GEOS-Chem Newsletter, Winter 2014/2015
05 Dec 2014

The GEOS-Chem Support Team
geos-chem-support@as.harvard.edu

Registration for IGC7 is now open!

We would like to invite you to register for the 7th International GEOS-Chem Meeting (aka IGC7), which will take place at Harvard University from May 4-7, 2015. You will find the latest information about the IGC7 meeting at our web site igc7.geos-chem.org.

We are in the process of collating lodging information, and will post that on igc7.geos-chem.org shortly.

We hope to see you in May 2015!

GEOS-Chem Steering Committee News

Decmeber 2014 GCSC telecon

Please take a moment to read the minutes from the Dec 5th GCSC telecon. The next GCSC telecon will take place before the IGC7 meeting (date TBD, Feb-Mar 2015).

GEOS-Chem Support Team News

Please welcome our newest GCST members!

We would like to welcome two new GEOS-Chem Support Team members:

- **Matt Yannetti** joined the GCST on October 1, 2014. Matt has a Bachelor's degree in Computer Science and a Master's degree in Earth and Atmospheric Science, both from Georgia Tech.
- **Lizzie Lundgren** will join the GCST on December 8, 2014. Lizzie has undergraduate degrees in both Physics and Environmental Engineering Science from MIT, plus a Master's degree in Civil and Environmental Engineering, also from MIT.

Matt and Lizzie will be working on GEOS-Chem software development as well as user support.

Please submit new GEOS-Chem publications to Junwei Xu

GCST team member Junwei Xu is now coordinating the GEOS-Chem publications list on ResearcherId. If you have any GEOS-Chem publications that you wish to add to the list, please contact Junwei at `jn231250[at]dal.ca`.

GEOS-Chem v10-01 news

GEOS-Chem v10-01e is now approved!

We are pleased to announce that GEOS-Chem v10-01e was approved on 01 Dec 2014. The major update in v10-01e is the HEMCO emissions component, which radically changes the way in which emissions are treated in GEOS-Chem. HEMCO greatly simplifies the process of combining and overwriting different emissions inventories. The implementation of HEMCO has also allowed us to remove many legacy modules and routines from GEOS-Chem (a complete list may be found here).

The validation process for v10-01e took longer than we had anticipated. Along the way we found and fixed many minor issues. We also gave HEMCO some additional functionality, such as the capability to read concentration data (kg/m3) and prod/loss data (kg/m3/s).

The GCSC identified a couple of unusual features that were evident in the v10-01e 1-year benchmark plots. But we were able to account for these features after a thorough investigation. Here is our summary report:

Christoph Keller wrote:

> As of November 2014, there were two open questions regarding the [GEOS-Chem v10-01e] 1-yr benchmark (at least to my knowledge):
>
> 1. The ozone blob forming over Europe in January [identified by Shiliang Wu]
> 2. The much lower concentration of NITs and SO4s in v10-01e compared to v10-01c.
>
> I think I have an explanation for both points, and I would argue that the new code makes more sense physically:
>
> 1. The ozone blob is actually a good thing as it removes an issue that was introduced in GEOS-Chem v9-02r. As Melissa [Sulprizio] pointed out to me, ever since v9-02r we have been seeing a drastic decrease of ozone over Europe (and to a lesser extent East Asia) in January. This was because [the] PARANOX [ship plume model] could produce—under low O3 conditions—incredibly high O3 deposition values (division of a high number by a very, very small number). This would basically create a black hole for ozone at the surface that would then remove more and more ozone from the entire column. Melissa did a number of sensitivity studies where she ran v10-01c with and without PARANOX, and those plots show this effect very nicely. The new code does not allow such a behavior anymore, which explains the higher ozone concentrations in January compared to v10-01e.
>
> 2. NITs and SO4s [(i.e. nitrate and sulfate on sea salt aerosol)] are closely linked to sulfate chemistry and thus alkalinity. Alkalinity is directly obtained from the total mass of emitted sea salt aerosol, e.g. we approximate alkalinity = SSA. Since SSA is a 2D field but alkalinity is 3D, the sea salt aerosol needs to be distributed vertically. In the old code, alkalinity was set to the total mass of emitted sea salt aerosol for every grid box that was below the PBL (zero otherwise). In other words, if the PBL was at level 11, the integrated alkalinity over this column was 10 times higher than the total mass of emitted

> sea salt aerosols in the same column. I found that not very realistic and therefore changed the alkalinity calculation so that it takes into account the fraction of the PBL. I also checked in with Fabien Paulot who has been working on this and he also thinks that this is the better way of doing it. After all, it ensures that the total alkalinity equals the total amount of emitted sea salt aerosols. This explains the much lower values for NITs and SO4s that we see in v10-01e compared to v10-01c.

Daniel Jacob wrote:

> After some clever detective work, Melissa and Christoph found the causes of the benchmark changes in v10-01-e for O3 (January mid-latitudes), NITs, and SO4s. It turns out that in both cases the implementation of HEMCO corrected bugs in previous versions of the model. PARANOX had a bug that caused ozone loss to go crazy under very low ozone conditions, and sea salt alkalinity in the MBL column was not added correctly. v10-01e corrects these bugs and seems fairly neutral in terms of comparison to observations.
>
> There being no further areas of concern, v10-01e is approved. UCX is very consistent with Trop-only in the 1-month benchmarks and in past 1-year benchmarks, so I don't see the need to run a [UCX] 1-year benchmark for [v10-01e]. We will run UCX benchmarks more systematically once we get stratospheric metrics—Dylan Jones is working on those.

More about the HEMCO emissions component

HEMCO can now read in the following types of data files:

- Emissions data (kg/m2/s)
- Concentration data, e.g. OH or O3 oxidant (kg/m3)
- Prod/loss data (kg/m3/s)
- Regional masks (unitless)
- Scale factors (unitless)

As of this writing, HEMCO is compatible with most of GEOS-Chem's specialty simulations:

Feature	Status
Full-chemistry	Completed
SOA simulations	Completed
UCX simulations	Completed
Aerosol-only simulations	Completed
CH4 simulation	Completed
CO2 simulation	Completed
O3 (tagged) simulation	Completed
POPs simulation	Completed
Rn-Pb-Be simulation	Completed
Diagnostic archiving	HEMCO now tracks emissions data for the ND01, ND06, ND07, ND08, ND11, ND13, ND28, ND29, ND30, ND32, ND34, ND36, ND46, ND56, and ND63 time-averaged diagnostics.

CO (tagged) simulation	We still need to do the following: • Create mask files used to define tagged CO regions • Create a HEMCO configuration file for the CO simulation
Hg simulation	Oxidant fields are now read in via HEMCO, but we still need to: • Convert data files from binary (bpch) to netCDF format • Create mask files used to define tagged Hg regions
TOMAS simulations	We still need to do the following: • Track emissions of size-resolved dust with HEMCO • Track emissions of size-resolved sea salt with HEMCO • Track emissions of size-resovled carbon aersols w/ HEMCO • Better interface HEMCO w/ TOMAS diagnostic outputs

Lastly, The journal article describing HEMCO has now been published in Geoscientific Model Development (aka GMD). The citation is:

- C. A. Keller, M. S. Long, R. M. Yantosca, A. M. Da Silva, S. Pawson, and D. J. Jacob, *HEMCO v1.0: A versatile, ESMF-compliant component for calculating emissions in atmospheric models*, Geosci. Model Devel. , **7**, 1409-1417, 2014. [Link to Article]

Updates still "in the pipeline" for GEOS-Chem v10-01

The table below lists some of the important features that have been added into GEOS-Chem v10-01, and those which are next in line.

Version	Date	Features
v10-01e	12/01/14	• Incorporating the Harvard-NASA Emissions Component (HEMCO) • Removed many modules that were rendered obsolete by HEMCO • Code updates for nested 0.25^o x 0.3125^o CH grid • Reactivating dust tracers in TOMAS • Fix for ND61 diagnostic in TOMAS
v10-01f	TBD	• RRTMG online radiative transfer in GEOS-Chem • Introduction of flexible precision to GEOS-Chem • Correct ALD2 photolysis in FAST-JX v7.0 • Read 2D input data for individual NOx species in UCX simulations • Add bug fixes & updates to tagged CO simulation
v10-01g	TBD	• Two-way coupling between global and nested GEOS-Chem grids
v10-01h	TBD	• Adding more emissions options to HEMCO, including: ◦ FINN biomass burning emissions) ◦ Updates to ship NOx chemistry via PARANOX ◦ Update EMEP emissions for 2008-2010 ◦ MASAGE NH3 inventory ◦ HTAP and hi-res Asian emissions ◦ Improved temporal resolution of anthropogenic CO2 sources ◦ Replace CASA CO2 biosphere climatology with year-specific fluxes ◦ UNEP 2010 anthropogenic Hg emissions

		○ Historical anthropogenic atmospheric emissions of Hg ○ NEI 2008 emissions with hourly resolution Science ○ Update to MEGAN 2.1 biogenic emsissions ○ Update to EDGAR 4.2 anthropogenic emissions ○ NEI 2008 Hg emission ○ QFED emissions

Most of the emissions updates in v10-01h can be done very quickly, as they just involve telling HEMCO where to find the various netCDF input files. For expediency's sake, we may also consider folding v10-01g and v10-01h into a single 1-month benchmark.

GEOS-Chem simulations on the 0.25° x 0.3125° China and SE Asia nested grids

Yuxuan Wang and colleagues have given GEOS-Chem the capability to perform simulations using GEOS-FP meteorology at the 0.25° x 0.3125° resolution for the nested-grid China region.

Yuxuan Wang wrote:

> I'm happy to report that we've finally implemented all the bug fixes and necessary input files to successfully run the 0.25 nested_CH full chemistry simulation on the Tsinghua platform. We've done a 10-day test run and the simulation results look reasonable, so I thought it's a good time to submit our code updates to the GC Support Team so that they can be incorporated in the next code release.
>
> Developers for China and Southeast Asia windows: Patrick Kim (Harvard), Libao Chai (Tsinghua), Yu Yao (Tsinghua) and Yuxuan Wang (Tsinghua and Texas A&M Galveston).

We have added Yuxuan's source code updates into GEOS-Chem v10-01e. We are still in the process of converting the data files to netCDF for input via HEMCO. We will have this done before the v10-01 public release.

Correcting errors in the stratospheric Br_y data files

Johan Schmidt has provided us with updated stratospheric Br_y data files for input into GEOS-Chem. These updates correct problems with the current data files.

Johan Schmidt wrote:

> Just to follow up on the bug in the external data used to prescribe Strat Bry in GEOS-Chem.
>
> There are two bugs:
>
> 1) Serious: Br_y fields used for 4° x 5° GEOS-4 runs is too high by a factor of 4.
> The problem can be removed by copying the corrected data files:
> `Bry_Stratosphere_*.bpch.geos4.4x5` to the

`GEOS_4x5/bromine_201205/CCM_stratosphere_Bry/` folder.

2) Less serious: Night Br_y data is equal to day Br_y data. This bug is found in the GEOS-5 2^o x 2.5^o, GEOS-5 1^o x 1^o and all GEOS-4 Strat Br_y fields. The problem is not present GEOS-5 4x5 Strat Br_y fields. We now have corrected daytime and nighttime strat Bry fields as netCDF files for the 2^o x 2.5^o horizontal grid and GEOS-5/MERRA/GEOS_FP vertical grid.

We will apply these updates to GEOS-Chem v10-01f.

Two-way nesting capability for GEOS-Chem

Jintai Lin's group at Peking University has implemented two-way nesting into GEOS-Chem. This lets the tracer concentrations that are computed on the nested grid feed back into the coarse simulation that is used to generate the transport boundary conditions.

Jintai Lin wrote:

> The two-way coupled GEOS-Chem is finally available, based on v9-02. We have provided Git patches to the GC Support Team. I think the code is ready to be incorporated in v10.1.
>
> In short, we developed a two-way coupler, PKUCPL, to implement the coupled process. Specifically, the global model (2^o x 2.5^o or 4^o x 5^o) outputs results to be LBCs of nested models, and the three nested models (at 0.5^o x 0.667^o; covering Asia, North America and Europe) output high-resolution results to update global model simulation within respective nested domains. Users can choose to couple any numbers of nested models that best suit their purposes, though.
>
> My PhD student, Yingying Yan, did the hard work on two-way coupling, and will help with questions/inquiries from users. Please feel free to contact us if you have any two-way related issues.

We will add this feature into GEOS-Chem v10-01g.

Other fundamental software development efforts

We are adding a flexible precision definition into GEOS-Chem

Flexible precision is a language feature that was introduced in Fortran-90. It allows you to define GEOS-Chem's floating-point variables with 4 bytes (aka REAL*4) or with 8 bytes (aka REAL*8) when you compile the code.

To implement flexible precision into GEOS-Chem, we first defined a precision parameter `fp`, as follows:

```
MODULE PRECISION_MOD

  IMPLICIT NONE
  PRIVATE
```

```
#if defined( USE_REAL8 )

  ! Use 8-byte floating point precision when asked for it
  INTEGER, PARAMETER, PUBLIC :: fp = KIND( 0.d0 )

#else

  ! Use 4-byte floating point precision by default
  INTEGER, PARAMETER, PUBLIC :: fp = KIND( 0.0  )

#endif

END MODULE PRECISION_MOD
```

This `fp` parameter will be set to specify either an 8-byte precision or a 4-byte precision, depending on whether or not you set the C-preprocessor switch `USE_REAL8` at compile time.

Once we defined the `fp` parameter, we used it to declare floating-point variables. Throughout the GEOS-Chem code, we changed declarations such as:

```
REAL*8   :: A, B   ! 8-byte floating point
```

With

```
USE PRECISION_MOD
. . .
REAL(fp) :: A, B   ! Flexible precision
```

We also replaced constant declarations such as:

```
A = 0.0d0
```

with

```
A = 0.0E+00_fp
```

Having this flexible precision definition in GEOS-Chem will allow us to interface with the NASA GEOS-5 GCM more efficiently. The GEOS-5 GCM uses 4-byte floating point variables, but GEOS-Chem has traditionally used 8-byte variables. When we run GEOS-Chem within the GEOS-5 GCM, we have to copy the 4-byte variables from the GCM into GEOS-Chem's 8-byte variables. This copying procedure incurs a 20% performance penalty. By giving GEOS-Chem the ability to use 4-byte variables, we can simply point to the GCM's variables, thus eliminating the performance bottleneck.

Matt Yannetti has been doing the lion's share of the work implementing the flexible precision. We expect to incorporate flexible precision into GEOS-Chem v10-01f.

NOTE: Some floating-point variables in GEOS-Chem may still have to be declared with 8-byte precision in order to ensure computational accuracy or to avoid numerical errors (e.g. overflow, underflow, infinity, etc.)

Removal of ISORROPIA ATE package

Jessica Kunke (Harvard) and Sebastian Eastham (MIT) are working on replacing the current ISORROPIA aerosol thermodynamical equilibrium package with the ATE package from the MOZAIC model. ISORROPIA uses legacy coding practices—such as common block storage—that are incompatible with high-performance computing environments. The work has just begun.

Implementation of Flexchem with KPPA solver into GEOS-Chem

John Linford (ParaTools, Inc.) and Mike Long have been working towards the implementation of KPP-Accelerated (KPPA) into GEOS-Chem. KPPA is a newer, faster version of the previous KPP chemical solver software. Preliminary tests show significant speedup (50-70%) over the existing SMVGEAR chemical solver software. The work is ongoing.

Making units consistent throughout GEOS-Chem

GEOS-Chem has traditionally carried tracers in units of kg, and then has converted them to v/v mixing ratio and molec/cm3 in other areas of the code as needed. But these unit conversions rely upon the surface area of each grid cell. GEOS-Chem currently computes surface areas assuming a Cartesian grid. On the other hand, the latest generation of Earth System Models—including NASA's GEOS-5 GCM—often use cubed-sphere grids. This renders the assumption of a Cartesian grid invalid.

We are planning to recode GEOS-Chem such that tracer and species concentrations are carried as mass per unit volume, such as kg/m3. Doing so will:

- Remove the burden of having to know what the surface area of each grid box is,
- Remove several commonly-repeated multiplications and divisions, thus increasing computational efficiency,
- Facilitate coupling GEOS-Chem with the GEOS-5 GCM and other Earth System Models

Lizzie Lundgren will begin working on this project shortly.

GEOS-FP met field update

Available met data at Dalhousie and Harvard

GEOS-Chem Support Team member Junwei Xu has been archiving the GEOS-FP met fields on the `rain.ucis.dal.ca` server at Dalhousie University. You may now download GEOS-FP data from April 2012 thru October 2014 for the following horizontal grids:

- 4^o x 5^o global
- 2^o x 2.5^o global
- 0.25^o x 0.3125^o nested grids:
    a. CH: China (70^o E – 140^o E; 15^o N – 55^o N)
    b. EU: Europe (15^o W – 40^o E; 32.75^oN – 61.25^o N)
    c. NA: North America (130^o W – 60^o W; 9.75^oN – 60^o N)
    d. SE: Southeast Asia (75^o E – 130^o E; 10^o N – 30^o N)

We are also storing the 4^o x 5^o global, 2^o x 2.5^o global and 0.25^o x 0.3125^o NA nested-grid data on the `ftp.as.harvard.edu` server at Harvard University.

Archiving GEOS-FP data at 0.5^o x 0.625^o global resolution

The Nested Model Working Group, along with the Dalhousie University group, is planning to develop the degraded nested-grid capability over East Asia and North America. To facilitate this effort, Junwei Xu at has begun downscaling the GEOS-FP met data from its original 0.25^o x 0.3125^o resolution to 0.5^o x 0.625^o resolution. Downscaled met fields for July 2014 – October 2014 are available on the Dalhousie FTP server (`rain.ucis.dal.ca`).

High-performance computing with GEOS-Chem

The ESMF/MPI version of GEOS-Chem now compiles and runs!

We have succeeded in compiling and running GEOS-Chem in a high-performance computing (HPC) environment. GEOS-Chem can now utilize the Earth System Modeling Framework (ESMF) with Message-Passing Interface (MPI) parallelization to run on distributed computing architectures. It also uses the MPI-enabled NASA/GMAO finite-volume dynamics core (FVdycore) as the transport operator.

Bob Yantosca wrote:

> I've finally gotten the ESMF/MPI standalone GEOS-Chem (code-name: "Mega-Chem") to build in one step on the Harvard Odyssey supercomputer!
>
> The GEOS-Chem base code that I used was cloned from the Dev branch of our GEOS-Chem "bleeding edge" repository. This was updated last at 12:14 PM on Friday Nov 21, so it truly is the latest and greatest code. It is essentially the same version that we used to run the 1-yr benchmarks plus a few minor updates….The important thing to note is that the GEOS-Chem code base we used to build the Mega-Chem is EXACTLY the same code base that we use for our traditional GEOS-Chem simulations.
>
> Even better, I created an installer package that automatically downloads the latest GEOS-Chem code plus the ESMF+MAPL+FVdycore code to your computer in a single step. After you download, you can compile Mega-Chem by adding in the hpc target to the traditional compilation commands, i.e.:
>
> ```
> make j4 MET=geosfp GRID=4x5 TRACEBACK=y hpc
> ```
>
> The first time you compile Mega-Chem, the process will build and install ESMF, MAPL, and the FVdycore libraries, plus all of the GEOS-Chem code. The next time you compile Mega-Chem, it doesn't have to rebuild the ESMF+MAPL+FVdycore because that's already been done.

So I think we are in a good position to start some Mega-Chem test simulations on Odyssey. The download and installation process is also simple enough so that other power-users can take the code and start testing on their end.

At present, we have gotten the HPC-enabled GEOS-Chem to run on both the Odyssey computer at Harvard, as well as the Discover supercomputer at NASA. Further testing is ongoing.

The Grid-Independent GEOS-Chem is now running in the NASA GEOS-5 DAS!

Christoph Keller and NASA Global Modeling and Assimilation Office (GMAO) personnel are currently running GEOS-Chem's chemistry modules plus the HEMCO emissions component within the NASA GEOS-DAS system. The GEOS-DAS system consists of the NASA GEOS-5 GCM plus a data assimilation module that ingests real-world observations from satellites, sondes, and various other sources.

In this configuration, GEOS-Chem is currently getting its meteorological input data directly from the GEOS-DAS. It uses these meteorological fields as input to its chemical mechanism, which computes concentrations of ozone and other trace gases. The computed ozone field is then fed back into the GEOS-DAS so that it can influence the computation of the assimilated meteorology for the next timestep. Further testing is ongoing.

Bob Yantosca
on behalf of the entire GEOS-Chem Support Team
05 Dec 2014