Atmos. Chem. Phys. Discuss., 8, S3759–S3773, 2008
www.atmos-chem-phys-discuss.net/8/S3759/2008/

Atmospheric
Chemistry
and Physics
Discussions

Interactive comment on “A computationally-efficient secondary organic aerosol module for three-dimensional air quality models” *by* P. Liu and Y. Zhang

P. Liu and Y. Zhang

Received and published: 16 June 2008

Reply to Referee #1

This paper presents an analysis of how CPU time can be reduced by relaxing stringency criteria for an aerosol model, along with a statistical fitting of the UNIFAC approach. Such an exercise is valuable, and the CPU speed-up obtained with a parameterized UNIFAC was interesting, but this paper has however a number of major shortcomings if the goal is to provide techniques for use in 3D atmospheric models.

Most importantly, the test cases considered are too few and too simple. A simple box model is used, with no emissions, no deposition, and a very long (24h) integration period. This is not an appropriate way of testing solvers designed for 3-D models, in which the chemistry is essentially re-started after each advection time step (e.g. 5-20mins). One long simulation leads to a rather gentle development of the chemical processes, and highly favours some types of solvers (especially Gear-type) over others which are better suited to rapid-restarts. The paper by Saylor and Ford (AE, 1995) stressed this problem, and suggest approaches to provide tests which are more appropriate for 3D models.

Reply:

There seems to be some misunderstanding regarding the integration time step for gas-phase chemistry, aerosol dynamics and thermodynamics used in our study. The time step we used is 1 minute, rather than 24-hr. The total simulation period is 24 hours. A time step of 1 minute for chemistry integration is among the smallest time steps used in current 3-D models. While a smaller time step can be used for box model studies, its use in 3-D models is usually too expensive to afford, so the use of 1-minute as a time step to derive the SOA parameterization is adequate in our study.

This has been clarified in the revised paper.

Another major problem of this paper is the very restricted number of test cases used, only four, and their lack of consideration of non-typical conditions. Real 3D modeling exposes the numerical scheme to a huge number of different concentration regimes, and it is important to try to cover the extremes of these regimes as well as more typical situations. It is the extremes which result in numerical instability or convergence problems. A numerical scheme which works fine for the "typical sunny day" as discussed in this paper may well run into severe trouble for less typical situations, for example an urban plume moving into a forested area in the middle of a heat-wave. For example, Hesstvedt explored 21 scenarios, with a factor 100 range in both NOx and HC levels, and found very different accuracy for different cases.

Reply:

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper

Note that the parameterization of activity coefficients in our study was obtained by accounting for a broad range of physical conditions and chemical concentrations of independent SOA precursor species (i.e., temperature ranges from 253 to 313 K; mole fraction ranges from 0 to 1), with a total of 1.16 Œ 108 simulations are conducted (120 temperatures Œ ˜106 mole fractions of 6 hydrophilic species). For MADRID2_FAST evaluation, we agree with the reviewer that extreme, non-typical conditions should be included in the test conditions to provide stress tests.

To address comments from the reviewer, two additional scenarios and two low temperatures for free tropospheric conditions have been added in our evaluation of MADRID 2_FAST:

(1) the rural condition with a very high isoprene mixing ratio (10 ppb) (2) the forest condition with an extremely low NOx (100 ppt) and a very high isoprene mixing ratios (10 ppb). (3) In addition, the temperature low limit has been extended from 273 K to 253 K for all 6 test conditions.

The above additions increase the total number of test simulations from 100 to 210 (= 7 temperatures Œ 5 RHs Œ 6 conditions).

In addition, we’ve tested 3-D CMAQ-MADRID 2_FAST against CMAQ-MADRID 2 for 24-hr over the contiguous U.S. and the results appear to be realistic. The spatial distributions of daily mean absolute and percentage differences of simulated hourly O3 and 24-hr average PM2.5 by the two simulations show the negligible differences. For example, the absolute differences for O3 concentrations are from -0.02 to 0.02 ppb (with percentage differences of -0.07 to 0.1%), and the absolute differences for PM2.5 concentrations are from -0.60 to 0.06 μg m-3 (with percentage differences of -6.1 to 1.5%). In our 3-D test simulations, the NOx mixing ratios range from less than 1 ppb to 139 ppb, and the ratios of NOx over VOCs vary from less than 1 to several hundreds or thousands. The temperature ranges from 259 to 318 at surface, and 205 to 231 K at the top of model layers. These values cover a broad range of

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper

atmospheric meteorological and chemical conditions. While the 3-D model test supplements the box model test presented in our current paper, it will be included in a separate paper, as stated in the conclusion section.

The above points have been added in page 12, first paragraph, page 19, lines 8-11. The test results under the new conditions have been added in Figures 3, 4, and 5.

In this context, the maximum isoprene concentration of 2 ppb is very low and certainly not sufficient to constrain a numerical scheme. Even for regional modelling in the U.S. much higher isoprene and BVOC levels need to be considered. For example, de Gouw et al (JGR, 2005) estimated isoprene concentrations in OA source areas to be of order 10 ppb. The recent review of Heald et al. (ACP, 2008) found daytime mean total observed organic carbon concentrations of between 4 to 456 ug/m3 over North America - 3D models need to cope with this full range!

Other conditions which need to be addressed in 3D models include those associated with the free troposphere, where most concentration levels are low, but where low temperatures encourage condensation of SOA - conditions far removed from those presented in this paper.

As a final example, global models need to deal with SOA formation over for example the Amazon, where NOx concentrations are less than 100 ppt but isoprene can easily be several ppb.

Reply:

As we mentioned in our reply to previous comments, we agree with the reviewer that we should include extreme, non-typical conditions in our test of MADRID 2_FAST. While the reported U.S. ambient air concentrations of isoprene range from 1 to 21 ppb, they are generally less than 10 ppb (NTP, 2005). Two new extreme scenarios have been added to account for very high BVOC conditions: the rural condition with an isoprene mixing ratio of 10 ppb and the forest condition with an extremely low NOx mixing ratio of 100 ppt and a very high isoprene mixing ratio of 10 (ppb). In addition, we have expanded the temperature low limit from 273 K to 253 K for free tropospheric conditions in box model evaluation testing. As expected, significant speedup was obtained using MADRID2_FAST under both extreme conditions. Under the rural conditions with an isoprene concentration of 10 ppb, MADRID 2_FAST reduces the CPU time by 96.2% (from 212.23 to 7.96 seconds), with the average deviations of major species from the benchmark (using SMVGEAR solver) of -11.1% to 12.6%. Under the forest condition, MADRID 2_FAST reduces the CPU time by 95.6% (from 179.02 to 7.95 seconds), with the average deviations of major species from the benchmark (using SMVGEAR solver) of -6.8% to 7.4%. These results have been included in the revised Figures 3 and 4 and relevant text sections. These results have been presented in the revised paper, see Figure 3 and relevant discussions.

The one-day 3-D CMAQ-MADRID 2_FAST simulation covers a full range of temperatures and concentrations in troposphere. For example, the maximum mixing ratio of isoprene is up to 21 ppb over the entire U.S., and most of 1-day mean values of ISOP are around 2 ppb in the southeast, and 0.3 ppb in California. The maximum concentration of organic matter (OM) is up to 26 μg m-3, most of 1-day mean values of OC are from 3 to 8 g m-3 along the west coast and in the southeast.

p7093, line 5. It is stated that the gas-phase chemistry requires 2.1s CPU versus 199.3s for the aerosol phase. This sounds like a very ineffective aerosol-phase solver to me, so the authors might be better advised to look for very different methods rather than trying to relax stringency in search of CPU savings. UNIFAC is a rather extreme choice for application within an aerosol dynamics model, and the authors should at least discuss alternative possibilities.

Reply:

It is well known that aerosol simulation with all major physical and chemical processes accounted for is indeed much more expensive than simulating gaseous species and chemistry. Calculation of activity coefficients is one of the most expensive processes. The current approaches used in aerosol modeling range from the simplest method by assuming an ideal solution with a value of 1 for all SOA precursors (e.g., CMAQ) to simple parameterizations such as the multicomponent Wilson equation or NRTL equation, or to the most sophisticated methods such as the use of UNIFAC (as for CMAQ-MADRID). Among alternative approaches for activity coefficient calculations, UNIFAC turns out to be the most accurate method. For example, Bowman and Melton (2004) have inter-compared several models of calculating activity coefficients, and their results show that the UNIFAC predictions are in general agreement with measured partitioning data, with reasonable bases. Despite its significant computational burden, UNIFAC has better performance than other activity coefficient models for multicomponent aerosol mixtures (Bowman and Melton, 2004). In addition, UNIFAC has been widely used in aerosol modeling. It is therefore a valuable practice to improve the computational efficiency of UNIFAC, as what we have done and focused on in this study. Compared with other alternative methods for activity coefficient calculations, the accuracy of UNIFAC provides a basis for its selection for parameterization.

In addition to alternative methods for activity coefficient calculations, there exist a number of aerosol models with various degrees of complexities in their treatments of aerosol microphysics and chemistry, ranging from simple aerosol yield approach to the most sophisticated approach that is based on mechanistic partitioning theories (e.g., MADRID 2) Different methods have been used in those models in simulating gas/particle partitioning and SOA formation. These methods and models all have their own merits and limitations (see some reviews on them in Pun et al. (2002), Zhang et al. (2004), and Zhang (2008)). Depending on the specific objectives, one can choose any of them. For a SOA model with complexities similar to MADRID 2, other methods of speed up have been explored by other groups (e.g., the use of reduced gas-phase mechanism and a fast, convergent iteration technique for solving the coupled gas-aerosol system by Tulet et al. (2006)). Additional speedup methods can certainly be explored in future study.

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper

The speedup methods we explored in this study is not just trying to relax stringency, but include several additional methods, such as conditionally activating organic-inorganic interactions and parameterization of UNIFAC activity coefficient calculation, the latter approach turns out to be the most effective speedup method.

To address the reviewer’s comment, the above points have been added in our discussions, see page 15, first paragraph.

Also, if the gas-phase consumes so little CPU, why bother presenting results from QSSA and SMVGEAR? Aren’t they irrelevant? What is the CPU split between partitioning and UNIFAC, with different test cases? This is crucial information which I didn’t find.

Reply:

Evaluating gas-phase chemistry solver is not only relevant but also important in terms of both accuracy and computational efficiency. There are two main reasons to include test of various gas-phase chemistry solvers such as QSSA, SMVGEAR, and ROS3 in our study:

(1) An accurate yet computationally-efficient solver for gas-phase mechanism is very important for 3-D model simulations that cover various conditions over thousands of grid cells and a long simulation period. Such a solver will yield appreciable computational benefits for 3-D simulations that may not be so obvious from box model simulations. QSSA, ROS3 and SMVGEAR have been used widely in solving gas-phase mechanisms, and it is therefore beneficial to identify one solver that provides the best compromise between computational efficiency and numerical accuracy for 3-D simulations, as well as one solver that can be used as a benchmark for box model testing to evaluate the deviation of the speedup from the benchmark without speedup.

(2) The objective of the development of MADRID 2_FAST in this study is to speed up 3-D simulations. As the first step toward this objective, we would need to test the

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper

three gas-phase chemistry solvers to identify appropriate solvers that can be used to perform simulations with both gases and aerosols before and after speedup. QSSA is the fastest solver that has been used widely for 3-D simulations; but it is the least accurate solver. SMVGEAR lies in the other extreme, i.e., most accurate but the slowest, and has been mostly used for box model simulations. ROS3 provides a good compromise in between and has been widely used in 3-D simulations, and therefore was selected as the solver used in the gas-aerosol simulations for speedup test in this study. SMVGEAR was selected as the benchmark solver to evaluate the accuracy of each speedup method.

These reasons have now been indicated in page 9, lines 1-3 from the bottom, and page 10, lines 1-3.

The CPU cost of UNIFAC is much more expensive than that of partitioning calculation, by factors of 19.2 to 20.2 under the four test conditions. This info has also been added into the revised version, see page 11, lines 8-9 from the bottom.

The presentation of results is also rather unsatisfactory. The type of accuracy vs CPU time plots shown e.g. in Sandu et al. (1996) provide very valuable information on the behaviour of different solvers. Here the authors present just a few sentences on accuracy loss, and one figure showing average percent deviations. (And this figure has a caption which isn’t explained - what do the different symbols mean? Readers should not need to search the text to find out.)

Most statements are made of the form, "reduction in some strigency requirement by X% reduces CPU by Y%". Very little information is presented on the actual costs in terms of accuracy, except some rather vague sentences. It isn’t even clear if the authors are citing average percentage deviations, or peak errors. In many cases, the accuracy obtained sounds unacceptable to me. If a scheme gets major species wrong by 5%, presumably it gets some of the "minor" species wrong by very significant margins. These minor species are often of importance for air pollution modelling.

Other papers dealing with numerical methods tend to aim at something like 1% accuracy.

p7092, Line 22 I am puzzled by the poor performance of the QSSA and Rosenbrock schemes when they are run with most stringent error-tolerance, with 4.2% for ROS3and 5.6% for QSSA. This suggests to me a poor setup of these solvers, since they should normally perform very well when short time steps are used. (This paper doesn’t say what timesteps were used for these solvers - this is rather crucial information). For example, Sandu et al. (1996) found SDA values of 4 (0.01% accuracy!) for ROS3 at the most stringent settings. Hesstvedt et al (1978) found maximum deviations for ozone of -1.76% for a 30 second time step were reduced to 0.23% for a 5 second timestep for QSSA. Further reductions in time-step would have given even better results.

Reply:

The work done by Sandu et al. (1996) evaluates the accuracy of various gas-phase chemistry solvers using the number of accurate digits for the error (NAD) that is derived based on the modified root mean square norm of the relative error. We agree that the NAD vs. CPU time plots in his paper is very valuable for the type of study that he did. But this type of plots cannot be made for the results with various speedup methods in our study. Sandu et al. (1996) obtained the NAD vs. CPU curve by running each solver at a fixed time step (15-min or 1-hr, which is the typical transport time used for 3-D models) with a desirable NAD. The desirable NADs and associated CPU costs were obtained by tuning error tolerances within each solver in each simulation, resulting in multiple simulations as shown in each curve. So, his results provide information on the CPU cost at a the desired level of accuracy. In our case, we tested various speed-up methods for SOA simulation, which is somewhat different from Sandu et al.’s work. For the speed up of the numerical solver for partitioning of organic compounds, we tested various internal parameters such as TOLF, TOLX, and MAXIT (see Figure 1) and identify the best combination (i.e., solver 2) through evaluation of accuracy and CPU for each simulation. This approach is essentially the same as Sandu et al. (1996), although we did not make the specific type of NAD vs. CPU plots. For other speedup methods such as turning-off the inorganic-organic interactions and parameterizing activity coefficients, the box model is run only one time under the R_HBG condition, which gives one point, rather than one set of results for making a curve. So, we cannot make the type of plots in Sandu et al. (1996) for various speedup methods. We, however, made similar CPU speedup vs. percent deviation plots (see Figure 4) for 210 test simulations using MADRID 2_FAST, which provides valuable information regarding MADRID 2_FAST’s capability in speedup at a reasonable compromise of accuracy under a variety of atmospheric conditions. This evaluation approach suite well for the system we studied.

Various criteria/measures have been used to evaluate accuracy vs. speedup for atmospheric models, depending on the degree of complexities of the systems evaluated and the objectives of the speedup. The work done by Sandu et al. (1996) evaluates accuracy and CPU cost for various gas-phase chemistry solvers only, which is a relatively simple system. The accuracy evaluation was performed for gas-phase chemistry systems with 38-86 gas-phase species using NADs by ruling out species concentrations below 100 molecules cm-3. The accuracy is given in terms of NAD, but no specific criteria were provided to help judge at which NAD and CPU ratio the performance of the solver is acceptable. In Sandu et al. (1997), similar evaluation was extended to a system with coupled gas and aqueous-phase chemistry, which is more complicated than the gas-phase chemistry system. The acceptable accuracy criteria is set to be within 1% error (NAD= 2). Hesstvedt et al. (1978) evaluated the accuracy of QSSA against GEAR for gas-phase chemistry system with 60 chemical species using percentage deviations for 9 species under two test conditions. They found that in most cases the percentage deviations are within 1%, although they could be as high as 6.5 -12.3% (e.g., for NO and NO2). The coupled gas-aerosol system that we studied is much more complicated than Sandu et al. (1996, 1997a, b) and Hesstvedt et al. (1978). It simulates 361 gas-phase chemical reactions involving 191 species and includes all detailed aerosol thermodynamic and dynamic processes. Our speedup focuses on secondary organic aerosol model, which is one of the most uncertain research areas, and the biases in simulated fine organic aerosols from the current 3-D model simulations are mostly in the range of 40-70% (e.g., Zhang et al. (2004, 2006)). It is not surprising to see relatively larger errors with various speedup methods given more complexities and larger uncertainties for the system we studied, as compared with other studies. As indicated in section 3.5 and shown in Figures 3 and 4, the percentage deviations for most species are within 10%. Larger deviations (10-34%) occur for only a few species whose concentrations are very small (e.g., concentrations of NO are on the order of 10-20 ppb). Such large deviations cannot be used to judge the general computational accuracy, as the very small concentrations have no atmospheric significance. Sandu et al. (1996, 1997a, b) indeed excluded all concentrations below (100 or 1.0 molecules/cm3) in their accuracy evaluation, whereas they were not excluded from our evaluation. We believe that inclusion of all species offers a complete picture of the overall model performance, although they cannot be used to judge the general computational accuracy (i.e., at least we will get an idea how worse the deviations could be for such species).

We use a time step of 1-minute, which is among the smallest time steps used in current 3-D models. While a smaller time step can be used for box model studies to yield better results, the use of such a small time step (e.g., 5 sec) is currently not realistic for 3-D model simulations, as it is too expensive to afford. Therefore, the use of 1-minute as a time step to derive the SOA parameterization is adequate in our study. The time step for integration has been added in the revised paper.

In addition to Figure 3, we actually included Figure 4 (now Figure 5 in the revised paper) to show correlation in concentrations of several species that are most sensitive to speedup from MADRID 2_Fast vs. the benchmark from all test simulations. Unfortunately, the original Figure 4 was inadvertently left out from the APCD paper during the editing process (although it was indeed included in our submitted version), so the reviewer did not see this figure and thought that we only include one figure, which was

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper

not true.

There are no mistakes in setting up QSSA and ROS3 simulations. The deviations of 4.2% for ROS3 and 5.6% for QSSA are instantaneous values occur only for HO2 radical at one time step, the percentage deviations averaged over the simulation period for HO2 radical are 1.7% for ROS3 and 2.5% for QSSA. HO2 which is a short-lived species and may have larger error as compared to other long-lived species. The range of absolute percentage deviation for most species $<$1% for both solvers under all the test conditions. Specifically, the maximum percentage deviations for O3 are 0-0.3% for R_HBG, 0% for R_LBG, -0.8˜0.6% for U_HBG, 0-0.5% for U_LBG, 0-0.3% for F_HBG, 0-0.3% for R_EHBG.

These results are indeed quite consistent with those of Sandu et al. (1996, 1997a, b) and Hesstvedt et al. (1978).

There are several major differences in the system we examined and the approaches we used for speedup, as compared with other studies indicated by the reviewer. It is not fair to use the same accuracy criteria (1%) for very different systems and different speedup methods, although our results are indeed less than 1% for most species, consistent with those of Sandu et al. (1996, 1997a, b) and Hesstvedt et al. (1978). Given those differences, we believe that our evaluation of accuracy of MADRID 2_FAST is rigorous and comprehensive, as we evaluate all the simulated species without exclusion (191 gas species and 14 PM species) from all 210 simulations. Figures 2-5 provide a synthesis of our evaluation. Figure 2 indeed shows the actual absolute values of CPU cost. It is redundant to state them again in the text, instead, we present them in terms of percentage reduction in the text, which complements the figure presentations. Note that each of the plots in Figure 3 shows average percentage deviations for 26 major species from multiple simulations. For example, Figure 3 (a) shows 7 simulations at a fixed RH of 10% but 7 temperatures under the rural high biogenic VOC conditions, and Figure (c) shows 5 simulations at a fixed T of 253 K but 5 RHs under the same condition. So, Figures 3 (a)-(d) in fact show the results from a total of 20 simulations selected from a total of 35 simulations for each condition. Figure 4 shows the speedup factor vs. average percentage deviations for O3 and PM2.5 from all the 210 simulations, which is a similar type of plots shown in Sandu et al. (1996), as suggested by the reviewer. Figure 4 shows the deviation and correlation of four species that are highly sensitive to SOA module speedup between benchmark and the simulation with MADRID 2_FAST from a total of 210 simulations. We believe that Figures 3-5 represent the overall accuracy performance of the speed-up methods and provide a synthesis of all simulations, it is a better method than showing results from individual test simulations.

In sum, the following changes have been made to address the reviewer’s comments:

(1) We have evaluated MADRID 2_FAST against MADRID 2 using typical statistical measures including correlation coefficient, mean bias, root mean square error, normalized mean bias and error. Those statistical measures are consistent with the NAD used in Sandu et al. (1996). A table summarizing performance statistics for major species under the rural high biogenic VOCs conditions has been added, see Table 5.

(2) We have added similar sets of plots in Figure 3 under the new, non-typical rural very high biogenic VOC conditions with 10 ppb isoprene to show the model performance from additional stress tests, see Figures 3 (e)-(h).

(3) Different symbols used in Figure 3 represent simulation results under each test condition with a specific temperature (T) and relative humidity (RH), this has been indicated in the figure caption.

(4) We have added a new figure to show CPU speedup vs. percent deviation plots (see Figure 4) for 210 test simulations, which provides valuable information regarding MADRID 2_FAST’s capability in speedup at a reasonable compromise of accuracy under a variety of atmospheric conditions.

(5) We have incorporated the above points in the revised paper, see sections 3.4 and

3.5.

All in all, I was convinced that relaxing stringency reduced CPU, but I had little idea of whether the scheme was usable in real atmospheric conditions.

Reply:

Relaxing stringency of the solver is only one of methods tested in this study to speed up MADRID 2, and other methods include turning on organic-inorganic interactions when hydrophilic SOA is significant and a parameterization of activity coefficient calculations, which are different from speeding up the numerical solvers. Among the three methods, a parameterization of activity coefficients is the most efficient method for speedup. As stated previously, MADRID 2_FAST with the optimal speed-up methods developed in this study has been tested in 3-D simulations over the contiguous U.S. and gives realistic results. More detailed 3-D test results will be presented in a separate paper.

References

Bowman, F. M., and Melton, J. A.: Effect of activity coefficient models on predictions of secondary organic aerosol partitioning, J. Aero. Sci., 35, 1415-1438, 2004.

Hesstvedt, E., Hov, Ö., and Isaksen, I. S. A.: Quasi-steady-state approximations in air pollution modeling: a comparison of two numerical schemes for oxidant prediction, Int. J. Chem. Kinet., 10, 971-994, 1978.

National Toxicology Program (NTP), Public Health Service, U.S. Department of Health and Human Services: 11th Report on Carcinogens, 2005. (http://ntp.niehs.nih.gov/index.cfm?objectid=32BA9724-F1F6-975E-7FCE50709CB4C932)

Sandu, A., Potra, F. A., Carmichael, G. R., and Damian, V.: Efficient implementation of fully implicit methods for atmospheric chemical kinetics, J. Comput. Phys., 129, 101-110, 1996.

Sandu, A., Verwer, J. G., Loon, M. V., Carmichael, G. R., Potra, F. A., Dabdub, D., and Seinfeld, J. H.: Benchmarking stiff ODE solvers for atmospheric chemistry problems-I. Implicit vs explicit, Atmos. Environ., 31, 3151-3166, 1997.

Tulet, P., Grini, A., Griffin, R. J., and Petitcol, S.: ORILAM-SOA: A computationally efficient model for predicting secondary organic aerosols in three-dimensional atmospheric models, J. Geophys. Res., 111, D23208, doi:10.1029/2006JD007152, 2006.

Zhang, Y., B. Pun, K. Vijayaraghavan, S.-Y. Wu, C. Seigneur, S. Pandis, M. Jacobson, A. Nenes and J.H. Seinfeld, 2004, Development and Application of the Model of Aerosol Dynamics, Reaction, Ionization and Dissolution (MADRID), J. Geophys. Res., 109, D01202, doi:10.1029/2003JD003501.

Zhang, Y., P. Liu, A. Queen, C. Misenis, B. Pun, C. Seigneur, and S.-Y. Wu, 2006, A Comprehensive Performance Evaluation of MM5-CMAQ for Summer 1999 Southern Oxidants Study Episode, Part-II. Gas and Aerosol Predictions, Atmospheric Environment, 40, 4839-4855.

Interactive comment on Atmos. Chem. Phys. Discuss., 8, 7085, 2008.

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper