

GEOS-Chem “Classic” 12.0.0 timing results

Bob Yantosca
GEOS-Chem Support Team
16 Aug 2018

Timing test configurations

We set up 7-day timing runs with these GEOS-Chem “Classic” 12.0.0 configurations:

1. GNU Fortran v7.1 + fast-math option + bpch diagnostics
2. GNU Fortran v7.1 + fast-math option + netCDF diagnostics
3. GNU Fortran v7.1 + default build + bpch diagnostics
4. GNU Fortran v7.1 + default build + netCDF diagnostics
5. Intel Fortran 17.0.4 + default build + bpch diagnostics
6. Intel Fortran 17.0.4 + default build + netCDF diagnostics

For each configuration 1-6, we submitted jobs using 4, 8, 12, 16, 24, and 32 cores.

NOTE: “fast-math” was activated by hardcoding “-ffast-math” to the makefile.

=====

GEOS-CHEM TIMERS

Timer name	DD-hh:mm:ss.SSS	Total Seconds
GEOS-Chem	: 00-01:35:17.875	5717.875
Initialization	: 00-00:00:06.500	6.500
Timesteps	: 00-01:35:11.000	5711.000
HEMCO	: 00-00:46:57.062	2817.062
All chemistry	: 00-00:27:26.750	1646.750
=> Gas-phase chem	: 00-00:18:12.500	1092.500
=> FAST-JX photolysis	: 00-00:02:22.125	142.125
=> All aerosol chem	: 00-00:06:17.500	377.500
=> Strat chem	: 00-00:00:13.562	13.562
=> Unit conversions	: 00-00:00:11.062	11.062
Transport	: 00-00:05:57.437	357.438
Convection	: 00-00:05:08.125	308.125
Boundary layer mixing	: 00-00:05:28.562	328.562
Dry deposition	: 00-00:00:13.250	13.250
Wet deposition	: 00-00:02:03.500	123.500
RRTMG	: >>>> THE TIMER DID NOT RUN <<<<	
All diagnostics	: 00-00:00:27.375	27.375
=> HEMCO diagnostics	: 00-00:00:00.062	0.062
=> Binary punch diagnostics	: 00-00:00:27.312	27.312
Reading met fields	: 00-00:00:22.562	22.562
Reading restart file	: 00-00:00:02.437	2.438
Writing restart file	: 00-00:00:01.625	1.625
Input	: 00-00:16:25.062	985.062
Output	: 00-00:00:28.625	28.625
Finalization	: 00-00:00:00.312	0.312

Time spent per operation


We used the output of the GEOS-Chem timers to determine the approximate amount of time that was spent in each operation.

The GEOS-Chem timers are turned on by compiling with option TIMERS=1. Results are printed at the end of the simulation log file.

Profiling with e.g. the TAU performance analyzer will always be more accurate. But the GEOS-Chem timers are sufficient to determine if there are operations that do not scale well.

Scalability across cores

Scalability: 7-day timing runs with GEOS-Chem 12.0.0




Takeaways:

netCDF diags incur more overhead than BPCH diags (I/O via library)

gfortran + fast-math + netCDF diags is comparable to ifort + netCDF up to about 24 cores

ALSO NOTE: A few runs were likely affected by disk issues on Odyssey:
Gfortran + netCDF @ 16 cores
Gfortran + bpch @ 12 cores
Gfortran + bpch + fast math @ 16 cores

Scalability: 7-day timing runs with GEOS-Chem 12.0.0 (truncated Y-axis)




Here is a view of the same plot on the preceding slide.

The Y-axis range has been truncated to better show the differences between runs.

Time spent in each operation:
Reference run: v11-02c

Time spent in each GEOS-Chem operation

1-month GEOS-Chem "Classic" simulations @ 4 x 5



In January 2018, we performed a similar set of timing tests (1-month runs) with GEOS-Chem "Classic" v11-02c.

Gas-phase chemistry (KPP+ FAST-JX) sped up by about a factor of 3.5 in going from 6 to 30 cores.

But other operations scaled less well:

Transport sped up by ~ 2X

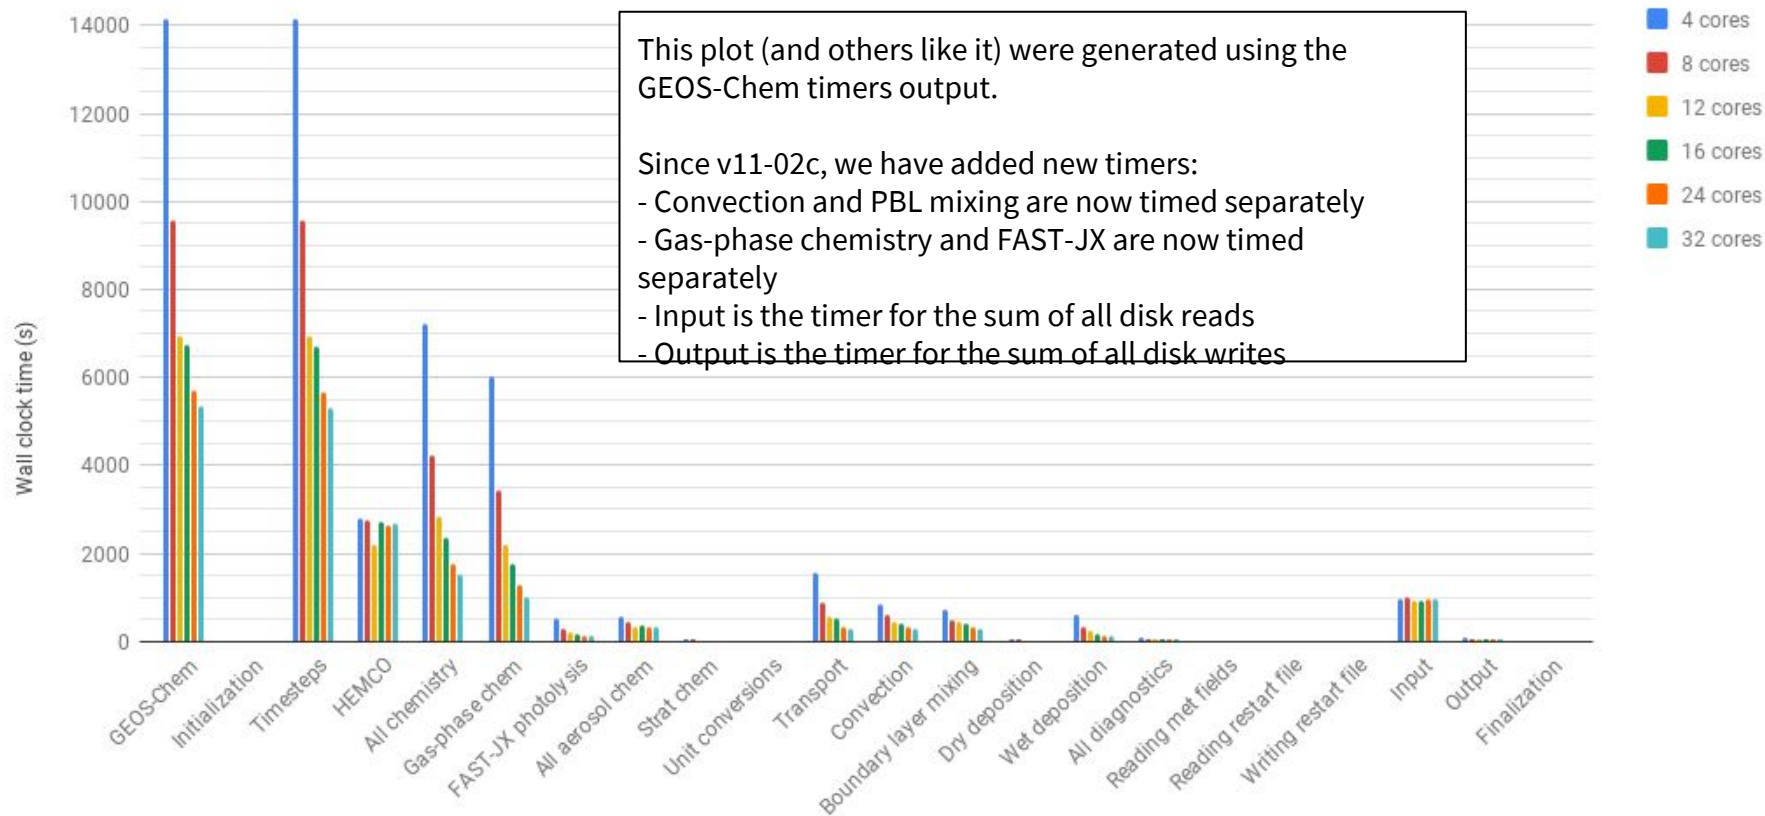
Convection sped up by ~2.5X

Wet deposition did not scale at all


HEMCO and Diagnostics are disk I/O intensive, and I/O is not parallelized, so they did not scale.

Time spent in each operation:
GNU Fortran v7.1 + fast math + bpch diagnostics


Time spent in each operation: gfortran + bpch diagnostics + fast math



Time spent in each operation: gfortran + bpch diagnostics + fast math (truncated Y-axis)




Percent speedup when switching from 4 to 16 cores and from 4 to 32 cores (gfortran + bpch diags + fast math)




We will show the same plots for the other timing test configurations...

Time spent in each operation:
GNU Fortran v7.1 + fast math + netCDF diagnostics


Time spent in each operation: gfortran + netCDF diagnostics + fast math



Time spent in each operation: gfortran + netCDF diagnostics + fast math (truncated Y-axis)




Percent speedup when switching from 4 to 16 cores and from 4 to 32 cores (gfortran + netCDF diags + fast math)




Time spent in each operation:
GNU Fortran v7.1 + default build + bpch diagnostics


Time spent in each operation: gfortran + bpch diagnostics



Time spent in each operation: gfortran + bpch diagnostics (truncated Y-axis)




Percent speedup when switching from 4 to 16 cores and from 4 to 32 cores (gfortran + bpch diagnostics)




Time spent in each operation:
GNU Fortran v7.1 + default build + netCDF diagnostics


Time spent in each operation: gfortran + netCDF diagnostics



Time spent in each operation: gfortran + netCDF diagnostics (truncated Y-axis)




Percent speedup when switching from 4 to 16 cores and from 4 to 32 cores (gfortran + netCDF diagnostics)




Time spent in each operation:
Intel Fortran 17.0.4 + default build + bpch diagnostics


Time spent in each operation: ifort + bpch diagnostics



Time spent in each operation: ifort + bpch diagnostics (truncated Y-axis)




Percent speedup when switching from 4 to 16 cores and from 4 to 32 cores (ifort + bpch diagnostics)




Time spent in each operation:
Intel Fortran 17.0.4 + default build + netCDF diagnostics

Time spent in each operation: ifort + netCDF diagnostics



Time spent in each operation: ifort + netCDF diagnostics (truncated Y-axis)



Percent speedup when switching from 4 to 16 cores and from 4 to 32 cores (ifort + netCDF diagnostics)

