

A variable time-step algorithm for air quality models

M. Talat Odman, Yongtao Hu

School of Civil and Environmental Engineering, Georgia Institute of Technology 311 Ferst Drive, Atlanta, GA 30332-0512

ABSTRACT

In current air quality models, distinct process operators are applied sequentially to pollutant concentration fields. A common time step is used to synchronize all the processes. Usually, the characteristic time for advection, which is equal to the grid length divided by the wind speed, is selected as the common step. Since the same time step is used everywhere in the domain, the maximum wind speed and minimum grid length determine the step size. This leads to computational inefficiency in cells where process characteristic times are much longer than the time step.

A variable time-step algorithm was developed that allows each grid cell to have its own time step. Concentrations in cells with shorter time steps are updated using fluxes from cells with longer time steps. Fluxes from cells with shorter time steps to cells with longer time steps are kept in reservoirs. Concentrations in cells with longer time steps remain constant until the time levels are synchronized. At the time of synchronization the mass in each reservoir is added to the corresponding cell.

A two-dimensional implementation of the algorithm that uses the same time step in each vertical column is described. $PM_{2.5}$ estimates obtained by using variable time steps are, on average, within 3% of those obtained by using a single time step. Larger differences are observed for $PM_{2.5}$ components, especially for sulfate, which is 12% higher in winter. The differences in light extinction are also within 3% and those in ozone are within 1%. The computation time decreased by 50% in a winter episode largely due to the economy realized in aerosol equilibrium calculations. The time saved by this algorithm can be spent in increasing the process detail in air quality models or improving their computational accuracy.

Keywords:
Time-splitting
Process splitting
Particulate matter
Visibility
Ozone

Article History:
Received: 31 March 2010
Revised: 04 September 2010
Accepted: 07 September 2010

Corresponding Author:
M. Talat Odman
Tel: +1-404-894-2783
Fax: +1-404-894-8266
E-mail: talat.odman@ce.gatech.edu

© Author(s) 2010. This work is distributed under the Creative Commons Attribution 3.0 License.

doi: 10.5094/APR.2010.030

1. Introduction

Air quality models (AQMs) are based on the atmospheric transport and chemistry equation:

$$\frac{\partial c}{\partial t} + \nabla \cdot (uc) = \nabla \cdot (K\nabla c) + R(c) + S \quad (1)$$

where c is a vector of pollutant concentrations, u is the wind field, K represents parameterized atmospheric turbulence, R denotes chemical production (or loss) which is a nonlinear function of c and S includes various sources (e.g., emissions) and sinks (e.g., deposition). The dependence of the variables on the coordinates x and time t is not shown here for simplicity. Both u and K are given (usually provided by a prognostic meteorological model) so that the problem is linear with respect to the transport part. Characteristic times differ from one process to another. In particular, the range of characteristic times for chemical reactions in R spans several orders of magnitude.

After spatial discretization of Equation (1) a semi-discrete system of the following form is obtained:

$$\frac{\partial w}{\partial t} = F(w) \quad (2)$$

The variables of vector w consist of c and some other parameters and F is a vector function of w . The computational power required to solve this system is enormous due to the stiffness caused by the wide range of characteristic times. In addition, various numerical difficulties associated with special requirements of each transport and chemistry process must be dealt with. Therefore, in AQMs, Equation (2) is divided into smaller pieces. A common approach is process splitting (Blom and Verwer, 2000) where F is split into functions representing different processes:

$$\frac{\partial w}{\partial t} = F(w) \equiv F_A(w) + F_D(w) + F_R(w) \quad (3)$$

The function F_A is the advection dominated horizontal transport, F_D is the diffusion dominated vertical transport, and F_R contains chemical reaction terms. Terms for aerosol and cloud processes can also be added to the right hand side of Equation (3). Dry deposition is generally treated as a boundary condition of F_D . Emissions are usually made part of F_D or F_R . If they are made part of F_D then fresh emission plumes would be diffused before they had a chance to react with radicals in the environment. This is the choice made in the Community Multiscale Air Quality (CMAQ) model (Byun and Schere, 2006). On the other hand, if made part of F_R they would participate in chemical reactions before they had a chance to diffuse into the surrounding air. It is obvious that emissions should be made part of the process with shorter characteristic time but the relative magnitudes of the characteristic

times may differ from place to place and time to time. The Urban-to-Regional Multiscale (URM) model (Boylan et al., 2002) solves the problem by combining F_D and F_R into a single diffusion-chemistry function.

Splitting drastically reduces the computational resources required by Equation (2). It also allows using custom-built numerical solvers for each piece or process. It is much easier to deal with process-specific problems individually rather than trying to develop a general solver. For example, the advection operator is usually made nonlinear to achieve positivity either through filtering or flux-limiting. The only disadvantage of splitting is that it introduces an error into the solution (Lanser and Verwer, 1999).

Splitting methods are classified as first or second order based on the order of the splitting time step, Δt , in their error terms. Most splitting methods, including the one used in CMAQ, are first order (i.e., their splitting error term is second order in Δt). They advance the solution by Δt in time by applying process operators consecutively as follows:

$$w^{n+1} = \Phi_R(t^n; \Delta t) \Phi_D(t^n; \Delta t) \Phi_A(t^n; \Delta t) w^n \tag{4}$$

where Φ_A , Φ_D and Φ_R are the integrators for F_A , F_D and, F_R in Equation (3) respectively; for simplicity other processes are not shown. Strang splitting (Strang, 1968), which is believed to be second order, is also popular in AQMs. It advances the solution in time by the following sequence of operators:

$$w^{n+1} = \Phi_A(t^{n+\frac{1}{2}}; \frac{\Delta t}{2}) \Phi_D(t^{n+\frac{1}{2}}; \frac{\Delta t}{2}) \Phi_R(t^n; \Delta t) \Phi_D(t^n; \frac{\Delta t}{2}) \Phi_A(t^n; \frac{\Delta t}{2}) w^n \tag{5}$$

Note that the transport operators Φ_A and Φ_D are applied for one half of Δt , symmetrically around the chemistry operator Φ_R . Sportisse (2000) recently argued that, for R linear in c , Strang splitting is only first order unless the stiff operator Φ_R is applied last.

Processes with characteristic times shorter than the splitting time step can be advanced in several sub-time steps but this “sub-cycling” is internal to the process operators. It is the splitting time step that determines the frequency by which different processes with differing characteristic times are “synchronized” or “coupled.” Typically, the characteristic time for advection, which is equal to the grid size divided by the wind speed, is selected as the splitting time step (splitting will be dropped hereafter). Current AQMs use a single time step for the entire domain. In uniform grid AQMs the maximum wind speed in the domain limits the time step. In variable grid AQMs the characteristic time of a cell with relatively small grid size and high wind speed would be used as the global time step, Δt . Usually, a large fraction of the grid cells would have characteristic advection times larger than Δt due to either relatively low local wind speeds or larger grid sizes. Time stepping those cells with a Δt much smaller than their characteristic time step does not make the results more accurate; therefore, using a single time step for the entire domain is computationally inefficient.

A variable time-step algorithm is described in this paper. Its accuracy and computational efficiency were evaluated in comparison to using a single time step in a uniform grid AQM. Computing time saved by this algorithm can be used for increasing the details of processes modeled in AQMs, adding new processes, or improving computational accuracy. For example, AQMs can be coupled with the dynamics models and the feedbacks of aerosols on radiation and clouds can be modeled. Alternatively, the grid resolution can be increased for more accurate representation of emission plumes or cumulus convection. Some AQMs confronted the grid resolution issue by using variable grids that are either

static (Boylan et al., 2002; Park et al., 2004) or dynamic (Odman et al., 2002; Constantinescu et al., 2008). If the grid sizes in a variable grid differ by orders of magnitude, and, worse yet, smallest grid cells are collocated with strongest winds, the inefficiency of using a single global time step becomes unbearable even for the most powerful computers. This variable time-step algorithm has been the enabling technology for the adaptive grid version of CMAQ, which uses a dynamic variable grid (Garcia-Menendez et al., 2010).

2. Methodology

The variable time-step algorithm (VARTSTEP) is developed with the objective of enabling local time steps in AQMs. VARTSTEP allows each grid cell to advance by its own time step. The version of the algorithm described here is two-dimensional: it uses the same time step for an entire vertical column. Extension of the algorithm to third dimension should be straightforward. Here, VARTSTEP is implemented in the CMAQ model. CMAQ uses a global time step, Δt , determined as follows:

$$\Delta t = s \frac{\Delta x}{u^{max}} \tag{6}$$

In Equation (6), Δx is the uniform grid size, u^{max} is the maximum wind speed in the domain, and s is a factor of safety equal to 0.75. Starting with Version 4.3, the user can specify an altitude (as a vertical layer) above which wind speeds are not considered in determining the time step. Horizontal advection is applied multiple times (i.e., sub-cycled) above that altitude with a sub-time step that guarantees stability. This feature is preserved in the VARTSTEP version of CMAQ, which will be referred to as VARTSTEP-CMAQ hereafter.

The two-dimensional (2D) VARTSTEP assigns every vertical column i its own time step, Δt_i , called the local time step hereafter to differentiate from the global time step. The local time step must satisfy two conditions. The first condition is:

$$\frac{u_i^{max} \Delta t_i}{\Delta x_i} \leq 1 \tag{7}$$

where u_i^{max} is the maximum wind speed in vertical column i (up to the user-specified altitude) and Δx_i is the horizontal grid size for that vertical column (subscript i is used considering that grid size may be non-uniform). This is the well known Courant, Friedrichs and Lewy (CFL) stability condition for advection (Anderson et al., 1984). In process splitting, this condition also assures that fresh emissions are not transported by a horizontal distance longer than the grid length before other processes are applied at least once. The second condition requires that Δt_i be an integer multiple of the global time step, Δt , and an integer divisor of the output time step. The last part assures synchronization of processes among all the grid cells before the results are written out. Note that the global time step must be determined first, before the second condition can be used to determine the local time steps. The following example illustrates how the second condition is applied. Assuming an output time step of 60 min, if the global time step is 5 min, the local time steps can be 5, 10, 15, 20, 30, or 60 min. On the other hand, 25, 35, 40, 45, 50, or 55 min cannot be used as local time steps, although they are multiples of 5 min, because they do not divide 60 min evenly.

The model clock advances in increments equal to the global time step, Δt . When the clock strikes an integer multiple of the local time step, Δt_i , (i.e., $t = N \times \Delta t_i$) VARTSTEP applies process operators to the concentrations of vertical column i for the duration of Δt_i . In other words, instead of advancing the solution in m steps with the global time step Δt as:

$$\begin{aligned}
 w^{n+m} &= \Phi_R(t^{n+m-1}; \Delta t) \Phi_A(t^{n+m-1}; \Delta t) \Phi_D(t^{n+m-1}; \Delta t) \dots \\
 &\Phi_R(t^{n+k}; \Delta t) \Phi_A(t^{n+k}; \Delta t) \Phi_D(t^{n+k}; \Delta t) \dots \\
 &\Phi_R(t^n; \Delta t) \Phi_A(t^n; \Delta t) \Phi_D(t^n; \Delta t) w^n
 \end{aligned}
 \tag{8}$$

where $1 \leq k \leq m-2$, VARTSTEP advances it in one step with the local time step $\Delta t_i = m \times \Delta t$ as:

$$w^{n+m} = \Phi_R(t^n; m\Delta t) \Phi_A(t^n; m\Delta t) \Phi_D(t^n; m\Delta t) w^n
 \tag{9}$$

Since there is a single local time step for each vertical column, there are no difficulties involved in doing this with the vertical transport operator, Φ_D . The chemical reaction operator, Φ_R , does not pose any problems either since it is applied to one grid cell at a time. However, the horizontal transport operator, Φ_A , requires special attention since neighboring grid cells in the horizontal may have different time steps. Figure 1 illustrates how the transport between two cells might be handled. If $\Delta t_2 < \Delta t_1$, any horizontal flux (advective or diffusive) from Cell 1 must be passed to Cell 2 when it is time to update Cell 2 concentrations (i.e., when $t = N \times \Delta t_2$). This situation is illustrated by the arrow marked "Pass" in Figure 1. On the other hand, if $\Delta t_2 > \Delta t_1$ and $t \neq N \times \Delta t_2$ any flux from Cell 1 to Cell 2 is directed to a reservoir. This is shown in Figure 1 by the arrow marked "Store". When the time comes for updating Cell 2 concentrations (i.e., when $t = N \times \Delta t_2$) not only any possible flux from Cell 1 is passed to Cell 2, but the mass accumulated in the reservoir is also added to Cell 2. This is represented by the arrow marked "Flush" in Figure 1.

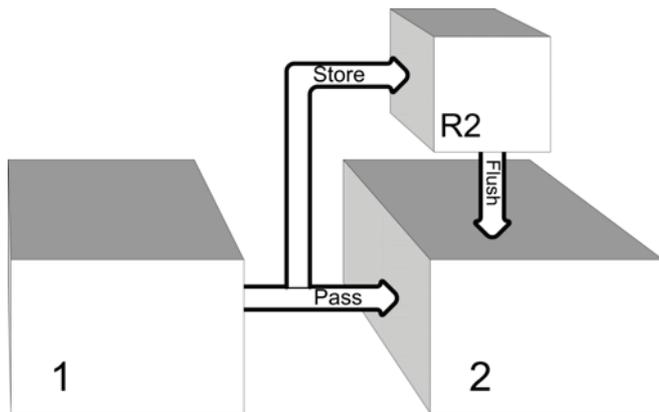


Figure 1. Concept of a reservoir used in VARTSTEP algorithm. The flux from Cell 1 is either passed directly to Cell 2 or stored in a reservoir which is flushed when it is time to update Cell 2 concentrations.

In practice, there is little computational gain in trying to calculate horizontal fluxes for the duration of each cell's local time step. There are also mass conservation concerns related to calculating the same flux at a grid cell interface twice, using one time step at the outflow side and another time step at the inflow side. It is much easier to calculate horizontal transport fluxes every global time step. In most transport schemes used in current AQMs, the calculation of a flux into or out of a cell does not only involve that cell's concentration but the concentrations of neighboring cells as well. Those concentrations may have been updated at a different time than the cell's own concentration. Horizontal transport fluxes are calculated using the latest updated concentration for each cell. The influxes are added to and outfluxes are subtracted from the cells' reservoirs. When the concentration of a cell is being updated, every local time step, the cell's reservoir is also flushed. This assures that all fluxes are eventually added to the appropriate cell and that mass is conserved. The process that advances a grid cell's concentration in time by the cell's local time step is illustrated in Figure 2.

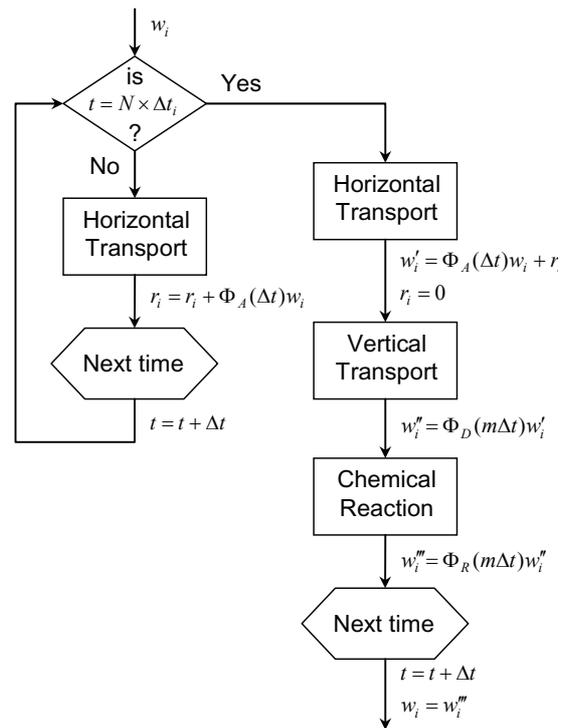


Figure 2. Flowchart of the VARTSTEP algorithm showing the sequence of the processes for advancing w_i , the solution for grid cell i , by one local time step $\Delta t_i = m\Delta t$.

3. Major Implementation Issues

In this section, the two-dimensional (2D) implementation of VARTSTEP in CMAQ is described. The discussion is limited to general issues and is intended to be useful for implementation in other air quality models. Remarks are made to aid a three-dimensional (3D) implementation in the future.

The first order splitting of Equation (4) is implemented in CMAQ as consecutive calls to various process modules inside a time step loop. Each module calculates the changes in pollutant concentrations resulting from a particular process (or processes) during the time step and applies them to the concentration array, CGRID. The specific order in which the process modules are called is as follows: vertical diffusion, which includes emissions and dry deposition; horizontal advection; vertical advection; horizontal diffusion; cloud processes; gas-phase chemistry; and aerosol processes. The CGRID array is four-dimensional consisting of 3D concentration fields for each and every pollutant species. The time step is set to the minimum characteristic time for advection among all grid cells.

In VARTSTEP-CMAQ, the time step loop advances the clock time by the global time step, TSTEP. In this 2D implementation of VARTSTEP, the local time steps assigned to each vertical column are stored in LOCSTEP, a 2D array with the dimensions of the horizontal grid. In a 3D implementation every grid cell would be assigned its own time step, hence, LOCSTEP would be a 3D array. LOCSTEP passes the local time step information to the process modules. Processes other than horizontal advection and diffusion are applied to vertical columns for the duration of the local time step. The decision to update pollutant concentrations is made based on whether the local time step evenly divides the clock time. The potential for computational time savings lies in processes that are applied less frequently and for periods longer than the global time step. In VARTSTEP-CMAQ, this is the case for vertical diffusion, cloud processes, gas-phase chemistry, and aerosol processes.

The local time steps of neighboring cells in a horizontal plane can be different; therefore, the implementation of VARTSTEP for horizontal transport is more complex compared to vertical transport where all the cells in a vertical column share the same local time step. In a 3D implementation, vertical transport, including transport due to convective clouds, would have to be treated in a way similar to the way horizontal transport is treated here. In VARTSTEP-CMAQ, the horizontal fluxes in and out of grid cells must be computed every global time step even though the cell concentrations are updated once per local time step. Calculation of fluxes that will not be used immediately may be unnecessary in other models. However, in CMAQ, flux parameters, namely wind velocity for advective flux and eddy diffusivity for diffusive flux, are evaluated at the middle of the time step through linear interpolation between two meteorological input records, typically one hour apart. Evaluating flux parameters at different times may lead to mass conservation problems. For example, suppose the wind velocity at the interface of two cells has one value when evaluated at the middle of the local time step of the first cell but another value when evaluated at the middle of the local time step of the second cell. This would yield a different flux leaving (or entering) the first cell than the flux entering (or leaving) the second cell. Having different fluxes on opposite sides of the cell interfaces would violate mass conservation. If the flux parameters were held constant during each output time step, horizontal transport could also be applied for the duration of the local time step.

VARTSTEP requires another array to store the fluxes into or out of cells whose local time steps are longer than the global time step and whose concentrations are not being updated immediately. This array, FGRID, fulfills the duty of the reservoirs described in Section 2. In CMAQ, horizontal transport processes (advection and diffusion) are applied to all horizontal planes, from the surface to the top, through a loop over vertical layers before CGRID is updated and handed over to another process. To be consistent with this structure, and, considering the possibility of a 3D implementation in the future, we chose to store the fluxes of each pollutant in a three-dimensional array. Further, since the fluxes of each pollutant must be kept in a different reservoir, we made the FGRID array four-dimensional, where the fourth dimension is indexed by pollutant species, just like the CGRID array. The FGRID array is the only significant burden of VARTSTEP on memory; therefore, if additional memory slightly larger than the size of the CGRID array is made available, full advantage can be taken from the algorithm's computational efficiency. Note that if the loops over vertical layers are taken out of horizontal process modules and these modules are called within a layers loop, then only the horizontal fluxes of one layer would have to be stored in a 2D array for each pollutant; therefore, FGRID can be a 3D array.

For mass conservation, horizontal transport schemes are usually in flux form; therefore, the calculation of horizontal fluxes is routine in most AQMs. In CMAQ, the preferred advection scheme is the Piecewise Parabolic Method of Collella and Woodward (1984). In this scheme, advective fluxes are calculated by fitting piecewise parabolas to the cell concentrations. Horizontal advection modules of CMAQ were modified not to update cell concentrations but to store already calculated lateral fluxes in the FGRID array. In CMAQ, advective outfluxes are limited such that they do not generate negative concentrations. No flux is allowed to remove more than the pollutant mass content of a cell. In VARTSTEP-CMAQ some of the cell content may already be committed to other outfluxes (e.g., advective flux in the orthogonal direction, advective flux in the same direction during a previous global time step, or diffusive flux) or influxes may have accumulated in the reservoir. Therefore, the outflux must be limited by the cell content plus the fluxes (positive or negative) stored in FGRID up to that point. The horizontal diffusion module was also modified to store the fluxes calculated by central differencing without updating the concentrations. In addition, flux

limiting was added to this module to take into account fluxes that are already committed.

When the local time step evenly divides the clock time, the cell concentrations in CGRID must be updated using the horizontal fluxes stored in FGRID. This can be done after the calls to horizontal transport modules. In VARTSTEP-CMAQ, instead of creating a separate module for this, the task was assigned to the vertical advection module, which already contains an update of CGRID based on LOCSTEP. This results in computational time savings by reducing the total number of conditional operations. However, the original process order had to be changed such that horizontal diffusion takes place before vertical advection. It should also be noted that VARTSTEP-CMAQ uses a vertical advection scheme that adjusts vertical velocities for strict mass conservation (Hu et al., 2006).

Vertical diffusion, which includes emissions and dry deposition, is the first process in CMAQ and it is followed by horizontal transport. The reason for this is most likely the desire to start each process cycle with fresh emissions and give diffusion priority because of its relatively short characteristic times. In VARTSTEP-CMAQ, if vertical diffusion were applied when the local time step evenly divides the clock time (i.e., $t = N \times \Delta t_i$), in cells with local time steps longer than the global time step, horizontal transport would start before emissions were injected and vertically diffused. On the other hand, there is no reason for applying vertical diffusion every global time step just to remain faithful to the original process order. Therefore, vertical diffusion was applied to each vertical column for the duration of the local time step at the beginning of the process cycle. Note that gas-phase chemistry, aerosol and cloud processes are applied when the local time step evenly divides the clock time, effectively taking place after horizontal transport. This would not be an issue in models where vertical diffusion and emissions are applied after horizontal transport.

The last important implementation issue is the time interpolation of meteorological parameters and emissions. In CMAQ, these parameters are evaluated at the middle of the global time step. The reading of input values and interpolation takes place for all the grid cells at once. Using this global interpolator may have some undesirable consequences. For example, assuming a 10 min global time step and hourly meteorological inputs, between 07:00 and 08:00, the interpolator would interpolate the parameters to 07:05, 07:15, 07:25, and so on. For a concentration update at 7:30, a cell with a local time step of 30 min should have its parameters at 07:15, which is the middle of the time step from 7:00 to 7:30, but the last available parameters from the global interpolator would be at 7:25. Having meteorological parameters and emission rates at a different time can violate mass conservation. For example, the total amount of emissions into cells with longer local time steps may be over/under-estimated. To avoid such consequences, a new interpolator was developed here for VARTSTEP-CMAQ that can evaluate time-dependent parameters for each grid cell at a different time. This way all the cell parameters can be evaluated at the middle of the local time step of that cell.

4. Results and Discussion

VARTSTEP was verified using a rotating-cone test. A right circular cone was introduced into a rotational wind field on a uniform grid (see the Supporting Material, SM, Figure S1a). The distance from the axis of rotation to the axis of the cone was equal to the diameter of the base circle; therefore, wind speeds increased by a factor of 3 across the base circle. Hence, the local time steps in the grid cells supporting the cone varied by a factor of 3 according to Equation (7). After one full rotation, the cone rotated by VARTSTEP was practically identical to the cone rotated by using a global time step (see the SM, Figures S1b and S1c). The small magnitude of the difference (see the SM, Figure S1d)

indicated that VARTSTEP is a viable alternative to the current time stepping algorithms in AQMs. Mass conservation was also tested and no mass conservation errors were found after one full rotation of the cone (see the SM, Figure S2).

To evaluate VARTSTEP-CMAQ, air quality simulations were performed in the Southeastern U.S. during a winter episode (1–20 January 2002) and a summer episode (12–27 July 2001). Version 4.3 of CMAQ was used for calculating the ambient fine particulate matter ($PM_{2.5}$) and ozone (O_3) concentrations. The horizontal grid resolution was 12 km over the domain shown in Figure 3. The global time steps were between the user-specified minimum of 5 min and the CFL condition imposed maximum of 7.5 min during the winter episode and 10 min during the summer episode. In addition to these simulations that served as a “benchmark”, a second set of simulations was conducted with VARTSTEP-CMAQ.

All computations were performed on a single processor 2.8 GHz Intel CPU with dedicated 4 gigabytes of memory. The total computation time for the simulation with VARTSTEP-CMAQ was 50% less than the benchmark CMAQ simulation for the winter episode, while the savings were only 13% for the summer episode (Table 1). This difference in savings is largely due to the load reduction of the aerosol module from 60% for the winter episode to 23% for the summer episode in CMAQ simulations. In Southeastern U.S., sulfate is the dominant PM species in summer, while nitrate is almost negligible. On the other hand, nitrate levels can exceed sulfate levels in winter. Nitrate equilibrium is much more complex than sulfate's: it does not only depend on the availability of nitric acid but also on ammonia and sulfate, as well as temperature and relative humidity. The large CPU time spent in the aerosol module in winter is due to the computational intensity of nitrate equilibrium, as well as shorter time steps dictated by stronger winds. With VARTSTEP, the times spent in the vertical diffusion, clouds, and chemistry modules were all reduced, but almost all of the benefits are realized in the aerosol module. The gas-phase chemistry module of CMAQ uses an implicit chemical kinetics solver whose computation time does not respond to changes in the time step. A different solver with larger overhead could have benefited more from less frequent calls to the chemistry module.

Since horizontal transport modules are called every global time step, no benefit of VARTSTEP should be expected. The increase in the computation time of the horizontal advection module is due to the flux limiting that accounts for committed fluxes in FGRID. The increase in the computation time of the horizontal diffusion module is due to the computation and storage of diffusive fluxes needed by VARTSTEP; those fluxes were not

needed in the original module. The VARTSTEP burden in vertical advection is due to the update of concentrations.

Now that we have seen how VARTSTEP can reduce the computation time, the question is whether it can produce reliable air quality results. Daily average $PM_{2.5}$ concentrations for 17 January 2002 resulting from the two simulations are compared in Figure 3. Both simulations produced a similar $PM_{2.5}$ distribution over the Southeastern U.S. with a peak near Pensacola, Florida. The value of the peak is $65.8 \mu g m^{-3}$ according to the benchmark simulation and $62.9 \mu g m^{-3}$ according to the VARTSTEP-CMAQ simulation. $PM_{2.5}$ distributions were also similar for the summer episode, VARTSTEP producing slightly lower values over the Southeast (see the SM, Figure S3).

Daily average observations of $PM_{2.5}$ and its composition are available through the Interagency Monitoring of Protected Visual Environments (IMPROVE) network. The network covers areas designated as “Class-I” for visibility protection such as national parks and wilderness areas. There are 22 such observation sites in the region shown in Figure 3. Eight of these sites are clustered along the Southern Appalachian Mountains, six sites are distributed along the coastline, and the remaining eight sites are scattered throughout the domain. The objective here is not to evaluate modeled concentrations vis-à-vis observations, but to perform a detailed comparison of the VARTSTEP-CMAQ results with benchmark CMAQ results at these 22 sites where performance is of utmost concern. Figure 4 compares the two sets of model results for daily average $PM_{2.5}$ concentrations at these 22 sites during the 1–20 January 2002 and 12–27 July 2001 periods. Note that all comparisons provided here are for the model layer directly above the ground.

There is strong correlation $R^2 > 0.98$ between the benchmark and the results obtained with VARTSTEP-CMAQ with the exception of a few outliers. VARTSTEP-CMAQ results are higher than the benchmark, on average, by about 1 to 3%. Since visibility is the primary concern in this study and since different components of $PM_{2.5}$ affect visibility differently, it is important to accurately model not just the total $PM_{2.5}$ but its components as well. Sulfate and nitrate particles, in the presence of water vapor, and elemental carbon can be important visibility degradation agents; other particles are relatively less important (Malm et al., 2000). Daily average concentrations of sulfate, nitrate, ammonium, soil, elemental carbon, and organic carbon components of $PM_{2.5}$ were compared at the 22 Class-I areas. The scatter plots in Figure 5 (1–20 January 2002) and 6 (12–27 July 2001) have the concentrations calculated by VARTSTEP-CMAQ on the y -axis and those calculated by CMAQ Version 4.3 on the x -axis.

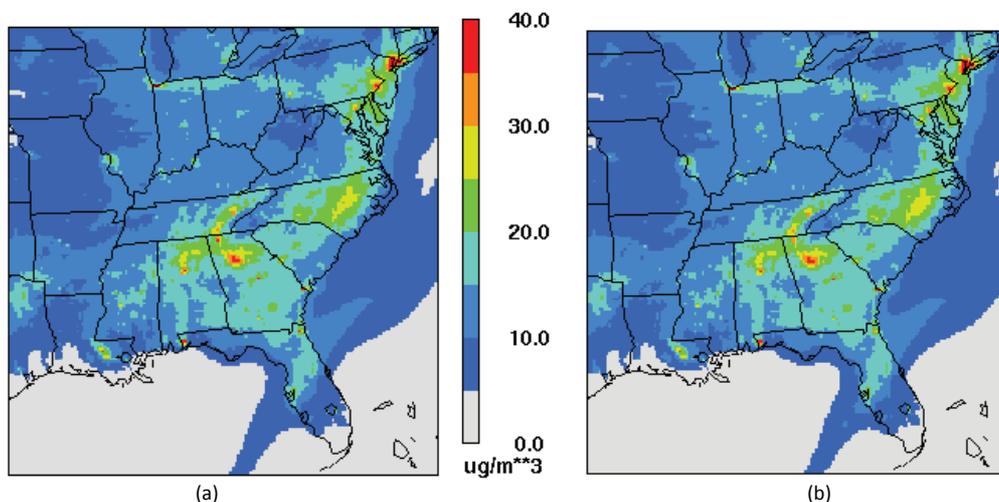


Figure 3. Daily average $PM_{2.5}$ concentrations ($\mu g m^{-3}$) over the Southeastern U.S. on 17 January 2002 as calculated by the benchmark CMAQ simulation (a) and the VARTSTEP-CMAQ simulation (b).

Table 1. CPU^a times per simulated day broken down by process for CMAQ (Version 4.3) and VARTSTEP–CMAQ, and the fraction of the total VARTSTEP savings attributed to each process

Module	CMAQ (s)		VARTSTEP–CMAQ (s)		Fraction of Savings (or Burden) (%)	
	Winter episode	Summer episode	Winter episode	Summer episode	Winter episode	Summer episode
Vertical Diffusion	4 148	3 226	1 850	1 398	7.51	52.5
Horizontal Advection	7 027	4 327	8 762	5 164	(5.68)	(24.0)
Horizontal Diffusion	481	391	1 517	1 159	(3.39)	(22.1)
Vertical Advection	1 027	733	2 126	1 426	(3.59)	(19.9)
Clouds	555	503	529	476	0.09	0.76
Chemistry	8 824	9 770	8 061	8 729	2.50	29.9
Aerosols	37 098	6 099	5 469	2 974	103	89.8
Other	2 470	1 663	2 726	1 906	(0.8)	(6.95)
TOTAL	61 630	26 712	31 040	23 232	100	100

^a Single processor 2.8 GHz Intel CPU with dedicated 4 GB memory

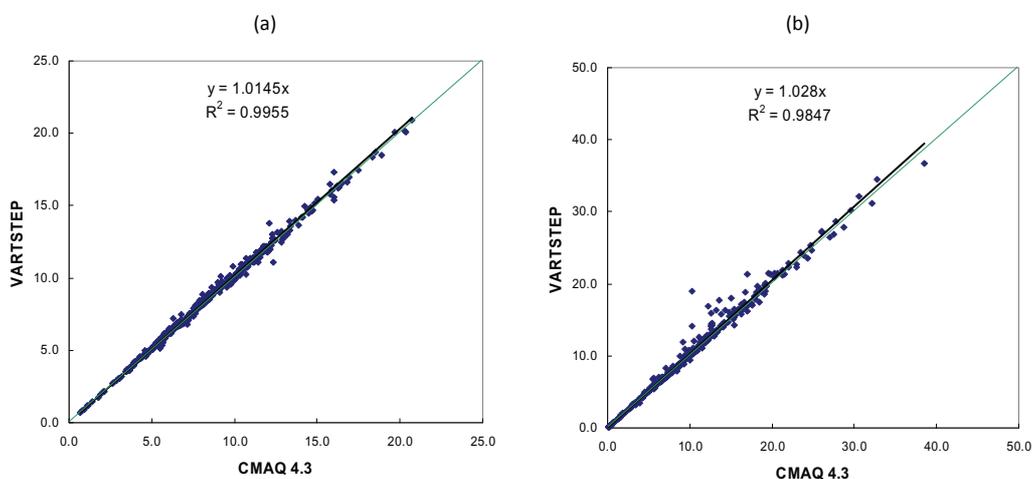


Figure 4. Comparison of daily average $PM_{2.5}$ concentrations ($\mu\text{g m}^{-3}$) calculated by VARTSTEP–CMAQ (y -axis) with the benchmark (x -axis) at 22 Class-I areas in the Southeastern U.S. during 1–20 January 2002 (a) and 12–27 July 2001 (b).

The two model results are highly correlated: R^2 is at least 0.94 (in the case of wintertime sulfate) and 0.98 or higher in the case of nitrate, ammonium soils, elemental carbon, and organic carbon for both episodes. For the winter episode, while sulfate concentrations calculated by VARTSTEP–CMAQ are, on average, 12% higher, the concentrations of nitrate are 5.0% lower than the benchmark (Figure 5). Sulfate is a smaller component of $PM_{2.5}$ than nitrate in winter when relative humidity is also low; therefore, compared to summertime, both sulfate and nitrate have a much smaller impact on visibility. Soil, the second largest component of $PM_{2.5}$, has no bias and ammonium is higher by 2.7%. Elemental carbon, an important contributor to light extinction, has no bias while organic carbon, which is the fourth largest component of $PM_{2.5}$, is 1.1% higher.

For the summer episode, the concentrations of all the $PM_{2.5}$ components calculated by VARTSTEP–CMAQ are highly correlated with the benchmark concentrations (Figure 6). Sulfate is the largest component of $PM_{2.5}$ during this period and, with high summertime relative humidity, it has the largest contribution to light extinction. Its concentrations are, on average, 3.6% higher. Organic carbon, which is the second largest component of $PM_{2.5}$, is higher by 4.6%. Ammonium and soil have almost no bias. Nitrate, the fifth largest summertime $PM_{2.5}$ component, is 5.5% lower. Elemental carbon, the smallest component, but an important visibility degradation agent, is lower by 3.5%.

To see the total impact of $PM_{2.5}$ components on visibility, light extinction, B_{ext} (Mm^{-1}), was calculated using the formula in Malm et al. (2000):

$$B_{\text{ext}} = 3f(RH)\left([SO_4^{2-}] + [NO_3^-] + [NH_4^+]\right) + 4[OC] + 10[EC] + [Soil] + 0.6[CM] + B_{\text{Rayleigh}} \quad (10)$$

where $[SO_4^{2-}]$, $[NO_3^-]$, $[NH_4^+]$, $[OC]$, $[EC]$, and $[Soil]$ are the concentrations ($\mu\text{g m}^{-3}$) of $PM_{2.5}$ associated with sulfate, nitrate, ammonium, organic carbon, elemental carbon, and soils, respectively; $[CM]$ is the concentration of coarse PM; $f(RH)$ is a dimensionless relative humidity adjustment factor, which is different for each Class-I area and varies by the time of the year; and B_{Rayleigh} is the Rayleigh scattering (10 Mm^{-1}). Daily average extinction coefficients derived from CMAQ and VARTSTEP–CMAQ concentrations were compared at the 22 Class-I areas (Figure 7). R^2 is 0.99 for both the 1–20 January 2002 and the 12–27 July 2001 episodes. On average, VARTSTEP–CMAQ light extinction is 1.6% higher than the benchmark during the winter episode and 2.5% higher during the summer episode.

Daily maximum 8-hr ozone concentrations were also compared at the 22 Class-I areas (Figure 8) to see the impact of oxidant concentrations on the $PM_{2.5}$ results above. The correlation between the benchmark and the results obtained with VARTSTEP–

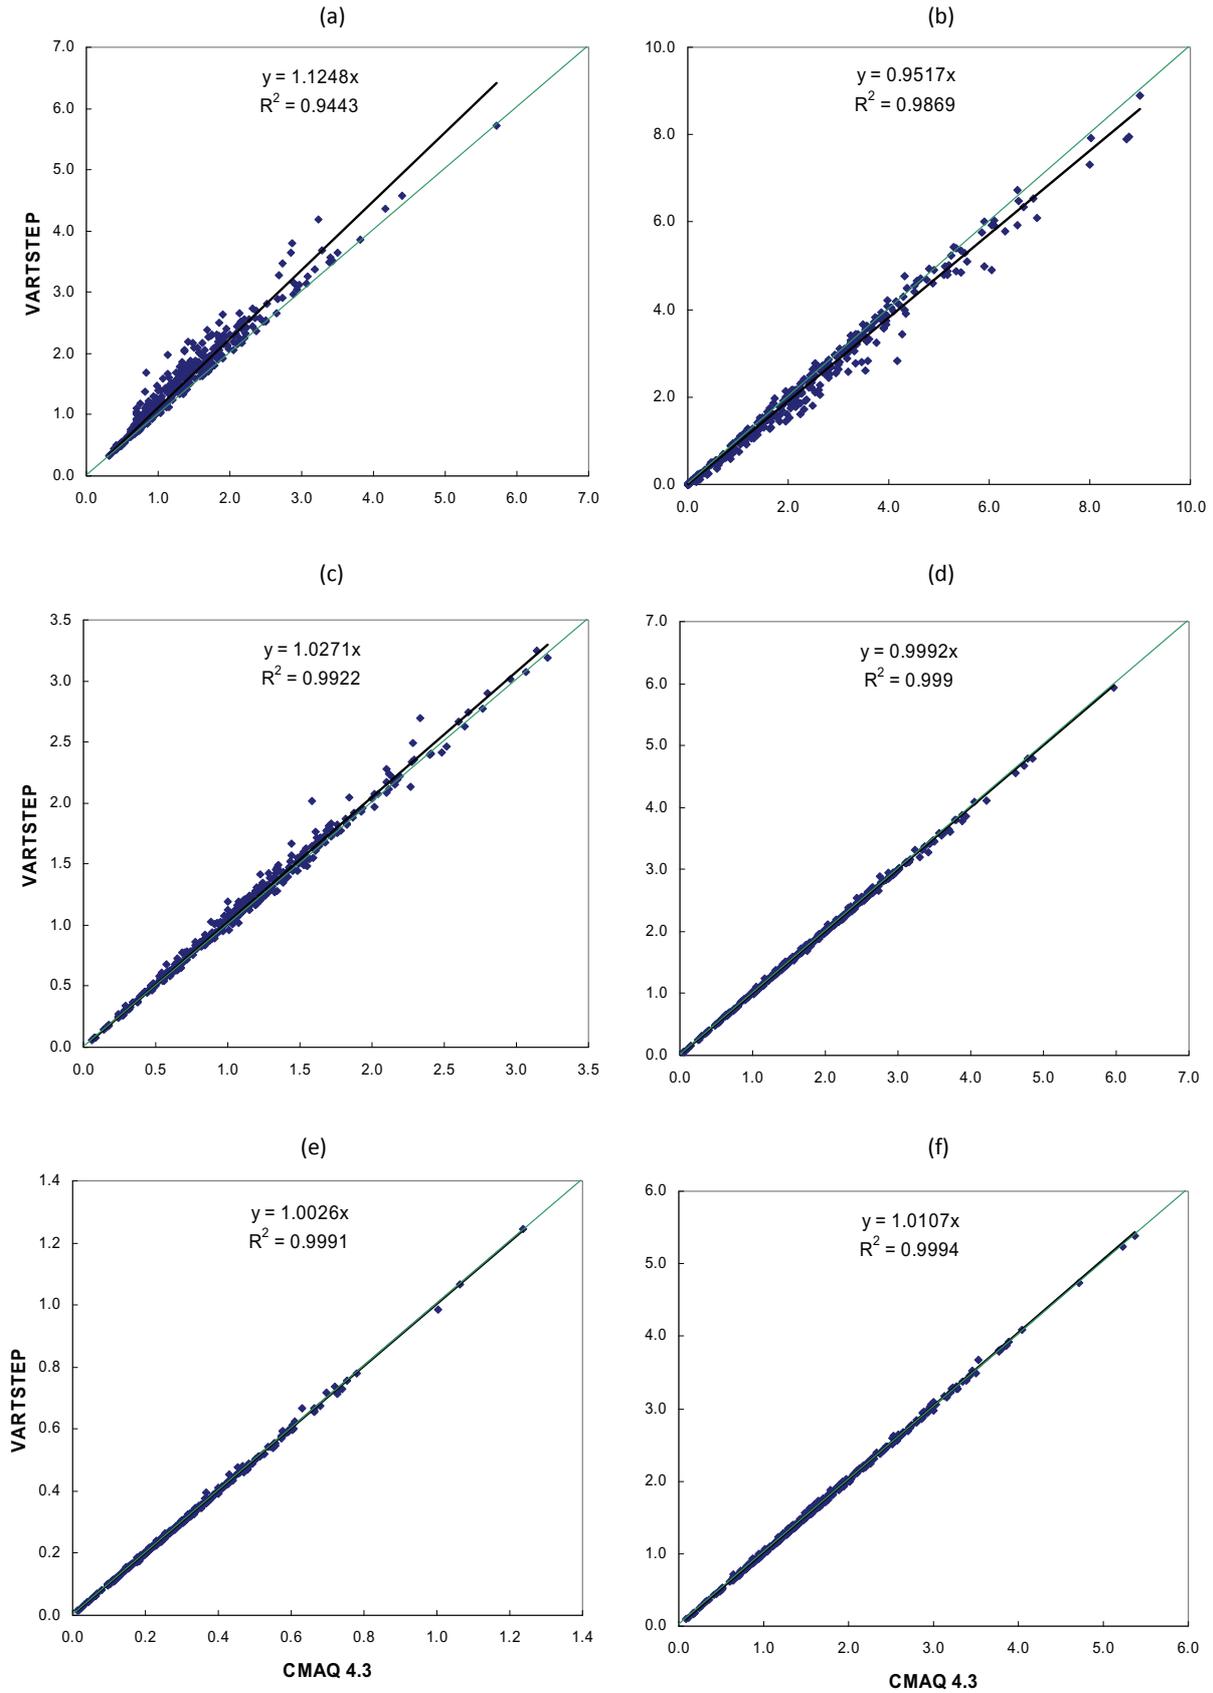


Figure 5. Comparison of daily average concentrations ($\mu\text{g m}^{-3}$) calculated by VARTSTEP–CMAQ (y-axis) with the benchmark (x-axis) for various components of $\text{PM}_{2.5}$ at 22 Class-I areas in the Southeastern U.S. during 1–20 January 2002: (a) Sulfate, (b) Nitrate, (c) Ammonium, (d) Soil, (e) Elemental Carbon, (f) Organic Carbon.

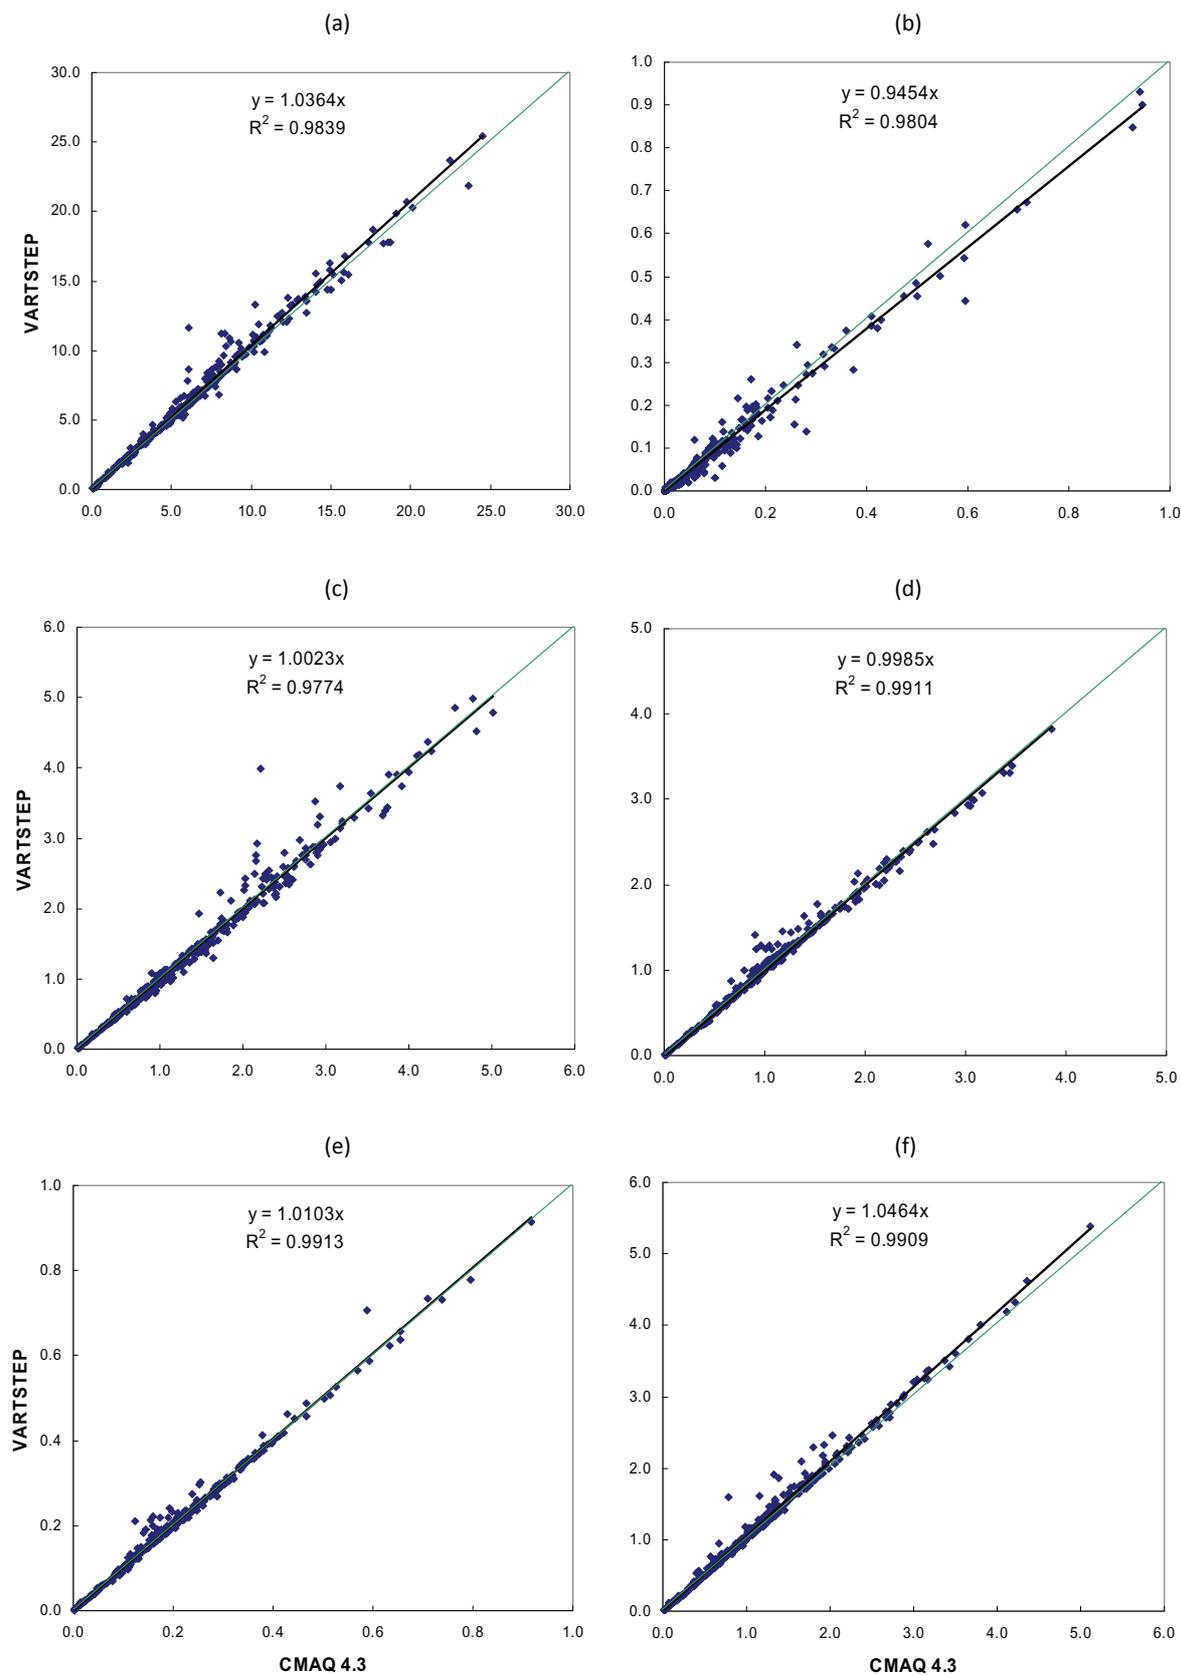


Figure 6. Comparison of daily average concentrations ($\mu\text{g m}^{-3}$) calculated by VARTSTEP-CMAQ (y-axis) with the benchmark (x-axis) for various components of $\text{PM}_{2.5}$ at 22 Class-I areas in the Southeastern U.S. during 12–27 July 2001: (a) Sulfate, (b) Nitrate, (c) Ammonium, (d) Soil, (e) Elemental Carbon, (f) Organic Carbon.

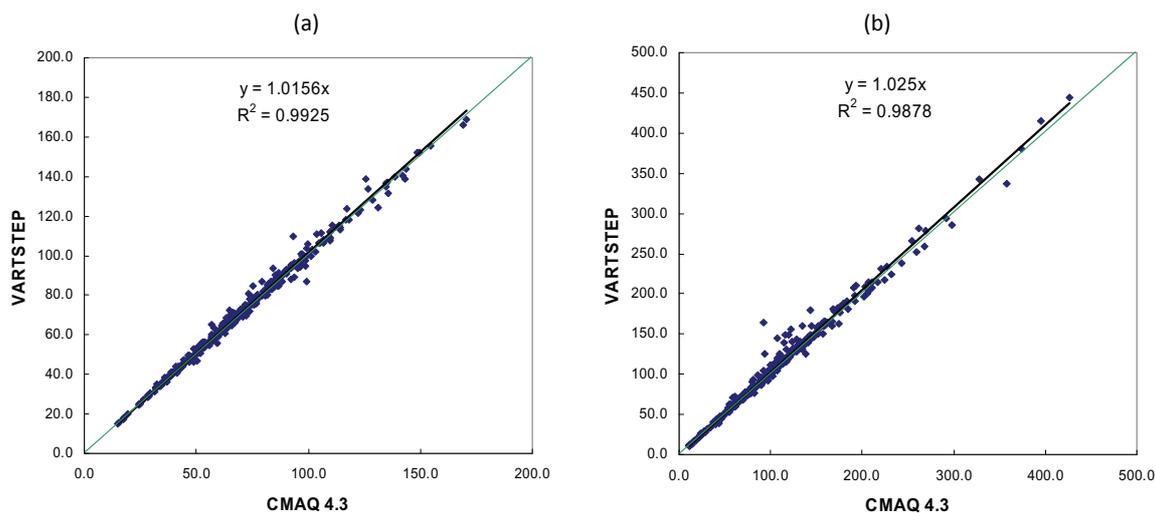


Figure 7. Comparison of daily average light extinction (Mm^{-1}) derived from VARTSTEP-CMAQ (y -axis) with the benchmark (x -axis) at 22 Class-I areas in the Southeastern U.S. during 1–20 January 2002 (a) and 12–27 July 2001 (b).

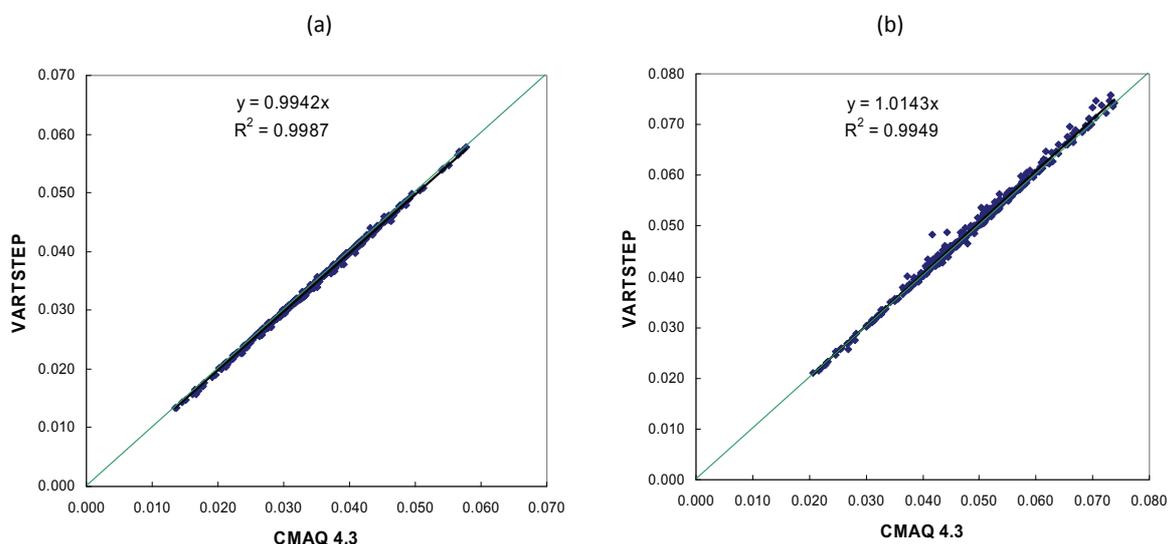


Figure 8. Comparison of daily maximum 8-hr O_3 concentrations (ppm) calculated by VARTSTEP-CMAQ (y -axis) with the benchmark (x -axis) at 22 Class-I areas in the Southeastern U.S. during 1–20 January 2002 (a) and 12–27 July 2001 (b).

CMAQ is strong for both the summer ($R^2 = 1.00$) and the winter episodes ($R^2 = 0.99$). On average, the VARTSTEP-CMAQ ozone concentrations are 1% lower in winter and 1% higher in summer compared to the benchmark.

The higher sulfate and lower nitrate concentrations calculated by VARTSTEP-CMAQ can be explained as follows. With VARTSTEP, the time steps are longer and different processes are coupled less frequently. Therefore, turbulent mixing and dry deposition, which are both performed in the vertical diffusion module, can operate longer on freshly emitted pollutants (emissions are also part of vertical diffusion) before those pollutants have a chance to take part in photochemical reactions in the chemistry module or change phase through condensation (or evaporation) in the aerosol module. The majority of SO_2 emissions are from elevated sources while almost all of the NH_3 emissions are from ground-level sources. In CMAQ (and VARTSTEP-CMAQ), aerosol equilibrium is such that free ammonia first neutralizes sulfate before any ammonium nitrate is formed. Enhanced mixing in VARTSTEP-CMAQ brings more SO_2 in contact with ammonia; therefore, more sulfate is formed whereas the amount of nitrate is diminished. This effect of enhanced mixing is more pronounced in the shallower wintertime boundary layer.

5. Conclusions

Current AQMs use a single, global time step to synchronize different processes in all of the grid cells. When the local characteristic time for various processes is longer than this global time step, computations are inefficient because a result with similar accuracy can be obtained by using a longer time step. The variable time-step algorithm, VARTSTEP, was developed to improve computational efficiency by using local time steps that are more in tune with characteristic process times of each grid cell. A two-dimensional version of the algorithm was implemented in CMAQ; it reduced the computation time by as much as 50% (13% for the less computationally intensive summertime episode). In CMAQ the characteristic time for advection is used to synchronize various processes. Since the CMAQ grid is uniform, all of the variability in local time steps comes from the wind speeds. In a non-uniform grid model such as the adaptive grid version of CMAQ (García-Menéndez et al., 2010) the variability would be amplified by varying grid lengths. With larger variability in local time steps, VARTSTEP would provide larger speedups compared to the single, global time step alternative.

The differences between concentration estimates of VARTSTEP–CMAQ and benchmark CMAQ are within 3% for PM_{2.5} but can be larger for certain PM_{2.5} components, especially sulfate and nitrate. Due to enhanced mixing of elevated sulfur emissions with ground-level ammonia emissions, and since free ammonia prefers to neutralize sulfate over nitrate in CMAQ, VARTSTEP results in higher sulfate and lower nitrate concentrations. This is more pronounced in the shallower wintertime boundary layer, when sulfate calculated by VARTSTEP–CMAQ is 12% higher. According to Morris et al. (2005), CMAQ underestimates sulfate by approximately the same amount during the same winter period. This does not necessarily mean that VARTSTEP–CMAQ is a better model than CMAQ. In fact, CMAQ should be more representative of the real atmosphere since it couples the processes, which happen simultaneously in reality, more frequently. However, a possible existing error in CMAQ, for example limited vertical mixing in the boundary layer, may be compensating for the less frequent coupling, leading to the right amount of mixing between the elevated SO₂ and ground-level ammonia emissions in VARTSTEP–CMAQ. More analysis is needed to investigate the exact reason for this favorable result with VARTSTEP–CMAQ.

In the Southeastern U.S., sulfate is a much smaller component of PM_{2.5} in winter (compared to summer); therefore, along with 5.0% lower nitrate, the PM_{2.5} calculated by VARTSTEP is only 1.5% higher than the benchmark. Since different PM_{2.5} components have very different contributions to light extinction, and since sulfate and nitrate are two important contributors, the results for visibility may be very different even though total PM_{2.5} is very similar. Light extinctions were calculated using the IMPROVE formula (Malm et al., 2000). The light extinction calculated from VARTSTEP–CMAQ is 1.6% higher than the benchmark in winter and 2.5% higher in summer. Finally, for ozone, the differences are within 1% in both winter and summer. Clearly, VARTSTEP–CMAQ can produce PM_{2.5}, visibility, and ozone result very similar to CMAQ with greater computational efficiency. The same assessment can be made for most PM_{2.5} components with a few exceptions, sulfates and nitrates in particular. In wintertime source apportionments, VARTSTEP–CMAQ is likely to attribute greater air quality impacts than CMAQ to SO₂ emitting sources, such as coal-fired power plants.

Acknowledgements

This research was supported in part by the Visibility Improvement State and Tribal Association of the Southeast (VISTAS) and by the U.S. Department of Defense, through the Strategic Environmental Research and Development Program (SERDP).

Supporting Material Available

Rotating cone test: (a) Initial cone (b) after 1 rotation using global time step, (c) after 1 rotation with VARTSTEP, (d) difference created by VARTSTEP (Figure S1); Total mass in the grid cells after each time step of the rotating cone test (Figure S2); Daily average PM_{2.5} concentrations over the Southeastern U.S. on 16 July 2001 as calculated by the benchmark CMAQ simulation and the VARTSTEP–CMAQ simulation (Figure S3). This information is available free of charge via the Internet at <http://www.atmospolres.com>.

References

- Anderson, D.A., Tannehill, J.C., Pletcher, R.H., 1984. Computational Fluid Mechanics and Heat Transfer, Hemisphere Publishing Corp, New York, pp. 70-77.
- Blom, J.G., Verwer, J.G., 2000. A comparison of integration methods for atmospheric transport-chemistry problems. *Journal of Computational and Applied Mathematics* 126, 381-396.
- Boylan, J.W., Odman, M.T., Wilkinson, J.G., Russell, A.G., Doty, K.G., Norris, W.B., McNider, R.T., 2002. Development of a comprehensive, multiscale “one atmosphere” modeling system: application to the Southern Appalachian Mountains. *Atmospheric Environment* 36, 3721-3734.
- Byun, D., Schere, K.L., 2006. Review of the governing equations, computational algorithms, and other components of the Models-3 Community Multiscale Air Quality (CMAQ) modeling system. *Applied Mechanics Reviews* 59, 51-77.
- Collela, P., Woodward, P.R., 1984. The piecewise parabolic method (PPM) for gas-dynamical simulations. *Journal of Computational Physics* 54, 174-201.
- Constantinescu, E.M., Sandu, A., Carmichael, G.R., 2008. Modeling atmospheric chemistry and transport with dynamic adaptive resolution. *Computational Geosciences* 12, 133-151.
- Garcia-Menendez, F., Yano, A., Hu, Y., Odman, M.T., 2010. An adaptive grid version of CMAQ for improving the resolution of plumes. *Atmospheric Pollution Research* 1, 239-249.
- Hu, Y., Odman, M.T., Russell, A.G., 2006. Mass conservation in the Community Multiscale Air Quality model. *Atmospheric Environment* 40, 1199-1204.
- Lanser, D., Verwer, J.G., 1999. Analysis of operator splitting for advection-diffusion-reaction problems from air pollution modeling. *Journal of Computational and Applied Mathematics* 111, 201-216.
- Malm, W.C., Pitchford, M.L., Scruggs, M., Sisler, J.F., Ames, R., Copeland, S., Gebhart, G.A., Day, D.E., 2000. *Spatial and Seasonal Patterns and Temporal Variability of Haze and its Constituents in the United States, Report III*, ISSN: 0737-5352-47, Colorado State University: Fort Collins, CO, pp. 3.1-3.38.
- Morris, R.E., McNally, D.E., Tesche, T.W., Tonnesen, G., Boylan, J.W., Brewer, P., 2005. Preliminary evaluation of the Community Multiscale Air Quality Model for 2002 over the Southeastern United States. *Journal of the Air and Waste Management Association* 55, 1694-1708.
- Odman, M.T., Khan, M.N., Srivastava, R.K., McRae, D.S., 2002. Initial application of the adaptive grid air pollution model, in *Air Pollution Modeling and its Application XV*, edited by C. Borrego and G. Schayes, Academic/Plenum Publishers, New York, pp. 319-328.
- Park, R.J., Pickering, K.E., Allen, D.J., Stenchikov, G.L., Fox-Rabinovitz, M.S., 2004. Global simulation of tropospheric ozone using the University of Maryland Chemical Transport Model (UMD-CTM): 2. Regional transport and chemistry over the central United States using a stretched grid. *Journal of Geophysical Research D: Atmospheres* 109, art. no. D09303.
- Sportisse, B., 2000. An analysis of operator splitting techniques in the stiff case. *Journal of Computational Physics* 161, 140-168.
- Strang, G., 1968. On the construction and comparison of difference schemes. *SIAM Journal of Numerical Analysis* 5, 506-517.