A comparison of integration methods for atmospheric transport-chemistry problems

J.G. Blom, J.G. Verwer

Modelling, Analysis and Simulation (MAS)

MAS-R9910 April 30, 1999

Report MAS-R9910
ISSN 1386-3703

CWI
P.O. Box 94079
1090 GB Amsterdam
The Netherlands

P.O. Box 94079, 1090 GB Amsterdam (NL)
Kruislaan 413, 1098 SJ Amsterdam (NL)
Telephone +31 20 592 9333
Telefax +31 20 592 4199

A Comparison of Integration Methods for Atmospheric Transport-Chemistry Problems

J.G. Blom and J.G. Verwer
CWI
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

ABSTRACT
This paper is devoted to the time integration of atmospheric transport-chemistry problems. Due to the large number of species and the 3D nature off-the-shelf solvers are not feasible. This has led to the use of special techniques. Most popular is operator splitting. This paper presents a comparison between standard operator splitting, source splitting and approximate matrix factorization. All methods under consideration are comparable in costs measured step wise. The comparison is directed at real-life problems. For that purpose a regional air pollution model is used.

1991 Mathematics Subject Classification: Primary: 65M06, 65M20. Secondary: 65Y20.
1991 Computing Reviews Classification System: G.1.8 and G.1.1,J.2.
Keywords and Phrases: Numerical time integration, air pollution problems.
Note: Work carried out under project MAS1.1 'Numerical algorithms for air quality modeling'.
Keywords and Phrases: Time integration; advection-diffusion-reaction; air pollution.

1. Introduction

We consider the time integration of stiff systems of type

$$\begin{aligned} \frac{\partial c}{\partial t} + \nabla \cdot (\underline{u} c) &= \nabla \cdot (\underline{K}\, \nabla c) + R\,(c)\,, \\ & c = c(\underline{x}, t), \quad c \in \mathbb{R}^m, \quad \underline{x} \in \Omega \subset \mathbb{R}^3, \end{aligned} \tag{1.1}$$

describing transport and chemistry in the atmosphere [7, 19]; $\underline{u}$ is a wind field and the diffusion term represents parametrized atmospheric turbulence; $\underline{u}$ and $\underline{K}$ are supposed to be given (off-line modeling), so that the problem is linear with respect to the transport part; c is a vector of m concentrations of trace gases. The reaction term R introduces stiffness into the problem as the range of characteristic reaction times in the atmosphere is huge. The coefficients and the reaction term R are allowed to depend on the spatial variable $\underline{x}$ and time t. As a rule this dependence is suppressed in our notation. Boundary conditions for (1.1) will be specified only when explicitly needed.

We assume the method of lines (MOL) approach, i.e., the PDE system with its boundary conditions is discretized in space on an Eulerian grid to yield a huge ODE system which then needs to be integrated in time. An extensive survey devoted to many different aspects of time integration of atmospheric problems (splitting, numerical advection, stiff solvers, high performance computing) can be found in [16]. The purpose of the current paper is to present a more specific comparison between standard operator splitting, source splitting and approximate matrix factorization implemented in a Rosenbrock method. In [17] we have applied these methods to a constructed test model. That model, however, was found too simple for drawing sufficiently reliable conclusions on how the methods would perform in actual practice. The current comparison is therefore directed at a more realistic regional air pollution model.

By way of introduction we start in Section 2 with a few preliminaries on integrating large-scale atmospheric transport-chemistry problems. In Section 3 we define the methods used in the comparison.

Section 4 describes the model. Results of the comparison are presented in Section 5. Section 6 summarizes our main conclusions.

2. PRELIMINARIES

Mass conservation All spatial terms are discretized in flux form to conserve mass for semi-discrete solutions. To avoid wiggles and negative concentrations, advection is discretized with the flux-limited, third-order upwind scheme from [8] surveyed in [16]. In each coordinate direction, the diffusion term is discretized on the standard 3-point stencil. When Runge-Kutta formulas are used for integrating transport terms, the property of mass conservation of the semi-discrete system carries over. Locally, molecular mass is conserved in the chemical kinetics system $\dot{c} = R(c)$. Implicit and linearly implicit ODE solvers applied to $\dot{c} = R(c)$ conserve mass when they work with the true analytic Jacobian matrix $R'(c)$. All methods we compare use a variant of the Runge-Kutta-Rosenbrock method (3.1) and conserve mass.

Positivity Positivity (non-negativity) is essential for a stable chemistry solution. Maintaining positivity in the integration of transport terms renders no serious difficulties, in contrast to the chemistry integration. The only solid positive method we know of is implicit Euler. However, Euler is only first order consistent and it requires an iterative technique, e.g. modified Newton, which not always converges for large step sizes. We prefer to avoid this and therefore favor a non-iterative stiff ODE solver, viz. the Rosenbrock method (3.1). This method performs very well, but does not guarantee positivity. We enforce positivity by clipping (negative concentrations are put to zero). Clipping creates mass errors. These errors are minor if clipping occurs only occasionally, which is the case for the Rosenbrock method.

Accuracy and stability Air pollution models require low accuracies, roughly 1 to 10% for output species. So low order methods (splitting) are suitable. Stability is a major concern though, in particular for the stiff chemistry solution. With the Rosenbrock method (3.1) tropospheric gas-phase chemistry can be handled with step sizes up to about 15 to 30 minutes, constant in time and space over the grid. Much larger step sizes are out of the question due to the photochemistry which results in temporal gradients at sunrise and sunset. These gradients move and oscillate over the grid. Sudden emissions and the characteristic reaction times for main output species also limit the step size. The range of reaction times is huge, from milliseconds and shorter to years (e.g. OH radical to CH_4).

Efficiency In spite of the tremendous increase in computer speed during the last years, speed is still a vital factor (see [3, 16] for HPCN aspects). Global and large-scale regional air pollution models can require excessive CPU times, caused by large numbers of species (between 20 and 100), large numbers of points in the 3D grids (from thousands to a few million) and long time spans (from weeks to years). The chemistry computation is normally most expensive (always more than 50% of CPU time) and must be carried out in a manner such that at the level of the numerical algebra only box-models are solved. The main reason is that one then can exploit sparsity of the chemistry Jacobian. It is stressed that splitting and approximate matrix factorization enable this. Even coupling the chemistry solution to 1D vertical diffusion, leading to band matrices only, but also to loss of sparsity, becomes too expensive when the number of species is large [15].

Splitting Operator splitting is in vogue already for a long time. In the atmospheric modeling field it is the standard way of solving the 3D transport-chemistry problem since the paper by McRae, Goodin and Seinfeld [10]. The basic idea of operator splitting is to treat processes like advection, diffusion and chemistry on their own in numerical time-stepping, so as to enable an easy use of well prepared, tailored solvers for these different subprocesses. A disadvantage of this method is that a discontinuity in the concentrations occurs at every time step taken to solve the chemistry process. In general this will result in stiff transients and thus in a laborious solution of the chemistry part. An alternative splitting which avoids these discontinuities is source splitting (see page 4 for more details). Furthermore, splitting gives rise to splitting errors which come on top of integration errors [9]. But, as said above, some form of splitting is of major importance to achieve high efficiency in the chemistry integration. Comparison with alternative approaches of comparable stepwise efficiency is therefore of clear interest.

3. THE INTEGRATION METHODS

Consider an arbitrary ODE system $\dot{y} = f(y)$. Let $\gamma = 1 + \frac{1}{2}\sqrt{2}$ and A an approximating matrix for the Jacobian. In [17] we studied variants of the second-order Rosenbrock method

$$\begin{aligned} y_{n+1} &= y_n + \frac{1}{2}k_1 + \frac{1}{2}k_2, \\ (I - \gamma\tau A)\,k_1 &= \tau f(y_n), \\ (I - \gamma\tau A)\,k_2 &= \tau f(y_n + k_1) - 2\gamma\tau A k_1, \end{aligned} \tag{3.1}$$

for integrating the various subsystems of 3D semi-discrete air pollution models. We use this method also in the current comparison, the main reason being that it is very suitable for stiff atmospheric chemistry. In addition, it allows adjusting A to the system at hand, while maintaining second-order consistency (W-method). With A the Jacobian $f'(y_n)$, it is L-stable. With A the zero matrix, the explicit trapezoidal rule

$$y_{n+1} = y_n + \frac{1}{2}\tau f(y_n) + \frac{1}{2}\tau f(y_n + \tau f(y_n)) \tag{3.2}$$

is obtained. This explicit method offers favorable stability and positivity properties for advection when combined with flux-limited, third-order upwind. Stability and positivity is guaranteed for CFL number 0.5. However, this theoretical bound is rather restrictive. Experiments have shown that a CFL number of 0.67 is in practice sufficient for the solution values to remain positive [8]. So, the Rosenbrock W-method offers interesting choices within the framework of splitting. Most interesting, however, is that (3.1) can also be applied directly to the full semi-discrete system if the choice for A is based on the idea of approximate matrix factorization.

We denote the semi-discrete system obtained by spatial discretization of (1.1) by

$$\dot{w} = F(w) \equiv F_A(w) + F_D(w) + F_R(w), \tag{3.3}$$

where the vector function F is split into functions F_A, F_D and F_R such that

- $\dot{w} = F_A(w)$ contains all advection terms extended with horizontal diffusion terms. These are kept together since both allow explicit time stepping for stability. Observe that F_A is nonlinear due to the flux-limiting; $F_A(w)$ decouples into m subsystems, one for each species.

- $\dot{w} = F_D(w)$ contains only vertical diffusion terms. This system usually requires implicit time stepping for stability; $F_D(w)$ is linear and decouples into m subsystems and each subsystem decouples over the horizontal grid. Since diffusion is discretized on a 3-point stencil only tridiagonal implicitness is encountered.

- $\dot{w} = F_R(w)$ contains all chemical reactions with emission and deposition. Of importance is that $\dot{w} = F_R(w)$ is decoupled over the grid. So per grid cell we encounter a stiff sparse nonlinear system of dimension m.

We now define the actual methods for system (3.3). There are four of them, two of order 1 and two of order 2. With respect to stability all methods have more or less the same characteristics for air pollution models: the critical step size is the same as in the explicit trapezoidal rule applied to the advection part only. For specific details see [16, 17].

Let $\Phi_A(t_n;\tau)$ denote the integrator for $F_A(w)$ stepping from t_n to t_{n+1}. Introduce similar notations $\Phi_D(t_n;\tau), \Phi_R(t_n;\tau)$ for F_D, F_R.

(I) Method (I) is the first-order operator-splitting method

$$w_{n+1} = \Phi_R(t_n;\tau)\,\Phi_D(t_n;\tau)\,\Phi_A(t_n;\tau)w_n, \tag{3.4}$$

where Φ_A is defined by the explicit trapezoidal rule (3.2) and Φ_D and Φ_R by the original second-order Rosenbrock method (3.1) using the true Jacobian.

(II) Method (II) is the second-order Strang version of (I), i.e.

$$w_{n+1} = \Phi_A(t_{n+1/2}; \frac{\tau}{2})\Phi_D(t_{n+1/2}; \frac{\tau}{2})\Phi_R(t_n; \tau)\, \Phi_D(t_n; \frac{\tau}{2})\, \Phi_A(t_n; \frac{\tau}{2}) w_n. \tag{3.5}$$

In both methods the initial value for the chemistry integration is in general not a result of the previous chemistry step. So at each splitting step the computed concentrations are 'discontinuous' for the chemistry integration, resulting in stiff transients. These transients are an artifact of the splitting and may complicate the numerical chemistry solution due to the nonlinearity. Methods (III) and (IV) avoid this artifact.

(III) Method (III) is a source-splitting method. Source splitting circumvents solution discontinuities for the stiff chemistry integration by treating transport as a piecewise constant source. That is, at successive split intervals, (3.3) is approximated by

$$\frac{d\tilde{w}}{dt} = F_R(\tilde{w}) + \frac{v_{n+1} - \tilde{w}(t_n)}{\tau}, \quad t_n \le t \le t_{n+1}, \tag{3.6}$$

where v_{n+1} is the solution at $t = t_{n+1}$ of the initial value problem

$$\dot{v} = F_T(v) \equiv F_A(v) + F_D(v), \quad v(t_n) = \tilde{w}(t_n).$$

Source splitting yields first-order consistency in τ and one has basically the same freedom as in standard splitting for one's favorite combination of algorithms. Our method (III) uses for the transport problem $\dot{v} = F_T(v)$ the Rosenbrock W-method

$$\begin{aligned} v_{n+1} &= v_n + \frac{1}{2}k_1 + \frac{1}{2}k_2, \\ (I - \gamma\tau A)\, k_1 &= \tau F_T(v_n) \\ (I - \gamma\tau A)\, k_2 &= \tau F_T(v_n + k_1) - 2\gamma\tau A\, k_1, \end{aligned} \tag{3.7}$$

with $A = F_D'(v_n)$, which is linearly implicit for the vertical diffusion and explicit for advection. The chemistry system (3.6) is again solved with one step of the Rosenbrock method (3.1) using the true Jacobian matrix.

(IV) Method (IV) is the Rosenbrock W-method

$$\begin{aligned} w_{n+1} &= w_n + \frac{1}{2}k_1 + \frac{1}{2}k_2, \\ (I - \gamma\tau A)\, k_1 &= \tau F(w_n) \\ (I - \gamma\tau A)\, k_2 &= \tau F(w_n + k_1) - 2\gamma\tau A\, k_1, \end{aligned} \tag{3.8}$$

applied to the full system (3.3) where $I - \gamma\tau A = (I - \gamma\tau F_D'(w_n))\,(I - \gamma\tau F_R'(w_n))$. Hence we factorize $I - \gamma\tau A$ approximately, treating advection explicitly and vertical diffusion and chemistry linearly implicitly. In a sense we split at the numerical algebra level, maintaining the computational advantages of standard splitting. Method (3.8) is second order consistent. Just like method (III) it avoids the problem of stiff transients in the solution of the chemistry. It differs from the previous three splitting methods in that it is consistent for stationary problems ($w_{n+1} = w_n$).

Computational costs We use the same step size τ over the grid, which is attractive for parallel implementations. For all methods τ is limited by a CFL condition (CFL number 0.67) since advection is computed explicitly; τ can further be limited by the nonlinear chemistry. For efficiency it is desirable to integrate with step sizes ranging from 15 to 30 minutes, say.

Generally speaking the costs of the chemistry computation will be dominant. In solving the chemistry problem we exploit sparsity with optimized routines for the LU-decomposition and the backsolve [12]. In principal all methods require per time step one LU-decomposition and two backsolves.

However, for method (IV) this would require the storage of the Jacobian for all grid points, which will be too costly in many cases. For this method either the Jacobians should be recalculated and decomposed or the implementation should be 'blocked' with as many grid points in a block as Jacobians will fit in memory. Methods (I) and (III) require also for the transport part roughly the same computations and storage. Method (II) seems to be the most costly since every time step two advection and two diffusion operators have to be computed. On the other hand, if the step size is limited by the CFL restriction for the advection, the step size can be taken twice as large as for the other three methods.

4. LOTOS: a LOng Term Ozone Simulation model

To compare the different time integration methods in a real-life setting, we implemented them in a three-dimensional regional dispersion model called LOTOS-HPCN that we are developing jointly with the TNO institute for Environment, Energy and Process Innovation. This new 3D model should replace TNO's existing operational LOTOS model (see [5, 11]), which is used for a variety of environmental studies related to air pollution with emphasis on ozone simulations in the troposphere. Both models are driven by analyzed meteorological data (off-line model) and by an emission data base. The domain is part of a shell around the earth. In horizontal direction the boundary surfaces are aligned with longitude and latitude coordinates. The main difference between the old and the new model lies in the vertical coordinate. The old LOTOS model has 4 physically determined layers, of which three are prognostic and a diagnostic surface layer, and a domain top of approx. 2 km. The vertical coordinate of the new LOTOS model is based on the 31-layer ECMWF hybrid coordinate system, which is terrain following on the surface of the earth and has equal-pressure layers at the top of the domain (approx. 20 km). Such a hybrid coordinate system means that the physical domain is defined by space and time dependent input variables: the orography of the earth and the surface pressure. To avoid problems with boundary conditions system (1.1) is solved in a boundary-conforming curvilinear coordinate system, which means that the computational domain is fixed and rectangular. The system of PDEs becomes after transformation slightly more complex but is in principal of the same type as (1.1). On the fixed rectangular grid the transformed problem can then be discretized and solved using standard numerical techniques. For a more elaborate description of the model and the transformation we refer to [4].

Since the LOTOS model will be used to simulate different scenario's it should be flexible both with respect to the physics and the chemistry and with respect to the resolution. E.g., it should be easy to add (parametrized) processes like cumulus convection and the replacement of a chemistry model should create no implementational overhead. For the latter we make use of a chemical preprocessor (KPP [1, 6]), which generates the necessary computer code from the kinetic equations. Different scenario's will lead to a largely different computational complexity: the number of variables can range from hundred thousand to hundred million and the simulation time can be weeks or years. Therefore the computational model is intended to run on different computer platforms like (a cluster of) workstations, massively parallel architectures and vector/parallel supercomputers. To avoid divergence of different implementations aimed at different computer platforms it is highly recommendable to have one implementation of LOTOS. The experiments with a benchmark code (see [2]) on various platforms show that it is possible to have a really transportable code without loosing efficiency if the setup is as simple as possible.

As yet the ozone simulation in LOTOS-HPCN is done only with gas-phase chemistry, viz. a model of CBM-IV type with 26 species and 55 reactions. Nevertheless, the concentration values vary strongly in time (day/night rhythm) and space (emissions, land/sea). Emissions are area and point sources. Both are modeled as source terms. Point sources do occur at a height up to the fourth vertical grid level. Currently the model contains no wet deposition so that only dry deposition is modeled. Although it would perhaps be more natural physically speaking to model dry deposition as a Neumann boundary condition (cf. [4]), in the current model it is implemented as a linear ODE only operating in the lowest vertical grid boxes. All source and sink terms are thought of as being a part of the chemistry operator.

In contrast with the expectations in [14], the placement of the boundary conditions appeared not to be significant in our experiments. Apart from the surface of the earth all necessary concentration values outside the LOTOS model are given. These concentrations are zonal and monthly-averaged values.

The eigenvalues of the chemistry Jacobian $R'(c)$ range typically from $O(-10^3)$ to $O(-10^{-8})$ (min^{-1}). Hence for step sizes τ between 15 and 30 minutes the chemistry computation is highly stiff and requires a robust, stable stiff solver.
The vertical diffusion (m^2/sec) varies between 0 and 10–100 (depending on the time of day). Since the vertical resolution in the high diffusion region is between 200–400 m it is advisable to integrate the vertical diffusion operator also implicitly.

The experiments for this paper are done on a relatively small model: the horizontal domain ranges from 10°W to 60°E and from 35°N to 70°N (Europe) and is divided in 70 × 70 grid cells. In the vertical direction only the lower 8 layers are used giving a domain top of approximately 2 km which is comparable with the old LOTOS model. But even for this small model a straightforward time integration with ROS2 of the semi-discrete system without any form of splitting would imply that every time step two linear systems with a dimension of more than a million would have to be solved.

5. TEST RESULTS

Theoretical results on the error made by splitting operators in time integration are of limited value for system (1.1). Assumptions like commuting operators or linearity which are required to derive these results are not fulfilled and the question is whether theoretical results found can be reproduced under realistic conditions. E.g., in [9] it is proven for nonlinear operators, that there is no split error if the operators commute, which is the case if the windfield $\underline{u}$, the diffusion coefficient $\underline{K}$, and the chemistry operator R are independent of $\underline{x}$, and if R is linear in c. This is of course very unrealistic for air pollution models. On the other hand, the advection and chemistry operator do commute if the windfield is divergence free and if the chemistry is space independent, which is often the case over large areas.

The sequential order in which the operators are solved in a splitting scheme can also be important. In [13, 18] it is shown for a model where the stiff operator (chemistry) is linear that this stiff operator should be the one to finish a split step. In [13, 18] it is also shown that the Strang splitting method (II) may suffer from order reduction from two to one. Therefore it is not obvious that the Strang splitting method (II) will give better accuracy than method (I). But again it is not clear what this implies for actual cases.

For the new LOTOS-HPCN model we have data for one week in the summer of 1994. Meteo data is available on a 3-hourly basis and is linearly interpolated in time. The wind during this week was not very strong: the step size resulting from the CFL restriction varies between 15–20 minutes (cf. Fig. 1). The reaction coefficients, emissions and deposition velocities are frozen during one (split) time integration step.

The reference solution in the plots given below is computed with method (IV) with a very small step size (10 seconds) without clipping of negative values. Method (II) gives the same solution for a split step of 10 seconds. In the actual test runs negative values that resulted from the chemistry computation were cut off. In the sequel method (I) is denoted by 'Split', method (II) by 'Strang', method (III) by 'Source Split', and method (IV) by '$\mathrm{ROS2_W}$'. All tests were done with the wind-dependent variable step size τ given in Fig. 1, $\mathrm{Dt_{adv}}$, and half this time step. Since method (II) takes only half advection time steps 'Strang' is also run with a split step τ of $2\mathrm{Dt_{adv}}$.

In the plots the solid line denotes the reference solution. Approximations with the various step sizes are given by:

dashed line: $\tau = 0.5\mathrm{Dt_{adv}}$,

dash-dotted line: $\tau = \mathrm{Dt_{adv}}$,

dotted line (only for method 'Strang'): $\tau = 2\mathrm{Dt_{adv}}$.

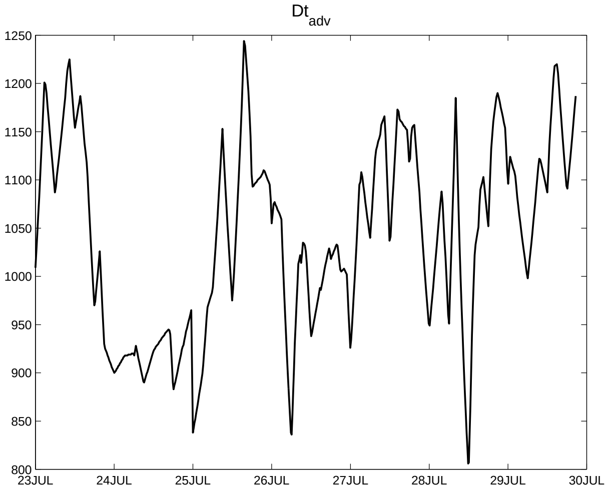

Figure 1: Advection step size in seconds imposed by CFL number 0.67.

Since LOTOS is an ozone simulation model we start with examining the influence of operator splitting on ozone concentrations. In Fig. 2 the time history of the area average of ozone in the surface layer (approx. 50 m) is shown and one can see that there is no significant difference between the concentration values computed by the various methods.

Zooming in (see Fig. 3) shows that method (IV), $\mathrm{ROS2_W}$ is almost exact for both time steps, as is Strang splitting using $\tau = 0.5\mathrm{Dt_{adv}}$. The first-order methods are clearly less accurate. But all results are less than 5% off from the reference solution.

Also important for these kind of models is the simulation of $\mathrm{NO_x}$ ($\mathrm{NO_2}$+NO). In Fig. 4 and Fig. 5 one can clearly see that the first-order methods (I) and (III) are not accurate enough. For the allowed time step the first-order methods commit errors of approximately 10%. Again $\mathrm{ROS2_W}$ is almost plot exact. Also Strang gives good results. For a split step of $2\mathrm{Dt_{adv}}$ the results for Strang are less accurate as can be expected, but more important is the local instable behavior in the last two days. This is presumably due to a too large step size (approx. 40 minutes) for the chemistry during sunrise, since a run with fixed split steps of half an hour gave good results. But whether this is really unstable behavior of ROS2 or whether too much mass is gained because of large negative concentration values is unknown. In any case, it shows that for ROS2 the chemistry step size should be limited when computing e.g. only in lower layers or in very calm periods. Note that even here the diurnal behavior is nicely simulated.

But apart from this local instability our results show that Strang-type operator splitting of second order is not so bad after all in real-life dispersion models. Although we showed here the area averages, the ozone and nitrogen oxide concentrations in specific points do not behave differently.

The question is if there is a situation where split errors do show up. Naturally, one should expect the largest error for rapidly varying radicals (species with a very small reaction time). So looking for a worst-case scenario we investigated the behavior of the radical $\mathrm{N_2O_5}$ near an emission peak in England. Here the chemistry is truly space dependent and so the theory [9] says that the chemistry operator will not commute with advection and with vertical diffusion. Indeed, Fig. 6 shows that the two methods (III) and (IV) that are continuous in the stiff chemistry follow the true solution well, although the first-order 'Source Split' method is less accurate and misses the peak at July 24 completely for the time step sizes given by $\mathrm{Dt_{adv}}$. On the other hand, the two operator-splitting methods (I) and (II), shown in the upper two plots of Fig. 7, give completely wrong values. Peaks are in the wrong place and have

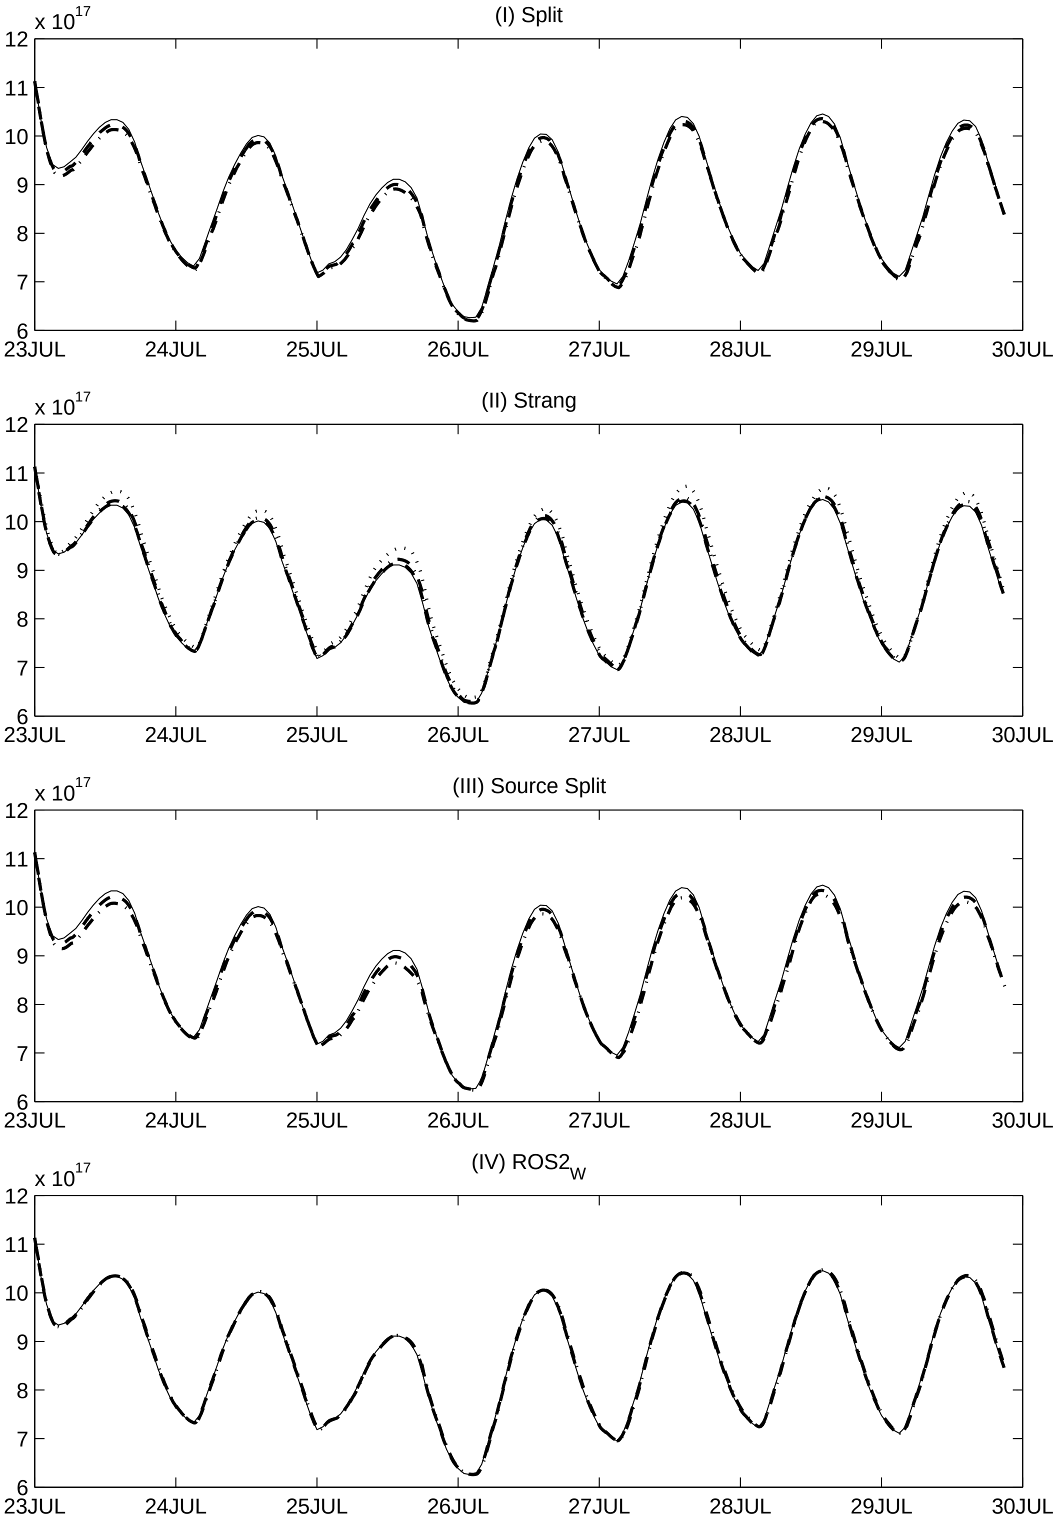

Figure 2: Area average of ozone in the surface layer (# molec./m^3).

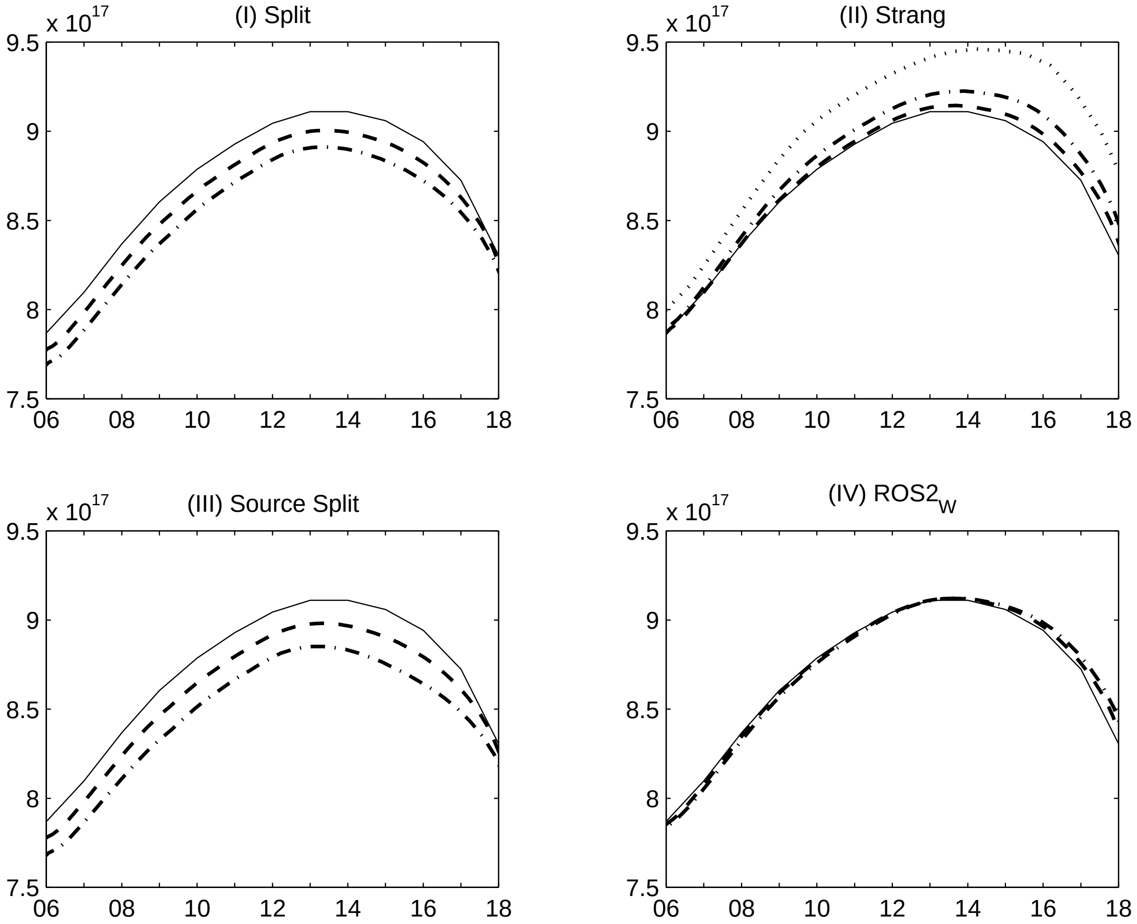

Figure 3: Area average of ozone in the surface layer, 25 July.

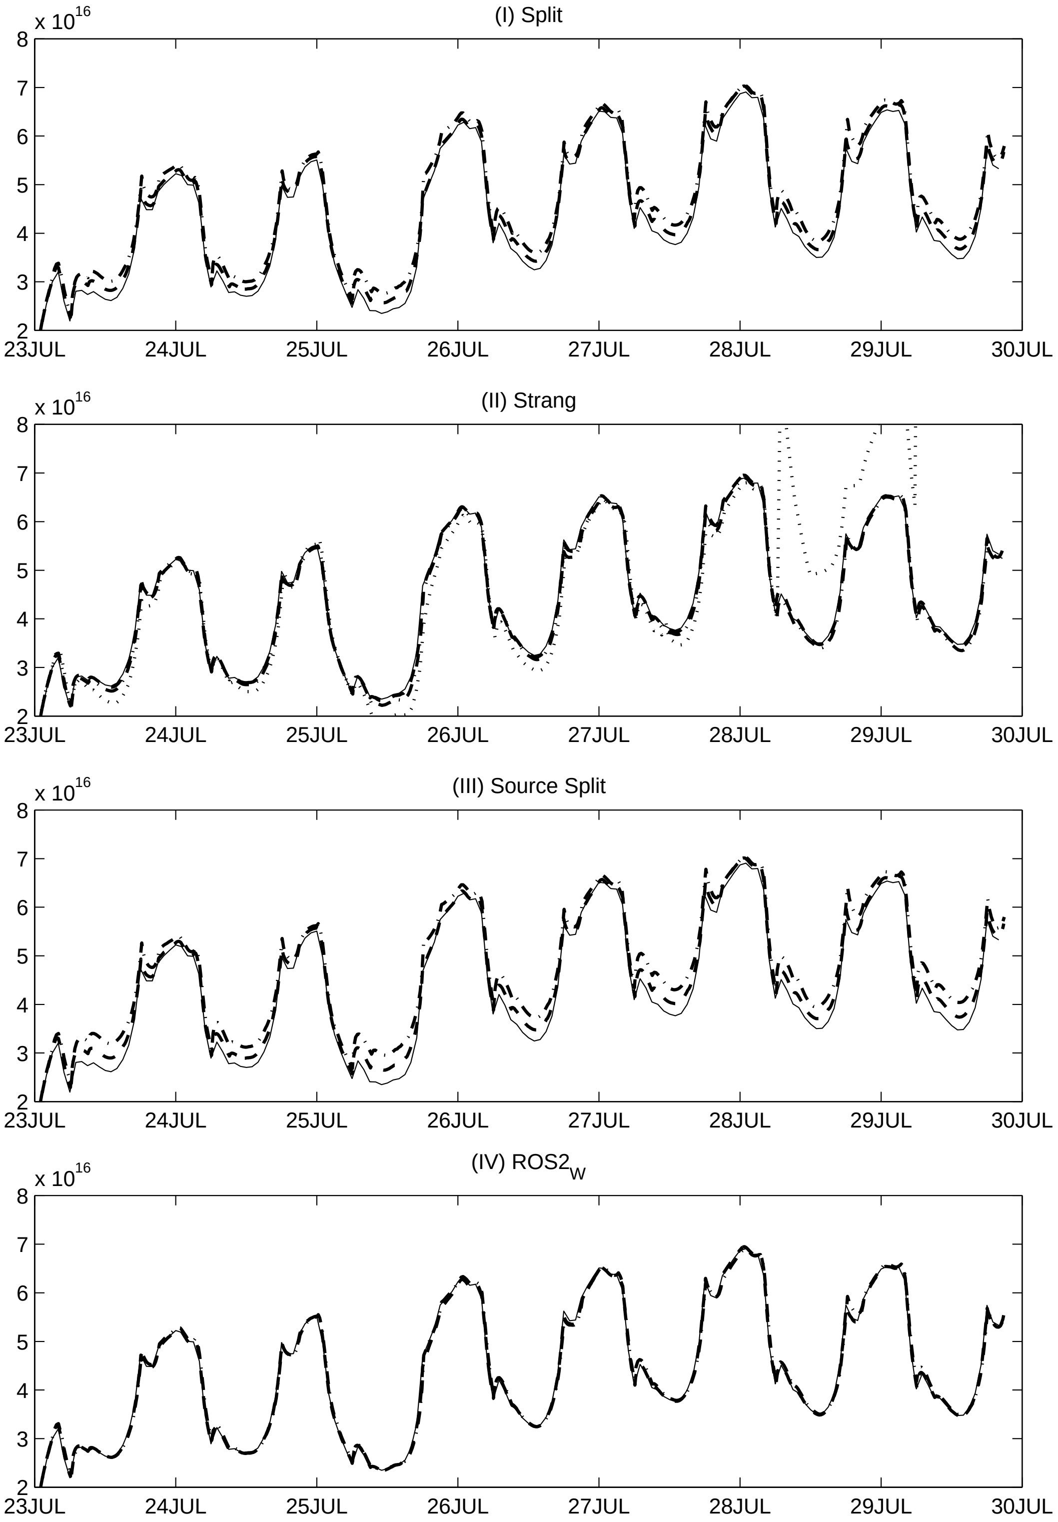

Figure 4: Area average of NO_x in the surface layer.

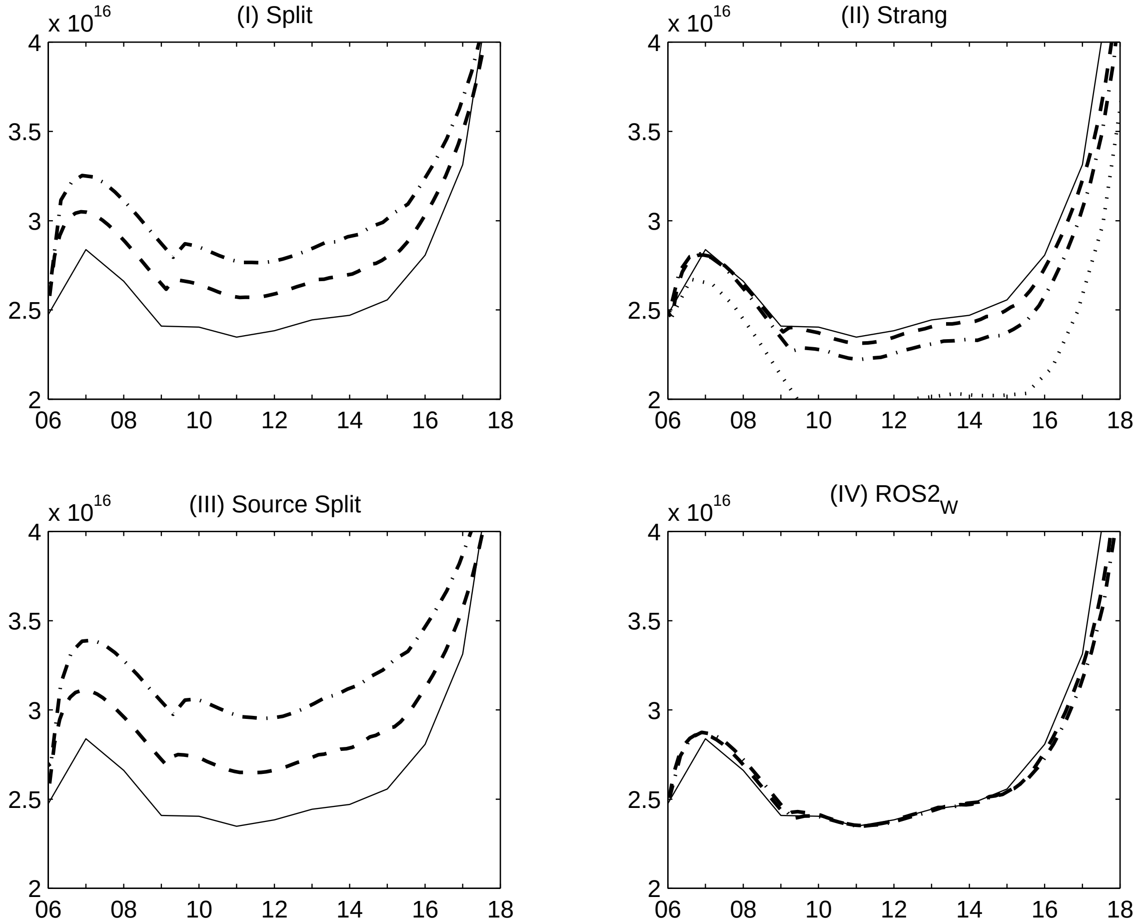

Figure 5: Area average of $\mathrm{NO_x}$ in the surface layer, 25 July.

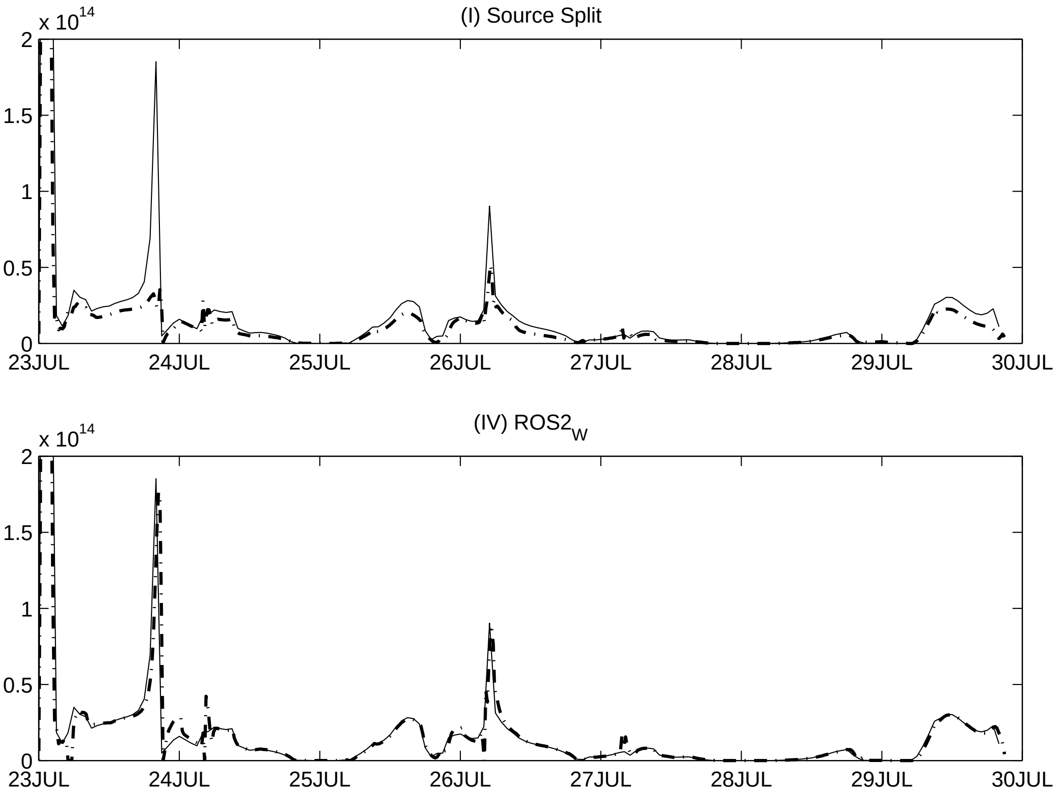

Figure 6: Radical N_2O_5 in box at surface layer with strong emission. Continuous chemistry integration.

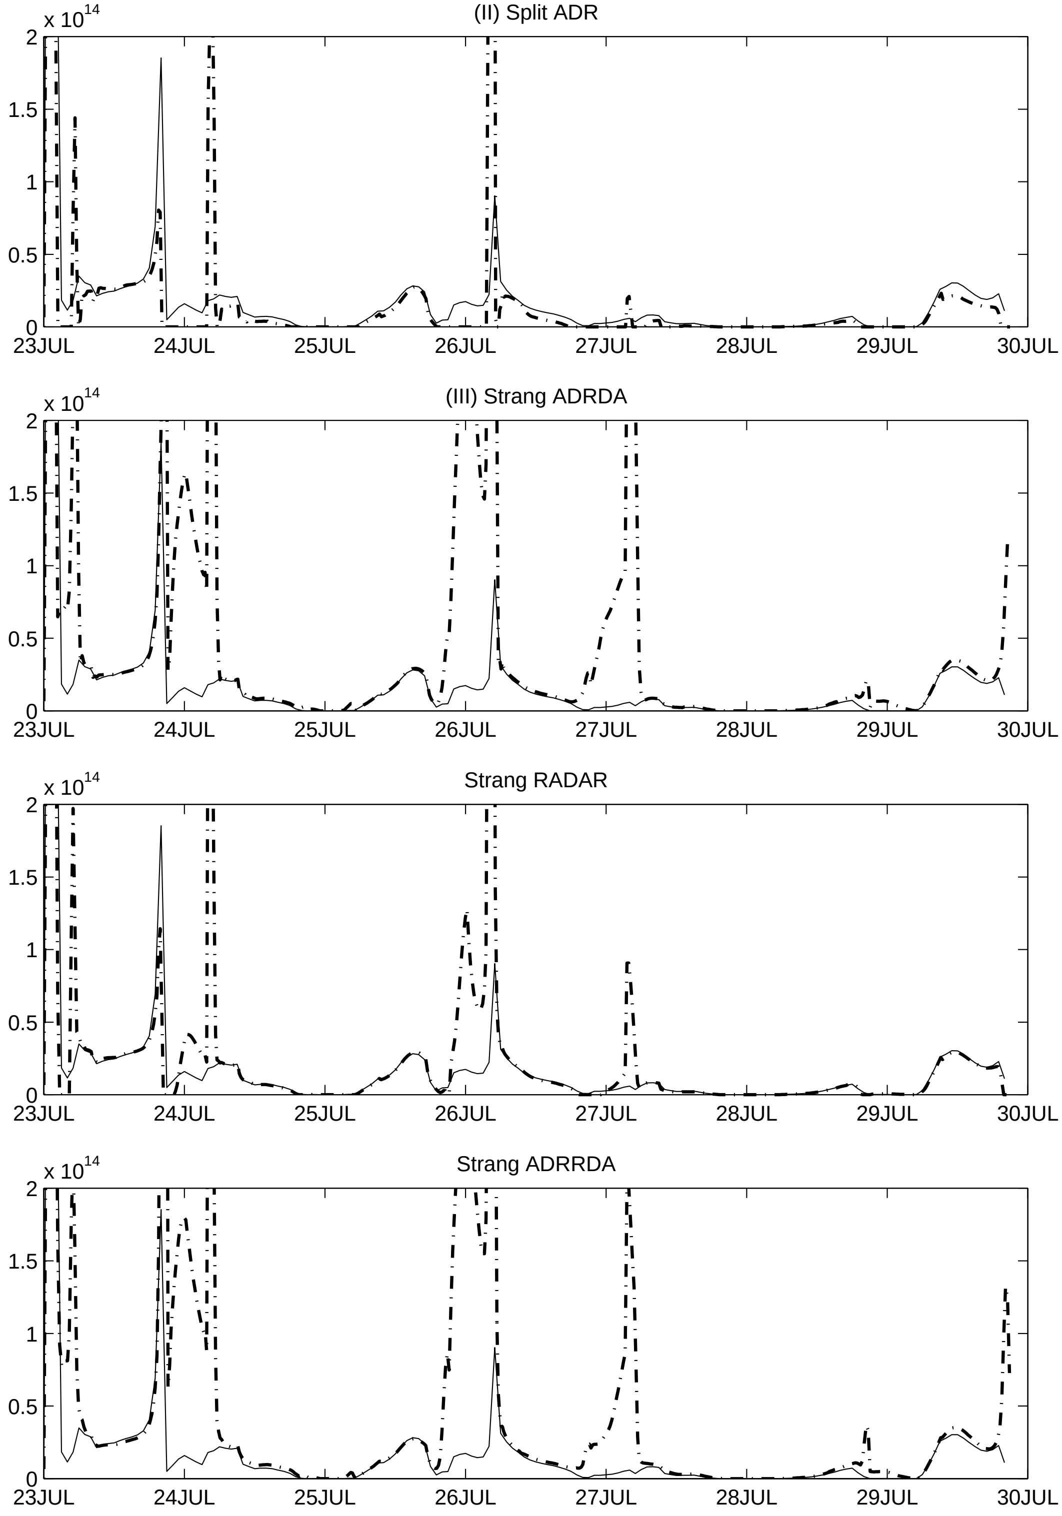
(II) Split ADR

(III) Strang ADRDA

Strang RADAR

Strang ADRRDA

Figure 7: Radical N_2O_5 in box at surface layer with strong emission. Operator splitting.

the wrong height. At first sight it seems perhaps strange that the first-order Split gives better results than the second-order Strang. But the linear theory in [13, 18] tells that the error should decrease if the stiff operator is at the end of the split step, which is the case for method (I), where the operator order is advection-diffusion-chemistry (ADR), but not for method (II) (ADRDA). To check this we rearranged the order in the Strang splitting into chemistry-advection-diffusion-advection-chemistry (RADAR). The results in the third picture of Fig. 7 show that it is indeed better but still not as good as Split. Note, that the improvement is not caused by the more accurate solution of chemistry. The last picture in Fig. 7 shows that taking two chemistry integration steps (ADRRDA) in stead of one has no large influence on the results.

Although it is nice to see a confirmation of theory, one should bear in mind that the sequence RADAR is more expensive than the usual sequence ADRDA because two chemistry steps are needed. This can of course be partly anticipated by a 'staggered' implementation, but this makes the code much more complex. We also stress that the split errors we observe are local in time, in accordance with the results of the model study in [18]. The errors do not accumulate in time, which we owe to the good performance of the chemistry solver ROS2. This solver is L-stable and able to eliminate the large errors for the radicals within one integration step.

Remains whether in actual practice the correct simulation of radicals is very important. Looking at the results for ozone and nitrogen oxide one would perhaps think that this is not the case for this type of air pollution models. On the other hand, currently only gas-phase chemistry is involved and it is known that radicals have a large influence on atmospheric aerosol processes. In the near future an aerosol module will be incorporated in the LOTOS-HPCN model. Then a more decisive judgment can be given on the impact of operator splitting and the numerically most accurate method $\mathrm{ROS2}_{\mathrm{W}}$.

6. SUMMARY AND CONCLUDING REMARKS

In this paper we have discussed time integration aspects for atmospheric transport-chemistry problems from a practical point of view. We assume a method of lines approach in which the system of PDEs (1.1) is discretized in space and the resulting system of ODEs should be solved in time. The stiff chemistry which is part of these models requires implicit time integration. Since in practical applications the order of this system lies in the mega to giga range, it is impossible to solve the complete system of ODEs with an off-the-shelf implicit time solver. In the air pollution field the standard way to solve these systems is *operator splitting*, integrating the different physical and chemical processes and subgrid parametrizations separately and sequentially. The advantage is clear: not only are the systems to solve much smaller, but the time integration method can also be tailored to the operator to be integrated. The disadvantage is twofold: on top of the (controlled) error made in the separate processes comes an extra splitting error which in general is hoped to be comparable. Moreover, the solutions of the separate processes have no physical meaning which is most clearly felt in the time integration of the chemistry, because due to the splitting the initial condition for the chemical ODE system is in general far from chemical equilibrium resulting in artificial stiff transients.

An alternative to operator splitting for making time integration feasible is *source splitting*. This technique avoids discontinuities in the chemistry solution by incorporating all other operators as a source term in the chemistry integration.

These splitting methods share the problem of implementing the boundary conditions. It is not always clear with which operator which boundary condition(s) should be integrated in time. The last method $\mathrm{ROS2}_{\mathrm{W}}$ avoids this problem, as well as the artificial stiff transients for the chemistry. Here the splitting is at the linear algebra level: the linear systems are solved with an *approximate Jacobian* to make the solution process feasible.

We have compared these time integration techniques in a real-life 3D air pollution model LOTOS-HPCN. Operator splitting has been implemented in a straightforward way, resulting in a first-order method, and in a symmetric way giving second order. The four time integration techniques are all based on a second-order Rosenbrock method, which maintains its second-order consistency if used with an approximate Jacobian. All are comparable in computational costs and have analogous vectorization and parallelization possibilities.

Our main conclusions based on these experiments are:

- First-order time integration is not accurate enough. For time steps of 15–20 *min.* the second-order Strang-type operator splitting and in particular the second-order Rosenbrock method with approximate factorization give accurate results for important species like ozone and nitrogen oxides. The first-order operator splitting and the source-splitting method show deviations of over 10%.
- Splitting errors are not clearly seen in major species. But for very fast varying species (radicals) operator-splitting methods are not capable of giving even a qualitative idea of the evolution of the solution. The order in which the operators are handled is of importance but to resolve radicals operator splitting is not the way to go. In the current model the correct simulation of radicals appears not to be significant for long-term ozone simulation. But it should be remembered that radicals are important in aerosol processes and thus can influence the ozone formation.
- Splitting at the linear algebra level (method (IV)) is numerically speaking by far the best option. However, the implementation of method (IV) is more complex than of Strang operator splitting. Which of these two second-order methods is preferable depends on the importance of a correct simulation of the time evolution of radical species.

Finally, it should be noted that for these types of models the investigated methods are nearly optimal qua computational complexity. The step sizes taken in the time integration are determined by the wind velocities resulting in steps varying between 15 and 20 *min.* It is not likely that a designated chemistry solver can take much larger time steps. Since at the computational level advection, diffusion and chemistry are decoupled and since all processes themselves are solved efficiently, there is probably no room for much improvement.

Acknowledgement This work was supported by the Stichting HPCN, with financial support from the Ministry of Economic Affairs through the HPCN-TASC project 'HPCN for Environmental Applications'.

REFERENCES

1. `http://www.cgrer.uiowa.edu/people/vdamian/kpp/`.
2. `http://www.cwi.nl/~gollum/LOTOS/`.
3. J.G. Blom, W.M. Lioen, and J.G. Verwer. HPCN and air quality modeling. In P. Sloot, M. Bubak, and B. Hertzberger, editors, *High-Performance Computing and Networking. International Conference and Exhibition. Amsterdam, The Netherlands, April 1998*, volume 1401 of *Lecture Notes in Computer Science*, pages 141–150. Springer, Berlin, 1998.
4. J.G. Blom and M.G.M. Roemer. Description of the 3D LOTOS model. Part I: Dynamics. Report MAS-N9701, CWI, Amsterdam, 1997.
5. P.J.H. Builtjes. The LOTOS – Long Term Ozone Simulation – project. Summary report. TNO-rapport TNO-MW – R 92/240, TNO Milieuwetenschappen, Delft, 1992.
6. V. Damian-Iordache. KPP – chemistry simulation development environment. Master's thesis, University of Iowa, 1996.
7. T.E. Graedel and P.J. Crutzen. *Atmosphere, Climate and Change.* Scientific American Library, Freeman and Company, New York-Oxford, 1995.
8. W. Hundsdorfer, B. Koren, M. van Loon, and J.G. Verwer. A positive finite-difference advection scheme. *J. Comput. Phys.*, 117:35–46, 1995. Revision of CWI Report NM-R9309.

9. D. Lanser and J.G. Verwer. Analysis of operator splitting for advection-diffusion-reaction problems from air pollution modelling. Report MAS-R9805, CWI, Amsterdam, 1998. (to appear in *JCAM*).
10. G.J. McRae, W. R. Goodin, and J.H. Seinfeld. Numerical solution of the atmospheric diffusion equation for chemically reacting flows. *J. Comput. Phys.*, 45:1–42, 1982.
11. R.E. Morris. User's guide to the Regional Transport Model version III Photochemical Oxidant (PHOXA). Technical Report SYSAPP-87/098, Systems Applications, Inc., San Rafael, CA, 1987.
12. A. Sandu, F.A. Potra, G.R. Carmichael, and V. Damian. Efficient implementation of fully implicit methods for atmospheric chemical kinetics. *J. Comput. Phys.*, 129:101–110, 1996.
13. B. Sportisse. An analysis of operator splitting techniques in the stiff case. Report 98–127, ENPC-CERMICS, Champs sur Marne, France, 1998.
14. B. Sportisse, G. Bencteux, and P. Plion. Method of lines versus operator splitting for reaction-diffusion systems with fast chemistry. In *Proceedings of the International Conference on Air Pollution Modelling and Simulation, APMS'98*, pages 381–392, Champs-sur-Marne, France, 1998.
15. J.G. Verwer and J.G. Blom. On the coupled solution of diffusion and chemistry in air pollution models. In *Procs. ICIAM/GAMM 95,* ZAMM*, Issue 4: Applied Sciences, especially Mechanics*, pages 454–457. Akademie Verlag, 1996.
16. J.G. Verwer, W. Hundsdorfer, and J.G. Blom. Numerical time integration for air pollution models. Report MAS-R9825, CWI, Amsterdam, 1998.
17. J.G. Verwer, E.J. Spee, J.G. Blom, and W. Hundsdorfer. A second order Rosenbrock method applied to photochemical dispersion problems. *SIAM J. Sci. Comput.*, 20:1456–1480, 1999.
18. J.G. Verwer and B. Sportisse. A note on operator splitting in a stiff linear case. Report MAS-R9830, CWI, Amsterdam, 1998.
19. Z. Zlatev. *Computer Treatment of Large Air Pollution Models.* Kluwer Academic Publishers, 1995.