

28th June 2018

Dear Editor,

Re: Review of “GEOS-Chem High Performance (GCHP): A next-generation implementation of the GEOS-Chem chemical transport model for massively parallel applications” for *Geoscientific Model Development*

We thank the referees for taking the time to review our paper, and for their insightful comments. We have listed their reviews in full below, with a response given for each point. Taking into account the recommendations from each reviewer and the comment from the executive editor, we have modified the manuscript (attached), highlighting any changes which were made since the original submission. We also attach a “clean” copy of the updated manuscript. Where possible, line numbers are included in the responses which correspond to the “clean” updated manuscript. In addition to the points raised by the reviewers and by the editor, we have made two other minor sets of changes. Firstly, the lower section of figure 5 (AOD) was previously and erroneously showing a one-hour rather than one-day average. This has been updated, and does not change the findings of the paper. Secondly, some minor changes have also been made to the text to improve clarity, but again these changes are superficial and all highlighted in the attached manuscript. We hope that you will find our edits satisfactory and look forward to your comments.

Executive Editor’s Comment

Dear authors,

1. Please include the version number for GCHP in the title, and throughout the manuscript.

Response: We have changed the title and several references to GCHP to instead read as “GCHP v11-02c”.

2. The precise version of the code discussed in the manuscript must be made available. The current best practice is for this code to be uploaded to a public repository and a DOI assigned. The DOI should be cited in the manuscript.

Response: We have now uploaded the specific version of the code used here to a public repository (GitHub). This version of the code has been assigned a DOI, which is provided in the Code Availability statement.

Reviewer #1

Summary

I am not sure what to expect from these sort of papers as they do describe complex systems with only couple of pages of words. So if you are an outside reader then it is almost impossible to get a firm idea about the described system. It is also debatable whether these brief descriptions have any usefulness to other researcher? The model described in this

paper has mainly two components, (i) a chemistry module, which is local and should, in principle, have a near perfect scaling (without any MPI instruction), while the second component (ii) is advection, which is not local and requires halo filling (or communication on demand) for parallel implementations. The only possible reason for publishing such papers would be to provide a reference for the model in question.

Response: The paper is indeed by necessity too short to provide a complete description of the model. To address the reviewer's concern, we now provide links to the GCHP website where a detailed description of the model including a user's manual are available. As recognized by the reviewer, publication of this paper in GMD will provide an important reference for the very large community of GEOS-Chem users wishing to use GCHP. But it will also be useful to developers of other models who may want to follow a similar development strategy, and they can go to the GCHP website to learn more.

Detailed comments

1. I think the paper needs a table summarizing the definitions of all the acronyms (GEOS-Chem, CTM, GCHP, ect...) used, because I kept going forward and backward to look for their meanings in the text.

Response: A table has now been added to the introduction (Table 1) which lists every acronym used in the paper.

2. My understanding of atmospheric chemistry models is that they solve a system of coupled ODEs at each grid of the model and each grid point is, in principle, independent of its neighbors. In other words, the chemical model needs information on the grid only and therefore such process should scale perfectly and these models are ideal for parallelism. However, advection needs information about the characteristic (hyperbolic problems) and this part that needs effort to make it work with parallel implementation (MPI). This looks like a coding task of combining FV3 transport (Lin et al.) with an existing chemistry model (with OpenMP). I think, the whole description of the new system and how it differs from the original GCC could be improved. There should be more clarifications and a detailed description of what changed from the original code (probably a table listing all the components of GCHP and how it differs from the original GCC).

Response: One of the main features of GCHP is that it does not change the code from GCC, and in fact uses the exact same code to calculate local chemistry. We now stress this detail in section 2.2 ("GCHP v11-02c model architecture"): "Within each atmospheric domain, local terms are calculated by a standard copy of the GEOS-Chem Classic code, embedded in the model as described by Long et al. (2015). This copy of the GEOS-Chem code is identical to that used in GEOS-Chem Classic, such that all processes other than advection which are simulated in GCC are simulated identically in GCHP". We also now clarify that "The embedded copy of GEOS-Chem in GCHP is compiled without OpenMP shared-memory parallelization, resulting in a pure MPI implementation" (page 5, lines 3 to 9).

3. Page 1, Abstract, Line 2. I am not sure how large the system of chemical species is? It would be better to give an order of unknowns for a typical system, i.e., order 100-1000000?.

Response: We have added a clarification that these systems typically involve of the order of 100 - 1,000 chemical species (page 1, line 16).

4. The discussion at page 7 in relation to semi-Lagrangian, cube-sphere and lat-long grids is a bit subjective!

- a) I am not sure what the authors means by “inconsistency in technique” in the sentence “The problem can be mitigatedin ensuring mass conservation”?

Response: GCC uses a mix of Eulerian and semi-Lagrangian methods, depending on the local CFL number. However, we agree that this is a minor point only and does not add significantly to the discussion. As such we have changed this list to focus on the question of mass conservation (page 7, line 9).

- b) I think the whole passage “In an MPI environment, sharing the pole” cannot be justified. Semi-Lagrangian (SL) schemes on any grid with the right halo size can achieve near perfect scaling [see for example with up to 100000 cores SL scalability experiments in the paper: *High-performance high-resolution semi-Lagrangian tracer transport on a sphere. J.B. White III and J.J. Dongarra, Journal of Computational Physics, Vol. 230, pp. 6778-679 (2011)*] or up to 10000 cores scalability results in *Allen and Zerroukat, Journal of Computational Physics, Vol. 319, p. 44-60 (2016)*]. Certainly the only argument against SL is the lack of inherent conservation, but in terms of accuracy, stability, simplicity, computational cost, and scalability there is no valid arguments against SL.

Response: We agree that the problems here stem more from the use of a rectilinear lat-lon grid than from the use of the semi-Lagrangian method. We now clarify simply that the use of a semi-Lagrangian method when $CFL > 1$ results in increased domain halo size, and therefore increased communication (page 7, lines 10-12).

- c) Of course there is no perfect scheme and every approach has some advantages and disadvantages and one has to be objective about these things. For example the cube-sphere has the disadvantages of dealing with 6 panels and their orientations and the non-orthogonality of the grid and its associated grid-imprinting (for example see your figure 5 for C24 where it is very clear that the pattern of the tracer distribution is very much influenced by the grid). It would be interesting to show the equivalent pictures (blowup pictures at Figure 5 for C24) due to the original GCC!.

Response: We now include results from GCC simulations at two resolutions in our comparisons of simulated ozone and AOD. Accordingly, Figure 5 has been split into two. The new Figure 5 shows ozone at 4km when simulated at C24 and C180 (GCHP) as well as at $4^\circ \times 5^\circ$ and $2^\circ \times 2.5^\circ$ (GCC). Figure 6 now shows AOD for the same 4 simulations. Furthermore, we use 2-D histograms (Figure 7) to directly compare estimated ozone at 4 km between GCC (at $4^\circ \times 5^\circ$) to the results at both C24 and C180 with GCHP (regidded to $4^\circ \times 5^\circ$). We believe that this provides a convincing

demonstration of GCHP's ability to reproduce the capabilities of GCC while also highlighting some of the differences between the two architectures.

5. Page 8, 1st paragraph of section 2.3.2. In parallel domain decomposition, I would suggest to use the right nomenclature of halo-filling instead of "providing boundary conditions", because there is no real boundaries between processors.

Response: We have changed this line accordingly (page 8, line 10).

6. Page 8, end of 2nd paragraph of section 2.3.2. It is not clear what does mean the comment: "GCHP defaults instead to a simple global air mass correction also applied to tracers"?? My understanding is that the main purpose of using FV3 is to achieve mass conservation in an inherent way so I don't understand what is this correction is doing? Does it mean the inherent conservation is lost? Further clarifications are needed here.

Response: The greatest advantage of the FV3 model, as we perceive it, is to enable the execution of the model using a grid other than the rectilinear latitude-longitude grid previously in use. However, any CTM which uses time-averaged winds and interpolated surface pressure tendencies to estimate advection will inherently introduce a mass conservation error ([Jöckel et al. 2001](#)). Once mass flux information is available directly from the parent GCM or data assimilation system - in this case, the NASA GEOS DAS - this error is expected to be significantly reduced, but currently the standard meteorological products from all major forecasting centers do not include mass fluxes, instead reporting wind speeds. Until these fluxes are available, a mass correction is needed to ensure that tracer mass is conserved. In our case we have used a simple global mass correction, which has advantages and disadvantages as discussed in the aforementioned paper. We now clarify in the text that no pressure fixer has yet been designed for advection on the cubed sphere, but that GCHP is designed to accept mass fluxes once available (page 8, lines 22-25).

7. I don't see any use for to the bottom sub-figure (tim/1000colms) of figure 3.

Response: The bottom sub-figure is designed to show that the model's performance and scalability per computer core improve as the resolution is increased, due to the diminishing effect of the I/O overhead, and that for higher resolutions GCHP's performance per grid point is superior to that of GCC. This is now clarified explicitly in the paper on page 11, lines 8-11.

8. Section 4 (example simulation). It is good to show how the scalability of GCHP can be exploited to run higher-resolution simulations and the improvement that comes as results of that (Figure 5). However, I feel there are some figures missing in this section/paper. Although there is a lot figures comparing GCC and GCHP in terms of computational efficiency (scaling and wall clock time), for fidelity there should be at least a couple of figures comparing (at least visually) GCC and GCHP for a benchmark test to show that the solutions of the two models are, at least visually, in a good agreement.

Response: We agree completely that it is useful to compare the output of GCHP to that from GCC, in order to validate the results. As mentioned in a prior comment, we

have therefore expanded this section to include multiple comparisons between GCHP and GCC.

Reviewer #2

The manuscript presents the development of a high-performance computing capability (GCHP) for the GEOS-Chem CTM. The paper address relevant scientific modelling questions within the scope of GMD and presents a model with valid and clearly outlined methods.

It should in theory be possible for an independent scientist to construct a model that, while not necessarily numerically identical, will produce scientifically equivalent results. For the benchmarking, it should be possible for the results to be precisely reproduced.

General Comments:

The large number of similar acronyms is confusing to the reader. Consider summarising in a table and perhaps showing the relation between the different modeling/data components by means of a diagram to aid the reader.

Response: A table has now been added (Table 1) which clarifies all acronyms used in the table. Although Figure 1 shows the layout of GCHP diagrammatically, we have also clarified on page 5 (lines 5-9) the relationship between pre-existing simulation code from GEOS-Chem (all columnar chemistry and physics), and new simulation code which is exclusive to GCHP (advection).

Beyond the scaling comparison between GHCP and GCC and changes in resolution, a direct comparison between the two should be added (in the same resolution/configuration) to show the correctness but also the difference in accuracy and efficiency because of the new grid.

Response: We have now expanded section 4 to include multiple comparisons between GCC and GCHP results, using both ozone and AOD at 4°x5° and 2°x2.5°.

P.8 l.26: Any indication/reference on the error introduced by the correction?

Response: The effect of these corrections was not quantified during the model runs, but the effect of using “pressure fixers” or mass corrections to counter the long-standing issue of inconsistency between winds and surface pressure tendencies is explored in detail in Jöckel et al (2001). In light of this comment and one from reviewer #1, we have extended the discussion of this correction to stress the advantage of using mass fluxes directly as opposed to inferring them from winds and surface pressure (page 8, lines 23-25).

P. 12: Is the statement “The scalability of advection suffers from the additional communication overhead associated with reducing the domain size, as each domain must communicate a larger portion of its concentration data to its neighbors” referring to the communication vs computation time for each MPI process?

Response: That is correct. Advection is the only simulated process which requires communication between processes.

Code Availability is outdated and has to be updated. Also, license information should be added.

Response: The Code Availability section has been updated in line with the recommendations of the Executive Editor, including the addition of a DOI. GEOS-Chem and GCHP both use the MIT License, and this detail has also been added to the Code Availability section (page 17).

Minor Comments:

P.8 l.13: Remove “information from”

Response: This sentence has been rewritten in response to comments from Reviewer #1 for clarity.

P.8 l.18: “/” instead of “.”

Response: Corrected.

P.15 l.8: Remove “an unchanged copy of”

Response: Removed.

P.15 l.12 Please consider rephrasing to “GCHP was tested and show to scale from six cores up to at least 540 [...]”

Response: Rephrased accordingly.

Thank you again for considering this work for publication in *Geoscientific Model Development*, and we would again like to thank the reviewers for their careful comments.

Sincerely,

Sebastian D. Eastham

GEOS-Chem High Performance (GCHP v11-02c): A next-generation implementation of the GEOS-Chem chemical transport model for massively parallel applications

5 Sebastian D. Eastham^{1,2*}, Michael S. Long², Christoph A. Keller^{3,4}, Elizabeth Lundgren², Robert M. Yantosca², Jiawei Zhuang², Chi Li⁵, Colin J. Lee⁵, Matthew Yannetti², Benjamin M. Auer^{3,6}, Thomas L. Clune³, Jules Kouatchou^{3,6}, William M. Putman³, Matthew A. Thompson^{3,6}, Atanas L. Trayanov^{3,6}, Andrea M. Molod³, Randall V. Martin^{5,7}, Daniel J. Jacob²

¹Laboratory for Aviation and the Environment, Massachusetts Institute of Technology

²John A. Paulson School of Engineering and Applied Sciences, Harvard University

10 ³NASA Global Modeling and Assimilation Office

⁴Universities Space Research Association, Columbia, MD, USA

⁵Department of Physics and Atmospheric Science, Dalhousie University

⁶Science Systems and Applications, Inc., Lanham, MD, USA

⁷Smithsonian Astrophysical Observatory, Harvard-Smithsonian Center for Astrophysics

15 *Correspondence to:* Sebastian D. Eastham (seastham@mit.edu)

Abstract. Global modeling of atmospheric ~~composition chemistry~~ is a grand computational challenge because of the need to simulate large coupled systems of ~100-1000 chemical species interacting with transport on all scales. Off-line chemical transport models (CTMs), where the chemical continuity equations are solved using meteorological data as input, have ~~the advantages of simplicity usability advantages and reproducibility~~, and are important vehicles for developing atmospheric chemistry knowledge that can then be transferred to Earth system models. However, they have generally not been designed to take advantage of massively parallel computing architectures. Here we develop such a high-performance capability (GCHP) for GEOS-Chem, a CTM driven by ~~GEOS~~ meteorological data from the NASA Goddard Earth Observation System (GEOS) and used by hundreds of research groups worldwide. GCHP is a grid-independent implementation of GEOS-Chem using the Earth System Modeling Framework (ESMF) that permits the same standard model to ~~be run~~ operate in a distributed-memory framework for massive parallelization. ~~-, scalable from six cores on a single node up to hundreds of cores distributed across a network.~~ GCHP also allows GEOS-Chem to take advantage of the native GEOS cubed-sphere grid for greater accuracy and computational efficiency in simulating transport. GCHP enables GEOS-Chem simulations to be conducted with high computational scalability up to at least 500 cores, so that global simulations of stratosphere-troposphere oxidant-aerosol chemistry at C180 spatial resolution ($\sim 0.5^\circ \times 0.625^\circ$) or finer become routinely feasible.

20
25

1 Introduction

Atmospheric chemistry models are used to address a wide range of problems related to climate forcing, air quality, and atmospheric deposition. Simulations of oxidant and aerosol chemistry involve hundreds of chemically interacting species, coupled to transport on all scales. The computational demands are considerable, which has limited the inclusion of atmospheric chemistry in climate models (National Research Council, 2012). Off-line chemical transport models (CTMs), where meteorology is provided as input data from a parent global climate model (GCM) or atmospheric data assimilation system (DAS), are frequently used for reasons of simplicity, reproducibility, and ability to focus on chemical processes. The global GEOS-Chem CTM originally described by Bey et al. (2001), using meteorological input from the Goddard Earth Observation System (GEOS) DAS of the NASA Global Modeling and Assimilation Office (GMAO), is used by hundreds of atmospheric chemistry research groups worldwide (<http://www.geos-chem.org>). Increasing computational resources in the form of massively parallel architectures can allow GEOS-Chem users to explore more complex problems at higher grid resolutions, but this requires re-engineering of the model to take advantage of these architectures. Here we describe a high-performance version of GEOS-Chem (GCHP) engineered for this purpose, and we demonstrate its ability to access a new range of capability and scales for global atmospheric chemistry modeling.

15

Table 1. Acronyms used in this paper.

<u>Acronym</u>	Description
AOD	Aerosol Optical Depth
CFL	Courant-Friedrichs-Lewy number
CTM	Chemical Transport Model
DAS	Data Assimilation System
ESMF	Earth System Modeling Framework
FV3	Finite Volume advection on the Cubed sphere
GCC	GEOS-Chem Classic
GCHP	GEOS-Chem High Performance
GCM	Global Climate Model
GEOS	Goddard Earth Observation System
GMAO	NASA Global Modeling and Assimilation Office
HEMCO	Harvard-NASA Emissions COmponent
MAPL	Modeling and Analysis Prediction Layer
MPI	Message Passing Interface
UCX	Unified Chemistry eXtension

The original GEOS-Chem CTM (“GEOS-Chem Classic”, or GCC) was designed for shared-memory (OpenMP) parallelization. Detailed description of the model including a user’s manual is available on the GEOS-Chem website (<http://www.geos-chem.org>). Computation is distributed over a number of cores on a single node, with data held in shared arrays. But recent growth in computational power has taken the form of massively-parallel networked systems, where additional computational power is achieved by increasing the number of identical nodes rather than by improving the nodes themselves. This has placed a restriction on growth in the problem size and complexity which can be solved by a single instance of GCC. To take advantage of massively-parallel architectures, a new framework is needed which allows GEOS-Chem to use a distributed-memory model, where the computation is distributed across multiple coordinated nodes using a Message Passing Interface (MPI) implementation such as MVAPICH2 or OpenMPI.

An important first step in this evolution was the integration of GEOS-Chem as the online chemistry component within the GEOS DAS (Long et al., 2015). In order to ensure that the online and offline versions of GEOS-Chem were identical, GCC was modified ~~so that to use~~ the exact same code ~~is now used~~ in the independent CTM and in the DAS. Major modifications were required to make GEOS-Chem grid-independent and compatible with the GEOS Modeling and Analysis Prediction Layer (MAPL) (Suarez et al., 2007), an Earth System Modeling Framework (ESMF) (Hill et al., 2004) based software layer which handles communication between different components of the GEOS DAS. The GEOS-Chem code was adapted to accept an arbitrarily-sized horizontal set of atmospheric columns, with no requirements regarding adjacency of the columns or overall coverage of any particular set. All these changes were made “under the hood” in the standard GEOS-Chem Code. When GEOS-Chem is run as GCC, the set of columns is designated as a single block which covers the entire globe or a subset in a nested domain, and parallelization is achieved by internally running parallel loops over the columns. When GEOS-Chem is run as part of GEOS, MAPL internally splits the atmosphere into smaller domains, each of which contains a different set of atmospheric columns. These domains can then be distributed across multiple nodes, exploiting massively-parallel architectures. As a result of these changes, the same GEOS-Chem code can now be run either as a stand-alone, shared-memory offline CTM, or as a GCM component in the massively-parallel, distributed GEOS DAS. Any improvement in chemical modeling developed for the offline CTM is thus immediately available in the GEOS DAS version, which never becomes out of date and remains referenceable to the current version of ~~the model~~ GEOS-Chem.

In this work we take the next logical step of developing GCHP as a distributed-memory, MAPL-based implementation of the GEOS-Chem CTM. GCHP uses an identical copy of the GEOS-Chem Classic (GCC) code to provide the same high-fidelity atmospheric chemical simulation capabilities, allowing users to switch between GCC and GCHP implementations with confidence that they are using the same model. The exact same internal code is used in GCC shared-memory and GCHP distributed-memory applications. This closes the development loop between online and offline modeling. By sharing infrastructure code between GCHP and GEOS in the form of MAPL, offline modelers can now take advantage of modeling

advances which originate in the online model in the same way that GEOS benefits from advances in chemical modeling developed in the GEOS-Chem CTM (Nielsen et al., 2017). By way of example, GEOS was recently able to conduct a full-year 13-km resolution “nature run” with the current standard version of GEOS-Chem tropospheric chemistry (Keller et al., 2017; Hu et al., 2018). In return, GCHP incorporates the more efficient cubed sphere grid and FV3 advection code present in GEOS, and is capable of directly ingesting GEOS output in its native cubed sphere format.

2 Model description

2.1 Overview

Atmospheric chemistry models such as GEOS-Chem solve the 3-D chemical continuity equations for an ensemble of m coupled chemical species (Brasseur and Jacob, 2017). The continuity equation for the number density n_i [molecules cm^{-3}] of species i is expressed as

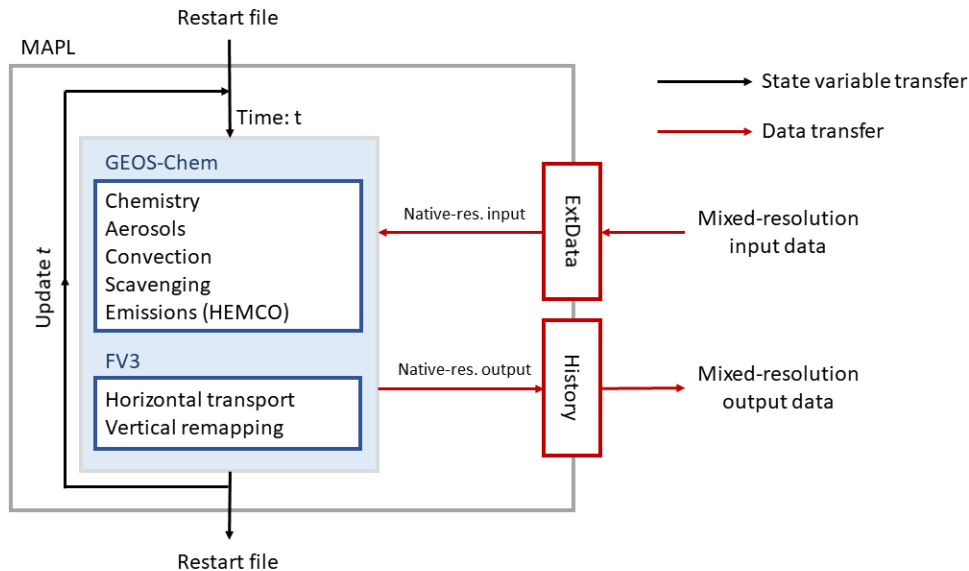
$$\frac{\partial n_i}{\partial t} = -\nabla \cdot (n_i \mathbf{v}) + s_i \quad (1)$$

where \mathbf{v} is the velocity vector [m s^{-1}], and s_i is the local net production and loss of species i [molecules $\text{cm}^{-3} \text{ s}^{-1}$]. In CTMs, \mathbf{v} is provided by archived output from a parent GCM or DAS, with subgrid-scale parameterized transport statistics (boundary layer mixing, deep convection) as additional CTM transport terms in Eqn. 1. From a computational standpoint, the local term s_i is grid-independent. However, the transport terms are grid-aware, as they move material between grid points. In GEOS-Chem, the atmosphere is split into independent columns, with each column made up of a number of discrete grid points (Long et al., 2015). Vertical processes (boundary layer mixing, deep convection) are then considered to be local in the sense that they are calculated independently for each column. In each column simulated by GEOS-Chem, the local term computes chemical evolution with a unified tropospheric-stratospheric mechanism (Eastham et al., 2014; Sherwen et al., 2016), convective transport (Wu et al., 2007), boundary layer mixing (Lin and McElroy, 2010), radiative transfer and photolysis (Prather, 2012), wet scavenging (Liu et al., 2001), dry deposition (Wang et al., 1998), particle sedimentation (Fairlie et al., 2007), and emissions (Keller et al., 2014).

The GEOS DAS meteorological fields used for input to GEOS-Chem are produced on a gnomonic cubed sphere grid (Putman and Lin, 2007) at a current horizontal resolution of C720 ($\sim 13 \text{ km} \times 13 \text{ km}$), with output provided operationally on a rectilinear grid at a resolution of $0.25^\circ \times 0.3125^\circ$. Currently, global oxidant-aerosol simulations with GEOS-Chem are effectively limited to $2^\circ \times 2.5^\circ$ resolution due to the prohibitive memory and time requirements of running a more finely-resolved simulation on a single node. In order to progress to finer resolutions, GEOS-Chem must be able to split the requirements for memory and computation across multiple nodes, and to ensure that communication between the different nodes is minimal and efficient. This is the role of GCHP.

2.2 GCHP v11-02c model architecture

The general software architecture of the GCHP model is shown in Figure 1. [Detailed description including a user's manual is available on the GCHP webpage of the GEOS-Chem website \(http://www.geos-chem.org\)](http://www.geos-chem.org). The GMAO-developed MAPL is included in the GCHP code download and is automatically built when compiling GCHP for the first time. MAPL initializes the model, establishes the atmospheric domain on each computational core, and handles model coordination and internal communication. Transport within and between each of the domains is calculated by the FV3 advection component. Within each atmospheric domain, local terms are calculated by a standard copy of the GEOS-Chem Classic code, embedded in the model as described by Long et al. (2015). [This copy of the GEOS-Chem code is identical to that used in GEOS-Chem Classic, such that all processes other than advection which are simulated in GCC are simulated identically in GCHP. GCHP v11-02c as presented here uses GEOS-Chem v11-02c. The embedded copy of GEOS-Chem in GCHP is compiled without OpenMP shared-memory parallelization, resulting in a pure MPI implementation.](#) Data input is handled through the External Data (ExtData) component, and output is handled through the History component. ExtData and History are structural components of MAPL (Long et al., 2015; Nielsen et al., 2017; Suarez et al., 2007).



15 **Figure 1. Connectivity of the major components of GCHP. The main time stepping loop is represented by the feedback loop from the model output back into the input.**

At initialization, a gridded representation of the atmosphere is generated by MAPL from user-specified input. GCHP can operate on any horizontal grid supported by MAPL as long as an appropriate advection scheme is available. Currently, the standard advection scheme in GCHP is the Putman and Lin FV3 scheme, which operates on a cubed sphere discretization, described in section 2.2. The initial state of the model is determined from a restart file, read by MAPL directly. During this stage, all relevant input data are also read into memory through the ExtData module. Data at any grid resolution are read

20

from NetCDF files in disk storage, and are regridded on the fly to the resolution at which the model is running. This allows data on either rectilinear latitude-longitude or gnomonic cubed sphere grids to be read in without requiring offline preprocessing. Data can be regridded using bilinear interpolation (used for wind fields), or first-order mass-conservative regridding (used for emissions and all other meteorological data). Conservative regridding is achieved using “tile files” generated analytically with the Tempest tool (Ullrich and Taylor, 2015). ~~Wind data are regridded by bilinear interpolation of the wind vector.~~ Additional regridding techniques are also available for special cases such as handling categorical (e.g. surface type) data. All constant fields are read in once, at the start of the simulation. For all time-varying fields, ExtData holds two samples in memory at all times: the previous sample (“left bracket”) and the upcoming sample (“right bracket”). All fields can either be held constant between samples, or smoothly interpolated between the two brackets.

10

Output is performed through the History component. Fields which are defined as “exports” within GCHP are tracked continuously by the History component. Any export can be requested by the user by adding it to an output collection in the HISTORY.rc input file, as either an instantaneous and/or time-averaged output. At each time step, the History component will acquire the current value of the field for each requested diagnostic and store it either at the native resolution or, if requested by the user, perform online regridding to a rectilinear latitude-longitude grid. This allows the user to decide the appropriate spatial and temporal resolution for their simulation output, independent of the resolution at which the simulation itself is conducted. All diagnostic quantities which are available in gridded form in GEOS-Chem Classic are automatically defined as exports in GCHP.

15

2.3 Grid discretization and transport


20 In GCHP the atmosphere is divided into independent atmospheric columns, with a subset of columns forming a single domain which is assigned to one of the computational cores. All local operations, such as chemistry, deposition, and emissions, are handled locally by components already present in the core GEOS-Chem code. The advection operator transfers mass between adjacent columns, requiring MPI-based data communication between them at domain boundaries. The amount and frequency of the communication depends on the chosen grid discretization and transport algorithm.

25 2.3.1 Grid discretization

GCHP inherits the equidistant gnomonic cubed sphere grid discretization used by the GEOS DAS (Putman and Lin, 2007). Cubed sphere grids split the surface of a sphere into six equal-sized faces. Each face is then subdivided into cells of approximately equal size, with each cell representing an atmospheric column. The equidistant gnomonic projection splits each cube edge into N equally-sized segments, connecting the opposing edges with great circle arcs in order to generate a regular mesh (see Figure 2). The grid resolution is referred to as CN , such as C48 for a grid with 48×48 atmospheric columns on each of the 6 faces. The grid cell spacing is approximately $10,000/N$ km, such that a C48 grid has a mean cell width of

30

~200 km. Each core is assigned a contiguous, rectangular set of columns on one of the 6 faces by MAPL, with the exact subdomain size determined based on the domain aspect ratio specified by the user at run time.



5 **Figure 2. Graphical description of the process used to generate a gnomonic cubed sphere grid. Sub-panels are numbered based on the textual description of the steps. The grids shown are C6 and, on the final frame, C24. Demonstration is available interactively at http://www.geos-chem.org/cubed_sphere.html.**

Cubed sphere grids offer several advantages over conventional rectilinear grids. The absolute cell size in a rectilinear grid decreases from the equator to the poles, for structural rather than scientific reasons, resulting in larger Courant-Friedrichs-Lewy (CFL) numbers at high latitudes. This reduces the minimum time step required for explicit Eulerian advection schemes to maintain stability. The problem can be mitigated by applying a semi-Lagrangian method when the CFL exceeds unity, at the expense of ~~additional computational overhead, inconsistency in technique, and difficulties in ensuring having to do non-physical mass conservation corrections-mass conservation~~. In an MPI environment, these issues also complicate domain decomposition for the purposes of distributing the grid between cores. The use of a semi-Lagrangian scheme results in tracer mass being transferred between grid cells which are not considered to be adjacent, increasing the size of the halo for each domain and therefore amount of communication necessary between cores. ~~Furthermore, any subdomain with an edge at a pole will be required to communicate with all other cores sharing the pole.~~

10
15

The cubed sphere grid helps to address these issues. The area ratio between the largest and smallest grid cells is ~ 2.3 , regardless of the resolution. There are no polar singularities, although advection across the edges and corners of the cube requires special considerations.

5

Vertically, the atmosphere is discretized into ~~a series of~~ hydrostatic, hybrid-sigma layers. The current GEOS DAS uses 72 layers ranging from the surface to 1 Pa at the upper edge.

2.3.2 Transport

Transport in GEOS-Chem is comprised of four operations: advection, moist convection, boundary-layer mixing, and aerosol settling. The latter three operations occur purely in-column and are unchanged between GCC and GCHP. However, advection must be grid-aware. Horizontal advection in GCHP is calculated on a layer-by-layer basis using the cubed sphere advection algorithm of Putman and Lin (2007). This algorithm is 4th-order accurate except at the six cube edges (2nd-order). Vertical advection is then calculated using a vertically-Lagrangian method (Lin, 2004). Prior to the advection step, each core requests ~~information-concentration data~~ from neighboring domains to ~~provide information on boundary conditions~~ fill the halo region. Advection is then calculated independently for each atmospheric domain.

15

Horizontal mass fluxes and CFL numbers are either supplied directly to the model or are calculated based on 3-hour average horizontal wind speed data and the instantaneous surface pressure at the start of the time step (time t). All fluxes are based on dry air mass and dry surface pressure. ~~_-~~ To ensure numerical stability, sub-stepping is implemented such that the internal advection timestep Δt is sub-divided into n number of sub-steps until the CFL is less than one. Horizontal advection is then performed n times. If the number of sub-steps n is greater than 1, changes to the air mass in each grid cell due to wind divergence are retained between sub-steps. The implied surface pressure resulting from changes in total column mass is also updated. However, horizontal mass fluxes are assumed to be constant over the time step, and no vertical remapping is performed between sub-steps. When mass fluxes have to be estimated offline from wind data, the simulated pressure can diverge from that in the meteorological archive (Jöckel et al., 2001). GCC solves this problem with the pressure fixer of Horowitz et al. (2003), which modifies calculated air mass fluxes to ensure the correct surface pressure tendency based on zonal totals. However, this approach corrupts the horizontal transport to some extent, and a pressure fixer has not yet been designed for transport on the cubed sphere. For archives where air mass fluxes are not explicitly available, including the meteorological data used for this work, GCHP defaults ~~instead to~~ using a simple global air mass correction ~~also applied to~~ tracers.

20
25
30

After the horizontal tracer advection loop is complete (time $t+\Delta t$), the total air mass in each vertical column will have changed, as will the vertical distribution. Vertical advection is calculated by remapping the deformed layers back to the

hydrostatic hybrid-eta grid defined by the surface pressure, as interpolated from the meteorological archive for the post-advection time ($t+\Delta t$). This ensures that the surface pressure accurately tracks that in the meteorological data.

2.4 Benchmarking

5 The standard benchmarking procedures applied to GEOS-Chem before each version release are also applied to GCHP, ensuring that the integrity of the model is maintained from version to version. Benchmarks involve a 1-year UCX (troposphere-stratosphere) oxidant-aerosol simulation with resolution of $4^\circ \times 5^\circ$ (GCC) or C48 (GHP), plus a 1-year simulation of the ^{222}Rn - ^{210}Pb - ^7Be system (Liu et al., 2001) for updates that may affect transport. Further documentation of benchmark procedures is available at <http://www.geos-chem.org>. Species concentrations and source/sink diagnostics from
10 the benchmark simulation are archived and compared to the previous model version and to selected climatological data. Results are inspected by the model developers and by the GEOS-Chem Steering Committee, which gives final approval. There are small differences between GCHP and GCC benchmarks ~~that we can relate~~due to differences in transport algorithm, but otherwise the two functionalities perform identically.

3 Model performance

15 We analyzed the performance of GCHP v11-02c by conducting simulations multiple grid resolutions (C24 to C180), each for a range of core counts (Table [42](#)). For low-resolution applications, performance is also compared to the maximum achievable performance using the GCC v11-02c shared-memory architecture. All simulations are for 1 month (July 2016) of troposphere-stratosphere oxidant-aerosol chemistry, including 206 species and 135 tracers, and using operational meteorological data from GEOS FP. The GCC v11-02c simulations use previously-regridded $2^\circ \times 2.5^\circ$ and $4^\circ \times 5^\circ$
20 meteorological fields, while the GCHP v11-02c simulations regrid the $0.25^\circ \times 0.3125^\circ$ fields to the cubed sphere on the fly through ExtData. A native-resolution cubed sphere GEOS data output stream is presently under development at GMAO and will benefit GCHP by reducing the need for regridding.

Table 24. Grid resolution and core counts for the performance test simulations.¹

GEOS-Chem implementation	Grid resolution	Number of grid cells	Number of cores used
GCHP	C24	250,000	6 – 216
GCHP	C48	1,000,000	6 – 540
GCHP	C90	3,500,000	12 – 540
GCHP	C180	14,000,000	90 – 540
GCC	4°×5°	240,000	6 – 30
GCC	2°×2.5°	940,000	6 – 30

All simulations were conducted on the Harvard Odyssey computational cluster. Both GCHP and GCC were compiled using the Intel Fortran compiler (v15.0.0), and MPI capabilities for GCHP were provided by OpenMPI (v1.10.3). Each node of the cluster has 32 Intel Broadwell 2.1 GHz cores sharing 128 GB of RAM, and all nodes are connected via Mellanox FDR Infiniband. Input and output data are stored using a Lustre parallel file system, accessible through the same Infiniband network fabric. All simulations were scheduled to enforce exclusive access to the nodes, preventing possible performance degradation due to sharing of node resources.

Figure 3 shows the total time taken to perform the simulation at each resolution, both in terms of wall time and in terms of the time per 1,000 simulated atmospheric columns in the model grid. Results using the conventional GEOS-Chem Classic (GCC) shared-memory platform at 4°×5° and 2°×2.5° resolutions are also shown for comparison.

¹ Both GCC and GCHP use v11-02c of the core GEOS-Chem code.




Figure 3. Wall time taken to perform 1-month GEOS-Chem simulations at different resolutions and with different numbers of cores. The upper panel shows the absolute time taken to complete each simulation at each resolution. The lower plot shows the wall time normalized by the number of atmospheric columns simulated at each resolution. Solid lines are for GCHP simulations (cubed sphere grids) and dashed lines are for GCC simulations (latitude-longitude grids). Grey lines on each plot show perfect scaling, corresponding to a 50% reduction in run time for each doubling of the number of cores.

At the lowest simulated resolution (C24), GCHP’s runtime exceeds that of GEOS-Chem Classic. This is predominantly due to overhead associated with file open operations, as the native-resolution meteorological data used to drive GCHP are opened and read independently for each field. This effect is clearly visible in the lower plot of Figure 3, where performance penalties due to this overhead result in significantly longer run times per 1,000 columns at C24 compared to C180. This can be addressed in the future through both structural changes and parallelization of the input.

After doubling the resolution to C48 ($2^\circ \times 2.5^\circ$), GCHP begins to out-perform GCC, with a reduced overall simulation time even at core counts which are currently accessible to GCC, despite the larger input requirements of GCHP. A 1-month simulation at C48 requires only 6 hours using 96 cores. GCHP scalability also improves as the model resolution increases. At C180, the reduction in simulation time for each doubling in the number of cores is approximately a factor of 1.6.


Two factors affect GCHP scalability: fixed costs, and overhead. Fixed costs are for operations which run on a fixed number of cores, regardless of the number of cores being dedicated to the simulation. Overhead is the need for additional

coordination and data communication between cores, which grows with the number of cores. Eventually the overhead of the additional cores exceeds the computational benefit, resulting in a performance plateau. This overhead includes one-off costs, such as the MPI interface initialization, which can be significant when running with a large number of cores but which can be reduced in relative terms by running longer simulations.

5

The scalability of each model component is shown in Figure 4. This shows the total time spent on each component at C48 and C180 resolution as a function of the total number of cores used, from 6 cores up to 540. We see that the dominant fixed cost at both C48 and C180 is input, which also dominates the overall cost for C48 with more than 48 cores. This is due to the serial nature of the current input code, overhead associated with file open operations, and the aforementioned use of native-resolution meteorological data for even low-resolution GCHP simulations. Output operations are a second fixed cost, being handled by a single core at all times. For these simulations, 22 3D fields were stored with hourly frequency, and output was a minor contributor to overall costs. Fixed costs can be converted into scalable costs by parallelizing the component in question, and this is a future work agenda.

10



15

Figure 4. Total wall time per component for a low (C48) and high (C180) resolution simulation. Simulations at C180 are limited to core counts of 90 or more across several nodes for the hardware used here, due to the high memory requirements of such high-resolution simulations.

Chemistry, advection, and convection all scale well with increasing core count. Chemistry is the most expensive process at both coarse (C48) and fine (C180) resolution, but has near-perfect scalability. Thus at C48 we see that input becomes the limiting process when the number of cores exceeds 48. Advection and other processes show more departure from perfect scalability, and may dominate the time requirement as the number of cores exceeds 600. The scalability of advection suffers from the additional communication overhead associated with reducing the domain size, as each domain must communicate a

20

larger proportion of its concentration data to its neighbors. More time is spent on communication relative to computation. However, wall time for advection does consistently fall with increasing core counts, an improvement compared to Long et al. (2015) where wall time increased as core counts exceeded 200 for a grid resolution equivalent to C48. We attribute this to the change from a latitude-longitude grid to the more scalable cubed sphere grid.

5

The remaining wall time is taken up by the “other” component, a mix of scalable and non-scalable processes. This includes the one-off cost of initializing the MPI interface, which grows non-linearly with the number of cores. At C180, these costs are still exceeded by scalable costs when running with 540 cores, so no plateau in performance is observed.

4 High resolution simulations with GCHP

10 The primary advantage of GCHP is the ability to perform high-resolution simulations of atmospheric modeling-chemistry at resolutions previously not available to the community. Figure 5 shows illustrative distributions of simulated ozone concentration at 4 km altitude ~~and aerosol optical depth (AOD)~~, simulated at C24 ($\approx 4 \times 5^\circ$) and C180 ($\approx 0.5^\circ \times 0.625^\circ$). Simulations were performed using 24 and 360 cores respectively. Results are also shown from simulations using GCC, at $4^\circ \times 5^\circ$ and $2^\circ \times 2.5^\circ$. Emissions are identical for all simulations.

15

Global-scale patterns in ozone concentration ~~and AOD~~ are not substantially affected by the increase in resolution, as these are determined by large-scale processes. The agreement between the two simulations illustrates the consistency of GCHP across scales. demonstrating the ability of GCHP to accurately simulate synoptic scale chemistry and transport at even low resolutions. However, increasing the horizontal resolution improves the ability of the model to capture the behavior of

20 intercontinental plumes (Eastham and Jacob, 2017). The consequences of this are visible in the ozone distributions over the Pacific and Atlantic. We also observe maxima in the coarse-resolution simulation which are not visible in the finer-resolution simulation, such as the peak in ozone concentration over Egypt. This suggests possible simulation biases at coarse resolution.

25 ~~Increasing the horizontal resolution also improves the ability of the model to resolve features at the scale of local air quality. This is especially evident in the simulated column AOD over northern India and Beijing.~~

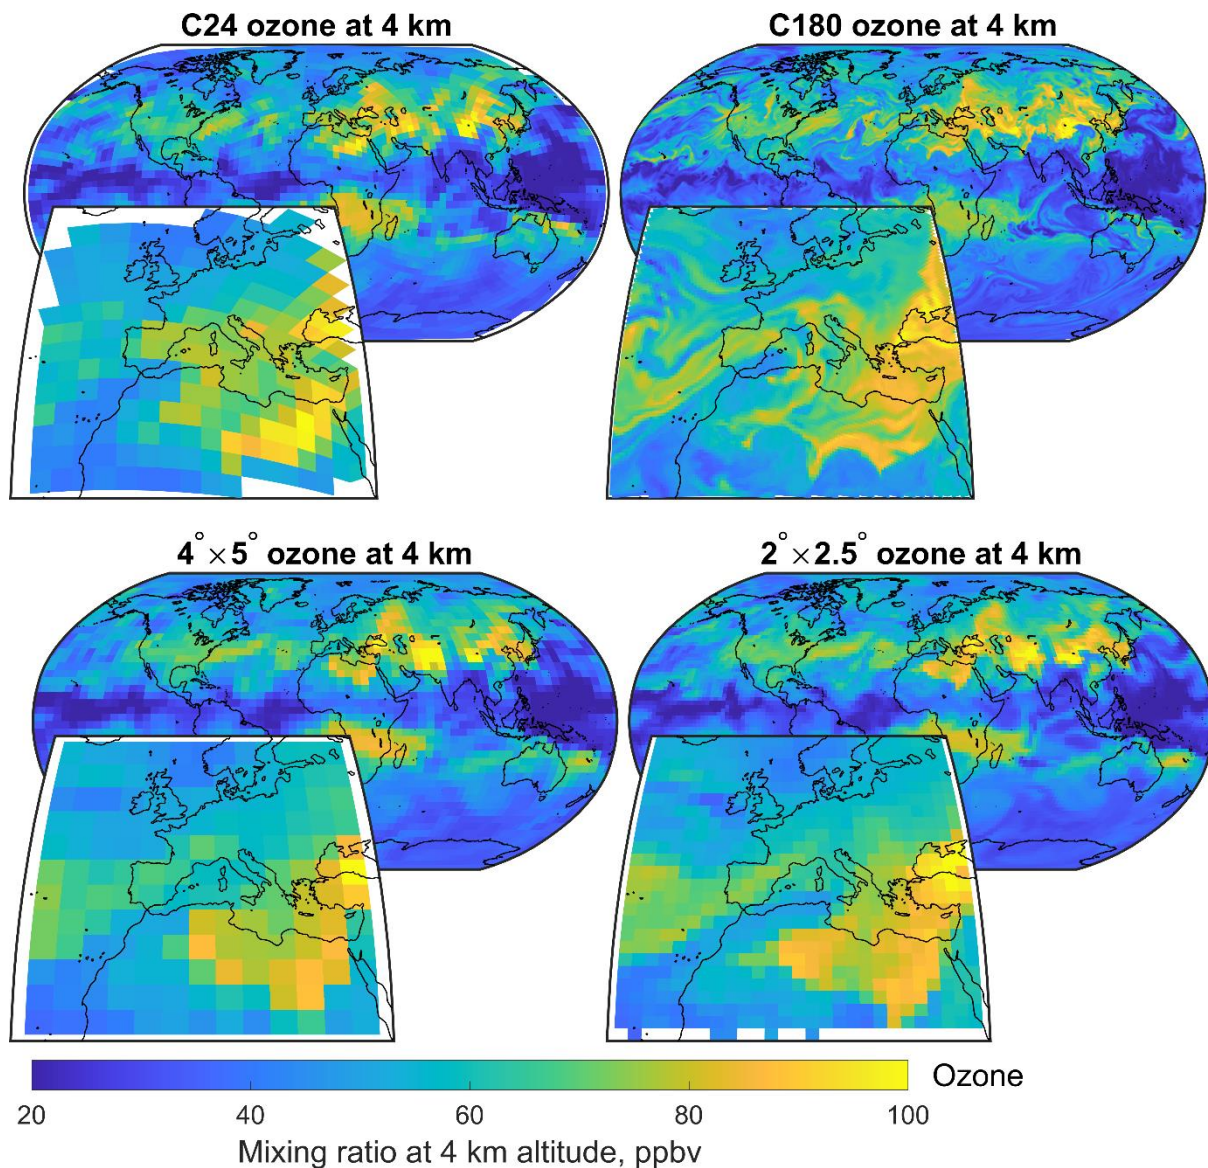
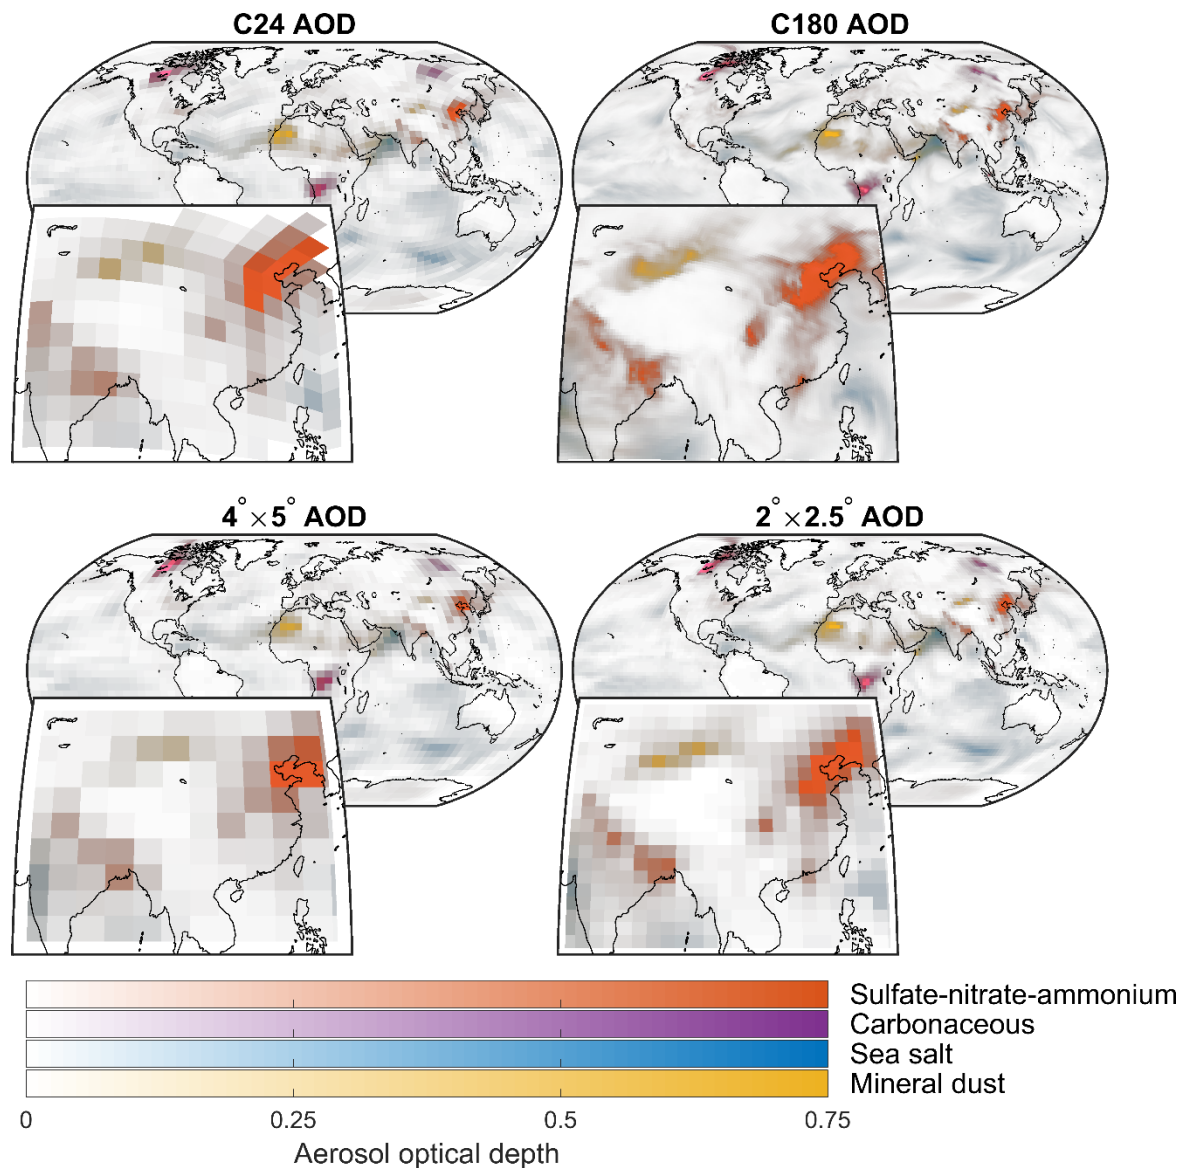


Figure 5. Simulated ozone concentrations at 4 km altitude for 23:00 on Model output from GCHP simulations at C24 (left) and C180 (right) for July 31st, 2016 after 1 month of initialization. The upper panels show 23:00 GMT ozone concentrations at 4 km altitude model output from GCHP simulations at C24 (left) and C180 (right), while the lower panels show model output from GCC simulations at $4^\circ \times 5^\circ$ and $2^\circ \times 2.5^\circ$, and the lower panels show daily average column aerosol optical depth (AOD) with each major aerosol class in GEOS-Chem represented using a different color. Calculated values in some regions exceeded the displayed limits. Zoom panels are also shown for ozone over Europe and AOD over East Asia.

Increasing the horizontal resolution also improves the ability of the model to resolve features at the scale of local air quality, as demonstrated by the plots of daily average aerosol optical depth (AOD) shown in Figure 6. This is especially evident in the simulated column AOD over northern India and Beijing. For both ozone and AOD, comparison to the results from GCC at $4^\circ \times 5^\circ$ shows that simulations with GCHP at coarse (C24) resolution are able to reproduce the same patterns, magnitude,

and variability as those observed in GCC. Some differences are observed, such as the region over Afghanistan which shows an elevated ozone mixing ratio in GCC compared to GCHP. However, the results from GCHP at C24 show the same patterns, magnitude, and variability as those at $4^\circ \times 5^\circ$.



5

Figure 6. As for Figure 5, but now showing daily-average aerosol optical depth for July 31st, 2016, after one month of initialization. Finally, we directly compare the simulated ozone data from GCHP and GCC using 2-D histograms. Figure 6 shows the ozone ratio in each $4^\circ \times 5^\circ$ grid cell, as simulated in GCC (horizontal axis) and as simulated in GCHP (vertical axis) at two different horizontal resolutions. All data are binned at a resolution of 1 ppbv ozone. In each case, GCHP results are

conservatively gridded to $4^\circ \times 5^\circ$ to provide a direct comparison, and data in the “polar cap” regions (top and bottom two $4^\circ \times 5^\circ$ latitude bands) are excluded. Fitting parameters are shown in white based on a reduced-major-axis (“geometric”) regression. Comparison of results at C24 to those at $4^\circ \times 5^\circ$ shows good agreement, with a correlation coefficient $r = 0.88$ and a slope of 0.95. At higher resolution, this agreement is slightly worsened ($r = 0.86$ and slope of 0.92), as smaller-scale processes and chemical non-linearity are resolved which could not be represented at the coarser resolution.

5

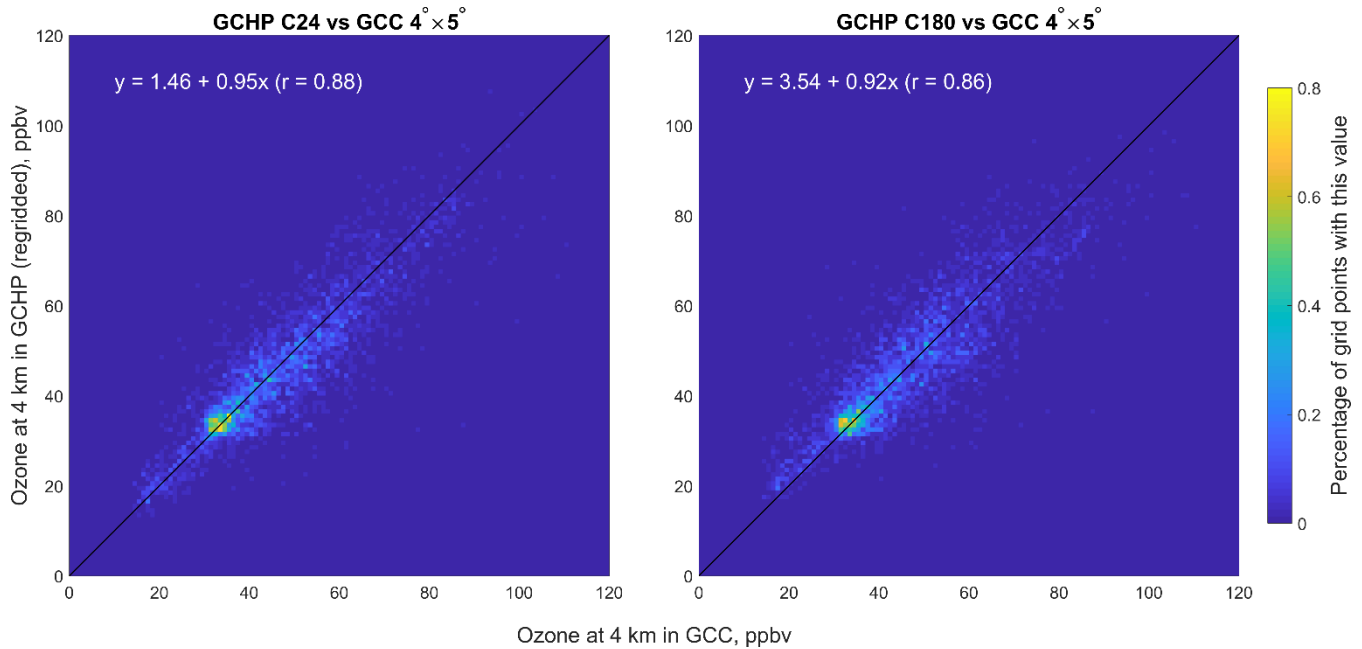


Figure 6. 2-D histograms of simulated ozone at 4 km altitude 23:00 on July 31st, 2016 after 1 month of initialization as calculated by GCC and GCHP. Each panel compares the simulated output from GCHP at a specific resolution with the simulated output from GCC at $4^\circ \times 5^\circ$. Data are binned at a resolution of 1 ppbv. GCHP data are conservatively regridded to $4^\circ \times 5^\circ$ prior to comparison. Data in the top and bottom two latitude bands are excluded, as GCC averages these points into two “polar caps”.

10

5 Summary

Models of atmospheric chemistry have grown continuously in resolution and complexity over the past decades to take advantage of increasing computational resources. The GEOS-Chem High Performance model (GCHP) is a next step in this growth, enabling the widely-used GEOS-Chem chemical transport model to exploit the computational speed and memory capacity of massively parallel architectures. In this manner we can achieve routine simulation of global stratosphere-troposphere oxidant-aerosol chemistry at unprecedented resolution and detail.

15

Detailed documentation of GCHP including a user’s manual is available on the GCHP website (URL reference). GCHP incorporates ~~an unchanged copy of~~ the existing GEOS-Chem shared-memory code into an ESMF-based framework

20

- (MAPL), enabling GEOS-Chem to be run in a distributed memory framework across multiple nodes while retaining all the features of the high-fidelity global chemical simulation. In addition to a new model framework, GCHP replaces the conventional rectilinear latitude-longitude grid with ~~a the~~ gnomonic cubed sphere grid of the NASA GEOS meteorological data used as input to GEOS-Chem. This provides greater computational accuracy and efficiency for transport calculations while removing an additional restriction on scalability. GCHP ~~has been tested and shown is scalable to scale~~ performs with high computational scalability from six cores up to at least 540 cores, completing a 1-month simulation of oxidant-aerosol chemistry in the troposphere and stratosphere (206 active species, 135 tracers) at a global resolution of C180 (~0.5°×0.625°) in 24 hours.
- 10 GCHP also provides a mechanism for ongoing improvement of modeling capability. With the base GEOS-Chem model, GCHP, and the GMAO GEOS atmospheric data assimilation system now all using an identical copy of the grid-independent GEOS-Chem code, GCHP closes the loop between online and offline modelers, allowing seamless propagation of model and framework improvements between all three. Future development opportunities range from improved parallelism in input operations to the direct ingestion of archived mass fluxes to further improve transport calculation accuracy.

15 **Code availability**

- GCHP has been openly available as part of the GEOS-Chem code since beta version release v11-02b in June 2017, and ~~will be was~~ part of the v11-02 public release in March 2018. Complete documentation and access to the GCHP code can be found at <http://www.geos-chem.org>. GCHP is an added functionality for GEOS-Chem users, who can choose to use either GCC or GCHP from the same code download. Both GCC and GCHP functionalities will be maintained in the standard GEOS-Chem model for the foreseeable future, recognizing that many users may not have access to the resources needed to use GCHP. For this work, GCHP v11-02c was used, a copy of which has been permanently archived (DOI: 10.5281/zenodo.1290835).
- 20

Acknowledgements

- This work was supported by the NASA Atmospheric Composition Modeling and Analysis Program (ACMAP) and the NASA Modeling, Analysis and Prediction (MAP) program. Sebastian D. Eastham was supported by the NOAA Climate and Global Change Postdoctoral Fellowship Program, administered by UCAR's Visiting Scientist Programs. Sebastian D. Eastham was also supported by a Harvard University Center for the Environment (HUCE) Postdoctoral Fellowship. The GEOS FP data used in this study/project were provided by the Global Modeling and Assimilation Office (GMAO) at NASA Goddard Space Flight Center. The computations were run on the Odyssey cluster supported by the FAS Division of Science, Research Computing Group at Harvard University. We are grateful to Compute Canada for hosting the Data Portal for GEOS-Chem to store and make available the GEOS meteorological fields used here. The authors would also like to thank
- 30

Kevin Bowman for providing analytical insight and additional computational resources with which to perform early testing of the GCHP model. Finally, we would like to thank Junwei Xu for assistance in processing and archiving meteorological data for input to the model.

Author contribution

- 5 MSL and AM designed the initial code infrastructure
RVM and DJJ provided project oversight and top-level design
SDE, MSL, CAK, EWL, MY, RY, JZ, MT, AT, TC, JK, WP, and CJL performed code development
SDE, EL, RVM, CL, and CJL adapted the meteorological archive
SDE, MSL, EWL, and JZ ran and debugged benchmark simulations
- 10 SDE, MSL, RY, and JZ performed scalability testing analysis
SDE, MSL, CAK, RY, EWL, JZ, RVM, and DJJ wrote the manuscript
All authors contributed to manuscript editing and revisions

References

- 15 Bey, I., Jacob, D. J., Yantosca, R. M., Logan, J. a., Field, B. D., Fiore, A. M., Li, Q., Liu, H. Y., Mickley, L. J. and Schultz, M. G.: Global modeling of tropospheric chemistry with assimilated meteorology: Model description and evaluation, *J. Geophys. Res.*, 106(D19), 23073, 2001.
- Brasseur, G. P. and Jacob, D. K.: *Modeling of Atmospheric Chemistry*, Cambridge University Press., 2017.
- Eastham, S. D. and Jacob, D. J.: Limits on the ability of global Eulerian models to resolve intercontinental transport of chemical plumes, *Atmos. Chem. Phys.*, 17, 2543–2553, 2017.
- 20 Eastham, S. D., Weisenstein, D. K. and Barrett, S. R. H.: Development and evaluation of the unified tropospheric–stratospheric chemistry extension (UCX) for the global chemistry-transport model GEOS-Chem, *Atmos. Environ.*, 89, 52–63, 2014.
- Fairlie, D. T., Jacob, D. J. and Park, R. J.: The impact of transpacific transport of mineral dust in the United States, *Atmos. Environ.*, 41(6), 1251–1266, 2007.
- 25 Hill, C., DeLuca, C., Balaji, Suarez, M. and Silva, A. D.: The architecture of the Earth System Modeling Framework, *Computing in Science Engineering*, 6(1), 18–28, 2004.
- Horowitz, L. W., Walters, S., Mauzerall, D. L., Emmons, L. K., Rasch, P. J., Granier, C., Tie, X., Lamarque, J.-F., Schultz, M. G., Tyndall, G. S., Orlando, J. J. and Brasseur, G. P.: A global simulation of tropospheric ozone and related tracers: Description and evaluation of MOZART, version 2, *J. Geophys. Res.*, 108(D24), doi:10.1029/2002JD002853, 2003.
- 30

- Jöckel, P., von Kuhlmann, R., Lawrence, M. G., Steil, B., Brenninkmeijer, C. A. M., Crutzen, P. J., Rasch, P. J. and Eaton, B.: On a fundamental problem in implementing flux-form advection schemes for tracer transport in 3-dimensional general circulation and chemistry transport models, *Q.J.R. Meteorol. Soc.*, 127(573), 1035–1052, 2001.
- 5 Keller, C., Auer, B., Hu, L., Jacob, D., Long, M., Nielsen, E., Pawson, S., da Silva, A. and Thompson, M.: High-resolution GEOS-5 nature run with tropospheric chemistry, 2017.
- Keller, C. A., Long, M. S., Yantosca, R. M., Da Silva, A. M., Pawson, S. and Jacob, D. J.: HEMCO v1.0: a versatile, ESMF-compliant component for calculating emissions in atmospheric models, *Geoscientific Model Development*, 7(4), 1409–1417, 2014.
- 10 Lin, J.-T. and McElroy, M. B.: Impacts of boundary layer mixing on pollutant vertical profiles in the lower troposphere: Implications to satellite remote sensing, *Atmos. Environ.*, 44(14), 1726–1739, 2010.
- Lin, S.-J.: A “Vertically Lagrangian” Finite-Volume Dynamical Core for Global Models, *Mon. Weather Rev.*, 132(10), 2293–2307, 2004.
- Liu, H., Jacob, D. J., Bey, I. and Yantosca, R. M.: Constraints from ²¹⁰Pb and ⁷Be on wet deposition and transport in a global three-dimensional chemical tracer model driven by assimilated meteorological fields, *J. Geophys. Res.*, 106(D11), 12109–12128, 2001.
- 15 Long, M. S., Yantosca, R., Nielsen, J. E., Keller, C. A., da Silva, A., Sulprizio, M. P., Pawson, S. and Jacob, D. J.: Development of a grid-independent GEOS-Chem chemical transport model (v9-02) as an atmospheric chemistry module for Earth system models, *Geoscientific Model Development*, 8(3), 595–602, 2015.
- 20 National Research Council: A National Strategy for Advancing Climate Modeling, The National Academies Press, Washington, DC., 2012.
- Nielsen, J. E., Pawson, S., Molod, A., Auer, B., da Silva, A. M., Douglass, A. R., Duncan, B., Liang, Q., Manyin, M., Oman, L. D., Putman, W., Strahan, S. E. and Wargan, K.: Chemical Mechanisms and Their Applications in the Goddard Earth Observing System (GEOS) Earth System Model, *J. Adv. Model. Earth Syst.*, 9(8), 3019–3044, 25 2017.
- Prather, M. J.: Fast-JX v7.0a, 2012. https://www.ess.uci.edu/group/prather/scholar_software
- Putman, W. M. and Lin, S. J.: Finite-volume transport on various cubed-sphere grids, *J. Comput. Phys.*, 227(1), 55–78, 2007.
- Sherwen, T., Evans, M. J., Carpenter, L. J., Andrews, S. J., Lidster, R. T., Dix, B., Koenig, T. K., Sinreich, R., Ortega, I., Volkamer, R., Saiz-Lopez, A., Prados-Roman, C., Mahajan, A. S. and Ordóñez, C.: Iodine’s impact on tropospheric oxidants: a global model study in GEOS-Chem, *Atmos. Chem. Phys.*, 16(2), 1161–1186, 2016.
- Suarez, M., Trayanov, A., Hill, C., Schopf, P. and Vikhliaev, Y.: MAPL: A High-level Programming Paradigm to Support More Rapid and Robust Encoding of Hierarchical Trees of Interacting High-performance Components, in *Proceedings of the 2007 Symposium on Component and Framework Technology in High-performance and Scientific Computing*, pp. 11–20, ACM, New York, NY, USA., 2007.
- 35

Ullrich, P. A. and Taylor, M. A.: Arbitrary-Order Conservative and Consistent Remapping and a Theory of Linear Maps: Part I, *Mon. Weather Rev.*, 143(6), 2419–2440, 2015.

Wang, Y., Jacob, D. J. and Logan, J. A.: Global simulation of tropospheric O₃-NO_x-hydrocarbon chemistry: 1. Model formulation, *J. Geophys. Res.*, 103(D9), 10713–10725, 1998.

- 5 Wu, S., Mickley, L. J., Jacob, D. J., Logan, J. A., Yantosca, R. M. and Rind, D.: Why are there large differences between models in global budgets of tropospheric ozone?, *J. Geophys. Res.*, 112(D5), D05302, 2007.

GEOS-Chem High Performance (GCHP v11-02c): A next-generation implementation of the GEOS-Chem chemical transport model for massively parallel applications

5 Sebastian D. Eastham^{1,2*}, Michael S. Long², Christoph A. Keller^{3,4}, Elizabeth Lundgren², Robert M. Yantosca², Jiawei Zhuang², Chi Li⁵, Colin J. Lee⁵, Matthew Yannetti², Benjamin M. Auer^{3,6}, Thomas L. Clune³, Jules Kouatchou^{3,6}, William M. Putman³, Matthew A. Thompson^{3,6}, Atanas L. Trayanov^{3,6}, Andrea M. Molod³, Randall V. Martin^{5,7}, Daniel J. Jacob²

¹Laboratory for Aviation and the Environment, Massachusetts Institute of Technology

²John A. Paulson School of Engineering and Applied Sciences, Harvard University

10 ³NASA Global Modeling and Assimilation Office

⁴Universities Space Research Association, Columbia, MD, USA

⁵Department of Physics and Atmospheric Science, Dalhousie University

⁶Science Systems and Applications, Inc., Lanham, MD, USA

⁷Smithsonian Astrophysical Observatory, Harvard-Smithsonian Center for Astrophysics

15 *Correspondence to:* Sebastian D. Eastham (seastham@mit.edu)

Abstract. Global modeling of atmospheric chemistry is a grand computational challenge because of the need to simulate large coupled systems of ~100-1000 chemical species interacting with transport on all scales. Off-line chemical transport models (CTMs), where the chemical continuity equations are solved using meteorological data as input, have usability advantages and are important vehicles for developing atmospheric chemistry knowledge that can then be transferred to Earth system models. However, they have generally not been designed to take advantage of massively parallel computing architectures. Here we develop such a high-performance capability (GCHP) for GEOS-Chem, a CTM driven by meteorological data from the NASA Goddard Earth Observation System (GEOS) and used by hundreds of research groups worldwide. GCHP is a grid-independent implementation of GEOS-Chem using the Earth System Modeling Framework (ESMF) that permits the same standard model to operate in a distributed-memory framework for massive parallelization. GCHP also allows GEOS-Chem to take advantage of the native GEOS cubed-sphere grid for greater accuracy and computational efficiency in simulating transport. GCHP enables GEOS-Chem simulations to be conducted with high computational scalability up to at least 500 cores, so that global simulations of stratosphere-troposphere oxidant-aerosol chemistry at C180 spatial resolution ($\sim 0.5^\circ \times 0.625^\circ$) or finer become routinely feasible.

20
25

1 Introduction

30 Atmospheric chemistry models are used to address a wide range of problems related to climate forcing, air quality, and atmospheric deposition. Simulations of oxidant and aerosol chemistry involve hundreds of chemically interacting species,

coupled to transport on all scales. The computational demands are considerable, which has limited the inclusion of atmospheric chemistry in climate models (National Research Council, 2012). Off-line chemical transport models (CTMs), where meteorology is provided as input data from a parent global climate model (GCM) or atmospheric data assimilation system (DAS), are frequently used for reasons of simplicity, reproducibility, and ability to focus on chemical processes. The global GEOS-Chem CTM originally described by Bey et al. (2001), using meteorological input from the Goddard Earth Observation System (GEOS) DAS of the NASA Global Modeling and Assimilation Office (GMAO), is used by hundreds of atmospheric chemistry research groups worldwide (<http://www.geos-chem.org>). Increasing computational resources in the form of massively parallel architectures can allow GEOS-Chem users to explore more complex problems at higher grid resolutions, but this requires re-engineering of the model to take advantage of these architectures. Here we describe a high-performance version of GEOS-Chem (GCHP) engineered for this purpose, and we demonstrate its ability to access a new range of capability and scales for global atmospheric chemistry modeling.

Table 1. Acronyms used in this paper.

Acronym	Description
AOD	Aerosol Optical Depth
CFL	Courant-Friedrichs-Lewy number
CTM	Chemical Transport Model
DAS	Data Assimilation System
ESMF	Earth System Modeling Framework
FV3	Finite Volume advection on the Cubed sphere
GCC	GEOS-Chem Classic
GCHP	GEOS-Chem High Performance
GCM	Global Climate Model
GEOS	Goddard Earth Observation System
GMAO	NASA Global Modeling and Assimilation Office
HEMCO	Harvard-NASA Emissions COmponent
MAPL	Modeling and Analysis Prediction Layer
MPI	Message Passing Interface
UCX	Unified Chemistry eXtension

15

The original GEOS-Chem CTM (“GEOS-Chem Classic”, or GCC) was designed for shared-memory (OpenMP) parallelization. Detailed description of the model including a user’s manual is available on the GEOS-Chem website (<http://www.geos-chem.org>). Computation is distributed over a number of cores on a single node, with data held in shared

arrays. But recent growth in computational power has taken the form of massively parallel networked systems, where additional computational power is achieved by increasing the number of identical nodes rather than by improving the nodes themselves. This has placed a restriction on growth in the problem size and complexity which can be solved by a single instance of GCC. To take advantage of massively parallel architectures, a new framework is needed which allows GEOS-Chem to use a distributed-memory model, where the computation is distributed across multiple coordinated nodes using a Message Passing Interface (MPI) implementation such as MVAPICH2 or OpenMPI.

An important first step in this evolution was the integration of GEOS-Chem as the online chemistry component within the GEOS DAS (Long et al., 2015). In order to ensure that the online and offline versions of GEOS-Chem were identical, GCC was modified to use the exact same code in the independent CTM and in the DAS. Major modifications were required to make GEOS-Chem grid-independent and compatible with the GEOS Modeling and Analysis Prediction Layer (MAPL) (Suarez et al., 2007), an Earth System Modeling Framework (ESMF) (Hill et al., 2004) based software layer which handles communication between different components of the GEOS DAS. The GEOS-Chem code was adapted to accept an arbitrarily-sized horizontal set of atmospheric columns, with no requirements regarding adjacency of the columns or overall coverage of any particular set. All these changes were made “under the hood” in the standard GEOS-Chem Code. When GEOS-Chem is run as GCC, the set of columns is designated as a single block which covers the entire globe or a subset in a nested domain, and parallelization is achieved by internally running parallel loops over the columns. When GEOS-Chem is run as part of GEOS, MAPL internally splits the atmosphere into smaller domains, each of which contains a different set of atmospheric columns. These domains can then be distributed across multiple nodes, exploiting massively parallel architectures. As a result of these changes, the same GEOS-Chem code can now be run either as a stand-alone, shared-memory offline CTM, or as a GCM component in the massively parallel, distributed GEOS DAS. Any improvement in chemical modeling developed for the offline CTM is thus immediately available in the GEOS DAS version, which never becomes out of date and remains referenceable to the current version of GEOS-Chem.

In this work we take the next logical step of developing GCHP as a distributed-memory, MAPL-based implementation of the GEOS-Chem CTM. GCHP uses an identical copy of the GEOS-Chem Classic (GCC) code to provide the same high-fidelity atmospheric chemical simulation capabilities, allowing users to switch between GCC and GCHP implementations with confidence that they are using the same model. The exact same internal code is used in GCC shared-memory and GCHP distributed-memory applications. This closes the development loop between online and offline modeling. By sharing infrastructure code between GCHP and GEOS in the form of MAPL, offline modelers can now take advantage of modeling advances which originate in the online model in the same way that GEOS benefits from advances in chemical modeling developed in the GEOS-Chem CTM (Nielsen et al., 2017). By way of example, GEOS was recently able to conduct a full-year 13-km resolution “nature run” with the current standard version of GEOS-Chem tropospheric chemistry (Hu et al.,

2018). In return, GCHP incorporates the more efficient cubed sphere grid and FV3 advection code present in GEOS, and is capable of directly ingesting GEOS output in its native cubed sphere format.

2 Model description

2.1 Overview

5 Atmospheric chemistry models such as GEOS-Chem solve the 3-D chemical continuity equations for an ensemble of m coupled chemical species (Brasseur and Jacob, 2017). The continuity equation for the number density n_i [molecules cm^{-3}] of species i is expressed as

$$\frac{\partial n_i}{\partial t} = -\nabla \cdot (n_i \mathbf{v}) + s_i \quad (1)$$

where \mathbf{v} is the velocity vector [m s^{-1}], and s_i is the local net production and loss of species i [molecules $\text{cm}^{-3} \text{s}^{-1}$]. In CTMs, \mathbf{v} is provided by archived output from a parent GCM or DAS, with subgrid-scale parameterized transport statistics (boundary layer mixing, deep convection) as additional CTM transport terms in Eqn. 1. From a computational standpoint, the local term s_i is grid-independent. However, the transport terms are grid-aware, as they move material between grid points. In GEOS-Chem, the atmosphere is split into independent columns, with each column made up of a number of discrete grid points (Long et al., 2015). Vertical processes (boundary layer mixing, deep convection) are then considered to be local in the sense that they are calculated independently for each column. In each column simulated by GEOS-Chem, the local term computes
15 chemical evolution with a unified tropospheric-stratospheric mechanism (Eastham et al., 2014; Sherwen et al., 2016), convective transport (Wu et al., 2007), boundary layer mixing (Lin and McElroy, 2010), radiative transfer and photolysis (Prather, 2012), wet scavenging (Liu et al., 2001), dry deposition (Wang et al., 1998), particle sedimentation (Fairlie et al., 2007), and emissions (Keller et al., 2014).

20 The GEOS DAS meteorological fields used for input to GEOS-Chem are produced on a gnomonic cubed sphere grid (Putman and Lin, 2007) at a current horizontal resolution of C720 ($\sim 13 \text{ km} \times 13 \text{ km}$), with output provided operationally on a rectilinear grid at a resolution of $0.25^\circ \times 0.3125^\circ$. Currently, global oxidant-aerosol simulations with GEOS-Chem are effectively limited to $2^\circ \times 2.5^\circ$ resolution due to the prohibitive memory and time requirements of running a more finely-resolved simulation on a single node. In order to progress to finer resolutions, GEOS-Chem must be able to split the
25 requirements for memory and computation across multiple nodes, and to ensure that communication between the different nodes is minimal and efficient. This is the role of GCHP.

2.2 GCHP v11-02c model architecture

The general software architecture of the GCHP model is shown in Figure 1. Detailed description including a user's manual is available on the GCHP webpage of the GEOS-Chem website (<http://www.geos-chem.org>). The GMAO-developed MAPL is

included in the GCHP code download and is automatically built when compiling GCHP for the first time. MAPL initializes the model, establishes the atmospheric domain on each computational core, and handles model coordination and internal communication. Transport within and between each of the domains is calculated by the FV3 advection component. Within each atmospheric domain, local terms are calculated by a standard copy of the GEOS-Chem Classic code, embedded in the model as described by Long et al. (2015). This copy of the GEOS-Chem code is identical to that used in GEOS-Chem Classic, such that all processes other than advection which are simulated in GCC are simulated identically in GCHP. GCHP v11-02c as presented here uses GEOS-Chem v11-02c, The embedded copy of GEOS-Chem in GCHP is compiled without OpenMP shared-memory parallelization, resulting in a pure MPI implementation. Data input is handled through the External Data (ExtData) component, and output is handled through the History component. ExtData and History are structural components of MAPL (Long et al., 2015; Nielsen et al., 2017; Suarez et al., 2007).

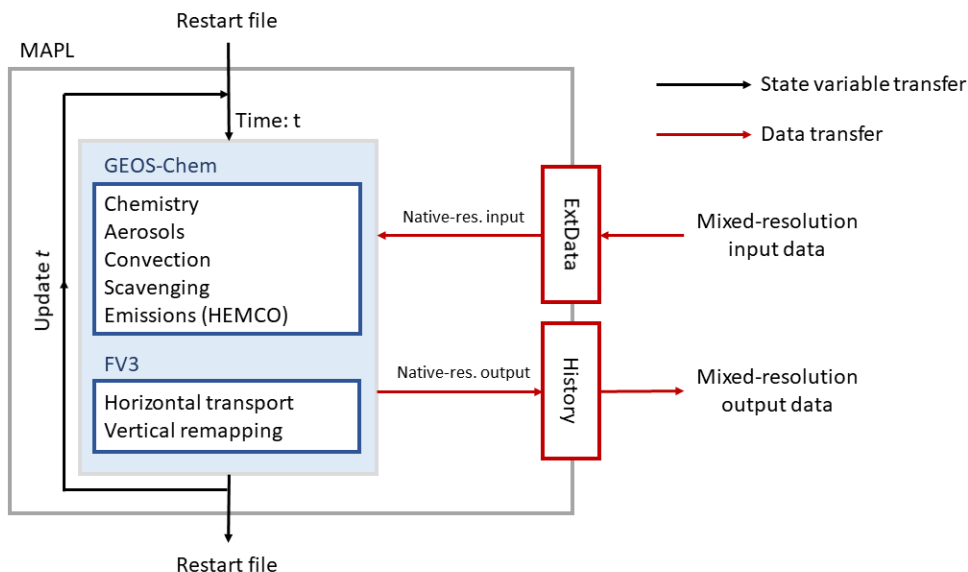


Figure 1. Connectivity of the major components of GCHP. The main time stepping loop is represented by the feedback loop from the model output back into the input.

At initialization, a gridded representation of the atmosphere is generated by MAPL from user-specified input. GCHP can operate on any horizontal grid supported by MAPL as long as an appropriate advection scheme is available. Currently, the standard advection scheme in GCHP is the Putman and Lin FV3 scheme, which operates on a cubed sphere discretization, described in section 2.2. The initial state of the model is determined from a restart file, read by MAPL directly. During this stage, all relevant input data are also read into memory through the ExtData module. Data at any grid resolution are read from NetCDF files in disk storage, and are regridded on the fly to the resolution at which the model is running. This allows data on either rectilinear latitude-longitude or gnomonic cubed sphere grids to be read in without requiring offline preprocessing. Data can be regridded using bilinear interpolation (used for wind fields), or first-order mass-conservative regridding (used for emissions and all other meteorological data). Conservative regridding is achieved using “tile files”

generated analytically with the Tempest tool (Ullrich and Taylor, 2015). Additional regridding techniques are also available for special cases such as handling categorical (e.g. surface type) data. All constant fields are read in once, at the start of the simulation. For all time-varying fields, ExtData holds two samples in memory at all times: the previous sample (“left bracket”) and the upcoming sample (“right bracket”). All fields can either be held constant between samples, or smoothly interpolated between the two brackets.

Output is performed through the History component. Fields which are defined as “exports” within GCHP are tracked continuously by the History component. Any export can be requested by the user by adding it to an output collection in the HISTORY.rc input file, as either an instantaneous and/or time-averaged output. At each time step, the History component will acquire the current value of the field for each requested diagnostic and store it either at the native resolution or, if requested by the user, perform online regridding to a rectilinear latitude-longitude grid. This allows the user to decide the appropriate spatial and temporal resolution for their simulation output, independent of the resolution at which the simulation itself is conducted. All diagnostic quantities which are available in gridded form in GEOS-Chem Classic are automatically defined as exports in GCHP.

2.3 Grid discretization and transport

In GCHP the atmosphere is divided into independent atmospheric columns, with a subset of columns forming a single domain which is assigned to one of the computational cores. All local operations, such as chemistry, deposition, and emissions, are handled locally by components already present in the core GEOS-Chem code. The advection operator transfers mass between adjacent columns, requiring MPI-based data communication between them at domain boundaries. The amount and frequency of the communication depends on the chosen grid discretization and transport algorithm.

2.3.1 Grid discretization

GCHP inherits the equidistant gnomonic cubed sphere grid discretization used by the GEOS DAS (Putman and Lin, 2007). Cubed sphere grids split the surface of a sphere into six equal-sized faces. Each face is then subdivided into cells of approximately equal size, with each cell representing an atmospheric column. The equidistant gnomonic projection splits each cube edge into N equally-sized segments, connecting the opposing edges with great circle arcs in order to generate a regular mesh (see Figure 2). The grid resolution is referred to as CN , such as C48 for a grid with 48×48 atmospheric columns on each of the 6 faces. The grid cell spacing is approximately $10,000/N$ km, such that a C48 grid has a mean cell width of ~ 200 km. Each core is assigned a contiguous, rectangular set of columns on one of the 6 faces by MAPL, with the exact subdomain size determined based on the domain aspect ratio specified by the user at run time.




Figure 2. Graphical description of the process used to generate a gnomonic cubed sphere grid. Sub-panels are numbered based on the textual description of the steps. The grids shown are C6 and, on the final frame, C24. Demonstration is available interactively at http://www.geos-chem.org/cubed_sphere.html.

5 Cubed sphere grids offer several advantages over conventional rectilinear grids. The absolute cell size in a rectilinear grid decreases from the equator to the poles, for structural rather than scientific reasons, resulting in larger Courant-Friedrichs-Lewy (CFL) numbers at high latitudes. This reduces the minimum time step required for explicit Eulerian advection schemes to maintain stability. The problem can be mitigated by applying a semi-Lagrangian method when the CFL exceeds unity, at the expense of having to do non-physical mass conservation corrections. In an MPI environment, these issues also

10 complicate domain decomposition for the purposes of distributing the grid between cores. The use of a semi-Lagrangian scheme results in tracer mass being transferred between grid cells which are not considered to be adjacent, increasing the size of the halo for each domain and therefore amount of communication necessary between cores.

The cubed sphere grid helps to address these issues. The area ratio between the largest and smallest grid cells is ~ 2.3 ,

15 regardless of the resolution. There are no polar singularities, although advection across the edges and corners of the cube requires special considerations.

Vertically, the atmosphere is discretized into hydrostatic, hybrid-sigma layers. The current GEOS DAS uses 72 layers ranging from the surface to 1 Pa at the upper edge.

2.3.2 Transport

5 Transport in GEOS-Chem is comprised of four operations: advection, moist convection, boundary-layer mixing, and aerosol settling. The latter three operations occur purely in-column and are unchanged between GCC and GCHP. However, advection must be grid-aware. Horizontal advection in GCHP is calculated on a layer-by-layer basis using the cubed sphere advection algorithm of Putman and Lin (2007). This algorithm is 4th-order accurate except at the six cube edges (2nd-order). Vertical advection is then calculated using a vertically-Lagrangian method (Lin, 2004). Prior to the advection step, each core
10 requests concentration data from neighboring domains to fill the halo region. Advection is then calculated independently for each atmospheric domain.

Horizontal mass fluxes and CFL numbers are either supplied directly to the model or are calculated based on 3-hour average horizontal wind speed data and the instantaneous surface pressure at the start of the time step (time t). All fluxes are based on
15 dry air mass and dry surface pressure. To ensure numerical stability, sub-stepping is implemented such that the internal advection timestep Δt is sub-divided into n number of sub-steps until the CFL is less than one. Horizontal advection is then performed n times. If the number of sub-steps n is greater than 1, changes to the air mass in each grid cell due to wind divergence are retained between sub-steps. The implied surface pressure resulting from changes in total column mass is also updated. However, horizontal mass fluxes are assumed to be constant over the time step, and no vertical remapping is
20 performed between sub-steps. When mass fluxes have to be estimated offline from wind data, the simulated pressure can diverge from that in the meteorological archive (Jöckel et al., 2001). GCC solves this problem with the pressure fixer of Horowitz et al. (2003), which modifies calculated air mass fluxes to ensure the correct surface pressure tendency based on zonal totals. However, this approach corrupts the horizontal transport to some extent, and a pressure fixer has not yet been designed for transport on the cubed sphere. For archives where air mass fluxes are not explicitly available, including the
25 meteorological data used for this work, GCHP defaults to using a simple global air mass correction.

After the horizontal tracer advection loop is complete (time $t+\Delta t$), the total air mass in each vertical column will have changed, as will the vertical distribution. Vertical advection is calculated by remapping the deformed layers back to the hydrostatic hybrid-eta grid defined by the surface pressure, as interpolated from the meteorological archive for the post-
30 advection time ($t+\Delta t$). This ensures that the surface pressure accurately tracks that in the meteorological data.

2.4 Benchmarking

The standard benchmarking procedures applied to GEOS-Chem before each version release are also applied to GCHP, ensuring that the integrity of the model is maintained from version to version. Benchmarks involve a 1-year UCX (troposphere-stratosphere) oxidant-aerosol simulation with resolution of $4^{\circ}\times 5^{\circ}$ (GCC) or C48 (GHP), plus a 1-year simulation of the ^{222}Rn - ^{210}Pb - ^7Be system (Liu et al., 2001) for updates that may affect transport. Further documentation of benchmark procedures is available at <http://www.geos-chem.org>. Species concentrations and source/sink diagnostics from the benchmark simulation are archived and compared to the previous model version and to selected climatological data. Results are inspected by the model developers and by the GEOS-Chem Steering Committee, which gives final approval. There are small differences between GCHP and GCC benchmarks due to differences in transport algorithm, but otherwise the two functionalities perform identically.

3 Model performance

We analyzed the performance of GCHP v11-02c by conducting simulations multiple grid resolutions (C24 to C180), each for a range of core counts (Table 2). For low-resolution applications, performance is also compared to the maximum achievable performance using the GCC v11-02c shared-memory architecture. All simulations are for 1 month (July 2016) of troposphere-stratosphere oxidant-aerosol chemistry, including 206 species and 135 tracers, and using operational meteorological data from GEOS FP. The GCC v11-02c simulations use previously-regridded $2^{\circ}\times 2.5^{\circ}$ and $4^{\circ}\times 5^{\circ}$ meteorological fields, while the GCHP v11-02c simulations regrid the $0.25^{\circ}\times 0.3125^{\circ}$ fields to the cubed sphere on the fly through ExtData. A native-resolution cubed sphere GEOS data output stream is presently under development at GMAO and will benefit GCHP by reducing the need for regridding.

20

Table 2. Grid resolution and core counts for the performance test simulations.¹

GEOS-Chem implementation	Grid resolution	Number of grid cells	Number of cores used
GCHP	C24	250,000	6 – 216
GCHP	C48	1,000,000	6 – 540
GCHP	C90	3,500,000	12 – 540
GCHP	C180	14,000,000	90 – 540
GCC	4°×5°	240,000	6 – 30
GCC	2°×2.5°	940,000	6 – 30

All simulations were conducted on the Harvard Odyssey computational cluster. Both GCHP and GCC were compiled using the Intel Fortran compiler (v15.0.0), and MPI capabilities for GCHP were provided by OpenMPI (v1.10.3). Each node of the cluster has 32 Intel Broadwell 2.1 GHz cores sharing 128 GB of RAM, and all nodes are connected via Mellanox FDR Infiniband. Input and output data are stored using a Lustre parallel file system, accessible through the same Infiniband network fabric. All simulations were scheduled to enforce exclusive access to the nodes, preventing possible performance degradation due to sharing of node resources.

Figure 3 shows the total time taken to perform the simulation at each resolution, both in terms of wall time and in terms of the time per 1,000 simulated atmospheric columns in the model grid. Results using the conventional GEOS-Chem Classic (GCC) shared-memory platform at 4°×5° and 2°×2.5° resolutions are also shown for comparison.

¹ Both GCC and GCHP use v11-02c of the core GEOS-Chem code.




Figure 3. Wall time taken to perform 1-month GEOS-Chem simulations at different resolutions and with different numbers of cores. The upper panel shows the absolute time taken to complete each simulation at each resolution. The lower plot shows the wall time normalized by the number of atmospheric columns simulated at each resolution. Solid lines are for GCHP simulations (cubed sphere grids) and dashed lines are for GCC simulations (latitude-longitude grids). Grey lines on each plot show perfect scaling, corresponding to a 50% reduction in run time for each doubling of the number of cores.

At the lowest simulated resolution (C24), GCHP's runtime exceeds that of GEOS-Chem Classic. This is predominantly due to overhead associated with file open operations, as the native-resolution meteorological data used to drive GCHP are opened and read independently for each field. This effect is clearly visible in the lower plot of Figure 3, where performance penalties due to this overhead result in significantly longer run times per 1,000 columns at C24 compared to C180. This can be addressed in the future through both structural changes and parallelization of the input.


After doubling the resolution to C48 ($2^\circ \times 2.5^\circ$), GCHP begins to out-perform GCC, with a reduced overall simulation time even at core counts which are currently accessible to GCC, despite the larger input requirements of GCHP. A 1-month simulation at C48 requires only 6 hours using 96 cores. GCHP scalability also improves as the model resolution increases. At C180, the reduction in simulation time for each doubling in the number of cores is approximately a factor of 1.6.

Two factors affect GCHP scalability: fixed costs and overhead. Fixed costs are for operations which run on a fixed number of cores, regardless of the number of cores dedicated to the simulation. Overhead is the need for additional coordination and

data communication between cores, which grows with the number of cores. Eventually the overhead of the additional cores exceeds the computational benefit, resulting in a performance plateau. This overhead includes one-off costs, such as the MPI interface initialization, which can be significant when running with a large number of cores but which can be reduced in relative terms by running longer simulations.

5

The scalability of each model component is shown in Figure 4. This shows the total time spent on each component at C48 and C180 resolution as a function of the total number of cores used, from 6 cores up to 540. We see that the dominant fixed cost at both C48 and C180 is input, which also dominates the overall cost for C48 with more than 48 cores. This is due to the serial nature of the current input code, overhead associated with file open operations, and the aforementioned use of native-resolution meteorological data for even low-resolution GCHP simulations. Output operations are a second fixed cost, being
 10 handled by a single core at all times. For these simulations, 22 3D fields were stored with hourly frequency, and output was a minor contributor to overall costs. Fixed costs can be converted into scalable costs by parallelizing the component in question, and this is a future work agenda.



15

Figure 4. Total wall time per component for a low (C48) and high (C180) resolution simulation. Simulations at C180 are limited to core counts of 90 or more across several nodes for the hardware used here, due to the high memory requirements of such high-resolution simulations.

Chemistry, advection, and convection all scale well with increasing core count. Chemistry is the most expensive process at
 20 both coarse (C48) and fine (C180) resolution, but has near-perfect scalability. Thus at C48 we see that input becomes the limiting process when the number of cores exceeds 48. Advection and other processes show more departure from perfect scalability, and may dominate the time requirement as the number of cores exceeds 600. The scalability of advection suffers from the additional communication overhead associated with reducing the domain size, as each domain must communicate a

larger proportion of its concentration data to its neighbors. More time is spent on communication relative to computation. However, wall time for advection does consistently fall with increasing core counts, an improvement compared to Long et al. (2015) where wall time increased as core counts exceeded 200 for a grid resolution equivalent to C48. We attribute this to the change from a latitude-longitude grid to the more scalable cubed sphere grid.

5

The remaining wall time is taken up by the “other” component, a mix of scalable and non-scalable processes. This includes the one-off cost of initializing the MPI interface, which grows non-linearly with the number of cores. At C180, these costs are still exceeded by scalable costs when running with 540 cores, so no plateau in performance is observed.

4 High resolution simulations with GCHP

10 The primary advantage of GCHP is the ability to perform simulations of atmospheric chemistry at resolutions previously not available to the community. Figure 5 shows illustrative distributions of simulated ozone concentration at 4 km altitude, simulated at C24 ($\approx 4 \times 5^\circ$) and C180 ($\approx 0.5^\circ \times 0.625^\circ$). Simulations were performed using 24 and 360 cores respectively. Results are also shown from simulations using GCC, at $4^\circ \times 5^\circ$ and $2^\circ \times 2.5^\circ$. Emissions are identical for all simulations. Global-scale patterns in ozone concentration are not substantially affected by the increase in resolution, as these are
15 determined by large-scale processes. The agreement between the two simulations illustrates the consistency of GCHP across scales. However, increasing the horizontal resolution improves the ability of the model to capture the behavior of intercontinental plumes (Eastham and Jacob, 2017). The consequences of this are visible in the ozone distributions over the Pacific and Atlantic. We also observe maxima in the coarse-resolution simulation which are not visible in the finer-
20 resolution simulation.




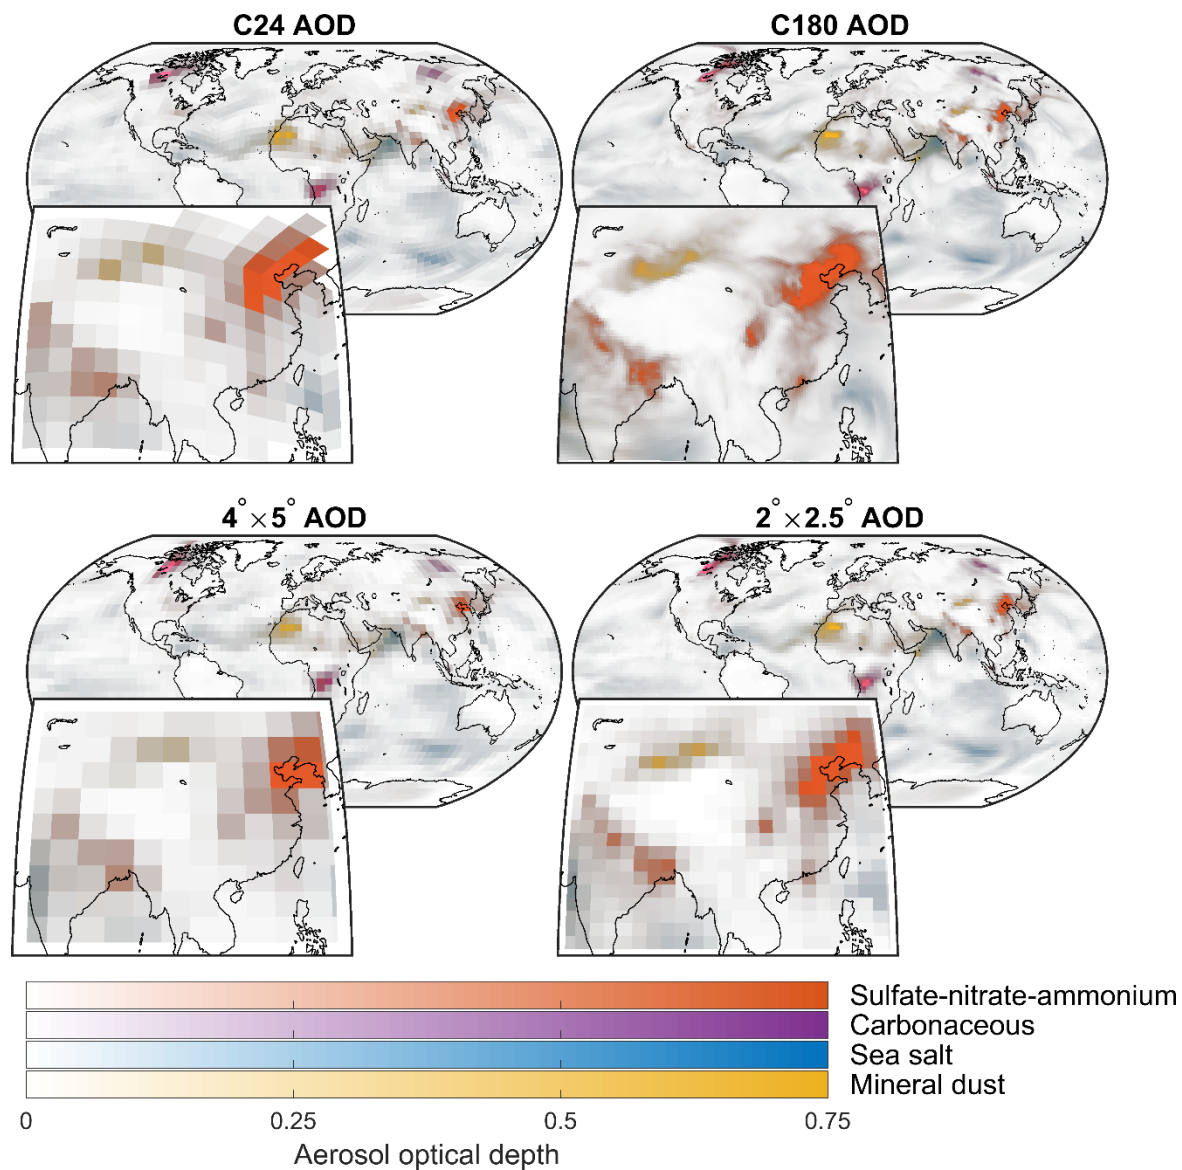
Figure 5. Simulated ozone concentrations at 4 km altitude for 23:00 on July 31st, 2016 after 1 month of initialization. The upper panels show model output from GCHP simulations at C24 (left) and C180 (right), while the lower panels show model output from GCC simulations at 4°×5° and 2°×2.5°. Calculated values in some regions exceeded the displayed limits. Zoom panels are also shown for Europe.

5

Increasing the horizontal resolution also improves the ability of the model to resolve features at the scale of local air quality, as demonstrated by the plots of daily average aerosol optical depth (AOD) shown in Figure 6. This is especially evident in the simulated column AOD over northern India and Beijing. For both ozone and AOD, comparison to the results from GCC at 4°×5° shows that simulations with GCHP at coarse (C24) resolution are able to reproduce the same patterns, magnitude, and variability as those observed in GCC. Some differences are observed, such as the region over Afghanistan which shows

10

an elevated ozone mixing ratio in GCC compared to GCHP. However, the results from GCHP at C24 show the same patterns, magnitude, and variability as those at $4^\circ \times 5^\circ$.



5 **Figure 6.** As for Figure 5, but now showing daily-average aerosol optical depth for July 31st, 2016, after one month of initialization.

Finally, we directly compare the simulated ozone data from GCHP and GCC using 2-D histograms. Figure 6 shows the ozone ratio in each $4^\circ \times 5^\circ$ grid cell, as simulated in GCC (horizontal axis) and as simulated in GCHP (vertical axis) at two different horizontal resolutions. All data are binned at a resolution of 1 ppbv ozone. In each case, GCHP results are conservatively gridded to $4^\circ \times 5^\circ$ to provide a direct comparison, and data in the “polar cap” regions (top and bottom two

4°×5° latitude bands) are excluded. Fitting parameters are shown in white based on a reduced-major-axis (“geometric”) regression. Comparison of results at C24 to those at 4°×5° shows good agreement, with a correlation coefficient $r = 0.88$ and a slope of 0.95. At higher resolution, this agreement is slightly worsened ($r = 0.86$ and slope of 0.92), as smaller-scale processes and chemical non-linearity are resolved which could not be represented at the coarser resolution.

5




Figure 6. 2-D histograms of simulated ozone at 4 km altitude 23:00 on July 31st, 2016 after 1 month of initialization as calculated by GCC and GCHP. Each panel compares the simulated output from GCHP at a specific resolution with the simulated output from GCC at 4°×5°. Data are binned at a resolution of 1 ppbv. GCHP data are conservatively regridded to 4°×5° prior to comparison. Data in the top and bottom two latitude bands are excluded, as GCC averages these points into two “polar caps”.

10

5 Summary

Models of atmospheric chemistry have grown continuously in resolution and complexity over the past decades to take advantage of increasing computational resources. The GEOS-Chem High Performance model (GCHP) is a next step in this growth, enabling the widely-used GEOS-Chem chemical transport model to exploit the computational speed and memory capacity of massively parallel architectures. In this manner we can achieve routine simulation of global stratosphere-troposphere oxidant-aerosol chemistry at unprecedented resolution and detail.

15

Detailed documentation of GCHP including a user’s manual is available on the GCHP website (URL reference). GCHP incorporates the existing GEOS-Chem shared-memory code into an ESMF-based framework (MAPL), enabling GEOS-Chem to be run in a distributed memory framework across multiple nodes while retaining all the features of the high-fidelity

20

global chemical simulation. In addition to a new model framework, GCHP replaces the conventional rectilinear latitude-longitude grid with the gnomonic cubed sphere grid of the NASA GEOS meteorological data used as input to GEOS-Chem. This provides greater computational accuracy and efficiency for transport calculations while removing an additional restriction on scalability. GCHP performs with high computational scalability up to at least 540 cores, completing a 1-month simulation of oxidant-aerosol chemistry in the troposphere and stratosphere (206 active species, 135 tracers) at a global resolution of C180 ($\sim 0.5^\circ \times 0.625^\circ$) in 24 hours.

GCHP also provides a mechanism for ongoing improvement of modeling capability. With the base GEOS-Chem model, GCHP, and the GMAO GEOS atmospheric data assimilation system now all using an identical copy of the grid-independent GEOS-Chem code, GCHP closes the loop between online and offline modelers, allowing seamless propagation of model and framework improvements between all three. Future development opportunities range from improved parallelism in input operations to the direct ingestion of archived mass fluxes to further improve transport calculation accuracy.

Code availability

GCHP has been openly available as part of the GEOS-Chem code since beta version release v11-02b in June 2017, and was part of the v11-02 public release in March 2018. Complete documentation and access to the GCHP code can be found at <http://www.geos-chem.org>. GCHP is an added functionality for GEOS-Chem users, who can choose to use either GCC or GCHP from the same code download. Both GCC and GCHP functionalities will be maintained in the standard GEOS-Chem model for the foreseeable future, recognizing that many users may not have access to the resources needed to use GCHP. For this work, GCHP v11-02c was used, a copy of which has been permanently archived (DOI: [10.5281/zenodo.1290835](https://doi.org/10.5281/zenodo.1290835)).

Acknowledgements

This work was supported by the NASA Atmospheric Composition Modeling and Analysis Program (ACMAP) and the NASA Modeling, Analysis and Prediction (MAP) program. Sebastian D. Eastham was supported by the NOAA Climate and Global Change Postdoctoral Fellowship Program, administered by UCAR's Visiting Scientist Programs. Sebastian D. Eastham was also supported by a Harvard University Center for the Environment (HUCE) Postdoctoral Fellowship. The GEOS FP data used in this study/project were provided by the Global Modeling and Assimilation Office (GMAO) at NASA Goddard Space Flight Center. The computations were run on the Odyssey cluster supported by the FAS Division of Science, Research Computing Group at Harvard University. We are grateful to Compute Canada for hosting the Data Portal for GEOS-Chem to store and make available the GEOS meteorological fields used here. The authors would also like to thank Kevin Bowman for providing analytical insight and additional computational resources with which to perform early testing

of the GCHP model. Finally, we would like to thank Junwei Xu for assistance in processing and archiving meteorological data for input to the model.

Author contribution

- MSL and AM designed the initial code infrastructure
- 5 RVM and DJJ provided project oversight and top-level design
- SDE, MSL, CAK, EWL, MY, RY, JZ, MT, AT, TC, JK, WP, and CJL performed code development
- SDE, EL, RVM, CL, and CJL adapted the meteorological archive
- SDE, MSL, EWL, and JZ ran and debugged benchmark simulations
- SDE, MSL, RY, and JZ performed scalability testing analysis
- 10 SDE, MSL, CAK, RY, EWL, JZ, RVM, and DJJ wrote the manuscript
- All authors contributed to manuscript editing and revisions

References

- 15 Bey, I., Jacob, D. J., Yantosca, R. M., Logan, J. a., Field, B. D., Fiore, A. M., Li, Q., Liu, H. Y., Mickley, L. J. and Schultz, M. G.: Global modeling of tropospheric chemistry with assimilated meteorology: Model description and evaluation, *J. Geophys. Res.*, 106(D19), 23073, 2001.
- Brasseur, G. P. and Jacob, D. K.: *Modeling of Atmospheric Chemistry*, Cambridge University Press., 2017.
- Eastham, S. D. and Jacob, D. J.: Limits on the ability of global Eulerian models to resolve intercontinental transport of chemical plumes, *Atmos. Chem. Phys.*, 17, 2543–2553, 2017.
- 20 Eastham, S. D., Weisenstein, D. K. and Barrett, S. R. H.: Development and evaluation of the unified tropospheric–stratospheric chemistry extension (UCX) for the global chemistry-transport model GEOS-Chem, *Atmos. Environ.*, 89, 52–63, 2014.
- Fairlie, D. T., Jacob, D. J. and Park, R. J.: The impact of transpacific transport of mineral dust in the United States, *Atmos. Environ.*, 41(6), 1251–1266, 2007.
- 25 Hill, C., DeLuca, C., Balaji, Suarez, M. and Silva, A. D.: The architecture of the Earth System Modeling Framework, *Computing in Science Engineering*, 6(1), 18–28, 2004.
- Horowitz, L. W., Walters, S., Mauzerall, D. L., Emmons, L. K., Rasch, P. J., Granier, C., Tie, X., Lamarque, J.-F., Schultz, M. G., Tyndall, G. S., Orlando, J. J. and Brasseur, G. P.: A global simulation of tropospheric ozone and related tracers: Description and evaluation of MOZART, version 2, *J. Geophys. Res.*, 108(D24), doi:10.1029/2002JD002853, 2003.
- 30 Jöckel, P., von Kuhlmann, R., Lawrence, M. G., Steil, B., Brenninkmeijer, C. A. M., Crutzen, P. J., Rasch, P. J. and Eaton, B.: On a fundamental problem in implementing flux-form advection schemes for tracer transport in 3-

- dimensional general circulation and chemistry transport models, *Q.J.R. Meteorol. Soc.*, 127(573), 1035–1052, 2001.
- Keller, C., Auer, B., Hu, L., Jacob, D., Long, M., Nielsen, E., Pawson, S., da Silva, A. and Thompson, M.: High-resolution GEOS-5 nature run with tropospheric chemistry, 2017.
- 5 Keller, C. A., Long, M. S., Yantosca, R. M., Da Silva, A. M., Pawson, S. and Jacob, D. J.: HEMCO v1.0: a versatile, ESMF-compliant component for calculating emissions in atmospheric models, *Geoscientific Model Development*, 7(4), 1409–1417, 2014.
- Lin, J.-T. and McElroy, M. B.: Impacts of boundary layer mixing on pollutant vertical profiles in the lower troposphere: Implications to satellite remote sensing, *Atmos. Environ.*, 44(14), 1726–1739, 2010.
- 10 Lin, S.-J.: A “Vertically Lagrangian” Finite-Volume Dynamical Core for Global Models, *Mon. Weather Rev.*, 132(10), 2293–2307, 2004.
- Liu, H., Jacob, D. J., Bey, I. and Yantosca, R. M.: Constraints from ²¹⁰Pb and ⁷Be on wet deposition and transport in a global three-dimensional chemical tracer model driven by assimilated meteorological fields, *J. Geophys. Res.*, 106(D11), 12109–12128, 2001.
- 15 Long, M. S., Yantosca, R., Nielsen, J. E., Keller, C. A., da Silva, A., Sulprizio, M. P., Pawson, S. and Jacob, D. J.: Development of a grid-independent GEOS-Chem chemical transport model (v9-02) as an atmospheric chemistry module for Earth system models, *Geoscientific Model Development*, 8(3), 595–602, 2015.
- National Research Council: A National Strategy for Advancing Climate Modeling, The National Academies Press, Washington, DC., 2012.
- 20 Nielsen, J. E., Pawson, S., Molod, A., Auer, B., da Silva, A. M., Douglass, A. R., Duncan, B., Liang, Q., Manyin, M., Oman, L. D., Putman, W., Strahan, S. E. and Wargan, K.: Chemical Mechanisms and Their Applications in the Goddard Earth Observing System (GEOS) Earth System Model, *J. Adv. Model. Earth Syst.*, 9(8), 3019–3044, 2017.
- Prather, M. J.: Fast-JX v7.0a, 2012. https://www.ess.uci.edu/group/prather/scholar_software
- 25 Putman, W. M. and Lin, S. J.: Finite-volume transport on various cubed-sphere grids, *J. Comput. Phys.*, 227(1), 55–78, 2007.
- Sherwen, T., Evans, M. J., Carpenter, L. J., Andrews, S. J., Lidster, R. T., Dix, B., Koenig, T. K., Sinreich, R., Ortega, I., Volkamer, R., Saiz-Lopez, A., Prados-Roman, C., Mahajan, A. S. and Ordóñez, C.: Iodine’s impact on tropospheric oxidants: a global model study in GEOS-Chem, *Atmos. Chem. Phys.*, 16(2), 1161–1186, 2016.
- 30 Suarez, M., Trayanov, A., Hill, C., Schopf, P. and Vikhliav, Y.: MAPL: A High-level Programming Paradigm to Support More Rapid and Robust Encoding of Hierarchical Trees of Interacting High-performance Components, in *Proceedings of the 2007 Symposium on Component and Framework Technology in High-performance and Scientific Computing*, pp. 11–20, ACM, New York, NY, USA., 2007.
- Ullrich, P. A. and Taylor, M. A.: Arbitrary-Order Conservative and Consistent Remapping and a Theory of

Linear Maps: Part I, Mon. Weather Rev., 143(6), 2419–2440, 2015.

Wang, Y., Jacob, D. J. and Logan, J. A.: Global simulation of tropospheric O₃-NO_x-hydrocarbon chemistry: 1. Model formulation, J. Geophys. Res., 103(D9), 10713–10725, 1998.

5 Wu, S., Mickley, L. J., Jacob, D. J., Logan, J. A., Yantosca, R. M. and Rind, D.: Why are there large differences between models in global budgets of tropospheric ozone?, J. Geophys. Res., 112(D5), D05302, 2007.