

1 **Development of a Grid-Independent GEOS-Chem Chemical Transport Model (v9-02) as an**  
2 **atmospheric chemistry module for Earth System Models.**

3  
4 M.S. Long ([mlong@seas.harvard.edu](mailto:mlong@seas.harvard.edu))

5 School of Engineering and Applied Sciences, Harvard University, Cambridge, MA USA

6 R. Yantosca ([yantasca@seas.harvard.edu](mailto:yantosca@seas.harvard.edu))

7 School of Engineering and Applied Sciences, Harvard University, Cambridge, MA USA

8 J. E. Nielsen ([jon.e.nielsen@nasa.gov](mailto:jon.e.nielsen@nasa.gov))

9 Global Modeling and Assimilation Office, NASA Goddard Space Flight Center, Greenbelt, MD;

10 Science Systems and Applications Inc., Lanham, MD

11 C.A. Keller ([ckeller@seas.harvard.edu](mailto:ckeller@seas.harvard.edu))

12 School of Engineering and Applied Sciences, Harvard University, Cambridge, MA USA

13 A. da Silva ([arlindo.m.dasilva@nasa.gov](mailto:arlindo.m.dasilva@nasa.gov))

14 Global Modeling and Assimilation Office, NASA Goddard Space Flight Center, Greenbelt, MD

15 M.P. Sulprizio ([mpayer@seas.harvard.edu](mailto:mpayer@seas.harvard.edu))

16 School of Engineering and Applied Sciences, Harvard University, Cambridge, MA USA

17 S. Pawson ([Steven.Pawson-1@nasa.gov](mailto:Steven.Pawson-1@nasa.gov))

18 Global Modeling and Assimilation Office, NASA Goddard Space Flight Center, Greenbelt, MD

19 D. J. Jacob ([djacob@seas.harvard.edu](mailto:djacob@seas.harvard.edu))

20 School of Engineering and Applied Sciences, Harvard University, Cambridge, MA USA

21

22 ***Abstract***

23 The GEOS-Chem global chemical transport model (CTM), used by a large atmospheric chemistry  
24 research community, has been re-engineered to also serve as an atmospheric chemistry module for Earth  
25 System Models (ESMs). This was done using an Earth System Modelling Framework (ESMF) interface  
26 that operates independently of the GEOS-Chem scientific code, permitting the exact same GEOS-Chem  
27 code to be used as an ESM module or as a stand-alone CTM. In this manner, the continual stream of  
28 updates contributed by the CTM user community is automatically passed on to the ESM module, which  
29 remains state-of-science and referenced to the latest version of the standard GEOS-Chem CTM. A major  
30 step in this re-engineering was to make GEOS-Chem grid-independent, i.e., capable of using any  
31 geophysical grid specified at run time. GEOS-Chem data “sockets” were also created for communication  
32 between modules and with external ESM code. The grid-independent, ESMF-compatible GEOS-Chem is  
33 now the standard version of the GEOS-Chem CTM. It has been implemented as an atmospheric chemistry  
34 module into the NASA GEOS-5 ESM. The coupled GEOS-5/GEOS-Chem system was tested for  
35 scalability and performance with a tropospheric oxidant-aerosol simulation (120 coupled species, 66  
36 transported tracers) using 48-240 cores and MPI distributed-memory parallelization. Numerical  
37 experiments demonstrate that the GEOS-Chem chemistry module scales efficiently for the number of  
38 cores tested, with no degradation as the number of cores increases. Although inclusion of atmospheric  
39 chemistry in ESMs is computationally expensive, the excellent scalability of the chemistry module means  
40 that the relative cost goes down with increasing number of cores in a massively parallel environment.

41

## 42 1. Introduction

43 Global modelling of atmospheric chemistry involves solution of the 3-D continuity equations for the  
44 concentrations of chemical species including the effects of emissions, transport, chemistry, and  
45 deposition. This is commonly done with Chemical Transport Models (CTMs) driven by input  
46 meteorological data and surface boundary conditions. CTMs are relatively simple computational tools  
47 because the chemical continuity equations are solved without coupling to atmospheric dynamics. They are  
48 adequate for many applications and play a central role in advancing knowledge of atmospheric chemistry.  
49 However, there is also increasing demand for atmospheric chemistry to be implemented as a coupled  
50 module in Earth System Models (ESMs) that represent the ensemble of processes affecting the Earth  
51 system. Here we describe a software framework through which the state-of-science GEOS-Chem CTM  
52 can be implemented seamlessly as a module in ESMs, so that the stand-alone CTM and the ESM module  
53 use exactly the same code. We describe the deployment of this capability in the NASA Goddard Earth  
54 Observing System (GEOS) developed at NASA's Global Modelling and Assimilation Office (GMAO).

55 GEOS-Chem (<http://www.geos-chem.org>) is a shared-memory parallel (OpenMP) global 3-D Eulerian  
56 CTM driven by assimilated meteorological data (Bey et al., 2001). It is used by over 100 research groups  
57 worldwide for a wide range of applications including simulation of tropospheric oxidants (Mao et al.,  
58 2013), aerosols (Fairlie et al., 2007; Jaeglé et al., 2011; Park et al., 2004; Trivitayanurak et al., 2008),  
59 carbon gases (Nassar et al., 2010; Wang et al., 2004; Wecht et al., 2014), mercury (Holmes et al., 2010;  
60 Selin et al., 2008), and stratospheric chemistry (Eastham et al., 2014; Murray et al., 2012). Development  
61 of GEOS-Chem is based on core principles of open-source code development, modular structure, nimble  
62 approach to innovation, strong version control, rigorous quality assurance (QA), extensive documentation,  
63 and user support. The large user base permits extensive model diagnosis and generates a continual stream  
64 of new developments to maintain the model at the forefront of the science. Implementation of these  
65 developments in the standard GEOS-Chem code can be done quickly and efficiently because of the  
66 simplicity of the code and the common interests of the user community. Maintaining state-of-science  
67 capability is more challenging in ESMs because of complexity of managing the central code and the need  
68 for dialogue across research communities to prioritize model development. On the other hand, CTMs such  
69 as GEOS-Chem have more difficulty staying abreast of high-performance computing (HPC) technology  
70 because of limited software engineering resources.

71 Here we present a re-engineered standard version of the GEOS-Chem CTM capable of serving as a  
72 flexible atmospheric chemistry module for ESMs. A key innovation is that GEOS-Chem is now grid-  
73 independent, i.e., it can be used with any geophysical grid. The same standard GEOS-Chem code can be  
74 integrated into ESMs through the Earth System Modelling Framework (ESMF, Hill et al., 2004)  
75 interface, or used as before as a stand-alone CTM driven by assimilated meteorological data. The re-

76 engineered grid-independent flexibility has been integrated into the standard open-code version of the  
 77 GEOS-Chem CTM. The exact same scientific code in the GEOS-Chem CTM now serves as atmospheric  
 78 chemistry module in the GEOS-5 ESM of the NASA Global Modeling and Assimilation Office (GMAO)  
 79 (Molod et al., 2012). Scientific updates to the GEOS-Chem CTM contributed by its user community and  
 80 incorporated in the standard model following QA are automatically integrated into the GEOS-5 ESM, so  
 81 that the ESM effortlessly remains state-of-science and traceable to the latest standard version of GEOS-  
 82 Chem.

83

## 84 2. *Grid-Independent GEOS-Chem Model Description*

85 The GEOS-Chem CTM consists of four modules executing operations for chemistry and dry deposition,  
 86 emissions, wet deposition, and transport (Fig. 1). GEOS-Chem solves the general Eulerian form of the  
 87 coupled continuity equations for  $m$  chemical species with number density vector  $\mathbf{n} = (n_1, \dots, n_m)^T$

88

$$89 \quad \frac{\partial n_i}{\partial t} = -\nabla \cdot (n_i \mathbf{U}) + P_i(\mathbf{n}) - L_i(\mathbf{n}) \quad i \in [1, m] \quad (1)$$

90

91 Here  $\mathbf{U}$  is the wind vector (including sub-grid components parameterized as turbulent diffusion and  
 92 convection). and  $P_i(\mathbf{n})$  and  $L_i(\mathbf{n})$  are the local production and loss rates of species  $i$  including terms to  
 93 describe chemical reactions, aerosol microphysics, emissions, precipitation scavenging, and dry  
 94 deposition. In GEOS-Chem, as in all 3-D CTMs, equation (1) is solved by operator splitting to separately  
 95 and successively apply concentration updates over finite time steps from a transport operator

96

$$97 \quad \frac{\partial n_i}{\partial t} = -\nabla \cdot (n_i \mathbf{U}) \quad i \in [1, m] \quad (2)$$

98

99 and a local operator (commonly called chemical operator)

100

$$101 \quad \frac{dn_i}{dt} = P_i(\mathbf{n}) - L_i(\mathbf{n}) \quad i \in [1, m] \quad (3)$$

102

103 The transport operator includes no coupling between species, while the chemical operator has no spatial  
 104 coupling. The transport operator is further split into 1-D advection operators, a convection operator, and a  
 105 boundary layer mixing operator. Operator splitting breaks down the multi-dimensionality of the coupled  
 106 system (1) and enables numerical solution by finite differencing. The chemical operator in GEOS-Chem  
 107 is further split into chemistry and dry deposition, emissions, and wet deposition modules for  
 108 computational convenience. Gravitational settling of particles is treated as part of the chemical operator.

109 Wet deposition from sub-grid convective precipitation cannot be decoupled from convective transport  
110 (Balkanski et al., 1993) and is treated as part of convection in the transport operator.

111 The transport operators in the standard GEOS-Chem CTM are applied on fixed latitude-longitude grids  
112 (e.g. Wu et al. 2007). When integrated into an ESM, GEOS-Chem does not need to calculate its own  
113 transport; this is done separately in the ESM as part of the simulation of atmospheric dynamics, where  
114 transport of chemical species is done concurrently with transport of meteorological variables. Thus the  
115 ESM only uses GEOS-Chem to solve the chemical operator (3) over specified time steps. The GEOS-  
116 Chem chemical operator must in turn be able to accommodate any ESM grid and return concentration  
117 updates on that grid.

118 The chemical operator has no spatial dimensionality (0-D) and could in principle be solved  
119 independently for all grid points of the ESM. However, grouping the grid points by column is more  
120 efficient as it permits simultaneous calculation of radiative transfer, precipitation scavenging,  
121 gravitational settling, and vertically distributed emissions for all grid points within the column. Thus we  
122 take a 1-D vertical column as the minimum set of grid points to be handled by a call to the chemical  
123 operator. Chemical operator updates for a given column can be completed without information from  
124 neighboring columns. Solving for the chemical operator column by column reduces memory overhead  
125 and facilitates scalable single program, multiple data (SPMD; Cotronis and Dongarra, 2001)  
126 parallelization in a distributed computing environment using the Message Passing Interface (MPI). It may  
127 sometimes be preferable to apply the chemical operator to ensembles of columns, grouped independent of  
128 geography, to balance the computational burden and achieve performance gains (Long et al., 2013).

129 Prior to this work, the horizontal grid of GEOS-Chem was defined at compile time from a limited  
130 selection of fixed latitude-longitude grids ( $1/4^\circ \times 5/16^\circ$ ,  $1/2^\circ \times 2/3^\circ$ ,  $1^\circ \times 1^\circ$ ,  $2^\circ \times 2.5^\circ$ ,  $4^\circ \times 5^\circ$ ) compatible with  
131 the advection module and offline meteorological fields. Our goal here was to re-engineer the existing  
132 GEOS-Chem code to accept any horizontal grid defined at runtime. The horizontal grid would be able to  
133 span the entire global domain, represent a single column to be calculated on a single compute node, or  
134 represent any collection of columns defined by their location. This permits use of the same scientific code  
135 for stand-alone CTM and coupled ESM applications.

136

## 137 *2.1 Code Modularization and Structure*

138 In order for the GEOS-Chem code to permit run-time horizontal grid definition, much of the  
139 FORTRAN-77 code base was updated to Fortran-90. This included extensive conversion of static to  
140 dynamically-allocatable arrays, and introduction of pointer-based derived data types. Data flow into,  
141 through, and out of GEOS-Chem's routines was reconfigured to use derived-type objects passed to  
142 routines as arguments in place of publicly-declared global-scope variables. This permitted the bundling of

143 data structures with similar functionality into common interfaces (data “sockets”) that simplify module  
144 communication within GEOS-Chem and coupling to external components through the ESMF interface  
145 (see Section 2.2). Three sockets are defined: a meteorology and physics socket, a chemistry socket, and  
146 an input options socket. The meteorology and physics socket provides data defining geophysical state  
147 variables and arrays. This includes temperature, pressure, humidity, wind fields, and many others. The  
148 chemistry socket provides data structures for chemical species including indexing, species names, and  
149 concentrations. The input options socket provides runtime information such as calendar, grid dimensions,  
150 diagnostic definitions, and locations of offline information stored on disk. Together, these sockets  
151 incorporate all of the quantities and fields necessary for coupling to and driving modules within GEOS-  
152 Chem.

153 The GEOS-Chem code includes specific conditional-compilation flags to accommodate the ESMF  
154 interface and permit coupling with external data streams. These flags do not interfere with GEOS-Chem’s  
155 scientific operation and are used exclusively in grid, I/O, and utility operations. There are three flags  
156 invoked as C-preprocessor statements: *ESMF\_*, *EXTERNAL\_GRID*, and *EXTERNAL\_FORCING*. Code  
157 bounded by these flags is neither compiled nor executed unless the specific flag is enabled at compile  
158 time. The *ESMF\_* flag bounds code specific for the ESMF. The *EXTERNAL\_GRID* flag bounds code that  
159 allows GEOS-Chem to operate on an externally defined and initialized grid (e.g. by an ESM). The  
160 *EXTERNAL\_FORCING* flag bypasses GEOS-Chem’s internal, offline data I/O operations necessary for  
161 CTM operation, and replaces them with ESMF-based I/O. Users do not need to have the ESMF installed  
162 in order to run GEOS-Chem as a stand-alone CTM. The system reverts to the standard GEOS-Chem  
163 CTM code relying on the legacy module interface when compiled without these flags enabled.

164 The recently developed Harvard-NASA Emissions Component HEMCO ([http://wiki-geos-  
165 chem.org/HEMCO/](http://wiki-geos-chem.org/HEMCO/)) is used for emission calculations (Keller et al., 2014). HEMCO is a Fortran-90  
166 based, ESMF compliant, highly customizable module that uses base emissions and scale factors from a  
167 library of emission inventories to construct time-dependent emission field arrays. Emission inventories  
168 and scale factors are selected by the user in a HEMCO-specific configuration file. Emission inventories  
169 for different species and source types need not be of the same grid dimensions or domain.

170 The redesign of GEOS-Chem’s data structures was meant to simplify coupling of GEOS-Chem with  
171 any ESM regardless of its ESMF compatibility. In the absence of an ESMF interface, users would be  
172 required to engineer a specific interface for their ESM. However, GEOS-Chem’s data sockets and  
173 conditional-compilation flags facilitate this task by having all input and output data structures and  
174 associated methods conveniently located in a few specific modules.

175 As with all modifications to the publicly-available GEOS-Chem source code, changes made for ESM  
176 coupling and grid-independence were subject to rigorous QA by conducting prescribed 1-month and 1-

177 year test simulations as benchmarks ([http://acmg.seas.harvard.edu/geos/geos\\_benchmark.html](http://acmg.seas.harvard.edu/geos/geos_benchmark.html)), and  
178 comparing results to the benchmarks of the previous model version. Our changes were not expected to  
179 modify any aspect of the benchmark simulation results and we verified that they did not. Results from the  
180 benchmark simulations for version 9-02k can be found here: [http://wiki.seas.harvard.edu/geos-](http://wiki.seas.harvard.edu/geos-chem/index.php/GEOS-Chem_v9-02_benchmark_history#v9-02k)  
181 [chem/index.php/GEOS-Chem\\_v9-02\\_benchmark\\_history#v9-02k](http://wiki.seas.harvard.edu/geos-chem/index.php/GEOS-Chem_v9-02_benchmark_history#v9-02k).

182

183

## 184 2.2 ESMF Interface

185 We made GEOS-Chem ESMF-compatible for interfacing with external ESMs. The ESMF is an open-  
186 source software application programming interface that provides a standardized high-performance  
187 software infrastructure for use in ESM design. It facilitates HPC, portability, and interoperability in Earth  
188 science applications (Collins et al., 2005).

189 GEOS-Chem is executed within the ESMF as a gridded component. The gridded component is the basic  
190 element of an ESMF-based program, and is defined as a set of discrete scientific and computational  
191 functions that operate on a geophysical grid. Likewise, other components of the Earth system are  
192 implemented as gridded components (e.g. atmospheric dynamics, ocean dynamics, terrestrial  
193 biogeochemistry, etc.).

194 Each gridded component consists of a routine establishing ESMF-specific services, and *Initialize*, *Run*,  
195 and *Finalize* operations methods for gridded component execution by the ESMF. The *Initialize* method is  
196 executed once at the beginning of the simulation and initializes component-specific runtime parameters.  
197 The *Run* method interfaces local data structures with ESMF States (see below) and executes the  
198 component code (GEOS-Chem in our case). The *Finalize* method wraps up code execution, closes any  
199 remaining open files, finalizes I/O and profiling processes, and flushes local memory.

200 Gridded components exchange information with each other through States. A State is an ESMF derived  
201 type that can contain multiple types of gridded and non-gridded information (Collins et al., 2005; Suarez  
202 et al., 2013). An ESMF gridded component is associated with an *Import* State and an *Export* State. The  
203 *Import* State provides access to data created by other gridded components. The *Export* State contains data  
204 that a component generates and makes available to other components. In the ESMF-enabled GEOS-  
205 Chem, data are passed into and out of the GEOS-Chem gridded component via interfacing an appropriate  
206 State with a corresponding GEOS-Chem data socket (Figure 1), making these data available within  
207 GEOS-Chem or to other ESM gridded components (see Section 2.1).

208 The ESMF was implemented within GEOS-Chem as an independent layer that operates on top of the  
209 CTM code. It includes code for interfacing with and executing GEOS-Chem as an ESMF gridded  
210 component. When coupling GEOS-Chem to an ESM, the GEOS-Chem transport modules are excluded

211 and only those modules necessary to solve Eq. (3) are used. Coupling specifically to the GEOS-5 ESM  
212 required an adaptation of GEOS-Chem’s ESMF interface for the GMAO’s Modelling, Analysis and  
213 Prediction Layer (MAPL) extension (Suarez et al., 2013). MAPL is otherwise not required for GEOS-  
214 Chem.

215

### 216 *3 Implementation, Performance, and Scalability*

217 The ESMF-enabled GEOS-Chem was embedded within the NASA GEOS-5 ESM (version Ganymed-  
218 4.0). The GEOS-5 ESM is the forward model of the GEOS-5 atmospheric data assimilation system  
219 (GEOS-DAS) (Ott et al., 2009; Rienecker et al., 2008). The system is built upon on an ESMF framework,  
220 and uses a combination of distributed memory (MPI) and, in some cases, hybrid distributed/shared  
221 memory parallelization. The dynamical core used here is based on Lin (2004), and operates on horizontal  
222 grid resolutions ranging from  $2^\circ \times 2.5^\circ$  to  $0.25^\circ \times 0.3125^\circ$ , with 72 vertical layers up to 0.01 hPa. Ocean  
223 surface and sea-ice boundaries are prescribed. The land and snow interfaces are based on Koster et al.  
224 (2000) and Stieglitz et al. (2001), respectively. For the coupled simulations, GEOS-5 ESM native  
225 dynamics and moist physics are applied to the GEOS-Chem chemical tracers.

226 All coupled GEOS-5/GEOS-Chem simulations were performed on the Discover system at the NASA  
227 Goddard Space Flight Center ([http://www.nccs.nasa.gov/discover\\_front.html](http://www.nccs.nasa.gov/discover_front.html)), using 12-core (dual hex-  
228 core) 2.8 GHz Intel Xeon Westmere (X5660) compute nodes equipped with 24 GB RAM, and an  
229 Infiniband DDR interconnect using the Intel compiler suite (v. 13.1.1) and MVAPICH2 (v. 1.8.1). GEOS-  
230 Chem’s shared-memory (OpenMP) parallelization was disabled.

231 The coupled GEOS-5/GEOS-Chem system was tested on  $2^\circ \times 2.5^\circ$  and  $0.5^\circ \times 0.625^\circ$  grids with a standard  
232 oxidant-aerosol simulation using 120 chemical species of which 66 are transported (“chemical tracers”).  
233 Radical species with very short chemical lifetimes are not transported. The chemistry module used the  
234 RODAS-3 (4-stage, order 3(2), stiffly accurate) solver with self-adjusting internal time step (Hairer and  
235 Wanner, 1996) as part of the Kinetics Pre-processor (KPP, Eller et al., 2009; Sandu and Sander, 2006).  
236 KPP was implemented with its supplied linear algebra (BLAS Level-1) routines in place. The  $2^\circ \times 2.5^\circ$   
237 simulation used a time step of 1800 seconds for all operations. For the  $0.5^\circ \times 0.625^\circ$  simulation, chemistry  
238 and system-operation time steps were both 450 seconds. Dynamics, physics, and radiation time steps were  
239 900 seconds. For both simulations, the atmosphere used 72 vertical hybrid-sigma (pressure) levels.  
240 Simulations were run for 31 days initialized on July 1, 2006. All chemical tracers were initialized from  
241 output of a GEOS-Chem CTM (v9-02) simulation.

242 The  $2^\circ \times 2.5^\circ$  coupled simulations were used to test scalability of the coupled system and for comparison  
243 to the GEOS-Chem CTM. We conducted simulations with 48, 96, 144, 192, and 240 total MPI processes  
244 operating on  $12 \times 4$ ,  $12 \times 8$ ,  $12 \times 12$ ,  $16 \times 12$ , and  $16 \times 15$  (lat x lon) contiguous grid point subdomains,



245 respectively. This represents a set of five simulations  $j \in [1, 5]$ . For comparison, the offline GEOS-Chem  
246 CTM (v9-02) was run on 8 shared-memory processes at  $2^\circ \times 2.5^\circ$  resolution using 8-core 2.6 GHz Intel  
247 Xeon processors, reflecting a typical CTM set-up, using otherwise identical settings and initial chemical  
248 conditions as the coupled GEOS-5/GEOS-Chem simulations. Since GEOS-5 is a pure MPI application,  
249 each MPI process corresponds to a single processor core.

250 Figure 2 gives execution wall times for the total simulation and for the chemistry (GEOS-Chem) and  
251 dynamics gridded components. To analyze the performance and scalability results, we define the  
252 normalized scaling efficiency  $S$  for simulation  $j$  relative to simulation  $j-1$  as

253

$$254 \quad S = \left( \frac{W_{x,j-1} - W_{x,j}}{W_{x,j-1}} \right) \left( \frac{N_j}{N_j - N_{j-1}} \right) \quad (4)$$

255

256 where  $W_{x,j}$  is the wall time for component  $x$ , and  $N_j$  is the number of cores allocated to the simulation.  $S$   
257 measures how efficiently the addition of computational resources speeds up execution. For example, a  
258 value of 0.9 indicates that a doubling of computational resources decreases wall time by a factor of 1.8. A  
259 value of zero means no speed-up. A negative value means slow-down, as might result from increasing  
260 I/O. Results for 48 cores ( $j = 1$ ) are given relative to the 8-process GEOS-Chem CTM simulation ( $j = 0$ ),  
261 which uses different shared-memory processes and a different transport code for chemical tracers only.  
262 The two simulations are not strictly comparable but results serve to benchmark the performance of the  
263 GEOS-5/GEOS-Chem system against the GEOS-Chem CTM.

264 We find that the scaling efficiency for the chemistry module (GEOS-Chem) in the GEOS-5/GEOS-  
265 Chem system is  $0.78 \pm 0.10$  for the range of cores tested. This represents excellent performance, with no  
266 decline as the number of cores increases, reflecting the independent nature of the chemistry calculation  
267 for individual columns. For that reason, we expect the excellent scalability of the chemistry module to  
268 extend to any number of cores. Scaling efficiency of the dynamics component decreases with increasing  
269 number of cores and becomes negative above 192. This reflects the small number of gridpoints allocated  
270 to individual cores increasing relative cost of communicating between processes versus operating within  
271 local memory, as well as a greater inter-nodal communication associated with additional chemical tracers.  
272 The results further suggest that the chemistry module would remain efficient for simulations beyond the  
273 range of values tested.

274 The  $0.5^\circ \times 0.625^\circ$  resolution simulation was used to examine the performance of the GEOS-5/GEOS-  
275 Chem system when operating on a finer grid resolution than permitted by the GEOS-Chem CTM using  
276 shared-memory OpenMP parallelization. The higher resolution also increases the problem size, permitting  
277 the efficient use of more computing power. For this simulation, the horizontal grid was decomposed into

278 24x25 lat/lon blocks over 600 cores. The  $0.5^\circ \times 0.625^\circ$  resolution simulation completed 0.35 simulation  
279 years per wall day.

280 About 20% of the wall time spent on chemistry in the GEOS-5/GEOS-Chem system was spent copying  
281 and flipping the vertical dimension of chemical tracer arrays between the GEOS-5 ESM and GEOS-  
282 Chem. This would be overcome to a large extent by linking GEOS-Chem tracer arrays to the ESMF using  
283 pointers, which access memory locations of preexisting variables directly. This cannot be done within the  
284 GEOS-5 ESM for two reasons: (1) GEOS-Chem stores concentrations in double-precision arrays, while  
285 the GEOS-5 system generally uses single precision. (2) GEOS-Chem indexes concentration arrays  
286 vertically from the surface of the Earth upward while the GEOS-5 system does the reverse. Such  
287 limitations are not intrinsic to GEOS-Chem and depend on the specific ESM to which GEOS-Chem is  
288 coupled; other ESMs may use different data precision and indexing. Further software engineering in  
289 GEOS-Chem could add flexibility in array definitions to accommodate different ESM configurations.

290 Figure 3 illustrates model results with 500 hPa  $O_3$  mixing ratios at 12 UT on July 15, 2006 for GEOS-  
291 5/GEOS-Chem simulations at  $2^\circ \times 2.5^\circ$  and  $0.5^\circ \times 0.625^\circ$  resolutions, and for the GEOS-Chem CTM using  
292 GEOS-5 assimilated meteorological data at  $2^\circ \times 2.5^\circ$  resolution. All three simulations are initialized from  
293 the same GEOS-Chem CTM fields at 0 UT on July 1, 2006, but have different meteorology because of  
294 differences in resolution and also because the CTM uses assimilated meteorological data while the  
295 GEOS-5/GEOS-Chem system in this implementation does not. The Figure demonstrates the fine structure  
296 of chemical transport that can be resolved with the  $0.5^\circ \times 0.625^\circ$  resolution. The general patterns are  
297 roughly consistent between simulations and are reasonable compared to satellite and sonde observations  
298 (Zhang et al., 2010). A scatterplot comparing output from the different simulations (Figure 4) shows that  
299 they have comparable results. Figures 3 and 4 are intended to illustrate the GEOS-5/GEOS-Chem  
300 capability. A more thorough evaluation of GEOS-Chem's chemistry within the GEOS-5 system would  
301 require the use of the same meteorological data as the offline CTM, diagnosing the full ensemble of  
302 simulated chemical species, and investigating the effect of transport errors when using off-line  
303 meteorological fields in the CTM. This will be documented in a separate publication.

304

#### 305 4 *Summary*

306 We have presented a new grid-independent version of the GEOS-Chem chemical transport model  
307 (CTM) to serve as atmospheric chemistry module within Earth system models (ESMs) through the Earth  
308 System Modelling Interface (ESMF). The new GEOS-Chem version uses any grid resolution or geometry  
309 specified at runtime. The exact same standard GEOS-Chem code (freely available from [http://geos-  
310 chem.org](http://geos-chem.org)) supports both ESM and stand-alone CTM applications. This ensures that the continual stream

311 of innovation from the worldwide community contributing to the stand-alone CTM is easily incorporated  
312 into the ESM version. The GEOS-Chem ESM module thus always remains state-of-science.

313 We implemented GEOS-Chem as an atmospheric chemistry module within the NASA GEOS-5 ESM  
314 and performed a tropospheric oxidant-aerosol simulation (120 coupled chemical species, 66 transported  
315 tracers) in that fully coupled environment. Analysis of scalability and performance for 48 to 240 cores  
316 shows that the GEOS-Chem atmospheric chemistry module scales efficiently with no degradation as the  
317 number of cores increases, reflecting the independent nature of the chemical computation for individual  
318 grid columns. Although the inclusion of detailed atmospheric chemistry in an ESM is a major  
319 computational expense, chemistry operations become relatively more efficient as the number of cores  
320 increases due to their efficient scalability.

321

322 **Code Availability.** GEOS-Chem source code is freely available to the public. Source code may be  
323 downloaded by following instructions found at [http://wiki.seas.harverd.edu/geos-chem](http://wiki.seas.harvard.edu/geos-chem). At time of  
324 writing, this work used a modified version of GEOS-Chem version 9-02k as indicated in the text. All  
325 developments presented here are now included with the current GEOS-Chem version 10-01f.

326

327 **Acknowledgments.** This work was supported by the NASA Modelling, Analysis and Prediction (MAP)  
328 Program. The authors thank Ben Auer (NASA-GMAO) and Jack Yatteau (Harvard University) for  
329 technical assistance.

## References:

- Balkanski, Y. J., Jacob, D. J., Gardner, G. M., Graustein, W. C. and Turekian, K. K.: Transport and residence times of tropospheric aerosols inferred from a global three-dimensional simulation of  $^{210}\text{Pb}$ , *J. Geophys. Res. Atmospheres*, 98(D11), 20573–20586, doi:10.1029/93JD02456, 1993.
- 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), 23073–23, 2001.
- Collins, N., Theurich, G., DeLuca, C., Suarez, M., Trayanov, A., Balaji, V., Li, P., Yang, W., Hill, C. and Silva, A. da: Design and Implementation of Components in the Earth System Modeling Framework, *Int. J. High Perform. Comput. Appl.*, 19(3), 341–350, doi:10.1177/1094342005056120, 2005.
- Cotronis, Y. and Dongarra, J.: Recent Advances in Parallel Virtual Machine and Message Passing Interface: 8th European PVM/MPI Users' Group Meeting, Santorini/Thera, Greece, September 23-26, 2001. Proceedings, Springer., 2001.
- 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, doi:10.1016/j.atmosenv.2014.02.001, 2014.
- Eller, P., Singh, K., Sandu, A., Bowman, K., Henze, D. K. and Lee, M.: Implementation and evaluation of an array of chemical solvers in the Global Chemical Transport Model GEOS-Chem, *Geosci. Model Dev.*, 2(2), 89–96, 2009.
- 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, doi:10.1016/j.atmosenv.2006.09.048, 2007.
- Hairer, E. and Wanner, G.: Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems, Springer., 1996.
- Hill, C., DeLuca, C., Suarez, M., da Silva, A. and others: The architecture of the earth system modeling framework, *Comput. Sci. Eng.*, 6(1), 18–28, 2004.
- Holmes, C. D., Jacob, D. J., Corbitt, E. S., Mao, J., Yang, X., Talbot, R. and Slemr, F.: Global atmospheric model for mercury including oxidation by bromine atoms, *Atmos Chem Phys*, 10(24), 12037–12057, doi:10.5194/acp-10-12037-2010, 2010.

Jaeglé, L., Quinn, P. K., Bates, T. S., Alexander, B. and Lin, J.-T.: Global distribution of sea salt aerosols: new constraints from in situ and remote sensing observations, *Atmospheric Chem. Phys.*, 11(7), 3137–3157, doi:10.5194/acp-11-3137-2011, 2011.

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, *Geosci. Model Dev. Discuss.*, 7(1), 1115–1136, doi:10.5194/gmdd-7-1115-2014, 2014.

Koster, R. D., Suarez, M. J., Ducharme, A., Stieglitz, M. and Kumar, P.: A catchment-based approach to modeling land surface processes in a general circulation model: 1. Model structure, *J. Geophys. Res. Atmospheres*, 105(D20), 24809–24822, doi:10.1029/2000JD900327, 2000.

Lin, S.-J.: A “vertically Lagrangian” finite-volume dynamical core for global models, *Mon. Weather Rev.*, 132(10), 2293–2307, 2004.

Long, M. S., Keene, W. C., Easter, R., Sander, R., Kerkweg, A., Erickson, D., Liu, X. and Ghan, S.: Implementation of the chemistry module MECCA (v2.5) in the modal aerosol version of the Community Atmosphere Model component (v3.6.33) of the Community Earth System Model, *Geosci. Model Dev.*, 6(1), 255–262, doi:10.5194/gmd-6-255-2013, 2013.

Mao, J., Paulot, F., Jacob, D. J., Cohen, R. C., Crouse, J. D., Wennberg, P. O., Keller, C. A., Hudman, R. C., Barkley, M. P. and Horowitz, L. W.: Ozone and organic nitrates over the eastern United States: Sensitivity to isoprene chemistry, *J. Geophys. Res. Atmospheres*, 118(19), 2013JD020231, doi:10.1002/jgrd.50817, 2013.

Molod, A., Takacs, L., Suarez, M., Bacmeister, J., Song, I.-S. and Eichmann, A.: The GEOS-5 atmospheric general circulation model: Mean climate and development from MERRA to Fortuna. [online] Available from: <http://ntrs.nasa.gov/search.jsp?R=20120011790> (Accessed 16 June 2014), 2012.

Murray, L. T., Jacob, D. J., Logan, J. A., Hudman, R. C. and Koshak, W. J.: Optimized regional and interannual variability of lightning in a global chemical transport model constrained by LIS/OTD satellite data, *J. Geophys. Res. Atmospheres*, 117(D20), D20307, doi:10.1029/2012JD017934, 2012.

Nassar, R., Jones, D. B. A., Suntharalingam, P., Chen, J. M., Andres, R. J., Wecht, K. J., Yantosca, R. M., Kulawik, S. S., Bowman, K. W., Worden, J. R., Machida, T. and Matsueda, H.: Modeling global atmospheric CO<sub>2</sub> with improved emission inventories and CO<sub>2</sub> production from the oxidation of other carbon species, *Geosci Model Dev Discuss*, 3(3), 889–948, doi:10.5194/gmdd-3-889-2010, 2010.

- Ott, L. E., Bacmeister, J., Pawson, S., Pickering, K., Stenchikov, G., Suarez, M., Huntrieser, H., Loewenstein, M., Lopez, J. and Xueref-Remy, I.: Analysis of Convective Transport and Parameter Sensitivity in a Single Column Version of the Goddard Earth Observation System, Version 5, General Circulation Model, *J. Atmospheric Sci.*, 66(3), 627–646, doi:10.1175/2008JAS2694.1, 2009.
- Park, R. J., Jacob, D. J., Field, B. D., Yantosca, R. M. and Chin, M.: Natural and transboundary pollution influences on sulfate-nitrate-ammonium aerosols in the United States: Implications for policy, *J. Geophys. Res. Atmospheres*, 109(D15), D15204, doi:10.1029/2003JD004473, 2004.
- Rienecker, M. M., Suarez, M. J., Todling, R., Bacmeister, J., Takacs, L., Liu, H. C., Gu, W., Sienkiewicz, M., Koster, R. D., Gelaro, R., Stajner, I. and Nielsen, J. E.: The GEOS-5 Data Assimilation System- Documentation of Versions 5.0.1, 5.1.0, and 5.2.0. [online] Available from: <http://ntrs.nasa.gov/search.jsp?R=20120011955> (Accessed 16 June 2014), 2008.
- Sandu, A. and Sander, R.: Technical note: Simulating chemical systems in Fortran90 and Matlab with the Kinetic PreProcessor KPP-2.1, *Atmos Chem Phys*, 6(1), 187–195, doi:10.5194/acp-6-187-2006, 2006.
- Selin, N. E., Jacob, D. J., Yantosca, R. M., Strode, S., Jaeglé, L. and Sunderland, E. M.: Global 3-D land-ocean-atmosphere model for mercury: Present-day versus preindustrial cycles and anthropogenic enrichment factors for deposition, *Glob. Biogeochem. Cycles*, 22(2), GB2011, doi:10.1029/2007GB003040, 2008.
- Stieglitz, M., Ducharne, A., Koster, R. and Suarez, M.: The Impact of Detailed Snow Physics on the Simulation of Snow Cover and Subsurface Thermodynamics at Continental Scales, *J. Hydrometeorol.*, 2(3), 228–242, 2001.
- Suarez, M., Trayanov, A., da Silva, A. and Chakraborty, P.: MAPL Manual, [online] Available from: [http://geos5.org/wiki/images/f/fa/MAPL\\_UsersGuide.pdf](http://geos5.org/wiki/images/f/fa/MAPL_UsersGuide.pdf), 2013.
- Trivitayanurak, W., Adams, P. J., Spracklen, D. V. and Carslaw, K. S.: Tropospheric aerosol microphysics simulation with assimilated meteorology: model description and intermodel comparison, *Atmos Chem Phys*, 8(12), 3149–3168, doi:10.5194/acp-8-3149-2008, 2008.
- Wang, J. S., Logan, J. A., McElroy, M. B., Duncan, B. N., Megretskaia, I. A. and Yantosca, R. M.: A 3-D model analysis of the slowdown and interannual variability in the methane growth rate from 1988 to 1997, *Glob. Biogeochem. Cycles*, 18(3), GB3011, doi:10.1029/2003GB002180, 2004.

Wecht, K. J., Jacob, D. J., Frankenberg, C., Jiang, Z. and Blake, D. R.: Mapping of North American methane emissions with high spatial resolution by inversion of SCIAMACHY satellite data, *J. Geophys. Res. Atmospheres*, 119(12), 2014JD021551, doi:10.1002/2014JD021551, 2014.

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. Atmospheres*, 112(D5), D05302, doi:10.1029/2006JD007801, 2007.

Figure 1. Coupling between the GEOS-Chem CTM (dashed beige box) and an ESM (blue box). The schematic shows how the coupling is managed through the ESMF, and utilizes only the GEOS-Chem components bound by the ESM box: Transport modules in the GEOS-Chem CTM are bypassed and replaced by the ESM transport modules through the atmospheric dynamics simulation .

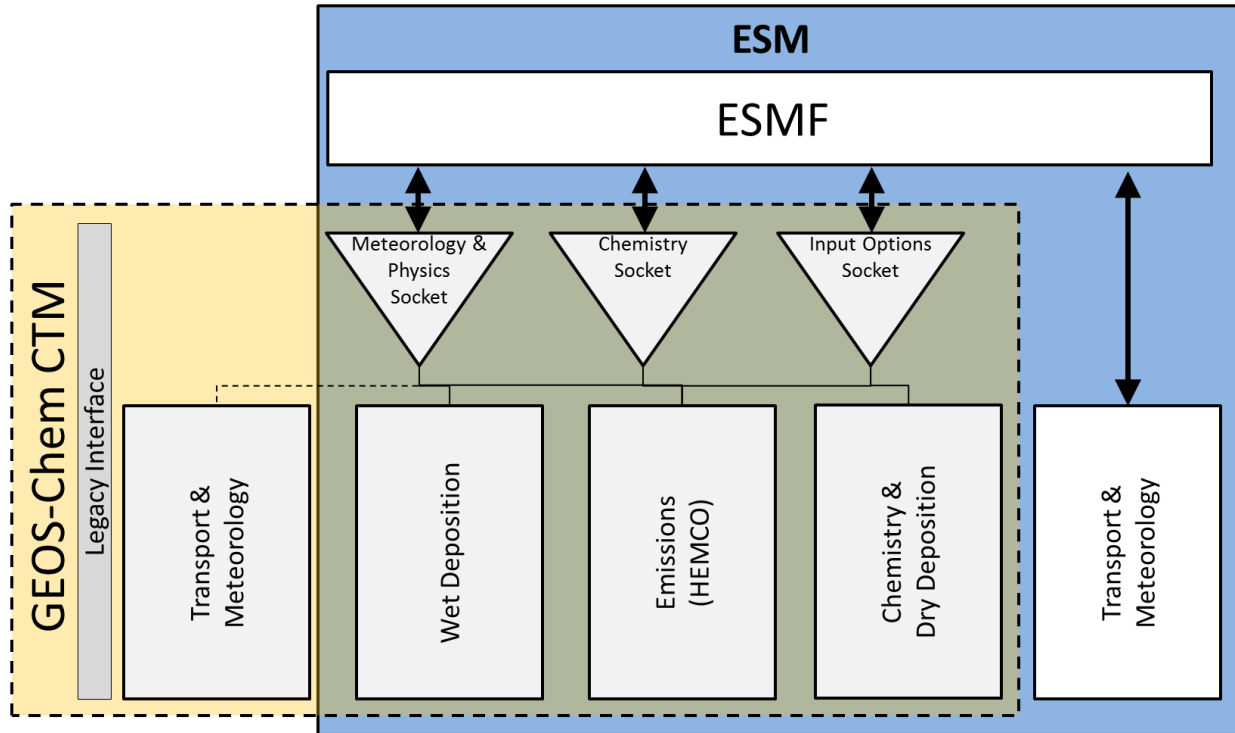




Figure 2. Performance and scalability of the GEOS-5/GEOS-Chem system for a 1-month test simulation including detailed oxidant-aerosol tropospheric chemistry at  $2^\circ \times 2.5^\circ$  horizontal resolution. Top panel: total and stacked wall-times for the chemical operator (GEOS-Chem), dynamics, and other routines versus number of processor cores. Bottom panel: Scaling efficiency (Eq. 4) for chemistry, dynamics, and the full GEOS-5/GEOS-Chem system. Values shown for 48 cores are relative to the 8-process shared-memory GEOS-Chem CTM.




Figure 3. Instantaneous 500 hPa ozone mixing ratios ( $\text{nmol mol}^{-1}$ ) at 12 UT on July 15, 2006, for CTM and ESM implementations of GEOS-Chem. Top panel: GEOS-Chem CTM at  $2^\circ \times 2.5^\circ$  resolution driven by GEOS-5 assimilated meteorological data with  $0.5^\circ \times 0.67^\circ$  resolution. Middle panel: GEOS-5/GEOS-Chem ESM at  $2^\circ \times 2.5^\circ$  resolution. Bottom panel: GEOS-5/GEOS-Chem ESM at  $0.5^\circ \times 0.625^\circ$  resolution. All three simulations are initialized with the same GEOS-Chem CTM fields at 0 UT on July 1, 2006, but the ESM as implemented here does not include meteorological data assimilation.

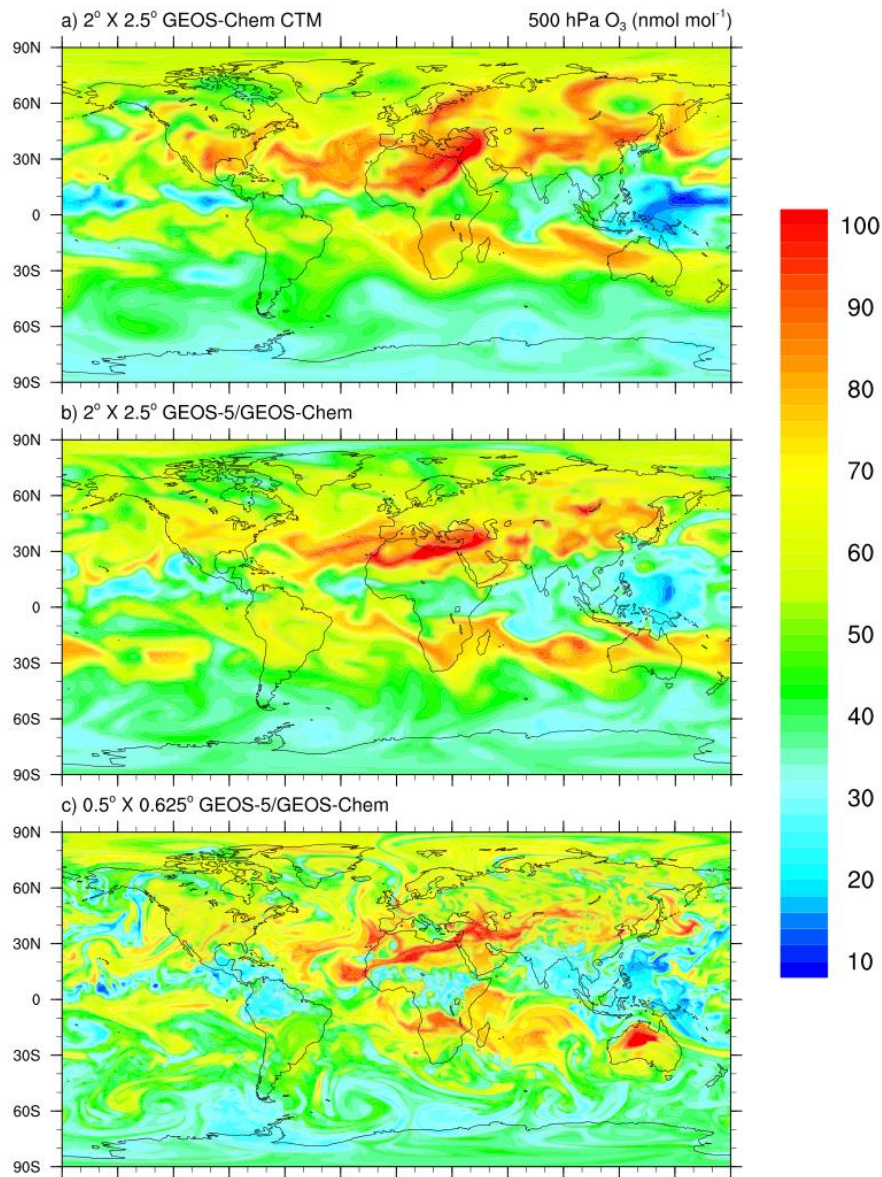


Figure 4. Comparison of instantaneous 500 hPa ozone mixing ratios ( $\text{nmol mol}^{-1}$ ) at 12 UT on July 15, 2006 in the stand-alone GEOS-Chem simulation at  $2^\circ \times 2.5^\circ$  horizontal resolution and the coupled GEOS-5/GEOS-Chem simulation at  $2^\circ \times 2.5^\circ$  (red) and  $0.5^\circ \times 0.625^\circ$  (blue) resolutions. The  $0.5^\circ \times 0.625^\circ$  results are regridded to  $2^\circ \times 2.5^\circ$  resolution, and each point represents a  $2^\circ \times 2.5^\circ$  grid square. The reduced-major-axis regression parameters and the 1:1 line are also shown.

