



# Improved Advection, Resolution, Performance, and Community Access in the New Generation (Version 13) of the High Performance GEOS-Chem Global Atmospheric Chemistry Model (GCHP)

5 Randall V. Martin<sup>1</sup>, Sebastian D. Eastham<sup>2</sup>, Liam Bindle<sup>1</sup>, Elizabeth W. Lundgren<sup>3</sup>, Thomas L. Clune<sup>4</sup>,  
Christoph A. Keller<sup>4,5,†</sup>, William Downs<sup>3,&</sup>, Dandan Zhang<sup>1</sup>, Robert A. Lucchesi<sup>4</sup>, Melissa P. Sulprizio<sup>3</sup>,  
Robert M. Yantosca<sup>3</sup>, Yanshun Li<sup>1</sup>, Lucas Estrada<sup>3</sup>, William M. Putman<sup>4</sup>, Benjamin M. Auer<sup>4</sup>, Atanas  
L. Trayanov<sup>4</sup>, Steven Pawson<sup>5</sup>, Daniel J. Jacob<sup>3</sup>

<sup>1</sup>Department of Energy, Environmental & Chemical Engineering, Washington University in St. Louis, St. Louis, MO, USA

<sup>2</sup>Laboratory for Aviation and the Environment, Massachusetts Institute of Technology, Cambridge, MA, USA

10 <sup>3</sup>John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA

<sup>4</sup>Global Modeling and Assimilation Office, NASA Goddard Space Flight Center, Greenbelt, MD, USA

<sup>5</sup>Universities Space Research Association, Columbia, MD, USA.

<sup>†</sup>Now at Morgan State University, MD, USA

<sup>&</sup>Now at Rosenstiel School of Marine and Atmospheric Science, University of Miami, FL, USA

15

*Correspondence to:* Randall V. Martin ([rvmartin@wustl.edu](mailto:rvmartin@wustl.edu))

**Abstract.** We describe a new generation of the high-performance GEOS-Chem (GCHP) global model of atmospheric composition developed as part of the GEOS-Chem version 13 series. GEOS-Chem is an open-source grid-independent model that can be used online within a meteorological simulation or off-line using archived meteorological data. GCHP is an  
20 offline implementation of GEOS-Chem driven by NASA Goddard Earth Observing System (GEOS) meteorological data for massively parallel simulations. Version 13 offers transformational advances in GCHP for ease of use, computational performance, versatility, resolution, and accuracy. Specific improvements include (a) stretched-grid capability for higher resolution in user-selected regions, (b) easier build with a build system generator (CMake) and a package manager (Spack),  
(c) software containers to enable immediate model download and configuration on local computing clusters, (d) better  
25 parallelization to enable simulation on thousands of cores, (e) multi-node cloud capability, and (f) more accurate transport with new native cubed-sphere GEOS meteorological archives including air mass fluxes at hourly temporal resolution with spatial resolution up to C720 (~12 km). The C720 data are now part of the operational GEOS Forward Processing (GEOS-FP) output stream, and a C180 (~50 km) consistent archive for 1998-present is now being generated as part of a new GEOS-IT data stream. Both of these data streams are continuously being archived by the GEOS-Chem Support Team for access by  
30 GCHP users. Directly using horizontal air mass fluxes rather than inferring from wind data significantly reduces global mean error in calculated surface pressure and vertical advection.



## 1 Introduction

35 Atmospheric chemistry and composition are central drivers of climate change, air quality, and biogeochemical cycling. They are a next frontier for Earth system model (ESM) development (NRC, 2012). Modeling of atmospheric chemistry is a grand scientific and computational challenge because of the need to simulate hundreds of gaseous and aerosol chemical species stiffly coupled to each other and interacting with transport on all scales. There is considerable demand for high-resolution atmospheric chemistry models from a broad community of researchers and stakeholders with interest in simulating a range of problems at local to global scales. But software engineering complexity and computational cost have been major barriers to access.

40 Atmospheric chemistry models solve the 3-D continuity equations for an ensemble of reactive and coupled gaseous/aerosol chemical species with terms to describe emissions, transport, chemistry, aerosol microphysics, and deposition (Brasseur and Jacob, 2017). The model may be integrated “online” within a meteorological model or ESM, with the chemical continuity equations solved together with the equations of atmospheric dynamics, or “offline” as a chemical transport model (CTM) where the chemical continuity equations are solved using external meteorological data as input. The online approach has the  
45 advantage of more accurately coupling chemical transport to dynamics, and has specific application to the study of aerosol-chemistry-climate interactions. It also enables consistent chemical and meteorological data assimilation. The offline approach has advantages of accessibility, cost, portability, reproducibility, and straightforward application to inverse modeling. The broad atmospheric chemistry community can easily access an offline CTM for reusable applications that advance atmospheric chemistry knowledge, but access to an online model is more limited and complicated. Ideally, the same  
50 state-of-the-art model must be able to operate both online and offline.

The GEOS-Chem atmospheric chemistry model ([geos-chem.org](http://geos-chem.org)) delivers this joint online-offline capability. GEOS-Chem is an open-source global 3-D model of atmospheric composition used by hundreds of research groups around the world for a wide range of applications. It simulates tropospheric and stratospheric oxidant-aerosol chemistry, aerosol microphysics, carbon gases, mercury, and other species (e.g., Eastham et al., 2014; Friedman et al., 2014; Kodros and Pierce, 2017; Li et al., 2017; Shah et al., 2021; Wang et al., 2021). GEOS-Chem has been developed and managed continuously for the past 20  
55 years (starting with Bey et al. (2001)) as a grass-roots community effort. The online version is part of the Goddard Earth Observation System (GEOS) of the NASA Global Modeling and Assimilation Office (GMAO) (Long et al., 2015; Hu et al., 2018; Keller et al., 2021) and has been implemented in other climate and meteorological models as well (Lin et al., 2020; Lu et al., 2020; Feng et al., 2021). The offline version uses exactly the same scientific code and is driven by GEOS  
60 meteorological data or by other meteorological fields (Murray et al., 2021). The offline GEOS-Chem has wide appeal among atmospheric chemists because it is a comprehensive, cutting-edge, open-source, well-documented modeling resource that is easy to use and modify but also has strong central management, version control, and user support through a GEOS-Chem Support Team (GCST) based at Harvard University and at Washington University.



The standard offline version of GEOS-Chem (“GEOS-Chem Classic”) is designed for easy use and a simple code base but  
65 relies on shared-memory parallelization and a rectilinear longitude-latitude grid, limiting its flexibility and scalability for  
high-resolution applications in modern High Performance Computing (HPC) environments. A high-performance version of  
GEOS-Chem (GCHP) was developed by Eastham et al. (2018) to address this limitation. GCHP is a grid-independent  
implementation of GEOS-Chem using Message Passing Interface (MPI) distributed-memory parallelization enabled through  
the Earth System Modeling Framework (ESMF, earthsystemmodeling.org) and the Modeling Analysis and Prediction Layer  
70 (MAPL), in the same way as the GEOS system (Suarez et al., 2007; Long et al., 2015; Eastham et al., 2018; Hu et al., 2018).  
GCHP operates on atmospheric columns as its basic computation units, with grid information specified at runtime through  
ESMF. Chemical transport is simulated using a finite volume advection code (FV3), allowing GEOS-Chem simulations to be  
performed on the native GEOS cubed-sphere grid (Putman and Lin, 2007), but the scientific code base is otherwise the same  
as GEOS-Chem Classic. GCHP enables GEOS-Chem simulations to be conducted with high computational scalability on up  
75 to a thousand cores (Eastham et al., 2018; Zhuang et al., 2020), so that global simulations of stratosphere–troposphere  
oxidant–aerosol chemistry with very high resolution become feasible.

Here we describe development of a new generation of GCHP (version 13) for improved community access, performance,  
resolution, and accuracy. Section 2 provides background on GEOS-Chem and GCHP. The MAPL coupler and GEOS system  
are described in Section 3. Section 4 provides a high-level overview of developments in the GCHP version 13 series. A  
80 stretched grid capability is described in Section 5. Advances in ESM coupling and software collaboration are presented in  
Section 6. In Section 7, improvements to GCHP performance and portability are described including (a) easier build system,  
(b) software containers to facilitate download and configuration, (c) better parallelization, and (d) capability for multi-node  
simulations on the cloud. Section 8 presents extension of GCHP capabilities to utilize fine-scale cubed-sphere meteorology  
at the global scale, and describes two new cubed-sphere data streams. A performance demonstration in Section 9 is followed  
85 by a section on future needs and opportunities. This new generation of GCHP is extensively documented on our GCHP Read  
The Docs site (<https://gchp.readthedocs.io/en/latest/>).

## 2 GEOS-Chem and GCHP

GEOS-Chem simulates the evolution of atmospheric composition by solving the system of coupled continuity equations for  
an ensemble of  $m$  species (gases or aerosols) with concentration vector  $\mathbf{n} = (n_1, \dots, n_m)^T$ :

$$90 \quad \frac{\partial n_i}{\partial t} = -\nabla \cdot (n_i \mathbf{U}) + P_i(\mathbf{n}) - L_i(\mathbf{n}) + E_i - D_i \quad i \in [1, m] \quad (1)$$

Here  $\mathbf{U}$  is the wind vector (including subgrid components parameterized as boundary layer mixing and wet convection);  
 $P_i(\mathbf{n})$  and  $L_i(\mathbf{n})$  are the local production and loss rates of species  $i$  from chemistry and/or aerosol microphysics, which depend  
on the concentrations of other species; and  $E_i$  and  $D_i$  represent emissions and deposition. Equation (1) is solved by operator  
splitting of the transport and local components over finite time steps. The local operator,



95

$$\frac{dn_i}{dt} = P_i(\mathbf{n}) - L_i(\mathbf{n}) + E_i - D_i \quad i \in [1, m] \quad (2)$$

includes no transport terms and thus reduces to a system of coupled first-order ordinary differential equations (ODEs). We refer to it as the GEOS-Chem chemical module even though it also includes terms for emission, deposition, and aerosol microphysics.

100 GEOS-Chem includes routines to conduct all of the operations in equation (1). The simulations can be conducted either offline or online. The offline mode uses archived meteorological data, including  $\mathbf{U}$  and other variables, to solve equation (1). This includes transport modules for grid-resolved advection, boundary layer mixing, and wet convection. The online mode uses the GEOS-Chem chemical module (equation (2)) to solve for the local evolution of chemical species within a meteorological model where transport of the chemical species is done independently as part of the model dynamics. The GEOS-Chem transport modules are then disabled.

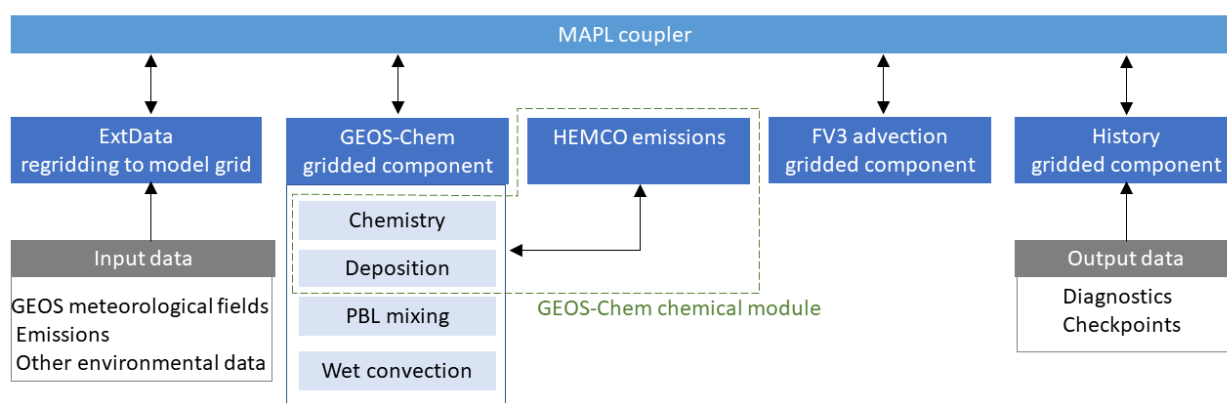
105 The standard offline implementation of GEOS-Chem uses NASA GEOS meteorological archives as input, currently either from the Modern-Era Retrospective analysis for Research and Applications, Version 2 (MERRA-2) for 1980 to present, or from the GEOS Forward Processing (GEOS-FP) product generated in near-real-time. In GEOS-Chem Classic, first described by Bey et al. (2001), the model provides a choice of rectilinear latitude-longitude Eulerian grids with shared-memory parallelization. The coding architecture is simple but efficient parallelization is limited to a single node with tens of cores.  
110 GEOS-Chem Classic can be used in principle at the native resolutions of MERRA-2 ( $0.5^\circ \times 0.625^\circ$ ) or GEOS-FP ( $0.25^\circ \times 0.3125^\circ$ ), but global simulations are limited in practice to  $2^\circ \times 2.5^\circ$  or  $4^\circ \times 5^\circ$  horizontal resolution because of the inefficient parallelization and prohibitive single-node memory requirements. Native-resolution simulations can be conducted for regional/continental domains (Zhang et al., 2015; Li et al., 2021), with boundary conditions from an independently conducted coarse-resolution global simulation. However, with simulated atmospheric chemistry continuously increasing in  
115 computational complexity and the performance of individual computational nodes relatively stagnant, the restrictions of running on a single node increasingly force users to choose between speed, resolution, and accuracy.

GCHP, first described by Eastham et al. (2018) evolved the offline implementation of GEOS-Chem to a grid-independent formulation with MPI distributed-memory parallelization. The grid-independent formulation of GEOS-Chem, originally developed by Long et al. (2015) for online applications, enables the model to operate on any horizontal grid specified at run  
120 time. The model solves for the chemical module (equation (2)) on 1-D vertical columns of the user-specified grid, and passes the updated concentrations at each time step to the transport modules. In GCHP, this grid-independent formulation is exploited in an offline mode with the MAPL coupler and ESMF to operate GEOS-Chem on the native cubed-sphere of the GEOS meteorological model. MAPL/ESMF delivers the MPI capability allowing for efficient parallelization on up to a thousand cores (Eastham et al., 2018; Zhuang et al., 2020) and enables global simulations at the native resolution of the  
125 GEOS meteorological data. At the same time, GCHP can still be run on a single node with similar performance as GEOS-Chem Classic for low-resolution applications.



Figure 1 shows the general architecture of GCHP. MAPL couples the different components of the model (gridded components), providing and receiving inputs, and handles parallelization. Meteorological and other data are read through the ExtData module and re-gridded as needed to the desired cubed-sphere resolution. Advection on the cubed-sphere is done with the offline FV3 module of Putman and Lin (2007). GEOS-Chem updates the chemical concentrations over model time steps in 1-D columns corresponding to the model grid, including subgrid vertical transport from boundary layer mixing and wet convection. Model output diagnostics are archived through the History module. GCHP is written in Fortran with the option to use either Intel or GNU compilers. Beyond the NetCDF libraries required for GEOS-Chem, GCHP's additional dependencies (external standalone libraries) are an MPI implementation and ESMF.

130



135

**Figure 1: Schematic of GCHP architecture. The model consists of four gridded components (ExtData, GEOS-Chem, FV3, History) exchanging information through the MAPL coupler. The HEMCO emissions module communicates directly with the GEOS-Chem gridded component in the current GCHP architecture, but it can also be used as a separate gridded component in other model architectures (Lin et al., 2021). The GEOS-Chem gridded component includes planetary boundary layer (PBL) mixing and wet convective transport of species as governed by the GEOS meteorological fields passed through MAPL. ‘Chemistry’ also includes aerosol microphysical processes for which the continuity equations are analogous. The GEOS-Chem chemical module as defined in the text and illustrated in the Figure includes emissions, chemistry, and deposition and would be the unit passed to a meteorological model or ESM in online applications.**

140

### 145 3 MAPL and GEOS

#### 3.1 MAPL overview

150

MAPL is an infrastructure layer that leverages ESMF to provide services that simplify the process of coupling model components and enforce certain consistency conventions across components. In particular, MAPL provides high-level interfaces that allow developers of gridded components to readily specify the imports, exports, and internal states for their components as well as to hierarchically incorporate “child” components. The “generic” layer in MAPL translates the high-level specifications to register initialize/run/finalize methods with ESMF, allocate storage, create ESMF fields and states,



and enable the use of shared pointers wherever possible to reduce memory and performance overheads. This generic layer additionally provides common services across components such as checkpoint and restart.

155 MAPL also provides two highly-configurable ESMF components: ExtData and History which manage spatially distributed input and output respectively as described in Section 2. MAPL automatically aggregates all component exports to make them available to the History component for output and sends any unsatisfied imports to ExtData for input. ExtData and History have the capability to automatically regrid to and from the model and component grid with a variety of temporal sampling and horizontal interpolation options.

160 MAPL also fills some gaps in ESMF functionality, though the nature of those gaps continually evolves as both frameworks advance. Currently MAPL provides at least one regridding method not yet available in ESMF, namely the ability to regrid horizontal fluxes in an exact manner for integral grid resolution ratios.

A major performance bottleneck in the original version of GCHP as described in Eastham et al. (2018) was in the reading of input data. The current version of MAPL includes optimizations to the ExtData layer used for input with elimination of redundant actions and use of multiple cores on a single node for data input, thereby reducing the input computational cost.  
165 GCHP timing tests with this new capability are presented in section 7.5.

### 3.2. GEOS system

The GEOS system of NASA GMAO provides meteorological inputs needed by GEOS-Chem including wind and pressure information, humidity and precipitation data, as well as surface quantities such as soil moisture, friction velocity and skin temperature. The full list of meteorological input data used by GEOS-Chem can be found at  
170 [http://wiki.seas.harvard.edu/geos-chem/index.php/List\\_of\\_GEOS-FP\\_met\\_fields](http://wiki.seas.harvard.edu/geos-chem/index.php/List_of_GEOS-FP_met_fields).

The GEOS-FP and MERRA-2 data products used to drive GEOS-Chem are generated by the GEOS ESM and Data Assimilation System (DAS), consisting of a suite of modular model components connected through the ESMF/MAPL software interface (Todling and Akkraoui, 2018). GEOS-FP (Lucchesi, 2017) uses the most recent validated version of the GEOS ESM system and produces meteorological and aerosol analyses and forecasts in near real time. Currently (version  
175 5.27.1), it runs on a cubed-sphere grid with a horizontal resolution of C720 (approx.  $12 \times 12$  km<sup>2</sup>) where the resolution of the cubed-sphere output is indicated by CN, and N is the number of grid boxes on one edge of one face of the cubed-sphere. Thus the total number of cells in one model level is  $6N^2$ . The outputs from this system have been conventionally archived on a latitude-longitude grid with a horizontal resolution of  $0.25^\circ \times 0.3125^\circ$ , incurring loss of resolution and accuracy in vector fields as presented in Section 8. MERRA-2 is a meteorological and aerosol reanalysis from 1979 to present produced with a  
180 stable version of GEOS (Gelaro et al., 2017). MERRA-2 simulations are conducted at a lower horizontal resolution than GEOS-FP (C180 vs. C720), and all MERRA-2 fields are archived on a latitude-longitude grid with a horizontal resolution of  $0.5^\circ \times 0.625^\circ$ . For both GEOS-FP and MERRA-2, traditional archival has been at one-hour temporal resolution for surface quantities and three-hour temporal resolution for 3D quantities such as winds. Winds are defined at the center of the grid cell (A-Grid staggering using the notation introduced by Arakawa and Lamb (1977)), while mass fluxes are defined at the center



185 of the relevant grid edges in 2D contexts and interfaces of the discrete volumes of the grid in 3D contexts (C-Grid  
staggering) as further described in Section 8.2. New cubed-sphere archives including mass fluxes are described in Section  
8.3. The GEOS-Chem Support Team historically reprocessed GEOS data into specific input formats including coarser  
resolution and nested domains for use by GEOS-Chem Classic. The FlexGrid option implemented in GEOS-Chem version  
12.4 enabled the generation of coarse-grid and custom nested data on the fly at run time (Shen et al., 2021b) but other  
190 reprocessing of the native fields was still required for GCHP. Access to pre-generated coarse-grid archives ( $2^\circ \times 2.5^\circ$  and  
 $4^\circ \times 5^\circ$ ) and pre-cut nested domains is still supported for GEOS-Chem Classic, but the reprocessing can now be skipped for  
GCHP as described in Section 8.4.

#### 4 Overview of new capabilities in the GCHP version 13 series

Table 1 contains an overview of the new capabilities for GCHP that have been implemented as part of the version 13 series  
195 for improved community access, performance, resolution, and accuracy.

**Table 1: Overview of new capabilities in the GCHP version 13 series**

Feature	Section <sup>a</sup>
Stretched-grid capability for higher resolution in user-selected regions	5
Advances in ESM coupling and software collaboration	6
Easier build with a build system generator (CMake) and a package manager (Spack)	7.1 & 7.2
Software containers to enable immediate model download and configuration on local computing clusters	7.3
Improved error and output diagnostics	7.4
Better parallelization to enable simulations on thousands of cores	7.5
Multi-node cloud capability	7.6
More accurate transport through use of mass fluxes on the cubed-sphere grid	8.1 & 8.2
New hourly native cubed-sphere GEOS meteorological archives	8.3
Direct ingestion of GEOS meteorological archives	8.4

<sup>a</sup> Section of the manuscript where the new capability is discussed

200 Community access was facilitated by improving the build system through a build system generator (CMake) and a package  
manager (Spack), by offering software containers, by improving error and output diagnostics, and by developing a multi-  
node cloud capability. Use of the CMake build system generator described in section 7.1 (a) improved the robustness of the  
build, (b) improved the maintainability of the build system, and (c) made building GCHP easier for users. Use of the Spack





package manager described in section 7.2 eased the installation of GCHP by specifying precisely how to build GCHP for  
205 different versions, configurations, platforms, and compilers. Use of containers enabled immediate download and  
configuration for cloud environments and for local environments that support containers, as described in section 7.3.  
Improved error and output diagnostics facilitate debugging and evaluation, as described in section 7.4. The ability for GCHP  
to directly use GEOS meteorological archives opened up new opportunities for near-real-time simulations as presented in  
section 8.4.

210 Performance was improved through better parallelization as described in Section 7.5, enabling efficient simulations on  
thousands of cores. For example, a full-chemistry simulation at C360 (~25 km) now achieves 20 model days per day of  
actual time on 2304 cores. The improved parallelization was achieved by updating the MAPL software to take advantage of  
improvements in input efficiency that eliminated the previous computational bottleneck as described in Section 3.1.

Resolution was improved through the generation of hourly GEOS archives for advection variables, with resolution up to  
215 cubed-sphere C720 (~12 km) and development of a stretched-grid capability for regional refinement. The GEOS-FP C720  
advection archive began production on March 11, 2021 and is continuing operationally. Hourly archiving (instead of three-  
hourly previously) of the advection variables (air mass fluxes, specific humidity, Courant numbers, and surface pressure)  
significantly reduces transport errors associated with transient (eddy and convective) advection (Yu et al., 2018). The cubed-  
sphere archive is most critical for advection variables since they increase the accuracy of the transport simulation. An in-  
220 progress GEOS-IT archive for the period 1998-present includes cubed-sphere archives of all meteorological variables at  
hourly C180 resolution as described in Section 8.3. The stretched-grid capability was described in Bindle et al. (2021) and is  
summarized in Section 5.

Accuracy was improved by directly ingesting mass fluxes instead of winds as described in Section 8.1, by conducting  
simulations directly on the native cubed-sphere grid of the meteorology, and through high- resolution meteorological  
225 archives. The use of mass fluxes is particularly important for accurate vertical transport in the stratosphere where weak  
vertical motion increases susceptibility to errors from use of winds. Conducting simulations directly on the cubed-sphere  
reduces errors from regriding to and from a latitude-longitude grid to and from cubed-sphere grid, and from restaggering to  
and from the center of a grid cell to and from the center of a grid edge as described in section 8.2.

## 5 Stretched grid

230 A limitation of the original version of GCHP was the absence of a grid-refinement capability over regions of specific  
interest. GEOS-Chem Classic has a nested-grid capability to allow native-resolution simulations over regional or continental  
domains with dynamic boundary conditions from the global simulation (Wang et al., 2003), and these domains can be  
defined at runtime with the FlexGrid facility (Li et al., 2021). This is not possible in GCHP because there is not yet a  
mechanism to specify boundary conditions in a non-global domain and because FlexGrid is only for latitude-longitude grids.  
235 Bindle et al. (2021) implemented grid stretching as a means for regional grid refinement in GCHP. Grid-stretching in GCHP






uses a modified Schmidt (1977) transform (Harris et al., 2016) to "stretch" the cubed-sphere grid for all input data through ExtData to create a refinement. The user has nimble control over the refinement location and strength using three runtime parameters; the *stretch-factor* controls the refinement strength, and the *target-latitude* and *target-longitude* control the refinement location.

240 Several recent developments enabled the implementation of the stretched-grid capability in GCHP. Harris et al. (2016) developed the stretched-grid capability for the FV3 advection code used in GCHP. The MAPL framework added support for that capability. An archive of state-dependent emissions at native resolution was developed for GEOS-Chem, producing consistent emissions regardless of the model grid (Weng et al., 2020; Meng et al., 2021).

245 Figure 2 shows a visualization of the stretched grid. A key advantage of grid-stretching compared to other refinement techniques, such as nesting, is the smoothness of the transition from the region of interest to the global background. Stretching does not change the logical structure (topology) of the grid, and two-way coupling is inherent; this means stretching can be implemented without major structural changes to the model or the need for a component to couple the simulation across distinct model grids.



250

**Figure 2: Visualization of grid stretching for two refinement scenarios using stretch factors ( $S$ ) on a C16 cubed-sphere. Resolution factor is the relative change of grid box edge length induced by stretching. Adapted from Bindle et al. (2021).**

## 6 Advances in ESM coupling and software collaboration

Here we describe restructuring of GCHP and its interfaces with MAPL and its dependencies in the version 13 series. These developments were motivated by the needs for tighter coupling of GCHP with the parent ESM (GEOS), for coordinated development of MAPL between the GCHP and GMAO development teams, and for reduction in GCHP build time.

The original version of GCHP (Eastham et al., 2018) was implemented as a code repository that users manually inserted into the GEOS-Chem code base. It contained a copy of MAPL and other GMAO software libraries that were frozen during the



initial development process and contained no GMAO development history. Ongoing improvements to the MAPL  
260 infrastructure by GMAO were not regularly or easily propagated to the GCHP code, while bug fixes and MAPL  
enhancements made by GCHP developers were not easily propagated back to GMAO. This disconnect resulted in  
divergence of code, difficulty updating GEOS-Chem in GEOS, and limitations on the progress of GCHP capabilities.

We restructured GCHP in version 13.0.0 to address these issues. We implemented independently maintained code bases such  
as MAPL and HEMCO as Git submodules that contained version history information. We replaced in GCHP the existing  
265 version of MAPL and its dependencies with the latest stable version releases, thereby expanding infrastructure capabilities  
for GCHP such as updates necessary for simulations on a stretched grid. We also developed a system for seamless version  
updates between GCHP and GEOS code bases by using forks of GMAO software repositories as Git submodules for  
straightforward merging of code updates via GitHub pull requests while retaining all version history.

GCHP 13.0.0 also changed how GCHP interfaces with the ESMF library. GCHP originally contained a copy of ESMF  
270 without version history and users were required to build ESMF from scratch with every new GCHP download. This caused  
build times substantially longer than necessary given ESMF in GCHP rarely changed. To reduce build time, we restructured  
GCHP to use ESMF as an external library. Users now may download ESMF from its public repository  
(earthsystemmodeling.org), build it locally, and use the same build for GCHP or any other ESMF-based applications. ESMF  
can even be built as a system-wide module, enabling all users on a system to use a centrally-maintained copy in the same  
275 way that components such as compilers, NetCDF, or MPI are treated. This is beneficial for the following reasons: (a) Allows  
greater flexibility for ESMF version updates, (b) Cuts initial GCHP build time in half, (c) Provides greater transparency in  
ESMF via original Git history, (d) Reflects that ESMF is a separate project from GCHP, with its own model development  
activities and support team, and (e) Leverages the advantages of centrally-maintained libraries in modern HPC systems,  
allowing science-focused users to build and run GCHP with minimal effort.

Overall the GCHP 13.0.0 restructuring with common version control repositories enables version updates of GMAO libraries  
such as MAPL to be seamless and improves software collaboration between GCHP and GMAO developers by ensuring that  
future improvements in either GMAO or the GCHP community are immediately available to both sets of developers. As an  
example, the update to a recent version of MAPL increased optimization of the ExtData layer used for inputs to enable  
efficient parallelization to extend from hundreds to thousands of cores, while GCHP updates to MAPL such as bug fixes in  
285 the stretched grid feature have enhanced GMAO capabilities. Updating GCHP to share its infrastructure code with GMAO  
via common version control repositories successfully achieves synergistic model development. GCHP's use of forks as Git  
submodules allows simple merging between GEOS-Chem and GEOS-ESM GitHub repositories; we utilize GitHub pull  
requests and issues for cross-communication and collaboration between GCHP and GMAO developers. Finally, use of  
GitHub issues and notifications improves transparency and communication with GCHP users, with all code exchanges and  
290 issues publicly viewable and searchable.



## 7 Improvements to GCHP performance and portability

Here we describe efforts to reduce the difficulty of compiling and running GCHP by improving the build system, the ease of installation through a package manager and software containers, diagnostics, parallelization, and the multi-node cloud capability. Previously, compiling GCHP was more difficult than compiling GEOS-Chem Classic because of the additional dependencies and complex interface. These surplus codes increased the time required to compile the GCHP model and increased the code brittleness by introducing points of failure. We addressed these issues through a build system generator (CMake) which simplifies building GCHP once its dependencies are satisfied, a package manager (Spack) which automates the process of acquiring missing dependencies, and software containers which can sidestep the entire process in environments that support containers.

### 300 7.1 CMake

Maintaining a portable and easy-to-use build system is challenging in the context of high performance computing (HPC) software because of the diversity in HPC environments. Environment differences between clusters include different combinations of dependencies in the software stack such as versions and families of compilers, differences in the build-time options of those dependencies such as library support extensions, and differences in system administration such as the paths to installed software. In practice, these environment differences translate to different compiler options. The Make build system previously used in GCHP was brittle and laborious to maintain due to its need for detailed customization to accommodate differences between clusters. Compared to Make, CMake has a more formal structure for organizing projects and specifying build properties; this facilitates the organization of GCHP's build files and interoperability of GCHP's build files with those of internal dependencies (dependencies which are built on-the-fly during the GCHP build). The interoperability of CMake-based projects allowed us to leverage existing build files for MAPL, developed and maintained at the GMAO, for building MAPL within the GCHP build.

To address these issues we implemented a build system generator (CMake, [cmake.org](http://cmake.org)) in GCHP to (a) improve the robustness of the build, (b) improve the maintainability of the build system, and (c) make building GCHP easier for end-users. Build system generators like CMake are specifically designed to generate a build system (build scripts) for the system, according to the compute environment. This new build system follows the canonical build procedure for CMake-based builds: a configuration step, a build step, and an install step. During the configuration step, the user executes CMake in a build directory; CMake inspects the environment and generates a set of Make build scripts that build the model. The build step is the familiar compile step where the user runs Make, and the compiler command sequence is executed. The install step is used to port the built executable into the users experiment directory. In addition to a more robust build, benefits of the new build system include more readable build logs and fewer environment variables.



## 7.2 Spack

Our next step was to ease the installation of GCHP by specifying precisely how to build GCHP for different versions, configurations, platforms, and compilers through an instruction set (i.e., “recipes”) for GCHP dependencies which can be built using Spack (Gamblin et al., 2015; <http://spack.io/>). Spack is an innovative package manager designed to ease installation of scientific software, by automating the process of building from public repositories all of the dependencies necessary for GCHP if any are missing from the target machine (C compiler, Fortran compiler, NetCDF-C, NetCDF-Fortran, MPI implementation, and ESMF). The flexibility of Spack facilitates implementing numerous build options that can handle a diversity of compilers and environments, thus ensuring that most users can build a functioning copy of GCHP dependencies in a single step, without requiring the user to understand the details of configuring GCHP for their environment. This activity includes testing the GCHP code with multiple versions of: GNU and Intel compilers; ESMF; multiple MPI implementations (e.g. OpenMPI, MVAPICH2, and MPICH); and the NetCDF libraries. Working configurations are implemented as publicly-available Spack packages to enable new users to install GCHP without concern for conflicts between different versions of different dependencies. The new CMake library is a part of this Spack package.

As part of this effort, we developed a Spack recipe that in a single command allows users to download from the Spack GitHub repository all GCHP dependencies and build GCHP. Users can modify this command to provide to GCHP specific compile-time build options such as whether to include a specific radiative transfer model (RRTMG; Iacono et al., 2008). Spack also provides syntax for specifying different release versions and compiler specifications for packages and their dependencies. Since GCHP can be built without any proprietary software, open source compilers are sufficient. The GCHP Spack package is maintained by the GEOS-Chem Support Team.

Spack is most useful on systems where few of GCHP’s required libraries already exist, e.g. new scientific computing clusters, cloud environments, or container creation. Spack itself only requires a basic C/C++ compiler and a Python installation (since Spack is written in Python) to begin building GCHP’s dependencies. These are usually available as standard in most modern Linux environments.

Users can manually specify any existing libraries on their system through Spack configuration files to avoid redundant installations of GCHP dependencies. This is a required setup step for using existing job schedulers such as Slurm on a user’s system. Additionally, Spack’s install command includes an option to only install package dependencies without installing the package itself. This option allows users to build and load all dependencies while retaining the ability to modify GCHP source code locally before compiling.

## 7.3 Containers

Both the improvements to the build system and to the installation process are beneficial to most users on most platforms. For HPC clusters that support containers, and for GCHP users in cloud environments, the GEOS-Chem Support Team now maintains pre-built software containers (Kurtzer et al., 2017; Reid and Randles, 2017) containing GCHP and its



dependencies. Software containers provide collections of prebuilt libraries to users that allow GCHP and its software environment to be moved smoothly between cloud platforms and local clusters, so that the identical compute environment can be executed on any machine. A software container encapsulates a compute environment (the operating system, installed libraries and software, system files, and environment variables) allowing the compute environment to be downloaded and executed virtually on other machines, but without the performance penalty associated with emulating hardware.

We created scripts to automatically generate containers for every new GCHP release, which include GCHP and its dependencies (built using Spack). Users can run one of these containers through Docker or Singularity (which natively supports running Docker images). Singularity is often preferred for running HPC applications like GCHP because it does not require elevated user privileges.

Software containers are particularly useful for quickly setting up GCHP environments. The only requirements for running one of these containers are the container software (e.g. Singularity) and an existing MPI installation on a user's system. With these requirements met, the only steps needed to run a GCHP container are downloading the container, downloading GEOS-Chem input data, and creating a run directory.

The main drawback of containers is that many HPC environments do not support their use. Running GCHP with container-based virtualization also results in a 5% to 15% performance decrease compared to an identical build of GCHP run natively on a system. This slowdown results from both additional overhead from using Singularity or Docker and from a lack of system-optimized fabric libraries in the container images.

#### 7.4 Error and output diagnostics

Here we describe workflow improvements related to error logging and output. A major challenge for MPI Fortran software is the lack of a standard solution for error logging that exists for other languages. Thus errors in GCHP were difficult to diagnose and debug. To address these challenges, MAPL was extended to include pFlogger (<https://github.com/Goddard-Fortran-Ecosystem/pFlogger>), an MPI-aware Fortran logging system analogous to Python's "logging" package (<https://docs.python.org/3/howto/logging.html>). MAPL users initialize pFlogger with a configuration file (YAML) read at runtime that controls how various diagnostic log messages are activated, annotated, and ultimately routed to files. Per-component log verbosity can then be set to activate fine-grained debugging diagnostics or to suppress everything except serious error conditions. MAPL error messages are all routed through pFlogger and can be optionally annotated to include the MPI process rank and component name and/or split into a separate file for each MPI process.

Output diagnostics that were straightforward in GEOS-Chem Classic are more challenging in GCHP due to distribution of information across processors. A common application of GEOS-Chem has been to compare simulated performance with observations from aircraft campaigns or with monthly means of observations, because differences between the simulation and observation can identify deficiencies in the model or in scientific understanding of the atmosphere. However, GCHP originally could only output data which covered either the entire global domain or a contiguous subdomain. Samples along aircraft tracks needed to be extracted during post-processing, meaning that users had to store unnecessary data in the interim,



suffering both a performance penalty and a data storage penalty. To facilitate comparisons with observations, we add 1D output capability that allows the user to sample a collection of diagnostics according to a 1-dimensional time series of geographic coordinates, and monthly average diagnostics that account for the variable duration of each month.

## 7.5 Parallelization improvement

390 The original version of GCHP (Eastham et al., 2018) was well parallelized for simulations on up to several hundred cores (Eastham et al., 2018), and up to 1152 cores on the AWS cloud using Intel-MPI or the elastic fabric adapter (EFA) for internode communication, but suffered from a bottleneck in data input that would significantly degrade performance on a larger number of cores as described in section 3.1. Here we assess the parallelization of GCHP version 13.

We conduct 7-day timing tests on four HPC clusters: Pleiades (NASA), Amazon Web Services (AWS) EC2, Compute1  
 395 (Washington University), and Cannon (Harvard University). We focus on typical resolutions at which GCHP is run: C48, C90, C180, and C360. All four clusters use an identical model configuration, except for the number of physical cores per node. The architecture of each cluster is summarized in Table 2. We also compare MPI options and Fortran compilers.

**Table 2: Summary of architectures used to evaluate GCHP performance.**


	Pleiades	AWS EC2 <sup>1</sup>	Compute1	Cannon
CPU	Intel® Xeon® E5-2680v4	Intel® Xeon® Platinum 8124M	Intel® Xeon® Gold 6154	Intel® Xeon® Platinum 8268
Physical cores per socket	14	18	18	24
Sockets per node	2	2	2	2
Clock speed	2.4 GHz	3.00 GHz	3.00 GHz	3.50 GHz
Microarchitecture	Broadwell	Skylake	Skylake	Cascade Lake
L2/L3 cache size	256K/35840K	1024K/25344K	1024K/25344K	1024K/36608K
Interconnect	InfiniBand	AWS EFA	Infiniband	Infiniband
Storage	Lustre	AWS EBS	IBM GFPS	Lustre

400 <sup>1</sup>AWS EC2 instances used c5n.18xlarge instances.





Figure 3 shows timing test actual “wall” times. Variability across clusters reflects the effects of different architectures on performance. Tests at C180 resolution exhibit excellent scalability, with near ideal speedup across all systems up to at least a thousand cores. Tests at C90 and C48 similarly exhibit good scalability, albeit with some degradation when using several hundred cores; such large core counts at those coarse resolutions result in excessive internode communication for advection relative to computation within the node (Long et al., 2015). Tests at C360 resolution conducted on Pleiades demonstrate excellent scalability to 2304 cores, achieving 20 model days per wall day. The performance of different Fortran compilers depends on architecture, with better performance using Intel on Pleiades and better performance using GNU on Compute1. Performance on AWS is better using IntelMPI than OpenMPI at this time.




410 **Figure 3: Timing test results for GCHP version 13 at variable resolutions on multiple platforms. Grey lines indicate ideal scaling. The Fortran compiler and MPI type are indicated in parentheses, with the latter abbreviated as IMPI (IntelMPI) and OMPI (OpenMPI).**





Figure 4 shows a component-wise breakdown of wall times on the Cannon cluster with GNU compilers. Chemistry is the dominant contributor to runtime, as previously shown by Eastham et al. (2018). At C90 resolution with 192 cores, the GEOS-Chem gridded component (dominated by chemistry) accounts for 84% of the total wall time. This could be addressed in future improvements to the chemistry solver including adaptive reduction of the mechanism (Shen et al., 2021a; Shen et al., 2021b) and smart load balancing to distribute the computationally expensive sunrise/sunset gridboxes across cores and nodes (Zhuang et al., 2020). After chemistry, the next most time consuming component is advection (13%) which also scales well, albeit with some reduction in performance at high core counts that increase inter-processor communication. Data input now contributes insignificantly to the total wall time. For example, at C90 resolution with 192 cores, data input accounted for only 2.4% of the total wall time. The improvements to the MAPL input server described in section 3.2 resolve the input bottleneck that impaired the original GCHP version.



425 **Figure 4: Component-wise breakdown of the total wall times for timing tests on the Cannon cluster. Grey lines indicate ideal scaling.**



## 7.6. Cloud capability

Cloud computing is desirable for broad community access, for having a common platform where model results can be intercompared, and for dealing with surges in demand that may overwhelm local systems. Zhuang et al. (2020) deployed  
430 GCHP on the AWS cloud for easy user access and demonstrated efficient scalability with performance comparable to the NASA Pleiades supercomputer. In doing so they solved the long-standing problem of inefficient inter-node communication in the cloud (Roloff et al., 2017; Salaria et al., 2017) by using the new EFA technology now available on the AWS cloud. Zhuang et al. (2020) demonstrated the efficient scalability of GCHP on the AWS cloud on up to 1152 cores.

The basic form of a multi-node cluster on the AWS cloud as described by Zhuang et al. (2020), uses a single Elastic Block  
435 Store (EBS) volume as a temporary shared storage for all nodes. The software environment for the main node and all compute nodes is created from an Amazon Machine Image (AMI). The user logs into a main node via SSH and submits jobs to compute nodes via a job scheduler. The compute nodes form an Auto Scaling group that automatically adjusts the number of nodes based on the jobs in the scheduler queue. After finishing the computation, the user archives select data to persistent data storage (S3) and subsequently terminates the entire cluster.

440 Subsequent to Zhuang et al. (2020), EFA errors at AWS disabled the GCHP cloud capability. We identified specific conditions that cause the failure; we removed settings from the default configuration which were typically extraneous and leading to the failure. We advise users of the specific settings that will result in a failure in the user documentation. Work is in progress to fully restore the capability. Nonetheless, given the stability of the interim solution, GCHP benchmark simulations to assess model fidelity are now routinely being conducted by the GEOS-Chem Support Team on the AWS  
445 cloud.

## 8. Development and application of GEOS cubed-sphere archives

The GEOS operational meteorological archives (GEOS-FP and MERRA-2) have historically been provided only on a rectilinear latitude-longitude grid, rather than on the native cubed-sphere grid. This was intended to facilitate general georeferencing use of the GEOS data but is a drawback for GCHP because of its need to convert the latitude-longitude data  
450 back to the cubed-sphere grid during input at runtime leading to errors through regridding and restaggering. In addition, the previous operational archives included only horizontal winds rather than air mass fluxes, so that advection in GCHP required a pressure fixer to reconcile changes in air convergence and surface pressure (Jöckel et al., 2001; Horowitz et al., 2003). Here we describe the capabilities to directly use (a) mass fluxes instead of winds and (b) data on the cubed-sphere instead of latitude-longitude grid. We then describe two new archives: (a) an operational hourly archive at C720 (~12 km) resolution  
455 and (b) an hourly long-term archive at C180 resolution over 1998-present. We begin with an assessment of mass flux archival on the cubed-sphere. We then describe the data streams being generated, and their archival by the GEOS-Chem Support Team for access by GCHP users.



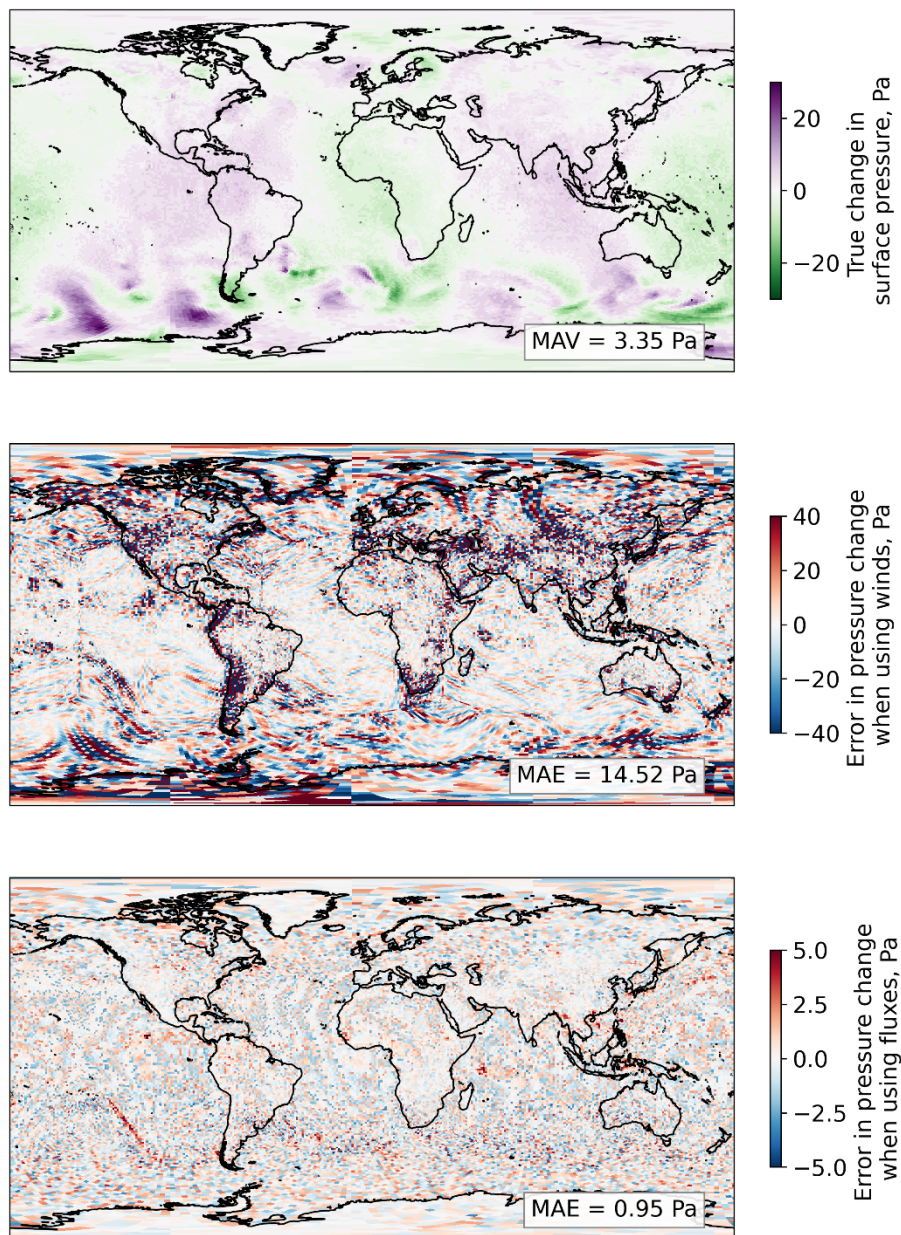
### 8.1. Mass fluxes versus winds

Standard meteorological archives include time-averaged horizontal winds and changes in surface pressure over the averaging  
460 time period of the archive, typically a few hours. A long-standing source of error in offline models has been the need to use  
the archived wind speeds to estimate the air mass fluxes between cells. As the pressure changes over the averaging time  
period, the instantaneous wind carries variable mass that is not captured by the wind speed average. In other words, the  
convergence computed from the time-averaged winds is not consistent with the archived change in surface pressure.  
Perfectly correcting for this error is impossible (Jöckel et al., 2001) although can be compensated for in offline models such  
465 as GEOS-Chem Classic by adjusting the winds with a so-called pressure fixer (Prather et al., 1987; Horowitz et al., 2003).  
But it can result in large error in vertical mass transport, which is inferred from the horizontal winds and the change in  
surface pressure. The problem can be solved by including air mass fluxes as part of the meteorological archive, but to our  
knowledge this had not previously been done for operational meteorological data products.

Figure 5 illustrates the error in surface pressure as computed from air mass convergence using either winds or air mass fluxes  
470 archived from a test GEOS C90 archive over a 5-minute timestep. The top panel shows the surface pressure tendency from  
the archive. The middle panel shows the error in this quantity when computed from the archived mass fluxes. The bottom  
panel shows the error when computed from the archived winds. We find that using mass fluxes directly rather than inferring  
them from wind data reduces the mean absolute error in the surface pressure tendency from 15 Pa to 1.0 Pa. Remaining  
errors reflect differences from water evaporation and precipitation that are implicitly included in the pressure tendency  
475 derived from the meteorological data.

The use of air mass fluxes in the meteorological archive requires a new approach for regridding. Mass fluxes are defined  
across grid cell edges, rather than at the cell center or averaged over the cell, with basis vectors that change across faces of  
the cubed-sphere. Thus if a simulation must be performed at a coarser resolution than the input data, typical regridding  
strategies such as area-conserving averaging or bilinear interpolation are not appropriate. Instead, for simulations performed  
480 at a resolution which is an integer divisor of the native-resolution data (e.g. C90 or C180 for C360), fluxes are summed. This  
is because the total flux across the edges of a grid cell at coarse resolution is the sum of the fluxes across the coincident  
edges of grid cells in the native-resolution data. Fluxes across cell edges which are not coincident are ignored, as these  
correspond to “internal” fluxes. We address this integer regridding need through a new capability for MAPL as noted in  
Section 3.1.

485 A related source of error in the original version of GCHP arose from the treatment of moisture in air mass fluxes. The  
original version of GCHP computed dry air mass fluxes for advection from winds and “dry pressures”, which needed to be  
estimated from the surface pressure and specific humidities supplied by GMAO. To reduce this error source, we implement  
into GCHP the capability to use total air mass fluxes for advection directly, thus eliminating the need for conversion.



490 **Figure 5: Illustration of the error in surface pressure change when computed from air mass convergence in an offline model using**  
**archived air mass fluxes or winds. Top: true change in surface pressure over a 5-minute time step as computed in a GEOS**  
**meteorological simulation at C90 resolution for July 1, 2019. Mean absolute value (MAV) is inset. Middle: error in the pressure**  
**change when computed using the archived air mass fluxes from that GEOS simulation. Mean absolute error (MAE) is inset.**  
**Bottom: error when the pressure change is computed from the archived winds. Note change in scale.**

495



## 8.2. Regridding and restaggering

Another source of error is the regridding and restaggering of advection data vector fields to latitude-longitude winds from the cubed-sphere mass fluxes and vice versa, operations that do not preserve the divergence of the vector field. Here we refer to wind as the advection data in an unstaggered grid formation (A-grid) with basis vectors North and East, and mass flux as the advection data in a staggered grid formation (C-grid) with local basis vectors which are perpendicular to the interfaces of the simulation grid-cells. Regridding changes the collocated grids of the vector components (i.e., A-grid) from latitude-longitude to cubed-sphere. Restaggering changes the grids of the vector components themselves; in an A-grid the grids of the vector components are collocated and identical to the simulation grid, but in a C-grid the grids of the vector components are distinct and are located at the interfaces of the discrete volumes of the simulation grid. Conceptually, the difference between a vector field on an A-grid and a C-grid is familiar as the distinction between wind (air flow in the North and East directions, defined at one location) and mass flux (air exchange between the finite volumes of the simulation grid, not defined at a single location).

To evaluate the effects of a C-grid cubed-sphere advection data (i.e., mass fluxes) versus A-grid latitude-longitude advection data (i.e., winds), we compare calculations of vertical air mass fluxes,  $J_z$ . Vertical mass fluxes are expected to be particularly sensitive to errors because they are computed from the convergence of horizontal mass fluxes. We use advection input data archives on a C180 cubed-sphere with C-grid mass fluxes and a  $0.5^\circ \times 0.625^\circ$  latitude-longitude grid with A-grid winds. Both archives were generated by the same GEOS simulation, which had a native grid of C180. All quantities on the A-grid are defined at the center of the grid cell, including both components of the wind vector, while on the C-grid the two components of the air mass flux vector are evaluated at the center of the relevant cell edge. We compare three alternative calculations of vertical mass fluxes:

$J_z(MF_{CS})$  is the vertical mass flux computed using native C180 C-grid mass fluxes; the C-grid mass fluxes are neither regridded nor restaggered.

$J_z(Wind_{CS})$  is the vertical mass flux computed using C180 A-grid winds; the original C-grid air mass fluxes are converted to winds and restaggered from C- to A-grid, and then restaggered from A- to C-grid when they are loaded in GCHP. The operations involve restaggering but no regridding.

$J_z(Wind_{LL})$  is the vertical mass flux computed using  $0.5^\circ \times 0.625^\circ$  A-grid winds; the original C-grid air mass fluxes are converted to winds on the latitude-longitude A-grid, and regridded from  $0.5^\circ \times 0.625^\circ$  to C180 and then restaggered from A- to C-grid when they are loaded in GCHP. The operations involve both restaggering and regridding.

The operations performed to the input data for  $J_z(MF_{CS})$ ,  $J_z(Wind_{CS})$ , and  $J_z(Wind_{LL})$  are summarized in Table 3.

Figure 6 compares the vertical mass flux calculations in the lower troposphere (near 900 hPa), mid-troposphere (near 500 hPa), and mid-stratosphere (near 50 hPa) for a 5-minute timestep at a nominal time (2017-03-01 12:30). In the troposphere,  $J_z(Wind_{LL})$  and  $J_z(Wind_{CS})$  both exhibit dampened upward and downward motion compared to  $J_z(MF_{CS})$ , as well as spurious noise. The dampening and noise in  $J_z(Wind_{CS})$  is significantly less than in  $J_z(Wind_{LL})$ , which is consistent with the extra





regridding operations done to the  $J_z(Wind_{LL})$  input data. The comparison of  $J_z(Wind_{CS})$  and  $J_z(MF_{CS})$  in the right column of Figure 6 demonstrates that restagging, even on the native grid, weakens vertical advection (slope=0.92). In the stratosphere where vertical motion is weak, both  $J_z(Wind_{LL})$  and  $J_z(Wind_{CS})$  are dominated by noise, reinforcing the importance of mass fluxes for vertical transport processes in the stratosphere.

**Table 3: Operations applied to GEOS advection data for input to GCHP<sup>a</sup>.**

	MF <sub>CS</sub>	Wind <sub>CS</sub>	Wind <sub>LL</sub>
GEOS data	None	$C^b \rightarrow A^c$ restagging (S) <sup>d</sup>	C → A restagging (S)
egress operations		Change of basis (M) <sup>e</sup>	Change of basis (M) CS → LL regrid (S)
GCHP data	None	Change of basis (M)	LL → CS regrid (S)
ingress operations		A → C restagging (S)	Change of basis (M) A → C restagging (S)

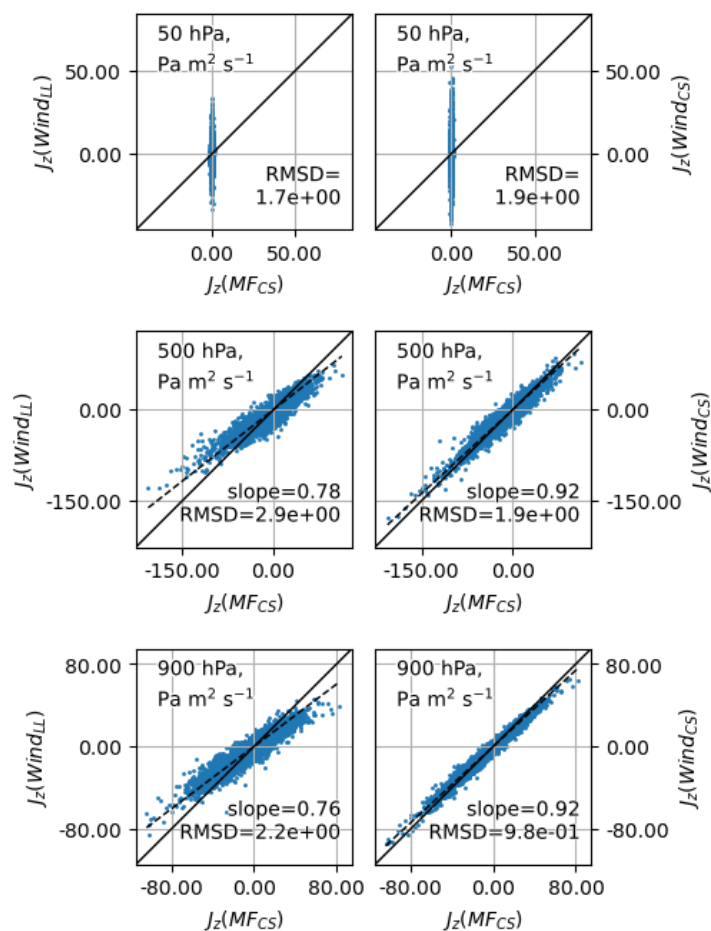
<sup>a</sup>The GEOS native data are cubed-sphere mass fluxes on the C-grid (MF<sub>CS</sub>) but are then converted in the standard archive to latitude-longitude winds on the A-grid (Wind<sub>LL</sub>). The GCHP model reconverted these Wind<sub>LL</sub> data to MF<sub>CS</sub> for input.

<sup>b</sup>Quantities on the C-grid are defined at the center of the relevant cell edge.

<sup>c</sup>Quantities on the A-grid are defined at the center of the grid cell.

<sup>d</sup>Operations that are a systematic source of error are marked with (S) and operations with machine-precision are marked with (M).

<sup>e</sup>Basis vectors differ for winds versus mass fluxes.



545 **Figure 6:** Comparison of vertical mass flux calculations at 50 hPa, 500 hPa, and 900 hPa in the global GCHP domain using different input fields for a 5-minute timestep at an example time (2017-March 1 12:30:00). Each point represents a grid cell at the corresponding pressure.  $J_z(MFCs)$  is the vertical mass flux using native C180 C-grid mass fluxes.  $J_z(WindCs)$  is the vertical mass flux using C180 A-grid winds.  $J_z(WindLL)$  is the vertical mass flux calculated using  $0.5^\circ \times 0.625^\circ$  A-grid winds. Note different scales for the different rows of panels.

### 8.3 Archive descriptions

550 Given the importance of cubed-sphere air mass fluxes for accuracy in off-line advection computations, and the previously noted need for higher temporal resolution to avoid smoothing of eddy and convective motions (Yu et al., 2018), two new cubed-sphere archives with hourly resolution are now being generated at GMAO as part of the GEOS-FP and GEOS-IT data streams. The generation of a new cubed-sphere GEOS-FP meteorological archive as a manageable operational product at GMAO is however a challenging task due to the additional output costs on top of the computationally intensive GEOS system. Despite the C720 resolution of the GEOS-FP system, the current operational archive is produced at  $0.25^\circ \times 0.3125^\circ$  resolution (corresponding to C360) with 3-hourly 3D fields including winds, because of output limitations.

555





We overcome this operational hurdle in GEOS-FP by limiting the hourly production of C720 output to the advection variables, and having those archived by the GEOS-Chem Support Team on the Washington University cluster. The cubed-sphere archive is most critical for advection variables. Other meteorological variables can be conservatively regridded from the operational  $0.25^{\circ} \times 0.3125^{\circ}$  archive. Advection requires only two 3D variables in the hydrostatic atmosphere of the GEOS system, namely the horizontal air mass fluxes and Courant numbers (to determine the number of substeps in the FV3 advection calculation), and 2D surface pressure. Currently the specific humidity is also archived, to allow accurate conversion between dry and total mass mixing ratios. This operational production of hourly C720 advection output has been ongoing in GEOS-FP since March 11, 2021, and this output is continuously being archived by the GEOS-Chem Support Team.

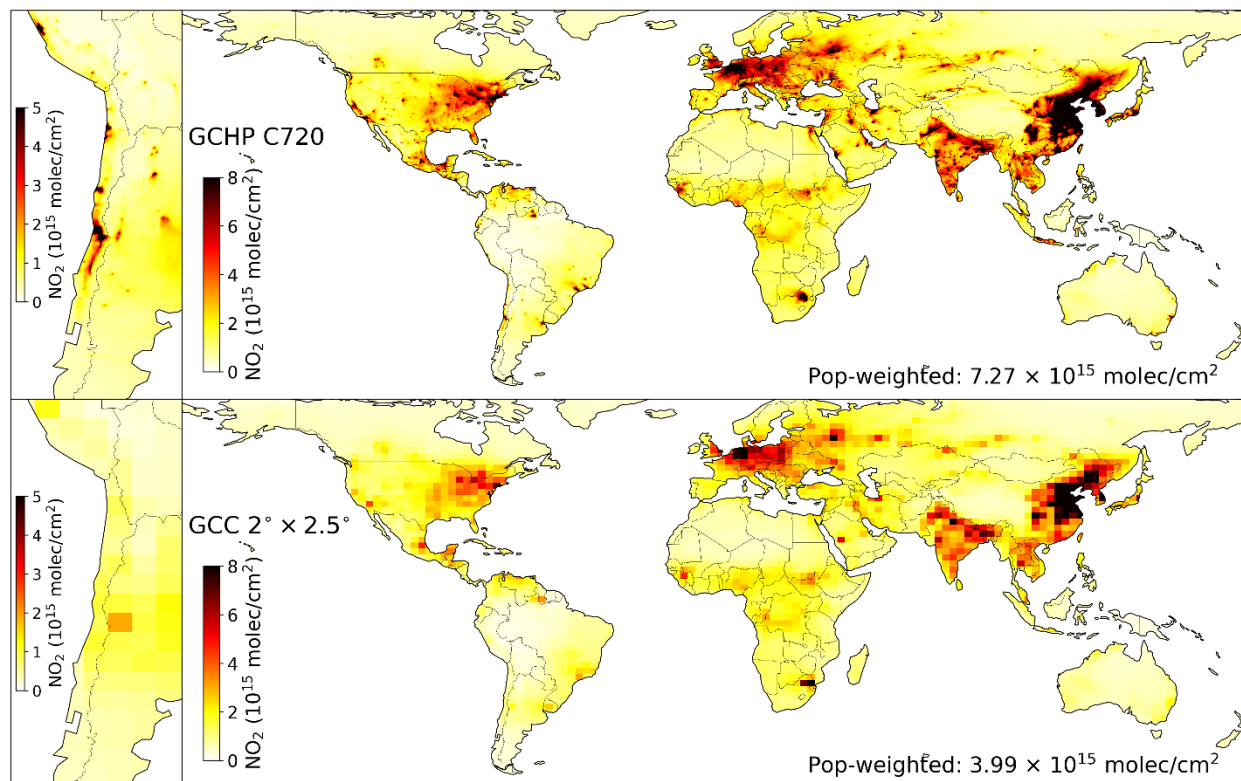
GMAO is also generating an hourly C180 full cubed-sphere GEOS-IT archive for all variables for the period 1998-present. This GEOS-IT archive will offer long-term meteorological consistency akin to the MERRA-2 archive, but on the cubed-sphere using GEOS-5.29. Both mass fluxes and winds are being archived. Two-dimensional products are also being provided on a latitude-longitude grid. This offline GEOS simulation offers the capability to archive the entire cubed-sphere dataset without the constraints of an operational system. Completion of the entire 24+ year archive expected by early 2023.

#### 8.4. Direct ingestion of GMAO meteorological data

GEOS-Chem has historically required reprocessing the GEOS meteorological data from GMAO into suitable GEOS-Chem input files. This reprocessing included modifying certain fields such as cloud optical depth into formats expected by GEOS-Chem, regridding data to coarser resolution as required by GEOS-Chem Classic, extracting regional data for pre-defined nested simulations, and flipping the vertical dimension of the arrays. We have developed the capability for GCHP to directly use the GMAO meteorological archive without modification and this will be an option in the standard model (version 13.4.0). This capability not only reduces effort, data duplication, and possible errors, but also facilitates simulations at near real time that directly read the operational post-processing and forecast data produced by GMAO.

## 9 Demonstration of technical performance

We bring together the developments described above to demonstrate the technical performance offered by GCHP. Figure 7 shows a GCHP simulation of tropospheric  $\text{NO}_2$  columns for April 9-15, 2021 using mass fluxes from the new hourly C720 GEOS-FP operational archive. Pronounced heterogeneity is apparent in tropospheric  $\text{NO}_2$  column concentrations, with clear enhancements over major urban and industrial regions. The attributes of high resolution are apparent for example along western South America, where the C720 resolution resolves distinct urban areas of Chile, Argentina, and Peru that were not evident at coarser resolution. Global population-weighted  $\text{NO}_2$  column concentrations simulated at C720 are twice those at  $2^{\circ} \times 2.5^{\circ}$ , indicating the importance of the high resolution offered by GCHP for air quality assessments. Table 4 contains statistics describing the simulations shown in Fig. 7, as conducted on the Pleiades cluster. The GCHP full chemistry simulation on 224 million grid boxes using 2,904 cores achieved a throughput of 5.2 days/day.



590 **Figure 7: Simulations of tropospheric NO<sub>2</sub> columns using GCHP at C720 resolution (top) and GEOS-Chem Classic (GCC) at 2° x 2.5° resolution (bottom). Both simulations used version 13.2.1 for the period April 9-15, 2021 following a one-week spinup.**

**Table 4: Characteristics of GCHP and GEOS-Chem Classic (GCC) simulations<sup>a</sup> in Fig. 7**

GEOS-Chem simulation	Total number of grid cells	Operator durations (minutes)	Number of physical cores	Total wall time (days)	Throughput (days/day)
GCC 2° x 2.5°	943,488	C20T10	36	0.27	52.4
GCHP C720	223,948,800	C10T5	2,904	2.69	5.20

<sup>a</sup>Simulations for April 2-15 were conducted on the Pleiades cluster. See Table 2 for the cluster architecture.

<sup>b</sup>Operator durations are represented as CcTt where *c* is the chemical operator duration and *t* is the transport operator duration.

## 10 Future needs and opportunities

The developments described above and now made available through the GEOS-Chem version 13 series increase the accessibility, accuracy, and capabilities of GCHP, but also highlight future opportunities for improvement. We identify here



595 four key opportunities to 1) further improve GCHP accessibility including on the cloud, 2) develop a tool for GCHP integration of satellite observations, 3) increase GCHP computational performance, and 4) modularize GCHP components.

**Improve GCHP accessibility including on the cloud.** There are four main areas where GCHP accessibility could be improved to benefit users. (a) Current GCHP configuration files are complicated with 12 input files, 10 file formats, redundant specification, and platform-specific settings. The need remains to simplify the process of configuring a GCHP simulation by consolidating the number of user-facing configuration files, eliminating overlap, and reducing the number of file formats. (b) The meteorological and emission input data for GCHP are extensive with over one million files available. It is challenging for users to identify and retrieve a minimum set of files needed for their simulation. This issue could be addressed with a cataloging system such as STAC (stacspe.org) which is specifically designed for cataloging Earth system data. (c) Analyzing GCHP output is currently impeded by its large data volumes. The next generation of file formats for Earth systems data such as Zarr (zarr.readthedocs.io) offers opportunities to efficiently index GCHP output data during analysis. (d) The process of setting up GCHP on the cloud is labor intensive. This could be addressed with automated pipelines for environment creation, input data synchronization, execution, and continuous testing. These developments would facilitate user exploitation of the full resources of GCHP for simulations of atmospheric composition.

**Develop tool for GCHP integration of satellite observations.** Quantitative analyses of satellite observations with a CTM require observational operators that mimic the orbit tracks, sampling schedule, and retrieval characteristics of individual satellite instruments. Developing these observational operators is presently done in an ad hoc way in the GEOS-Chem community, resulting in duplications of effort and representing an obstacle for the exploitation of satellite data. A general facility to which the community could readily contribute would allow users to select satellite orbit tracks and instrument scan characteristics, and to apply instrument vertical sensitivity such as through air mass factors and averaging kernels. This open-source library of observational operators would increase the utility of GCHP for interpretation and assimilation of satellite observations.

**Increase GCHP computational performance.** Chemistry is the most time-consuming component of GCHP calculations and remains a major barrier to the inclusion of atmospheric composition in ESMs. Two general bottlenecks currently impede performance, in GCHP and other atmospheric composition models: a) unnecessarily detailed chemical calculations in regions of simpler chemistry such as the background troposphere or stratosphere, and b) idled processors awaiting completion by a few processors of lengthy calculations at sunrise/sunset. The first bottleneck could be addressed by applying recent developments in adaptive chemical solvers for greater efficiency, and the second through smart load balancing that more efficiently allocates processors across grid boxes, thus enabling high-resolution global simulations with complex chemistry.

**Modularize GCHP components.** GCHP consists of a number of operators computing emissions, transport, radiation, chemistry, and deposition. Modularization of these operators will facilitate exchange of code with other models for both scientific benefit and good software engineering practice. This has already been done with the emissions component (HEMCO), which is now adopted in the NASA GOCART, NCAR CESM, and NOAA GFS models (Lin et al., 2021). There



630 is a strong need to generalize this practice to other GCHP modules. This will avoid redundancy and promote interoperability  
with other atmospheric composition models used by the research community, including in particular in the NASA GEOS  
system.

*Code Availability.* GCHP is publicly available at [www.geos-chem.org](http://www.geos-chem.org) with documentation at [gchp.readthedocs.io](http://gchp.readthedocs.io). The latest  
635 GCHP version (13.3.4) is available at [zenodo.org/record/5764877#.Ydr54\\_7MI2w](https://zenodo.org/record/5764877#.Ydr54_7MI2w).

*Data Availability.* GEOS-FP output is publicly available at <https://fluid.nccs.nasa.gov/weather/> and from an archive  
maintained by the GEOS-Chem Support Team at <http://geoschemdata.wustl.edu>.

*Author contributions.* RVM, DJJ, and SDE conceptualized the project. RVM, DJJ, SDE, TLC, CAK, and SP conducted  
640 funding acquisition. SDE, LB, EWL, TLC, CAK, WD, DZ, RAL, MPS, RMY, YL, LE, WMP, BMA, and ALT contributed  
to development of GCHP version 13 code and input data. LB, SDE, EWL, and DZ conducted data analysis and visualization.  
RVM, DJJ, SDE, LB, EWL, TLC, CAK, and WD wrote the manuscript with contributions from all coauthors.

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

645 *Acknowledgements.* This work was supported by NASA AIST Grant 80NSSC20K0281. Resources supporting this work  
were provided by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS)  
Division at Ames Research Center, and by local computing platforms at Washington University and Harvard University.

## References

- 650 Bey, I., D. J. Jacob, R. M. Yantosca, J. A. Logan, B. D. Field, A. M. Fiore, Q. Li, H. Y. Liu, L. J. Mickley, and M. G.  
Schultz (2001), Global modeling of tropospheric chemistry with assimilated meteorology: Model description and  
evaluation, *J. Geophys. Res.*, *106*, 23073-23096.
- Bindle, L., R. V. Martin, M. J. Cooper, E. W. Lundgren, S. D. Eastham, B. M. Auer, T. L. Clune, H. Weng, J. Lin, L. T.  
Murray, J. Meng, C. A. Keller, W. M. Putman, S. Pawson, and D. J. Jacob (2021), Grid-stretching capability for the  
655 GEOS-Chem 13.0.0 atmospheric chemistry model, *Geosci. Model Dev.*, *14*(10), 5977-5997, doi:10.5194/gmd-14-  
5977-2021.
- Brasseur, G. P., and D. J. Jacob (2017), *Modeling of Atmospheric Chemistry*, Cambridge University Press, Cambridge,  
doi:DOI: 10.1017/9781316544754.
- Eastham, S. D., M. S. Long, C. A. Keller, E. Lundgren, R. M. Yantosca, J. Zhuang, C. Li, C. J. Lee, M. Yannetti, B. M.  
660 Auer, T. L. Clune, J. Kouatchou, W. M. Putman, M. A. Thompson, A. L. Trayanov, A. M. Molod, R. V. Martin,



- and D. J. Jacob (2018), GEOS-Chem High Performance (GCHP v11-02c): a next-generation implementation of the GEOS-Chem chemical transport model for massively parallel applications, *Geosci. Model Dev.*, *11*(7), 2941-2953, doi:10.5194/gmd-11-2941-2018.
- 665 Eastham, S. D., D. K. Weisenstein, and S. R. H. Barrett (2014), Development and evaluation of the unified tropospheric–stratospheric chemistry extension (UCX) for the global chemistry-transport model GEOS-Chem, *Atmos. Environ.*, *89*, 52-63, doi:<https://doi.org/10.1016/j.atmosenv.2014.02.001>.
- Feng, X., H. Lin, T. M. Fu, M. P. Sulprizio, J. Zhuang, D. J. Jacob, H. Tian, Y. Ma, L. Zhang, X. Wang, Q. Chen, and Z. Han (2021), WRF-GC (v2.0): online two-way coupling of WRF (v3.9.1.1) and GEOS-Chem (v12.7.2) for modeling regional atmospheric chemistry–meteorology interactions, *Geosci. Model Dev.*, *14*(6), 3741-3768, doi:10.5194/gmd-14-3741-2021.
- 670 Friedman, C. L., Y. Zhang, and N. E. Selin (2014), Climate Change and Emissions Impacts on Atmospheric PAH Transport to the Arctic, *Environmental Science & Technology*, *48*(1), 429-437, doi:10.1021/es403098w.
- Gamblin, T., M. LeGendre, M. R. Collette, G. L. Lee, A. Moody, B. R. de Supinski, and S. Futral (2015), The Spack package manager: bringing order to HPC software chaos, paper presented at ACM.
- 675 Gelaro, R., W. McCarty, M. J. Suárez, R. Todling, A. Molod, L. Takacs, C. A. Randles, A. Darmenov, M. G. Bosilovich, R. Reichle, K. Wargan, L. Coy, R. Cullather, C. Draper, S. Akella, V. Buchard, A. Conaty, A. M. da Silva, W. Gu, G.-K. Kim, R. Koster, R. Lucchesi, D. Merkova, J. E. Nielsen, G. Partyka, S. Pawson, W. Putman, M. Rienecker, S. D. Schubert, M. Sienkiewicz, and B. Zhao (2017), The Modern-Era Retrospective Analysis for Research and Applications, Version 2 (MERRA-2), *J. Climate*, *30*(14), 5419-5454, doi:10.1175/jcli-d-16-0758.1.
- 680 Harris, L. M., S.-J. Lin, and C. Tu (2016), High-Resolution Climate Simulations Using GFDL HiRAM with a Stretched Global Grid, *29*(11), 4293-4314, doi:10.1175/jcli-d-15-0389.1.
- Horowitz, L. W., S. Walters, D. L. Mauzerall, L. K. Emmons, P. J. Rasch, C. Granier, X. Tie, J.-F. Lamarque, M. G. Schultz, G. S. Tyndall, J. J. Orlando, and G. P. Brasseur (2003), A global simulation of tropospheric ozone and related tracers: Description and evaluation of MOZART, version 2, *J. Geophys. Res.*, *108*, 4784, , doi:10.1029/2002JD002853.
- 685 Hu, L., C. A. Keller, M. S. Long, T. Sherwen, B. Auer, A. Da Silva, J. E. Nielsen, S. Pawson, M. A. Thompson, A. L. Trayanov, K. R. Travis, S. K. Grange, M. J. Evans, and D. J. Jacob (2018), Global simulation of tropospheric chemistry at 12.5&thinsp;km resolution: performance and evaluation of the GEOS-Chem chemical module (v10-1) within the NASA GEOS Earth system model (GEOS-5 ESM), *Geosci. Model Dev.*, *11*(11), 4603-4620, doi:10.5194/gmd-11-4603-2018.
- 690 Iacono, M. J., J. S. Delamere, E. J. Mlawer, M. W. Shephard, S. A. Clough, and W. D. Collins (2008), Radiative forcing by long-lived greenhouse gases: Calculations with the AER radiative transfer models, *Journal of Geophysical Research: Atmospheres*, *113*(D13), doi:<https://doi.org/10.1029/2008JD009944>.



- 695 Jöckel, P., R. von Kuhlmann, M. G. Lawrence, B. Steil, C. A. M. Brenninkmeijer, P. J. Crutzen, P. J. Rasch, and B. Eaton  
(2001), On a fundamental problem in implementing flux-form advection schemes for tracer transport in 3-  
dimensional general circulation and chemistry transport models, *127*(573), 1035-1052,  
doi:doi:10.1002/qj.49712757318.
- Keller, C. A., K. E. Knowland, B. N. Duncan, J. Liu, D. C. Anderson, S. Das, R. A. Lucchesi, E. W. Lundgren, J. M. Nicely,  
E. Nielsen, L. E. Ott, E. Saunders, S. A. Strode, P. A. Wales, D. J. Jacob, and S. Pawson (2021), Description of the  
700 NASA GEOS Composition Forecast Modeling System GEOS-CF v1.0, *Journal of Advances in Modeling Earth  
Systems*, *13*(4), e2020MS002413, doi:<https://doi.org/10.1029/2020MS002413>.
- Kodros, J. K., and J. R. Pierce (2017), Important global and regional differences in aerosol cloud-albedo effect estimates  
between simulations with and without prognostic aerosol microphysics, *Journal of Geophysical Research:  
Atmospheres*, *122*(7), 4003-4018, doi:10.1002/2016JD025886.
- 705 Kurtzer, G. M., V. Sochat, and M. W. Bauer (2017), Singularity: Scientific containers for mobility of compute, *PLOS ONE*,  
*12*(5), e0177459, doi:10.1371/journal.pone.0177459.
- Li, C., R. V. Martin, A. van Donkelaar, B. L. Boys, M. S. Hammer, J.-W. Xu, E. A. Marais, A. Reff, M. Strum, D. A.  
Ridley, M. Crippa, M. Brauer, and Q. Zhang (2017), Trends in Chemical Composition of Global and Regional  
Population-Weighted Fine Particulate Matter Estimated for 25 Years, *Environmental Science & Technology*,  
710 *51*(19), 11185-11195, doi:10.1021/acs.est.7b02530.
- Li, K., D. J. Jacob, H. Liao, Y. Qiu, L. Shen, S. Zhai, K. H. Bates, M. P. Sulprizio, S. Song, X. Lu, Q. Zhang, B. Zheng, Y.  
Zhang, J. Zhang, H. C. Lee, and S. K. Kuk (2021), Ozone pollution in the North China Plain spreading into the late-  
winter haze season, *Proceedings of the National Academy of Sciences*, *118*(10), e2015797118,  
doi:10.1073/pnas.2015797118.
- 715 Lin, H., X. Feng, T. M. Fu, H. Tian, Y. Ma, L. Zhang, D. J. Jacob, R. M. Yantosca, M. P. Sulprizio, E. W. Lundgren, J.  
Zhuang, Q. Zhang, X. Lu, L. Zhang, L. Shen, J. Guo, S. D. Eastham, and C. A. Keller (2020), WRF-GC (v1.0):  
online coupling of WRF (v3.9.1.1) and GEOS-Chem (v12.2.1) for regional atmospheric chemistry modeling – Part  
1: Description of the one-way model, *Geosci. Model Dev.*, *13*(7), 3241-3265, doi:10.5194/gmd-13-3241-2020.
- Lin, H., D. J. Jacob, E. W. Lundgren, M. P. Sulprizio, C. A. Keller, T. M. Fritz, S. D. Eastham, L. K. Emmons, P. C.  
720 Campbell, B. Baker, R. D. Saylor, and R. Montuoro (2021), Harmonized Emissions Component (HEMCO) 3.0 as a  
versatile emissions component for atmospheric models: application in the GEOS-Chem, NASA GEOS, WRF-GC,  
CESM2, NOAA GEFS-Aerosol, and NOAA UFS models, *Geosci. Model Dev.*, *14*(9), 5487-5506,  
doi:10.5194/gmd-14-5487-2021.
- Long, M. S., R. Yantosca, J. E. Nielsen, C. A. Keller, A. da Silva, M. P. Sulprizio, S. Pawson, and D. J. Jacob (2015),  
725 Development of a grid-independent GEOS-Chem chemical transport model (v9-02) as an atmospheric chemistry  
module for Earth system models, *Geosci. Model Dev.*, *8*(3), 595-602, doi:10.5194/gmd-8-595-2015.





- Lu, X., L. Zhang, T. Wu, M. S. Long, J. Wang, D. J. Jacob, F. Zhang, J. Zhang, S. D. Eastham, L. Hu, L. Zhu, X. Liu, and M. Wei (2020), Development of the global atmospheric chemistry general circulation model BCC-GEOS-Chem v1.0: model description and evaluation, *Geosci. Model Dev.*, 13(9), 3817-3838, doi:10.5194/gmd-13-3817-2020.
- 730 Lucchesi, R. (2017), File specification for GEOS-5 FP. GMAO Office Note No. 4 (version1.1) *Rep.*, 61 pp.
- Meng, J., R. V. Martin, P. Ginoux, M. Hammer, M. P. Sulprizio, D. A. Ridley, and A. van Donkelaar (2021), Grid-independent high-resolution dust emissions (v1.0) for chemical transport models: application to GEOS-Chem (12.5.0), *Geosci. Model Dev.*, 14(7), 4249-4260, doi:10.5194/gmd-14-4249-2021.
- Murray, L. T., E. M. Leibensperger, C. Orbe, L. J. Mickley, and M. Sulprizio (2021), GCAP 2.0: A global 3-D chemical-transport model framework for past, present, and future climate scenarios, *Geosci. Model Dev. Discuss.*, 2021, 1-52, doi:10.5194/gmd-2021-144.
- 735 National Research Council (2012), *A National Strategy for Advancing Climate Modeling*, National Academics Press, Washington DC.
- Putman, W. M., and S.-J. Lin (2007), Finite-volume transport on various cubed-sphere grids, *Journal of Computational Physics*, 227(1), 55-78, doi:<https://doi.org/10.1016/j.jcp.2007.07.022>.
- 740 Reid, P., and T. Randles (2017), Charliecloud: Unprivileged containers for user-defined software stacks in hpc." *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. ACM*.
- Roloff, E., M. Diener, L. P. Gasparly, and P. O. A. Navaux (2017), HPC Application Performance and Cost Efficiency in the Cloud, Proc. - 2017 25th Euromicro Int. Conf. Parallel, Distrib. Network-Based Process. PDP 2017, doi:10.1109/PDP.2017.59.
- 745 Salaria, S., K. Brown, H. Jitsumoto, and S. Matsuoka (2017), Evaluation of HPC-Big Data Applications Using Cloud Platforms, Proc. 17th IEEE/ACM Int. Symp. Clust. Cloud Grid Comput., 1053–1061, doi:10.1109/CCGRID.2017.143.
- Schmidt, F. (1977), Variable fine mesh in spectral global models, *Beitr. Phys. Atmos.*, 50, 211-217.
- 750 Shah, V., D. J. Jacob, C. P. Thackray, X. Wang, E. M. Sunderland, T. S. Dibble, A. Saiz-Lopez, I. Černušák, V. Kellö, P. J. Castro, R. Wu, and C. Wang (2021), Improved Mechanistic Model of the Atmospheric Redox Chemistry of Mercury, *Environmental Science & Technology*, doi:10.1021/acs.est.1c03160.
- Shen, L., D. J. Jacob, M. Santillana, K. Bates, J. Zhuang, and W. Chen (2021a), A machine learning-guided adaptive algorithm to reduce the computational cost of atmospheric chemistry in Earth System models: application to GEOS-Chem versions 12.0.0 and 12.9.1, *Geosci. Model Dev. Discuss.*, 2021, 1-17, doi:10.5194/gmd-2020-425.
- 755 Shen, L., D. Zavala-Araiza, R. Gautam, M. Omara, T. Scarpelli, J. Sheng, M. P. Sulprizio, J. Zhuang, Y. Zhang, Z. Qu, X. Lu, S. P. Hamburg, and D. J. Jacob (2021b), Unravelling a large methane emission discrepancy in Mexico using satellite observations, *Remote Sensing of Environment*, 260, 112461, doi:<https://doi.org/10.1016/j.rse.2021.112461>.





- 760 Suarez, M., A. Trayanov, C. Hill, P. Schopf, and Y. Vikhliav (2007), MAPL: a high-level programming paradigm to support more rapid and robust encoding of hierarchical trees of interacting high-performance components, paper presented at Proceedings of the 2007 symposium on Component and framework technology in high-performance and scientific computing - CompFrame '07, ACM.
- Todling, R., and A. E. Akkraoui (2018), The GMAO Hybrid Ensemble-Variational Atmospheric Data Assimilation System: Version 2.0Rep., 184 pp, NASA GSFC, Greenbelt, MD.
- 765 Wang, X., D. J. Jacob, W. Downs, S. Zhai, L. Zhu, V. Shah, C. D. Holmes, T. Sherwen, B. Alexander, M. J. Evans, S. D. Eastham, J. A. Neuman, P. R. Veres, T. K. Koenig, R. Volkamer, L. G. Huey, T. J. Bannan, C. J. Percival, B. H. Lee, and J. A. Thornton (2021), Global tropospheric halogen (Cl, Br, I) chemistry and its impact on oxidants, *Atmos. Chem. Phys.*, 21(18), 13973-13996, doi:10.5194/acp-21-13973-2021.
- 770 Weng, H., J. Lin, R. Martin, D. B. Millet, L. Jaeglé, D. Ridley, C. Keller, C. Li, M. Du, and J. Meng (2020), Global high-resolution emissions of soil NO<sub>x</sub>, sea salt aerosols, and biogenic volatile organic compounds, *Scientific Data*, 7(1), 148, doi:10.1038/s41597-020-0488-5.
- Yu, K., C. A. Keller, D. J. Jacob, A. M. Molod, S. D. Eastham, and M. S. Long (2018), Errors and improvements in the use of archived meteorological data for chemical transport modeling: an analysis using GEOS-Chem v11-01 driven by GEOS-5 meteorology, *Geosci. Model Dev.*, 11(1), 305-319, doi:10.5194/gmd-11-305-2018.
- 775 Zhang, L., L. Liu, Y. Zhao, S. Gong, X. Zhang, D. K. Henze, S. L. Capps, T.-M. Fu, Q. Zhang, and Y. Wang (2015), Source attribution of particulate matter pollution over North China with the adjoint method, *Environmental Research Letters*, 10(8), 084011, doi:10.1088/1748-9326/10/8/084011.
- Zhuang, J., D. J. Jacob, H. Lin, E. W. Lundgren, R. M. Yantosca, J. F. Gaya, M. P. Sulprizio, and S. D. Eastham (2020), Enabling High-Performance Cloud Computing for Earth Science Modeling on Over a Thousand Cores: Application to the GEOS-Chem Atmospheric Chemistry Model, *Journal of Advances in Modeling Earth Systems*, 12(5), e2020MS002064, doi:<https://doi.org/10.1029/2020MS002064>.

785