CMAQ 5.3 PARALLEL PERFORMANCE FOR A 101-DAY SIMULATION**

George Delic*
HiPERiSM Consulting, LLC, P.O. Box 569, Chapel Hill, NC 27514, USA

1. INTRODUCTION

This presentation reports on implementation of a thread parallel sparse matrix solver, FSparse [1], in the Chemistry Transport Model (CTM) of CMAQ and also the addition of thread parallelism in the horizontal advection (HADV) science process. In this report performance results of the original U.S. EPA JSparse [2] and FSparse versions are presented. This report includes results with CMAQ for each of Euler-backward (EBI), Rosenbrock (ROS3), and SMV Gear (GEAR) algorithms in the CTM.

2. TEST BED ENVIRONMENT

2.1 Hardware

The hardware systems chosen were the platforms at HiPERiSM Consulting, LLC, shown in Table 2.1. Nodes 20 and 21 host two Intel E5v3 CPUs with 16 cores and each node has four Intel Phi co-processor many integrated core (MIC) cards [3] with 60 and 59 cores, respectively (but not applied in this report). These are the base nodes of a heterogeneous cluster that includes a HP blade server [4] hosting nodes 27 to 34 with dual 4-core Intel E5640 CPUs and nodes 35 to 40 with dual 6-core Intel X5670 CPUs. The total core count of this heterogeneous cluster is 192. For the standard U.S. EPA version 4x6=24 MPI processes are launched across a combination of these nodes. This cluster allows for comparison of runtimes and numerical precision for species in the FSparse hybrid (MPI + OpenMP) parallel versions of CMAQ with the original EPA version.

2.2 Compilers

Results reported here implemented the Intel Parallel Studio® suite (release 17.6, [3]), with compiler options for a heterogeneous cluster that enable OpenMP threads and instruction level vector processing.

2.3 Episode studied

This report used the benchmark test data available in the CMAQ 5.3 download for part of an annual episode. This episode was for the date range 2015-12-22 to 2016-03-31 (101 days), using the cb6r3_ae7_aq mechanism with 147 active species and 329 reactions. For day/night chemistry this results in 1400/1348 non-zero entries in the Jacobian matrix. The episode was run on a 299 X 459 CONUS (12US1) domain at 12 Km grid spacing and 35 vertical layers for a total of 4,803,435 grid cells. In this report, due to runtime constraints, only 24 MPI processes (NP) were used in all three CTM versions with 8 threads (omp8) in the OpenMP case.

Table 2.1. CPU platforms at HiPERiSM Consulting, LLC

Platform	Node20-21 (each node)	Node27-34 (each node)	Node35-40 (each node)
Operating system	OpenSuSE 13.2	OpenSuSE 42.3	OpenSuSE 42.3
Processor	Intel™ x86-64 (E5-2698v3)	Intel™ x86-64 (E5640)	Intel™ x86-64 (X5670)
Coprocessor	4 x Intel Phi 7120/5120	NA	NA
Peak Gflops / CPU (SP/DP)	~589 (SP)	~ 43 (DP)	~ 70 (DP)
Power consumption	135 Watts	80 Watts	95 Watts
Cores per processor	16	4	6
Power per core	8.44 Watts	20 Watts	29 Watts
Processor count	2	2	2
Total core count	32	8	12
Clock	2.3 GHz	2.67 GHz	2.93 GHz
Bandwidth	68 GB/sec	25.6 GB/sec	32 GB/sec
Bus speed	2133 MHz	2933 MHz	3200 MHz
L1 cache	16x32 KB	4x32 KB	6x32 KB
L2 cache	16x256 KB	4x256 KB	6x256 KB
L3 cache	40 MB	12 MB	12 MB

In the following the performance metric introduced to assess parallel performance in the

* *Corresponding author:* George Delic, george@hiperism.com.

** Additional results are added as the simulation progresses.

MPI and OpenMP modified code is *Speedup* defined as the gain in runtime over the standard U.S. EPA version.

2.4 Interconnect fabric

Results reported here used the heterogeneous cluster consisting of node 20 & 21 and the HP blade. The blade chassis has an internal switch connecting node27 to node40 and uplinks all blades to the 10GigE switch to join all nodes together.

For MPI traffic in cluster mode bandwidth is via an Infiniband (IB) fabric with a (theoretical) limit of 40G bits/sec.

3. RESULTS FOR TWO CMAQ MODELS

3.1 Performance profile of CMAQ

For a 101 day simulation with the EBI solver a profile of time consumed by science process is shown Fig. 3.1. The dominant science processes in CMAQ are the CTM (CHEM), horizontal advection (HADV), vertical diffusion (VDIFF), and aerosol (AERO). The EPA version is compared with the FSparse version for 8 threads as identified in the legend. The fraction of total time (percent) for each science process is shown in Fig. 3.2. In the OpenMP case, as time in CHEM and HADV decreases, the fraction of time in the other science processes correspondingly increases. With 24 MPI processes (as used here), it is evident that the horizontal advection (HADV) science process dominates the fraction of wall clock time in both EPA and FSparse versions of CMAQ.

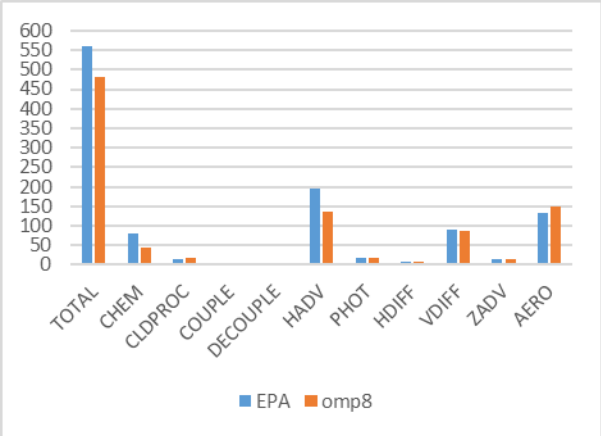

Fig 3.1 Wall clock time (hours) by science process for the U.S. EPA (EPA) and FSparse versions of the EBI solver of CMAQ for 24 MPI processes and OpenMP thread count of 8 (omp8), for a total of 101 simulation days.

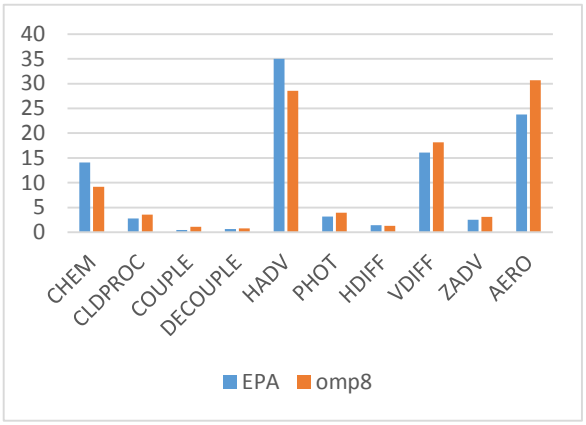

Fig 3.2: Fraction of wall clock time (percent) by science process for the U.S. EPA (EPA) and FSparse versions of the EBI solver of CMAQ for 24 MPI processes and OpenMP thread count of 8 (omp8), for a total of 101 simulation days.

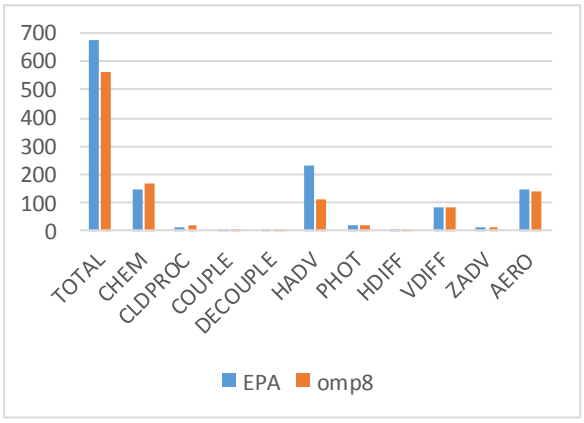

Fig 3.3 Wall clock time (hours) by science process for the U.S. EPA (EPA) and FSparse versions of the ROS3 solver of CMAQ for 24 MPI processes and OpenMP.

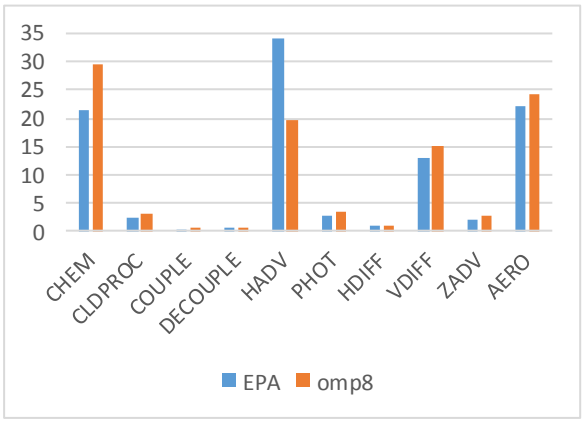

Fig 3.4: Fraction of wall clock time (percent) by science process for the U.S. EPA (EPA) and FSparse versions of the ROS3 solver of CMAQ for 24 MPI processes and OpenMP thread count of 8 (omp8), for a total of 101 simulation days.

Figs. 3.3 and 3.4 show corresponding results for the ROS3 solver case.

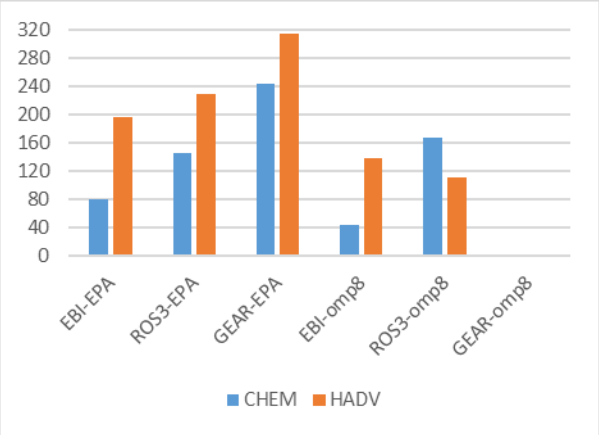

Fig 3.5: Wall clock time (hours) for science processes CHEM and HADV for the U.S. EPA (EPA) and FSparse versions of the EBI and ROS3 solvers of CMAQ for 24 MPI processes and OpenMP thread counts of 8 (omp8) for a total of 101 simulation days.

Both CHEM and HADV in the OpenMP threaded version show reduction in wall clock time, except for CHEM in ROS3. This is the focus of Fig. 3.5 where the EBI and ROS3 solver results have completed for EPA and Fsparse versions. The anomalous result of CHEM in ROS3 could be due to oversubscription of cores by two of the MPI processes and will be investigated in the future.

3.2 Wall clock time performance

Table 3.1. Total wall clock time (hours) and speedup of the FSparse OpenMP 8 thread version over EPA with 24 MPI processes for a 101 day simulation.

CTM version	Wall clock time for 101 day simulation and average speedup		
	EPA time (hours)	OpenMP time (hours)	Average Speeup
GEAR	886	---	1.25
ROS3	678	571	1.19
EBI	567	487	1.16

Table 3.2. Daily wall clock time statistics (in hours) for U.S. EPA and FSparse OpenMP 8 thread version of CMAQ with 24 MPI processes for a 101 day simulation.

	EPA version			Fsparse version		
	EBI	ROS3	GEAR	EBI	ROS3	GEAR
MIN	4.6	5.6	6.7	3.8	4.8	--
MAX	6.9	8.3	11.9	6.0	7.1	--
MEAN	5.6	6.7	8.8	4.8	5.7	--

Table 3.1 shows wall clock time, for 24 MPI processes in a 101 day simulation. Fig. 3.6 shows a time series of wall clock time in hours for 101 days in this simulation with the EPA version of CMAQ for all three solvers in the CTM. Some statistics for this series in summarized in Table

3.2. Results for the FSparse versions of GEAR in the CTM are pending.

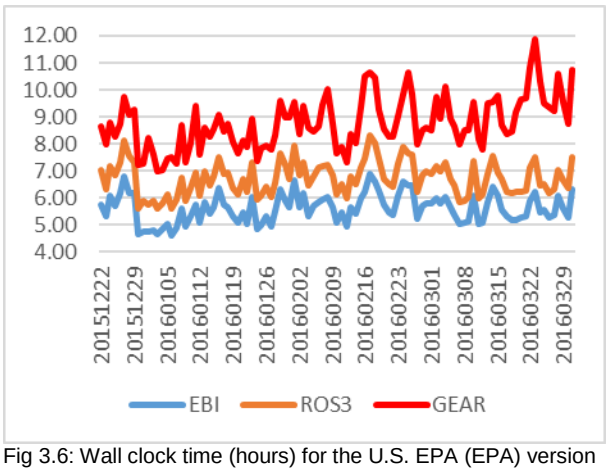

Fig 3.6: Wall clock time (hours) for the U.S. EPA (EPA) version of the EBI, ROS3, and GEAR solvers of CMAQ for 24 MPI processes for 101 simulation days.

3.3 FSparse speedup versus EPA

While there is variability over the 101 simulation days, in the results of Table 3.1 (so far) the FSparse version has a speedup over the U.S. EPA version of 1.16 (EBI), 1.20 (ROS3), and 1.25 (GEAR). Looking into detail, Table 3.3 shows numerical values for the two OpenMP threaded science processes (CHEM and HADV in Fig. 3.5).

Table 3.3. Total wall clock time (hours) in CHEM and HADV with speedup of the FSparse OpenMP 8 thread version over EPA with 24 MPI processes for a 101 day simulation.

CTM version	Science process	Wall clock time for 101 day simulation and average speedup		
		EPA time (hours)	OpenMP time (hours)	Average Speeup
EBI	CHEM	79.1	44.2	1.79
	HADV	196.6	137.4	1.43
ROS3	CHEM	145.5	167.3	0.87
	HADV	229.3	111.1	2.06
GEAR	CHEM	243.9	---	---
	HADV	314.1	---	---

4. NUMERICAL RESULTS

The absolute error in predictions of species concentrations were compared for EBI, ROS3, and GEAR versions of the CTM solver. Eleven species were compared in a metric that used a histogram displaying the fraction (percent) of the total population in the sample of each species concentration absolute error in the first layer of the CONUS grid. The sample size is 137241 values. The boundary values of the histogram bins are normalized (i.e. divided by) the maximum value of the concentration for each of the respective species listed

here: O3, CO, NH3, HNO3, NO3, NO, NO2, H2O2, SO2, HO2, and N2O5. The lower bound of each histogram bin is shown in the following figures. The concentration values were for the end of a 24 hour scenario on 12/31/2015.

4.1 EBI version of CMAQ

Species concentrations predicted by the EBI CTM solver in JSparse and FSparse versions were compared for the eleven species listed above. The absolute error of the difference was sorted into a frequency histogram shown in Fig. 4.1.

The comparison shows that the agreement is two significant figures relative to the maximum value for each species, with the exception of O3. Possible explanations for this discrepancy are (a) the accuracy of EBI algorithm itself, or (b) changes in treating conversion from single to double precision in the CTM of CMAQ.

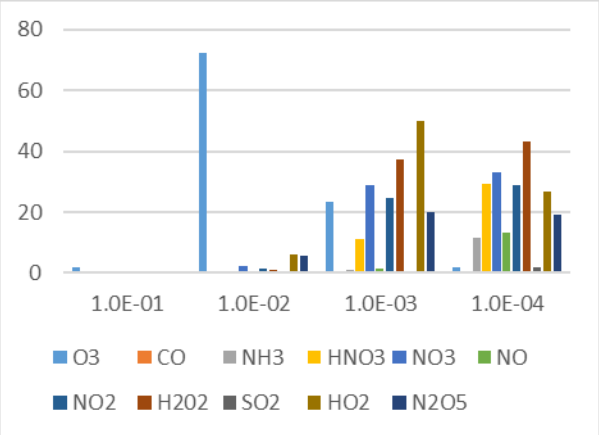

Fig.4.1. Frequency data for the absolute error in comparing concentration predictions of the EBI solver in the EPA JSparse and FSparse versions of CMAQ for the species shown in the legend. The histogram bins are in decreasing decades from left to right. The vertical scale shows the percent of the total sample falling into the corresponding bin.

4.2 ROS3 version of CMAQ

Species concentrations predicted by the ROS3 CTM solver in JSparse and FSparse versions were compared for the eleven species listed above. The absolute error of the difference was sorted into a frequency histogram shown in Fig. 4.2. These results are superior to EBI and could be due to the difference in the algorithms and the respective error tolerances.

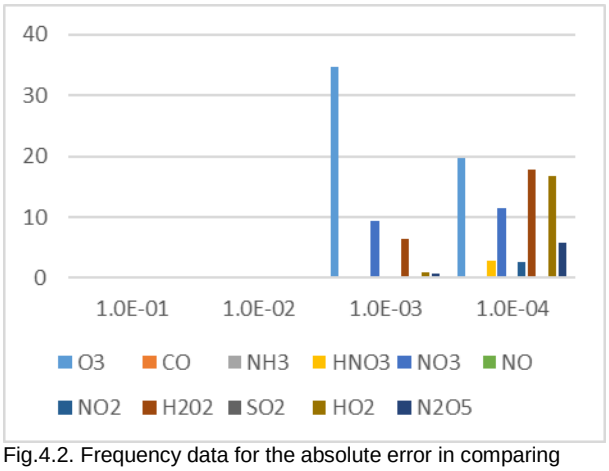

Fig.4.2. Frequency data for the absolute error in comparing concentration predictions of the ROS3 solver in the EPA JSparse and FSparse versions of CMAQ for the species shown in the legend. The histogram bins are in decreasing decades from left to right. The vertical scale shows the percent of the total sample falling into the corresponding bin.

4.3 EPA version of CMAQ

Species concentrations were compared between the three solvers: EBI, ROS3 and, GEAR in the CTM of CMAQ with the standard U.S. EPA versions. The eleven species listed above were compared as the difference between predictions of EBI, ROS3, and GEAR solvers, respectively. The comparison was for all 137241 values of concentrations in the first layer of the grid.

Fig. 4.3 compares differences of predictions between EBI and ROS3 in the CTM of the U.S. EPA CMAQ version. While Fig. 4.4 shows the same comparison for ROS3 versus GEAR. The major difference is the prediction for O3 concentrations.

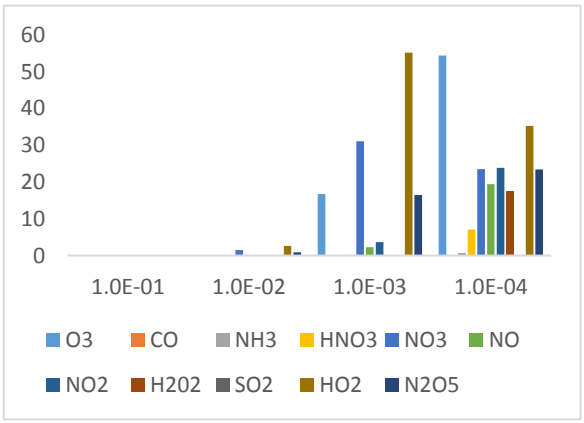

Fig.4.3. Frequency data for the absolute error in comparing concentration predictions of EBI and ROS3 solvers in the EPA JSparse version of CMAQ for the species shown in the legend. The histogram bins are in decreasing decades from left to right.

The vertical scale shows the percent of the total sample falling into the corresponding bin.

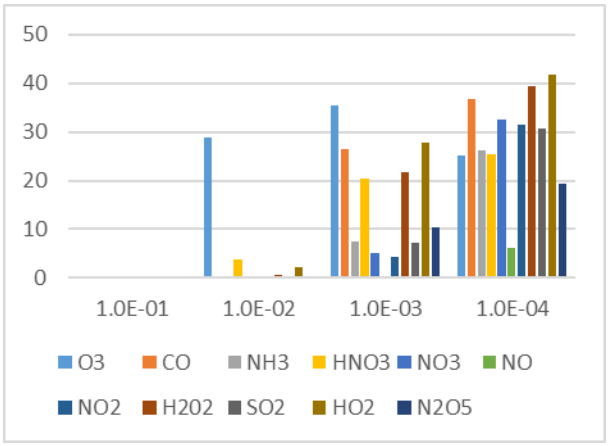

Fig. 4.4. Frequency data for the absolute error in comparing concentration predictions of ROS3 and GEAR solvers in the EPA JSparse version of CMAQ for the species shown in the legend. The histogram bins are in decreasing decades from left to right. The vertical scale shows the percent of the total sample falling into the corresponding bin.

4.4 Precision in CMAQ

In section 4.3 species concentrations were compared for all three CTM solvers in CMAQ with the U.S. EPA version. Generally, the precision in species concentration values is two significant figures in all three solvers. The exception is O3 where nearly 30 percent of the sample values may have less than this (see Fig.4.4). Keeping in mind that this may be occurring where the value of the concentration is not large. To investigate this further scatter plots of the absolute error are shown in Figs. 4.5 to 4.7. These respectively compare the absolute error in predictions of EBI versus ROS3, EBI versus GEAR, and ROS3 versus GEAR. Fig. 4.5 shows good agreement to two significant figures, but this deteriorates somewhat in Figs.4.6 and 4.7. In other words better than this precision of agreement between all three algorithms may not be possible. There remains the question: which is the superior CTM algorithm? It is the GEAR algorithm that is viewed as superior to the other two.

The same comparison metric shows that the difference between the EPA version and the OpenMP version is within this tolerance. As a footnote, while the ROS3 algorithm uses an absolute error tolerance of 10^{-7}, GEAR requires 10^{-9}, but 10^{-8} was used because in CMAQ 5.3 it does not always converge. This contradicts our experience with previous versions of U.S. EPA's CMAQ releases (see Section 6 in [5]). This needs further investigation.

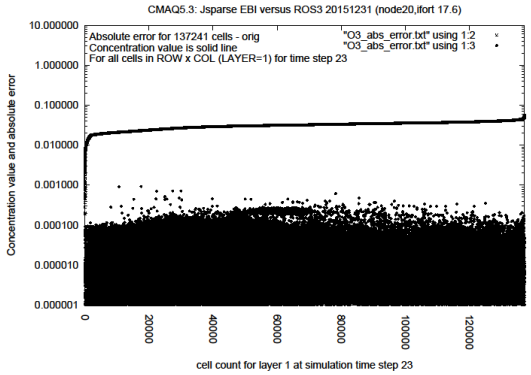

Fig 4.5: For the JSparse EBI versus ROS3 solver of CMAQ this shows the O3 species concentration absolute error (scattered points) and concentration value (solid line) for 137241 values in layer 1 of the CONUS domain. The ranking is in increasing concentration value from left to right

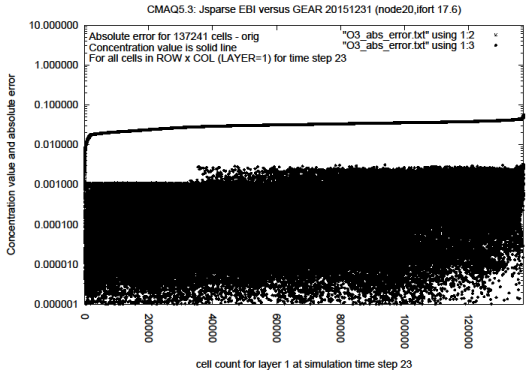

Fig 4.6: For the JSparse EBI versus GEAR solver of CMAQ this shows the O3 species concentration absolute error (scattered points) and concentration value (solid line) for 137241 values in layer 1 of the CONUS domain. The ranking is in increasing concentration value from left to right.

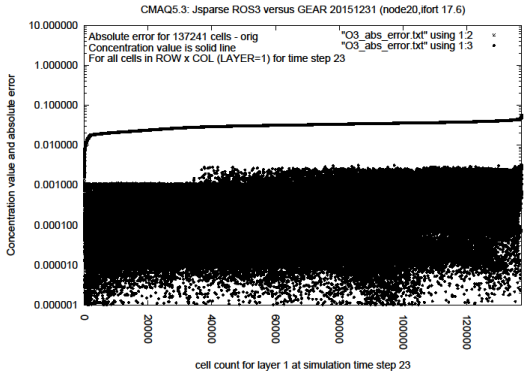

Fig 4.7: For the JSparse ROS3 versus GEAR solver of CMAQ this shows the O3 species concentration absolute error (scattered points) and concentration value (solid line) for

137241 values in layer 1 of the CONUS domain. The ranking is in increasing concentration value from left to right.

5. SUMMARY OF KEY POINTS

5.1 Speedup

- FSparse OpenMP speedup over the U.S. EPA version of CMAQ ranges from 1.16 (EBI) to 1.25 (GEAR).

5.2 Numerical precision

- The three algorithms, EBI, ROS3, and GEAR, can produce different precision in numerical values for species concentrations, but use different error tolerances.
- The ROS3 and GEAR algorithms are superior to EBI and should be implemented in its place.
- For reasons of wall clock time expense ROS3 may be preferred to GEAR in production use.

6. LESSONS LEARNED

6.1 Benefits of the FSparse method

Comparing runtime performance for CMAQ 5.3 in the OpenMP parallel version with the U.S. EPA release showed benefits such as:

- Thread speedup with 8 threads
- The CTM with the Rosenbrock algorithm in the FSparse version delivers wall clock times similar to the Euler-backward algorithm in the standard U.S. EPA release of CMAQ.
- Numerical values of predicted species concentration that are within the error tolerance inherent in the algorithms

6.2 Next steps

A continuation of this work would include:

- Completion of the remaining days of the 2016 CONUS scenario and update of results.
- Investigation of the anomalous ROS3 speedup in CHEM.
- Inspection of numerical accuracy in all three CTM algorithms, especially GEAR.

7. CONCLUSIONS

This report has described an analysis of CMAQ 5.3 behavior in the standard U.S. EPA release and a new thread parallel version of CMAQ suitable for the Euler-backward, Rosenbrock and, SMV Gear solvers.

The new FSparse version of CMAQ offers layers of parallelism not available in the standard U.S. EPA release and is portable across multi-core hardware and compilers that support thread parallelism.

The SMV Gear algorithm may be used for high precision predictions, despite the expense in wall clock time. However, as a compromise in wall clock time demand, a general recommendation is that CMAQ should be used with the FSparse version of the Rosenbrock solver in the CTM in production scenarios.

Updates to this report will be posted at [6] as more results complete.

REFERENCES

[1] Delic, G., Modern Environmental Science and Engineering, Vol. 5, Nr.9, 2019, pp. 775-791. Full text available at: https://www.researchgate.net/publication/338581080_A_Thread_Parallel_Sparse_Matrix_Chemistry_Algorithm_for_the_Community_Multiscale_Air_Quality_Model

[2] Jacobson, M. and Turco, R.P., (1994), Atmos. Environ. 28, 273-284.

[3] INTEL: Intel Corporation, http://www.intel.com

[4] https://en.wikipedia.org/wiki/HP_BladeSystem

[5] G. Delic, 2018, 17th Annual CMAS Conference, https://www.researchgate.net/publication/328074959_CMAQ_521_PARALLEL_PERFORMANCE_WITH_MPI_AND_OPENMP

[6] G. Delic, 2020, https://www.researchgate.net/project/Sparse-Matrix-Solvers-in-Air-Quality-Models-on-Parallel-Processors

ACKNOWLEDGEMENTS

The author gratefully acknowledges help from Kristen Foley (U.S. EPA), Ed Anderson (GDIT), and Elizabeth Adams (UNC) in providing model data and resolving implementation issues.