Technical Memorandum #5

REMSAD Platform Intercomparison Experiments

Prepared by Northeast States for Coordinated
Air Use Management (NESCAUM)

For the MANE-VU Regional Planning
Organization

March 21, 2003

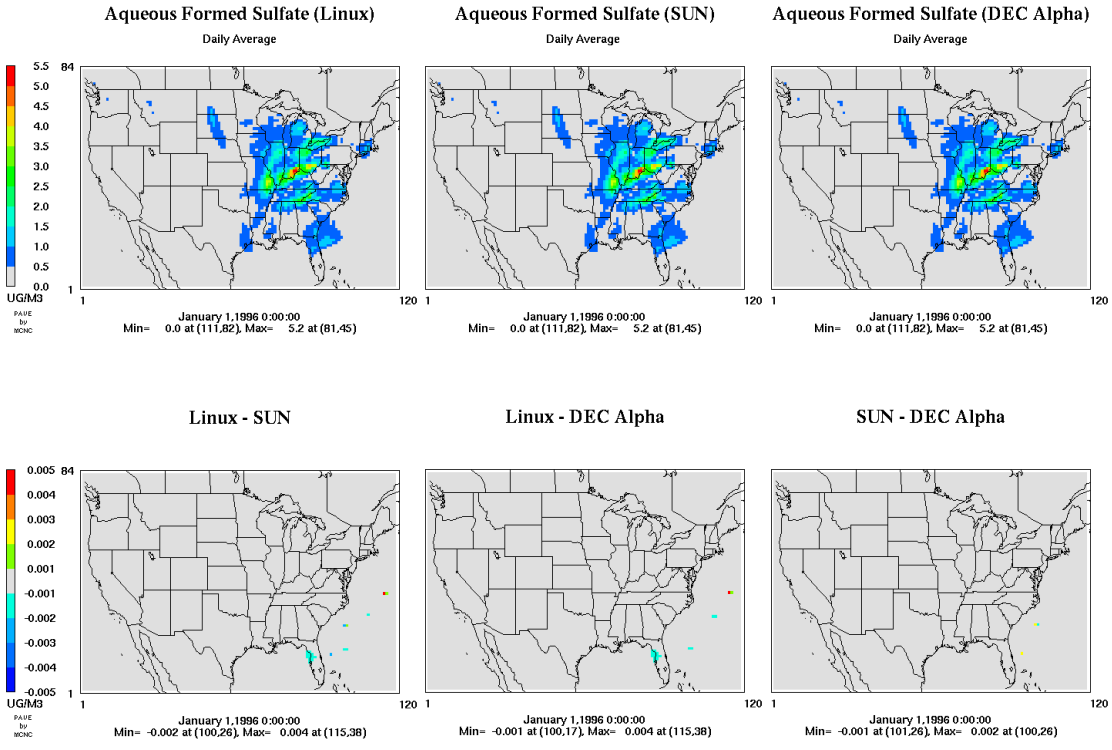

NESCAUM
Northeast States for Coordinated Air Use Management
101 Merrimac Street, 10th Floor
Boston, MA 02114

TEL: 617-367-8540
FAX: 617-742-9162
http://www.nescaum.org

Technical Memorandum #5

REMSAD Platform Intercomparison Experiments

March 21, 2003

Submitted to the United States Environmental Protection Agency, Region III in partial fulfillment of requirements under EPA Agreement X-983607-01-1 to the Ozone Transport Commission

Project Manager

Gary Kleiman

Editors

John Graham, Gary Kleiman

Principal Contributors (NESCAUM)

Rawlings Miller, NESCAUM
Kevin Civerolo, NY DEC
John Graham, NESCAUM

USEPA Project Officer

Russ S. Bowen (USEPA Region III)

Acknowledgements

We gratefully acknowledge the funding support provided by the United States Environmental Protection Agency (USEPA) under cooperative agreement with the Ozone Transport Commission.

We thank Sharon Douglas of SAI/ICF Consulting for providing us with REMSAD output developed on a DEC platform for comparison and Tom Myers of ICF Consulting for valuable assistance.
.

Table of Contents

Units, Species, Acronyms

Acronyms

ATDM – Aerosol and Toxics Deposition Model
IMPROVE – Interagency Monitoring of Protected Visual Environments
MANE-VU – Mid-Atlantic/Northeast Visibility Union
MARAMA – Mid Atlantic Regional Air Management Association
NESCAUM – Northeast States for Coordinated Air Use Management
NET – National Emission Trends Database
REMSAD – Regional Modeling System for Aerosols and Deposition
USEPA – United States Environmental Protection Agency

Chemical Species

ASO4 – sulfate formed through aqueous processes
GSO4 – sulfate formed through gaseous processes
NH4N – ammonium associated with nitrate particles
NH4S – ammonium associated with sulfate particles
PEC – primary elemental carbon
$PM_{2.5}$ – all particulate matter up to 2.5 µm in diameter
PMFINE – primary particulate matter up to 2.5 µm in diameter (excluding PEC and POA)
PNO3 – particulate nitrate
POA – primary organic aerosol
SOA – secondary organic aerosol

Units/Symbols

Mass
µg – micrograms (0.000001 x g; 10^{-6} g)
ng – nanograms (10^{-9} g)

Concentration
$\mu g/m^3$ – micrograms per cubic meter
ng/m^3 – nanograms per cubic meter

Executive Summary

As part of state efforts to comply with the requirements of U.S. Evironmental Protection Agency's (USEPA) 1999 Regional Haze Rule, the Mid-Atlantic/Northeast Visbility Union (MANE-VU) has been developing technical expertise for regional-scale air quality modeling. A variety of modeling systems and platforms are available for conducting assessments of pollutant contributions to visibility impairment at Federal Class I areas which are the subject of the haze rule. Previous work has demonstrated that the REMSAD modeling system may be useful for conducting year-long simulations with adequate complexity and sophistication to accurately simulate annual aggregate visibility statistics.

The present work describes a series of platform intercomparison experiments which were intended to demonstrate portability of the REMSAD modeling system between various UNIX, LINUX or Windows based platforms that may be used for MANE-VU modeling in the future. Working within the constraints of a one-day simulation and given the limited number of grid cells exhibiting significant differences between platforms/compilers, this study concludes that REMSAD, when ported to a SUN, Windows NT or Linux platform, is suitable for regulatory investigation of regional haze, particulate matter or other topics where differences of less than 1 ng/m^3 (roughly 10,000 times smaller than the proposed standard for PM) are required. This work also suggests that a different precision threshold should be considered for gas phase species and perhaps coarse particles that have significantly higher ambient mass concentrations, and thus greater average differences between platforms.

Additional investigations are required to determine whether the conclusions drawn here are supported when longer simulations are conducted. The suitability of the proposed precision thresholds for individual species also warrants further study; however, given current photochemical grid modeling accuracy, which can be as poor as a factor of two for many species, the observed cross-platform precision is not a priority for further investigation.

I. Introduction

As part of state efforts to comply with the requirements of U.S. Evironmental Protection Agency's (USEPA) 1999 Regional Haze Rule, The Mid-Atlantic/Northeast Visbility Union (MANE-VU) has been developing technical expertise for regional-scale air quality modeling. A variety of modeling systems and platforms are available for conducting assessments of pollutant contributions to visibility impairment at Federal Class I areas which are the subject of the haze rule.

Chemical Transport Models (CTMs) have traditionally been developed on UNIX based systems. This has, in some cases, prevented the development of a wider user base among agencies with limited computational resources. Recently, the processing capabilities of personal computers have evolved such that they can compete and, in some cases, outperform sophisticated air quality models run on UNIX based workstations. However, these models are not always directly portable from platform to platform. In addition, variations in processors, operating systems and compilers or actual coding may affect model results.

To assess the issues associated with platform portability, NESCAUM has conducted a direct comparison of results obtained with identical model inputs and identical model versions running on four separate platforms. The results of this portability study demonstrate that the model performs similarly on all platforms. Application of these models for regulatory purposes will require that results derived on a non-native platform be consistent – within a prescribed threshold – to results obtained by running the model on its native platform. The air quality modeling community has not reached a consensus regarding this prescribed threshold; however, 1 part per trillion (ppt) for ozone and 1 ng/m^3 for particles have been suggested as appropriate values to explore.[1]

Previous work has demonstrated that the Regulatory Modeling System for Aerosols and Deposition (REMSAD) may be useful for conducting year-long simulations with adequate complexity and sophistication to accurately simulate annual aggregate visibility statistics (Wayland, 1999). This memorandum details a platform portability study of REMSAD, a USEPA approved alternative regulatory model for $PM_{2.5}$, regional haze, and toxic, nitrogen, and acid deposition. This study was performed by Northeast States for Coordinated Air Use Management (NESCAUM) and the New York Department of Environmental Conservation (NYDEC). Results from four separate platforms are compared: the original Systems Applications International (SAI) DEC Alpha system, NYDEC's SUN OS system and NESCAUM's Windows NT and Linux PC systems[2]. The comparison shows little difference amongst the platforms.

In general, cross-platform model performance was best for particulate species with somewhat larger differences found for gaseous species. The differences on average are close to zero and normally distributed. The present analysis confirms that the REMSAD model performs on three different platforms within the proposed thresholds for all species except carbon monoxide and ozone. This study

[1] These values were suggested as appropriate at the LADCO/TNRCC Linux Computing Workshop, Austin, Texas, March 25-27, 2002.

[2] Throughout this report, SAI/DEC Alpha, NYDEC/SUN, and NESCAUM/Linux and NESCAUM/Windows NT are used as shorthand for both operating system and computational hardware; SAI/DEC Alpha refers to a DEC Alpha running Compaq Tru64 UNIX V5.1A, NYDEC/ SUN refers to a SUN Ultra 2 running SUN OS v5.7, NESCAUM/Linux refers to an Intel x86 PC running Linux RedHat v7.2, and NESCAUM/Windows NT refers to an Intel x86 PC running Windows NT 2000.

does not, however, address model accuracy and comparison of results to other photochemical models and to observations is required to establish overall model performance.

II. Methodology

REMSAD[3] version 6.3 model code was ported to NYDEC and NESCAUM where it was installed on a SUN workstation (in the case of NYDEC) and on a Windows NT and Linux PC (NESCAUM). Steps in this process included compiling and linking the source code and altering provided scripts; however, it is important to note that identical source code was the basis for both installations. The source code was obtained from ICF Consulting/Systems Applications International (SAI), the copyright owner and a contractor to USEPA responsible for developing the REMSAD modeling system. SAI has developed the REMSAD system on a DEC Alpha workstation as its native platform.

SAI provided the necessary meteorological, boundary condition, and initialization files for a 1 day simulation of January 1, 1996 to conduct the benchmarking experiment. These data are available on the REMSAD website.[4] To determine REMSAD's portability, model results from the SUN, Windows NT and Linux environments were compared with those provided by SAI's DEC Alpha system.

The model was run on each platform using identical inputs. In addition to the required input files, SAI also provided a C shell script that allows REMSAD to be run using identical command line entries. Included in that script are a prescribed set of model flags that had been used for the one-day simulation and specification of the input/output files used. Table 1 lists the flag settings.

SAI provided thirteen species concentration binary output files representing the concentrations within the first model layer averaged for January 1, 1996. These species include: sulfate formed through the aqueous phase, sulfate formed through the gaseous phase, particulate nitrate, coarse particles, primary fine particles, ammonium associated with nitrate, ammonium associated with sulfate, nitrogen oxide, nitrogen dioxide, carbon monoxide and ozone. To obtain similar output binary files the extract program was used. This program post-processes and calculates 24-hour average values from the REMSAD output file, atdm.0101.test.avrg, which contains hourly averaged species concentrations. These post-processed files are in binary format and require a graphical package such as the Package for Analysis and Visualization of Environmental data (PAVE) and view. NESCAUM installed the Linux beta version, PAVE 2.1, to conduct the graphical comparisons.

[3] REMSAD is used here to refer to the ATDM (the aerosol transport deposition model) within the REMSAD framework. Other components of REMSAD are referenced separately.
[4] The REMSAD modeling homepage provided by SAI is located at: http://remsad.saintl.com

Table 1. REMSAD model parameters and inputs used in this exercise.

Model Parameters / Inputs	Values	Notes
Simulation Day	January 1, 1996	
Coarse Grid	36 km.	Lat/long coordinates
Fine Nesting	NO	
Model Boundaries (Western longitude, Eastern longitude, Southern latitude, Northern latitude)	-126.00, -66.00, 24.00, 52.00 degrees	
Vertical Resolution	12 layers	
Meteorological Input files	MM5	
Chemistry Mechanism Invoked	Micro CM-IV	
OTHER FLAGS	**Values**	**Notes**
LDEPN, LWET	True	Wet and dry deposition calculated
LSED	True	Species-dependent settling velocities are calculated for pm depositions
LTXCHM	False	No toxic chemistry invoked
LAREA, LPTS	True	Area and point source emissions included
LSTAGR	False	Horizontal wind components are defined at cell centers
LO3	False	Do not need Ozone and OH files
DTSTEP	30.0 minutes	Maximum time step for advection calculations
DLONG, DLAT	0.500, 0.333 degrees	East-west grid spacing, north-south grid spacing

III. Compiling and Linking Code

REMSAD version 6.3 source code was downloaded from SAI's modeling web site. In addition to the programs (written in the Fortran programming language) there were script files for compiling the source code and executing the program. SAI's script files contain commands that are executed directly within the DEC operating system's C shell environment and thus revisions are necessary in order for successful execution on a non-native platform that does not support the C shell. A listing of these files is contained in Table 2.

Table 2. Operating system executable files requiring alteration for successful execution on a non-native platform

REMSAD:

makefile → *compiles and links REMSAD source code*
all.test.new.job → *C shell script that runs the REMSAD executable for a one-day simluation*

EXTRACT:

makefile → *compiles and links Extract source code*
xyex.daily.pms.all.job → *C shell script that runs the Extract executable to create smaller output files formatted for use with graphical viewers such as PAVE*

The source code consists of approximately 142 subroutines which are linked together during the compilation process. SAI supplied the makefile for compiling and linking the code. NESCAUM used the Portland Groups Inc. (PGI) Fortran compiler v3.1 on the Windows NT PC and PGI v3.3-2 on the Linux PC. Given the use of both a different platform and Fortran compiler, it was necessary to alter the makefile beyond the basic path changes suggested on the SAI website.

No significant changes were necessary to run the script "all.test.new.job" on the Linux platform; however, due to the limited choice of shell languages on the Windows NT platform, translation of the provided C shell script to the Bash shell language was required. NESCAUM had already completed this exercise in 2001 when installing REMSAD version 4.0, making the translation a second time trivial.

The Extract program, which is a post-processor developed to manipulate the large output binary files, was made available by SAI. A second makefile was also available to compile and link the extract program along with scripts for running the program. As with the REMSAD makefile, NESCAUM

altered the Extract makefile for both the Windows NT and Linux systems and then translated the scripts used by the Windows NT from the C shell to the Bash shell language.

Platform specific details regarding necessary changes for installing REMSAD on the Windows NT, Linux and SUN OS computers are provided in the following sections.

A. Linux PC[5]

REMSAD and Extract executables were created using the provided makefiles on the Linux platform. No compiler errors were discovered so the only changes necessary involved the compiler call. The REMSAD and the Extract source code is all written in Fortran77, however, three subroutines within the REMSAD code were identified which included Fortran90 statements.[6] To allow the REMSAD source code to compile, PGI Fortran90 was used. The Extract code was compiled using PGI Fortran77. To account for these compiler differences, the "f77" call contained in each makefile was switched to "pgf90" and "pgf77" for the REMSAD and Extract source code respectively. The following compiler flags were used (see footnotes for details):

- REMSAD source code: -O [7] –Mbyteswapio [8] –Mextend [9]
- Extract source code: -O –Mbyteswapio.

NESCAUM chose to compile the source code with Fortran90 as opposed to altering the code to allow for compiling with Fortran77 (see footnote 6). However, after viewing the initial results of the study, NESCAUM altered the source code as discussed in footnote 6 and then invoked the Fortran77 compiler with the same compiler flags as shown above. The results of these exercises are provided (See Appendix C). An estimate of model variation resulting from different compilers can be obtained by

[5]The PC specifications: dual 2.2 GHz Intel Xeon Processors with 512K level 2 Cache, Intel 860 Chipset with 400 MHz system bus, 2 GB dual-channel PC800 RDRAM memory, and 2 73 GB 10,000 RPM Ultra 160 SCSI hard drives.

[6] Three subroutines (rchem.f, rchem2.f, rchemf.f) use the term "cycle" to end an "If" statement. This term is a Fortran90 term though compatible with some Fortan77 compilers. For instance, SAI found their Fortran77 compiler understood the statement while New York DEC (using SUN Workshop 5 compiler) and NESCAUM (using PGI compiler) found it necessary to either compile with a Fortran90 compiler or to change the function to a compatible Fortran77 statement. Kevin Civerolo at NY DEC changed the statement to a Fortran77 statement by replacing the term "cycle" with a "block if" statement. NESCAUM ran the one-day simulation both ways after confirming the second approach with Tom Myers at ICF Consulting/SAI (March, 2002). A comparison of the results is presented later in the document.

[7] The "O" flag refers to the default optimization level.

[8] This flag is compatible with "-convert big_endian". Originally we had thought this flag was not necessary as guidance for UAM-V stated this flag was "DEC specific and are … available or necessary on systems other than DEC" (http://uamv.saintl.com/faq.htm). It is believed Linux and Windows NT were not considered as platform options for model simulations at the time the above guidance was written. When the source code was compiled without this flag, the executable created was unable to read from the binary land use input file, remsad.36km_lu.full.bin, reporting an unformatted sequential access error with record 1. The executable obtained when either the Mbyteswapio or byteswapio flag was used did not report any difficulties reading from the binary input files nor differences in the result files.

[9] This flag is compatible with "-extend_source". It was necessary for us to use the "Mextend" flag which tells the compiler to accept 132 column source code (i.e. not the generally accepted limit of 72 columns). Without this flag, the code did not compile. (http://uamv.saintl.com/faq.htm;http://www.ibiblio.org/pub/languages/fortran/ch3-1.html, http://www.llnl.gov/icc/lc/asci/fpe/fpe.options.html)

comparing the NESCAUM output shown in Appendix A and C. To preserve SAI's code free from changes, the results from the Fortran90 version are used for comparison purposes to other platforms.

B. Windows NT[10]

Installing REMSAD onto the Windows NT platform was more challenging than for the Linux system. However, prior work installing REMSAD version 4.0 onto a Windows NT system provided helpful experience for this undertaking as similar compiler errors were encountered during the installation of both versions 4.0 and 6.3. In order to overcome these obstacles, the following steps were taken:

- REMSAD has a common memory block file called com3.cmd. Windows NT already uses a file with that name which prevented downloading the model's file onto the system. The REMSAD version of this file was renamed com33.cmd and all relevant include statements in the source code were changed from "include com3.cmd" to "include com33.cmd".

- The variable "spec" located in the rdobs.f file was changed from an "integer" declaration to "character*4". This was necessary for compiling.[11]

- The variable "h1plus" was added to the end of the rpmares subroutine call located in remeql3.f. This was necessary for compiling.[12]

- The script files were converted from C shell to Bash shell language.

- The compiler options in the makefile were changed to "-O –Mbyteswapio -Mextend". The compiler call was switched from "f77" to "pgf90". See Linux description above for more description on these flags.

Extract Program:

- The variable MSUB was added to the XTRACT.INC include file: "CHARACTER*4 MSUB". This was necessary for compiling.[13]

[10] 850 MHz PC with a Pentium III processor.
[11] Suggested via electronic mail by Tom Myers at ICF Consulting (March, 2002).
[12] This was confirmed by Tom Myers at ICF Consulting (March, 2002) with the additional comment: "h1plus is no longer used in the code". Hence, this change was not necessary when compiling on the Linux platform.
[13] This was suggested through correspondence with Brian Timin at OAQPS, US. EPA (Jan, 2001). This was confirmed by Tom Myers at ICF Consulting (March, 2002).

- The makefile for the extract program was changed to reflect our compiler brand. The compiler options were switched to "-O –Mbyteswapio". The compiler call was switched from "f77" to "pgf77".

C. SUN Workstation[14]

NYDEC did not encounter significant installation issues except as noted in footnote 6. The Fortran compiler used was SUN Workshop 5.(f77 Version 5.0, f90 Version 2.0). The following command line was used to compile the code:

- **REMSAD source code: -O –e -c**
- **Extract source code: -O -c**

Table 3. Summary table of compilers and options

Platform/System	Compiler	Command	Flags/Options
DEC Alpha (SAI)	Compaq Fortran77	f77	-O –static[1] –switch fe_ioworst[1] -convert big_endian –non_shared[1] –extend_source
Linux PC (NESCAUM)	PGI v3.3-2	pgf90[1]	-O –Mbyteswapio -Mextend
Windows NT (NESCAUM)	PGI v3.1	pgf90	-O –Mbyteswapio -Mextend
SUN (NYDEC)	Sun Workshop 5	f77	-O –e -c

The Windows NT run required 51 minutes of processing time to simulate the one-day comparison while the Linux PC required 30 minutes of processing time. The SUN system required approximately **1.5** hours. As expected, the chemistry routines account for over 80% of the processing time.

IV. Results

Results are presented for the concentrations of thirteen species located within the first layer of the model and averaged over the full day. No differences were discernable between the Linux PC and

[14] SUN Ultra 2 workstation (450 MHz, 4 Gb memory, OS Version 5.7)

Windows NT (i.e. Linux values minus Windows NT values were zero for all grid cells and species to within a precision of 0.001 ng/m^3.) Hence, for the purposes of this study, the Linux results are interchangeable with the Windows NT. The species mass output, the difference plots and the relative percent difference plots for Linux, SUN and DEC Alpha runs are included in Appendices A and B.

Absolute differences were calculated for each of 9,676 grid cells (82 by 118) in the spatial domain for each of the 13 species of interest. This involved subtracting the 24-hour average value calculated for each grid cell and species on each platform from the value calculated on the other platforms. For the two ammonium species and particulate nitrate, relatively small to negligible differences (always less than or equal to 1 ng/m^3) were calculated between 24-hour average values calculated on the different platforms. Slightly larger (less than 4 ng/m^3) differences were calculated for a minimal number of grid cells for the two sulfate species and nitrogen oxide. Differences in coarse particles were generally less than 6 ng/m^3, but ranged as high as 150 ng/m^3 in a few grid cells. Primary fine particles had differences of up to about 20 ng/m^3 with the primary organic fraction contributing almost half of that difference. Differences up to 6-7 ng/m^3 were found for nitrogen dioxide.

By far, the greatest discrepancies were observed between the gas phase species, particularly carbon monoxide and ozone (nearly 9 μg/m^3). This is not surprising given that the ambient concentrations of these gases are significantly higher (on a mass basis) than the concentrations of the particulate matter components calculated. These differences would translate to about 6 to 7 ppb of CO and about 1ppb of Ozone, which is small relative to ambient concentrations of these gases of about 100-200 ppb and 40-50 ppb respectively. Despite the fact that differences are small relative to ambient, it does raise the question of why the same model on different platforms would produce differences of this magnitude.

Tables 4, 5 and 6 summarize the distribution of absolute differences calculated between platforms. In each of these tables (focusing on the differences between the Linux versus Sun, Linux versus DEC and Sun versus DEC platforms respectively) the range of differences are listed. A second column contains the value of three times the standard deviation, or the "3σ level". The 3σ level gives the range within which 99.7% of all values fall. The difference between this value and the absolute range demonstrates that the absolute range is driven by a few outliers rather than typical values of the grid cell difference. Finally, in a third column, are listed the number of grid cells whose difference lies outside the threshold of 1 ng/m^3. While this is a somewhat arbitrary threshold, it does point out those species with a larger number of differences that are potentially of concern.

The figures in Appendix A do not show the full extent of the differences found across the coarse grid,[15] due to the choice of color scale. These figures use a color scale specifically selected to highlight the spatial distribution of the extremes. Closer examination of the discrepancies in absolute difference show that the values are normally distributed about zero. Standard deviations were calculated for each difference pairing in order to determine the 3σ levels presented in the tables. The average 3σ level for all species (except carbon monoxide and ozone) are 0.1, 0.8 and 0.8 ng/m^3 for Linux –Sun, Linux – DEC Alpha and SUN – DEC Alpha, respectively (see Tables 4-6 for a complete listing). Since the 3σ values are less than the threshold for significant differences (1 ng/m^3), the disparities for these eleven species

[15] The coarse grid contains 119 cells East to West and 83 cells South to North; however, the first row and column provide boundary conditions and do not contain true values leaving a 118 by 82 cell grid for a total of 9,676 cells with calculated values.

are considered minimal. However, the same cannot be said for carbon monoxide and ozone, as their 3σ levels are on the order of 1 $\mu g/m^3$ for CO and 0.1$\mu g/m^3$ for O_3 for all pairings.

The reproducibility of local minima and maxima is particularly important for in a model designed to serve regulatory purposes since attainment and nonattainment designations (as well as reasonable progress goals) are generally based on the extreme values of air quality metrics. In general, the model predicts similar magnitude and location of local minima and maxima for all species across platforms. For some species, the absolute difference in the maximum value or minimum value (across the whole domain) was non-zero, but still less than 1 ng/m^3.

To place the absolute differences in perspective relative to a species' typical mass concentration, the relative percent difference for each pairing was calculated by dividing the absolute difference by the average grid cell value and multiplying the result by 100. This calculation was also made for each species and in each of the 9,676 grid cells in the geographic domain. Minimum and maximum values of the relative percent difference for each species (corresponding to the minimum and maximum values of the absolute difference range) are listed in Table 7. Here, all species but nitrogen oxide have relative percent differences less than 10 percent.[16] Ozone and carbon monoxide, the species with the greatest absolute difference, by this measure are now comparable with the other species. However, this table does not give any indication of the number of grid cells affected.

The most significant differences were clustered over the Atlantic Ocean and Florida. Fewer maximum grid cell differences were seen in the Maryland area and a scattering of individual cells were observed across the grid. These outliers do not tend to occur in the areas of maximum species concentration, with the exception of ozone, off the Atlantic coast of Florida.

The study was extended to include a comparison of the effects of changing compilers to PGI Fortran77 from PGI Fortran90. This comparison was restricted to runs simulated on the Linux PC. Appendix C provides representative figures of those species demonstrating differences and Table 8 summarizes the absolute and relative range of the differences. The species include aqueous formed sulfate, carbon monoxide, primary organic aerosols, ammonium associated with sulfate, nitrogen oxide and ozone. The remaining species displayed no significant differences. Nitrogen oxide, ozone and carbon monoxide represent the species with the largest absolute and/or relative differences. These values are comparable with those found in Tables 4-7. As with the comparison illustrated in Appendix A, the majority of the differences shown are located over the Atlantic Ocean off the South Carolina and Florida coast.

[16] The maximum relative differences for nitrogen oxide found are located in the Atlantic ocean off the coast of South Carolina/Georgia where the NO values are very low (approximately 0.001 $\mu g/m^3$).

Table 4 Results of a one-day January 1, 1996 REMSAD v6.3 simulation comparison between versions ported to a Linux PC and a SUN Workstation. This table shows the range of the absolute differences between model output from the two platforms. Values were rounded to the nearest hundredth of a nanogram. Model differences within a ± 1.00 ng/m^3 benchmark threshold are considered acceptable. Model cross-platform performance is, in general, worse for the gaseous species.

Species	REMSAD I.D.	Linux - SUN		
		Absolute Range (ng/m^3)	3 σ (ng/m^3)	# of cells exceeding threshold value*
Sulfate Formed in the Aqueous Phase	**ASO4**	-2.26 to 4.40	0.22	7
Sulfate Formed in the Gaseous Phase	**GSO4**	-0.12 to 0.09	0.01	0
Particulate Nitrate	**PNO3**	-0.86 to 0.87	0.06	0
Coarse Particles	**PMCOARS**	-2.10 to 0.27	0.07	1
Secondary Organic Aerosols	**SOA**	-0.03 to 0.13	0.01	0
Primary Organic Aerosols	**POA**	-0.47 to 1.42	0.08	2
Primary Fine Particles	**PMFINE**	-0.55 to 1.55	0.11	3
Ammonium associated with Sulfate	**NH4S**	-1.05 to 0.92	0.10	1
Ammonium associated with Nitrate	**NH4N**	-0.25 to 0.25	0.02	0
Nitrogen Oxide	**NO**	-1.09 to 2.52	0.10	1
Nitrogen Dioxide	**NO2**	-6.40 to 7.89	0.60	15
Carbon Monoxide	**CO**	-8363.62 to 6424.97	1032.16	123
Ozone	**O3**	-1919.82 to 2545.04	146.77	198

* A threshold value of 1.3 ng/m^3 (equivalent to 1ppt) was used for NO; 2.1 ng/m^3 (equivalent to 1ppt) for NO_2; 2.1 ng/m^3 (equivalent to 1ppt) for ozone and 1.3 µg/m^3 (equivalent to 1ppb) for CO.

Table 5. Results of a one-day January 1, 1996 REMSAD v6.3 simulation comparison between versions provided by SAI using a DEC Alpha and a version ported to a Linux PC. This table shows the range of the absolute differences between model output from the two platforms. Values were rounded to the nearest hundredth of a nanogram. Model differences within a ± 1.00 ng/m^3 benchmark threshold are considered acceptable. Model cross-platform performance is, in general, worse for the gaseous species.

Species	REMSAD I.D.	Linux - DEC		
		Absolute Range (ng/m^3)	3 σ (ng/m^3)	# of cells exceeding threshold value*
Sulfate Formed in the Aqueous Phase	**ASO4**	-1.16 to 4.40	0.19	5
Sulfate Formed in the Gaseous Phase	**GSO4**	-1.57 to 0.07	0.09	3
Particulate Nitrate	**PNO3**	-0.86 to 0.66	0.07	0
Coarse Particles	**PMCOARS**	-149.64 to 1.01	5.59	318
Secondary Organic Aerosols	**SOA**	-7.83 to 0.13	0.28	6
Primary Organic Aerosols	**POA**	-10.19 to 1.42	0.67	46
Primary Fine Particles	**PMFINE**	-21.70 to 1.55	1.03	64
Ammonium associated with Sulfate	**NH4S**	-1.05 to 0.92	0.10	1
Ammonium associated with Nitrate	**NH4N**	-0.25 to 0.19	0.02	0
Nitrogen Oxide	**NO**	-1.09 to 2.52	0.10	1
Nitrogen Dioxide	**NO2**	-6.36 to 6.56	0.56	14
Carbon Monoxide	**CO**	-8770.36 to 6016.64	1040.05	130
Ozone	**O3**	-2222.11 to 332.48	76.94	143

* A threshold value of 1.3 ng/m^3 (equivalent to 1ppt) was used for NO; 2.1 ng/m^3 (equivalent to 1ppt) for NO_2; 2.1 ng/m^3 (equivalent to 1ppt) for ozone and 1.3 µg/m^3 (equivalent to 1ppb) for CO.

Table 6. Results of a one-day January 1, 1996 REMSAD v6.3 simulation comparison between versions provided by SAI using a DEC Alpha and a version ported to a Sun Workstation. This table shows the range of the absolute differences between model output from the two platforms. Values were rounded to the nearest hundredth of a nanogram. Model differences within a ± 1.00 ng/m^3 benchmark threshold are considered acceptable. Model cross-platform performance is, in general, worse for the gaseous species.

Species	REMSAD I.D.	SUN - DEC		
		Absolute Range (ng/m^3)	3 σ (ng/m^3)	# of cells exceeding threshold value*
Sulfate Formed in the Aqueous Phase	**ASO4**	-0.83 to 2.19	0.10	2
Sulfate Formed in the Gaseous Phase	**GSO4**	-1.56 to 0.07	0.09	3
Particulate Nitrate	**PNO3**	-0.83 to 0.024	0.06	0
Coarse Particles	**PMCOARS**	-149.67 to 1.01	5.56	318
Secondary Organic Aerosols	**SOA**	-7.83 to 0.08	0.28	6
Primary Organic Aerosols	**POA**	-10.19 to 0.27	0.67	44
Primary Fine Particles	**PMFINE**	-21.70 to 0.43	1.02	61
Ammonium associated with Sulfate	**NH4S**	-0.80 to 0.63	0.04	0
Ammonium associated with Nitrate	**NH4N**	-0.24 to 0.07	0.02	0
Nitrogen Oxide	**NO**	-0.26 to 1.14	0.05	0
Nitrogen Dioxide	**NO2**	-7.91 to 1.74	0.32	3
Carbon Monoxide	**CO**	-6513.40 to 5330.16	796.55	89
Ozone	**O3**	-2494.33 to 2153.91	140.08	141

* A threshold value of 1.3 ng/m^3 (equivalent to 1ppt) was used for NO; 2.1 ng/m^3 (equivalent to 1ppt) for NO_2; 2.1 ng/m^3 (equivalent to 1ppt) for ozone and 1.3 µg/m^3 (equivalent to 1ppb) for CO.

Table 7 Results of a one-day January 1, 1996 simulation when REMSAD v6.3 was ported to a Linux PC and SUN Workstation compared to those provided by SAI using a DEC Alpha. This table shows the relative percent difference (RPD) between two models where RPD is defined as (A-B)/[(A+B)/2]. On a relative basis, model cross-platform performance of the gaseous species is similar to other species.

Species	REMSAD I.D.	Linux - SUN Relative % Difference	Linux – DEC Alpha Relative % Difference	SUN - DEC Alpha Relative % Difference
Sulfate Formed in the Aqueous Phase	**ASO4**	-1.2 to 1.5%	-1.0 to 1.5%	-0.3 to 1.4%
Sulfate Formed in the Gaseous Phase	**GSO4**	-0.5 to 0.5%	-0.6 to 0.5%	-0.6 to 0.3%
Particulate Nitrate	**PNO3**	-8.3 to 3.8%	-8.3 to 3.7%	-5.1 to 1.9%
Coarse Particles	**PMCOARS**	-0.1 to 0.2%	-1.1 to 0.2%	-1.1 to 0.04%
Secondary Organic Aerosols	**SOA**	-0.1 to 0.5%	-0.4 to 0.5%	-0.4 to 0.02%
Primary Organic Aerosols	**POA**	-1.9 to 6.6%	-1.9 to 6.6%	-0.4 to 0.03%
Primary Fine Particles	**PMFINE**	-1.9 to 5.8%	-1.9 to 5.8%	-0.7 to 1.3%
Ammonium associated with Sulfate	**NH4S**	-3.7 to 2.9%	-3.7 to 2.9%	-3.7 to 2.9%
Ammonium associated with Nitrate	**NH4N**	-8 to 4%	-8 to 4%	-5 to 2%
Nitrogen Oxide	**NO**	-20.3 to 29.1%	-18.6 to 33.9%	-3.6 to 10.9
Nitrogen Dioxide	**NO2**	-8 to 8%	-8 to 10%	-1 to 2%
Carbon Monoxide	**CO**	-7 to 7%	-7 to 6%	-7 to 5%
Ozone	**O3**	-2.2 to 2.8%	-2.6 to 0.4%	-2.7 to 2.5%

Table 8. Differences found when REMSAD v6.3 is run with PGI Fortran90 (Linux90) compiler versus PGI Fortran77 (Linux77) compiler. The PGI Fortran77 on Linux results are also compared to the DEC platform which uses a different Fortran Compiler, a different operating system, and different hardware.

Species	REMSAD I.D.	Linux90 – Linux77		Linux77 – DEC Alpha	
		Absolute ($\mu g/m^3$)	Relative (%)	Absolute ($\mu g/m^3$)	Relative (%)
Sulfate Formed in the Aqueous Phase	**ASO4**	-0.0010 to 0.0019	-0.54 to 0.73	-0.0020 to 0.0044	-1.0 to 1.5
Primary Organic Aerosols	**POA**	-0.0002 to 0.0002	-0.5 to 0.7	-0.0102 to 0.0014	-1.9 to 6.8
Ammonium associated with Sulfate	**NH4S**	-0.0000 to 0.0002	-0.09 to 0.7	-0.001 to 0.0009	-3.6 to 2.9
Nitrogen Oxide	**NO**	-0.001 to 0.003	-5.2 to 20.9	-0.001 to 0.003	-12.5 to 27.0
Carbon Monoxide	**CO**	-8.7 to 4.7	-6.4 to 5.2	-5.7 to 6.4	-5.9 to 7.1
Ozone	**O3**	-2.181 to 0.403	-2.6 to 0.5	-0.168 to 0.137	-0.20 to 0.16

V. Discussion/Conclusion

The porting of REMSAD version 6.3 source code from its native environment to an alternative computer platform has been shown to result in discrepancies which arise from a combination of differences in operating system, hardware, and source code compilers. It is difficult to de-couple the effects of changing compilers from the combined effects due to switching operating systems and hardware, however, results suggest that the choice of compiler may be the most significant contributor. Based on the current analysis, the gaseous species nitrogen oxide, ozone and carbon monoxide appear to be most sensitive to platform/compiler differences. This may be related to the differences in ambient concentrations rather than the particular phase of individual species, but may, nonetheless, have implications for the success of true "one atmosphere" modeling.

Observed differences are greatest between the native model run on a DEC Alpha workstation and the other three platforms tested, however, some significant discrepancies occur between all platforms. While generally small, all differences must be understood given that identical model inputs were used on the same version of the model. It should be noted, however, that most of the significant discrepancies (those greater than 1 ng/m^3 for particulate phase species and those greater than 1ppt for NO, NO_2, and ozone; 1ppb for CO) occur in a relatively few grid cells. The vast majority of 9,676 grid cells that were simulated to have ambient PM concentrations within a benchmark threshold of 1 ng/m^3 to those values produced on other platforms.

While these exercises are useful for confirming sufficient precision between platforms to conduct regulatory demonstrations for regional haze and PM issues, additional work – including comparison of results to other models and observations – is needed to determine the models accuracy which is expected to be a much greater contributor to overall uncertainty.

A one-day study is not sufficient for testing the potential propagation of small errors into more significant discrepancies as a model is run for longer periods of time. The limited duration of the simulation compared here also has implications for the accurate calculation of secondary aerosol species. These secondary components result primarily from gas to particle conversion in the atmosphere on timescales that are often longer than a single day. Although modeling multiple days (on the order of 10 to 15 days) would provide additional information, such runs would require additional multi-day simulations by the model developer on the native DEC Alpha platform.

The present analysis is also limited in that it compares values averaged over the first modeled day. In general, spin up time is required to allow a model to stabilize relative to the initial conditions. Ideally, a follow-up study would focus on hourly results or even instantaneous results obtained after the model has been run for a substantially longer simulation period. Unfortunately, this would also require the availability of multi-day model inputs provided by the model developer.

Despite these shortcomings, considerable knowledge was gained through the exercises described in this memorandum and it serves as a framework for evaluating the precision of other photochemical grid models – and future versions of REMSAD – across platforms. Based on the one-day simulation results, we find similar numerical results across platforms with some species performing significantly better than others. The Linux platform demonstrates superior computational efficiency, executing the 24-hour simulation in one third the time of the SUN and half the time of the Windows NT platform. As various regions and stakeholders prepare to run models on different platforms, further tests may be necessary to determine if a platform dependent model bias exists, but the present study does not find any bias greater than 10^{-6} $\mu g/m^3$.

The results presented in Appendix C of this memorandum indicate that compiler choice certainly effects the numerical stability of model simulations. The exact mechanism that leads to observed differences remains unknown, however, these results indicate that greater differences result from the use of alternative compilers than differences in operating system. Differences found when REMSAD version 6.3 is ported to a Windows NT and Linux platform using the same Fortran compiler are significantly smaller than differences found with different operating systems and different compilers. One plausible explanation for the observed differences is that the Fortran library routines developed by different vendors contain slight differences in algorithms. REMSAD source code calling such functions might return slightly different values each time the functions are called with the potential to compound these slight differences as computation proceeds.

Working within the constraints of a one-day simulation and given the limited number of grid cells exhibiting significant differences between platforms/compilers, this study concludes that REMSAD version 6.3, when ported to a SUN, Windows NT or Linux platform, is suitable for regulatory investigation of regional haze, particulate matter or other topics where differences of less than 1 ng/m^3 (roughly 10,000 times smaller than the proposed standard for PM) are required.

References

Wayland, *REMSAD 1990 base case simulation: model performance evaluation annual average statistics*, USEPA Office of Air Quality Planning and Standards, Research Triangle Park, NC, 1999.