CMAQ 5.3 PARALLEL PERFORMANCE FOR CY 2016**

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) and CTM science processes. In this report performance results of the original (Legacy) U.S. EPA JSparse [2] and FSparse versions are presented. This report includes results with CMAQ for Euler-backward (EBI) and Rosenbrock (ROS3) 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 upgraded from the previous year's report. Currently nodes 20 and 21 host two Intel CPUs (E5-2698v3) with 16 cores and each node has four Intel Phi co-processor many integrated core (MIC) cards [3] with 60 and 59 cores, respectively. The upgrade to the cluster added a dual CPU (2683v4) 16-core node 22 and dual CPU (E5-2699v3) 18 core node53 and node54. The HP blade server [4] (hosting nodes 27 to 40) was upgraded with dual 6-core Intel E5670 CPUs. The total core count used on this heterogeneous cluster is 328 (up from 192 in last year's report). The upgraded cluster now hosts CMAQ 5.3 with 4x6=24 MPI processes launched across a combination of these nodes. This allows comparison of runtimes and 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 an annual episode (2016) of 376 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 for 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 24 MPI processes were used in both CMAQ versions with 12 threads (omp12) in the OpenMP case using 288 cores.

Table 2.1. CPU platforms at HiPERiSM Consulting, LLC

Platform	Node20-22,53-54 (each node)	Node27-40 (each node)
Operating system	OpenSuSE 13.2	OpenSuSE 42.3
Processor	Intel™ x86-64 (E5-2698v3)	Intel™ x86-64 (X5670)
Coprocessor	4 x Intel Phi 7120/5120	NA
Peak Gflops / CPU (SP/DP)	~589 (SP)	~ 70 (DP)
Power consumption	135 Watts	95 Watts
Cores per processor	16	6
Power per core	8.44 Watts	29 Watts
Processor count	2	2
Total core count	32	12
Clock	2.3 GHz	2.93 GHz
Bandwidth	68 GB/sec	32 GB/sec
Bus speed	2133 MHz	3200 MHz
L1 cache	16x32 KB	6x32 KB
L2 cache	16x256 KB	6x256 KB
L3 cache	40 MB	12 MB

In the following the performance metric introduced to assess parallel performance in the MPI and OpenMP modified code is *Speedup*

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

** This is an update of the previous year's report.

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 the nodes listed in Section 2.1. The HP blade chassis has an internal switch connecting node27 to node40 and uplinks all blades to the 10GigE switch to join all other 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 376 day simulation with the EBI solver a profile of percent of time consumed by science process is shown Fig. 3.1. Dominant science processes in CMAQ are the CTM (CHEM), horizontal advection (HADV), and aerosol (AERO). The EPA version is compared with the FSparse version for 12 OpenMP threads (as identified in the legend) in CHEM and HADV. The total time (hours) for these two is shown in Fig. 3.2 where the average OpenMP speedup is 2.2 and 1.4 respectively. With 24 MPI processes (as used here), it is evident that the horizontal advection (HADV) and AERO science processes dominate the fraction of wall clock time in both EPA and FSparse versions of CMAQ. However, AERO has too much scalar code to be thread parallelized but VDIFF could be investigated further.

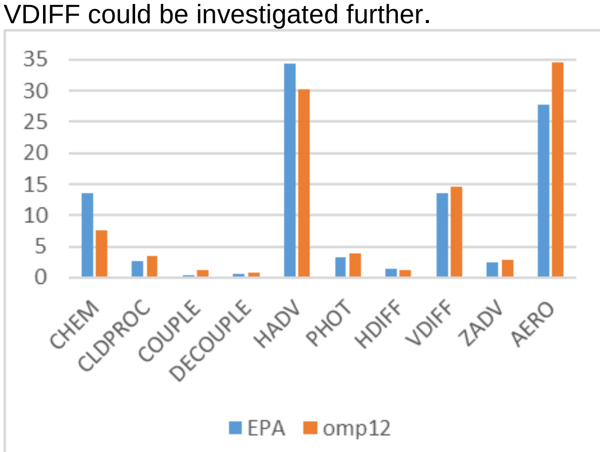

Fig 3.1 Fraction of total wall clock time (percent) by science process for the U.S. EPA (EPA) and FSparse versions of the EBI solver of CMAQ for 24 MPI processes and an OpenMP thread count of 12 (omp12), for a total of 376 simulation days.

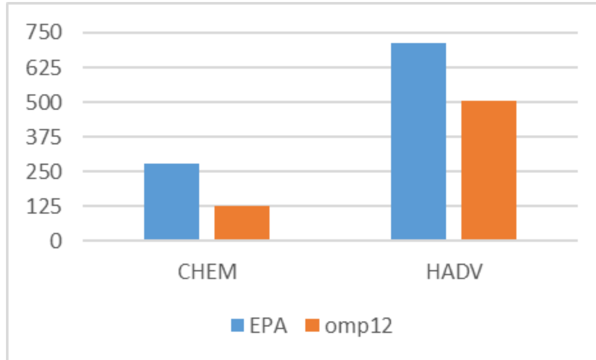

Fig 3.2: Wall clock time (hours) for the CHEM and HADV science processes for the U.S. EPA (EPA) and FSparse versions of the EBI solver of CMAQ for 24 MPI processes and an OpenMP thread count of 12 (omp12), for a total of 376 simulation days.

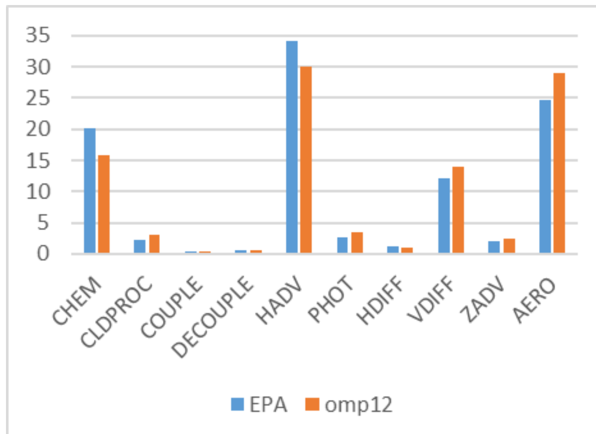

Fig 3.3. Fraction of total wall clock time (percent) by science process for the U.S. EPA (EPA) and FSparse versions of the ROS3 solver of CMAQ for 24 MPI processes and an OpenMP thread count of 12 (omp12), for a total of 376 simulation days.

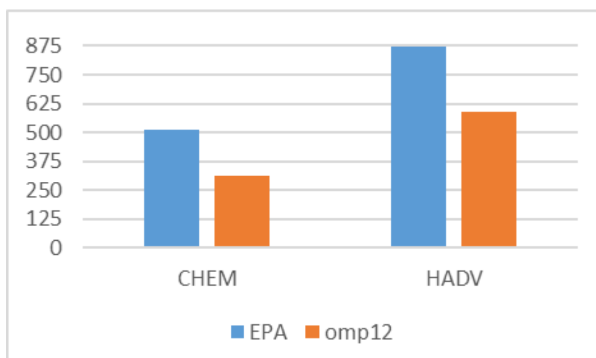

Fig 3.4. Wall clock time (hours) for the CHEM and HADV science processes for the U.S. EPA (EPA) and FSparse versions of the ROS3 solver of CMAQ for 24 MPI processes and an OpenMP thread count of 12 (omp12), for a total of 376 simulation days.

Figs. 3.3 and 3.4 show corresponding results for the ROS3 solver case and the latter shows an average speed up of 1.65 and 1.48, respectively.

.

3.2 Wall clock time performance

Both CHEM and HADV in the OpenMP threaded version show reductions in wall clock time. Table 3.1 shows wall clock time for 24 MPI processes in a 376 day simulation. The average speedup in both ROS3 and EBI solvers is shown in the last column. Results for the FSparse versions of GEAR in the CTM are pending completion of the full CY2016 simulation.

Table 3.1. Total wall clock time (hours) and speedup of the FSparse OpenMP 12 thread version over the legacy EPA version with 24 MPI processes for a 376 day simulation.

CTM version	Wall clock time for 376 day simulation and average speedup		
	EPA time (hours / days)	OpenMP time (hours / days)	Average Speeup
ROS3	2565 / 107	1984 / 83	1.29
EBI	2093 / 87	1684 / 70	1.24

4. FSparse speedup versus EPA

4.1 Average over 376 days

For the 376-day simulation there are 117238 calls to both CHEM and HADV science procedures and one way of displaying this amount of detail is with the Probability Density Function (PDF). This is constructed by selecting bins in the speedup values and counting the number of samples in each bin – in other words a histogram. The area under the corresponding curve in the PDF is then the sample size within that speedup bin.

The detailed behavior in each CTM solver is described in the following two sections

4.2 EBI speedup profile

The previous section discussed average speed up in the OpenMP version of CMAQ over the CY2016 simulation. However, in the 117238 calls to CHEM and HADV over 376 days of the simulation there is a distribution of the speed up value and these have been collected in histograms or Probability Distribution Functions (PDF). Fig. 4.1 shows the histogram for CHEM in 117238 calls over 376 days of simulation for the EBI solver in CMAQ. This shows speedup on the horizontal axis and fraction of the sample in the vertical axis. An integral under the curve shows the fraction of the whole sample in a specific range of speed up values. Thus 94% of calls have a speedup between 2.1 and 2.4.

Fig. 4.2 shows the corresponding histogram for HADV in 117238 calls over 376 days of simulation for the EBI solver in CMAQ. This distribution differs from that of CHEM and also shows speedup on the horizontal axis and fraction of the sample in the vertical axis. Here some 71% of all calls have a speedup between 1.0 and 1.6, while 25% of all calls have a speedup between 1.6 and 4.5 (and higher).

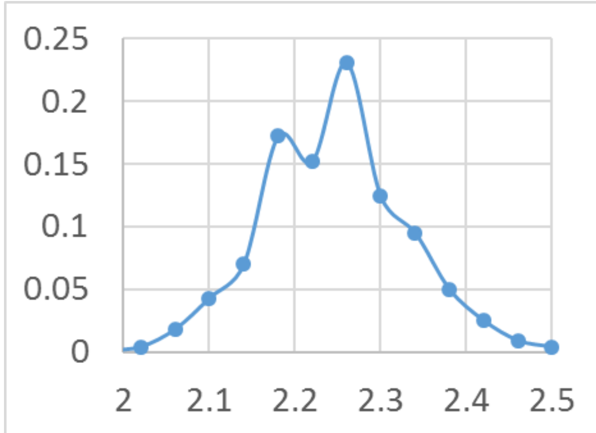

Fig 4.1: Speedup PDF distribution for 117238 calls to science process CHEM for the EBI solver of CMAQ with 24 MPI processes and an OpenMP thread count of 12 (omp12).

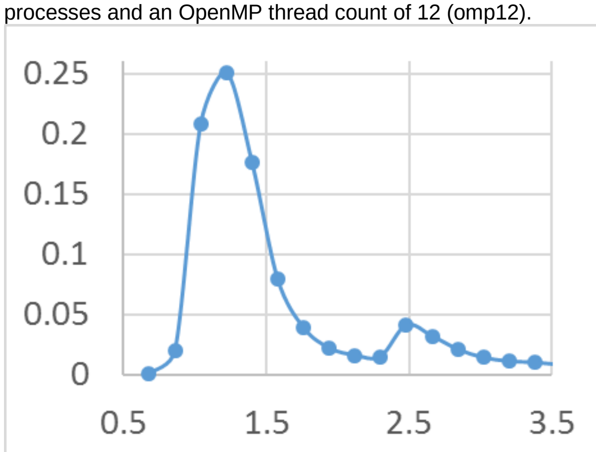

Fig 4.2: Speedup PDF distribution for 117238 calls to science process HADV for the EBI solver of CMAQ with 24 MPI processes and an OpenMP thread count of 12 (omp12).

4.3 ROS3 speedup profile

Fig. 4.3 shows the histogram for CHEM in 117238 calls over 376 days of simulation for the ROS3 solver in CMAQ. This shows speedup on the horizontal axis and fraction of the sample in the vertical axis. Thus 99% of calls have a speedup between 1.5 and 1.9.

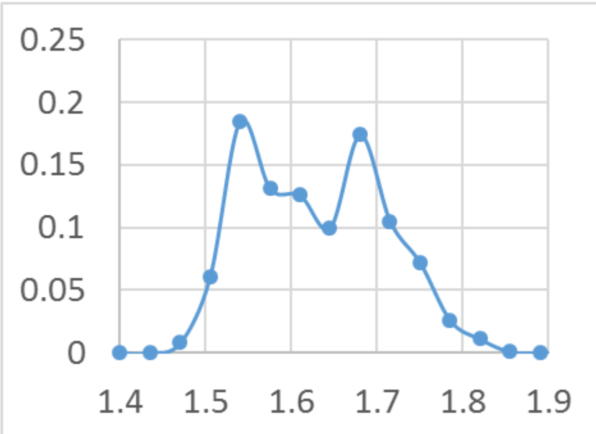

Fig 4.3: Speedup PDF distribution for 117238 calls to science process CHEM for the ROS3 solver of CMAQ with 24 MPI processes and an OpenMP thread count of 12 (omp12).

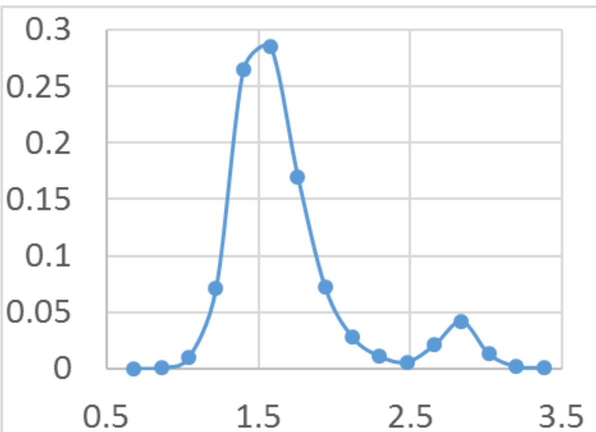

Fig 4.4: Speedup PDF distribution for 117238 calls to science process HADV for the ROS3 solver of CMAQ with 24 MPI processes and OpenMP thread counts of 12 (omp12).

Fig. 4.4 shows the corresponding histogram for HADV in 117239 calls over 376 days of simulation for the ROS3 solver in CMAQ. This shows speedup on the horizontal axis and fraction of the sample in the vertical axis. Thus 63% of calls have a speedup between 1.0 and 1.6, while 37% of all calls have a speedup between 1.6 and 3.4.

5. SUMMARY OF KEY POINTS

5.1 Average speedup

FSparse OpenMP average speedup for a 376-day simulation over the U.S. EPA version of CMAQ was 1.24 and 1.29 for the EBI and Rosenbrock solvers, respectively.

5.2 Speedup profile

The speedup profiles in the thread enabled science procedures ranged from 0.5 to 4.6 with the majority of the samples well above 1.5.

5.3 Next steps

A continuation of this work would include completion of whole year simulation of the 2016 CONUS scenario with the GEAR CTM solver in CMAQ and inspection of numerical accuracy in all three CTM algorithms.

6. 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 and Rosenbrock chemistry solvers in CMAQ 5.3.

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.

Updates to this report will be posted at future CMAS meetings.

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

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.