namely, GPU-HADVPPM, is built to improve the CAMx
 173 performance.



174

175 **Figure 1.** The computation performance of the modules in the CAMx model. (a) Computation and
 176 communication profiling of P0 and P1. (b) Overhead proportions of P1. The top four most time-
 177 consuming modules are chemistry, diffusion, horizontal advection and vertical advection.

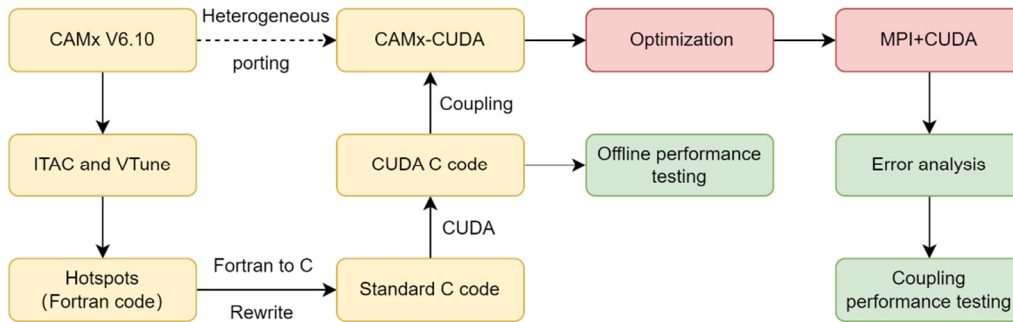
178

179 2.3. Porting scheme introduction

180 The CAMx-CUDA heterogeneous scheme is shown in Figure 2. The second time-
 181 consuming hadvppm program in the CAMx model was selected to implement
 182 heterogeneous porting. To map the hadvppm program to the GPU, the Fortran code was
 183 converted to standard C code. Then, the CUDA programming language, which was
 184 tailor-made for NVIDIA, was added to convert the standard C code into CUDA C for

185 data-parallel execution on the GPU, as GPU-HADVPPM. It prepared the input data for
 186 GPU-HADVPPM by constructing random numbers, and tested its offline performance
 187 on the GPU platform.

188 After coupling the GPU-HADVPPM to the CAMx model, the advection module
 189 code was optimized according to the characteristics of the GPU architecture to improve
 190 the overall computational efficiency on the CPU-GPU heterogeneous platform. Then,
 191 the multi-CPU core and multi-GPU card acceleration algorithm was adopted to improve
 192 the parallel extensibility of heterogeneous computing. Finally, the coupling
 193 performance test is implemented after verifying the different CAMx model simulation
 194 results.



195
 196 **Figure 2.** Heterogeneous porting scheme of the CAMx-CUDA model.

197 2.4. Hardware components and software environment of the testing system

198 The experiments are conducted on two GPU clusters, K40m and V100. The
 199 hardware components and software environment of the two clusters are listed in Table
 200 1. The K40m cluster is equipped with two 2.5 GHz 16-core Intel Xeon E5-2682 v4
 201 CPU processors and one NVIDIA Tesla K40m GPU card on each node. The NVIDIA
 202 Tesla K40m GPU has 2880 CUDA cores with 12 GB of memory. The V100 cluster
 203 contains two 2.7 GHz 24-core Intel Xeon Platinum 8168 processors and eight NVIDIA
 204 Tesla V100 GPU cards with 5120 CUDA cores and 16 GB memory on each card.

205 **Table 1.** Configurations of GPU cluster.

Hardware components	
CPU	GPU

K40m cluster	Intel Xeon E5-2682 v4 CPU @2.5 GHz, 16 cores	NVIDIA Tesla K40m, 2880 CUDA cores, 12GB memory
V100 cluster	Intel Xeon Platinum 8168 CPU @2.7 GHz, 24 cores	NVIDIA Tesla V100, 5120 CUDA cores, 16GB memory
Software environment		
	Compiler and MPI	Programming Model
K40m cluster	Intel-2021.4.0	CUDA-10.2
V100 cluster	Intel-2019.1.144	CUDA-10.0

206 For Fortran and standard C programming, Intel Toolkit (including compiler and
 207 MPI library) version 2021.4.0 and version 2019.1.144 are employed for compiling on
 208 an Intel Xeon E4-2682 v4 CPU and Intel Xeon Platinum 8168 CPU, respectively. Then,
 209 CUDA version 10.2 and version 10.0 are employed on an NVIDIA Tesla K40m GPU
 210 and NVIDIA Tesla V100 GPU. CUDA (NVIDIA, 2020) is an extension of the C
 211 programming language that offers direct programming of the GPUs. In CUDA
 212 programming, a kernel is actually a subroutine that can be executed on the GPU. The
 213 underlying code in the kernel is divided into a series of threads, each with a unique "ID"
 214 number that simultaneously process different data through a single-instruction multiple-
 215 thread (SIMT) parallel mode. These threads are grouped into equal-sized thread blocks,
 216 which are organized into a grid.

217 **3. Porting and optimization of the CAMx advection module on a heterogeneous** 218 **platform**

219 **3.1. Mapping the HADVPPM scheme to the GPU**

220 **3.1.1. Manual code translation from Fortran to standard C**

221 As the CAMx V6.10 code was written in Fortran 90, we rewrote the hadvppm
 222 program from Fortran to CUDA C. As an intermediate conversion step, we refactor the
 223 original Fortran code using standard C. During the refactoring, some of the
 224 considerations are listed in Table 2:

225 (1) The subroutine name refactored with standard C must be followed by an

226 underscore identifier, which can only be recognized when Fortran calls.

227 (2) In the Fortran language, the parameters are transferred by a memory address
228 by default. In the case of mixed programming in Fortran and standard C, the parameters
229 transferred by Fortran are processed by the pointer in standard C.

230 (3) Variable precision types defined in standard C must be strictly consistent with
231 those in Fortran.

232 (4) Some built-in functions in Fortran are not available in standard C, and need to
233 be defined in the standard C macro definitions.

234 (5) For multidimensional arrays, Fortran and standard C follow a column-major
235 and row-major order, and in-memory read and write, respectively;

236 (6) Array subscripts in Fortran and standard C are indexed from any integer and 0,
237 respectively.

238 **Table 2.** Some considerations during Fortran to C refactoring.

	Fortran code	C code
Function name	<i>subroutine hadvppm()</i>	<i>void hadvppm()</i>
Parameter passing	<i>hadvppm(nn, dt,dx, con,vel, area,areav, flxarr,mynn)</i>	<i>hadvppm(int *nn, float *dt, float *dx, float *con, float *vel, float *area, float *areav, float *flxarr, int *mynn)</i>
Variable precision	<i>real(kind=8) x</i>	<i>double x</i>
Built-in functions	<i>max</i>	<i>#define Max(a, b) ((a)>(b)?(a):(b))</i>
Memory read and write for multidimensional array	Column-major	Row-major
Array subscript index	Starting from any integer	Starting from 0

239

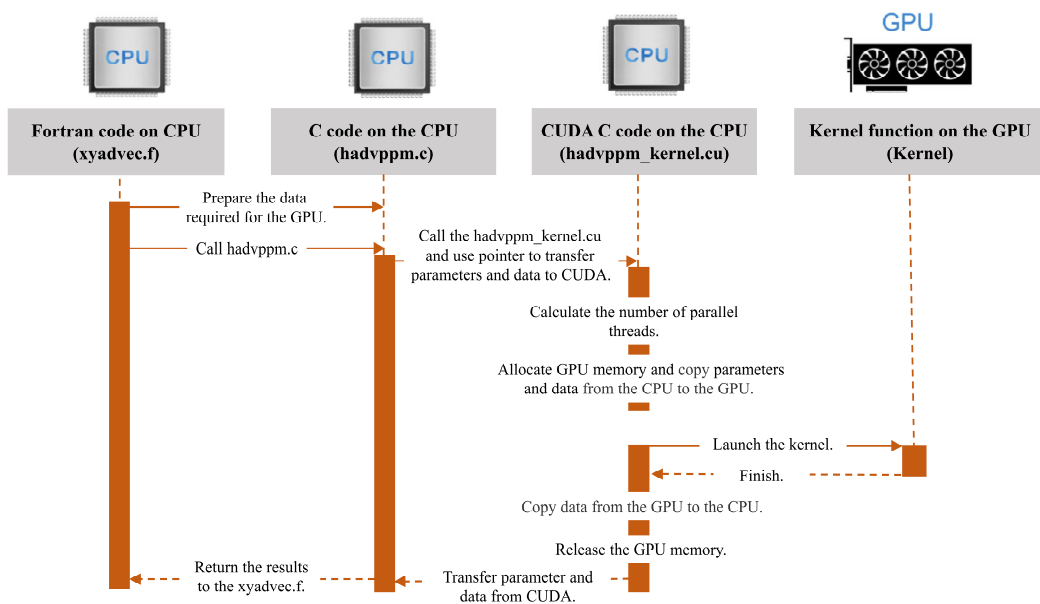
240 3.1.2. Converting standard C code into CUDA C

241 After refactoring the Fortran code of the hadvppm program with standard C,

242 CUDA was used to convert the C code into CUDA C to make it computable on the
 243 GPU. A standard C program using CUDA extensions distributes a large number of
 244 copies of the kernel functions into available multiprocessors and executes them
 245 simultaneously on the GPU.

246 Figure 3 shows the GPU-HADVPPM implementation process. As mentioned in
 247 Sect. 2.1, the xyadvec program calls the hadvppm program to solve the horizontal
 248 advection function. Since the rewritten CUDA program cannot be called directly by the
 249 Fortran program (xyadvec.f), we add an intermediate subroutine (hadvppm.c) as an
 250 interface to transfer the parameters and data required for GPU computing from the
 251 xyadvec Fortran program to the hadvppm_kernel CUDA C program.

252 A CUDA program automatically uses numerous threads on the GPU to execute
 253 kernel functions. Therefore, the hadvppm_kernel CUDA C program first calculates the
 254 number of parallel threads according to the array dimension. Then, the GPU memory
 255 is allocated, and the parameters and data are copied from the CPU to the GPU. As the
 256 CUDA program launches a large number of parallel threads to execute kernel functions
 257 simultaneously, the computation results will be copied from the GPU back to the CPU.
 258 Finally, the GPU memory is released, and the data computed on the GPU are returned
 259 to the xyadvec program via the hadvppm C program.



260

261 **Figure 3.** The calling and computation process of the GPU-HADVPPM on the CPU-GPU

262 heterogeneous platform.

263 **3.2. Coupling and optimization of the GPU-HADVPPM scheme on a single GPU**

264 After the hadvppm program was rewritten with standard C and CUDA, the
265 implementation process of the HADVPPM scheme was loaded from the CPU to the
266 GPU. Then, we coupled the GPU-HADVPPM to the CAMx model. For ease of
267 description, we will refer to this original heterogeneous version of CAMx as CAMx-
268 CUDA V1.0. In CAMx-CUDA V1.0, four external loops are nested when the hadvppm
269 C program is called by the xyadvec program. This will result in widespread data
270 transfers from the CPU to the GPU over the PCIe bus within a time step, making the
271 computation of CAMx-CUDA V1.0 inefficient.

272 Therefore, we optimize the xyadvec Fortran program to significantly reduce the
273 frequency of data transmission between the CPU and GPU, increase the amount of data
274 computation on the GPU, and improve the total computing efficiency of the CAMx on
275 the CPU-GPU heterogeneous platforms. In the original CAMx-CUDA V1.0, four
276 external loops outside the hadvppm C program, and several one-dimensional arrays, are
277 computed before calling the hadvppm C program. Then, the CPU will frequently launch
278 the GPU and transfer data to it within a time step. When the code optimization is
279 completed, the three- or four-dimensional arrays required for a GPU computation
280 within a time step will be sorted before calling the hadvppm C program, and then the
281 CPU will package and transfer the arrays to the GPU in batches. An example of the
282 xyadvec Fortran program optimization is shown in Figure S1.

283 The details of the four different versions are shown in Table 3. In CAMx-CUDA
284 V1.0, the Fortran code of the HADVPPM scheme was rewritten using standard C and
285 CUDA, and the xyadvec program was not optimized. The dimensions of the cld
286 variable array transmitted to the GPU in the X and Y directions are 157 and 145 in this
287 case, respectively. In CAMx-CUDA V1.1 and CAMx-CUDA V1.2, the cld variable
288 transmitted from the CPU to GPU is expanded to two (approximately 23,000 numbers)
289 and four dimensions (approximately 27.4 million numbers) by optimizing the xyadvec

290 Fortran program and hadvppm_kernel CUDA C program, respectively.

291 The order in which the data are accessed in GPU memory affects the
 292 computational efficiency of the code. In CAMx-CUDA V1.3 of Table 4, we further
 293 optimized the order in which the data are accessed in GPU memory based on the order
 294 in which they are stored in memory and eliminated the unnecessary assignment loops
 295 that were added due to the difference in memory read order between Fortran and C.

296 As described in Sect. 2.4, a thread is the basic unit of parallelism in CUDA
 297 programming. The thread structure is organized into a three-level hierarchy. The highest
 298 level is a grid, which consists of three-dimensional thread blocks. The second level is a
 299 block, which also consists of three-dimensional threads. The built-in CUDA variable
 300 *threadIdx.x* determines a unique thread "ID" number inside a thread block. Similarly,
 301 the built-in variables *blockIdx.x* and *blockIdx.y* determine which block to execute on,
 302 and the size of the block is determined by using the built-in variable *blockDim.x*. For
 303 the two-dimensional horizontal grid points, many threads and blocks can be organized
 304 so that each CUDA thread computes the results for different spatial positions
 305 simultaneously.

306 Before CAMx-CUDA V1.4, the loops for the three-dimensional spatial grid points
 307 (i,j,k) are replaced by index computations using only the thread index ($i = threadIdx.x$
 308 $+ blockIdx.x * blockDim.x$) to use the thread indices to simultaneously compute the grid
 309 point in the x or y direction. To take full advantage of the thousands of threads in the
 310 GPU, we implement thread and block indices ($i = threadIdx.x + blockIdx.x * blockDim.x$;
 311 $j = blockIdx.y$) to simultaneously compute all the horizontal grid points (i,j) in CAMx-
 312 CUDA V1.4. This is permitted because there are no interactions among the horizontal
 313 grid points.

314 **Table 3.** The details of different CAMx-CUDA versions during optimization.

Version	Major revisions	Amount of data computation on GPU
CAMx-CUDA V1.0	The Fortran code of the HADVPPM subroutine was rewritten using standard C and CUDA, and <i>xyadvec.f</i> was not optimized.	157 and 145 in the x direction and y direction for the <i>cld</i> variable, respectively.
CAMx-CUDA V1.1	Optimize <i>xyadvec.f</i> and	157×145,

	<i>hadvppm_kernel.cu</i> to expand the dimension of the array transmitted to the GPU from 1-dimensional to 2-dimensional.	approximately 23,000 numbers for the c2d variable.
CAMx-CUDA V1.2	Based on the CAMx-CUDA V1.1, the dimension of the array transmitted to the GPU is extended from 2 to 4 dimensions.	157×145×14×86, approximately 27.4 million numbers for the c4d variable.
CAMx-CUDA V1.3	Based on the CAMx-CUDA V1.2, the order of GPU memory access is optimized and unnecessary assignment loops are eliminated.	157×145×14×86, approximately 27.4 million numbers for the c4d variable.
CAMx-CUDA V1.4	Based on the CAMx-CUDA V1.3, using thread and block indices ($i = threadIdx.x + blockIdx.x*blockDim.x; j = blockIdx.y$).	157×145×14×86, approximately 27.4 million numbers for the c4d variable.

315

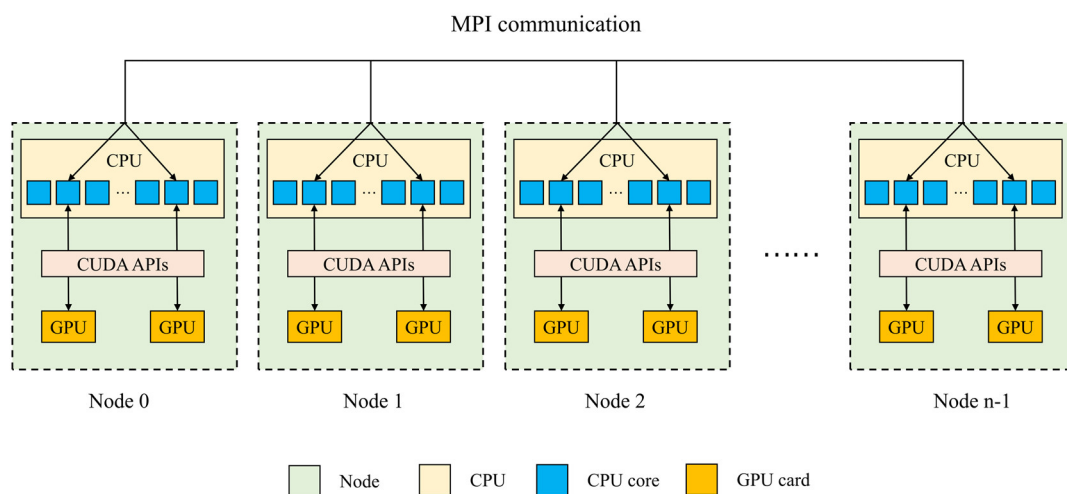
316 **3.3. MPI+CUDA acceleration algorithm of CAMx-CUDA on multiple GPUs**

317 Generally, superlarge clusters have thousands of compute nodes. The current
318 CAMx V6.10, implemented by adopting MPI communication technology, typically
319 runs on dozens of compute nodes. Once the GPU-HADVPPM is coupled into the
320 CAMx, it also has to run on multiple compute nodes that are equipped with one or more
321 GPUs on each node. To make full use of multicore and multi-GPU supercomputers, and
322 further improve the overall computational performance of CAMx-CUDA, we adopt a
323 parallel architecture with an MPI+CUDA hybrid paradigm; that is, the collaborative
324 computing strategy of multiple CPU cores and multiple GPU cards is adopted during
325 the operation of the CAMx-CUDA model. Adopting this strategy, the GPU-HADVPPM
326 can run on multiple GPUs, and the Fortran code of the other modules in the CAMx-
327 CUDA model can run on multiple CPU cores.

328 As shown in Figure 4, after the simulated region is subdivided by MPI, a CPU
329 core is responsible for the computation of a subregion. To improve the total
330 computational performance of the CAMx-CUDA model, we further used the NVIDIA
331 CUDA library to obtain the number of GPUs per node and then used the MPI process
332 ID and remainder function to determine the GPU ID to be launched by each node.

333 Finally, we used the NVIDIA CUDA library, `cudaSetDevice`, to configure a GPU card
334 for each CPU core.

335 According to the benchmark performance experiments, the parallel design of
336 CAMx adopts the master-slave mode, and P0 is responsible for inputting and outputting
337 data. If two processes (P0 and P1) were launched, only the P1 and its configured GPU
338 participate in integration.



339

340 **Figure 4.** An example of parallel architecture with an MPI+CUDA hybrid paradigm on multiple
341 GPUs.

342 4. Experimental results

343 The validation and evaluation of porting the HADVPPM scheme from the CPU to
344 the GPU platform were conducted using offline and coupling performance experiments.
345 First, we validated the results between the different CAMx versions, and then the offline
346 performance of the GPU-HADVPPM on a single GPU was tested by offline
347 experiments. Finally, coupling performance experiments illustrate its potential in three
348 dimensions with varying chemical regimes. In Sect.4.2 and Sect.4.4, the CAMx
349 versions of the HADVPPM scheme written in Fortran, standard C and CUDA C are
350 named F, C and CUDA C, respectively.

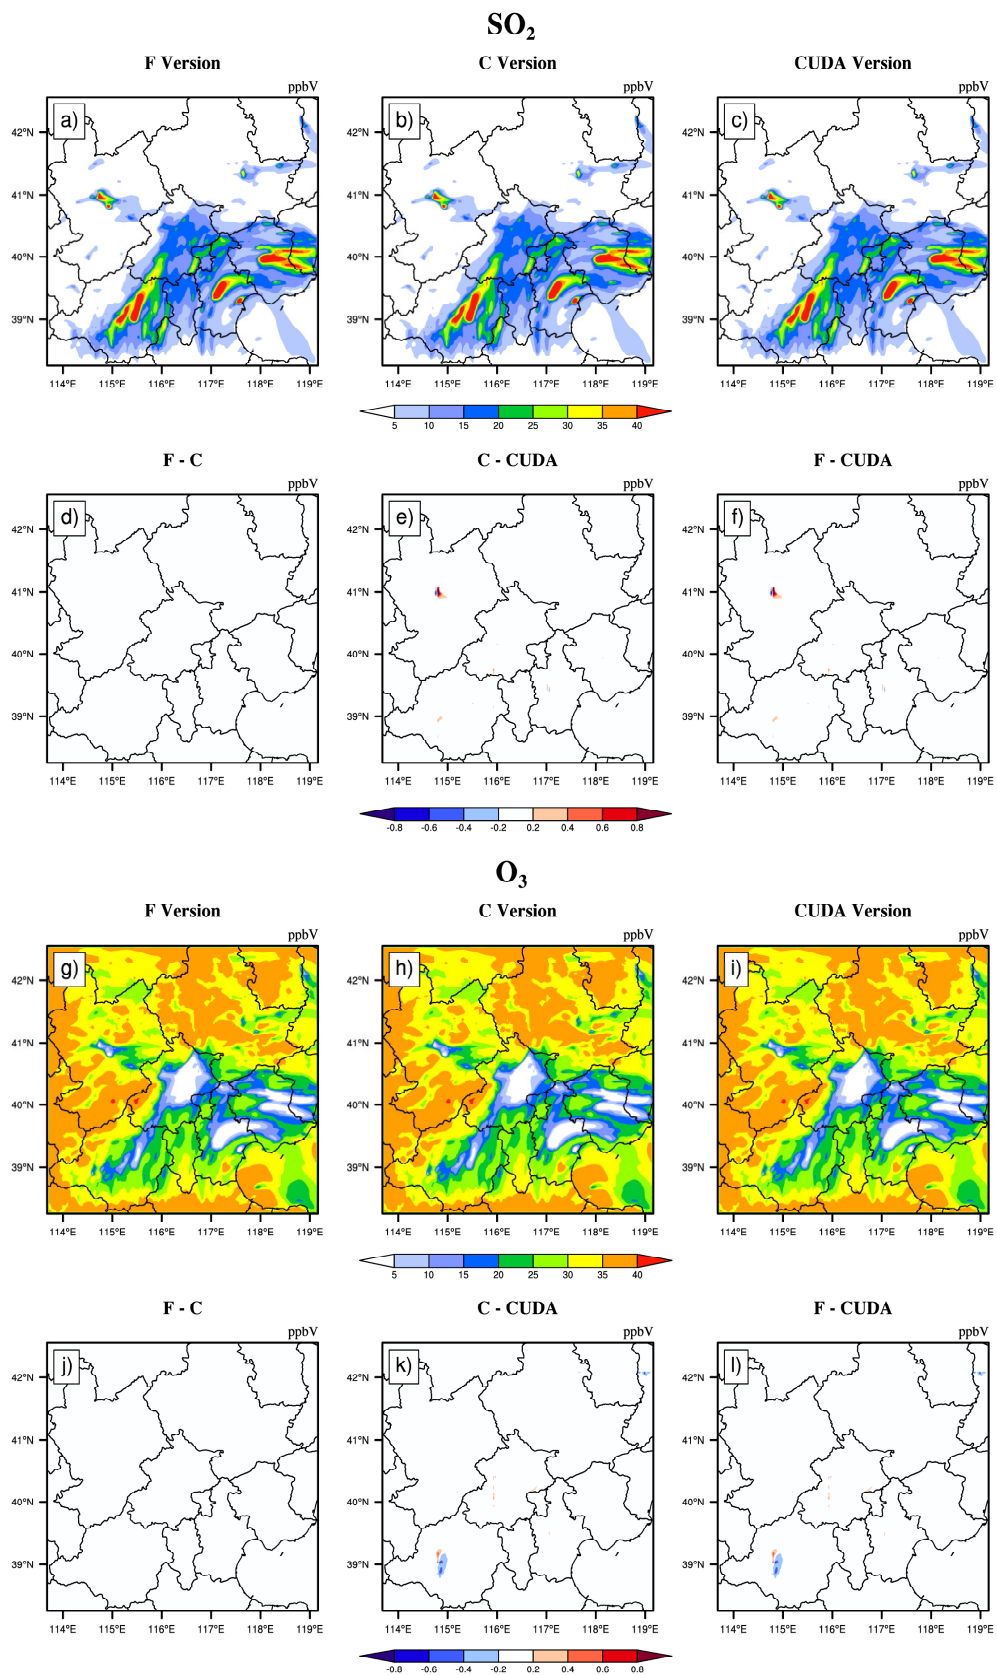
351 **4.1. Experimental setup**

352 The test case is a 48 h simulation covering Beijing, Tianjin and part of the Hebei
353 Province region. The horizontal resolution is 3 km with 145×157 grid boxes. The
354 model adopted 14 vertical layers. The simulation started at 12:00 UTC on 01 November
355 2020 and ended at 12:00 UTC on 03 November 2020. The meteorological fields driving
356 the CAMx model were provided by the weather research and forecasting (WRF;
357 Skamarock et al., 2008) model. The sparse matrix operator kernel emission (SMOKE;
358 Houyoux and Vukovich, 1999) version 2.4 model is used to provide gridded emission
359 data for the CAMx model. The emission inventories (Sun et al., 2022) include the
360 regional emissions in East Asia that were obtained from the transport and chemical
361 evolution over the Pacific (TRACE-P; Streets et al., 2003; Streets et al., 2006) project,
362 30-min (approximately 55.6 km at mid-latitude) spatial resolution Intercontinental
363 chemical transport experiment-Phase B (INTEX-B; Zhang et al., 2009) and the updated
364 regional emission inventories in North China. The physical and chemical numerical
365 methods selected during the CAMx model integration are listed in Table S2.

366 **4.2. Error analysis**

367 The hourly concentrations of different CAMx simulations (Fortran, C, and CUDA
368 C versions) are compared to verify the usefulness of the CUDA C version of CAMx for
369 numerical precision for scientific usage. Here, we chose six major species, i.e., SO₂, O₃,
370 NO₂, CO, H₂O₂ and PSO₄, after 48 h of integration to verify the results. Due to the
371 differences in programming languages and hardware, the simulation results are affected
372 during the porting process. Figures 5~7 present the spatial distributions of SO₂, O₃, NO₂,
373 CO, H₂O₂ and PSO₄, as well as the absolute errors (AEs) of their concentrations from
374 different CAMx versions. The species' spatial patterns of the three CAMx versions are
375 visually very similar. Between the Fortran and C versions, especially, the AEs in all the
376 grid boxes are in the range of ± 0.01 ppbV (the unit of PSO₄ is $\mu\text{g} \cdot \text{m}^{-3}$). During the
377 porting process, the primary error comes from converting standard C to CUDA C, and

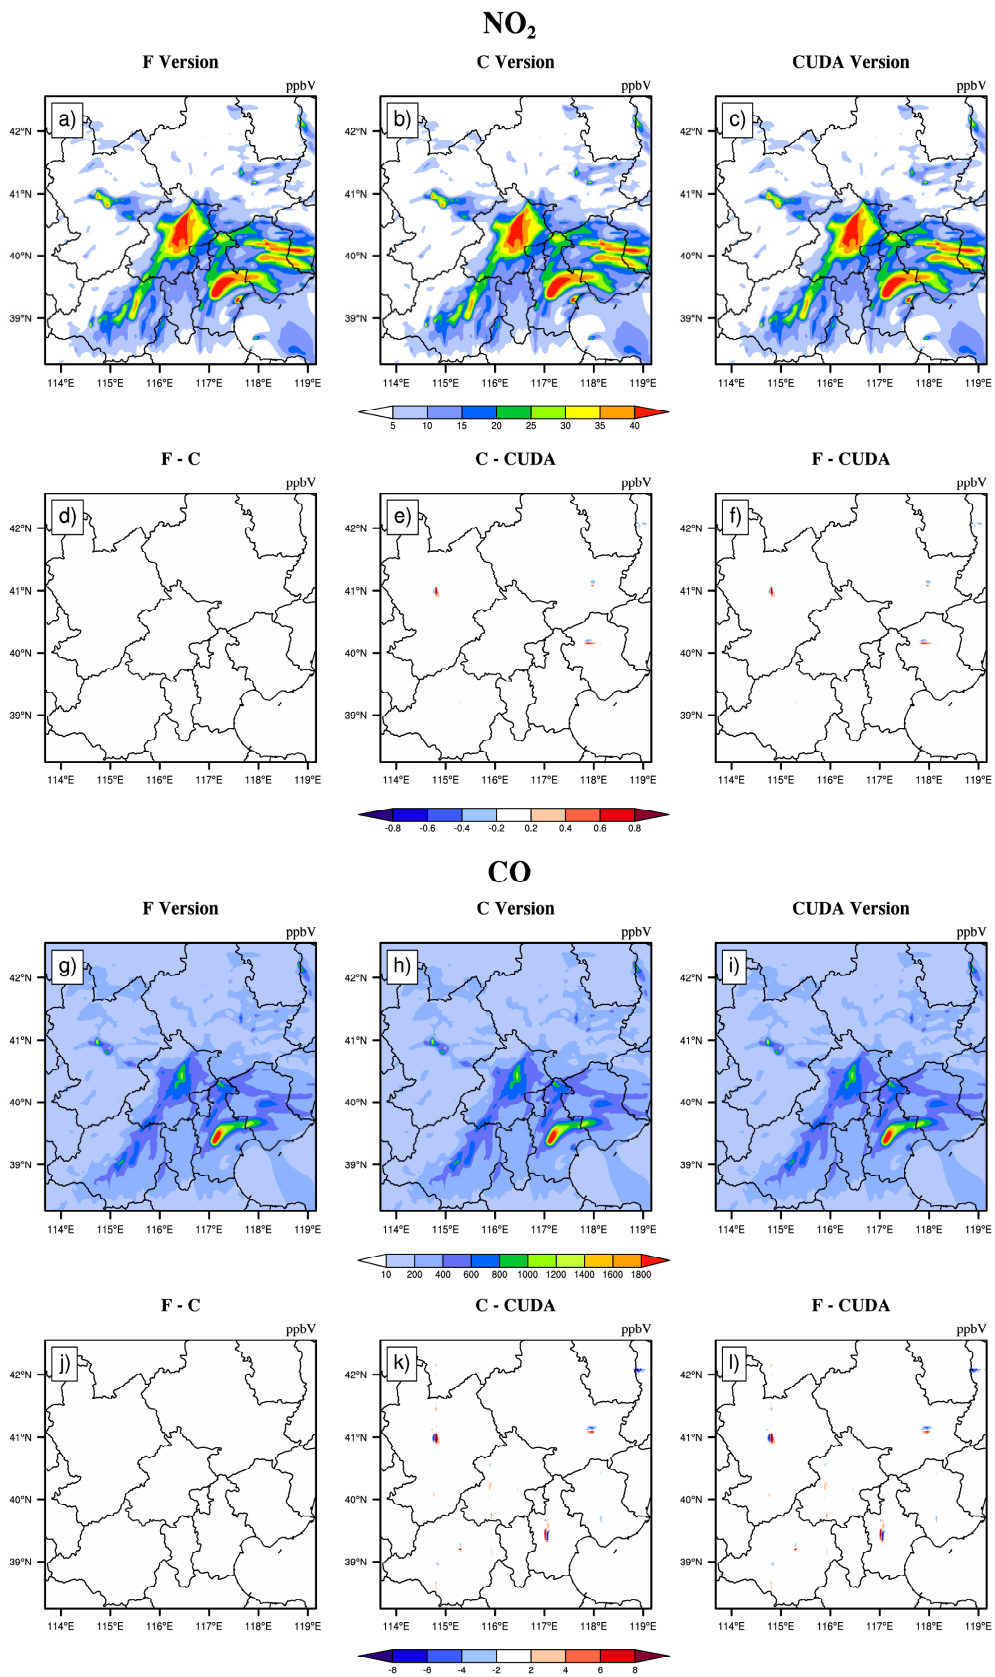
378 the main reason is related to the hardware difference between the CPU and GPU. Due
379 to the slight difference in data operation and accuracy between the CPU and GPU
380 (NVIDIA,2023), the concentration variable of the hadvppm program appears to have
381 minimal negative values (approximately $-10^{-9}\sim-10^{-4}$) when integrated on the GPU.
382 To allow the program to continue running, we forcibly replace these negative values
383 with 10^{-9} . It is because these negative values are replaced by positive values that the
384 simulation results are biased. In general, for SO_2 , O_3 , NO_2 , H_2O_2 and PSO_4 , the AEs in
385 the majority of the grid boxes are in the range of ± 0.8 ppbV or $\mu\text{g} \cdot \text{m}^{-3}$ between the
386 standard C and CUDA C versions; for CO, because its background concentration is
387 higher, the AEs of the standard C and CUDA C versions are outside that range, and fall
388 into the range of -8 and 8 ppbV in some grid boxes and shows more obvious AEs than
389 the other species.



390

391 **Figure 5.** SO₂ and O₃ concentrations outputted by the CAMx model for the Fortran, standard C, and

392 CUDA C versions. Panels (a) and (g) are from the Fortran versions. Panels (b) and (h) are from the
393 standard C versions. Panels (c) and (i) are from the CUDA C versions. Panels (d) and (j) are the
394 output concentration differences of the Fortran and standard C versions. Panels (e) and (k) are the
395 output concentration differences of the standard C and CUDA C versions. Panels (f) and (l) are the
396 output concentration differences of the Fortran and CUDA C versions.

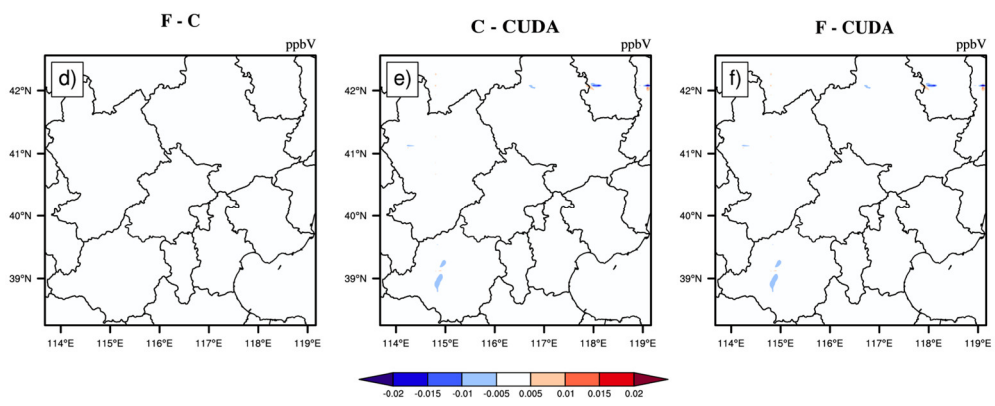
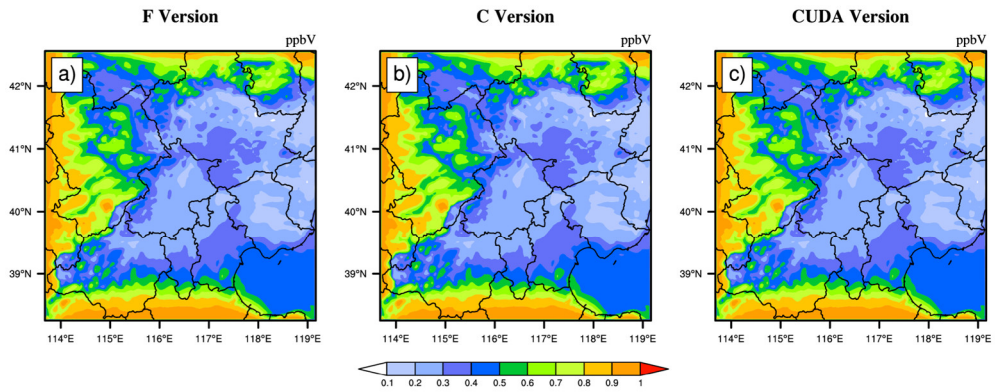


397

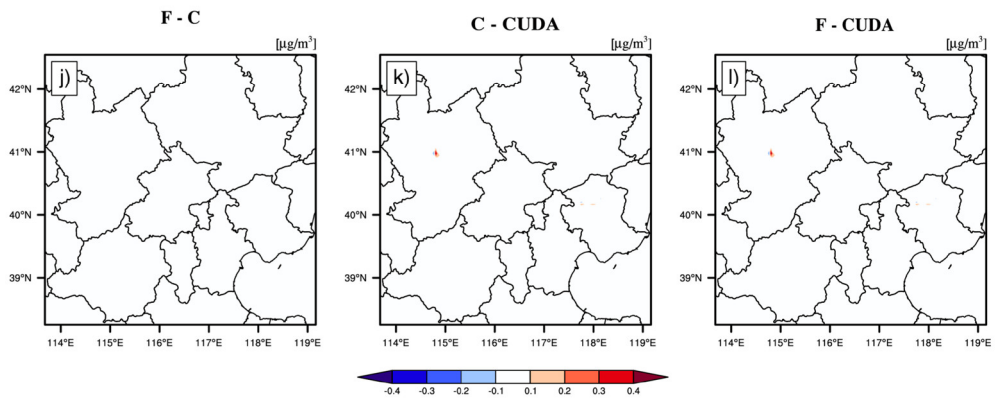
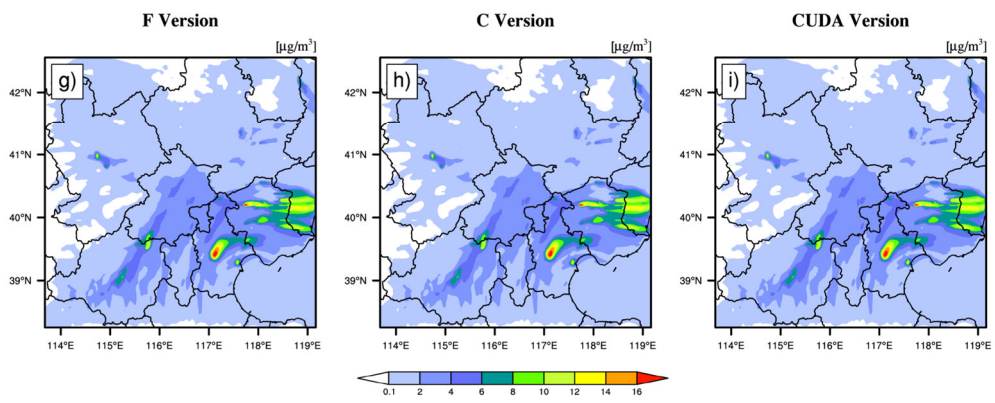
398 **Figure 6.** NO₂ and CO concentrations outputted by the CAMx model for the Fortran, standard C,

399 and CUDA C versions. Panels (a) and (g) are from the Fortran versions. Panels (b) and (h) are from
400 the standard C versions. Panels (c) and (i) are from the CUDA C versions. Panels (d) and (j) are the
401 output concentration differences of the Fortran and standard C versions. Panels (e) and (k) are the
402 output concentration differences of the standard C and CUDA C versions. Panels (f) and (l) are the
403 output concentration differences of the Fortran and CUDA C versions.

H₂O₂

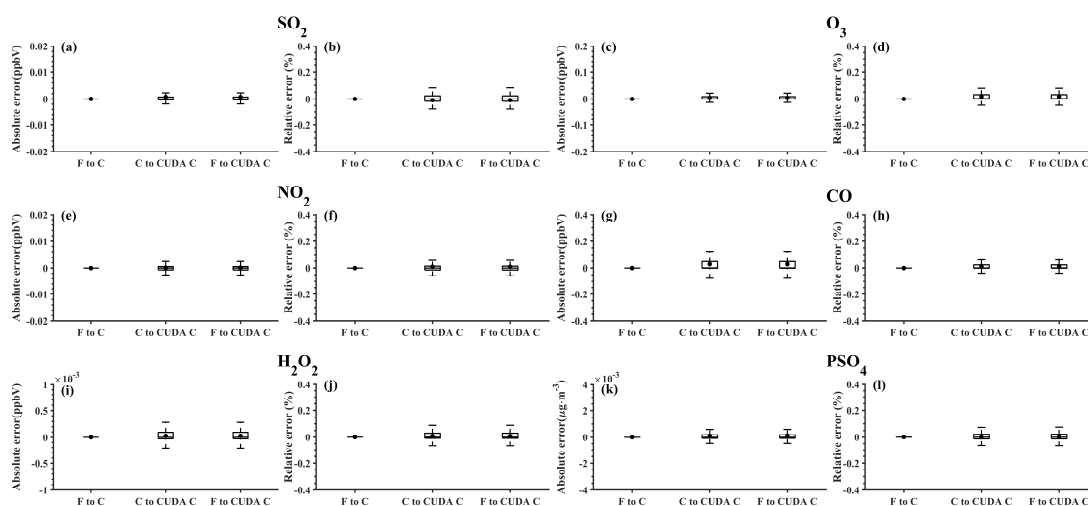


PSO₄



405 **Figure 7.** H₂O₂ and PSO₄ concentrations output by the CAMx model for the Fortran, standard C,
 406 and CUDA C versions. Panels (a) and (g) are from the Fortran versions. Panels (b) and (h) are from
 407 the standard C versions. Panels (c) and (i) are from the CUDA C versions. Panels (d) and (j) are the
 408 output concentration differences of Fortran and standard C versions. Panels (e) and (k) are the output
 409 concentration differences of the standard C and CUDA C versions. Panels (f) and (l) are the output
 410 concentration differences of the Fortran and CUDA C versions.

411 Figure 8 shows the boxplot of the AEs and relative error (REs) in all the grid boxes
 412 for the six species during the porting process. As described above, the AEs and REs
 413 introduced by Fortran to the standard C code refactoring process are significantly small,
 414 and the primary error comes from converting standard C to CUDA C. Statistically, the
 415 average AEs (REs) of SO₂, O₃, NO₂, CO, H₂O₂ and PSO₄ were -0.0009 ppbV (-0.01%),
 416 0.0004 ppbV (-0.004%), 0.0005 ppbV (0.008%), 0.03 ppbV (0.01%), 2.1×10^{-5}
 417 ppbV (-0.01%) and $0.0002 \mu\text{g} \cdot \text{m}^{-3}$ (0.0023%), respectively, between the Fortran
 418 and CUDA C versions. In terms of the time series, the regionally averaged time series
 419 of the three versions are almost consistent (as shown in Figure S2), and the maximum
 420 AEs for the above six species are 0.001 ppbv, 0.005 ppbv, 0.002 ppbv, 0.03 ppbv, 0.0001
 421 ppbv and $0.0002 \mu\text{g} \cdot \text{m}^{-3}$, respectively, between the Fortran and CUDA C versions.



422 **Figure 8.** The distributions of absolute errors and relative errors for SO₂, O₃, NO₂, CO, H₂O₂ and
 423 PSO₄ in all of the grid boxes after 48 hours of integration.
 424

425 Figure 9 presents the regionally averaged time series and the AEs of SO₂, O₃, NO₂,
 426 CO, H₂O₂ and PSO₄. The time series between the different versions is almost consistent,

427 and the maximum AEs for the above six species are 0.001 ppbv, 0.005 ppbv, 0.002 ppbv,
428 0.03 ppbv, 0.0001 ppbv and $0.0002 \mu\text{g} \cdot \text{m}^{-3}$, respectively, between the Fortran and
429 CUDA C versions.

430 It is difficult to verify the scientific applicability of the results from the CUDA C
431 version because the programming language and hardware are different between the
432 Fortran and CUDA C versions. Here, we used the evaluation method of Wang et al.
433 (2021a) to compute the root mean square errors (RMSEs) of SO₂, O₃, NO₂, CO, H₂O₂
434 and PSO₄ between the Fortran and CUDA C versions, which are 0.0007 ppbV, 0.001
435 ppbV, 0.0002 ppbV, 0.0005 ppbV, 0.00003 ppbV and $0.0004 \mu\text{g} \cdot \text{m}^{-3}$, respectively,
436 much smaller than the spatial variation of the whole region, which is 7.0 ppbV
437 (approximately 0.004%), 9.7 ppbV (approximately 0.003%), 7.4 ppbV (approximately
438 0.003%), 142.2 ppbV (approximately 0.006%), 0.2 ppbv (approximately 0.015%) and
439 $1.7 \mu\text{g} \cdot \text{m}^{-3}$ (approximately 0.004%). The bias between CUDA C and the Fortran
440 version of the above six species is negligible compared with their own spatial changes,
441 and the results of the CUDA C version are generally acceptable for research purposes.
442

443 **4.3. Offline performance comparison of GPU-HADVPPM**

444 As described in Sect. 4.2, we validate that the CAMx model result of the CUDA
445 C version is generally acceptable for scientific research. We tested the offline
446 performance of the HADVPPM and GPU-HADVPPM schemes on 1 CPU core and 1
447 GPU card. There are 7 variables input into the HADVPPM program, which are nn, dt,
448 dx, con, vel, area and areav, and their specific meanings are shown in Table S1.


449 First, we use the random_number function in Fortran to create random single-
450 precision floating-point numbers of different sizes for the above 7 variables, and then
451 transmit these random numbers to the hadvppm Fortran program and hadvppm_kernel
452 CUDA C program for computation. Finally, we test the offline performance of the
453 HADVPPM and GPU-HADVPPM on the CPU and GPU platforms. During the offline
454 performance experiments, we used two different CPUs and GPUs described in Sect.
455 2.4., and the experimental results are shown in Figure 9.

456 On the CPU platform, the wall time of the hadvppm Fortran program does not
457 change significantly when the data size is less than 1000. With the increase in the data
458 size, its wall time increases linearly. When the data size reaches 10^7 , the wall time of
459 the hadvppm Fortran program on the Intel Xeon E5-2682v4 and Intel Platinum 8168
460 CPU platforms is 1737.3 ms and 1319.0 ms, respectively. On the GPU platform, the
461 reconstructed and extended CUDA C program implements parallel computation of
462 multiple grid points by executing a large number of kernel function copies, so the
463 computational efficiency of the hadvppm_kernel CUDA C code on it is significantly
464 improved. In the size of 10^7 random numbers, the hadvppm_kernel CUDA C program
465 takes only 12.1 ms and 1.6 ms to complete the computation on the NVIDIA Tesla K40m
466 and NVIDIA Tesla V100 GPU.

467 Figure 9. (b) shows the speedup of HADVPPM and GPU-HADVPPM on the CPU
468 platform and GPU platform under different data sizes. When mapping the HADVPPM
469 scheme to the GPU, the computational efficiency under different data sizes is not only
470 significantly improved, but the larger the data size is, the more obvious the acceleration
471 effect of the GPU-HADVPPM. For example, in the size of 10^7 random numbers, the
472 GPU-HADVPPM achieved a 1113.6x and 845.4x acceleration on the NVIDIA Tesla
473 V100 GPU, respectively, compared to the two CPU platforms. Although the K40m
474 GPU's single-card computing performance is slightly lower than that of the V100 GPU,
475 GPU-HADVPPM can also achieve up to a 143.3x and 108.8x acceleration.

476 As described in Sect. 3.2, the thread is the most basic GPU unit for parallel
477 computing. Each dimension of the three-dimensional block can contain a maximum
478 number of threads of 1024, 1024 and 64. Each dimension of the three-dimensional grid
479 can contain a maximum number of blocks of $2^{31} - 1$, 65535, and 65535. It is
480 theoretically possible to distribute a large number of copies of kernel functions into tens
481 of billions of threads for parallel computing without exceeding the GPU memory. In
482 the offline performance experiments, the GPU achieved up to 10 million threads of
483 parallel computing, while the CPU can only use serial cyclic computation. Therefore,
484 GPU-HADVPPM achieves a maximum acceleration of approximately 1100x without

485 I/O. In addition to this study, the GPU-based SBU-YLIN scheme in the WRF model
 486 can achieve a 896x acceleration compared to the Fortran implementation running on
 487 the CPU (Mielikainen et al., 2012b).



488
 489 **Figure 9.** The offline performance of the HADVPPM and GPU-HADVPPM scheme on the CPU
 490 and GPU. The unit of the wall times for the offline performance experiments is in milliseconds (ms).

491 4.4. Coupling performance comparison of GPU-HADVPPM with different GPU 492 configurations

493 4.4.1. CAMx-CUDA on a single GPU

494 The offline performance results show that the larger the data size is, the more
 495 obvious the acceleration effect of the GPU-HADVPPM scheme. After coupling the
 496 GPU-HADVPPM to CAMx without changing the advection module algorithm, the
 497 overall computational efficiency of the CAMx-CUDA model is extremely low, and it
 498 takes approximately 621 minutes to complete a one-hour integration on the V100
 499 cluster. Therefore, according to the optimization scheme in Sect. 3.2, by optimizing the
 500 algorithm of the xyadvec Fortran program, we gradually increase the size of the data
 501 transmitted and reduce the data transmission frequency between the CPU and GPU.
 502 When the data transmission frequency between the CPU and GPU is reduced to 1 within
 503 one time step, we further optimize the GPU memory access order on the GPU card,
 504 eliminate unnecessary assignment loops before kernel functions are launched and use

505 the thread and block indices.

506 Table 4 lists the total elapsed time for different versions of the CAMx-CUDA
507 model during the optimization, as described in Section 3.2. Since the xyadvec program
508 in CAMx-CUDA V1.0 is not optimized, it is extremely computationally inefficient
509 when starting two CPU processes and configuring a GPU card for P1. On the K40m
510 and V100 clusters, it takes 10829 seconds and 37237 seconds, respectively, to complete
511 a 1-hour simulation.

512 By optimizing the algorithm of the xyadvec Fortran program and hadvppm_kernel
513 CUDA C program, the data transmission frequency between the CPU and GPU was
514 decreased, and the overall computing efficiency was improved after GPU-HADVPPM
515 was coupled to the CAMx-CUDA model. In CAMx-CUDA V1.2, the data transmission
516 frequency between CPU-GPU within one time step is reduced to 1, the elapsed time on
517 the two heterogeneous clusters is 1207 seconds and 548 seconds, respectively, and the
518 speedup is 9.0x and 68.0x compared to CAMx-CUDA V1.0.

519 The GPU memory access order can directly affect the overall GPU-HADVPPM
520 computational efficiency on the GPU. In CAMx-CUDA V1.3, we optimized the
521 memory access order of the hadvppm_kernel CUDA C program on the GPU and
522 eliminated the unnecessary assignment loops before the kernel functions were launched,
523 which further improved the CAMx-CUDA model's computational efficiency, resulting
524 in 12.7x and 94.8x speedups.


525 Using thread and block indices to simultaneously compute the horizontal grid
526 points can greatly improve the computational efficiency of the GPU-HADVPPM and
527 thus reduce the overall elapsed time of the CAMx-CUDA model. CAMx-CUDA V1.4
528 further reduces the elapsed time by 378 seconds and 103 seconds on the K40m cluster
529 and V100 cluster, respectively, compared with CAMx-CUDA V1.3 and achieves up to
530 a 29.0x and 128.4x speedup compared with CAMx-CUDA V1.0.

531 **Table 4.** Total elapsed time for different versions of CAMx-CUDA during the optimization. The
532 unit of elapsed time for experiments is in seconds (s).

Versions	K40m cluster	V100 cluster
----------	--------------	--------------

	Elapsed Time	Speedup	Elapsed Time	Speedup
CAMx-CUDA V1.0	10829	1.0	37237	1.0
CAMx-CUDA V1.1	1403	7.7	1082	34.4
CAMx-CUDA V1.2	1207	9.0	548	68.0
CAMx-CUDA V1.3	751	12.7	393	94.8
CAMx-CUDA V1.4	373	29.0	290	128.4

533 In terms of the single module computational efficiency of HADVPPM and GPU-
534 HADVPPM, we further tested the computational performance of the Fortran version of
535 HADVPPM on the CPU, C version of HADVPPM on the CPU, and the CUDA C
536 version of GPU-HADVPPM in CAMx-CUDA V1.4 (GPU-HADVPPM V1.4) on the
537 GPU using `system_clock` functions in the Fortran language and `cudaEvent_t` in CUDA
538 programming. The specific results are shown in Figure 10. On the K40m cluster, it takes
539 37.7 seconds and 51.4 seconds to launch the Intel Xeon E5-2682 v4 CPU to run the
540 Fortran and C version HADVPPM, respectively, and the C version is 26.7% slower
541 than the Fortran version. After the CUDA technology was used to convert the C code
542 into CUDA C, the CUDA C version took 29.6 seconds to launch an NVIDIA Telsa
543 K40m GPU to run GPU-HADVPPM V1.4, with a 1.3x and 1.7x acceleration. On the
544 V100 cluster, the Fortran, C, and CUDA C versions are computationally more efficient
545 than those on the K40m cluster. It takes 30.1 seconds and 45.2 seconds to launch the
546 Intel Xeon Platinum 8168 CPU to run the Fortran and C version HADVPPM, and 1.6
547 seconds to run the GPU-HADVPPM V1.4 using an NVIDIA V100 GPU. The
548 computational efficiency of the CUDA C version is 18.8x and 28.3x higher than the
549 Fortran and C versions, respectively.




550

551 **Figure 10.** The elapsed time of the Fortran version HADVPPM on the CPU, the C version
 552 HADVPPM on the CPU and the CUDA C version GPU-HADVPPM V1.4 on the GPU. The unit is
 553 in seconds (s).

554 4.4.2. CAMx-CUDA on multiple GPUs

555 To make full use of the multicore and multi-GPUs in the heterogeneous cluster,
 556 the MPI+CUDA acceleration algorithm was implemented to improve the total
 557 computational performance of the CAMx-CUDA model. Two different compile flags
 558 were implemented in this study before comparing the computational efficiency of
 559 CAMx-CUDA V1.3 and V1.4 on multiple GPUs, namely, *-mieee-fp* and *-fp-model*
 560 *precise*. The *-mieee-fp* compile flag comes from the *Makefile* of the official CAMx
 561 version, which uses the IEEE standard to compare the floating-point numbers. Its
 562 computational accuracy is higher but the efficiency is slower. The *-fp-model precise*
 563 compile flag controls the balance between the precision and efficiency of the floating-
 564 point calculations, and it can force the compiler to use the vectorization of some
 565 calculations under value safety. The experimental results show that the *-fp model*
 566 *precise* compile flag is 41.4% faster than *-mieee-fp*, and the AEs of the simulation
 567 results are less than ± 0.05 ppbv (Figure S3). Therefore, the *-fp model precise* compile
 568 flag is implemented when comparing the computational efficiency of CAMx-CUDA

569 V1.3 and V1.4 on multiple GPU cards. Figure 11 shows the total elapsed time and
 570 speedup of CAMx-CUDA V1.3 and V1.4 on the V100 cluster. The total elapsed time
 571 decreases as the number of CPU cores and GPU cards increases. When starting 8 CPU
 572 cores and 8 GPU cards, the speedup of CAMx-CUDA V1.4 is increased from 3.9x to
 573 4.5x compared with V1.3, and the computational efficiency is increased by 35.0%.



574

575 **Figure 11.** The total elapsed time and speedup of CAMx-CUDA V1.3 and V1.4 on multiple
 576 GPUs. The unit of elapsed time for experiments is seconds (s).

577 5. Conclusions and discussion

578 GPU accelerators are playing an increasingly important role in high-performance
 579 computing. In this study, a GPU acceleration version of the PPM solver (GPU-
 580 HADVPPM) of horizontal advection for an air quality model was developed, which
 581 runs on GPU accelerators using the standard C programming language and CUDA
 582 technology. The offline performance experimental results showed that the K40m and
 583 V100 GPU can achieve up to a 845.4x and 1113.6x speedup, respectively, and the larger
 584 the data input to the GPU, the more obvious the acceleration effect. After coupling the
 585 GPU-HADVPPM to the CAMx model, a series of optimization measures were taken,
 586 including reducing the CPU-GPU communication frequency, increasing the data

587 computation size on the GPU, optimizing the GPU memory access order and using
588 thread and block indices to improve the overall computing performance of the CAMx-
589 CUDA model. Using a single GPU card, the optimized CAMx-CUDA V1.4 model
590 improved the computing efficiency by 29.0x and 128.4x on the K40m cluster and the
591 V100 cluster, respectively. In terms of the single-module computational efficiency of
592 GPU-HADVPPM, it achieved a 1.3x and 18.8x speedup on an NVIDIA Tesla K40m
593 GPU and NVIDIA Tesla V100 GPU, respectively. To make full use of the multicore and
594 multi-GPU supercomputers and further improve the total computational performance
595 of the CAMx-CUDA model, a parallel architecture with an MPI+CUDA hybrid
596 paradigm was presented. After implementing the acceleration algorithm, the total
597 elapsed time decreased as the number of CPU cores and GPU cards increased, and it
598 achieved up to a 4.5x speedup when launching 8 CPU cores and 8 GPU cards compared
599 with 2 CPU cores and 2 GPU cards.

600 However, the current approach has some limitations, which are as follows:

601 1) We currently implement thread and block coindexing to compute horizontal
602 grid points in parallel. Given the CAMx Model 3-dimensional grid computing
603 characteristics, in the future, 3-dimensional thread and block coindexing will be
604 considered to compute 3-dimensional grid points in parallel.

605 2) The communication bandwidth of data transfer is one of the main issues
606 restricting the computing performance of the CUDA C codes on the GPUs. This
607 restriction holds true not only for GPU-HADVPPM but also for the WRF module
608 (Mielikainen et al., 2012b; Mielikainen et al., 2013b; Huang et al., 2013). In this study,
609 the data transmission efficiency between the CPU and GPU is improved only by
610 reducing the communication frequency. In the future, more technologies, such as
611 pinned memory (Wang et al., 2016), will be considered to resolve the communication
612 bottleneck between the CPUs and GPUs.

613 3) To further improve the overall computational efficiency of the CAMx model,
614 the heterogeneous porting scheme proposed in this study will be considered to conduct
615 the heterogeneous porting of other CAMx modules in the future.

616

617 *Code and data availability.* The source codes of CAMx version 6.10 are available at
618 <https://camx-wp.azurewebsites.net/download/source/>(last access: 24 March 2023,
619 ENVIRON,2022). The dataset related to this paper and the CAMx-CUDA codes are
620 available online via ZENODO (<http://doi.org/10.5281/zenodo.7765218>; Cao et
621 al.,2023).

622

623 *Author contributions.* KC conducted the simulation and prepared the materials. QZW,
624 LLW and LNW planned and organized the project. KC, QZW and XT refactored and
625 optimized the codes. LLW, NW, HQC and DQL collected and prepared the data for the
626 simulation. KC, QZW, XT and LNW participated in the discussion.

627

628 *Competing interests.* The authors declare that they have no conflicts of interest.

629

630 **Acknowledgements.** The National Key R&D Program of China (2020YFA0607804 &
631 2017YFC0209805), the National Supercomputing Center in Zhengzhou Innovation
632 Ecosystem Construction Technology Special Program (Grant No. 201400210700) and
633 the Beijing Advanced Innovation Program for Land Surface funded this work. The
634 research is supported by the High Performance Scientific Computing Center (HSCC)
635 of Beijing Normal University and the National Supercomputing Center in Zhengzhou.

636

637 **References**

638 Bleichrodt, F., Bisseling, R. H., and Dijkstra, H. A.: Accelerating a barotropic ocean
639 model using a GPU, *Ocean Modelling*, 41, 16-21, 10.1016/j.ocemod.2011.10.001,
640 2012.

641 Cao, K., Wu, Q., Wang, L., Wang, N., Cheng, H., Tang, X., Li, D., and Wang, L.: The
642 dataset of the manuscript "GPU-HADVPPM V1.0: high-efficient parallel GPU

643 design of the Piecewise Parabolic Method (PPM) for horizontal advection in air
644 quality model (CAMx V6.10)", ZENODO,
645 <https://doi.org/10.5281/zenodo.7765218>, 2023.

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

649 ENVIRON: User Guide for Comprehensive Air Quality Model with Extensions
650 Version 6.1, available at: [https://camx-wp.azurewebsites.net/Files/CAMxUser](https://camx-wp.azurewebsites.net/Files/CAMxUserGuide_v6.10.pdf)
651 [sGuide_v6.10.pdf](https://camx-wp.azurewebsites.net/Files/CAMxUserGuide_v6.10.pdf) (last access: 19 December 2022), 2014

652 Govett, M., Rosinski, J., Middlecoff, J., Henderson, T., Lee, J., MacDonald, A., Wang,
653 N., Madden, P., Schramm, J., and Duarte, A.: Parallelization and Performance of
654 the NIM Weather Model on CPU, GPU, and MIC Processors, *Bulletin of the*
655 *American Meteorological Society*, 98, 2201-2213, [10.1175/bams-d-15-00278.1](https://doi.org/10.1175/bams-d-15-00278.1),
656 2017.

657 Houyoux, M. R. and Vukovich, J. M.: Updates to the Sparse Matrix Operator Kernel
658 Emissions (SMOKE) Modeling System and Integration with Models-3,
659 Huang, B., Mielikainen, J., Plaza, A. J., Huang, B., Huang, A. H. L., and Goldberg, M.
660 D.: GPU acceleration of WRF WSM5 microphysics, *High-Performance*
661 *Computing in Remote Sensing*, [10.1117/12.901826](https://doi.org/10.1117/12.901826), 2011.

662 Huang, B., Huang, M., Mielikainen, J., Huang, B., Huang, H. L. A., Goldberg, M. D.,
663 and Plaza, A. J.: On the acceleration of Eta Ferrier Cloud Microphysics Scheme in
664 the Weather Research and Forecasting (WRF) model using a GPU, *High-*
665 *Performance Computing in Remote Sensing II*, [10.1117/12.976908](https://doi.org/10.1117/12.976908), 2012.

666 Huang, M., Huang, B., Chang, Y.-L., Mielikainen, J., Huang, H.-L. A., and Goldberg,
667 M. D.: Efficient Parallel GPU Design on WRF Five-Layer Thermal Diffusion
668 Scheme, *IEEE Journal of Selected Topics in Applied Earth Observations and*
669 *Remote Sensing*, 8, 2249-2259, [10.1109/jstars.2015.2422268](https://doi.org/10.1109/jstars.2015.2422268), 2015.

670 Huang, M., Huang, B., Mielikainen, J., Huang, H. L. A., Goldberg, M. D., and Mehta,
671 A.: Further Improvement on GPU-Based Parallel Implementation of WRF 5-Layer

672 Thermal Diffusion Scheme, 2013 International Conference on Parallel and
673 Distributed Systems, 10.1109/icpads.2013.126, 2013.

674 Jiang, J., Lin, P., Wang, J., Liu, H., Chi, X., Hao, H., Wang, Y., Wang, W., and Zhang,
675 L.: Porting LASG/ IAP Climate System Ocean Model to Gpus Using OpenAcc,
676 IEEE Access, 7, 154490-154501, 10.1109/access.2019.2932443, 2019.

677 Mielikainen, J., Huang, B., Huang, H.-L. A., and Goldberg, M. D.: GPU Acceleration
678 of the Updated Goddard Shortwave Radiation Scheme in the Weather Research
679 and Forecasting (WRF) Model, IEEE Journal of Selected Topics in Applied Earth
680 Observations and Remote Sensing, 5, 555-562, 10.1109/jstars.2012.2186119,
681 2012a.

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

686 Mielikainen, J., Huang, B., Huang, H. L. A., Goldberg, M. D., and Mehta, A.: Speeding
687 Up the Computation of WRF Double-Moment 6-Class Microphysics Scheme with
688 GPU, Journal of Atmospheric and Oceanic Technology, 30, 2896-2906,
689 10.1175/jtech-d-12-00218.1, 2013a.

690 Mielikainen, J., Huang, B., Wang, J., Allen Huang, H. L., and Goldberg, M. D.:
691 Compute unified device architecture (CUDA)-based parallelization of WRF
692 Kessler cloud microphysics scheme, Computers & Geosciences, 52, 292-299,
693 10.1016/j.cageo.2012.10.006, 2013b.

694 NVIDIA: CUDA C++ Programming Guide Version 10.2, available at:
695 [https://docs.nvidia.com/cuda/archive/10.2/pdf/CUDA_C_Programming_Guide.p](https://docs.nvidia.com/cuda/archive/10.2/pdf/CUDA_C_Programming_Guide.pdf)
696 [df](https://docs.nvidia.com/cuda/archive/10.2/pdf/CUDA_C_Programming_Guide.pdf) (last access: 19 December 2022), 2020

697 NVIDIA: Floating Point and IEEE 754 Compliance for NVIDIA GPUs. Release 12.1,
698 available at: <https://docs.nvidia.com/cuda/floating-point/#differences-from-x86>
699 (last access: 18 May 2023), 2023.

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

701 Source Code Documentation and Validation, 1996.

702 Price, E., Mielikainen, J., Huang, M., Huang, B., Huang, H.-L. A., and Lee, T.: GPU-
703 Accelerated Longwave Radiation Scheme of the Rapid Radiative Transfer Model
704 for General Circulation Models (RRTMG), IEEE Journal of Selected Topics in
705 Applied Earth Observations and Remote Sensing, 7, 3660-3667,
706 10.1109/jstars.2014.2315771, 2014.

707 Skamarock, W. C., Klemp, J. B., Dudhia, J., Gill, D. O., Barker, D.M., Duda, M. G.,
708 Huang, X. Y., Wang, W., and Powers, J. G.: A Description of the Advanced
709 Research WRF Version3 (No.NCAR/TN-475CSTR), University Corporation for
710 Atmospheric Research, <https://doi.org/10.5065/D68S4MVH>, NCAR, 2008.

711 Streets, D. G., Zhang, Q., Wang, L., He, K., Hao, J., Wu, Y., Tang, Y., and Carmichael,
712 G. R.: Revisiting China's CO emissions after the Transport and Chemical
713 Evolution over the Pacific (TRACE-P) mission: Synthesis of inventories,
714 atmospheric modeling, and observations, Journal of Geophysical Research:
715 Atmospheres, 111, <https://doi.org/10.1029/2006JD007118>, 2006.

716 Streets, D. G., Bond, T. C., Carmichael, G. R., Fernandes, S. D., Fu, Q., He, D., Klimont,
717 Z., Nelson, S. M., Tsai, N. Y., Wang, M. Q., Woo, J. H., and Yarber, K. F.: An
718 inventory of gaseous and primary aerosol emissions in Asia in the year 2000,
719 Journal of Geophysical Research: Atmospheres, 108,
720 <https://doi.org/10.1029/2002JD003093>, 2003.

721 Sun, Y., Wu, Q., Wang, L., Zhang, B., Yan, P., Wang, L., Cheng, H., Lv, M., Wang, N.,
722 and Ma, S.: Weather Reduced the Annual Heavy Pollution Days after 2016 in
723 Beijing, Sola, 18, 135-139, 10.2151/sola.2022-022, 2022.

724 Wahib, M. and Maruyama, N.: Highly optimized full GPU-acceleration of non-
725 hydrostatic weather model SCALE-LES, 2013 IEEE International Conference on
726 Cluster Computing (CLUSTER), 23-27 Sept. 2013, 1-8,
727 10.1109/CLUSTER.2013.6702667,

728 Wang, P., Jiang, J., Lin, P., Ding, M., Wei, J., Zhang, F., Zhao, L., Li, Y., Yu, Z., Zheng,
729 W., Yu, Y., Chi, X., and Liu, H.: The GPU version of LASG/IAP Climate System

730 Ocean Model version 3 (LICOM3) under the heterogeneous-compute interface for
731 portability (HIP) framework and its large-scale application, *Geosci. Model Dev.*,
732 14, 2781-2799, 10.5194/gmd-14-2781-2021, 2021a.

733 Wang, Y., Guo, M., Zhao, Y., and Jiang, J.: GPUs-RRTMG_LW: high-efficient and
734 scalable computing for a longwave radiative transfer model on multiple GPUs,
735 *The Journal of Supercomputing*, 77, 4698-4717, 10.1007/s11227-020-03451-3,
736 2021b.

737 Wang, Z., Wang, Y., Wang, X., Li, F., Zhou, C., Hu, H., and Jiang, J.: GPU-
738 RRTMG_SW: Accelerating a Shortwave Radiative Transfer Scheme on GPU,
739 *IEEE Access*, 9, 84231-84240, 10.1109/access.2021.3087507, 2016.

740 Xiao, H., Lu, Y., Huang, J., and Xue, W.: An MPI+OpenACC-based PRM scalar
741 advection scheme in the GRAPES model over a cluster with multiple CPUs and
742 GPUs, *Tsinghua Science and Technology*, 27, 164-173,
743 10.26599/TST.2020.9010026, 2022.

744 Xu, S., Huang, X., Oey, L. Y., Xu, F., Fu, H., Zhang, Y., and Yang, G.: POM.gpu-v1.0:
745 a GPU-based Princeton Ocean Model, *Geoscientific Model Development*, 8,
746 2815-2827, 10.5194/gmd-8-2815-2015, 2015.

747 Zhang, Q., Streets, D. G., Carmichael, G. R., He, K. B., Huo, H., Kannari, A., Klimont,
748 Z., Park, I. S., Reddy, S., Fu, J. S., Chen, D., Duan, L., Lei, Y., Wang, L. T., and
749 Yao, Z. L.: Asian emissions in 2006 for the NASA INTEX-B mission, *Atmos.*
750 *Chem. Phys.*, 9, 5131-5153, 10.5194/acp-9-5131-2009, 2009.