Parallel Performance Analysis of Two Infrastructure Frameworks for GMI Chemistry

Jules Kouatchou, Thomas Clune, Benjamin Auer, J. Eric Nielsen and Andrea Molod

NASA Goddard Space Flight Center, Greenbelt, Maryland, USA

A generalized Chemistry Transport Model (CTM) named the Goddard Earth Observing System CTM (GEOS CTM), developed to facilitate the integration of several Chemistry modules within a single code base [3], will be evaluated here in terms of computational performance. GEOS CTM was designed to leverage the infrastructure development of the GEOS General Circulation Model (GCM) and the GEOS Chemisty Climate Model (CCM) [5]. It relies on the Earth System Modeling Framework (ESMF) [2] and includes the advection only segment (AdvCore) of the Putman and Lin dynamical core [6] based on a cubed-sphere grid and as well as a software component (called ExtData) that has the ability to read external files at any resolution (horizontal and temporal) and from any source (as long as the data are on a latitude-longitude grid). In particular, the GEOS CTM experiments to be discussed here are configured to run with the software components that comprise the Global Modeling Initiative (GMI) CTM (Emission, Deposition, Diffusion, Convection and Chemistry). GMI has been in general use in the community for a long time [1], and is driven by a multi-dimensional flux form semi-Lagrangian advection scheme (named LL-TpCore) [4] on a latitude-longitude grid. The use of such a grid formulation has shown some limitations at high spatial resolutions and with increasing number of transported tracers in GMI. The GMI configuration of GEOS CTM (that we now refer to as GEOS CTM) contains the components (Emission, Deposition, Diffusion, Convection and Chemistry) that originated in GMI. GMI has 124 tracers (of which 69 are advected), whereas GEOS CTM has 122 tracers (72 of which are advected). GEOS CTM has fewer tracers because the GMI Chemistry is also used in the GEOS CCM which does not require the full set of tracers (as in GMI). Overall, GMI and GEOS CTM share 121 identical tracers and associated chemical reactions.

We carry out a series of one-day numerical integrations of GMI (at $1^\circ \times 1.25^\circ$horizontal resolution or 288×181 grid points) and GEOS CTM (at C90 horizontal resolution or $90 \times 90 \times 6$ grid points) to analyze the parallel performance of both implementations.

We are interested in comparing the computational speed of GEOS CTM with respect to GMI and in particular the performance of their most computationally intensive components. All the experiments described here were performed on Intel Xeon Haswell processor nodes where each node has 28 cores (2.6 GHz each) and 128 Gb of available memory.

In Figure 1, we plot the total wall clock times for both models as the number of processors varies. The integration with GEOS CTM is faster than that of GMI, and the wall clock time steadily decreases as the number of processors increases. If we examine the distribution of the overall wall clock time, we note that Advection and Chemistry take together up to 90% of the time in GMI and up to 76% in GEOS CTM. We plot the wall clock times (in milliseconds) per tracer per grid point for Advection (Figure 2) and Chemistry (Figure 4) in both models. We also present the parallel scaling of each of the two components. Figure 2 shows that AdvCore is at least four times faster than LL-TpCore and AdvCore has a more favorable parallel scalability (Figure 3). The clustering of grid points at the poles in LL-TpCore combined with the parallel domain decomposition approach demand significant data communication between processors covering areas around the poles [7]. This results in an increased overhead and load imbalances. The quasi-uniform nature of the cubed-sphere grid in AdvCore eliminates such problems at the poles. In addition, the fact that data are in single precision in AdvCore and double precision in LL-TpCore, makes AdvCore less computationally demanding. We expect AdvCore's relative parallel efficiency to further improve, as one continues to increase the GEOS CTM horizontal resolution.

Chemistry in GEOS CTM costs less per tracer per grid point (Figure 4). However, as shown in Figure 5 the scaling is better with GMI because it has more grid points that contribute to a better load balancing in its Chemistry.

As the number of processors increases, LL-TpCore slowly becomes the more dominant component in terms of run-time in GMI and quickly reaches a limit on its parallel efficiency. At some given model resolution, using more processors in GMI will no longer provides any performance improvement because of the communication overhead in LL-TpCore. AdvCore in GEOS CTM does not suffer from have such issues.

We have observed that, the ExtData component, critical in the GEOS CTM implementation, contributes to an increase of the overall wall clock time. The reason for this is at least in part because of the overheads for opening separate files for each required meteorological field as well for automatic regridding. GMI avoids these run-time overheads (at the loss of some flexibility) by requiring data to be prepared in advance at the model resolution and aggregated into a single file for a given time slice. As shown in Table 1, the time spent in ExtData remains basically flat regardless of the number of processors used. As the number of processor increases, ExtData might be the most dominant component. To limit the number of file openings by ExtData, we have started the process of modifying the code by opening each external file only once. Instead of having at least 50 file openings (one for each meteorological related field), there are now only 11 (one for each file collection). We conducted additional experiments with the modified code and recorded the ExtData wall clock times. The results in Table 1 show at least a 30% reduction of time reading external data files. We are currently exploring other options to further decrease the ExtData time requirement.

Our experiments show that the most time consuming components of GMI (Advection and Chemistry) are faster in GEOS CTM. Overall, GEOS CTM is a better alternative to GMI as far as the parallel performance is concerned. The advantage of GEOS CTM is not limited to its computational speed. It does not require any pre-processing (for regridding) of input data files, can be run without any restart input file, and has the ability to integrate various CTMs using the same executable.

References

[1] B. Duncan, S. Strahan, Y. Yoshida, S. Steenrod, and N. Livesey. Model study of the cross-tropopause transport of biomass burning pollution. *Atmos Chem Phys.*, 7:3713–3736, 2007.

[2] C. Hill, C. DeLuca, V. Balaji, M. Suarez, and A. da Silva. The architecture of the Earth System Modeling Framework. *Computing in Science and Engineering*, 11(6):18–28, 2004.

[3] J. Kouatchou, A. Molod, J. E. Nielsen, B. Auer, W. Putman, and T. Clune. GEOS-5 CTM User's Guide. GMAO Office Note 7, NASA Goddard Space Flight Center, Greenbelt, Maryland, 2015.

[4] S.-J. Lin and R. B. Rood. Multidimensional flux form semi-lagrangian transport schemes. *Mon. Wea. Rev.*, 124, 1996.

[5] A. Molod, L. Takacs, M. Suarez, and J. Bacmeister. Development of the GEOS-5 atmospheric general circulation model: evolution from MERRA to MERRA2. *Geosci. Model Dev.*, 8:1339–1356, 2015.

[6] W. M. Putman and S.-J. Lin. Finite-volume transport on various cubed-sphere grids. *J. Comp. Phys.*, 227:55–78, 2007.

[7] A. Staniforth and J. Thuburn. Horizontal grids for global weather and climate prediction models: a review. *Q. J. R. Meteorol. Soc.*, 138, 2012.

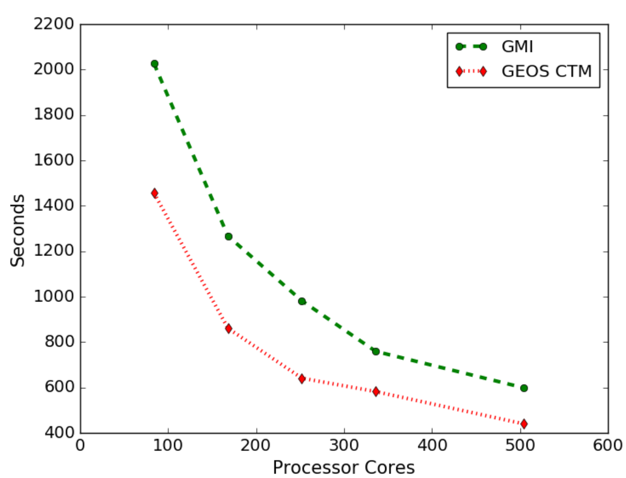

Figure 1: Overall wall clock times of GMI and GEOS CTM when the number of processor varies.

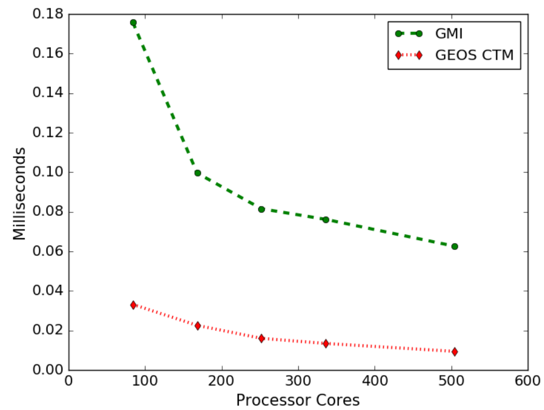

Figure 2: Wall clock times (per tracer per grid point) of Advection for GMI and GEOS CTM when the number of processors varies.

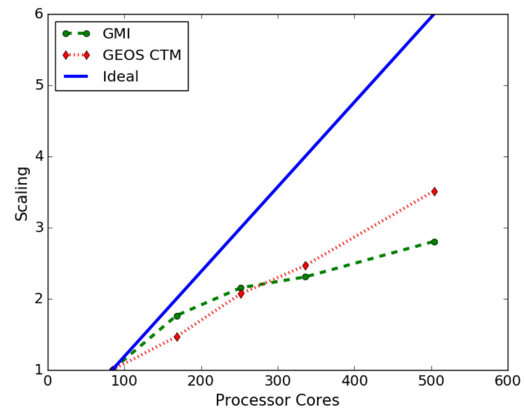

Figure 3: Scaling performance of Advection for GMI and GEOS CTM.

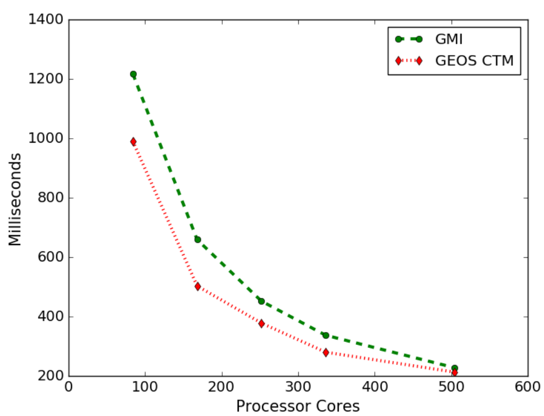

Figure 4: Wall clock times (per tracer per grid point) of Chemistry for GMI and GEOS CTM when the number of processors varies.

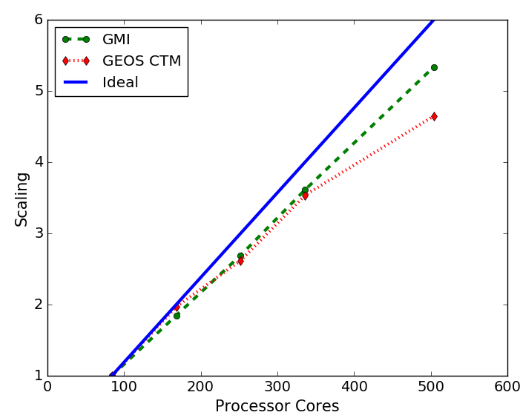

Figure 5: Scaling performance of Chemistry for GMI and GEOS CTM.

	Number of Cores				
	84	168	252	336	504
Original ExtData	134	131	133	181	160
Modified ExtData	89	84	94	91	99

Table 1: Wall clock times (in seconds) obtained by running a GEOS CTM using the original and modified versions of ExtData with various processor cores.